Improved approximate rips filtrations with shifted integer lattices and cubical complexes

Aruni Choudhary[1] · Michael Kerber[2] · Sharath Raghvendra[3]

Abstract

Rips complexes are important structures for analyzing topological features of metric spaces. Unfortunately, generating these complexes is expensive because of a combinatorial explosion in the complex size. For n points in $\mathbb{R}^d$, we present a scheme to construct a 2-approximation of the filtration of the Rips complex in the L_∞-norm, which extends to a $2d^{0.25}$-approximation in the Euclidean case. The k-skeleton of the resulting approximation has a total size of $n2^{O(d \log k + d)}$. The scheme is based on the integer lattice and simplicial complexes based on the barycentric subdivision of the d-cube. We extend our result to use cubical complexes in place of simplicial complexes by introducing *cubical maps* between complexes. We get the same approximation guarantee as the simplicial case, while reducing the total size of the approximation to only $n2^{O(d)}$ (cubical) cells. There are two novel techniques that we use in this paper. The first is the use of *acyclic carriers* for proving our approximation result. In our application, these are maps which relate the Rips complex and the approximation in a relatively simple manner and greatly reduce the complexity of showing the approximation guarantee. The second technique is what we refer to as *scale balancing*, which is a simple trick to improve the approximation ratio under certain conditions.

✉ Aruni Choudhary
arunich@inf.fu-berlin.de

Michael Kerber
kerber@tugraz.at

Sharath Raghvendra
sharathr@vt.edu

[1] Institut für Informatik, Freie Universität Berlin, Berlin, Germany

[2] Institut für Geometrie, Technische Universität Graz, Graz, Austria

[3] Department of Computer Science, Virginia Tech, Blacksburg, VA, USA

1 Introduction

Context. Persistent homology (Carlsson 2009; Edelsbrunner and Harer 2010; Edelsbrunner et al. 2002) is a technique to analyze data sets using topological invariants. The idea is to build a multi-scale representation of data sets and to track its homological changes across the scales.

A standard construction for the important case of point clouds in Euclidean space is the *Vietoris-Rips complex* (usually abbreviated as simply the *Rips complex*): for a scale parameter $\alpha \geq 0$, it is the collection of all subsets of points with diameter at most α. When α increases from 0 to ∞, the Rips complexes form a *filtration*, an increasing sequence of nested simplicial complexes whose homological changes can be computed and represented in terms of a *barcode*.

The computational drawback of Rips complexes is their sheer size: the k-skeleton of a Rips complex (that is, where only subsets of size at most $k+1$ are considered) for n points consists of $\Theta(n^{k+1})$ simplices because every $(k+1)$-subset joins the complex for a sufficiently large scale parameter. This size bound makes barcode computations for large point clouds infeasible even for low-dimensional homological features[1]. This difficulty motivates the question of what we can say about the barcode of the Rips filtration without explicitly constructing all of its simplices.

We address this question using approximation techniques. The space of barcodes forms a metric space: two barcodes are close if similiar homological features occur on roughly the same range of scales. More precisely, the bottleneck distance is used as a distance metric between barcodes. The first approximation scheme by Sheehy (2013) constructs a $(1+\varepsilon)$-approximation of the k-skeleton of the Rips filtration using only $n(\frac{1}{\varepsilon})^{O(\lambda k)}$ simplices for arbitrary finite metric spaces, where λ is the doubling dimension of the metric. Further approximation techniques for Rips complexes (Dey et al. 2014) and the closely related *Čech complexes* (Botnan and Spreemann 2015; Cavanna et al. 2015; Kerber and Sharathkumar 2013) have been derived subsequently, all with comparable size bounds. More recently, we constructed an approximation scheme (Choudhary et al. 2019) for the Čech filtrations of n points in $\mathbb{R}^d$ that had size $n \left(\frac{1}{\varepsilon}\right)^{O(d)} 2^{O(d \log d + dk)}$ for the k-skeleton, improving the size bound from previous work.

In Choudhary et al. (2017b), we constructed an approximation scheme for Rips filtration in Euclidean space that yields a worse approximation factor of only $O(d)$, but uses only $n2^{O(d \log k + d)}$ simplices. In Choudhary et al. (2017b), we also show a lower bound result on the size of approximations: for any $\varepsilon < 1/\log^{1+c} n$ with some constant $c \in (0, 1)$, any ε-approximate filtration has size $n^{\Omega(\log \log n)}$.

---

[1] An exception are point clouds in $\mathbb{R}^2$ and $\mathbb{R}^3$, for which *alpha complexes* (Edelsbrunner and Harer 2010) are an efficient alternative.

There has also been work on using cubical complexes to compute persistent homology, such as in Wagner et al. (2012). Cubical complexes are typically smaller than their simplicial counterparts, simply because they avoid triangulations. However, to our knowledge, there has been no attempt to utilize them in computing approximations of filtrations. Also, while there are efficient methods to compute persistence for simplicial complexes connected with simplicial maps (Dey et al. 2014; Kerber and Schreiber 2017), we are not aware of such counterparts for cubical complexes.

Our contributions. For the Rips filtration of n points in $\mathbb{R}^d$ with distances taken in the L_∞-norm, we present a 2-approximation whose k-skeleton has size at most $n6^{d-1}(2k+4)(k+3)! \left\{ \begin{matrix} d \\ k+2 \end{matrix} \right\} = n2^{O(d \log k + d)}$ where $\left\{ \begin{matrix} a \\ b \end{matrix} \right\}$ denotes Stirling numbers of the second kind. This translates to a $2d^{0.25}$-approximation of the Rips filtration in the Euclidean metric and hence improves the asymptotic approximation quality of our previous approach (Choudhary et al. 2017b) with the same size bound. Our scheme gives the best size guarantee over all previous approaches.

On a high level, our approach follows a straightforward approximation scheme: given a scaled and appropriately shifted integer grid on $\mathbb{R}^d$, we identify those grid points that are close to the input points and build an approximation complex using these grid points. The challenge lies in how to connect these grid points to a simplicial complex such that close-by grid points are connected, while avoiding too many connections to keep the size small. Our approach first selects a set of *active faces* in the cubical complex defined over the grid, and defines the approximation complex using the barycentric subdivision of this cubical complex.

We also describe an output-sensitive algorithm to compute our approximation. By randomizing the aforementioned shifts of the grids, we obtain a worst-case running time of $n2^{O(d)} \log \Delta + 2^{O(d)} M$ in expectation, where Δ is the *spread* of the point set (that is, the ratio of the diameter to the closest distance of two points) and M is the size of the approximation.

Additionally, this paper makes the following technical contributions:

– We follow the standard approach of defining a sequence of approximation complexes and establishing an *interleaving* between the Rips filtration and the approximation. We realize our interleaving using *chain maps* connecting a Rips complex at scale α to an approximation complex at scale $c\alpha$, and vice versa, with $c \geq 1$ being the approximation factor. Previous approaches (Choudhary et al. 2017b; Dey et al. 2014; Sheehy 2013) used *simplicial maps* for the interleaving, which induce an elementary form of chain maps and are therefore more restrictive.

The explicit construction of such maps can be a non-trivial task. The novelty of our approach is that we avoid this construction by the usage of *acyclic carriers* (Munkres 1984). In short, carriers are maps that assign subcomplexes to subcomplexes under some mild extra conditions. While they are more flexible, they still certify the existence of suitable chain maps, as we exemplify in Sect. 2. We believe that this technique is of general interest for the construction of approximations of cell complexes.

– We exploit a simple trick that we call *scale balancing* to improve the quality of approximation schemes. In short, if the aforementioned interleaving maps from and to the Rips filtration do not increase the scale parameter by the same amount, one can simply multiply the scale parameter of the approximation by a constant. Concretely, given maps

$$\phi_\alpha : \mathcal{R}_\alpha \to \mathcal{X}_\alpha \qquad \psi_\alpha : \mathcal{X}_\alpha \to \mathcal{R}_{c\alpha}$$

interleaving the Rips complex $\mathcal{R}_\alpha$ and the approximation complex $\mathcal{X}_\alpha$, we can define $\mathcal{X}'_\alpha := \mathcal{X}_{\alpha/\sqrt{c}}$ and obtain maps

$$\phi'_\alpha : \mathcal{R}_\alpha \to \mathcal{X}'_{\sqrt{c}\alpha} \qquad \psi_\alpha : \mathcal{X}'_\alpha \to \mathcal{R}_{\sqrt{c}\alpha}$$

which improves the interleaving from c to $\sqrt{c}$. While it has been observed that the same trick can be used for improving the worst-case distance between Rips and Čech filtrations,[2] our work seems to be the first to make use of it in the context of approximations.

– We extend our approximation scheme to use cubical complexes instead of simplicial complexes, thereby achieving a marked reduction in size complexity. To connect the cubical complexes at different scales, we introduce the notion of *cubical maps*, which is a simple extension of simplicial maps to the cubical case. While we do not know of an algorithm that can compute persistence for the case of cubical complexes with cubical maps, we believe that this is a first step towards advocating the use of cubical complexes as approximating structures.

Our technique can be combined with dimension reduction techniques in the same way as in Choudhary et al. (2017b) (see Theorems 19, 21, and 22 therein), with improved logarithmic factors. We state the main results in the paper, while omitting the technical details.

Updates from the conference version. An earlier version of this paper appeared at the 25th European Symposium on Algorithms (Choudhary et al. 2017a). In that version, we achieved a $3\sqrt{2}$-approximation of the L_∞ Rips filtration and correspondingly, a $3\sqrt{2}d^{0.25}$-approximation of the L_2 case. In this version, we improve the weak interleaving of Choudhary et al. (2017a) to a strong interleaving to get improved approximation factors. We expand upon the details of scale balancing, among other proofs that were missing from the conference version. We add the case of cubical complexes in this version.

There is a subtle yet important distinction between the approximation complexes used in the conference version and the current result. In the conference version, our simplicial complex was built using only *active* faces, while the current version uses both *active* and *secondary* faces (please see Sect. 4 for definitions). This makes it easier to relate the simplicial and the cubical complexes in the current version. On the other hand the complexes are different, hence the associated proofs have been adapted accordingly.

---

[2] Ulrich Bauer, private communication.

Outline. We start by explaining the relevant topological concepts in Sect. 2. We give details of the integer grids that we use in Sect. 3. In Sect. 4 we present our approximation scheme that uses the barycentric subdivision, and present the computational aspects in Sect. 5. The extension to cubical complexes is presented in Sect. 6. We discuss practical aspects of our scheme and conclude in Sect. 7. Some details of the strong interleaving from Sect. 4 are deferred to Appendix A.

2 Preliminaries

We briefly review the essential topological concepts needed. More details are available in standard references (see Bubenik et al. 2015; Chazal et al. 2009; Edelsbrunner and Harer 2010; Hatcher 2002; Munkres 1984).

Simplicial complexes. A *simplicial complex* K on a finite set of elements S is a collection of subsets $\{\sigma \subseteq S\}$ called *simplices* such that each subset $\tau \subset \sigma$ is also in K. The dimension of a simplex $\sigma \in K$ is $k := |\sigma| - 1$, in which case σ is called a *k-simplex*. A simplex τ is a *sub-simplex* of σ if $\tau \subseteq \sigma$. We remark that, commonly a sub-simplex is called a "face" of a simplex, but we reserve the word "face" for a different structure. For the same reason, we do not introduce the common notation of of "vertices" and "edges" of simplicial complexes, but rather refer to 0- and 1-simplices throughout. The *k-skeleton* of K consists of all simplices of K whose dimension is at most k. For instance, the 1-skeleton of K is a graph defined by its 0-simplices and 1-simplices.

Given a point set $P \subset \mathbb{R}^d$ and a real number $\alpha \geq 0$, the *(Vietoris-)Rips* complex on P at scale α consists of all simplices $\sigma = (p_0, \cdots, p_k) \subseteq P$ such that $diam(\sigma) \leq \alpha$, where $diam$ denotes the diameter. In this work, we write $\mathcal{R}_\alpha$ for the Rips complex at scale 2α with the Euclidean metric, and $\mathcal{R}_\alpha^\infty$ when using the metric of the L_∞-norm. In either way, a Rips complex is an example of a *flag complex*, which means that whenever a set $\{p_0, \cdots, p_k\} \subseteq P$ has the property that every 1-simplex $\{p_i, p_j\}$ is in the complex, then the k-simplex $\{p_0, \cdots, p_k\}$ is also in the complex.

A related complex is the Čech *complex* of P at scale α, which consists of simplices of P for which the radius of the minimum enclosing ball is at most α. We do not study Čech complexes in this paper, but we mention them briefly while showing a connection with the Rips complex later in this section.

A simplicial complex K' is a *subcomplex* of K if $K' \subseteq K$. For instance, $\mathcal{R}_\alpha$ is a subcomplex of $\mathcal{R}_{\alpha'}$ for $0 \leq \alpha \leq \alpha'$. Let L be a simplicial complex. Let $\hat{\varphi}$ be a map which assigns a vertex of L to each vertex of K. A *simplicial map* is a map $\varphi : K \to L$ induced by a vertex map $\hat{\varphi}$, such that for every simplex $\{p_0, \cdots, p_k\}$ in K, the set $\{\hat{\varphi}(p_0), \cdots, \hat{\varphi}(p_k)\}$ is a simplex of L. For K' a subcomplex of K, the inclusion map $inc : K' \to K$ is an example of a simplicial map. A simplicial map is completely determined by its action on the 0-simplices of K.

Chain complexes. A *chain complex* $\mathcal{C}_* = (\mathcal{C}_p, \partial_p)$ with $p \in \mathbb{Z}$ is a collection of abelian groups $\mathcal{C}_p$ and homomorphisms $\partial_p : \mathcal{C}_p \to \mathcal{C}_{p-1}$ such that $\partial_{p-1} \circ \partial_p = 0$. A simplicial complex K gives rise to a chain complex $\mathcal{C}_*(K)$ for a fixed base field $\mathcal{F}$: define $\mathcal{C}_p$ for $p \geq 0$ as the set of formal linear combinations of p-simplices in K over $\mathcal{F}$, and $\mathcal{C}_{-1} := \mathcal{F}$. The boundary of a k-simplex with $k \geq 1$ is the (signed) sum of

its sub-simplices of co-dimension one[3]; the boundary of a 0-simplex is simply set to 1. The homomorphisms ∂_p are then defined as the linear extensions of this boundary operator. Note that $\mathcal{C}_*(K)$ is sometimes called *augmented chain complex* of K, where the augmentation refers to the addition of the non-trivial group $\mathcal{C}_{-1}$.

