



HAL
open science

Convergence Rates for Persistence Diagram Estimation in Topological Data Analysis

Frédéric Chazal, Marc Glisse, Catherine Labruère, Bertrand Michel

► **To cite this version:**

Frédéric Chazal, Marc Glisse, Catherine Labruère, Bertrand Michel. Convergence Rates for Persistence Diagram Estimation in Topological Data Analysis. *Journal of Machine Learning Research*, 2015, 16, pp.3603-3635. <10.5555/2789272.2912112>. <hal-01284275>

HAL Id: hal-01284275

<https://hal.sorbonne-universite.fr/hal-01284275v1>

Submitted on 7 Mar 2016

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.



Distributed under a Creative Commons CC BY 4.0 - Attribution - International License

Convergence Rates for Persistence Diagram Estimation in Topological Data Analysis

Frédéric Chazal

FREDERIC.CHAZAL@INRIA.FR

Marc Glisse

MARC.GLISSE@INRIA.FR

*Inria Saclay - Île de France
1 rue Honoré d'Estienne d'Orves
Bâtiment Alan Turing
Campus de l'École Polytechnique
91120 Palaiseau, France*

Catherine Labruère

CLABRUE@U-BOURGOGNE.FR

*Université de Bourgogne
Institut de Mathématiques de Bourgogne, UMR CNRS 5584
9 Avenue Alain Savary, B.P. 47870
21078 Dijon, France*

Bertrand Michel

BERTRAND.MICHEL@UPMC.FR

*LSTA, Université Pierre et Marie Curie Paris 6
15-25, bureau 220
4, place Jussieu
75005 Paris, France*

Editor: Matthias Hein

Abstract

Computational topology has recently seen an important development toward data analysis, giving birth to the field of topological data analysis. Topological persistence, or persistent homology, appears as a fundamental tool in this field. In this paper, we study topological persistence in general metric spaces, with a statistical approach. We show that the use of persistent homology can be naturally considered in general statistical frameworks and that persistence diagrams can be used as statistics with interesting convergence properties. Some numerical experiments are performed in various contexts to illustrate our results.

Keywords: persistent homology, convergence rates, topological data analysis

1. Introduction

During the last decades, the wide availability of measurement devices and simulation tools has led to an explosion in the amount of available data in almost all domains of science, industry, economy and even everyday life. Often these data come as point clouds sampled in possibly high (or infinite) dimensional spaces. They are usually not uniformly distributed in the embedding space but carry some geometric structure (manifold or more general stratified space) which reflects important properties of the “systems” from which they have been generated. Moreover, in many cases data are not embedded in Euclidean spaces and come as (finite) sets of points with pairwise distance information. This often happens,

e.g. with social network or sensor network data where each sensor may not know its own position, but may evaluate its distance to the other sensors using the strength of the signal received from them. In such cases, data are given as matrices of pairwise distances between the observations, i.e. as (discrete) metric spaces. Again, although they come as abstract spaces, these data often carry specific topological and geometric structures.

1.1 Topological Data Analysis

A large amount of research has been done on dimensionality reduction, manifold learning and geometric inference for data embedded in Euclidean spaces and assumed to be concentrated around submanifolds; see for instance Wang (2012). However, the assumption that data lies on a manifold may fail in many applications. In addition, the strategy of representing data by points in Euclidean spaces may introduce large metric distortions as the data may lie in highly curved spaces. With the emergence of new geometric inference and algebraic topology tools, computational topology (Edelsbrunner and Harer, 2010) has recently seen an important development toward data analysis, giving birth to the field of Topological Data Analysis (TDA) (Carlsson, 2009) whose aim is to infer relevant, multiscale, qualitative and quantitative topological structures directly from the data. Topological persistence, more precisely *persistent homology* appears as a fundamental tool for TDA. Roughly, *homology* (with coefficient in a field such as, e.g., $\mathbb{Z}/2\mathbb{Z}$) associates to any topological space \mathbb{M} , a family of vector spaces (the so-called homology groups) $H_k(\mathbb{M})$, $k = 0, 1, \dots$, each of them encoding topological features of \mathbb{M} . The k^{th} Betti number of \mathbb{M} , denoted β_k , is the dimension of $H_k(\mathbb{M})$ and measures the number of k -dimensional features of \mathbb{M} : for example, β_0 is the number of connected components of \mathbb{M} , β_1 the number of independent cycles or “tunnels”, β_2 the number of “voids”, etc. (see Hatcher, 2001). Persistent homology provides a framework (Edelsbrunner et al., 2002; Zomorodian and Carlsson, 2005; Chazal et al., 2012a) and efficient algorithms to encode the evolution of the homology of families of nested topological spaces indexed by a set of real numbers that may often be seen as scales, such as the sublevel sets of a function, the union of growing balls, etc. The obtained multiscale topological information is then represented in a simple way as a barcode or persistence diagram; see Figure 4 and Section 2.3.

In TDA, persistent homology has found applications in many fields, including neuroscience (Singh et al., 2008), bioinformatics (Kasson et al., 2007), shape classification (Chazal et al., 2009b), clustering (Chazal et al., 2013), sensor networks (De Silva and Ghrist, 2007) or signal processing (Bauer et al., 2014). It is usually computed for a *filtered simplicial complex* built on top of the available data, i.e. a nested family of simplicial complexes whose vertex set is the data set (see Section 2.3). The obtained persistence diagrams are then used as “topological signatures” to exhibit and compare the topological structure underlying the data; see Figure 1. The relevance of this approach relies on stability results ensuring that close data sets, with respect to the Hausdorff or Gromov-Hausdorff distance, have close persistence diagrams (Cohen-Steiner et al., 2007; Chazal et al., 2009a, 2012a,b). However these results are not statistical and thus only provide heuristic or exploratory uses in data analysis.

The goal of this paper is to show that, thanks to recent results by Chazal et al. (2012a,b) that allow to consider persistence diagrams associated to infinite spaces, the use of persis-




Figure 1: A classical pipeline for persistence in TDA.

tent homology in TDA can be naturally considered in general statistical frameworks and persistence diagrams can be used as statistics with interesting convergence properties.

1.2 Contribution

In this paper we assume that the available data is the realization of a probability distribution supported on an unknown compact metric space. We consider the persistent homology of different filtered simplicial complexes built on top of the data. We study, with a minimax approach, the rate of convergence of the associated persistence diagrams to some well-defined persistence diagram associated to the support of the probability distribution. More precisely, we assume that we observe a set of n points $\widehat{\mathbb{X}}_n = \{X_1 \dots, X_n\}$ in a metric space (\mathbb{M}, ρ) , drawn i.i.d. from some unknown measure μ whose support is a compact set denoted $\mathbb{X}_\mu \subseteq \mathbb{M}$. We also assume that μ satisfies the so-called (a, b) -standard assumption for some constants $a, b > 0$: for any $x \in \mathbb{X}_\mu$ and any $r > 0$, $\mu(B(x, r)) \geq \min(ar^b, 1)$. The following theorem illustrates the kind of results we obtain under such assumption.

Theorem (4 in Section 3): *Let (\mathbb{M}, ρ) , $a > 0$ and $b > 0$ as above. Then for any measure μ satisfying the (a, b) -standard assumption*

$$\mathbb{E} \left[d_b(\text{dgm}(\text{Filt}(\mathbb{X}_\mu)), \text{dgm}(\text{Filt}(\widehat{\mathbb{X}}_n))) \right] \leq C \left(\frac{\ln n}{n} \right)^{1/b}$$

where the constant C only depends on a and b (not on \mathbb{M}). Assume moreover that there exists a non isolated point x in \mathbb{M} and consider any sequence $(x_n) \in (\mathbb{M} \setminus \{x\})^{\mathbb{N}}$ such that $\rho(x, x_n) \leq (an)^{-1/b}$. Then for any estimator $\widehat{\text{dgm}}_n$ of $\text{dgm}(\text{Filt}(\mathbb{X}_\mu))$:

$$\liminf_{n \rightarrow \infty} \rho(x, x_n)^{-1} \mathbb{E} \left[d_b(\text{dgm}(\text{Filt}(\mathbb{X}_\mu)), \widehat{\text{dgm}}_n) \right] \geq C'$$

where C' is an absolute constant.

Our approach relies on the general theory of persistence modules and our results follow from two recently proven properties of persistence diagrams (Chazal et al., 2012b, 2009a, 2012a).

