

Dimensionality reduction for k-distance applied to persistent homology

Shreya Arya, Jean-Daniel Boissonnat, Kunal Dutta, Martin Lotz

► To cite this version:

Shreya Arya, Jean-Daniel Boissonnat, Kunal Dutta, Martin Lotz. Dimensionality reduction for k-distance applied to persistent homology. Journal of Applied and Computational Topology, 2021, 5, pp.671-691. 10.1007/s41468-021-00079-x . hal-03412594

HAL Id: hal-03412594 https://amu.hal.science/hal-03412594

Submitted on 3 Nov 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers. L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Dimensionality reduction for *k*-distance applied to persistent homology

Shreya Arya¹ · Jean-Daniel Boissonnat² · Kunal Dutta³ (b) · Martin Lotz⁴

Received: 11 May 2020 / Revised: 19 June 2021 / Accepted: 27 September 2021 / Published online: 19 October 2021 © The Author(s) 2021

Abstract

Given a set *P* of *n* points and a constant *k*, we are interested in computing the persistent homology of the Čech filtration of *P* for the *k*-distance, and investigate the effectiveness of dimensionality reduction for this problem, answering an open question of Sheehy (The persistent homology of distance functions under random projection. In Cheng, Devillers (eds), 30th Annual Symposium on Computational Geometry, SOCG'14, Kyoto, Japan, June 08–11, p 328, ACM, 2014). We show that *any* linear transformation that preserves pairwise distances up to a $(1 \pm \varepsilon)$ multiplicative factor, must preserve the persistent homology of the Čech filtration up to a factor of $(1 - \varepsilon)^{-1}$. Our results also show that the Vietoris-Rips and Delaunay filtrations for the *k*-distance, as well as the Čech filtration for the approximate *k*-distance of Buchet et al. [J Comput

Kunal DuttaK.Dutta@mimuw.edu.pl; duttkunal@gmail.com

Shreya Arya shreya.arya14@gmail.com

Jean-Daniel Boissonnat jean-daniel.boissonnat@inria.fr

Martin Lotz martin.lotz@warwick.ac.uk

- ¹ Duke University, Durham, USA
- ² Université Côte d'Azur, INRIA, Sophia-Antipolis, France
- ³ Faculty of Mathematics, Informatics, and Mechanics, University of Warsaw, Warsaw, Poland
- ⁴ Mathematics Institute, University of Warwick, Coventry, UK

S. Arya, J.-D. Boissonnat, K. Dutta: The research leading to these results has received funding from the European Research Council (ERC) under the European Union's Seventh Framework Programme (FP/2007-2013) / ERC Grant Agreement No. 339025 GUDHI (Algorithmic Foundations of Geometry Understanding in Higher Dimensions). J.-D. Boissonnat: Supported by the French government, through the 3IA Côte d'Azur Investments in the Future project managed by the National Research Agency (ANR) with the reference number ANR-19-P3IA-0002. K. Dutta: Supported by the Polish NCN SONATA Grant no. 2019/35/D/ST6/04525.

Geom, 58:70–96, 2016] are preserved up to a $(1 \pm \varepsilon)$ factor. We also prove extensions of our main theorem, for point sets (*i*) lying in a region of bounded Gaussian width or (*ii*) on a low-dimensional submanifold, obtaining embeddings having the dimension bounds of Lotz (Proc R Soc A Math Phys Eng Sci, 475(2230):20190081, 2019) and Clarkson (Tighter bounds for random projections of manifolds. In Teillaud (ed) Proceedings of the 24th ACM Symposium on Computational Geom- etry, College Park, MD, USA, June 9–11, pp 39–48, ACM, 2008) respectively. Our results also work in the *terminal dimensionality reduction* setting, where the distance of any point in the original ambient space, to any point in *P*, needs to be approximately preserved.

Keywords Dimensionality reduction \cdot Johnson-Lindenstrauss lemma \cdot Topological data analysis \cdot Persistent homology $\cdot k$ -distance \cdot Distance to measure

Mathematics Subject Classification 55N31 · 68U05

1 Introduction

Persistent homology is one of the main tools to extract information from data in topological data analysis. Given a data set as a point cloud in some ambient space, the idea is to construct a filtration sequence of topological spaces from the point cloud, and extract topological information from this sequence. The topological spaces are usually constructed by considering balls around the data points, in some given metric of interest, as the open sets. However the usual distance function is highly sensitive to the presence of outliers and noise. One approach is to use distance functions that are more robust to outliers, such as the *distance-to-a-measure* and the related *k-distance* (for finite data sets), proposed recently by Chazal et al. (2011) Although this is a promising direction, an exact implementation can have significant cost in run-time. To overcome this difficulty, approximations of the *k*-distance have been proposed recently that led to certified approximations of persistent homology (Guibas et al. 2013; Buchet et al. 2016). Other approaches involve using kernels (Phillips et al. 2015) and de-noising algorithms (Buchet et al. 2018; Zhang 2013).

In all the above settings, the sub-routines required for computing persistent homology have exponential or worse dependence on the ambient dimension, and rapidly become unusable in real-time once the dimension grows beyond a few dozens - which is indeed the case in many applications, for example in image processing, neurobiological networks, and data mining (see e.g. Giraud 2014). This phenomenon is often referred to as the *curse of dimensionality*.

The Johnson-Lindenstrauss Lemma. One of the simplest and most commonly used mechanisms to mitigate this curse, is that of *random projections*, as applied in the celebrated Johnson-Lindenstrauss lemma (JL Lemma for short) (Johnson and Lindenstrauss 1984). The JL Lemma states that any set of *n* points in Euclidean space can be embedded into a space of dimension $O(\varepsilon^{-2} \log n)$ with $(1 \pm \varepsilon)$ distortion. Since the initial non-constructive proof of this fact by Johnson and Lindenstrauss (1984), several authors have given successive improvements, e.g., Indyk et al. (1997), Das-

gupta and Gupta (2003), Achlioptas (2001), Ailon and Chazelle (2009), Matoušek (2008), Krahmer and Ward (2011), and Kane and Nelson (2014). These address the issues of *efficient* construction and implementation, using random matrices that support fast multiplication. Dirksen (2016) gave a unified theory for dimensionality reduction using subgaussian matrices.

In a different direction, variants of the Johnson-Lindenstrauss lemma giving embeddings into spaces of lower dimension than the JL bound have been given under several specific settings. For point sets lying in regions of bounded Gaussian width, a theorem of Gordon (1988) implies that the dimension of the embedding can be reduced to a function of the Gaussian width, independent of the number of points. Sarlós (2006) showed that points lying on a d-flat can be mapped to $O(d/\varepsilon^2)$ dimensions independently of the number of points. Baraniuk and Wakin (2009) proved an analogous result for points on a smooth submanifold of Euclidean space, which was subsequently sharpened by Clarkson (2008) (see also Verma (2011)), whose version directly preserves geodesic distances on the submanifold. Other related results include those of Clarkson (2008) for sets of bounded doubling dimension and Alon and Klartag (2017) for general inner products, with additive error only. Recently, Narayanan and Nelson (2019), building on earlier results (Elkin et al. 2017; Mahabadi et al. 2018), showed that for a given set of points or *terminals*, using just one extra dimension from the Johnson-Lindenstrauss bound, it is possible to achieve dimensionality reduction in a way that preserves not only inter-terminal distances, but also distances between any terminal to any point in the ambient space.

Remark 1 Our results are based on the notion of weighted points, and as in most applications of the JL lemma, give a reduced dimensionality typically of the order of hundreds. This is very useful if the ambient dimensionality is much higher magnitude (e.g. 10^6). Moreover, some of the above-mentioned variants and generalizations such as for point sets having bounded Gaussian width or lying on a lower-dimensional submanifold, the reduced dimensionality is independent of the number of input points, which allows for still better reductions.

Dimension Reduction and Persistent Homology. The JL Lemma has also been used by Sheehy (2014) and Lotz (2019) to reduce the complexity of computing persistent homology. Both Sheehy and Lotz show that the persistent homology of a point cloud is approximately preserved under random projections (Sheehy 2014; Lotz 2019), up to a $(1 \pm \varepsilon)$ multiplicative factor, for any $\varepsilon \in [0, 1]$. Sheehy proves this for an *n*-point set, whereas Lotz's generalization applies to sets of bounded Gaussian width, and also implies dimensionality reductions for sets of bounded doubling dimension, in terms of the *spread* (ratio of the maximum to minimum interpoint distance). However, their techniques involve only the usual distance to a point set and therefore remain sensitive to outliers and noise as mentioned earlier. The question of adapting the method of random projections in order to reduce the complexity of computing persistent homology using the k-distance, is therefore a natural one, and has been raised by Sheehy (2014), who observed that "One notable distance function that is missing from this paper [i.e. (Sheehy 2014)] is the so-called distance to a measure or \ldots k-distance ... it remains open whether the k-distance itself is $(1 \pm \varepsilon)$ -preserved under random projection."