A *chain map* $\phi : \mathcal{C}_* \rightarrow D_*$ between chain complexes $\mathcal{C}_* = (\mathcal{C}_p, \partial_p)$ and $D_* = (D_p, \partial'_p)$ is a collection of group homomorphisms $\phi_p : \mathcal{C}_p \rightarrow D_p$ such that $\phi_{p-1} \circ \partial_p = \partial'_p \circ \phi_p$. For simplicial complexes K and L, we call a chain map $\phi : \mathcal{C}_*(K) \rightarrow \mathcal{C}_*(L)$ *augmentation-preserving* if ϕ_{-1} is the identity. A simplicial map $\varphi : K \rightarrow L$ between simplicial complexes induces an augmentation-preserving chain map $\bar{\varphi} : \mathcal{C}_*(K) \rightarrow \mathcal{C}_*(L)$ between the corresponding chain complexes. This construction is *functorial*, meaning that for φ the identity function on a simplicial complex K, $\bar{\varphi}$ is the identity function on $\mathcal{C}_*(K)$, and for composable simplicial maps φ, φ', we have that $\overline{\varphi \circ \varphi'} = \bar{\varphi} \circ \bar{\varphi'}$.

Homology. The *p-th homology group* $H_p(\mathcal{C}_*)$ of a chain complex is defined as $\ker \partial_p / \mathrm{im}\, \partial_{p+1}$. The *p*-th homology group of a simplicial complex K, $H_p(K)$, is the *p*-th homology group of its induced chain complex $\mathcal{C}_*(K)$. Note that this definition is commonly referred to as *reduced* homology, but we ignore this distinction and consider reduced homology throughout. $H_p(\mathcal{C}_*)$ is an $\mathcal{F}$-vector space because we have chosen our base ring $\mathcal{F}$ as a field. Intuitively, when the chain complex is generated from a simplicial complex, the dimension of the *p*-th homology group counts the number of *p*-dimensional holes in the complex. We write $H(\mathcal{C}_*)$ for the direct sum of all $H_p(\mathcal{C}_*)$ for $p \geq 0$.

A chain map $\phi : \mathcal{C}_* \rightarrow D_*$ induces a linear map $\phi^* : H(\mathcal{C}_*) \rightarrow H(D_*)$ between the homology groups. Again, this construction is functorial, meaning that it maps identity maps to identity maps, and it is compatible with compositions.

Acyclic carriers. We call a simplicial complex K *acyclic*, if K is connected and all homology groups $H_p(K)$ are trivial. For simplicial complexes K and L, an *acyclic carrier* Φ is a map that assigns to each simplex σ in K, a non-empty acyclic subcomplex $\Phi(\sigma) \subseteq L$, and whenever τ is a sub-simplex of σ, then $\Phi(\tau) \subseteq \Phi(\sigma)$. We say that a chain $c \in \mathcal{C}_p(K)$ is *carried* by a subcomplex K', if c takes value 0 except for *p*-simplices in K'. A chain map $\phi : \mathcal{C}_*(K) \rightarrow \mathcal{C}_*(L)$ is *carried by* Φ, if for each simplex $\sigma \in K$, $\phi(\sigma)$ is carried by $\Phi(\sigma)$. We state the *acyclic carrier theorem* (Munkres 1984, Thm 13.3), adapted to our notation:

Theorem 1 *Let $\Phi : K \rightarrow L$ be an acyclic carrier. Then,*

- *There exists an augmentation-preserving chain map $\phi : \mathcal{C}_*(K) \rightarrow \mathcal{C}_*(L)$ carried by Φ.*
- *If two augmentation-preserving chain maps $\phi_1, \phi_2 : \mathcal{C}_*(K) \rightarrow \mathcal{C}_*(L)$ are both carried by Φ, then $\phi_1^* = \phi_2^*$.*[4]

We remark that "augmentation-preserving" is crucial in the statement: without it, the trivial chain map (that maps everything to 0) turns the first statement trivial and easily leads to a counter-example for the second claim.

---

[3] To avoid thinking about orientations, it is often assumed that $\mathcal{F} = \mathbb{Z}_2$ is the field with two elements.

[4] In the language of Munkres (1984), this result is stated as the existence of a *chain homotopy* between ϕ_1 and ϕ_2. As evident from Munkres (1984), Theorem 12.4, this implies that the induced linear maps are the same.

Filtrations and towers. Let $I \subseteq \mathbb{R}$ be a set of real values which we refer to as *scales*. A *filtration* is a collection of simplicial complexes $(K_\alpha)_{\alpha \in I}$ such that $K_\alpha \subseteq K'_\alpha$ for all $\alpha \leq \alpha' \in I$. For instance, $(\mathcal{R}_\alpha)_{\alpha \geq 0}$ is a filtration which we call the *Rips filtration*. A *(simplicial) tower* is a sequence $(K_\alpha)_{\alpha \in J}$ of simplicial complexes with J being a discrete set (for instance $J = \{2^k \mid k \in \mathbb{Z}\}$), together with simplicial maps $\varphi_\alpha : K_\alpha \to K_{\alpha'}$ between complexes at consecutive scales. For instance, the Rips filtration can be turned into a tower by restricting to a discrete range of scales, and using the inclusion maps as φ. The approximation constructed in this paper will be another example of a tower.

We say that a simplex σ is *included* in the tower at scale α', if σ is not in the image of the map $\varphi_\alpha : K_\alpha \to K_{\alpha'}$, where α is the scale preceding α' in the tower. The *size* of a tower is the number of simplices included over all scales. If a tower arises from a filtration, its size is simply the size of the largest complex in the filtration (or infinite, if no such complex exists). However, this is not true in general for simplicial towers, because simplices can collapse in the tower and the size of the complex at a given scale may not take into account the collapsed simplices which were included at earlier scales in the tower.

Barcodes and Interleavings. A collection of vector spaces $(V_\alpha)_{\alpha \in I}$ connected with linear maps $\lambda_{\alpha_1,\alpha_2} : V_{\alpha_1} \to V_{\alpha_2}$ is called a *persistence module*, if $\lambda_{\alpha,\alpha}$ is the identity on V_α and $\lambda_{\alpha_2,\alpha_3} \circ \lambda_{\alpha_1,\alpha_2} = \lambda_{\alpha_1,\alpha_3}$ for all $\alpha_1 \leq \alpha_2 \leq \alpha_3 \in I$ for the index set I.

We generate persistence modules using the previous concepts. Given a simplicial tower $(K_\alpha)_{\alpha \in I}$, we generate a sequence of chain complexes $(\mathcal{C}_*(K_\alpha))_{\alpha \in I}$. By functoriality, the simplicial maps φ of the tower give rise to chain maps $\overline{\varphi}$ between these chain complexes. Using functoriality of homology, we obtain a sequence $(H(K_\alpha))_{\alpha \in I}$ of vector spaces with linear maps $\overline{\varphi}^*$, forming a persistence module. The same construction applies to filtrations as a special case.

Persistence modules admit a decomposition into a collection of intervals of the form $[\alpha, \beta]$ (with $\alpha, \beta \in I$), called the *barcode*, subject to certain tameness conditions. The barcode of a persistence module characterizes the module uniquely up to isomorphism. If the persistence module is generated by a simplicial complex, an interval $[\alpha, \beta]$ in the barcode corresponds to a homological feature (a "hole") that comes into existence at complex K_α and persists until it disappears at K_β.

Two persistence modules $(V_\alpha)_{\alpha \in I}$ and $(W_\alpha)_{\alpha \in I}$ with linear maps $\phi_{\cdot,\cdot}$ and $\psi_{\cdot,\cdot}$ are said to be *weakly (multiplicatively) c-interleaved* with $c \geq 1$, if there exist linear maps $\gamma_\alpha : V_\alpha \to W_{c\alpha}$ and $\delta_\alpha : W_\alpha \to V_{c\alpha}$, called *interleaving maps*, such that the diagram

$$\cdots \longrightarrow V_{\alpha c} \xrightarrow{\phi} V_{\alpha c^3} \longrightarrow \cdots \qquad (1)$$

$$\cdots \longrightarrow W_\alpha \xrightarrow{\psi} W_{\alpha c^2} \longrightarrow \cdots$$

commutes, that is, $\psi = \gamma \circ \delta$ and $\phi = \delta \circ \gamma$ for all $\{\ldots, \alpha/c^2, \alpha/c, \alpha, c\alpha, \ldots\} \in I$ (we have skipped the subscripts of the maps for readability). In such a case, the barcodes of the two modules are $3c$-approximations of each other in the sense of Chazal et al.

(2009). We say that two towers are *c-approximations* of each other if their persistence modules are *c*-approximations.

Under the more stringent conditions of *strong interleaving*, the approximation ratio can be improved. Two persistence modules $(V_\alpha)_{\alpha \geq 0}$ and $(W_\alpha)_{\alpha \geq 0}$ with respective linear maps $\phi_{\cdot,\cdot}$ and $\psi_{\cdot,\cdot}$ are said to be *(multiplicatively) strongly c-interleaved* if there exist a pair of families of linear maps $\gamma_\alpha : V_\alpha \to W_{c\alpha}$ and $\delta_\alpha : W_\alpha \to V_{c\alpha}$ for $c > 0$, such that Diagram (2) commutes for all $0 \leq \alpha \leq \alpha'$ (the subscripts of the maps are excluded for readability). In such a case, the persistence barcodes of the two modules are said to be *c*-approximations of each other in the sense of Chazal et al. (2009).

$$\begin{array}{cc}
\begin{array}{c}
V_{\frac{\alpha}{c}} \xrightarrow{\ \ \phi\ \ } V_{c\alpha'} \\
\gamma \searrow \quad \nearrow \delta \\
W_\alpha \xrightarrow{\psi} W_{\alpha'}
\end{array}
&
\begin{array}{c}
V_{c\alpha} \xrightarrow{\phi} V_{c\alpha'} \\
\delta \nearrow \quad \nearrow \delta \\
W_\alpha \xrightarrow{\psi} W_{\alpha'}
\end{array}
\end{array} \qquad (2)$$

$$\begin{array}{cc}
\begin{array}{c}
V_\alpha \xrightarrow{\phi} V_{\alpha'} \\
\delta \nearrow \quad \searrow \gamma \\
W_{\frac{\alpha}{c}} \xrightarrow{\ \ \psi\ \ } W_{c\alpha'}
\end{array}
&
\begin{array}{c}
V_\alpha \xrightarrow{\phi} V_{\alpha'} \\
\gamma \searrow \quad \searrow \gamma \\
W_{c\alpha} \xrightarrow{\psi} W_{c\alpha'}
\end{array}
\end{array}$$

Finally, we mention a special case that relates equivalent persistence modules (Carlsson and Zomorodian 2005; Goodman et al. 2017). Two persistence modules $\mathbb{V} = (V_\alpha)_{\alpha \in I}$ and $\mathbb{W} = (W_\alpha)_{\alpha \in I}$ that are connected through linear maps ϕ, ψ respectively are isomorphic if there exists an isomorphism $f_\alpha : V_\alpha \to W_\alpha$ for each $\alpha \in I$ for which the following diagram commutes for all $\alpha \leq \beta \in I$:

$$\begin{array}{ccccccc}
\cdots & \longrightarrow & V_\alpha & \xrightarrow{\ \phi\ } & V_\beta & \longrightarrow & \cdots \\
& & \downarrow{f_\alpha} & & \downarrow{f_\beta} & & \\
\cdots & \longrightarrow & W_\alpha & \xrightarrow{\ \psi\ } & W_\beta & \longrightarrow & \cdots
\end{array} \qquad (3)$$

Isomorphic persistence modules have identical persistence barcodes.

Scale balancing. Let $\mathbb{V} = (V_\alpha)_{\alpha \in I}$ and $\mathbb{W} = (W_\alpha)_{\alpha \in I}$ be two persistence modules with linear maps f_v, f_w, respectively. Let there be linear maps $\phi : V_{\alpha/\varepsilon_1} \to W_\alpha$ and $\psi : W_\alpha \to V_{\alpha\varepsilon_2}$ for $1 \leq \varepsilon_1, \varepsilon_2$ such that all $\alpha, \alpha/\varepsilon_1, \alpha\varepsilon_2 \in I$. Suppose that the following diagram commutes, for all $\alpha \in I$.

$$\begin{array}{ccccccc}
\cdots & \longrightarrow & W_\alpha & \xrightarrow{\quad f_w \quad} & W_{\alpha\varepsilon_1\varepsilon_2} & \longrightarrow & \cdots \\
& \phi \nearrow & & \searrow \psi & & \nearrow \psi & \\
\cdots \longrightarrow & V_{\alpha/\varepsilon_1} & & \xrightarrow{\ f_v\ } & V_{\alpha\varepsilon_2} & \longrightarrow & \cdots
\end{array} \qquad (4)$$

Let $\varepsilon := max(\varepsilon_1, \varepsilon_2)$. Then, by replacing ε_1, ε_2 by ε in Diagram (4), the diagram still commutes, so $\mathbb{V}$ is a 3ε-approximation of $\mathbb{W}$.

We define a new vector space $V'_{c\alpha} := V_\alpha$, where $c = \sqrt{\frac{\varepsilon_1}{\varepsilon_2}}$ and $c\alpha \in I$. This gives rise to a new persistence module, $\mathbb{V}' = (V_{c\alpha})_{\alpha \in I}$. The maps ϕ and ψ can then be interpreted as $\phi : V'_{\alpha / \sqrt{\varepsilon_1 \varepsilon_2}} \to W_\alpha$, or $\phi : V'_\alpha \to W_{\alpha\sqrt{\varepsilon_1\varepsilon_2}}$ and $\psi : W_\alpha \to V'_{\alpha\sqrt{\varepsilon_1\varepsilon_2}}$. Then, Diagram (4) can be re-interpreted as

$$\cdots \longrightarrow W_{\alpha\sqrt{\varepsilon_1\varepsilon_2}} \xrightarrow{\quad f_w \quad} W_{\alpha(\varepsilon_1\varepsilon_2)^{3/2}} \longrightarrow \cdots \tag{5}$$

$$\cdots \longrightarrow V_{\alpha'} \xrightarrow{\quad f_v \quad} V'_{\alpha\varepsilon_1\varepsilon_2} \longrightarrow \cdots$$

which still commutes. Therefore, $\mathbb{V}'$ is a $3\sqrt{\varepsilon_1\varepsilon_2}$-approximation of $\mathbb{W}$, which is an improvement over $\mathbb{V}$, since $\sqrt{\varepsilon_1\varepsilon_2} \le max(\varepsilon_1, \varepsilon_2)$. $\mathbb{V}$ and $\mathbb{V}'$ have the same barcode up to a scaling factor.