First, as \mathbb{X}_μ can be any compact metric space (possibly infinite), the filtered complex $\text{Filt}(\mathbb{X}_\mu)$ is usually not finite or even countable and the existence of its persistence diagram cannot be established from the “classical” persistence theory (Zomorodian and Carlsson, 2005; Edelsbrunner et al., 2002). In our setting, the existence of $\text{dgm}(\text{Filt}(\mathbb{X}_\mu))$ follows from

the general persistence framework introduced by Chazal et al. (2009a, 2012a). Notice that although this framework is rather abstract and theoretical, it does not have any practical drawback as only persistence diagrams of complexes built on top of finite data are computed. Second, a fundamental property of the persistence diagrams we are considering is their stability proven by Chazal et al. (2012b): the bottleneck distance between $\text{dgm}(\text{Filt}(\mathbb{X}_\mu))$ and $\text{dgm}(\text{Filt}(\widehat{\mathbb{X}}_n))$ is upper bounded by twice the Gromov-Hausdorff distance between \mathbb{X}_μ and $\widehat{\mathbb{X}}_n$. This result establishes a strong connection between our persistence estimation problem and support estimation problems. Upper bounds on the rate of convergence of persistence diagrams are then easily obtained using the same arguments as the ones usually used to obtain convergence results for support estimation with respect to the Hausdorff metric. We take advantage of this general remark to find rates of convergence of persistence diagrams in general metric spaces (Section 3) and also in the more classical case where the measure is supported in \mathbb{R}^d (Section 4). Using Le Cam's lemma, we also compute the corresponding lower bounds to check that the rates of convergence are optimal in the minimax sense.

1.3 Related Works

Although it is attracting more and more interest, the use of persistent homology in data analysis remains widely heuristic. There are relatively few papers establishing connections between persistence and statistics and, despite a few promising results, the statistical analysis of homology, persistent homology and more general topological and geometric features of data is still in its infancy.

One of the first statistical results about persistent homology has been given in a parametric setting, by Bubenik and Kim (2007). They show for instance that for data sampled on an hypersphere according to a von-Mises Fisher distribution (among other distributions), the persistence diagrams of the density can be estimated with the parametric rate $n^{-1/2}$. However assuming that both the support and the parametric family of the distribution are known are strong assumptions which are hardly met in practice.

Closely related to our approach, statistical analysis of homology and of persistent homology has also been proposed very recently by Balakrishnan et al. (2012); Fasy et al. (2014) in the specific context of manifolds, i.e. when the geometric structure underlying the data is assumed to be a smooth submanifold of an Euclidean space. In the first paper, the authors exhibit minimax rates of convergence for the estimation of the Betti numbers of the underlying manifold under different models of noise. This approach is also strongly connected to manifold estimation results obtained by Genovese et al. (2012b). Related lower bounds have also been recently obtained by Weinberger (2014) in a different and more restrictive setting. Our results are in the same spirit as Balakrishnan et al. (2012) but extend to persistent homology and allow us to deal with general compact metric spaces. In the second paper, the authors develop several methods to find confidence sets for persistence diagrams using subsampling methods and kernel estimators among other approaches. Although they tackle a different problem, it has some connections with the problem considered in the present paper that we briefly mention in Section 3.4.

Both Fasy et al. (2014) and our work start from the observation that persistence diagram inference is strongly connected to the better known problem of support estimation. As far

as we know, only few results about support estimation in general metric spaces have been given in the past. An interesting framework is proposed by De Vito et al. (2014): in this paper the support estimation problem is tackled using kernel methods. On the other hand, a large amount of literature is available for measure support estimation in \mathbb{R}^d ; see for instance the review by Cuevas (2009) for more details. Note that many results on this topic are given with respect to the volume of symmetric set difference (see for instance Biau et al., 2009, and references therein) while in our topological estimation setting we need convergence results for support estimation in Hausdorff metric.

The estimator $\widehat{X}_n = \{X_1, \dots, X_n\}$ and the Devroye and Wise (1980) estimator, $\widehat{S}_n = \bigcup_{i=1}^n \bar{B}(X_i, \varepsilon_n)$, where $\bar{B}(x, \varepsilon)$ denotes the closed ball centered at x with radius ε , are both natural estimators of the support. The use of \widehat{S}_n is particularly relevant when the convergence of the measure of the symmetric set difference is considered but does not provide better results than \widehat{X}_n in our Hausdorff distance setting. The convergence rate of \widehat{X}_n to the support of the measure with respect to the Hausdorff distance is given by Cuevas and Rodríguez-Casal (2004) in \mathbb{R}^d . Support estimation in \mathbb{R}^d has also been studied under various additional assumptions such as convexity assumptions (Dümbgen and Walther, 1996; Rodríguez-Casal, 2007; Cuevas et al., 2012) or through boundary fragments estimation (Korostelëv and Tsybakov, 1993; Korostelëv et al., 1995) just to name a few. Another classical assumption is that the measure has a density with respect to the Lebesgue measure. In this context, plug-in methods based on non parametric estimators of the density have been proposed by Cuevas and Fraiman (1997) and Tsybakov (1997). We consider persistence diagram estimation in the density framework of Singh et al. (2009) in Section 4 and show in this particular context that \widehat{X}_n allows us to define a persistence diagram estimator that reaches optimal rates of convergence in the minimax sense.

A few different methods have also been proposed for topology estimation in non-deterministic frameworks such as those based on deconvolution (Caillerie et al., 2011; Niyogi et al., 2011). Several recent attempts have also been made, with completely different approaches, to study persistence diagrams from a statistical point of view, such as Mileyko et al. (2011) who study probability measures on the space of persistence diagrams or Bubenik (2012) who introduces a functional representation of persistence diagrams, the so-called persistence landscapes, allowing means and variance of persistence diagrams to be defined. Notice that our results should easily extend to persistence landscapes.

The paper is organized as follows. Background notions and results on metric spaces, filtered simplicial complexes, and persistent homology that are necessary to follow the paper are presented in Section 2. The rates of convergence for the estimation of persistence diagrams in general metric spaces are established in Section 3. We also study these convergence rates in \mathbb{R}^d for a few classical problems in Section 4. Some numerical experiments illustrating our results are given in Section 5. All the technical proofs are given in Appendix.

2. Background

We first recall the required background about measured metric spaces and persistent homology.

2.1 Metric Measure Spaces

Recall that a metric space is a pair (\mathbb{M}, ρ) where \mathbb{M} is a set and $\rho : \mathbb{M} \times \mathbb{M} \rightarrow \mathbb{R}$ is a nonnegative map such that for any $x, y, z \in \mathbb{M}$, $\rho(x, y) = 0$ if and only if $x = y$, $\rho(x, y) = \rho(y, x)$ and $\rho(x, z) \leq \rho(x, y) + \rho(y, z)$. We denote by $\mathcal{K}(\mathbb{M})$ the set of all the compact subsets of \mathbb{M} . For a point $x \in \mathbb{M}$ and a subset $C \in \mathcal{K}(\mathbb{M})$, the distance $d(x, C)$ of x to C is the minimum over all $y \in C$ of $d(x, y)$. The Hausdorff distance $d_H(C_1, C_2)$ between two subsets $C_1, C_2 \in \mathcal{K}(\mathbb{M})$ is the maximum over all points in C_1 of their distance to C_2 and over all points in C_2 of their distance to C_1 :

$$d_H(C_1, C_2) = \max\left\{ \sup_{x \in C_1} d(x, C_2), \sup_{y \in C_2} d(y, C_1) \right\}.$$

Note that $(\mathcal{K}(\mathbb{M}), d_H)$ is a metric space and can be endowed with its Borel σ -algebra.

Two compact metric spaces (\mathbb{M}_1, ρ_1) and (\mathbb{M}_2, ρ_2) are *isometric* if there exists a bijection $\Phi : \mathbb{M}_1 \rightarrow \mathbb{M}_2$ that preserves distances, namely: $\forall x, y \in \mathbb{M}_1, \rho_2(\Phi(x), \Phi(y)) = \rho_1(x, y)$. Such a map Φ is called an *isometry*. One way to compare two metric spaces is to measure how far these two metric spaces are from being isometric. The corresponding distance is called the *Gromov-Hausdorff distance* (see for instance Burago et al., 2001). Intuitively, it is the infimum of their Hausdorff distance over all possible isometric embeddings of these two spaces into a common metric space.

Definition 1 *Let (\mathbb{M}_1, ρ_1) and (\mathbb{M}_2, ρ_2) be two compact metric spaces. The Gromov-Hausdorff distance $d_{GH}((\mathbb{M}_1, \rho_1), (\mathbb{M}_2, \rho_2))$ is the infimum of the real numbers $r \geq 0$ such that there exist a metric space (\mathbb{M}, ρ) and subspaces C_1 and C_2 in $\mathcal{K}(\mathbb{M})$ which are isometric to \mathbb{M}_1 and \mathbb{M}_2 respectively and such that $d_H(C_1, C_2) < r$. The Gromov-Hausdorff distance d_{GH} defines a metric on the space \mathcal{K} of isometry classes of compact metric spaces (see Burago et al., 2001, Theorem 7.3.30).*

Notice that when \mathbb{M}_1 and \mathbb{M}_2 are subspaces of a same metric space (\mathbb{M}, ρ) then $d_{GH}(\mathbb{M}_1, \mathbb{M}_2) \leq d_H(\mathbb{M}_1, \mathbb{M}_2)$.