Our Contribution In this paper, we combine the method of random projections with the k-distance and show its applicability in computing persistent homology. It is not very hard to see that for a given point set P, the random Johnson-Lindenstrauss mapping preserves the pointwise k-distance to P (Theorem 17). However, this is not enough to preserve intersections of balls at varying scales of the radius parameter, and thus does not suffice to preserve the persistent homology of Čech filtrations, as noted by Sheehy (2014) and Lotz (2019). We show how the squared radius of a set of weighted points can be expressed as a convex combination of pairwise squared distances. From this, it follows that the Čech filtration under the k-distance, will be preserved by any linear mapping that preserves pairwise distances.

Extensions Further, as our main result applies to any linear mapping that approximately preserves pairwise distances, the analogous versions for bounded Gaussian width, points on submanifolds of \mathbb{R}^{D} , terminal dimensionality reduction and others apply immediately. Thus, we give several extensions of our results. The extensions provide bounds which do not depend on the number of points in the sample. The first one, analogous to Lotz (2019), shows that the persistent homology with respect to the k-distance, of point sets contained in regions having bounded Gaussian width, can be preserved via dimensionality reduction, using an embedding with dimension bounded by a function of the Gaussian width. Another result is that for points lying in a lowdimensional submanifold of a high-dimensional Euclidean space, the dimension of the embedding preserving the persistent homology with k-distance depends linearly on the dimension of the submanifold. Both these settings are commonly encountered in high-dimensional data analysis and machine learning (see, e.g., the *manifold* hypothesis Fefferman et al. 2016). We mention that analogous to Narayanan and Nelson (2019), it is possible to preserve the k-distance based persistent homology while also preserving the distance from any point in the ambient space to every point (i.e., terminal) in P (and therefore the k-distance to P), using just one extra dimension.

Run-time and Efficiency In many other applications of the Johnson-Lindenstrauss dimensionality reduction, multiplying by a dense gaussian matrix is a significant overhead, and can seriously affect any gains resulting from working in a lower dimensional space. However, as is pointed out in Lotz (2019), in the computation of persistent homology the dimensionality reduction step is carried out only once for the n data points at the beginning of the construction. Having said that, it should still be observed that most of the recent results on dimensionality reduction using sparse subgaussian matrices (Ailon and Chazelle 2009; Kane and Nelson 2014; Krahmer and Ward 2011) can also be used to compute the k-distance persistent homology, with little to no extra cost.

Remark 2 It should be noted that the approach of using dimensionality reduction for the *k*-distance, is complementary to denoising techniques such as Buchet et al. (2018) as we do not try to remove noise, only to be more robust to noise. Therefore, it can be used in conjunction with denoising techniques, as a pre-processing tool when the dimensionality is high.

Outline The rest of this paper is organized as follows. In Sect. 2, we briefly summarize some basic definitions and background. Our theorems are stated and proven in Sect. 3.

Some applications of our results are derived in Sect. 4. We end with some final remarks and open questions in Sect. 5.

2 Preliminaries

We need a well-known identity for the variance of bounded random variables, which will be crucial in the proof of our main theorem. A short probabilistic proof of (1) is given in the "Appendix". Let A be a set of points $p_1, \ldots, p_l \in \mathbb{R}^m$. A point $b \in \mathbb{R}^m$ is a *convex combination* of the points in A if there exist non-negative reals $\lambda_1, \ldots, \lambda_l \geq 0$ such that $b = \sum_{i=1}^{l} \lambda_i p_i$ and $\sum_{i=1}^{l} \lambda_i = 1$. Let $b = \sum_{i=1}^{k} \lambda_i p_i$ be a convex combination of points $p_1, \ldots, p_k \in \mathbb{R}^D$. Then for

any point $x \in \mathbb{R}^D$

$$\sum_{i=1}^{k} \lambda_i \|x - p_i\|^2 = \|x - b\|^2 + \sum_{i=1}^{k} \lambda_i \|b - p_i\|^2.$$
(1)

In particular, if $\lambda_i = 1/k$ for all *i*, we have

$$\frac{1}{k}\sum_{i=1}^{k}\|x-p_i\|^2 = \|x-b\|^2 + \sum_{i=1}^{k}\frac{1}{k}\|b-p_i\|^2.$$
(2)

2.1 The Johnson–Lindenstrauss Lemma

The Johnson-Lindenstrauss Lemma Johnson and Lindenstrauss (1984) states that any subset of *n* points of Euclidean space can be embedded in a space of dimension $O(\varepsilon^{-2} \log n)$ with $(1 \pm \varepsilon)$ distortion. We use the notion of an ε -distortion map with respect to P (also commonly called a Johnson-Lindenstrauss map).

Definition 1 Given a point set $P \subset \mathbb{R}^D$, and $\varepsilon \in (0, 1)$, a mapping $f : \mathbb{R}^D \to \mathbb{R}^d$ for some d < D is an ε -distortion map with respect to P, if for all $x, y \in P$,

$$(1 - \varepsilon) \|x - y\| \le \|f(x) - f(y)\| \le (1 + \varepsilon) \|x - y\|.$$

A random variable X with mean zero is said to be *subgaussian* with *subgaussian* norm K if $\mathbb{E}\left[\exp\left(X^2/K^2\right)\right] \le 2$. In this case, the tails of the random variable satisfy

$$\mathbb{P}\left[|X| \ge t\right] \le 2\exp\left(-t^2/2K^2\right).$$

We focus on the case where the Johnson-Lindenstrauss embedding is carried out via random subgaussian matrices, i.e., matrices where for some given K > 0, each entry is an independent subgaussian random variable with subgaussian norm K. This case is general enough to include the mappings of Achlioptas (2001), Ailon and Chazelle (2009), Dasgupta and Gupta (2003), Indyk et al. (1997), and Matoušek (2008) (see Dirksen for a unified treatment Dirksen 2016).

Lemma 2 (JL Lemma) Given $0 < \varepsilon, \delta < 1$, and a finite point set $P \subset \mathbb{R}^D$ of size |P| = n. Then a random linear mapping $f : \mathbb{R}^D \to \mathbb{R}^d$ where $d = O(\varepsilon^{-2} \log n)$ given by $f(v) = \sqrt{\frac{D}{d}} Gv$ where G is a $d \times D$ subgaussian random matrix, is an ε -distortion map with respect to P, with probability at least $1 - \delta$.

Definition 3 For ease of recall, we shall refer to a random linear mapping $f : \mathbb{R}^D \to \mathbb{R}^d$ given by $f(v) = \sqrt{\frac{D}{d}} Gv$ where G is a $d \times D$ subgaussian random matrix, as a subgaussian ε -distortion map.

While in the version given here the dimension of the embedding depends on the number of points in P and subgaussian projections, the JL lemma has been generalized and extended in several different directions, some of which are briefly outlined below. The generalization of the results of this paper to these more general settings is straightforward.

Sets of Bounded Gaussian Width

Definition 4 Given a set $S \subset \mathbb{R}^D$, the *Gaussian width* of S is

$$w(S) := \mathbb{E}\left[\sup_{x \in S} \langle x, g \rangle\right],$$

where $g \in \mathbb{R}^{D}$ is a random standard *D*-dimensional Gaussian vector.

In several areas like geometric functional analysis, compressed sensing, machine learning, etc. the Gaussian width is a very useful measure of the width of a set in Euclidean space (see e.g. Foucart and Rauhut (2013) and the references therein). It is also closely related to the *statistical dimension* of a set (see e.g. (Vershynin 2018, Chapter 7). The following analogue of the Johnson Lindenstrauss lemma for sets of bounded Gaussian width was given in Lotz (2019). It essentially follows from a result of Gordon (1988).

Theorem 5 (Lotz 2019, Theorem 3.1) Given ε , $\delta \in (0, 1)$, $P \subset \mathbb{R}^D$, let $S := \{(x-y)/||x-y|| : x, y \in P\}$. Then for any $d \ge \frac{\left(w(S)+\sqrt{2\log(2/\delta)}\right)^2}{\varepsilon^2} + 1$, the function $f : \mathbb{R}^D \to \mathbb{R}^d$ given by $f(x) = (\sqrt{D/d}) Gx$, where G is a random $d \times D$ Gaussian matrix G, is a subgaussian ε -distortion map with respect to P, with probability at least $1 - \delta$.

The result extends to subgaussian matrices with slightly worse constants. One of the benefits of this version is that the set P does not need to be finite. We refer to Lotz (2019) for more on the Gaussian width in our context.

Submanifolds of Euclidean Space For point sets lying on a low-dimensional submanifold of a high-dimensional Euclidean space, one can obtain an embedding with a smaller dimension using the bounds of Baraniuk and Wakin (2009) or Clarkson (2008), which will depend only on the parameters of the submanifold.

Clarkson's theorem is summarised below.

Theorem 6 (Clarkson 2008) There exists an absolute constant c > 0 such that, given a connected, compact, orientable, differentiable μ -dimensional submanifold $M \subset \mathbb{R}^D$, and $\varepsilon, \delta \in (0, 1)$, a random projection map $f : \mathbb{R}^D \to \mathbb{R}^d$, given by $v \mapsto \sqrt{\frac{D}{d}} Gv$, where G is a $d \times D$ subgaussian random matrix, is an ε -distortion map with respect to P, with probability at least $1 - \delta$, for

$$d \ge c \left(\frac{\mu \log(1/\varepsilon) + \log(1/\delta)}{\varepsilon^2} + \frac{C(M)}{\varepsilon^2} \right),$$

where C(M) depends only on M.