This scaling trick also works when $\mathbb{V}$ and $\mathbb{W}$ are strongly interleaved. If we have the following commutative diagrams: (where we have skipped the maps for readability):

$$\begin{array}{cc}
W_\alpha \longrightarrow W_{\alpha'\varepsilon_1\varepsilon_2} & W_{\alpha\varepsilon_1} \to W_{\alpha'\varepsilon_1} \qquad (6) \\
V_{\alpha\varepsilon_2} \to V_{\alpha'\varepsilon_2} & V_\alpha \longrightarrow V_{\alpha'}
\end{array}$$

$$\begin{array}{cc}
W_{\alpha\varepsilon_1} \to W_{\alpha'\varepsilon_1} & W_\alpha \to W_{\alpha'} \\
V_\alpha \longrightarrow V_{\alpha'\varepsilon_1\varepsilon_2} & V_{\alpha\varepsilon_2} \to V_{\alpha'\varepsilon_2}
\end{array}$$

then $\mathbb{V}$ and $\mathbb{W}$ are $max(\varepsilon_1, \varepsilon_2)$-approximations of each other. By defining $\mathbb{V}'$ as before, the following diagrams

$$\begin{array}{cc}
W_\alpha \longrightarrow W_{\alpha d^2} & W_{\alpha d} \longrightarrow W_{\alpha' d} \quad (7) \\
V'_{\alpha d} \longrightarrow V'_{\alpha' d} & V'_\alpha \longrightarrow V'_{\alpha'}
\end{array}$$

$$\begin{array}{cc}
W_{\alpha d} \longrightarrow W_{\alpha' d} & W_\alpha \longrightarrow W_{\alpha'} \\
V'_\alpha \longrightarrow V'_{\alpha' d^2} & V'_{\alpha d} \longrightarrow V'_{\alpha' d}
\end{array}$$

commute for $d = c\varepsilon_2 = \sqrt{\varepsilon_1 \varepsilon_2}$, so we can improve a $\max(\varepsilon_1, \varepsilon_2)$-approximation to an $\sqrt{\varepsilon_1 \varepsilon_2}$-approximation.

We end the section by discussing a basic but important relation between Čech and Rips filtrations. It is well-known that for any $\alpha \geq 0, \mathcal{C}_\alpha \subseteq \mathcal{R}_\alpha \subseteq \mathcal{C}_{\sqrt{2}\alpha}$ (Edelsbrunner and Harer 2010). This gives a strong interleaving between the towers $(\mathcal{C}_\alpha)_{\alpha \geq 0}$ and $(\mathcal{R}_\alpha)_{\alpha \geq 0}$ with $\varepsilon_1 = 1$ and $\varepsilon_2 = \sqrt{2}$. Applying the scale balancing technique, we get that

Lemma 1 *The scaled Čech persistence module $(H(\mathcal{C}_{\sqrt[4]{2}\alpha}))_{\alpha \geq 0}$ and the Rips persistence module $(H(\mathcal{R}_\alpha))_{\alpha \geq 0}$ are $\sqrt[4]{2}$-approximations of each other.*

3 Shifted integer lattices

In this section, we take a look at simple modifications of the integer lattice.

We denote by $I := \{\alpha_s := \lambda 2^s \mid s \in \mathbb{Z}\}$ with $\lambda > 0$, a discrete set of scales. For each scale in I, we define grids which are scaled and translated (shifted) versions of the integer lattice.

Definition 1 (Scaled and shifted grids) For each scale $\alpha_s \in I$, we define the *scaled and shifted grid* G_{α_s} inductively as:

- For $s = 0$, G_{α_s} is simply the scaled integer grid $\lambda \mathbb{Z}^d$, where each basis vector has been scaled by λ.
- For $s \geq 0$, we choose an arbitrary point $O_{\alpha_s} \in G_{\alpha_s}$ and define

$$G_{\alpha_{s+1}} = 2\left(G_{\alpha_s} - O_{\alpha_s}\right) + O_{\alpha_s} + \frac{\alpha_s}{2}\left(\pm 1, \cdots, \pm 1\right), \tag{8}$$

  where the signs of the components of the last vector are chosen independently and uniformly at random (and the choice is independent for each s).
- For $s \leq 0$, we define

$$G_{\alpha_{s-1}} = \frac{1}{2}\left(G_{\alpha_s} - O_{\alpha_s}\right) + O_{\alpha_s} + \frac{\alpha_{s-1}}{2}\left(\pm 1, \cdots, \pm 1\right), \tag{9}$$

 where the last vector is chosen as in the case of $s \geq 0$.

Equations (8) and (9) are consistent at $s = 0$. A simple example of the above construction is the sequence of grids with $G_{\alpha_s} := \alpha_s \mathbb{Z}^d$ for even s, and $G_{\alpha_s} := \alpha_s \mathbb{Z}^d + \frac{\alpha_{s-1}}{2}(1, \cdots, 1)$ for odd s.

Next, we motivate the shifting of the grids. Let $\text{Vor}_{G_s}(x)$ denote the Voronoi cell of any point $x \in G_s$ with respect to the point set G_s. It is clear that the Voronoi cell is a cube of side length α_s centered at x. The shifting of the grids ensures that each $x \in G_{\alpha_s}$ lies in the Voronoi region of a unique $y \in G_{\alpha_{s+1}}$. Using an elementary calculation, we show a stronger statement:

Lemma 2 *Let $x \in G_{\alpha_s}, y \in G_{\alpha_{s+1}}$ be such that $x \in \text{Vor}_{G_{\alpha_{s+1}}}(y)$. Then,*

$$\text{Vor}_{G_{\alpha_s}}(x) \subset \text{Vor}_{G_{\alpha_{s+1}}}(y).$$

Fig. 1 G_{α_s} is represented by small disks (yellow), while $G_{\alpha_{s+1}}$ is represented by larger disks (green). Possible locations of x are indicated with their Voronoi regions. The Voronoi regions of the larger grid contain those of x
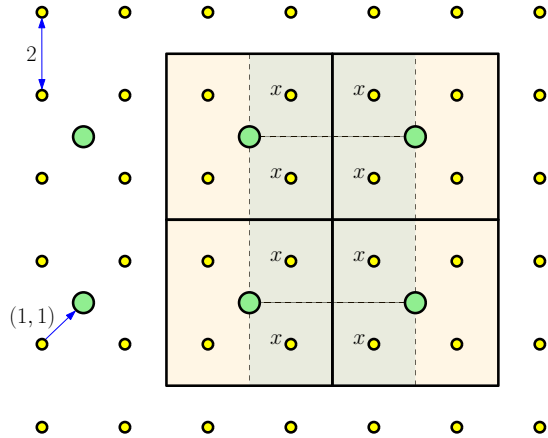
Proof Without loss of generality, we can assume that $\alpha_s = 2$ and x is the origin, using an appropriate translation and scaling. Also, we assume for the sake of simplicity that $G_{\alpha_{s+1}} = 2G_{\alpha_s} + (1, \cdots, 1)$; the proof is analogous for any other translation vector. In that case, it is clear that $y = (1, \cdots, 1)$. Since $G_{\alpha_s} = 2\mathbb{Z}^d$, the Voronoi region of x is the set $[-1, 1]^d$. Since $G_{\alpha_{s+1}}$ is a translated version of $4\mathbb{Z}^d$, the Voronoi region of y is the cube $[-1, 3]^d$, which covers $[-1, 1]^d$. The claim follows. For an example look to Fig. 1. $\qquad\square$

3.1 Cubical complex of $\mathbb{Z}^d$

The integer grid $\mathbb{Z}^d$ naturally defines a cubical complex, where each element is an axis-aligned, k-dimensional cube with $0 \leq k \leq d$. To define it formally, let $\square$ denote the set of all integer translates of faces of the unit cube $[0, 1]^d$, considered as a convex polytope in $\mathbb{R}^d$. We call the elements of $\square$ *faces of* $\mathbb{Z}^d$.

Each face has a dimension k; the 0-faces, or *vertices* are exactly the points in $\mathbb{Z}^d$. The *facets* of a k-face E are the $(k-1)$-faces contained in E. We call a pair of facets of E *opposite facets*, if they are disjoint. Naturally, these concepts carry over to scaled and shifted versions of $\mathbb{Z}^d$, so we define $\square_{\alpha_s}$ as the cubical complex defined by G_{α_s}.

We define a map $g_{\alpha_s} : \square_{\alpha_s} \to \square_{\alpha_{s+1}}$ as follows: for vertices of $\square_{\alpha_s}$, we assign to $x \in G_{\alpha_s}$ the (unique) vertex $y \in G_{\alpha_{s+1}}$ such that $x \in \text{Vor}_{G_{\alpha_{s+1}}}(y)$ (see Lemma 2). For a k-face f of $\square_{\alpha_s}$ with vertices $(p_1, \cdots, p_{2^k})$ in G_{α_s}, we set $g_{\alpha_s}(f)$ to be the convex hull of $\{g_{\alpha_s}(p_1), \cdots, g_{\alpha_s}(p_{2^k})\}$; the next lemma shows that this is a well-defined map. In this paper, we sometimes call g_{α_s} a *cubical map*, since it is a counterpart of simplicial maps for cubical complexes.

Lemma 3 *Let f be k-face of $\square_{\alpha_s}$ with vertices $\{p_1, \cdots, p_{2^k}\} \subset G_{\alpha_s}$. Then*

– *the set of vertices $\{g_{\alpha_s}(p_1), \cdots, g_{\alpha_s}(p_{2^k})\}$ form a face e of $\square_{\alpha_{s+1}}$.*
– *for every face $e_1 \subset e$, there is a face $f_1 \subset f$ such that $g_{\alpha_s}(f_1) = e_1$.*
– *if e_1, e_2 are any two opposite facets of e, then there exists a pair of opposite facets f_1, f_2 of f such that $g_{\alpha_s}(f_1) = e_1$ and $g_{\alpha_s}(f_2) = e_2$.*

Proof **First claim:** We prove the first claim by induction on the dimension of faces of G_{α_s}. Base case: for vertices, the claim is trivial using Lemma 2. Induction case: let the claim hold true for all $(k-1)$-faces of G_{α_s}. We show that the claim holds true for all k-faces of G_{α_s}.

Let f be a k-face of G_{α_s}. Let f_1 and f_2 be opposite facets of f, along the m-th coordinate. Let us denote the vertices of f_1 by $(p_1, \cdots, p_{2^{k-1}})$ and those of f_2 by $(p_{2^{k-1}+1}, \cdots, p_{2^k})$ taken in the same order, that is, p_j and $p_{2^{k-1}+j}$ differ in only the m-th coordinate for all $1 \leq j \leq 2^{k-1}$. By definition, all vertices of f_1 share the m-th coordinate, and we denote coordinate of these vertices by z. Then, the m-th coordinate of all vertices of f_2 equals $z + \alpha_s$. Then $g_{\alpha_s}(p_j)$ and $g_{\alpha_s}(p_{2^{k-1}+j})$ have the same coordinates, except possibly the m-th coordinate. By induction hypothesis, $e_1 = g_{\alpha_s}(f_1)$ and $e_2 = g_{\alpha_s}(f_2)$ are two faces of G_{s+1}. This implies that e_2 is a translate of e_1 along the m-th coordinate.

There are two cases: if e_1 and e_2 share the m-th coordinate, then $e_1 = e_2$ and therefore $g_{\alpha_s}(f) = e_1 = e_2 = e$, so the claim follows. On the other hand, if e_1 and e_2 do not share the m-th coordinate, then they are two faces of $\square_{\alpha_{s+1}}$ which differ in only one coordinate by α_{s+1}. So they are opposite facets of a co-dimension one face e of $G_{\alpha_{s+1}}$. Using induction, the claim follows.

Second claim: We prove the claim by induction over the dimension of e_1. Base case: e_1 is a vertex. The vertices of f in Voronoi region of e_1 form f_1. Since f is an axis parallel face and the Voronoi region is also axis-parallel, it is immediate that f_1 is a face of f. Assume that the claim is true up to dimension i. For e_1 a face of dimension $i + 1$, consider opposite facets e_a and e_b of e. By the induction claim, there exist faces $f_a, f_b \subset f$ that satisfy $g_{\alpha_s}(f_a) = e_a$, $g_{\alpha_s}(f_b) = e_b$. f_a and f_b are disjoint since otherwise $g_{\alpha_s}(f_a \cap f_b)$ would be common to both e_a and e_b, a contradiction. If e_a is a translate of e_b along the m-th coordinate, then f_a is also a translate of f_b along the same coordinate. Therefore f_a and f_b are opposite faces of a face f_1 and $g_{\alpha_s}(f_1) = e_1$.

Third claim: Without loss of generality, assume that x_1 is the direction in which e_2 is a translate of e_1. Using the second claim, let h denote the maximal face of f such that $g_{\alpha_s}(h) = e_1$. Clearly, $h \neq f$, since that would imply $g_{\alpha_s}(f) = e_1 = e$, which is a contradiction.

Suppose h has dimension less than $k - 1$. Let h' be the facet of f that contains h and has the same x_1 coordinates for all vertices. Then $g_{\alpha_s}(h') = e_1$, which contradicts the maximality of h.