2.2 Measure

Let μ be a probability measure on (\mathbb{M}, ρ) equipped with its Borel algebra. Let \mathbb{X}_μ denote the support of the measure μ , namely the smallest closed set with probability one. In the following of the paper, we will assume that \mathbb{X}_μ is compact and thus $\mathbb{X}_\mu \in \mathcal{K}(\mathbb{M})$. Also note that $(\mathbb{X}_\mu, \rho) \in \mathcal{K}$.

The main assumption we will need in the following of the paper provides a lower bound on the measure μ . We say that μ satisfies the *standard assumption* if there exist $a' > 0$, $r_0 > 0$ and $b > 0$ such that

$$\forall x \in \mathbb{X}_\mu, \forall r \in (0, r_0), \mu(B(x, r)) \geq a' r^b \tag{2.1}$$

where $B(x, r)$ denotes the open ball of center x and radius r in \mathbb{M} . This assumption is popular in the literature about set estimation (see for instance Cuevas, 2009) but it has generally been considered with $b = d$ in \mathbb{R}^d . Since \mathbb{X}_μ is compact, reducing the constant a' to a smaller constant a if necessary, we easily check that assumption (2.1) is equivalent to

$$\forall x \in \mathbb{X}_\mu, \forall r > 0, \mu(B(x, r)) \geq 1 \wedge ar^b \tag{2.2}$$




Figure 2: From left to right: the α sublevelset of the distance function to a point set \mathbb{X} in \mathbb{R}^2 , the α -complex, $\text{Cech}_\alpha(\mathbb{X})$ and $\text{Rips}_{2\alpha}(\mathbb{X})$. The last two include a tetrahedron.

where $x \wedge y$ denotes the minimum between x and y . We then say that μ satisfies the (a, b) -standard assumption.

2.3 Simplicial Complexes on Metric Spaces

The geometric complexes we consider in this paper are built on top of metric spaces and come as nested families indexed by a real parameter. Topological persistence is used to infer and encode the evolution of the topology of these families as the parameter grows. For a complete definition of these geometric filtered complexes built on top of metric spaces and their use in TDA, we refer to Chazal et al. (2012b), Section 4.2. Here we only give a brief reminder and refer to Figure 2 for illustrations. A simplicial complex \mathcal{C} is a set of simplices (points, segments, triangles, etc) such that any face from a simplex in \mathcal{C} is also in \mathcal{C} and the intersection of any two simplices of \mathcal{C} is a (possibly empty) face of these simplices. Notice that we do not assume such simplicial complexes to be finite. The complexes we consider in this paper can be seen as a generalization of neighborhood graphs in dimension larger than 1.

Given a metric space \mathbb{X} which will also serve as the vertex set, the *Vietoris-Rips complex* $\text{Rips}_\alpha(\mathbb{X})$ is the set of simplices $[x_0, \dots, x_k]$ such that $d_{\mathbb{X}}(x_i, x_j) \leq \alpha$ for all (i, j) . The *Čech complex* $\text{Cech}_\alpha(\mathbb{X})$ is similarly defined as the set of simplices $[x_0, \dots, x_k]$ such that the $k + 1$ closed balls $B(x_i, \alpha)$ have a non-empty intersection. Note that these two complexes are related by $\text{Rips}_\alpha(\mathbb{X}) \subseteq \text{Cech}_\alpha(\mathbb{X}) \subseteq \text{Rips}_{2\alpha}(\mathbb{X})$. Note also that these two families of complexes only depend on the pairwise distances between the points of \mathbb{X} .

When \mathbb{X} is embedded in some larger metric space \mathbb{M} , we can extend the definition of the Čech complex to the set of simplices $[x_0, \dots, x_k]$ such that the $k + 1$ closed balls $B(x_i, \alpha)$ have a non-empty intersection in \mathbb{M} (not just in \mathbb{X}). We can also define the *alpha-complex* or *α -complex* as the set of simplices $[x_0, \dots, x_k]$ such that, for some $\beta \leq \alpha$ that depends on the simplex, the $k + 1$ closed balls $B(x_i, \beta)$ and the complement of all the other balls $B(x, \beta)$ for




Figure 3: A torus \mathbb{T} filtered by its z -coordinate: $\text{Filt}_\alpha = \{P \in \mathbb{T} \mid P_z \leq \alpha\}$, its persistence barcode, and its persistence diagram.

$x \in \mathbb{X}$ have a non-empty intersection in \mathbb{M} . In the particular case where $\mathbb{M} = \mathbb{R}^d$, those two complexes have the same homotopy type (they are equivalent for our purposes) as the union of the balls $B(x, \alpha)$ for $x \in \mathbb{X}$, as in Figure 2, and the α -complex only contains simplices of dimension at most d . Note that the union of the balls $B(x, \alpha)$ is also the α -sublevel set of the distance to \mathbb{X} function $d(\cdot, \mathbb{X})$, and as a consequence, those filtrations thus provide a convenient way to study the evolution of the topology of the union of growing balls or the sublevel sets of $d(\cdot, \mathbb{X})$ (see Figure 2 and Section 5 for more examples).

There are several other families that we could also have considered, most notably witness complexes (Chazal et al., 2012b). Extending our results to them is straightforward and yields very similar results, so we will restrict to the families defined above in the rest of the paper.

All these families of complexes have the fundamental property that they are non-decreasing with α ; for any $\alpha \leq \beta$, there is an inclusion of $\text{Rips}_\alpha(\mathbb{X})$ in $\text{Rips}_\beta(\mathbb{X})$, and similarly for the Čech, and Alpha complexes. They are thus called *filtrations*. In the following, the notation $\text{Filt}(\mathbb{X}) := (\text{Filt}_\alpha(\mathbb{X}))_{\alpha \in A}$ denotes one of the filtrations defined above.

2.4 Persistence Diagrams

An extensive presentation of persistence diagrams is available in Chazal et al. (2012a). We recall a few definitions and results that are needed in this paper.

We first give the intuition behind persistence. Given a filtration as above, the topology of $\text{Filt}_\alpha(\mathbb{X})$ changes as α increases: new connected components can appear, existing connected components can merge, cycles and cavities can appear and can be filled, etc. Persistent homology is a tool that tracks these changes, identifies *features* and associates a *lifetime* to them. For instance, a connected component is a feature that is born at the smallest α such that the component is present in $\text{Filt}_\alpha(\mathbb{X})$, and dies when it merges with an older connected component. Intuitively, the longer a feature persists, the more relevant it is.




Figure 4: An α -complex filtration, the sublevelset filtration of the distance function, and their common persistence barcode (they are homotopy equivalent).

We now formalize the presentation a bit. Given a filtration as above, we can consider the \mathbb{Z}_2 -homology groups ¹ of the simplicial complexes and get a sequence of vector spaces $(H(\text{Filt}_\alpha(\mathbb{X})))_{\alpha \in \mathcal{A}}$, where the inclusions $\text{Filt}_\alpha(\mathbb{X}) \subseteq \text{Filt}_\beta(\mathbb{X})$ induce linear maps $H(\text{Filt}_\alpha(\mathbb{X})) \rightarrow H(\text{Filt}_\beta(\mathbb{X}))$. In many cases, this sequence can be decomposed as a direct sum of intervals, where an interval is a sequence of the form

$$0 \rightarrow \dots \rightarrow 0 \rightarrow \mathbb{Z}_2 \rightarrow \dots \rightarrow \mathbb{Z}_2 \rightarrow 0 \rightarrow \dots \rightarrow 0$$

(the linear maps $\mathbb{Z}_2 \rightarrow \mathbb{Z}_2$ are all the identity). These intervals can be interpreted as features of the (filtered) complex, such as a connected component or a loop, that appear at parameter α_{birth} in the filtration and disappear at parameter α_{death} . An interval is determined uniquely by these two parameters. It can be represented as a segment whose extremities have abscissae α_{birth} and α_{death} ; the set of these segments is called the barcode of $\text{Filt}(\mathbb{X})$. An interval can also be represented as a point in the plane, where the x -coordinate indicates the birth time and the y -coordinate the death time. The set of points (with multiplicity) representing the intervals is called the persistence diagram $\text{dgm}(\text{Filt}(\mathbb{X}))$. Note that the diagram is entirely contained in the half-plane above the diagonal Δ defined by $y = x$, since death always occurs after birth. Chazal et al. (2012a) show that this diagram is still well-defined even in cases where the sequence might not be decomposable as a finite sum of intervals, and in particular $\text{dgm}(\text{Filt}(\mathbb{X}))$ is well-defined for any compact metric space \mathbb{X} (Chazal et al., 2012b). Note that for technical reasons, the points of the diagonal Δ are considered as part of every persistence diagram, with infinite multiplicity. The most persistent features (supposedly the most important) are those represented by the longest bars in the barcode, i.e. the points furthest from the diagonal in the diagram, whereas points close to the diagonal can be interpreted as noise.