Terminal Dimensionality Reduction In a recent breakthrough result, Narayanan and Nelson (2019) showed that it is possible to $(1 \pm O(\varepsilon))$ -preserve distances from a set of *n terminals* in a high-dimensional space to *every point* in the space, using only one dimension more than the Johnson-Lindenstrauss bound.

A summarized version of their theorem is as follows. The derivation of the second statement is given in the "Appendix".

Theorem 7 (Narayanan and Nelson 2019, Theorem 3.2, Lemma 3.2) Given terminals $x_1, \ldots, x_n \in \mathbb{R}^D$ and $\varepsilon \in (0, 1)$, there exists a non-linear map $f : \mathbb{R}^D \to \mathbb{R}^{d'}$ with d' = d + 1, where $d = O\left(\frac{\log n}{\varepsilon^2}\right)$ is the bound given in Lemma 2, such that f is an ε -distortion map for any pairwise distance between $x_i, x_j \in P$, and an $O(\varepsilon)$ -distortion map for the distances between any pairs of points (x, u), where $x \in P$ and $u \in \mathbb{R}^D$. Further, the projection of f to its first d - 1 coordinates is a subgaussian ε -distortion map.

As noted in Narayanan and Nelson (2019), any such map must necessarily be nonlinear. Suppose not, then on translating the origin to be a terminal, it follows that the Euclidean norm of each point on the unit sphere around the origin must be $O(\varepsilon)$ preserved, which means that the dimension of any embedding given by a linear map would not be any less than the original dimension.

2.2 k-distance

The distance to a finite point set P is usually taken to be the minimum distance to a point in the set. For the computations involved in geometric and topological inference, however, this distance is highly sensitive to outliers and noise. To handle this problem of sensitivity, Chazal et al. in Chazal et al. (2011) introduced the *distance to a probability measure* which, in the case of a uniform probability on P, is called the *k-distance*.

Definition 8 (*k*-distance) For $k \in \{1, ..., n\}$ and $x \in \mathbb{R}^D$, the *k*-distance of *x* to *P* is

$$d_{P,k}(x) = \min_{S_k \in \binom{P}{k}} \sqrt{\frac{1}{k} \sum_{p \in S_k} \|x - p\|^2} = \sqrt{\frac{1}{k} \sum_{p \in \mathrm{NN}_P^k(x)} \|x - p\|^2}$$
(3)

where $NN_P^k(x) \subset P$ denotes the *k* nearest neighbours in *P* to the point $x \in \mathbb{R}^D$.

It was shown in Aurenhammer (1990), that the k-distance can be expressed in terms of weighted points and *power distance*. A weighted point \hat{p} is a point p of \mathbb{R}^{D} together with a (not necessarily positive) real number called its weight and denoted by w(p). The *power distance* between a point $x \in \mathbb{R}^{D}$ and a weighted point $\hat{p} = (p, w(p))$, denoted by $D(x, \hat{p})$ is $||x - p||^2 - w(p)$, i.e. the power of x with respect to a ball of radius $\sqrt{w(p)}$ centered at p. The distance between two weighted points $\hat{p}_i =$ $(p_i, w(i))$ and $\hat{p}_j = (p_j, w(j))$ is defined as $D(\hat{p}_i, \hat{p}_j) = ||p_i - p_j||^2 - w(i) - w(j)$. This definition encompasses the case where the two weights are 0, in which case we have the squared Euclidean distance, and the case where one of the points has weight 0, in which case, we have the power distance of a point to a ball. We say that two weighted points are *orthogonal* when their weighted distance is zero.

Let $B_{P,k}$ be the set of iso-barycentres of all subsets of k points in P. To each barycenter $b \in B_{P,k}$, $b = (1/k) \sum_i p_i$, we associate the weight $w(b) = -\frac{1}{k} \sum_i ||b - p_i||^2$. Note that, despite the notation, this weight does not only depend on b, but also on the set of points in P for which b is the barycenter. Writing $\hat{B}_{P,k} = \{\hat{b} = (b, w(b)), b \in B_{P,k}\}$, we see from (2) that the k-distance is the square root of a power distance (Aurenhammer 1990)

$$d_{P,k}(x) = \min_{\hat{b} \in \hat{B}_{P,k}} \sqrt{D(x, \hat{b})}.$$
 (4)

Observe that in general the squared distance between a pair of weighted points can be negative, but the above assignment of weights ensures that the *k*-distance $d_{P,k}$ is a real function. Since $d_{P,k}$ is the square root of a non-negative power distance, the α -sublevel set of $d_{P,k}$, $d_{P,k}([-\infty, \alpha])$, $\alpha \in \mathbb{R}$, is the union of $\binom{n}{k}$ balls $B(b, \sqrt{\alpha^2 + w(b)})$, $b \in B_{P,k}$. However, some of the balls may be included in the union of others and be redundant. In fact, the number of barycenters (or equivalently of balls) required to define a level set of $d_{P,k}$ is equal to the number of the non-empty cells in the *k*thorder Voronoi diagram of *P*. Hence the number of non-empty cells is $\Omega\left(n^{\lfloor (D+1)/2 \rfloor}\right)$ (Clarkson and Shor 1989) and computing them in high dimensions is intractable. It is then natural to look for approximations of the *k*-distance, as proposed in Buchet et al. (2016).

Definition 9 (Approximation) Let $P \subset \mathbb{R}^D$ and $x \in \mathbb{R}^D$. The approximate *k*-distance $\tilde{d}_{P,k}(x)$ is defined as

$$\tilde{d}_{P,k}(x) := \min_{p \in P} \sqrt{D(x, \hat{p})}$$
(5)

where $\hat{p} = (p, w(p))$ with $w(p) = -d_{P,k}^2(p)$, the negative of the squared *k*-distance of *p*.

In other words, we replace the set of barycenters with *P*. As in the exact case, $\tilde{d}_{P,k}$ is the square root of a power distance and its α -sublevel set, $\alpha \in \mathbb{R}$, is a union of balls, specifically the balls $B(p, \sqrt{\alpha^2 - d_{P,k}^2(p)}), p \in P$. The major difference with the exact case is that, since we consider only balls around the points of *P*, their number is *n* instead of $\binom{n}{k}$ in the exact case (compare Eq. (5) and Eq. (4)). Still, $\tilde{d}_{P,k}(x)$ approximates the *k*-distance (Buchet et al. 2016):

$$\frac{1}{\sqrt{2}} d_{P,k} \le \tilde{d}_{P,k} \le \sqrt{3} d_{P,k}.$$
(6)

We now make an observation for the case when the weighted points are barycenters, which will be useful in proving our main theorem.

Lemma 10 If $b_1, b_2 \in B_{P,k}$, and $p_{i,1}, \ldots, p_{i,k} \in P$ for i = 1, 2, such that $b_i = \frac{1}{k} \sum_{l=1}^{k} p_{i,l}$, and $w(b_i) = \frac{1}{k} \sum_{l=1}^{k} ||b_i - p_{i,l}||^2$ for i = 1, 2, then

$$D(\hat{b}_1, \hat{b}_2) = \frac{1}{k^2} \sum_{l,s=1}^k \|p_{1,l} - p_{2,s}\|^2.$$

Proof We have

$$D(\hat{b}_1, \hat{b}_2) = \|b_1 - b_2\|^2 - w(b_1) - w(b_2)$$

= $\|b_1 - b_2\|^2 + \frac{1}{k} \sum_{l=1}^k \|b_1 - p_{1,l}\|^2 + \frac{1}{k} \sum_{l=1}^k \|b_2 - p_{2,l}\|^2.$

Applying the identity (2), we get $||b_1 - b_2||^2 + \frac{1}{k} \sum_{l=1}^k ||b_2 - p_{2,l}||^2 = \frac{1}{k} \sum_{l=1}^k ||b_1 - p_{2,l}||^2$, so that

$$D(\hat{b}_{1}, \hat{b}_{2}) = \frac{1}{k} \sum_{l=1}^{k} \|b_{1} - p_{2,l}\|^{2} + \frac{1}{k} \sum_{l=1}^{k} \|b_{1} - p_{1,l}\|^{2}$$

$$= \frac{1}{k} \sum_{l=1}^{k} \|b_{1} - p_{2,l}\|^{2} + \frac{1}{k^{2}} \sum_{s=1}^{k} \sum_{l=1}^{k} \|b_{1} - p_{1,l}\|^{2}$$

$$= \frac{1}{k} \sum_{l=1}^{k} \left(\|b_{1} - p_{2,l}\|^{2} + \frac{1}{k} \sum_{s=1}^{k} \|b_{1} - p_{1,s}\|^{2} \right)$$

$$= \frac{1}{k} \sum_{l=1}^{k} \left(\frac{1}{k} \sum_{s=1}^{k} \|p_{1,s} - p_{2,l}\|^{2} \right) = \frac{1}{k^{2}} \sum_{l,s=1}^{k} \|p_{1,s} - p_{2,l}\|^{2}, \quad (7)$$

where in (7), we again applied (2) to each of the points $p_{2,s}$, with respect to the barycenter b_1 .