Therefore, the only possibility is that h is a facet f_1 of f such that $g_{\alpha_s}(f_1) = e_1$. Let f_2 be the opposite facet of f_1. From the proof of the first claim, it is easy to see that $g_{\alpha_s}(f_2) = e_2$. The claim follows. $\square$
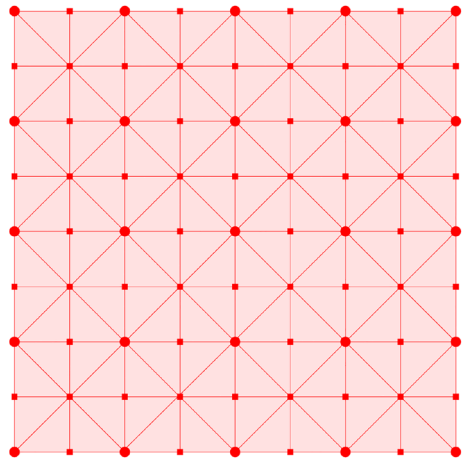
3.2 Barycentric subdivision

We discuss a special triangulation of $\square_{\alpha_s}$. A *flag* in $\square_{\alpha_s}$ is a set of faces $\{f_0, \cdots, f_k\}$ of $\square_{\alpha_s}$ such that

$$f_0 \subseteq \cdots \subseteq f_k.$$

Fig. 2 A portion of the grid in two dimensions. The dots are the grid points which form the 0-faces of the cubical complex
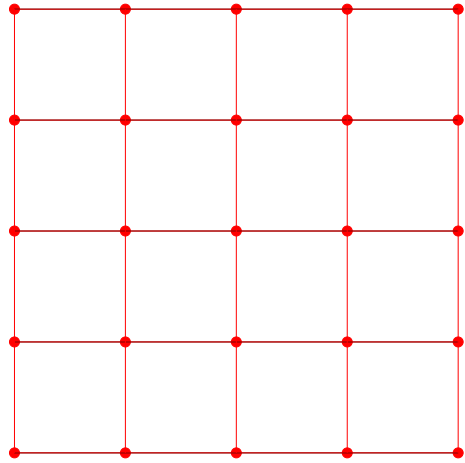
Fig. 3 The barycentric subdivision of the grid. The tiny squares are barycenters of the 1-faces and 2-faces of the cubical complex

The *barycentric subdivision* of $\square_{\alpha_s}$, denoted by sd_{α_s}, is the (infinite) simplicial complex whose simplices are the flags of $\square_{\alpha_s}$ (Munkres 1984).

In particular, the 0-simplices of sd_{α_s} are the faces of $\square_{\alpha_s}$. An equivalent geometric description of sd_{α_s} can be obtained by defining the 0-simplices as the barycenters of the faces in sd_{α_s}, and introducing a k-simplex between $(k+1)$ barycenters if the corresponding faces form a flag. For a simple example, see Figs. 2 and 3. It is easy to see that sd_{α_s} is a flag complex. Given a face f in $\square_{\alpha_s}$, we write $sd(f)$ for the subcomplex of sd_{α_s} consisting of all flags that are formed only by faces contained in f.

4 Approximation scheme with simplicial complexes

We define our approximation complex for a finite set of points in $\mathbb{R}^d$. Recall from Definition 1 that we can define a collection of scaled and shifted integer grids G_{α_s} over a collection of scales $I := \{\alpha_s = 2^s \mid s \in \mathbb{Z}\}$ in $\mathbb{R}^d$. To make the exposition simple, we define our complex in a slightly generalized form.

4.1 Barycentric spans

Fix some $s \in \mathbb{Z}$ and let V denote any non-empty subset of G_{α_s}.
Vertex span. We say that a face $f \in \square_{\alpha_s}$ is *spanned* by V, if the set of vertices $V(f) := f \cap V$

- is non-empty, and
- not contained in any facet of f.

Trivially, the vertices of $\square_{\alpha_s}$ which are spanned by V are precisely the points in V. Any face of $\square_{\alpha_s}$ which is not a vertex must contain at least two vertices of V in order to be spanned. We point out that the set of spanned faces of $\square_{\alpha_s}$ is *not* closed under taking sub-faces. For instance, if V consists of two antipodal points of a d-cube, the only faces spanned by V are the d-cube and the two vertices; all other faces of the d-cube contain at most one vertex and hence are not spanned.

It is simple to test whether any given k-face $f \in \square_{\alpha_s}$ is spanned by the set of points $V(f)$. Let $T \subseteq [1, \cdots, d]$ be the set of common coordinates of the points in $V(f)$. $V(f)$ spans f if and only if the standard basis vectors of $\mathbb{R}^d$ corresponding to T span f. T can be computed in $|V(f)|O(d) = O(2^k d)$ time by a linear scan of the coordinates. The coordinate directions spanned by f can also be found and compared with T within the same time bound.
Barycentric span. The *barycentric span* of V is the subcomplex of sd_{α_s} obtained by taking the union of the complete barycentric subdivisions of the maximal faces of $\square_{\alpha_s}$ that are spanned by V. The barycentric span of V is indeed a simplicial complex by definition. Moreover, the barycentric span is a flag complex. Then for any face $f \in \square_{\alpha_s}$, the barycentric span of $V(f)$ is either empty or acyclic.

Furthermore, for any non-empty subset $W \subseteq V$, the faces of $\square_{\alpha_s}$ that are spanned by W are also spanned by V. Consequently, the barycentric span of W is a subcomplex of the barycentric span of V.

4.2 Approximation complex

We denote by $P \subset \mathbb{R}^d$ a finite set of points. We define two maps:

- $a_{\alpha_s} : P \to G_{\alpha_s}$: for each point $p \in P$, we let $a_{\alpha_s}(p)$ denote the grid point in G_{α_s} that is closest to p, that is, $p \in \mathrm{Vor}_{G_{\alpha_s}}(a_{\alpha_s}(p))$. We assume for simplicity that this closest point is unique, which can be ensured using well-known methods (Edelsbrunner and Mücke 1990). We define the *active vertices of* G_{α_s} as

$$V_{\alpha_s} := \mathrm{im}\left(a_{\alpha_s}\right) = a_{\alpha_s}(P) \subset G_{\alpha_s},$$

that is, the set of grid points that have at least one point of P in their Voronoi cells.

– $b_{\alpha_s} : V_{\alpha_s} \to P$: the map b_{α_s} takes an active vertex of G_{α_s} to its closest point in P. By taking an arbitrary total order on P to resolve multiple assignments, we ensure that this assignment is unique.

Naturally, $b_{\alpha_s}(v)$ is a point inside $\mathrm{Vor}_{G_{\alpha_s}}(v)$ for any $v \in V_{\alpha_s}$. It follows that the map b_{α_s} is a section of a_{α_s}, that is, $a_{\alpha_s} \circ b_{\alpha_s} : V_{\alpha_s} \to V_{\alpha_s}$ is the identity on V_{α_s}. However, this is not true for $b_{\alpha_s} \circ a_{\alpha_s}$ in general.

Recall that the map $g_{\alpha_s} : \square_{\alpha_s} \to \square_{\alpha_{s+1}}$ takes grid points of G_{α_s} to grid points of $G_{\alpha_{s+1}}$. Using Lemma 2, it follows at once that:

Lemma 4 *For all $\alpha_s \in I$ and each $x \in V_{\alpha_s}$, $g_{\alpha_s}(x) = (a_{\alpha_{s+1}} \circ b_{\alpha_s})(x)$.*

Recall that $\mathcal{R}_\alpha^\infty$ denotes the Rips complex at scale α for the L_∞-norm. The next statement is a direct application of the the triangle inequality; let $diam_\infty()$ denote the diameter in the L_∞-norm.

Lemma 5 *Let $Q \subseteq P$ be a non-empty subset such that $diam_\infty(Q) \le \alpha_s$. Then, the set of grid points $a_{\alpha_s}(Q)$ is contained in a face of $\square_{\alpha_s}$.*

Equivalently, for any simplex $\sigma = (p_0, \cdots, p_k) \in \mathcal{R}_{\alpha_s/2}^\infty$ on P, the set of active vertices $\{a_{\alpha_s}(p_0), \cdots, a_{\alpha_s}(p_k)\}$ is contained in a face of $\square_{\alpha_s}$.

Proof We prove the claim by contradiction. Suppose that the set of active vertices $a_{\alpha_s}(Q)$ is not contained in a face of $\square_{\alpha_s}$. Then, there exists at least one pair of points $\{x, y\} \in Q$ such that $a_{\alpha_s}(x), a_{\alpha_s}(y)$ are not in a common face of $\square_{\alpha_s}$. By the definition of the grid G_{α_s}, the grid points $a_{\alpha_s}(x), a_{\alpha_s}(y)$ therefore have L_∞-distance at least $2\alpha_s$. Moreover, x has L_∞-distance less than $\alpha_s/2$ from $a_{\alpha_s}(x)$, and the same is true for y and $a_{\alpha_s}(y)$. By the triangle inequality, the L_∞-distance of x and y is more than α_s, which is a contradiction to the fact that $diam_\infty(Q) \le \alpha_s$. $\qquad\square$

We now define our approximation tower. For any scale α_s, we define $\mathcal{X}_{\alpha_s}$ as the barycentric span of the active vertices $V_{\alpha_s} \subset G_{\alpha_s}$. See Figs. 4, 5 and 6 for a simple illustration.

To simplify notation, we denote

– the faces of $\square_{\alpha_s}$ spanned by V_{α_s} as *active faces*, and
– the faces of active faces that are not spanned by V_{α_s} as *secondary faces*.

To complete the description of the approximation tower, we need to define simplicial maps of the form $\tilde{g}_{\alpha_s} : \mathcal{X}_{\alpha_s} \to \mathcal{X}_{\alpha_{s+1}}$, which connect the simplicial complexes at consecutive scales. We show that such maps are induced by g_{α_s}.

Lemma 6 *Let f be any active face of $\square_{\alpha_s}$. Then, $g_{\alpha_s}(f)$ is an active face of $\square_{\alpha_{s+1}}$.*

Proof Using Lemma 3, $e := g_{\alpha_s}(f)$ is a face of $\square_{\alpha_s}$. If e is a vertex, then it is active, because f contains at least one active vertex v, and $g_{\alpha_s}(v) = e$ in this case. If e is not a vertex, we assume for a contradiction that it is not active. Then, it contains a facet e_1 that contains all active vertices in e. Let e_2 denote the opposite facet of e_1 in e. By Lemma 3, f contains opposite facets f_1, f_2 such that $g_{\alpha_s}(f_1) = e_1$ and $g_{\alpha_s}(f_2) = e_2$. Since f is active, both f_1 and f_2 contain active vertices; in particular, f_2 contains an active vertex v. But then the active vertex $g_{\alpha_s}(v)$ must lie in e_2, contradicting the fact that e_1 contains all active vertices of e. $\qquad\square$
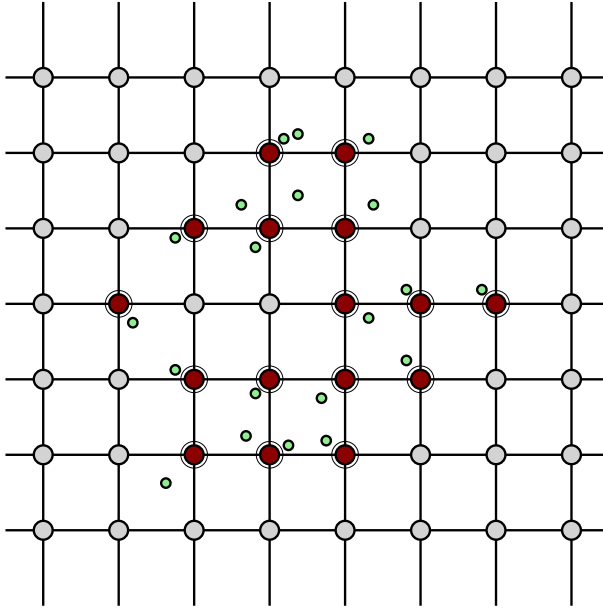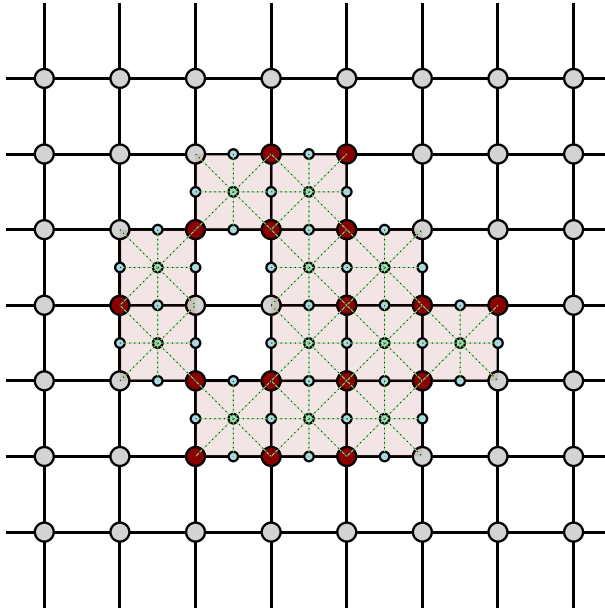
Fig. 4 A two-dimensional grid, shown along with its cubical complex. The green points (small dots) denote the points in P and the red vertices (encircled) are the active vertices (color figure online)
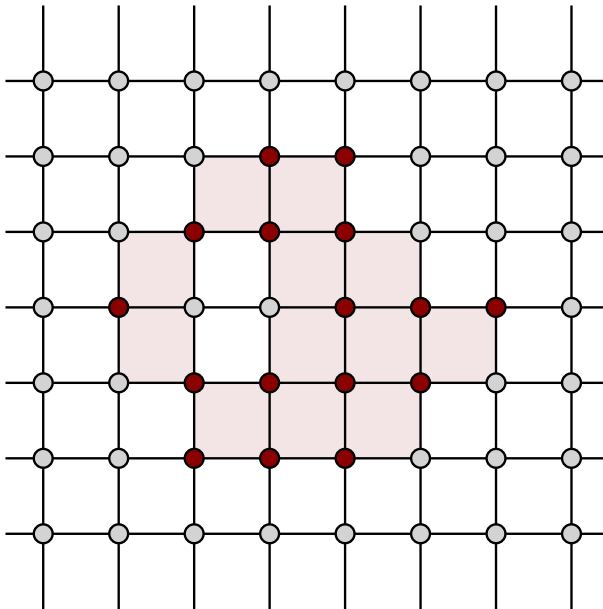
Fig. 5 The active faces are shaded. The closure of the active faces forms the cubical complex

Fig. 6 The generated approximation complex, whose vertices consist of those of the cubical complex and the blue vertices (small dots), which are the barycenters of active and secondary faces

As a result, g is well defined for each face $e \in \square_{\alpha_s}$, since there exists some active face $e' \in \square_{\alpha_s}$ with $e \subseteq e'$, and $g(e) \subseteq g(e')$. By definition, a simplex $\sigma \in \mathcal{X}_{\alpha_s}$ is a flag $(f_0 \subseteq \cdots \subseteq f_k)$ of faces in $\square_{\alpha_s}$. We set

$$\tilde{g}_{\alpha_s}(\sigma) := \left(g_{\alpha_s}(f_0), \cdots, g_{\alpha_s}(f_k) \right),$$

where $(g_{\alpha_s}(f_0) \subseteq \cdots \subseteq g_{\alpha_s}(f_k))$ is a flag of faces in $\square_{\alpha_{s+1}}$ by Lemma 6, and hence is a simplex in $\mathcal{X}_{\alpha_{s+1}}$. It follows that $\tilde{g}_s : \mathcal{X}_{\alpha_s} \to \mathcal{X}_{\alpha_{s+1}}$ is a simplicial map. This completes the description of the simplicial tower

$$\left(\mathcal{X}_{\alpha_s} \right)_{s \in \mathbb{Z}}.$$

4.3 Interleaving with the Rips module

First, we show that our tower is a constant-factor approximation of the the L_∞-Rips filtration of P. We then show the relation between our approximation tower and the Euclidean Rips filtration of P.