1. The notion of (simplicial) homology is a classical concept in algebraic topology that provides powerful tools to formalize and handle the notion of topological features of a simplicial complex in an algebraic way. For example the 0-dimensional homology group H_0 represents the 0-dimensional features, i.e. the connected components of the complex, H_1 represents the 1-dimensional features (cycles), H_2 represents the 2-dimensional features (cavities),... See Hatcher (2001) for an introduction to simplicial homology.




Figure 5: Two diagrams at bottleneck distance ε .

The space of persistence diagrams is endowed with a metric called the *bottleneck distance* d_b . Given two persistence diagrams, it is defined as the infimum, over all perfect matchings of their points, of the largest L^∞ -distance between two matched points, see Figure 5. The presence of the diagonal in all diagrams means we can consider partial matchings of the off-diagonal points, and the remaining points are matched to the diagonal. With more details, given two diagrams dgm_1 and dgm_2 , we can define a matching m as a subset of $\text{dgm}_1 \times \text{dgm}_2$ such that every point of $\text{dgm}_1 \setminus \Delta$ and $\text{dgm}_2 \setminus \Delta$ appears exactly once in m . The bottleneck distance is then:

$$d_b(\text{dgm}_1, \text{dgm}_2) = \inf_{\text{matching } m} \max_{(p,q) \in m} \|q - p\|_\infty.$$

Note that points close to the diagonal Δ are easily matched to the diagonal, which fits with their interpretation as irrelevant noise.

A fundamental property of persistence diagrams, proved by Chazal et al. (2012a), is their *stability*. If \mathbb{X} and $\tilde{\mathbb{X}}$ are two compact metric spaces then one has

$$d_b(\text{dgm}(\text{Filt}(\mathbb{X})), \text{dgm}(\text{Filt}(\tilde{\mathbb{X}}))) \leq 2d_{\text{GH}}(\mathbb{X}, \tilde{\mathbb{X}}). \tag{2.3}$$

Moreover, if \mathbb{X} and $\tilde{\mathbb{X}}$ are embedded in the same metric space (\mathbb{M}, ρ) then one has

$$d_b(\text{dgm}(\text{Filt}(\mathbb{X})), \text{dgm}(\text{Filt}(\tilde{\mathbb{X}}))) \leq 2d_{\text{GH}}(\mathbb{X}, \tilde{\mathbb{X}}) \leq 2d_{\text{H}}(\mathbb{X}, \tilde{\mathbb{X}}). \tag{2.4}$$

Notice that these properties are only metric properties: they do not involve here any probability measure on \mathbb{X} and $\tilde{\mathbb{X}}$.

3. Persistence Diagram Estimation in Metric Spaces

Let (\mathbb{M}, ρ) be a metric space. Assume that we observe n points $X_1 \dots, X_n$ in \mathbb{M} drawn i.i.d. from some unknown measure μ whose support is a compact set denoted \mathbb{X}_μ .

3.1 From Support Estimation to Persistence Diagram Estimation

The Gromov-Hausdorff distance allows us to compare \mathbb{X}_μ with compact metric spaces not necessarily embedded in \mathbb{M} . We thus consider (\mathbb{X}_μ, ρ) as an element of \mathcal{K} (rather than an

element of $\mathcal{K}(\mathbb{M})$). In the following, an *estimator* $\widehat{\mathbb{X}}$ of \mathbb{X}_μ is a function of $X_1 \dots, X_n$ that takes values in \mathcal{K} and which is measurable for the Borel algebra induced by d_{GH} .

Let $\text{Filt}(\mathbb{X}_\mu)$ and $\text{Filt}(\widehat{\mathbb{X}})$ be filtrations defined on \mathbb{X}_μ and $\widehat{\mathbb{X}}$. The statistical analysis of persistence diagrams proposed above starts from the following key fact: according to (2.3), for any $\varepsilon > 0$:

$$\mathbb{P}\left(d_b\left(\text{dgm}(\text{Filt}(\widehat{\mathbb{X}}_\mu)), \text{dgm}(\text{Filt}(\widehat{\mathbb{X}}))\right) > \varepsilon\right) \leq \mathbb{P}\left(d_{\text{GH}}(\mathbb{X}_\mu, \widehat{\mathbb{X}}) > 2\varepsilon\right) \quad (3.1)$$

where the probability corresponds to the product measure $\mu^{\otimes n}$. Our strategy then consists in finding an estimator of the support which is close to \mathbb{X}_μ for the d_{GH} distance. Note that this general strategy of estimating \mathbb{X}_μ in \mathcal{K} is not only of theoretical interest. Indeed, as mentioned in the introduction, in some cases the space \mathbb{M} is unknown and the observations $X_1 \dots, X_n$ are just known through their matrix of pairwise distances $\rho(X_i, X_j)$, $i, j = 1, \dots, n$. The use of the Gromov-Hausdorff distance then allows us to consider this set of observations as an abstract metric space of cardinality n independently of the way it is embedded in \mathbb{M} .

This general framework includes the more standard approach consisting in estimating the support by restraining the values of $\widehat{\mathbb{X}}$ to $\mathcal{K}(\mathbb{M})$. According to (2.4), in this case, for any $\varepsilon > 0$:

$$\mathbb{P}\left(d_b\left(\text{dgm}(\text{Filt}(\widehat{\mathbb{X}}_\mu)), \text{dgm}(\text{Filt}(\widehat{\mathbb{X}}))\right) > \varepsilon\right) \leq \mathbb{P}\left(d_{\text{H}}(\mathbb{X}_\mu, \widehat{\mathbb{X}}) > 2\varepsilon\right). \quad (3.2)$$

Using equations (3.1) and (3.2) the problem of persistence diagrams estimation reduces to the better known problem of estimating the support of a measure.

Let $\widehat{\mathbb{X}}_n := \{X_1, \dots, X_n\}$ be a set of independent observations sampled according to μ endowed with the restriction of the distance ρ to this set. This finite metric space is a natural estimator of the support \mathbb{X}_μ . In several contexts discussed in the following, $\widehat{\mathbb{X}}_n$ shows optimal rates of convergence for the estimation of \mathbb{X}_μ with respect to the Hausdorff and Gromov-Hausdorff distance. From (3.2) we will then obtain upper bounds on the rate of convergence of $\text{Filt}(\widehat{\mathbb{X}}_n)$. We also obtain the corresponding lower bounds to prove optimality.

In the next subsection, we tackle persistence diagram estimation in the general framework of abstract metric spaces. We will consider more particular contexts later in the paper.

3.2 Convergence of Persistence Diagrams

Cuevas and Rodríguez-Casal (2004) give the rate of convergence in Hausdorff distance of $\widehat{\mathbb{X}}_n$ for some probability measure μ satisfying an (a, b) -standard assumption on \mathbb{R}^d . In this section, we consider the more general context where μ is a probability measure satisfying an (a, b) -standard assumption on a metric space (\mathbb{M}, ρ) , with $b > 0$. We give below the rate of convergence of $\widehat{\mathbb{X}}_n$ in this context. The proof follows the lines of the proof of Cuevas and Rodríguez-Casal (2004, Theorem 3).

Theorem 2 *Assume that a probability measure μ on \mathbb{M} satisfies the (a, b) -standard assumption. Then, for any $\varepsilon > 0$:*

$$\mathbb{P}\left(d_{\text{H}}(\mathbb{X}_\mu, \widehat{\mathbb{X}}_n) > 2\varepsilon\right) \leq \frac{2^b}{a\varepsilon^b} \exp(-na\varepsilon^b) \wedge 1.$$

Moreover, there exist two constants C_1 and C_2 only depending on a and b such that

$$\limsup_{n \rightarrow \infty} \left(\frac{n}{\log n} \right)^{1/b} d_H(\mathbb{X}_\mu, \widehat{\mathbb{X}}_n) \leq C_1 \quad \text{almost surely,}$$

and

$$\lim_{n \rightarrow \infty} \mathbb{P} \left(d_H(\mathbb{X}_\mu, \widehat{\mathbb{X}}_n) \leq C_2 \left(\frac{\log n}{n} \right)^{1/b} \right) = 1.$$

Since $d_{\text{GH}}(\mathbb{X}_\mu, \widehat{\mathbb{X}}_n) \leq d_H(\mathbb{X}_\mu, \widehat{\mathbb{X}}_n)$ the above theorem also holds when the Gromov distance is replaced by the Gromov-Hausdorff distance. In practice this allows to consider $\widehat{\mathbb{X}}_n$ as an abstract metric space without taking care of the way it is embedded in the, possibly unknown, metric space \mathbb{M} .