2.3 Persistent homology

Simplicial Complexes and Filtrations Let V be a finite set. An (abstract) simplicial complex with vertex set V is a set K of finite subsets of V such that if $A \in K$ and $B \subseteq A$,

then $B \in K$. The sets in K are called the simplices of K. A simplex $F \in K$ that is strictly contained in a simplex $A \in K$, is said to be a *face* of A.

A simplicial complex *K* with a function $f : K \to \mathbb{R}$ such that $f(\sigma) \leq f(\tau)$ whenever σ is a face of τ is a filtered simplicial complex. The sublevel set of *f* at $r \in \mathbb{R}$, $f^{-1}(-\infty, r]$, is a subcomplex of *K*. By considering different values of *r*, we get a nested sequence of subcomplexes (called a filtration) of $K, \emptyset = K^0 \subseteq K^1 \subseteq$ $\dots \subseteq K^m = K$, where K^i is the sublevel set at value r_i .

The Čech filtration associated to a finite set P of points in \mathbb{R}^D plays an important role in Topological Data Analysis.

Definition 11 (Čech Complex) The Čech complex $\check{C}_{\alpha}(P)$ is the set of simplices $\sigma \subset P$ such that $\operatorname{rad}(\sigma) \leq \alpha$, where $\operatorname{rad}(\sigma)$ is the radius of the smallest enclosing ball of σ , i.e.

$$\operatorname{rad}(\sigma) = \min_{x \in \mathbb{R}^D} \max_{p_i \in \sigma} \|x - p_i\|.$$

When the threshold α goes from 0 to $+\infty$, we obtain the Čech filtration of P. $\check{C}_{\alpha}(P)$ can be equivalently defined as the nerve of the closed balls $\overline{B}(p, \alpha)$, centered at the points in P and of radius α :

$$\check{C}_{\alpha}(P) = \{ \sigma \subset P \mid \cap_{p \in \sigma} \overline{B}(p, \alpha) \neq \emptyset \}.$$

By the nerve lemma, we know that the union of balls $U_{\alpha} = \bigcup_{p \in P} \overline{B}(p, \alpha)$, and $\check{C}_{\alpha}(P)$ have the same homotopy type.

Persistence Diagrams Persistent homology is a means to compute and record the changes in the topology of the filtered complexes as the parameter α increases from zero to infinity. Edelsbrunner et al. (2002) gave an algorithm to compute the persistent homology, which takes a filtered simplicial complex as input, and outputs a sequence $(\alpha_{birth}, \alpha_{death})$ of pairs of real numbers. Each such pair corresponds to a topological feature, and records the values of α at which the feature appears and disappears, respectively, in the filtration. Thus the topological features of the filtration can be represented using this sequence of pairs, which can be represented either as points in the extended plane $\mathbb{R}^2 = (\mathbb{R} \cup \{-\infty, \infty\})^2$, called the *persistence diagram*, or as a sequence of barcodes (the *persistence barcode*) (see, e.g., Edelsbrunner and Harer (2010)). A pair of persistence diagrams \mathbb{G} and \mathbb{H} corresponding to the filtrations (G_{α}) and (H_{α}) respectively, are *multiplicatively* β -*interleaved*, ($\beta \geq 1$), if for all α , we have that $G_{\alpha/\beta} \subseteq H_{\alpha} \subseteq G_{\alpha\beta}$. We shall crucially rely on the fact that a given persistence diagram is closely approximated by another one if they are multiplicatively *c*-interleaved, with *c* close to 1 (see e.g. Chazal et al. (2016)).

The Persistent Nerve Lemma (Chazal and Oudot 2008) shows that the persistent homology of the Čech complex is the same as the homology of the α -sublevel filtration of the distance function.

The Weighted Case Our goal is to extend the above definitions and results to the case of the *k*-distance. As we observed earlier, the *k*-distance is a power distance in disguise.

Accordingly, we need to extend the definition of the Čech complex to sets of weighted points.

Definition 12 (Weighted Čech Complex) Let $\hat{P} = \{\hat{p}_1, ..., \hat{p}_n\}$ be a set of weighted points, where $\hat{p}_i = (p_i, w(i))$. The α -Čech complex of $\hat{P}, \check{C}_{\alpha}(\hat{P})$, is the set of all simplices σ satisfying

$$\exists x, \ \forall p_i \in \sigma, \ \|x - p_i\|^2 \le w(i) + \alpha^2 \quad \Leftrightarrow \quad \exists x, \ \forall p_i \in \sigma, \ D(x, \ \hat{p}_i) \le \alpha^2.$$

In other words, the α -Čech complex of \hat{P} is the nerve of the closed balls $\overline{B}(p_i, r_i^2 = w(i) + \alpha^2)$, centered at the p_i and of squared radius $w(i) + \alpha^2$ (if negative, $\overline{B}(p_i, r_i^2)$ is imaginary).

The notions of weighted Čech filtrations and their persistent homology now follow naturally. Moreover, it follows from (4) that the Čech complex $\check{C}_{\alpha}(P)$ for the *k*-distance is identical to the weighted Čech complex $\check{C}_{\alpha}(\hat{B}_{P,k})$, where $\hat{B}_{P,k}$ is, as above, the set of iso-barycenters of all subsets of *k* points in *P*.

In the Euclidean case, we equivalently defined the α -Čech complex as the collection of simplices whose smallest enclosing balls have radius at most α . We can proceed similarly in the weighted case. Let $\hat{X} \subseteq \hat{P}$. We define the *squared radius of* \hat{X} as

$$\operatorname{rad}^{2}(\hat{X}) = \min_{x \in \mathbb{R}^{D}} \max_{\hat{p}_{i} \in \hat{X}} D(x, \hat{p}_{i}),$$

and the weighted center or simply the *center* of \hat{X} as the point, noted $c(\hat{X})$, where the minimum is reached.

Our goal is to show that preserving smallest enclosing balls in the weighted scenario under a given mapping, also preserves the persistent homology. Sheehy (2014) and Lotz (2019), proved this for the unweighted case. Their proofs also work for the weighted case but only under the assumption that the weights stay unchanged under the mapping. In our case however, the weights need to be recomputed in $f(\hat{P})$. We therefore need a version of (Lotz 2019, Lemma 2.2) for the weighted case which does not assume that the weights stay the same under f. This is Lemma 16, which follows at the end of this section. The following lemmas will be instrumental in proving Lemma 16 and in proving our main result. Let $\hat{X} \subseteq \hat{P}$ and assume without loss of generality that $\hat{X} = {\hat{p}_1, ..., \hat{p}_m}$, where $\hat{p}_i = (p_i, w(i))$.

Lemma 13 $c(\hat{X})$ and $rad(\hat{X})$ are uniquely defined.

Proof of Lemma 13 The proof follows from the convexity of D (see Lemma 10). Assume, for a contradiction, that there exists two centers c_0 and $c_1 \neq c_0$ for \hat{X} . For convenience, write $r = \operatorname{rad}(\hat{X})$. By the definition of the center of \hat{X} , we have

$$\begin{aligned} \exists \hat{p}_0, \forall \hat{p}_i : D(c_0, \hat{p}_i) \le D(c_0, \hat{p}_0) = \|c_0 - p_0\|^2 - w(0) = r^2 \\ \exists \hat{p}_1, \forall \hat{p}_i : D(c_1, \hat{p}_i) \le D(c_1, \hat{p}_1) = \|c_1 - p_1\|^2 - w(1) = r^2. \end{aligned}$$

Consider $D_{\lambda}(\hat{p}_i) = (1 - \lambda)D(c_0, \hat{p}_i) + \lambda D(c_1, \hat{p}_i)$ and write $c_{\lambda} = (1 - \lambda)c_0 + \lambda c_1$. For any $\lambda \in (0, 1)$, we have

$$D_{\lambda}(\hat{p}_{i}) = (1 - \lambda)D(c_{0}, \hat{p}_{i}) + \lambda D(c_{1}, \hat{p}_{i})$$

= $(1 - \lambda)(c_{0} - p_{i})^{2} + \lambda(c_{1} - p_{i})^{2} - w(i)$
= $D(c_{\lambda}, \hat{p}_{i}) - c_{\lambda}^{2} + (1 - \lambda)c_{0}^{2} + \lambda c_{1}^{2}$
= $D(c_{\lambda}, \hat{p}_{i}) + \lambda(1 - \lambda)(c_{0} - c_{1})^{2}$
> $D(c_{\lambda}, \hat{p}_{i}).$

Moreover, for any *i*,

$$D_{\lambda}(\hat{p}_i) = (1-\lambda)D(c_0, \hat{p}_i) + \lambda D(c_1, \hat{p}_i) \le r^2.$$

Thus, for any i and any $\lambda \in (0, 1)$, $D(c_{\lambda}, \hat{p}_i) < r^2$. Hence c_{λ} is a better center than c_0 and c_1 , and r is not the minimal possible value for rad (\hat{X}) . We have obtained a contradiction.