We start by defining two acyclic carriers. First, we set $\lambda = 1$ and abbreviate $\alpha := \alpha_s = 2^s$ to simplify notation.

– $C_1^\alpha : \mathcal{R}_{\alpha/2}^\infty \to \mathcal{X}_\alpha$: for any simplex $\sigma = (p_0, \cdots, p_k)$ in $\mathcal{R}_{\alpha/2}^\infty$, we set $C_1^\alpha(\sigma)$ as the barycentric span of $U := \{a_s(p_0), \cdots, a_s(p_k)\}$, which is a subcomplex of $\mathcal{X}_\alpha$. Using Lemma 5, U lies in a maximal active face f of $\square_\alpha$, so that $C_1^\alpha(\sigma)$ is acyclic.

The barycentric span of any subset of U is a subcomplex of the barycentric span of U, so C_1^α is a carrier. Therefore, C_1^α is an acyclic carrier.

– $C_2^\alpha : \mathcal{X}_\alpha \to \mathcal{R}_\alpha^\infty$: let σ be any flag of $\mathcal{X}_\alpha$ and let E be the smallest active face of $\Box_\alpha$ that contains σ (we break ties by making use of an arbitrary global order $\succ$ on P)[5]. We collect all the points of P that map to vertices of E under the map a_α and set $C_2^\alpha(\sigma)$ as the simplex on this set of points. By an application of the triangle inequality, we see that the L_∞-diam of $C_2^\alpha(\sigma)$ is at most 2α, so $C_2^\alpha(\sigma) \in \mathcal{R}_\alpha^\infty$ and is acyclic. It is also clear that $C_2^\alpha(\tau) \subseteq C_2^\alpha(\sigma)$ for each $\tau \subseteq \sigma$, so C_2^α is an acyclic carrier.

Using the acyclic carrier theorem (Theorem 1), there exist augmentation-preserving chain maps

$$c_1^\alpha : \mathcal{C}_* \left(\mathcal{R}_{\alpha/2}^\infty \right) \to \mathcal{C}_* (\mathcal{X}_\alpha) \quad \text{and} \quad c_2^\alpha : \mathcal{C}_* (\mathcal{X}_\alpha) \to \mathcal{C}_* \left(\mathcal{R}_\alpha^\infty \right),$$

between the chain complexes, which are carried by C_1^α and C_2^α respectively, for each $\alpha \in I$. We obtain the following diagram of augmentation-preserving chain maps:

$$
\begin{array}{ccccccc}
\cdots & \longrightarrow & \mathcal{C}_*(\mathcal{R}_\alpha^\infty) & \xrightarrow{\ inc\ } & \mathcal{C}_*(\mathcal{R}_{2\alpha}^\infty) & \longrightarrow & \cdots \\
& & \Big\uparrow{c_2} & \nwarrow{c_1} & \Big\uparrow{c_2} & & \\
\cdots & \longrightarrow & \mathcal{C}_*(\mathcal{X}_\alpha) & \xrightarrow{\ \tilde{g}\ } & \mathcal{C}_*(\mathcal{X}_{2\alpha}) & \longrightarrow & \cdots
\end{array}
\tag{10}
$$

where inc corresponds to the chain map for inclusion maps, and $\tilde{g}$ denotes the chain map for the corresponding simplicial map g (we removed indices of the maps for readability).

The chain complexes give rise to a diagram of the corresponding homology groups, connected by the induced linear maps $c_1^*, c_2^*, inc^*, \tilde{g}^*$:

$$
\begin{array}{ccccccc}
\cdots & \longrightarrow & H(\mathcal{R}_\alpha^\infty) & \xrightarrow{\ inc^*\ } & H(\mathcal{R}_{2\alpha}^\infty) & \longrightarrow & \cdots \\
& & \Big\uparrow{c_2^*} & \nwarrow{c_1^*} & \Big\uparrow{c_2^*} & & \\
\cdots & \longrightarrow & H(\mathcal{X}_\alpha) & \xrightarrow{\ \tilde{g}^*\ } & H(\mathcal{X}_{2\alpha}) & \longrightarrow & \cdots
\end{array}
\tag{11}
$$

Lemma 7 *For all $\alpha \in I$, the linear maps in the lower triangle of Diagram (11) commute, that is,*

$$\tilde{g}^* = c_1^* \circ c_2^*.$$

---

[5] We define an order between the active faces of $\Box_\alpha$, using $\succ$: for each active face $F \in \Box_\alpha$, there are at least two points of P whose images under g_α are vertices of F; say $\{q_1 \succ q_2 \succ \cdots \succ q_m\} \subseteq P$ are the points that map to F. We assign to F the string of length n: $q_1 q_2 \cdots q_m \overbrace{\emptyset, \cdots, \emptyset}^{n-m}$. Each active face has a unique string associated to it. A total order on the faces is obtained by taking the lexicographic orders of the strings of each active face.

Proof We look at the corresponding triangle in Diagram (10). We show that the (augmentation-preserving) chain maps $\tilde{g}$ and $c_1 \circ c_2$ are both carried by an acyclic carrier $D : \mathcal{X}_\alpha \to \mathcal{X}_{2\alpha}$. The claim then follows from the acyclic carrier theorem.

Let $\sigma \in \mathcal{X}_\alpha$ be any flag and let $E \in \Box_\alpha$ denote the minimal active face containing σ. Let $\{q_1, \ldots, q_k\}$ be the active vertices of E. Let $\{p_1, \ldots, p_m\}$ be the set of points of P that map to $\{q_1, \ldots, q_k\}$ under the map a_α. Since the L_∞-diameter of $\{p_1, \ldots, p_m\}$ is at most 2α, using Lemma 5 we see that $\{a_{2\alpha}(p_1), \ldots, a_{2\alpha}(p_m)\}$ is a face of $\Box_{2\alpha}$. We set $D(\sigma)$ as the barycentric span of $\{a_{2\alpha}(p_1), \ldots, a_{2\alpha}(p_m)\}$. It follows that D is an acyclic carrier.

Further, $\{a_{2\alpha}(p_1), \ldots, a_{2\alpha}(p_m)\} = \{g_{2\alpha}(q_1), \ldots, g_{2\alpha}(q_k)\}$ from Lemma 2, so $D(\sigma)$ is the barycentric subdivision of $g_{2\alpha}(E)$. As a result $D = C_1 \circ C_2$ so that it carries $c_1 \circ c_2$. We show that D also carries the map $\tilde{g}$.

By definition, for each face $e \subseteq E$, $g(e) \subseteq g(E)$ and $\tilde{g}(sd(e)) \subseteq \tilde{g}(sd(E))$. This means that $\tilde{g}(\sigma)$ is contained in $g(E)$. This shows that $\tilde{g}(\sigma) \in C_1 \circ C_2(\sigma)$ implying that $\tilde{g}$ is carried by $C_1 \circ C_2$, as required. □

Lemma 8 *For all $\alpha \in I$, the linear maps in the upper triangle of Diagram (11) commute, that is,*

$$inc^* = c_2^* \circ c_1^*.$$

Proof The proof technique is analogous to the proof of Lemma 7. We define an acyclic carrier $D : \mathcal{R}_\alpha^\infty \to \mathcal{R}_{2\alpha}^\infty$ which carries inc and $c_2 \circ c_1$, both of which are augmentation-preserving.

Let $\sigma = (p_0, \cdots, p_k) \in \mathcal{R}_\alpha^\infty$ be any simplex. The set of active vertices

$$U := \{a_{2\alpha}(p_0), \cdots, a_{2\alpha}(p_k)\} \subset G_{2\alpha}$$

lie in a face f of $G_{2\alpha}$, using Lemma 5. We can assume that f is active, as otherwise, we argue about a facet of f that contains U. We set $D(\sigma)$ as the simplex on the subset of points in P, whose closest grid point in $G_{2\alpha}$ is any vertex of f. Using the triangle inequality we see that $D(\sigma) \in \mathcal{R}_{2\alpha}^\infty$, so D is an acyclic carrier. The vertices of σ are a subset of $D(\sigma)$, so D carries the map inc. Showing that D carries $c_2 \circ c_1$ requires further explanation.

Let δ be any simplex in $\mathcal{X}_{2\alpha}$ for which the chain $c_1(\sigma)$ takes a non-zero value. Since $c_1(\sigma)$ is carried by $C_1(\sigma)$, we have that $\delta \in C_1(\sigma)$, which is the barycentric span of U. Furthermore, for any $\tau \in C_1(\sigma)$, $C_2(\tau)$ is a simplex on the set of vertices $\{p \in P \mid a_{2\alpha}(p) \in V(f)\}$. It follows that $C_2(\tau) \subseteq D(\sigma)$. In particular, since c_2 is carried by C_2, $c_2(c_1(\sigma)) \subseteq D(\sigma)$ as well. □

Using Lemmas 7 and 8, we see that the two persistence modules $\left(H(\mathcal{X}_{\alpha_s})\right)_{s \in \mathbb{Z}}$ and $\left(H(\mathcal{R}_\alpha^\infty)\right)_{\alpha \geq 0}$ are weakly 2-interleaved.

With elementary modifications in the definition of $\mathcal{X}$ and $\tilde{g}$, we can get a tower of the form $(\mathcal{X}_\alpha)_{\alpha \geq 0}$. Furthermore, with minor changes in the interleaving arguments, we show that the corresponding persistence module is strongly 4-interleaved with the L_∞-Rips module. Using scale balancing, this result improves to a strong 2-interleaving

(see Lemma 16). Since the techniques used in the proof are very similar to the concepts used in this section, for the sake of brevity we defer all further details to Appendix A.

Using the strong stability theorem for persistence modules and taking scale balancing into account, we immediately get that:

Theorem 2 *The scaled persistence module $\left(H(\mathcal{X}_{2\alpha})\right)_{\alpha \geq 0}$ and the L_∞-Rips persistence module $\left(H(\mathcal{R}_\alpha^\infty)\right)_{\alpha \geq 0}$ are 2-approximations of each other.*

For any pair of points $p, p' \in \mathbb{R}^d$, it holds that

$$\|p - p'\|_\infty \leq \|p - p'\|_2 \leq \sqrt{d}\, \|p - p'\|_\infty.$$

This in turn shows that the L_2- and the L_∞-Rips filtrations are strongly $\sqrt{d}$-interleaved. Using the scale balancing technique for strongly interleaved persistence modules, we get:

Lemma 9 *The scaled persistence module $(H(\mathcal{R}_{\alpha/d^{0.25}}))_{\alpha \geq 0}$ and $(H(\mathcal{R}_\alpha^\infty))_{\alpha \geq 0}$ are strongly $d^{0.25}$-interleaved.*

Using Theorem 2, Lemma 9 and the fact that interleavings satisfy the triangle inequality (Bubenik and Scott 2014, Theorem 3.3), we see that the module $(H(\mathcal{X}_{2\alpha}))_{\alpha \geq 0}$ is strongly $2d^{0.25}$-interleaved with the scaled Rips persistence module $(H(\mathcal{R}_{\alpha/d^{0.25}}))_{\alpha \geq 0}$. We can remove the scaling in the Rips filtration simply by multiplying the scales on both sides with $d^{0.25}$ and obtain our final approximation result:

Theorem 3 *The module $\left(H(\mathcal{X}_{2\sqrt[4]{d}\alpha})\right)_{\alpha \geq 0}$ and the Euclidean Rips persistence module $\left(H(\mathcal{R}_\alpha)\right)_{\alpha \geq 0}$ are $2d^{0.25}$-approximations of each other.*

5 Computational complexity

In this section, we discuss the computational aspects of constructing the approximation tower. In Sect. 5.1 we discuss the size complexity of the tower. An algorithm to compute the tower efficiently is presented in Sect. 5.2.

Range of relevant scales. Set $n := |P|$ and let $CP(P)$ denote the closest pair distance of P. At scale $\alpha_0 := \frac{CP(P)}{3d}$ and lower, no two active vertices lie in the same face of the grid, so the approximation complex consists of n isolated 0-simplices. At scale $\alpha_m := diam(P)$ and higher, points of P map to active vertices of a common face (by Lemma 5), so the generated complex is acyclic. We inspect the range of scales $[\alpha_0, \alpha_m]$ to construct the tower, since the barcode is explicitly known for scales outside this range. For this, we set $\lambda = \alpha_0$ in the definition of the scales. The total number of scales is

$$\lceil \log_2 \alpha_m/\alpha_0 \rceil = \left\lceil \log_2 \frac{diam(P)3d}{CP(P)} \right\rceil = \lceil \log_2 \Delta + \log_2 3d \rceil = O(\log \Delta + \log d),$$

where $\Delta = \frac{diam(P)}{CP(P)}$ is the spread of the point set.