Using (3.1) and (2.4), we then derive from the previous result the following corollary for the convergence rate of the persistence diagram $\text{Filt}(\widehat{\mathbb{X}}_n)$ toward $\text{Filt}(\mathbb{X}_\mu)$.

Corollary 3 *Assume that the probability measure μ on \mathbb{M} satisfies the (a, b) -standard assumption, then for any $\varepsilon > 0$:*

$$\mathbb{P} \left(d_b \left(\text{dgm}(\text{Filt}(\mathbb{X}_\mu)), \text{dgm}(\text{Filt}(\widehat{\mathbb{X}}_n)) \right) > \varepsilon \right) \leq \frac{2^b}{a\varepsilon^b} \exp(-na\varepsilon^b) \wedge 1. \quad (3.3)$$

Moreover,

$$\limsup_{n \rightarrow \infty} \left(\frac{n}{\log n} \right)^{1/b} d_b \left(\text{dgm}(\text{Filt}(\mathbb{X}_\mu)), \text{dgm}(\text{Filt}(\widehat{\mathbb{X}}_n)) \right) \leq C_1 \quad \text{almost surely,}$$

and

$$\lim_{n \rightarrow \infty} \mathbb{P} \left(d_b \left(\text{dgm}(\text{Filt}(\mathbb{X}_\mu)), \text{dgm}(\text{Filt}(\widehat{\mathbb{X}}_n)) \right) \leq C_2 \left(\frac{\log n}{n} \right)^{1/b} \right) = 1$$

where C_1 and C_2 are the same constants as in Theorem 2.

3.3 Minimax Optimal Rate of Convergence

Let $\mathcal{P}(a, b, \mathbb{M})$ be the set of all the probability measures on the metric space (\mathbb{M}, ρ) satisfying the (a, b) -standard assumption on \mathbb{M} :

$$\mathcal{P}(a, b, \mathbb{M}) := \left\{ \mu \text{ on } \mathbb{M} \mid \mathbb{X}_\mu \text{ is compact and } \forall x \in \mathbb{X}_\mu, \forall r > 0, \mu(B(x, r)) \geq 1 \wedge ar^b \right\}.$$

The next theorem gives upper and lower bounds for the rate of convergence of persistence diagrams. The upper bound comes as a consequence of Corollary 3, while the lower bound is established using the so-called Le Cam's lemma (see Lemma 9 in Appendix).

Theorem 4 *Let (\mathbb{M}, ρ) be a metric space and let $a > 0$ and $b > 0$. Then:*

$$\sup_{\mu \in \mathcal{P}(a, b, \mathbb{M})} \mathbb{E} \left[d_b(\text{dgm}(\text{Filt}(\mathbb{X}_\mu)), \text{dgm}(\text{Filt}(\widehat{\mathbb{X}}_n))) \right] \leq C \left(\frac{\log n}{n} \right)^{1/b} \quad (3.4)$$

where the constant C only depends on a and b (not on \mathbb{M}). Assume moreover that there exists a non isolated point x in \mathbb{M} and consider any sequence $(x_n) \in (\mathbb{M} \setminus \{x\})^{\mathbb{N}}$ such that $\rho(x, x_n) \leq (an)^{-1/b}$. Then for any estimator $\widehat{\text{dgm}}_n$ of $\text{dgm}(\text{Filt}(\mathbb{X}_\mu))$:

$$\liminf_{n \rightarrow \infty} \rho(x, x_n)^{-1} \sup_{\mu \in \mathcal{P}(a,b,\mathbb{M})} \mathbb{E} \left[d_b(\text{dgm}(\text{Filt}(\mathbb{X}_\mu)), \widehat{\text{dgm}}_n) \right] \geq C'$$

where C' is an absolute constant.

Consequently, the estimator $\text{dgm}(\text{Filt}(\widehat{\mathbb{X}}_n))$ is minimax optimal on the space $\mathcal{P}(a, b, \mathbb{M})$ up to a logarithmic term as soon as we can find a non-isolated point in \mathbb{M} and a sequence (x_n) in \mathbb{M} such that $\rho(x_n, x) \sim (an)^{-1/b}$. This is obviously the case for the Euclidean space \mathbb{R}^d .

One classical method to obtain tight lower bounds with sup norm metrics is applying a Fano’s strategy based on several hypotheses (see for instance Tsybakov and Zaiats, 2009, Chapter 2). Applying this method is more difficult than it seems in our context. Indeed, the bottleneck distance makes tricky the construction of multiple hypotheses. However, in specific cases, we can obtain the matching lower bound with a more direct proof.

Theorem 5 Consider $(\frac{1}{2}, 1)$ -standard measures on the unit segment $[0, 1]$. For any estimator $\widehat{\text{dgm}}_n$ of $\text{dgm}(\text{Filt}(\mathbb{X}_\mu))$:

$$\liminf_{n \rightarrow \infty} \sup_{\mu \in \mathcal{P}(\frac{1}{2}, 1, [0,1])} \frac{n}{\log n} \mathbb{E} \left[d_b(\text{dgm}(\text{Filt}(\mathbb{X}_\mu)), \widehat{\text{dgm}}_n) \right] \geq C$$

where C is an absolute constant.

It should be straightforward to extend this to measures on the cube $[0, 1]^b$, as long as b is an integer, with a lower-bound of $C_b (\frac{\log n}{n})^{1/b}$. Note that this bound applies to the homology of dimension b . It is possible that lower-dimensional homology may be easier to estimate.

3.4 Confidence Sets for Persistence Diagrams

Corollary 3 can also be used to find confidence sets for persistence diagrams. Assume that a and b are known and let $\Psi : \eta \rightarrow \exp(-\eta)/\eta$. Then for $\alpha \in (0, 1)$,

$$B_{d_b} \left(\text{dgm}(\text{Filt}(\mathbb{X}_\mu)), \left[\frac{1}{na} \Psi^{-1} \left(\frac{\alpha}{n2^b} \right) \right]^{1/b} \right)$$

is a confidence region for $\text{dgm}(\text{Rips}(\mu(K)))$ of level $1 - \alpha$. Nevertheless, in practice the coefficients a and b can be unknown. In \mathbb{R}^d , the coefficient b can be taken equal to the ambient dimension d in many situations. Finding lower bounds on the coefficient a is a tricky problem that is out of the scope of the paper. Alternative solutions have been proposed recently by Fasy et al. (2014) and we refer the reader to this paper for more details.

4. Persistence Diagram Estimation in \mathbb{R}^k

In this section, we study the convergence rates of persistence diagram estimators for data embedded in \mathbb{R}^k . In particular we study two situations of interest proposed respectively by Singh et al. (2009) and Genovese et al. (2012b) in the context of measure support estimation. In the first situation the measure has a density with respect to the Lebesgue measure on \mathbb{R}^d whose behavior is controlled near the boundary of its support. In the second case, the measure is supported on a manifold. These two frameworks are complementary and provide realistic frameworks for topological inference in \mathbb{R}^d .

4.1 Minimax Optimal Persistence Diagram Estimation for Nonsingular Measures on \mathbb{R}^k

The paper by Singh et al. (2009) is a significant breakthrough for level set estimation through density estimation. It presents a fully data-driven procedure, in the spirit of Lepski's method, that is adaptive to unknown local density regularity and achieves a Hausdorff error control that is minimax optimal for a class of level sets with very general shapes. In particular, the assumptions of Singh et al. (2009) describe the smoothness of the density near the boundary of the support.

In this section, we propose to study persistence diagram inference in the framework of Singh et al. (2009) since this framework is very intuitive and natural. Nevertheless, we do not use the estimator of Singh et al. (2009) for this task since we only consider here the support estimation problem (and not the more general level set issue as in Singh et al., 2009). Indeed, we will see that the estimator \hat{X}_n has the optimal rate of convergence for estimating the support according to d_H , as well as for estimating the persistence diagram. We now recall the framework of Singh et al. (2009, Section 4.3) corresponding to support set estimation.