Lemma 14 Let I be the set of indices for which $D(c, \hat{p}_i) = \operatorname{rad}^2(\hat{X})$ and let $\hat{X}_I = \{\hat{p}_i, i \in I\}$. Then there exist $(\lambda_i > 0)_{i \in I}$ such that $c(\hat{X}) = \sum_{i \in I} \lambda_i p_i$ with $\sum_{i \in I} \lambda_i = 1$.

Proof of Lemma 14 We write for convenience $c = c(\hat{X})$ and $r = \operatorname{rad}(\hat{X})$ and prove that $c \in \operatorname{conv}(X_I)$ by contradiction. Let $c' \neq c$ be the point of $\operatorname{conv}(X_I)$ closest to c, and $\tilde{c} \neq c$ be a point on [cc']. Since $\|\tilde{c} - p_i\| < \|c - p_i\|$ for all $i \in I$, $D(\tilde{c}, \hat{p}_i) < D(c, \hat{p}_i)$ for all $i \in I$. For \tilde{c} sufficiently close to c, \tilde{c} remains closer to the weighted points \hat{p}_j , $j \notin I$, than to the $\hat{p}_i, i \in I$. We thus have

$$D(\tilde{c}, \hat{p}_i) < D(\tilde{c}, \hat{p}_i) < D(c, \hat{p}_i) = r^2.$$

It follows that *c* is not the center of \hat{X} , a contradiction.

Combining the above results with (Lotz 2019, Lemma 4.2) gives the following lemma.

Lemma 15 Let I, $(\lambda_i)_{i \in I}$ be as in Lemma 14. Then the following holds.

$$\operatorname{rad}^{2}(\hat{X}) = \frac{1}{2} \sum_{i \in I} \sum_{j \in I} \lambda_{i} \lambda_{j} D(\hat{p}_{i}, \hat{p}_{j}).$$

Proof of Lemma 15 From Lemma 14, and writing $c = c(\hat{X})$ for convenience, we have

$$\operatorname{rad}^{2}(\hat{X}) = \sum_{i \in I} \lambda_{i} (\|c - p_{i}\|^{2} - w(i)).$$

We use the following simple fact from (Lotz 2019, Lemma 4.5) (a probabilistic proof is included in the "Appendix", Lemma 25).

$$\sum_{i \in I} \lambda_i \|c - p_i\|^2 = \frac{1}{2} \sum_{i \in I} \sum_{j \in I} \lambda_i \lambda_j \|p_i - p_j\|^2.$$

D Springer

Substituting in the expression for $rad^2(\hat{X})$,

$$\operatorname{rad}^{2}(\hat{X}) = \frac{1}{2} \sum_{j \in I} \sum_{i \in I} \lambda_{j} \lambda_{i} ||p_{i} - p_{j}||^{2} - \frac{1}{2} \sum_{i \in I} 2\lambda_{i} w(i)$$

$$= \frac{1}{2} \sum_{i, j \in I} \lambda_{j} \lambda_{i} ||p_{i} - p_{j}||^{2} - \frac{1}{2} \sum_{i, j \in I} 2\lambda_{i} \lambda_{j} w(i) \quad (\text{since } \sum_{j \in I} \lambda_{j} = 1)$$

$$= \frac{1}{2} \sum_{i, j \in I} \lambda_{j} \lambda_{i} ||p_{i} - p_{j}||^{2} - \frac{1}{2} \sum_{i, j \in I} \lambda_{i} \lambda_{j} (w(i) + w(j))$$

$$= \frac{1}{2} \sum_{i, j \in I} \lambda_{i} \lambda_{j} (||p_{i} - p_{j}||^{2} - w(i) - w(j)) = \frac{1}{2} \sum_{i, j \in I} \lambda_{i} \lambda_{j} D(\hat{p}_{i}, \hat{p}_{j}).$$

Let $X \in \mathbb{R}^D$ be a finite set of points and \hat{X} be the associated weighted points where the weights are computed according to a weighting function $w : X \to \mathbb{R}^-$. Given a mapping $f : \mathbb{R}^D \to \mathbb{R}^d$, we define $\widehat{f(X)}$ as the set of weighted points $\{(f(x), w(f(x))), x \in X\}$. Note that the weights are recomputed in the image space \mathbb{R}^d .

Lemma 16 In the above setting, if f is such that for some $\varepsilon \in (0, 1)$ and for all subsets $\hat{S} \subseteq \hat{X}$ we have

$$(1 - \varepsilon) \operatorname{rad}^2(\hat{S}) \le \operatorname{rad}^2(\widehat{f(S)}) \le (1 + \varepsilon) \operatorname{rad}^2(\hat{S}),$$

then the weighted Čech filtrations of \hat{X} and $f(\hat{X})$ are multiplicatively $(1 - \varepsilon)^{-1/2}$ interleaved.

3 *ɛ*-distortion maps preserve *k*-distance Čech filtrations

For the subsequent theorems, we denote by *P* a set of *n* points in \mathbb{R}^D .

Our first theorem shows that for the points in *P*, the pointwise *k*-distance $d_{P,k}$ is approximately preserved by a random subgaussian matrix satisfying Lemma 2.

Theorem 17 Given $\varepsilon \in (0, 1]$, any ε -distortion map with respect to $P f : \mathbb{R}^D \to \mathbb{R}^d$, where $d = O(\varepsilon^{-2} \log n)$ satisfies for all points $x \in P$:

$$(1-\varepsilon)d_{P,k}^2(x) \le d_{f(P),k}^2(f(x)) \le (1+\varepsilon)d_{P,k}^2(x).$$

Proof of Theorem 17 The proof follows from the observation that the squared *k*-distance from any point $p \in P$ to the set *P*, is a convex combination of the squares of the Euclidean distances to the *k* nearest neighbours of *p*. Since the mapping in the JL Lemma 2 is linear and $(1 \pm \varepsilon)$ -preserves squared pairwise distances, their convex combinations also get $(1 \pm \varepsilon)$ -preserved.

As mentioned previously, the preservation of the pointwise *k*-distance does not imply the preservation of the Čech complex formed using the points in *P*. Nevertheless, the following theorem shows that this can always be done in dimension $O(\log n/\varepsilon^2)$.

Let $\hat{B}_{P,k}$ be the set of iso-barycenters of every *k*-subset of *P*, weighted as in Sect. 2.2. Recall from Sect. 2.3 that the weighted Čech complex $\check{C}_{\alpha}(\hat{B}_{P,k})$ is identical to the Čech complex $\check{C}_{\alpha}(P)$ for the *k*-distance. We now want to apply Lemma 16, for which the following theorem will be needed.

Theorem 18 (*k*-distance) Let $\hat{\sigma} \subseteq \hat{B}_{P,k}$ be a simplex in the weighted Čech complex $\check{C}_{\alpha}(\hat{B}_{P,k})$. Then, given $d \leq D$ such that there exists a ε -distortion map $f : \mathbb{R}^D \to \mathbb{R}^d$ with respect to P, it holds that

$$(1-\varepsilon)\operatorname{rad}^2(\hat{\sigma}) \le \operatorname{rad}^2(\widehat{f(\sigma)}) \le (1+\varepsilon)\operatorname{rad}^2(\hat{\sigma}).$$

Proof of Theorem 18 Let $\hat{\sigma} = {\hat{b}_1, \hat{b}_2, ..., \hat{b}_m}$, where \hat{b}_i is the weighted point defined in Sect. 2.3, i.e. $\hat{b}_i = (b_i, w(b_i))$ with $b_i \in B_{P,k}$ and $w(b_i) = -\frac{1}{k} \sum_{l=1}^k ||b_i - p_{il}||^2$, where $p_{i,1}, ..., p_{i,k} \in P$ are such that $b_i = \frac{1}{k} \sum_{j=1}^k p_{i,j}$. Applying Lemma 15 to $\hat{\sigma}$, we have that

$$\operatorname{rad}^{2}(\hat{\sigma}) = \frac{1}{2} \sum_{i,j \in I} \lambda_{i} \lambda_{j} D(\hat{b}_{i}, \hat{b}_{j}).$$
(8)

By Lemma 10, the distance between \hat{b}_i and \hat{b}_j is $D(\hat{b}_i, \hat{b}_j) = \frac{1}{k^2} \sum_{l,s=1}^k ||p_{i,l} - p_{j,s}||^2$. As this last expression is a convex combination of squared pairwise distances of points in *P*, it is $(1 \pm \varepsilon)$ -preserved by any ε -distortion map with respect to *P*, which implies that the convex combination $\operatorname{rad}^2(\hat{\sigma}) = \frac{1}{2} \sum_{i,j \in I} \lambda_i \lambda_j D(\hat{p}_i, \hat{p}_j)$ corresponding to the squared radius of σ in \mathbb{R}^D , will be $(1 \pm \varepsilon)$ -preserved.