5.1 Size of the tower

The size of a tower is the number of simplices that do not have a preimage, that is, the number of simplex inclusions in the tower. We start by counting the number of active faces used in the tower.

Lemma 10 *The number of active faces without pre-image in the tower is at most $n3^d$.*

Proof At scale α_0, there are n inclusions of 0-simplices in the tower, due to n active vertices. Using Lemma 2, g is surjective on the active vertices of $\square$ (for any scale). Hence, no further active vertices are added to the tower.

It remains to count the maximal active faces of dimension ≥ 1 without preimage. We will use a charging argument, charging the existence of such an active face to one of the points in P. We show that each point of P is charged at most $3^d - 1$ times, which proves the claim. For that, we first fix an arbitrary total order $\prec$ on P. Each active vertex on any scale has a non-empty subset of P in its Voronoi region; we call the maximal such point with respect to the order $\prec$ the *representative* of the active vertex.

For each active face f of dimension at least one, we define the *signature* of f as the set of representatives of the active vertices of f. If for any set of active vertices $u_1, \ldots, u_k$ we have that $v = g(u_1) = \cdots = g(u_k)$, then the representative of v is one of the representatives of $u_1, \ldots, u_k$, using Lemma 2. Therefore, the signatures of the active faces that are images of f under g are subsets of the signature of f. This implies that each maximal active face that is included has a unique maximal signature. We bound the number of maximal signatures to get a bound on the number of maximal active face inclusions. We charge the addition of each maximal signature to the lowest ordered point according to $\prec$.

Each signature contains representatives of active vertices from a face of $\square_\alpha$. Since each active vertex v has $3^d - 1$ neighboring vertices in the grid that lie in a common face, the representative p of v can be charged $3^d - 1$ times. There is a canonical isomorphism between the neighboring vertices of v at each scale. Then, for p to be charged more times, the image of v and some neighboring vertex u must be identical under g at some scale. But then, the representative of $g(v) = g(u)$ is not p anymore, since p was the lowest ranked point in its neighborhood, hence the representative changes when the Voronoi regions are combined. So, p could not have been charged in such a case. Therefore, each point $p \in P$ is indeed charged at most $3^d - 1$ times.

There are n active faces of dimension 0 and at most $n(3^d - 1)$ active faces of higher dimension. The upper bound is $n + n(3^d - 1) = n3^d$, as claimed. $\qquad\square$

Theorem 4 *The k-skeleton of the tower has size at most*

$$n6^{d-1}(2k+4)(k+3)! \begin{Bmatrix} d \\ k+2 \end{Bmatrix} = n2^{O(d \log k + d)},$$

where $\begin{Bmatrix} a \\ b \end{Bmatrix}$ denotes the Stirling number of the second kind.

Proof Each k-simplex that is included in the tower at any given scale α is a part of the barycentric subdivision of an active face that is also included at α. Therefore, we can account for the inclusion of this simplex by including the barycentric subdivision of its parent active face.

From Lemma 10 at most $n3^d$ active faces are included in the tower over all dimensions. We bound the number of k-simplices in the barycentric subdivision of a d-cube. Multiplying with $n3^d$ gives the required bound.

Let c be any d-cube of $\square_\alpha$. To count the number of flags of length $(m+1)$ contained in c that start with some vertex and end with c, we use similar ideas as in Edelsbrunner and Kerber (2012): first, we fix any vertex v of c and count the flags of the form $v \subseteq \cdots \subseteq c$. Every ℓ-face in c incident to v corresponds to a subset of ℓ coordinate indices, in the sense that the coordinates not chosen are fixed to the coordinates of v for the face. With this correspondence, a flag from v to c of length $(m+1)$ corresponds to an ordered m-partition of $\{1, \cdots, d\}$. The number of such partitions is known as $m!$ times the quantity $\left\{ \begin{matrix} d \\ m \end{matrix} \right\}$, which is the Stirling number of second kind (Rennie and Dobson 1969), and is upper bounded by $2^{O(d \log m)}$. Since c has 2^d vertices, the total number of flags $v \subseteq \cdots \subseteq c$ of length $(m+1)$ with any vertex v is hence $2^d m! \left\{ \begin{matrix} d \\ m \end{matrix} \right\}$.

We now count the number of flags of length $k+1$. Each such flag is $(k+1)$-subset of some flag of length $m = k + 3$ that start with a vertex and end with c. There are $2^d (k+2)! \left\{ \begin{matrix} d \\ k+2 \end{matrix} \right\}$ such flags and each of them has $\binom{k+3}{k+1} = (k+3)(k+2)/2$ subsets of size $(k+1)$. The number of $(k+1)$-flags is upper bounded by $2^d (k+2)! \left\{ \begin{matrix} d \\ k+2 \end{matrix} \right\} \frac{(k+3)(k+2)}{2} = 2^{d-1}(k+2)(k+3)! \left\{ \begin{matrix} d \\ k+2 \end{matrix} \right\}$. The k-skeleton has size at most

$$n3^d 2^{d-1}(k+2)(k+3)! \left\{ \begin{matrix} d \\ k+2 \end{matrix} \right\} = n6^{d-1}(2k+4)(k+3)! \left\{ \begin{matrix} d \\ k+2 \end{matrix} \right\}.$$

$\square$

5.2 Computing the tower

From Sect. 3, we know that $G_{\alpha_{s+1}}$ is built from G_{α_s} by making use of an arbitrary translation vector $(\pm 1, \cdots, \pm 1) \in \mathbb{Z}^d$. In our algorithm, we pick the components of this translation vector uniformly at random from $\{+1, -1\}$, and independently for each scale. The choice behind choosing this vector randomly becomes more clear in the next lemma.

From the definition, the cubical maps $g_{\alpha_s} : \square_{\alpha_s} \to \square_{\alpha_{s+1}}$ can be composed for multiple scales. For a fixed α_s, we denote by $g^{(j)} : \square_{\alpha_s} \to \square_{\alpha_{s+j}}$ the j-fold composition of g, that is,

$$g^{(j)} = g_{\alpha_{s+j-1}} \circ g_{\alpha_{s+j-2}} \circ \cdots \circ g_{\alpha_{s+1}} \circ g_{\alpha_s},$$

for $j \geq 1$.

Lemma 11 *For any k-face $f \in \square_{\alpha_s}$ with $1 \leq k \leq d$, let Y denote the minimal integer j such that $g^{(j)}(f)$ is a vertex, for a given choice of the randomly chosen translation vectors. Then, the expected value of Y satisfies*

$$\mathbb{E}[Y] \leq 3 \log k,$$

which implies that no face of $\square_{\alpha_s}$ survives more than $3 \log d$ scales in expectation.

Proof Without loss of generality, assume that the grid under consideration is $\mathbb{Z}^d$ and f is the k-face spanned by the vertices $\{\underbrace{\{0, 1\}, \cdots, \{0, 1\}}_{k}, 0, \cdots, 0\}$, so that the origin is a vertex of f. The proof for the general case is analogous.

Let $y_1 \in \{-1, 1\}$ denote the randomly chosen first coordinate of the translation vector, so that the corresponding shift is one of $\{-1/2, 1/2\}$.

- If $y_1 = 1$, then the grid G' on the next scale has some grid point with x_1-coordinate $1/2$. Clearly, the closest grid point in G' to the origin is of the form $(+1/2, \pm 1/2, \cdots, \pm 1/2)$, and thus, this point is also closest to $(1, 0, 0, \cdots, 0)$. The same is true for any point $(0, *, \cdots, *)$ and its corresponding point $(1, *, \cdots, *)$ on the opposite facet of f. Hence, for $y_1 = 1$, $g(f)$ is a face where all points have the same x_1-coordinate.
- On the other hand, if $y_1 = -1$, the origin is mapped to some point which has the form $(-1/2, \pm 1/2, \cdots, \pm 1/2)$ and $(1, 0, \cdots, 0)$ is mapped to $(3/2, \pm 1/2, \cdots, \pm 1/2)$, as one can directly verify. Hence, in this case, in $g(f)$, points do not all have the same x_1 coordinate.

We say that the x_1-coordinate *collapses* in the first case and *survives* in the second. Both events occur with the same probability $1/2$. Because the shift is chosen uniformly at random for each scale, the probability that x_1 did not collapse after j iterations is $1/2^j$.

f spans k coordinate directions, so it must collapse along each such direction to contract to a vertex. Once a coordinate collapses, it stays collapsed at all higher scales. As the random shift is independent for each coordinate direction, the probability of a collapse is the same along all coordinate directions that f spans. Using the union bound, the probability that $g^j(f)$ has not collapsed to a vertex is at most $k/2^j$. With Y as in the statement of the lemma, it follows that

$$P(Y \geq j) \leq k/2^j.$$

Hence,

$$\mathbb{E}[Y] = \sum_{j=1}^{\infty} j P(Y = j) = \sum_{j=1}^{\infty} P(Y \geq j)$$

$$\leq \log k + \sum_{c=1}^{\infty} \sum_{j=c \log k}^{(c+1) \log k} P(Y \geq j)$$

$$\leq \log k + \sum_{c=1}^{\infty} \sum_{j=c \log k}^{(c+1) \log k} P(Y \geq c \log k)$$

$$\leq \log k + \sum_{c=1}^{\infty} \log k \frac{k}{2^{c \log k}}$$

$$\leq \log k + \log k \sum_{c=1}^{\infty} \frac{1}{k^{c-1}}$$

$$\leq \log k + 2 \log k \leq 3 \log k.$$

$\square$

As a consequence of the lemma, the expected "lifetime" of k-simplices in our tower with $k > 0$ is rather short: given a flag $e_0 \subseteq \cdots \subseteq e_\ell$, the face e_ℓ will be mapped to a vertex after $O(\log d)$ steps, and so will be all its sub-faces, turning the flag into a vertex. It follows that summing up the total number of k-simplices with $k > 0$ over $\mathcal{X}_\alpha$ for all $\alpha \geq 0$ yields an upper bound of $n2^{O(d \log k + d)}$ as well.

Algorithm description

Recall that a simplicial map can be written as a composition of simplex inclusions and contractions of vertices (Dey et al. 2014; Kerber and Schreiber 2017). That means, given the complex $\mathcal{X}_{\alpha_s}$, to describe the complex at the next scale α_{s+1}, it suffices to specify

- which pairs of vertices in $\mathcal{X}_{\alpha_s}$ map to the same image under $\tilde{g}$, and
- which simplices in $\mathcal{X}_{\alpha_{s+1}}$ are included at scale $\mathcal{X}_{\alpha_{s+1}}$.

The input is a set of n points $P \subset \mathbb{R}^d$. The output is a list of *events*, where each event is of one of the three diam types:

- A *scale event* defines a real value α and signals that all upcoming events happen at scale α (until the next scale event).
- An *inclusion event* introduces a new simplex, specified by the list of vertices on its boundary (we assume that every vertex is identified by a unique integer).
- A *contraction event* is a pair of vertices (i, j) from the previous scale, and signifies that i and j are identified as the same from that scale.

In a first step, we estimate the range of scales that we are interested in. We compute a 2-approximation of $diam(P)$ by taking any point $p \in P$ and calculating $\max_{q \in P} \|p - q\|$. Then we compute $CP(P)$ using a randomized algorithm in $n2^{O(d)}$ expected time (Khuller and Matias 1995).

Next, we proceed scale-by-scale and construct the list of events accordingly. On the lowest scale, we simply compute the active vertices by point location for P in a cubical grid, and enlist n inclusion events (this is the only step where the input points are considered in the algorithm).

For the data structure, we use an auxiliary container S and maintain the invariant that whenever a new scale is considered, S consists of all simplices of the previous

scale, sorted by dimension. In S, for each vertex, we store an id and a coordinate representation of the active face to which it corresponds. Every ℓ-simplex with $\ell > 0$ is stored just as a list of integers, denoting its boundary vertices. We initialize S with the n active vertices at the lowest scale.

Let $\alpha < \alpha'$ be any two consecutive scales with $\square, \square'$ the respective cubical complexes and $\mathcal{X}, \mathcal{X}'$ the approximation complexes, with $\tilde{g} : \mathcal{X} \to \mathcal{X}'$ being the simplicial map connecting them. Suppose we have already constructed all events at scale α.

- First, we enlist the scale event for α'.
- Then, we enlist the contraction events. For that, we iterate through the vertices of $\mathcal{X}$ and compute their value under g, using point location in a cubical grid. We store the results in a list S' (which contains the simplices of $\mathcal{X}'$). If for a vertex j, $g(j)$ is found to be equal to $g(i)$ for a previously considered vertex i, we choose the minimal such i and enlist a contraction event for (i, j).
- We turn to the inclusion events:

  - We start with the case of vertices. Every vertex of $\mathcal{X}'$ is either an active face or a secondary face of $\square'$. Each active face must contain an active vertex, which is also a vertex of $\mathcal{X}'$. We iterate through the elements in S'. For each active vertex v encountered, we go over all faces of the cubical complex $\square'$ that contain v as a vertex, and check whether they are active. For every active face E encountered that is not in S' yet, we add it to S' and enlist an inclusion event of a new 0-simplex. Additionally, we go over each face of E, add it to S' and enlist a vertex inclusion event, thereby enumerating the secondary faces that are in E. At termination, all vertices of $\mathcal{X}'$ have been detected.
  - Next, we iterate over the simplices of S of dimension ≥ 1, and compute their image under $\tilde{g}$ using the pre-computed vertex map; we store the result in S'.
  - To find the simplices of dimension ≥ 1 included at $\mathcal{X}'$, we exploit our previous insight that they contain at least one vertex that is included at the same scale (see the proof of Theorem 4). Hence, we iterate over the vertices included in $\mathcal{X}'$ and find the included simplices inductively in dimension.
    Let v be the current vertex under consideration; assume that we have found all $(p-1)$-simplices in $\mathcal{X}'$ that contain v. Each such $(p-1)$-simplex σ is a flag of length p in $\square'$. We iterate over all faces e that extend σ to a flag of length $p+1$. If e is active, we have found a p-simplex in $\mathcal{X}'$ incident to v. If this simplex is not in S' yet, we add it and enlist an inclusion event for it. We also enqueue the simplex in our inductive procedure, to look for $(p+1)$-simplices in the next round. At the end of the procedure, we have detected all simplices in $\mathcal{X}'$ without preimage, and S' contains all simplices of $\mathcal{X}'$. We set $S \leftarrow S'$ and proceed to the next scale.