Let X_1, \dots, X_n be i.i.d. observations drawn from an unknown probability measure μ having density f with respect to the Lebesgue measure and defined on a compact set $\chi \subset \mathbb{R}^k$. Let X_f denote the support of μ , and let $G_0 := \{x \in \chi \mid f(x) > 0\}$. The boundary of a set G is denoted ∂G and for any $\varepsilon > 0$, $I_\varepsilon(G) := \bigcup_{x \mid B(x, \varepsilon) \subset G} B(x, \varepsilon)$ is the ε -inner of G . The two main assumptions of Singh et al. (2009) are the following:

[A]: The density f is upper bounded by $f_{\max} > 0$ and there exist constants $\alpha, C_a, \delta_a > 0$ such that for all $x \in G_0$ with $f(x) \leq \delta_a$, $f(x) \geq C_a d(x, \partial G_0)^\alpha$.

[B]: There exist constants $\varepsilon_0 > 0$ and $C_b > 0$ such that for all $\varepsilon \leq \varepsilon_0$, $I_\varepsilon(G_0) \neq \emptyset$ and $d(x, I_\varepsilon(G_0)) \leq C_b \varepsilon$ for all $x \in \partial G_0$.

We denote by $\mathcal{F}(\alpha)$ the set composed of all the densities on χ satisfying assumptions [A] and [B], for a fixed set of positive constants $C_a, C_b, \delta_a, \varepsilon_0, f_{\max}, p$ and α .

Assumption [A] describes how fast the density increases in the neighborhood of the boundary of the support: the smaller α , the easier the support may be possible to detect. Assumption [B] prevents the boundary from having arbitrarily small features (as for cusps). We refer to Singh et al. (2009) for more details and discussions about these two assumptions and their connections with assumptions in other works.

For persistence diagram estimation, we are interested in estimating the support X_f whereas the assumptions [A] and [B] involve the set G_0 . However, as stated in Lemma 11

(given in Appendix B.4), these two sets are here almost identical in the sense that $d_H(G_0, \mathbb{X}_f) = 0$. Moreover, it can be proved that under assumptions $[A]$ and $[B]$, the measure μ also satisfies the standard assumption with $b = \alpha + k$ (see Lemma 11). According to Proposition 4, the estimator $\text{dgm}(\text{Filt}(\widehat{\mathbb{X}}_n))$ thus converges in expectation to $\text{dgm}(\text{Filt}(\mathbb{X}_f))$ for d_b with a rate upper bounded by $(\log n/n)^{1/(k+\alpha)}$. We also show that this rate is minimax over the sets $\mathcal{F}(\alpha)$ by adapting the ideas of the proof given by Singh et al. (2009) for the Hausdorff lower bound.

Proposition 6 1. For all $n \geq 1$,

$$\sup_{f \in \mathcal{F}(\alpha)} \mathbb{E} \left[d_b(\text{dgm}(\text{Filt}(\mathbb{X}_f)), \text{dgm}(\text{Filt}(\widehat{\mathbb{X}}_n))) \right] \leq C \left(\frac{n}{\log n} \right)^{-1/(k+\alpha)}$$

where C is a constant depending only on $C_a, C_b, \delta_a, \varepsilon_0, f_{max}, p$ and α .

2. There exists $c > 0$ such that

$$\inf_{\widehat{\text{dgm}}_n} \sup_{f \in \mathcal{F}(\alpha)} \mathbb{E} \left[d_b(\text{dgm}(\text{Filt}(\mathbb{X}_f)), \widehat{\text{dgm}}_n) \right] \geq cn^{-1/(k+\alpha)}$$

for n large enough. The infimum is taken over all possible estimators $\widehat{\text{dgm}}_n$ of $\text{dgm}(\text{Filt}(\mathbb{X}_f))$ based on n observations.

Remark 7 The paper by Singh et al. (2009) is more generally about adaptive level set estimation. For this problem, Singh et al. define an histogram based estimator. Let A_j denote the collection of cells, in a regular partition of $\chi = [0, 1]^k$ into hypercubes of dyadic side length 2^{-j} . Their estimator \hat{f} is the histogram $\hat{f}(A) = \hat{P}(A)/\mu(A)$, where $\hat{P}(A) = \sum_{i=1 \dots n} 1_{X_i \in A}$. For estimating the level set $G_\gamma := \{x | f(x) \geq \gamma\}$, they consider the estimator

$$\hat{G}_{\gamma, j} = \bigcup_{A \in A_j \mid \hat{f}(A) > \gamma} A.$$

It is proved by Singh et al. (2009) that $\hat{G}_{\gamma, \hat{j}}$ achieves optimal rates of convergence for estimating the level sets, with \hat{j} chosen in a data driven way. Concerning support estimation, they also show that $\hat{G}_{0, j}$ achieves optimal rates of convergence for estimating G_0 . We have seen that in this context it is also the case for the estimator \mathbb{X}_n . Since no knowledge of α is required for this last estimator, we thus prefer to use this simpler estimator in this context.

4.2 Minimax Optimal Rates of Convergence of Persistence Diagram Estimation for Singular Measures in \mathbb{R}^D

We now consider the case where the support of μ is a smooth submanifold of \mathbb{R}^D . As far as we know, rates of convergence for manifold estimation, namely for the estimation of the support of a singular probability measure supported on a Riemannian manifold of \mathbb{R}^D , have only been studied recently by Genovese et al. (2012b,a). These papers assume several noise models, which all could be considered in the context of persistence diagram estimation. However, for the sake of simplicity, we only study here the problem where no additional noise is observed, which is referred as the *noiseless model* in the first of these two papers.

As before, upper bounds given by Genovese et al. (2012b) on the rates of convergence for the support estimation in Hausdorff distance directly provide upper bounds on the rates of convergence of the persistence diagram of the support. Before giving the rates of convergence we first recall and discuss the assumptions of Genovese et al. (2012b).

For any $r > 0$ and any set $A \subset \mathbb{R}^D$, let $A \oplus r := \bigcup_{a \in A} B(a, r)$. Let $\Delta(\mathbb{X}_\mu)$ be the largest r such that each point in $\mathbb{X}_\mu \oplus r$ has a unique projection onto \mathbb{X}_μ , this quantity has been introduced by Federer (1959), it is called reach or condition number in the literature.

For a fixed positive integer $k < D$, for some fixed positive constants b, B, κ and for a fixed compact domain χ in \mathbb{R}^D , Genovese et al. (2012b) define the set of probability measures $\mathcal{H} := \mathcal{H}(d, A, B, \kappa, \chi)$ on χ satisfying the two following assumptions:

- $[H_1]$ The support of the measure μ is a compact Riemannian manifold \mathbb{X}_μ (included in χ) of dimension k whose reach satisfies

$$\Delta(\mathbb{X}_\mu) \geq \kappa. \tag{4.1}$$

- $[H_2]$ The measure μ is assumed to have a density g with respect to k -dimensional volume measure vol_k on \mathbb{X}_μ , such that

$$0 < A \leq \inf_{y \in \mathbb{X}_\mu} g(y) \leq \sup_{y \in \mathbb{X}_\mu} g(y) \leq B < \infty. \tag{4.2}$$

These two assumptions can be easily connected to the standard assumption. Indeed, according to Niyogi et al. (2008) and using $[H_1]$, for all $r \leq \kappa$ there exists some constant $C > 0$ such that for any $x \in \mathbb{X}_\mu$, we have

$$\begin{aligned} vol_k(B(x, r) \cap \mathbb{X}_\mu) &\geq C \left(1 - \frac{r^2}{4\kappa^2}\right)^{k/2} r^k \\ &\geq C' r^k \end{aligned}$$

and the same holds for μ according to $[H_2]$. Under these two assumptions, μ satisfies the standard assumption with $b = k$. Thus, if we take $\hat{\mathbb{X}}_n$ for estimating the support \mathbb{X}_μ in this context, we obtain a rate of convergence upper bounded by $(\frac{\log n}{n})^{1/k}$ both for support and persistence diagram estimation. Nevertheless, this rate is not minimax optimal for estimating the support on the spaces \mathcal{H} as shown by Genovese et al. (2012b, Theorem 2). The correct minimax rate is $n^{-2/k}$ and the same is true for the persistence diagram estimation, as stated in the following proposition. However, the achievement of this optimal rate relies on a “theoretical” estimator proposed by Genovese et al. (2012b) that can not be computed in practice.

Proposition 8 *Assume that we observe an n -sample under the previous assumptions, then there exist two constants C and C' depending only on \mathcal{H} such that*

$$Cn^{-2/k} \leq \inf_{\widehat{\text{dgm}}_n} \sup_{\mu \in \mathcal{H}} \mathbb{E} \left[d_b(\text{dgm}(\text{Filt}(\mathbb{X}_\mu)), \widehat{\text{dgm}}_n) \right] \leq C'n^{-2/k} \tag{4.3}$$

where the infimum is taken over all the estimators of the persistence diagram.

5. Experiments

A series of experiments were conducted in order to illustrate the behavior of the persistence diagrams under sampling of metric spaces endowed with a probability measure and to compare the convergence performance obtained in practice with the theoretical results obtained in the previous sections.