Let $f : \mathbb{R}^D \to \mathbb{R}^d$ be an ε -distortion map with respect to P, from \mathbb{R}^D to \mathbb{R}^d , where d will be chosen later. By Lemma 15, the centre of $\widehat{f(\sigma)}$ is a convex combination of the points $(f(b_i))_{i=1}^m$. Let the centre $c(\widehat{f(\sigma)})$ be given by $c(\widehat{f(\sigma)}) = \sum_{i \in I} v_i D(\widehat{f(b_i)})$. where for $i \in I$, $v_i \ge 0$, $\sum_i v_i = 1$. Consider the convex combination of power distances $\sum_{i,j \in I} v_i v_j D(\widehat{b}_i, \widehat{b}_j)$. Since f is an ε -distortion map with respect to P, by Lemmas 10 and 2 we get

$$\frac{1}{2}(1-\varepsilon)\sum_{i,j\in I}\nu_i\nu_j D(\hat{b}_i,\hat{b}_j) \le \frac{1}{2}\sum_{i,j\in I}\nu_i\nu_j D(\widehat{f(b_i)},\widehat{f(b_j)}) = \operatorname{rad}^2(\widehat{f(\sigma)}).$$
(9)

On the other hand, since the squared radius is a minimizing function by definition, we get that

$$\operatorname{rad}^{2}(\hat{\sigma}) = \frac{1}{2} \sum_{i,j \in I} \lambda_{i} \lambda_{j} D(\hat{b}_{i}, \hat{b}_{j}) \leq \frac{1}{2} \sum_{i,j \in I} \nu_{i} \nu_{j} D(\hat{b}_{i}, \hat{b}_{j})$$

$$\leq \frac{1}{(1-\varepsilon)} \operatorname{rad}^{2}(f(\sigma)), \text{ by } (9).$$
(10)

🖉 Springer

$$\operatorname{rad}^{2}(\widehat{f(\sigma)}) = \frac{1}{2} \sum_{i,j \in I} \nu_{i} \nu_{j} D(\widehat{f(b_{i})}, \widehat{f(b_{j})}) \leq \frac{1}{2} \sum_{i,j \in I} \lambda_{i} \lambda_{j} D(\widehat{f(b_{i})}, \widehat{f(b_{j})}).$$
(11)

Combining the inequalities (9), (10), (11) gives

$$(1-\varepsilon)\operatorname{rad}^{2}(\widehat{\sigma}) \leq \operatorname{rad}^{2}(\widehat{f(\sigma)}) \leq \frac{1}{2} \sum_{i,j \in I} \lambda_{i} \lambda_{j} D(\widehat{f(b_{i})}, \widehat{f(b_{j})}) \leq (1+\varepsilon)\operatorname{rad}^{2}(\widehat{\sigma}).$$

where the final inequality follows by Lemma 2, since f is an ε -distortion map with respect to P. Thus, we have that

$$(1-\varepsilon)\operatorname{rad}^2(\hat{\sigma}) \leq \operatorname{rad}^2(\widehat{f(\sigma)}) \leq (1+\varepsilon)\operatorname{rad}^2(\hat{\sigma}),$$

which completes the proof of the theorem.

Theorem 19 (Approximate k-distance) Let \hat{P} be the weighted points associated with P, introduced in Definition 9 (Equ. 5). Let, in addition, $\hat{\sigma} \subseteq \hat{P}$ be a simplex in the associated weighted Čech complex $\check{C}_{\alpha}(\hat{P})$. Then an ε -distortion mapping with respect to P, $f : \mathbb{R}^D \to \mathbb{R}^d$ satisfies: $(1 - \varepsilon) \operatorname{rad}^2(\hat{\sigma}) \leq \operatorname{rad}^2(\widehat{f(\sigma)}) \leq (1 + \varepsilon) \operatorname{rad}^2(\hat{\sigma})$.

Proof of Theorem 19 Recall that, in Sect. 2.2, we defined the approximate *k*-distance to be $\tilde{d}_{P,k}(x) := \min_{p \in P} \sqrt{D(x, \hat{p})}$, where $\hat{p} = (p, w(p))$ is a weighted point, having weight $w(p) = -d_{P,k}^2(p)$. So, the Čech complex would be formed by the intersections of the balls around the weighted points in *P*. The proof follows on the lines of the proof of Theorem 18. Let $\hat{\sigma} = \{\hat{p}_1, \hat{p}_2, ..., \hat{p}_m\}$, where $\hat{p}_1, ..., \hat{p}_m$ are weighted points in \hat{P} , and let $c(\hat{\sigma})$ be the center of $\hat{\sigma}$. Applying again Lemma 15, we get

$$\operatorname{rad}^{2}(\hat{\sigma}) = \frac{1}{2} \sum_{i,j \in I} \lambda_{i} \lambda_{j} \|p_{i} - p_{j}\|^{2} + \sum_{i \in I} \lambda_{i} w(p_{i}) = \sum_{i,j \in I; i < j} \lambda_{i} \lambda_{j} \|p_{i} - p_{j}\|^{2} + \sum_{i \in I} \lambda_{i} w(p_{i}),$$

where $w(p) = d_{P,k}^2(p)$. In the second equality, we used the fact that the summand corresponding to a fixed pair of distinct indices i < j is being counted twice and that the contribution of the terms corresponding to indices i = j is zero. An ε -distortion map with respect to P preserves pairwise distances and the *k*-distance in dimension $O(\varepsilon^{-2} \log n)$. The result then follows as in the proof of Theorem 18.

Applying Lemma 16 to the theorems 18 and 19, we get the following corollary.

Corollary 20 The persistent homology for the Čech filtrations of P and its image f(P) under any ε -distortion mapping with respect to P, using the (i) exact k-distance, as well as the (ii) approximate k-distance, are preserved upto a multiplicative factor of $(1 - \varepsilon)^{-1/2}$.

However, note that the approximation in Corollary 20 (*ii*) is with respect to the *approximate* k-distance, which is itself an approximation of the k-distance by a distortion factor $3\sqrt{2}$, (i.e. bounded away from 1 - see(6)).

4 Extensions

As Theorem 18 applies to arbitrary ε -distortion maps, it naturally follows that many of the extensions and variants of the JL Lemma, e.g. discussed in Sect. 2.1, have their corresponding versions for the *k*-distance as well. In this section we elucidate some of the corresponding extensions of Theorem 18.

These can yield better bounds for the dimension of the embedding, stronger dimensionality reduction results, or easier to implement reductions in their respective settings.

The first result in this section, is for point sets contained in a region of bounded Gaussian width.

Theorem 21 Let $P \subset \mathbb{R}^D$ be a finite set of points, and define $S := \{(x - y)/||x - y|| : x, y \in P\}$. Let w(S) denote the Gaussian width of S. Then, given any $\varepsilon, \delta \in (0, 1)$, any subgaussian ε -distortion map from \mathbb{R}^D to \mathbb{R}^d preserves the persistent homology of the k-distance based Čech filtration associated to P, up to a multiplicative factor of $(1 - \varepsilon)^{-1/2}$, given that

$$d \ge \frac{\left(w(S) + \sqrt{2\log(2/\delta)}\right)^2}{\varepsilon^2} + 1.$$

Note that the above theorem is not stated for an arbitrary ε -distortion map. Also, since the Gaussian width of an *n*-point set is at most $O(\log n)$ (using e.g. the Gaussian concentration inequality, see e.g. (Boucheron et al. 2013, Sect. 2.5), Theorem 21 strictly generalizes Corollary 20.

Proof of Theorem 21 By Theorem 5, the scaled random Gaussian matrix $f : x \mapsto (\sqrt{D/d}) Gx$ is an ε -distortion map with respect to P, having dimension $d \ge \frac{(w(s)+\sqrt{2\log(2/\delta)})^2}{\varepsilon^2} + 1$. Now applying Theorem 18 to the point set P with the mapping f, immediately gives us that for any simplex $\hat{\sigma} \in \check{C}_{\alpha}(\hat{B}_{P,k})$, where $\check{C}_{\alpha}(\hat{B}_{P,k})$ is the weighted Čech complex with parameter α , the squared radius rad²($\hat{\sigma}$) is preserved up to a multiplicative factor of $(1 \pm \varepsilon)$. By Lemma 16, this implies that the persistent

homology for the Čech filtration is $(1 - \varepsilon)^{-1/2}$ -multiplicatively interleaved. For point sets lying on a low-dimensional submanifold of a high-dimensional Euclidean space, one can obtain an embedding having smaller dimension, using the bounds of Baraniuk and Wakin (2009) or Clarkson (2008), which will depend only on

the parameters of the submanifold.

Theorem 22 There exists an absolute constant c > 0 such that, given a finite point set *P* lying on a connected, compact, orientable, differentiable μ -dimensional submanifold $M \subset \mathbb{R}^D$, and $\varepsilon, \delta \in (0, 1)$, an ε -distortion map $f : \mathbb{R}^D \to \mathbb{R}^d$ preserves the persistent homology of the Čech filtration computed on *P*, using the *k*-distance, provided

$$d \ge c \left(\frac{\mu \log(1/\varepsilon) + \log(1/\delta)}{\varepsilon^2} + \frac{C(M)}{\varepsilon^2} \right),$$

where C(M) depends only on M.

Proof of Theorem. The proof follows directly, by applying the map in Clarkson's bound (Theorem 6) as the ε -distortion map in Theorem 18.

Next, we state the terminal dimensionality reduction version of Theorem 18. This is a useful result when we wish to preserve the distance (or k-distance) from any point in the ambient space, to the original point set.