This ends the description of the algorithm.

Theorem 5 *To compute the k-skeleton, the algorithm takes*

$$n2^{O(d)} \log \Delta + 2^{O(d)} M$$

time in expectation and M space, where M denotes the size of the tower. In particular, the expected time is bounded by

$$n2^{O(d)} \log \Delta + n2^{O(d \log k+d)}$$

and the space is bounded by $n2^{O(d \log k+d)}$.

Proof In the analysis, we ignore the costs of point locations in grids, checking whether a face is active, and searches in data structures S, since all these steps have negligible costs when appropriate data structures are chosen.

Computing the image of a vertex of $\mathcal{X}$ costs $O(2^d)$ time. Moreover, there are at most $n2^{O(d)}$ vertices altogether in the tower in expectation (using Lemma 10), so this bound in particular holds on each scale. Hence, the contraction events on a fixed scale can be computed in $n2^{O(d)}$ time. Finding new active vertices requires iterating over the cofaces of a vertex in a cubical complex. There are 3^d such cofaces for each vertex. This has to be done for a subset of the vertices in $\mathcal{X}'$, so the running time is also $n2^{O(d)}$. Further, for each new active face, we go over its $2^{O(d)}$ faces to enlist the secondary faces, so this step also consumes $n2^{O(d)}$ time. Since there are $O(\log \Delta + \log d)$ scales considered, these steps require $n2^{O(d)} \log \Delta$ over all scales.

Computing the image of $\tilde{g}$ for a fixed scale costs at most $O(2^d|\mathcal{X}|)$. M is the size of the tower, that is, the simplices without preimage, and I is the set of scales considered. The expected bound for $\sum_{\alpha \in I} |\mathcal{X}_\alpha| = O(\log dM)$, because every simplex has an expected lifetime of at most $3 \log d$ by Lemma 11. Hence, the cost of these steps is bounded by $2^{O(d)}M$.

In the last step of the algorithm, we find the simplices of $\mathcal{X}'$ included at α'. We consider a subset of simplices of $\mathcal{X}'$, and for each, we iterate over a collection of faces in the cubical complex of size at most $2^{O(d)}$. Hence, this step is also bounded by $2^{O(d)}|\mathcal{X}|$ per scale, and hence bounded $2^{O(d)}M$ as well.

For the space complexity, the auxiliary data structure S gets as large as $\mathcal{X}$, which is clearly bounded by M. For the output complexity, the number of contraction events is at most the number of inclusion events, because every contraction removes a vertex that has been included before. The number of inclusion events is the size of the tower. The number of scale events as described is $O(\log \Delta + \log d)$. However, it is simple to get rid of this factor by only including scale events in the case that at least one inclusion or contraction takes place at that scale. The space complexity bound follows. $\qquad \square$

5.3 Dimension reduction

When the ambient dimension d is large, our approximation scheme can be combined with dimension reduction techniques to reduce the final complexity, very similar to the application in Choudhary et al. (2017b). For a set of n points $P \subset \mathbb{R}^d$, we apply the dimension reduction schemes of Johnson-Lindenstrauss (JL) (Johnson et al. 1986), Matoušek (MT) (Matoušek 1990), and Bourgain's embedding (BG) (Bourgain 1985). We then compute the approximation on the lower-dimensional point set. We only state the main results in Table 1, leaving out the proofs since they are very similar to those from Choudhary et al. (2017b).

Table 1 Comparison of dimension reduction techniques: here the approximation ratio is for the Rips persistence module, and the size refers to the size of the k-skeleton of the approximation

technique	approximation ratio	size	runtime
JL	$O(\log^{0.25} n)$	$n^{O(\log k)}$	$n^{O(1)} \log \Delta + n^{O(\log k)}$
MT	$O((\log n)^{0.75} (\log \log n)^{0.25})$	$n^{O(1)}$	$n^{O(1)} \log \Delta$
BG + MT	$O((\log n)^{1.75} (\log \log n)^{0.25})$	$n^{O(1)}$	$n^{O(1)} \log \Delta$

6 Approximation scheme with cubical complexes

We extend our approximation scheme to use cubical complexes in place of simplicial complexes. We start by detailing a few aspects of cubical complexes.

6.1 Cubical complexes

We now briefly describe the concept of cubical complexes, essentially expanding upon the contents of Sect. 3.1. For a detailed overview of cubical homology, we refer to Kaczynski et al. (2004).

6.1.1 Definition

We define cubical complexes over the grids G_{α_s}. For any fixed α_s, the grids G_{α_s} defines a natural collection of cubes. An *elementary cube* γ is a product of intervals $\gamma = I_1 \times I_2 \times \cdots \times I_d$, where each interval is of the form $I_j = (x_j, x_j + m_j)$, such that the vertex $(x_1, \cdots, x_m) \in G_{\alpha_s}$ and each m_j is either 0 or α_s. That means, an (elementary) cube is simply a face of a d-cube of the grid. An interval I_j is said to be *degenerate* if $m_j = 0$. The dimension of γ is the number of non-degenerate intervals that defines it. We define the boundary of any interval as the two degenerate intervals that form its endpoints and denote this by $\partial(I_j) = (x_j, x_j) + (x_j + m_j, x_j + m_j)$. Taking the boundary of any fixed subset of the intervals defining γ consecutively gives a sum of faces of γ. A *cubical complex* of G_{α_s} is a finite collection of cubes of G_{α_s}.

We define chain complexes for the cubical case in the same way as in simplicial complexes. The chain complexes are connected by boundary homomorphisms, where the boundary of a cube is defined as:

$$\partial (I_1 \times \cdots \times I_d) = (\partial(I_1) \times I_2 \times \cdots \times I_d) + \cdots + (I_1 \times \cdots \times I_{d-1} \times \partial(I_d)),$$

where $(I_1 \times \cdots \times \partial(I_j) \times \cdots \times I_d)$ denotes the sum

$$(I_1 \times \cdots \times (x_i, x_i) \times \cdots \times I_d) + (I_1 \times \cdots \times (x_i + m_i, x_i + m_i) \times \cdots \times I_d).$$

It can be quickly verified that for each cube γ, $\partial \circ \partial(\gamma) = 0$ since each term appears twice in the expression and the addition is over $\mathbb{Z}_2$.

6.1.2 Cubical maps and induced homology

Let T_{α_s} and T_{α_t} denote the cubical complexes defined by the grids G_{α_s} and G_{α_t}, respectively, for $s \leq t$. We use the vertex map $g : G_{\alpha_s} \rightarrow G_{\alpha_t}$ to define a map between the cubical complexes. Note that if (a, b) are vertices of a cube of T_{α_s} that differ in one coordinate, then $(g(a), g(b))$ are vertices of a cube of T_{α_t} that differ in at most one coordinate. A *cubical map* is a map $f : T_{\alpha_s} \rightarrow T_{\alpha_t}$ defined using g, such that for each cube $\gamma = [a_1, b_1] \times \cdots \times [a_d, b_d]$ of T_{α_s}, $f(\gamma) := [g(a_1), g(b_1)] \times \cdots \times [g(a_d), g(b_d)]$ spans a cube of T_{α_t}. The cubical map can also be restricted to sub-complexes of T_{α_s} and T_{α_t}, provided that the image $f(\gamma)$ is well-defined.

Each cubical map also defines a corresponding continuous map between the underlying spaces of the respective complexes. Let $x \in |\gamma|$ be a point in γ. Then, the coordinates of x can be uniquely written as $x = [\lambda_1 a_1 + (1 - \lambda_1)b_1, \cdots, \lambda_d a_d + (1 - \lambda_d)b_d]$ where each $\lambda_i \in [0, 1]$. The image of x under the continuous extension of f is the point $[\lambda_1 g(a_1) + (1 - \lambda_1)g(b_1), \cdots, \lambda_d g(a_d) + (1 - \lambda_d)g(b_d)]$ in the cube $g(\gamma)$.

The cubical map f gives rise to a chain map $f_\# : C_p(T_{\alpha_s}) \rightarrow C_p(T_{\alpha_t})$ between the p-th chain groups of the complexes, for each $p \in [0, \cdots, d]$. For each cube γ, $f_\#(\gamma) = f(\gamma)$ if $dim(\gamma) = dim(f(\gamma))$ and 0 otherwise. For any chain $c = \sum_i \gamma_i$, the chain map is defined linearly $f_\#(c) = \sum_i f_\#(\gamma_i)$. It is simple to verify that $\partial \circ f_\# = f_\# \circ \partial$, so this gives a homomorphism between the chain groups.

Moving to the homology level, we get the respective homology groups $H(T_{\alpha_s})$ and $H(T_{\alpha_t})$ and the chain map from above induces a linear map between them. The concept of reduced homology and augmentation maps is also applicable to the cubical chain complexes. For a sequence of cubical complexes connected with cubical maps, this generates a persistence module.

Cubical filtrations and towers are defined in a similar manner to the simplicial case. A *cubical filtration* is a collection of cubical complexes $(T_\alpha)_{\alpha \in I}$ such that $T_\alpha \subseteq T'_\alpha$ for all $\alpha \leq \alpha' \in I$. A (cubical) tower is a sequence $(T_\alpha)_{\alpha \in J}$ of cubical complexes with J being an index set together with cubical maps between complexes at consecutive scales. A cubical tower can be written as a sequence of inclusions and contractions, where an inclusion refers to the addition of a cube and a contraction refers to collapsing a cube along a coordinate direction to either of the endpoints of the interval.

6.2 Description

We choose the simplest possible cubical complex to define our approximation cubical tower: for each scale α_s, we define the cubical complex U_{α_s} as the set of active faces and secondary faces spanned by V_{α_s}. Hence the cubical complex is closed under taking faces and is well-defined. See Fig. 5 for a simple example.

Recall from Sect. 4 that for each $s \in \mathbb{Z}$, U_{α_s} and $U_{\alpha_{s+1}}$ are related by a cubical map g_{α_s}, which gives rise to the cubical tower

$$\left(U_{\alpha_s}\right)_{s \in \mathbb{Z}}.$$

We extend this to a tower $(U_\alpha)_{\alpha \geq 0}$ by using techniques from Appendix A. In Sect. 4 we saw that the tower $(\mathcal{X}_\alpha)_{\alpha \geq 0}$ gives an approximation to the Rips filtration. The relation between the simplicial and cubical towers is trivial: $\mathcal{X}_{\alpha_s}$ is simply a triangulation of $|U_{\alpha_s}|$. Hence $\mathcal{X}_{\alpha_s}$ and U_{α_s} have the same homology (Munkres 1984). Moreover, the simplicial map is derived from an application of the cubical map. In particular, the continuous versions of both maps are the same. For any $0 \leq \alpha \leq \beta$, let

- $f_1 : H_*(U_\alpha) \rightarrow H_*(U_\beta)$ denote the homomorphism induced by the cubical map,
- $f_2 : H_*(\mathcal{X}_\alpha) \rightarrow H_*(\mathcal{X}_\beta)$ denote the homomorphism induced by the simplicial map, and
- $f_0 : H_*(|\mathcal{X}_\alpha| = |U_\alpha|) \rightarrow H_*(|\mathcal{X}_\beta| = |U_\beta|)$ denote the homomorphism induced by the common continuous map.

It is well-established that $f_1 = f_0$ (Kaczynski et al. 2004, Chapter. 6) and $f_2 = f_0$ (Munkres 1984, Chapter. 2). Therefore, we conclude that the persistence modules $\left(H(U_\alpha)\right)_{\alpha \geq 0}$ and $\left(H(\mathcal{X}_\alpha)\right)_{\alpha \geq 0}$ are persistence-equivalent. Combining this observation with the result of Theorem 3, we get

Theorem 6 *The scaled persistence modules*

- $\left(H(U_{2\alpha})\right)_{\alpha \geq 0}$ *and the* L_∞*-Rips module* $\left(H(\mathcal{R}_\alpha^\infty)\right)_{\alpha \geq 0}$ *are 2-approximations of each other, and*
- $\left(H(U_{2\sqrt[4]{d}\alpha})\right)_{\alpha \geq 0}$ *and the Rips module* $\left(H(\mathcal{R}_\alpha)\right)_{\alpha \geq 0}$ $2d^{0.25}$*-approximate each other.*

To compute the cubical tower, we simply re-use the algorithm for the simplicial case, with small changes:

- In the simplicial case, we used a container S to hold the simplices from the previous scale. We alter S to store the cubes from the previous scale. For each interval, we store an id and its coordinates. Each cube is stored as the set of ids of the intervals that define it.
- At each scale, we enumerate the image of the cubical map by computing the image of each interval, and then use this pre-computed map to compute the image of (≥ 1)-dimensional cubes.
- For the inclusions, we find all the active and secondary faces but do not compute the simplices. The inclusions in the cubical tower correspond exactly to the inclusions of active and secondary faces in the simplicial tower, so this enumerates all inclusions correctly.

From Lemma 10 at most $n3^d$ active faces are added to the tower. Hence at most $n3^d 3^d = n6^d$ active and secondary faces are added to the tower. Computing the tower takes time as in Theorem 5 by replacing M with the size bound. We conclude that:

Theorem 7 *The cubical tower has size at most $n6^d$ and takes at most $n6^d \log \Delta$ time in expectation to compute, where Δ is the spread of the point set.*

7 Discussion

7.1 Practicality

We now touch upon the practical aspects of our constructions. An implementation of our approximation scheme would be a tool that computes the (approximate) persistence barcode for any input data set. For any scheme to be useful in practice, it should be able to compute sufficiently close approximations using a reasonable amount of resources.

Our cubical tower consists of cubical complexes connected via cubical maps. To our knowledge, there are no algorithms to compute barcodes in this setting where the cubical maps are more than just trivial inclusions. As such, although our cubical scheme has exponentially lower theoretical guarantees compared to the simplicial tower, we can not hope to test it in practice unless the appropriate primitives are available. It could be an interesting research direction to develop this primitive and in particular investigate whether the techniques used in computing persistence barcodes for a simplicial tower allow a generalization to the cubical case.