5.1 Spaces and Data

We consider four different metric spaces, denoted \mathbb{M}_1 , \mathbb{M}_2 , \mathbb{M}_3 and \mathbb{M}_4 hereafter, that are described below.

\mathbb{M}_1 (**Lissajous curve in \mathbb{R}^2**): the planar curve with the parametric equations $x(t) = \sin(3t + \pi/2)$, $y(t) = \sin(2t)$, $t \in [0, 2\pi]$ (see Figure 6, left). Its metric is the restriction of the Euclidean metric in \mathbb{R}^2 and it is endowed with the push forward by the parametrization of the uniform measure on the interval $[0, 2\pi]$.

\mathbb{M}_2 (**sphere in \mathbb{R}^3**): the unit sphere in \mathbb{R}^3 (see Figure 6, center). Its metric is the restriction of the Euclidean metric in \mathbb{R}^3 and it is endowed with the uniform area measure on the sphere.

\mathbb{M}_3 (**torus in \mathbb{R}^3**): the torus of revolution in \mathbb{R}^3 with the parametric equations $x(u, v) = (5 + \cos(u)) \cos(v)$, $y(u, v) = (5 + \cos(u)) \sin(v)$ and $z(u, v) = \sin(u)$, $(u, v) \in [0, 2\pi]^2$ (see Figure 6, right). Its metric is the restriction of the Euclidean metric in \mathbb{R}^3 and it is endowed with the push forward by the parametrization of the uniform measure on the square $[0, 2\pi]^2$.

\mathbb{M}_4 (**rotating shape space**): for this space we used a 3D character from the SCAPE database (Anguelov et al., 2005) and considered all the images of this character from a view rotating around it. We converted these images in gray color and resized these images to $300 \times 400 = 120,000$ pixels (see Figure 7). Each is then identified with a point in $\mathbb{R}^{120,000}$ where the i^{th} coordinate is the level of gray of the i^{th} pixel. Moreover, we normalized these images by projecting them on the unit sphere in $\mathbb{R}^{120,000}$. The metric space \mathbb{M}_4 is the obtained subset of the unit sphere with the restriction of the Euclidean metric in $\mathbb{R}^{120,000}$. As it is parametrized by a circular set of views, it is endowed with the push forward of the uniform measure on the circle.

5.2 The Experiments

From each of the measured metric spaces \mathbb{M}_1 , \mathbb{M}_2 , \mathbb{M}_3 and \mathbb{M}_4 we sampled k sets of n points for different values of n from which we computed persistence diagrams for different geometric complexes (see Table 1). For \mathbb{M}_1 , \mathbb{M}_2 and \mathbb{M}_3 we have computed the persistence diagrams for the 1 or 2-dimensional homology of the α -complex built on top of the sampled sets. As α -complexes have the same homotopy type as the corresponding union of balls, these persistence diagrams are the ones of the distance function to the sampled point set (Edelsbrunner, 1995). So, for each n we computed the average bottleneck distance between the obtained diagrams and the persistence diagram of the distance to the metric space from




Figure 6: The spaces \mathbb{M}_1 , \mathbb{M}_2 and \mathbb{M}_3 .




Figure 7: Images sampled from the space \mathbb{M}_4 .

which the points were sampled. For \mathbb{M}_4 , as it is embedded in a very high dimensional space, computing the α -complex is practically out of reach. So we have computed the persistence diagrams for the 1-dimensional homology of the Vietoris-Rips complex built on top of the sampled sets. The obtained results are described and discussed below.

- **Results for \mathbb{M}_1 :** we approximated the 1-dimensional homology persistence diagram of the distance function to the Lissajous curve $\mathbf{dgm}(\mathbb{M}_1)$ by sampling \mathbb{M}_1 with 500,000 points and computing the persistence diagram of the corresponding α -complex. As the Hausdorff distance between our sample and \mathbb{M}_1 was of order 10^{-5} we obtained a sufficiently precise approximation of $\mathbf{dgm}(\mathbb{M}_1)$ for our purpose. $\mathbf{dgm}(\mathbb{M}_1)$ is represented in blue on the left of Figure 8. For each n , the average bottleneck distance between $\mathbf{dgm}(\mathbb{M}_1)$ and the persistence diagrams obtained for the $k = 300$ randomly sampled sets \mathbb{X}_n of size n has been used as an estimate $\hat{\mathbb{E}}$ of $\mathbb{E} \left[d_b(\mathbf{dgm}(C_\alpha(\mathbb{M}_1)), \mathbf{dgm}(C_\alpha(\hat{\mathbb{X}}_n))) \right]$ where C_α denotes the α -complex filtration. On Figure 8, right, $\log(\hat{\mathbb{E}})$ is plotted as a function of $\log(\log(n)/n)$. As expected, since the Lissajous curve is 1-dimensional, the points are close to a line of slope 1.
- **Results for \mathbb{M}_2 and \mathbb{M}_3 :** the persistence diagrams $\mathbf{dgm}(\mathbb{M}_2)$ and $\mathbf{dgm}(\mathbb{M}_3)$ of the distance functions to \mathbb{M}_2 and \mathbb{M}_3 are known exactly and are represented in blue on Figures 9 and 10, left, respectively. Notice that we considered the 2-dimensional homology for \mathbb{M}_2 and 1-dimensional homology for \mathbb{M}_3 . For $i = 2, 3$ and for each n , the average bottleneck distance between $\mathbf{dgm}(\mathbb{M}_i)$ and the persistence diagrams obtained for the $k = 100$ randomly sampled sets \mathbb{X}_n of size n has been used as an estimate $\hat{\mathbb{E}}$ of $\mathbb{E} \left[d_b(\mathbf{dgm}(C_\alpha(\mathbb{M}_i)), \mathbf{dgm}(C_\alpha(\hat{\mathbb{X}}_n))) \right]$ where C_α denotes the α -complex filtration. $\log(\hat{\mathbb{E}})$ is plotted as a function of $\log(\log(n)/n)$ on Figures 9 and 10, right. As expected, since the sphere and the torus are 2-dimensional, the points are close to a line of slope 1/2.
- **Results for \mathbb{M}_4 :** As in that case we do not know the persistence diagram of the Vietoris-Rips filtration built on top of \mathbb{M}_4 , we only computed the 1-dimensional homology persistence diagrams of the Vietoris-Rips filtrations built on top of 20 sets of 250 points each, randomly sampled on \mathbb{M}_4 . All these diagrams have been plotted on the same Figure 11, left. The right of Figure 11 represents a 2D embedding of one of the 250 points sampled data set using the Multidimensional Scaling algorithm (MDS). Since \mathbb{M}_4 is a set of images taken according a rotating point of view, it carries a cycle structure. This structure is reflected in the persistence diagrams that all have one point which is clearly off the diagonal. Notice also a second point off the diagonal which is much closer to it and that probably corresponds to the pinching in \mathbb{M}_4 visible at the bottom left of the MDS projection.

6. Discussion and Future Works

In previous works, the use of persistent homology in TDA has been mainly considered with a deterministic approach. As a consequence persistence diagrams were usually used as exploratory tools to analyze the topological structure of data. In this paper, we propose

Space	k (sampled sets for each n)	n range	Geometric complex
\mathbb{M}_1	300	[2100 : 100 : 3000]	α -complex
\mathbb{M}_2	100	[12000 : 1000 : 21000]	α -complex
\mathbb{M}_3	100	[4000 : 500 : 8500]	α -complex
\mathbb{M}_4	20	250	Vietoris-Rips complex

Table 1: Sampling parameters and geometric complexes where $[n_1 : h : n_2]$ denotes the set of integers $\{n_1, n_1 + h, n_1 + 2h, \dots, n_2\}$.




Figure 8: Convergence rate for the persistence diagram of the α -filtration built on top of points sampled on \mathbb{M}_1 . Left: in blue the persistence diagram $\text{dgm}(\mathbb{M}_1)$ of the distance to \mathbb{M}_1 (1-dimensional homology); in red a persistence diagram of the α -filtration built on top of $n = 2100$ points randomly sampled on \mathbb{M}_1 . Right: the x -axis is $\log(\log(n)/n)$ where n is the number of points sampled on \mathbb{M}_1 . The y -axis is the log of the estimated expectation of the bottleneck distance between the diagram obtained from an α -filtration built on top of n points sampled on \mathbb{M}_1 and $\text{dgm}(\mathbb{M}_1)$.