Theorem 23 Let $P \in \mathbb{R}^D$ be a set of *n* points. Then, given any $\varepsilon \in (0, 1]$, there exists a map $f : \mathbb{R}^D \to \mathbb{R}^d$, where $d = O\left(\frac{\log n}{\varepsilon^2}\right)$, such that the persistent homology of the *k*-distance based Čech filtration associated to *P* is preserved up to a multiplicative factor of $(1 - \varepsilon)^{-1/2}$, and the *k*-distance of any point in \mathbb{R}^D to *P*, is preserved up to a $(1 \pm O(\varepsilon))$ factor.

Proof The second part of the theorem follows immediately by applying Theorem 7, with the point set *P* as the set of terminals. By Theorem 7 (*ii*), the dimensionality reduction map of Narayanan and Nelson (2019) is an outer extension of a subgaussian ε -distortion map $\Pi : \mathbb{R}^D \to \mathbb{R}^{d-1}$. Now applying Theorem 18 to Π gives the first part of the theorem.

5 Conclusion and future work

k-Distance Vietoris-Rips and Delaunay filtrations Since the Vietoris-Rips filtration (Oudot 2015, Chapter 4) depends only on pairwise distances, it follows from Theorem 17 that this filtration with *k*-distances, is preserved upto a multiplicative factor of $(1 - \varepsilon)^{-1/2}$, under a Johnson-Lindenstrauss mapping. Furthermore, the *k*-distance Delaunay and the Čech filtrations (Oudot 2015, Chapter 4) have the same persistent homology. Corollary 20 (*i*) therefore implies that the *k*-distance Delaunay filtration of a given finite point set *P* is also $(1 - \varepsilon)^{-1/2}$ -preserved under an ε -distortion map with respect to *P*. Thus, Corollary 20 (*ii*) apply also to the approximate *k*-distance Vietoris-Rips and *k*-distance Delaunay filtrations.

Kernels. Other distance functions defined using kernels have proved successful in overcoming issues due to outliers. Using a result analogous to Theorem 17, we can show that random projections preserve the persistent homology for kernels up to a $C(1 - \varepsilon)^{-1/2}$ factor where *C* is a constant. We don't know if we can make C = 1 as for the *k*-distance.

Acknowledgements The authors would like to thank the referees for their comments and observations, which helped add more consequences to the main results and also improved the quality of the paper.

Declarations

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

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit http://creativecommons.org/licenses/by/4.0/.

Appendix

The following elementary lemma gives identity (1).

Lemma 24 Let $b = \sum_{i=1}^{k} \lambda_i p_i$ be a convex combination of points p_1, \ldots, p_k . Then for any point $x \in \mathbb{R}^D$,

$$\sum_{i=1}^{k} \lambda_i \|x - p_i\|^2 = \|x - b\|^2 + \sum_{i=1}^{k} \lambda_i \|b - p_i\|^2.$$
(12)

Proof Recall the following fundamental relation between the variance and expectation of a random variable.

Let $X \in \mathbb{R}^D$ be a random variable bounded in ℓ_2 . Then, by one characterization of the variance,

$$\mathbb{E}\left[\|X\|^2\right] = \sum_i \operatorname{Var}(X) + \|\mathbb{E}[X]\|^2.$$
(13)

Consider a point $x \in \mathbb{R}^D$, a set $P \subset \mathbb{R}^D$ of *k* points, and a probability distribution $\{\lambda_i\}_{i=1}^k$, along with a weighted sum $b = \sum_{i=1}^k \lambda_i p_i$. The random vector *Y* supported on *P*, with probability $\mathbb{P}[Y = p_i] = \lambda_i$, then satisfies $\mathbb{E}[Y] = b$. Define X := x - Y, so that $\mathbb{E}[X] = x - b$. Then

$$\mathbb{E}\left[\|X\|^2\right] = \sum_{i=1}^k \lambda_i \|x - p_i\|^2, \text{ and}$$
$$\sum_i \operatorname{Var}(X) = \mathbb{E}\left[\|X - \mathbb{E}[X]\|^2\right] = \mathbb{E}\left[\|Y - b\|^2\right] = \sum_{i=1}^k \lambda_i \|b - p_i\|^2.$$

Substituting in (13), the claim follows.

Lemma 25 (Lotz 2019, Lemma 4.2) Given a set $P = \{p_1, \ldots, p_l\} \subset \mathbb{R}^D$ of points, and a point $c \in \mathbb{R}^D$ such that $c = \sum_{i \in I} \lambda_i p_i$, where I is a subset of indices from [l], and $\lambda_i \geq 0$, with $\sum_{i \in I} \lambda_i = 1$. Then

$$\sum_{i \in I} \lambda_i \|c - p_i\|^2 = \frac{1}{2} \sum_{i, j \in I} \lambda_i \lambda_j \|p_i - p_j\|^2.$$

Deringer

Consequently,

$$\operatorname{rad}^{2}(P) = \sum_{i \in I} \lambda_{i} \|c - p_{i}\|^{2} = \frac{1}{2} \sum_{i, j \in I} \lambda_{i} \lambda_{j} \|p_{i} - p_{j}\|^{2}.$$

where $rad^{2}(P)$ is the squared radius of the minimum enclosing ball of the set P of points.

Proof The proof again follows directly from Eqn. (13). Suppose we choose two random points X_1 , X_2 independently from P, with the point p_i being chosen with probability λ_i . Then,

$$\mathbb{E}\left[\|X_1 - X_2\|^2\right] = \|\mathbb{E}\left[X_1 - X_2\right]\|^2 + Var(X_1 - X_2).$$

Evaluating, we get that

$$\mathbb{E}\left[\|X_1 - X_2\|^2\right] = \sum_{i,j \in I} \lambda_i \lambda_j \|p_i - p_j\|^2,$$

$$\|\mathbb{E}\left[X_1 - X_2\right]\|^2 = \|\mathbb{E}\left[X_1\right] - \mathbb{E}\left[X_2\right]\|^2 = 0, \text{ and}$$

$$Var(X_1 - X_2) = 2 \cdot Var(X_1) = 2 \sum_{i \in I} \lambda_i \|c - p_i\|^2,$$

where in the last line, we used the fact that X_1 and X_2 are independent. Substituting the above values in the variance identity (13), completes the proof.

A probabilistic proof of Lemma 10 is also provided below. Lemma 10 - a probabilistic proof. Consider the following random experiment: pick a random point X from p_1, \ldots, p_k according to the distribution $(\lambda_i)_{i=1}^k$ and another independently random point Y from q_1, \ldots, q_k according to $(\mu_i)_{i=1}^k$.

Using the law of total variance on the variable X - Y, conditioning on Y, we get that

$$Var(X - Y) = Var_Y(\mathbb{E}_X(X - Y|Y)) + \mathbb{E}_Y[Var_X(X - Y|Y)] \text{ or,}$$
$$\mathbb{E}\left[\|X - Y\|^2\right] = \|\mathbb{E}\left[X - Y\right]\|^2 + Var_Y(\mathbb{E}_X(X - Y|Y)) + \mathbb{E}_Y[Var_X(X - Y|Y)]$$
(14)

Let us consider the terms in the above equation one by one.

- 1. The LHS has $\mathbb{E}[||X Y||^2]$, which by the independence of X and Y is clearly equal to $\sum_{i,j=1}^k \lambda_i \mu_j ||p_i q_j||^2$.
- 2. In the RHS, the first term is $||\mathbb{E}[X Y]||^2 = ||b_1 b_2||^2$.
- 3. The second term is $Var_Y(\mathbb{E}_X(X Y)|Y)$, which is equal to $Var_Y(b_1 Y) = Var(Y) = \sum_{i=1}^k \mu_i ||b_2 q_i||^2$, where the last expression was evaluated directly from the definition of variance, i.e. $Var(Z) = \mathbb{E}[(Z \mathbb{E}[Z])^2]$, and that for constant *a*, Var(a Z) = Var(Z).

4. The final term is $\mathbb{E}_Y(Var_X(X-Y|Y))$. Conditioning on *Y*, the variance $Var_X(X-Y|Y) = Var(X)$, i.e. $\sum_{i=1}^k \lambda_i ||b_1 - p_i||^2$. Since this holds for each value of *Y*, we get that $\mathbb{E}_Y[Var_X(X-Y|Y)] = \mathbb{E}_Y[Var(X)] = \sum_{i=1}^k \lambda_i ||b_1 - p_i||^2$.