It makes more sense to inspect the simplicial tower. We saw in Theorem 4 that the size of the tower is $n6^{d-1}(2k+4)(k+3)!\left\{{d \atop k+2}\right\}$. Unfortunately, this bound is already too large so that the storage requirement of the Algorithm (Theorem 7) explodes exponentially. Let us assume a conservative bound of 1 Byte of memory requirement per simplex. For a point set in $d = 8$ dimensions and $k = 4$, the complexity bound is already at least 4000 Terabytes, before factoring in n. For a point set in $d = 10$ dimensions and $k = 5$, this explodes to 10^{20} Terabytes. While these are upper bounds, in practice the complexity will still need to be many orders of magnitude smaller to be feasile, which is unlikely. Even with conservative estimates our storage requirement is impractical.

Therefore we are not very hopeful that implementing the scheme in its current state will provide any useful insight for high dimensional approximations. Making it implementation-worthy demands more optimizations and tools at the algorithmic level. This is worth another Algorithmic engineering project in its own right. We plan to pursue this line of research in the future. Since our focus in this paper was geared towards theoretical aspects of approximations, we exclude experimental results in the current work. We hope that a more careful implementation-focussed approach may prove more practical.

On the other hand, the upper bound for the cubical case is simply $n6^d$. Even for $d = 10$, the storage requirement would be less than 100 Megabytes before factoring in n. This is far more attractive than the simplicial case. As such, it may make more sense to invest time and effort in developing tools to compute barcodes in the cubical setup.

7.2 Summary

We presented an approximation scheme for the Rips filtration, with improved approximation ratio, size and computational complexity than previous approaches for the

case of high-dimensional point clouds. In particular, we are able to achieve a marked reduction in the size of the approximation by using cubical complexes in place of simplicial complexes. This is in contrast to all other previous approaches that used simplicial complexes as approximating structures.

An important technique that we used in our scheme is the application of acyclic carriers to prove interleaving results. An alternative would to be explicitly construct chain maps between the Rips and the approximation towers; unfortunately, this make the interleaving analysis significantly more complex. While the proof of the interleaving in Sect. 4.3 is still technically challenging, it greatly simplifies by the usage of acyclic carriers. There is also no benefit in knowing the interleaving maps because they are only required for the analysis of the interleaving, and not for the actual computation of the approximation tower. We believe that this technique is of general interest for the construction of approximations of cell complexes.

Our simplicial tower is connected by simplicial maps; there are (implemented) algorithms to compute the barcode of such towers (Dey et al. 2014; Kerber and Schreiber 2017). It is also quite easy to adapt our tower construction to a streaming setting (Kerber and Schreiber 2017), where the output list of events is passed to an output stream instead of being stored in memory.

Declarations

A Strong interleaving for barycentric scheme

Recall that we build the approximation tower over the set of scales $I := \{\alpha_s = 2^s \mid s \in \mathbb{Z}\}$. The tower $(\mathcal{X}_\alpha)_{\alpha \in I}$ connected with the simplicial map $\tilde{g}$ can be extended to the set of scales $\{\alpha \geq 0\}$ with simple modifications:

– for $\alpha \in I$, we define $\mathcal{X}_\alpha$ in the usual manner. The map $\tilde{g}$ stays the same as before for complexes at such scales.

– for all $\alpha \in [\alpha_s, \alpha_{s+1})$, we set $\mathcal{X}_\alpha = \mathcal{X}_{\alpha_s}$, for any $\alpha_s \in I$. That means, the complex stays the same in the interval between any two scales of I, so we define $\tilde{g}$ as the identity within this interval.

These give rise to the tower $(\mathcal{X}_\alpha)_{\alpha \geq 0}$, that is connected with the simplicial map $\tilde{g}$. This modification helps in improving the interleaving with the Rips persistence module.

First, we extend the acyclic carriers C_1 and C_2 from before to the new case:

– $C_1^\alpha : \mathcal{R}_\alpha^\infty \rightarrow \mathcal{X}_{4\alpha}, \alpha > 0$: we define C_1 as before, simply changing the scales in the definition. It is straightforward to see that C_1 is still a well-defined acyclic carrier.
– $C_2^\alpha : \mathcal{X}_\alpha \rightarrow \mathcal{R}_\alpha^\infty, \alpha \geq 0$: this stays the same as before. It is simple to check that C_2 is still a well-defined acyclic carrier.

These give rise to augmentation-preserving chain maps between the chain complexes:

$$c_1^\alpha : C_* \left(\mathcal{R}_\alpha^\infty \right) \rightarrow C_* (\mathcal{X}_{4\alpha}) \qquad \text{and} \qquad c_2^\alpha : C_* (\mathcal{X}_\alpha) \rightarrow C_* \left(\mathcal{R}_\alpha^\infty \right),$$

using the acyclic carrier theorem as before (Theorem 1).

Lemma 12 *The diagram*

$$
\begin{array}{ccc}
C_*(\mathcal{R}_\alpha^\infty) & \xrightarrow{\ inc\ } & C_*(\mathcal{R}_{\alpha'}^\infty) \\
& \searrow{\scriptstyle c_1} & \quad\searrow{\scriptstyle c_1} \\
& C_*(\mathcal{X}_{4\alpha}) & \xrightarrow{\ \tilde{g}\ } C_*(\mathcal{X}_{4\alpha'})
\end{array}
\tag{12}
$$

commutes on the homology level, for all $0 \leq \alpha \leq \alpha'$.

Proof Consider the acyclic carrier $C_1 \circ inc : \mathcal{R}_\alpha^\infty \rightarrow \mathcal{X}_{4\alpha'}$. It is simple to verify that this carrier carries both $c_1 \circ inc$ and $\tilde{g} \circ c_1$, so the induced diagram on the homology groups commutes, from Theorem 1. □

Lemma 13 *The diagram*

$$
\begin{array}{ccc}
 & C_*(\mathcal{R}_\alpha^\infty) & \xrightarrow{\ inc\ } C_*(\mathcal{R}_{\alpha'}^\infty) \\
\nearrow{\scriptstyle c_2} & & \nearrow{\scriptstyle c_2} \\
C_*(\mathcal{X}_\alpha) & \xrightarrow{\ \tilde{g}\ } C_*(\mathcal{X}_{\alpha'}) &
\end{array}
\tag{13}
$$

commutes on the homology level, for all $0 \leq \alpha \leq \alpha'$.

Proof We construct an acyclic carrier $D : \mathcal{X}_\alpha \rightarrow \mathcal{R}_{\alpha'}^\infty$ which carries $inc \circ c_2$ and $c_2 \circ \tilde{g}$, thereby proving the claim (Theorem 1).

Consider any simplex $\sigma \in \mathcal{X}_\alpha$ and let $E \in \square_\alpha$ be the minimal active face of containing σ. We set $D(\sigma)$ as the simplex on the set of input points of P, which lie in the Voronoi regions of the vertices of $g(E)$. By the triangle inequality, $D(\sigma)$ is a simplex of $\mathcal{R}_{\alpha'}^\infty$, so that D is a well-defined acyclic carrier. It is straightforward to verify that D carries both $c_2 \circ \tilde{g}$ and $inc \circ c_2$. □

Lemma 14 *The diagram*

$$
\begin{array}{ccc}
& \mathcal{C}_*(\mathcal{R}_\alpha^\infty) \xrightarrow{\ inc\ } \mathcal{C}_*(\mathcal{R}_{\alpha'}^\infty) & \\
{\scriptstyle c_2}\nearrow & & \searrow{\scriptstyle c_1} \\
\mathcal{C}_*(\mathcal{X}_\alpha) \xrightarrow{\quad\quad \tilde{g} \quad\quad} & & \mathcal{C}_*(\mathcal{X}_{4\alpha'})
\end{array}
\tag{14}
$$

commutes on the homology level, for all $0 \le \alpha \le \alpha'$.

Proof The diagram is essentially the same as the lower triangle of Diagram 10, with a change in the scales. As a result, the proof of Lemma 7 also applies for our claim directly. □

Lemma 15 *The diagram*

$$
\begin{array}{ccc}
\mathcal{C}_*(\mathcal{R}_\alpha^\infty) \xrightarrow{\quad\quad inc \quad\quad} & & \mathcal{C}_*(\mathcal{R}_{4\alpha'}^\infty) \\
{\scriptstyle c_1}\searrow & & \nearrow{\scriptstyle c_2} \\
& \mathcal{C}_*(\mathcal{X}_{4\alpha}) \xrightarrow{\ \tilde{g}\ } \mathcal{C}_*(\mathcal{X}_{4\alpha'}) &
\end{array}
\tag{15}
$$

commutes on the homology level, for all $0 \le \alpha \le \alpha'$.

Proof The diagram can be re-interpreted as:

$$
\begin{array}{ccc}
\mathcal{C}_*(\mathcal{R}_\alpha^\infty) \xrightarrow{\quad\quad inc \quad\quad} & & \mathcal{C}_*(\mathcal{R}_{4\alpha'}^\infty) \\
{\scriptstyle \tilde{g}\circ c_1}\searrow & & \nearrow{\scriptstyle c_2} \\
& \mathcal{C}_*(\mathcal{X}_{4\alpha'}) &
\end{array}
\tag{16}
$$

The modified diagram is essentially the same as the upper triangle of Diagram 10, with a change in the scales and a replacement of c_1 with $\tilde{g} \circ c_1$, that is equivalent to the chain map at the scale α'. Hence, the proof of Lemma 8 also applies for our claim directly. □

Using Lemmas 12, 13, 14, 15, and the scale balancing technique for strongly interleaved persistence modules, it follows that

Lemma 16 *The persistence modules* $\big(H(\mathcal{X}_{2\alpha})\big)_{\alpha \ge 0}$ *and* $\big(H(\mathcal{R}_\alpha^\infty)\big)_{\alpha \ge 0}$ *are strongly 2-interleaved.*

References

Botnan, M., Spreemann, G.: Approximating persistent homology in Euclidean space through collapses. Appl. Algebra Eng. Commun. Comput. **26**(1–2), 73–101 (2015)

Bourgain, J.: On Lipschitz embedding of finite metric spaces in Hilbert space. Israel J. Math. **52**(1–2), 46–52 (1985)

Bubenik, P., Scott, J.A.: Categorification of persistent homology. Discrete Comput. Geom. **51**(3), 600–627 (2014)

Bubenik, P., de Silva, V., Scott, J.: Metrics for generalized persistence modules. Found. Comput. Math. **15**(6), 1501–1531 (2015)

Carlsson, G.: Topology and data. Bull. Am. Math. Soc. **46**, 255–308 (2009)

Carlsson, G., Zomorodian, A.: Computing persistent homology. Discrete Comput. Geom. **33**(2), 249–274 (2005)

Cavanna, N., Jahanseir, M., Sheehy, D.: A Geometric perspective on sparse filtrations. In: Proceedings of the 27th Canadian Conference on Computational Geometry (CCCG), pp. 116–121 (2015)

Chazal, F., Cohen-Steiner, D., Glisse, M., Guibas, L., Oudot, S.: Proximity of persistence modules and their diagrams. In: ACM Symposium on Computational Geometry (SoCG), pp. 237–246 (2009)

Choudhary, A., Kerber, M., Raghavendra, S.: Improved approximate rips filtrations with shifted integer lattices. In: Proceedings of the 25th Annual European Symposium on Algorithms (ESA), pp. 28:1–28:13 (2017)

Choudhary, A., Kerber, M., Raghavendra, S.: Polynomial-sized topological approximations using the permutahedron (extended version). Discrete Comput. Geom. (2017)

Choudhary, A., Kerber, M., Raghavendra, S.: Improved topological approximations by digitization. In: Proceedings of the Symposium on Discrete Algorithms (SODA), pp. 448:1–448:14 (2019)

Dey, T.K., Fan, F., Wang, Y.: Computing topological persistence for simplicial maps. In: Proceedings of the 30th Annual Symposium on Computational Geometry (SoCG), pp. 345–354 (2014)

Edelsbrunner, H., Harer, J.: Computational Topology—An Introduction. American Mathematical Society, New York (2010)

Edelsbrunner, H., Kerber, M.: Dual complexes of cubical subdivisions of $\mathbb{R}^n$. Discrete Comput. Geom. **47**(2), 393–414 (2012)

Edelsbrunner, H., Mücke, E.P.: Simulation of simplicity: a technique to cope with degenerate cases in geometric algorithms. ACM Trans. Gr. 66–104 (1990)

Edelsbrunner, H., Letscher, D., Zomorodian, A.: Topological persistence and simplification. Discrete Comput. Geom. **28**(4), 511–533 (2002)

Goodman, J.E., O'Rourke, J., Tóth, C.D. (eds.): Handbook of Computational Geometry. CRC Press, Boca Raton (2017)

Hatcher, A.: Algebraic Topology. Cambridge University Press, Cambridge (2002)

Johnson, W.B., Lindenstrauss, J., Schechtman, G.: Extensions of Lipschitz maps into Banach spaces. Israel J. Math. **54**(2), 129–138 (1986)

Kaczynski, T., Mischaikow, K., Mrozek, M.: Computational Homology. Applied Mathematical Sciences, Springer, New York (2004)

Kerber, M., Schreiber, H.: Barcodes of towers and a streaming algorithm for persistent homology. In: Proceedings of 33rd International Symposium on Computational Geometry (SoCG), pp. 57:1–57:15 (2017)

Kerber, M., Sharathkumar, R.: Approximate Čech complex in low and high dimensions. In: Algorithms and Computation—24th International Symposium (ISAAC), pp. 666–676 (2013)

Khuller, S., Matias, Y.: A simple randomized sieve algorithm for the closest-pair problem. Inf. Comput. **118**(1), 34–37 (1995)

Matoušek, J.: Bi-Lipschitz embeddings into low-dimensional Euclidean spaces. Commentationes Mathematicae Universitatis Carolinae (1990)

Munkres, J.R.: Elements of Algebraic Topology. Westview Press, Milton Park (1984)

Rennie, B.C., Dobson, A.J.: On stirling numbers of the second kind. J. Comb. Theory **7**(2), 116–121 (1969)

Sheehy, D.: Linear-size approximations to the Vietoris-rips filtration. Discrete Comput. Geom. **49**(4), 778–796 (2013)

Wagner, H., Chen, C., Vuçini, E.: Efficient Computation of Persistent Homology for Cubical Data, pp. 91–106. Springer, Berlin (2012)