Substituting the above expressions for the terms in (14), we get

$$\sum_{i,j=1}^{k} \lambda_{i} \mu_{j} \|p_{i} - q_{j}\|^{2} = \|b_{1} - b_{2}\|^{2} + \sum_{i=1}^{k} \mu_{i} \|b_{2} - q_{i}\|^{2} + \sum_{i=1}^{k} \lambda_{i} \|b_{1} - p_{i}\|^{2}.$$

References

- Achlioptas, D.: Database-friendly random projections. In Buneman, P. (ed.) Proceedings of the Twentieth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems, May 21-23, 2001, Santa Barbara, California, USA. ACM, 2001. https://doi.org/10.1145/375551.375608,
- Ailon, N., Chazelle, B.: The fast Johnson-Lindenstrauss transform and approximate nearest neighbors. SIAM J. Comput. 39(1), 302–322 (2009). https://doi.org/10.1137/060673096
- Alon, N., Klartag, B.: Optimal compression of approximate inner products and dimension reduction. In Umans, C. (ed.) 58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017, Berkeley, CA, USA, October 15-17, 2017, pp. 639–650. IEEE Computer Society, (2017). https://doi. org/10.1109/FOCS.2017.65,
- Aurenhammer, F.: A new duality result concerning Voronoi diagrams. Discret. Comput. Geom. 5, 243–254 (1990). https://doi.org/10.1007/BF02187788
- Baraniuk, R.G., Wakin, M.B.: Random projections of smooth manifolds. Found. Comput. Math. 9(1), 51–77 (2009). https://doi.org/10.1007/s10208-007-9011-z
- Boucheron, S., Lugosi, G., Massart, P.: Concentration Inequalities A Nonasymptotic Theory of Independence. Oxford University Press, Oxford (2013). https://doi.org/10.1093/acprof:oso/9780199535255. 001.0001
- Buchet, M., Chazal, F., Oudot, S.Y., Sheehy, D.R.: Efficient and robust persistent homology for measures. Comput. Geom. 58, 70–96 (2016). https://doi.org/10.1016/j.comgeo.2016.07.001
- Buchet, M., Dey, T.K., Wang, J., Wang, Y.: Declutter and resample: towards parameter free denoising. JoCG 9(2), 21–46 (2018). https://doi.org/10.20382/jocg.v9i2a3
- Chazal, F., Cohen-Steiner, D., Mérigot, Q.: Geometric inference for probability measures. Found. Comput. Math. 11(6), 733–751 (2011). https://doi.org/10.1007/s10208-011-9098-0
- Chazal, F., de Silva, V., Glisse, M., Oudot, S.Y.: The Structure and Stability of Persistence Modules. Springer Briefs in Mathematics. Springer, Berlin (2016). https://doi.org/10.1007/978-3-319-42545-0
- Chazal, F., Oudot, S.: Towards persistence-based reconstruction in Euclidean spaces. In Teillaud, M. (ed.) Proceedings of the 24th ACM Symposium on Computational Geometry, College Park, MD, USA, June 9–11, 2008, pp. 232–241. ACM, (2008). https://doi.org/10.1145/1377676.1377719
- Clarkson, K.L.: Tighter bounds for random projections of manifolds. In Teillaud, M. (ed.) Proceedings of the 24th ACM Symposium on Computational Geometry, College Park, MD, USA, June 9–11, 2008, pp. 39–48. ACM, (2008). https://doi.org/10.1145/1377676.1377685
- Clarkson, K.L., Shor, P.W.: Application of random sampling in computational geometry. II. Discret. Comput. Geom. 4, 387–421 (1989). https://doi.org/10.1007/BF02187740
- Dasgupta, S., Gupta, A.: An elementary proof of a theorem of Johnson and Lindenstrauss. Random Struct. Algorithms 22(1), 60–65 (2003). https://doi.org/10.1002/rsa.10073
- Dirksen, S.: Dimensionality reduction with subgaussian matrices: a unified theory. Found. Comput. Math. 16(5), 1367–1396 (2016). https://doi.org/10.1007/s10208-015-9280-x
- Edelsbrunner, H., Harer, J.: Computational Topology An Introduction. American Mathematical Society (2010). http://www.ams.org/bookstore-getitem/item=MBK-69
- Edelsbrunner, H., Letscher, D., Zomorodian, A.: Topological persistence and simplification. Discret. Comput. Geom. 28(4), 511–533 (2002). https://doi.org/10.1007/s00454-002-2885-2
- Elkin, M., Filtser, A., Neiman, O.: Terminal embeddings. Theor. Comput. Sci. 697, 1–36 (2017). https:// doi.org/10.1016/j.tcs.2017.06.021
- Fefferman, C., Mitter, S., Narayanan, H.: Testing the manifold hypothesis. J. Am. Math. Soc. 29(4), 983– 1049 (2016). https://doi.org/10.1090/jams/852

- Foucart, S., Rauhut, H.: A mathematical introduction to compressive sensing. Appl. Numer. Harmonic Anal. (2013). https://doi.org/10.1007/978-0-8176-4948-7
- Giraud, C.: Introduction to High-Dimensional Statistics. Chapman and Hall/CRC Monographs on Statistics and Applied Probability. Taylor and Francis, (2014). https://www.crcpress.com/Introduction-to-High-Dimensional-Statistics/Giraud/p/book/9781482237948
- Gordon, Y.: On Milman's inequality and random subspaces which escape through a mesh in Rⁿ. In Lindenstrauss, J., Milman, V.D. (eds.) *Geometric Aspects of Functional Analysis*, pp. 84–106, Berlin, Heidelberg (1988). https://doi.org/10.1007/BFb0081737
- Guibas, L.J., Morozov, D., Mérigot, Q.: Witnessed k-distance. Discret. Comput. Geom. 49(1), 22–45 (2013). https://doi.org/10.1007/s00454-012-9465-x
- Indyk P., Motwani, R., Raghavan, P., Vempala, S.S.: Locality-preserving hashing in multidimensional spaces. In Leighton, F.T., Shor, P.W. (eds.) Proceedings of the Twenty-Ninth Annual ACM Symposium on the Theory of Computing, El Paso, Texas, USA, May 4–6, 1997, pp. 618–625. ACM (1997). https:// doi.org/10.1145/258533.258656
- Johnson, W.B., Lindenstrauss, J.: Extensions of Lipschitz mappings into a Hilbert space. Contemp. Math. 26(189–206), 1 (1984). https://doi.org/10.1007/BF02764938
- Kane, D.M., Nelson, J.: Sparser Johnson–Lindenstrauss transforms. J. ACM 61(1), 4:1-4:23 (2014). https:// doi.org/10.1145/2559902
- Krahmer, F., Ward, R.: New and improved Johnson-Lindenstrauss embeddings via the restricted isometry property. SIAM J. Math. Anal. 43(3), 1269–1281 (2011). https://doi.org/10.1137/100810447
- Lotz, M.: Persistent homology for low-complexity models. Proc. R. Soc. A Math. Phys. Eng. Sci. 475(2230), 20190081 (2019). https://doi.org/10.1098/rspa.2019.0081
- Mahabadi, S., Makarychev, K., Makarychev, Y., Razenshteyn, I.P.: In Ilias, D., David, K., Monika, H. (eds.) Nonlinear dimension reduction via outer bi-lipschitz extensions, Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2018, Los Angeles, CA, USA, June 25–29, 2018, pp. 1088–1101. ACM (2018). https://doi.org/10.1145/3188745.3188828
- Matoušek, J.: On variants of the Johnson Lindenstrauss lemma. Random Struct. Algorithms 33(2), 142–156 (2008). https://doi.org/10.1002/rsa.20218
- Narayanan, S., Nelson, J.: Optimal terminal dimensionality reduction in Euclidean space. In Charikar, M., Cohen, E. (eds.) Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing, STOC 2019, Phoenix, AZ, USA, June 23–26, 2019, pages 1064–1069. ACM, (2019). https://doi.org/ 10.1145/3313276.3316307
- Oudot, S.Y.: Persistence Theory From Quiver Representations to Data Analysis. In Volume 209 of Mathematical Surveys and Monographs. American Mathematical Society (2015). http://bookstore.ams.org/ surv-209/
- Phillips, J.M., Wang, B., Zheng, Y.: Geometric inference on kernel density estimates. In Arge, L., Pach, J. (eds.) 31st International Symposium on Computational Geometry, SoCG 2015, June 22–25, 2015, Eindhoven, The Netherlands, volume 34 of LIPIcs, pp. 857–871. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, (2015). https://doi.org/10.4230/LIPIcs.SOCG.2015.857
- Sarlós, T.: In: Improved approximation algorithms for large matrices via random projections, In 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2006), 21–24 October 2006, Berkeley, California, USA, Proceedings, pp. 143–152. IEEE Computer Society (2006). https://doi.org/10.1109/ FOCS.2006.37
- Sheehy, D.R.: The persistent homology of distance functions under random projection. In Cheng, S.-W., Devillers, O. (eds.) 30th Annual Symposium on Computational Geometry, SOCG'14, Kyoto, Japan, June 08–11, 2014, p. 328. ACM, (2014). https://doi.org/10.1145/2582112.2582126
- Verma, N.: A note on random projections for preserving paths on a manifold. Technical report, UC San Diego (2011). https://csetechrep.ucsd.edu/Dienst/UI/2.0/Describe/ncstrl.ucsd_cse/CS2011-0971
- Vershynin, R.: High-Dimensional Probability: An Introduction with Applications in Data Science. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press (2018). https:// doi.org/10.1017/9781108231596
- Zhang, J.: Advancements of outlier detection: A survey. EAI Endorsed Trans. Scalable Inf. Syst. 1(1), e2 (2013). https://doi.org/10.4108/trans.sis.2013.01-03.e2

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