Figure 9: Convergence rate for the persistence diagram of the α -filtration built on top of points sampled on \mathbb{M}_2 . Left: in blue the persistence diagram $\text{dgm}(\mathbb{M}_2)$ of the distance to \mathbb{M}_2 (2-dimensional homology); in red a persistence diagram of the α -filtration built on top of $n = 12000$ points randomly sampled on \mathbb{M}_2 . Right: the x -axis is $\log(\log(n)/n)$ where n is the number of points sampled on \mathbb{M}_2 . The y -axis is the log of the estimated expectation of the bottleneck distance between the diagram obtained from an α -filtration built on top of n points sampled on \mathbb{M}_2 and $\text{dgm}(\mathbb{M}_2)$.




Figure 10: Convergence rate for the persistence diagram of the α -filtration built on top of points sampled on \mathbb{M}_3 . Left: in blue the persistence diagram $\text{dgm}(\mathbb{M}_3)$ of the distance to \mathbb{M}_3 (1-dimensional homology); in red a persistence diagram of the α -filtration built on top of $n = 14000$ points randomly sampled on \mathbb{M}_3 . Right: the x -axis is $\log(\log(n)/n)$ where n is the number of points sampled on \mathbb{M}_3 . The y -axis is the log of the estimated expectation of the bottleneck distance between the diagram obtain from α -filtration built on top of n points sampled on \mathbb{M}_3 and $\text{dgm}(\mathbb{M}_3)$.




Figure 11: Left: on the same figure the 1-dimensional homology persistence diagrams of the Vietoris-Rips filtration of 20 sets of 250 points sampled on M_4 . Right: the plot of the embedding of M_4 in \mathbb{R}^2 using MDS.

a rigorous framework to study the statistical properties of persistent homology and more precisely we give a general approach to study the rates of convergence for the estimation of persistence diagrams. The results we obtain open the door to a rigorous use of persistence diagrams in statistical framework. Our approach, consisting in reducing persistence diagram estimation to another more classical estimation problem (here support estimation) is based upon recently proven stability results in persistence theory that are very general.

In this paper, the persistence diagram of interest is the one of the support of the measure μ according which the data points are sampled. As a consequence, if the data points are sampled according to some perturbed measure ν whose support is not close to the one of μ then the estimator obviously non longer converges to the diagram of the support of μ . A first solution to overcome this problem is to plug denoising methods for support estimation (with respect to Hausdorff distance), such as deconvolution methods (Meister, 2009), to our approach.

Building on ideas developed by Chazal et al. (2011) and Caillerie et al. (2011), more satisfactory solutions have been recently proposed by Chazal et al. (2014a,b) that allow to infer persistent homology information from data corrupted by different kind of noise.

In another direction, an interesting representation of persistence diagrams as elements of a Hilbert space has recently been proposed by Bubenik (2012). Our results easily extend to this representation of persistence diagrams called *persistence landscapes*. Following this promising point of view, we also intend to adapt classical kernel-based methods with kernels carrying topological information.

Acknowledgments

The authors acknowledge the support of the European project CG-Learning EC contract No. 255827, the ANR projects GIGA (ANR-09-BLAN-0331-01) and TopData (ANR-13-BS01-0008), the ERC project GUDHI and the Google Faculty Research Award.

Appendix A. Lecam’s Lemma

The version of Lecam’s Lemma given below is from Yu (1997) (see also Genovese et al., 2012a). Recall that the total variation distance between two distributions P_0 and P_1 on a measured space $(\mathcal{X}, \mathcal{B})$ is defined by

$$\text{TV}(P_0, P_1) = \sup_{B \in \mathcal{B}} |P_0(B) - P_1(B)|.$$

Moreover, if P_0 and P_1 have densities p_0 and p_1 for the same measure λ on \mathcal{X} , then

$$\text{TV}(P_0, P_1) = \frac{1}{2} \ell_1(p_0, p_1) := \int_{\mathcal{X}} |p_0 - p_1| d\lambda.$$

Lemma 9 *Let \mathcal{P} be a set of distributions. For $P \in \mathcal{P}$, let $\theta(P)$ take values in a metric space (\mathbb{X}, ρ) . Let P_0 and P_1 in \mathcal{P} be any pair of distributions. Let X_1, \dots, X_n be drawn i.i.d. from some $P \in \mathcal{P}$. Let $\hat{\theta} = \hat{\theta}(X_1, \dots, X_n)$ be any estimator of $\theta(P)$, then*

$$\sup_{P \in \mathcal{P}} \mathbb{E}_{P^n} \rho(\theta, \hat{\theta}) \geq \frac{1}{8} \rho(\theta(P_0), \theta(P_1)) [1 - \text{TV}(P_0, P_1)]^{2n}.$$

Appendix B. Proofs

All the proofs of the paper are given in this section.

B.1 Proof of Theorem 2

The proof follows the lines of the proof of Cuevas and Rodríguez-Casal (2004, Theorem 3). The only point to be checked is that the covering number of \mathbb{X}_μ under the (a, b) -standard assumption can be controlled as when $b = d \in \mathbb{N}$, the rest of the proof being unchanged.

The covering number $\text{cv}(\mathbb{X}_\mu, r)$ of \mathbb{X}_μ is the minimum number of balls of radius r that are necessary to cover \mathbb{X}_μ :

$$\text{cv}(\mathbb{X}_\mu, r) = \min \left\{ k \in \mathbb{N}^* : \exists (x_1, \dots, x_k) \in (\mathbb{X}_\mu)^k \text{ such that } \mathbb{X}_\mu = \bigcup_{i=1}^k B(x_i, r) \right\}.$$

The packing number $\text{pk}(\mathbb{X}_\mu, r)$ is the maximum number of balls of radius r that can be packed in \mathbb{X}_μ without overlap:

$$\text{pk}(\mathbb{X}_\mu, r) = \max \left\{ \begin{array}{l} k \in \mathbb{N}^* : \exists (x_1, \dots, x_k) \in (\mathbb{X}_\mu)^k \text{ such that } B(x_i, r) \subset \mathbb{X}_\mu \\ \text{and, } \forall i \neq j, B(x_i, r) \cap B(x_j, r) = \emptyset \end{array} \right\}$$

The covering and packing numbers are related by the following inequalities (see for instance Massart, 2007, p. 71):

$$\text{pk}(\mathbb{X}_\mu, 2r) \leq \text{cv}(\mathbb{X}_\mu, 2r) \leq \text{pk}(\mathbb{X}_\mu, r). \tag{B.1}$$

Lemma 10 *Assume that the probability μ satisfies a standard (a, b) -assumption. Then for any $r > 0$ we have*

$$\text{pk}(\mathbb{X}_\mu, r) \leq \frac{1}{ar^b} \vee 1 \text{ and } \text{cv}(\mathbb{X}_\mu, r) \leq \frac{2^b}{ar^b} \vee 1.$$

Proof The result is trivial for $r \geq a^{-1/b}$. Let $r < a^{-1/b}$ and let $p = \text{pk}(\mathbb{X}_\mu, r)$, we choose a maximal packing $B_1 = B(x_1, r), \dots, B_p = B(x_p, r)$ of \mathbb{X}_μ . Since the balls of the packing are pairwise disjoint and μ is a probability measure we have $\sum_{i=1}^p \mu(B_i) \leq 1$. Using that $\mu(B_i) \geq ar^b$ we obtain that $par^b \leq \sum_{i=1}^p \mu(B_i) \leq 1$ from which we get the upper bound on $\text{pk}(\mathbb{X}_\mu, r)$. Since from (B.1) we have $\text{cv}(\mathbb{X}_\mu, r) \leq \text{pk}(\mathbb{X}_\mu, r/2)$ we immediately deduce the upper bound on $\text{cv}(\mathbb{X}_\mu, r)$. \blacksquare

B.2 Proof of Proposition 4

We first prove the upper bound.

B.2.1 UPPER BOUND

According to Corollary 3, thanks to Fubini we have

$$\mathbb{E} \left[d_b(\text{dgm}(\text{Filt}(\mathbb{X}_\mu)), \text{dgm}(\text{Filt}(\widehat{\mathbb{X}}_n))) \right] \leq \int_{\varepsilon > 0} \mathbb{P} \left[d_b(\text{dgm}(\text{Filt}(\mathbb{X}_\mu)), \text{dgm}(\text{Filt}(\widehat{\mathbb{X}}_n))) > \varepsilon \right] d\varepsilon$$

Let $\varepsilon_n = 4 \left(\frac{\log n}{an} \right)^{1/b}$. By bounding the probability inside this integral by one on $[0, \varepsilon_n]$, we find that:

$$\begin{aligned} \mathbb{E} \left[d_b(\text{dgm}(\text{Filt}(\mathbb{X}_\mu)), \text{dgm}(\text{Filt}(\widehat{\mathbb{X}}_n))) \right] &\leq \varepsilon_n + \int_{\varepsilon > \varepsilon_n} \frac{8^b}{a} \varepsilon^{-b} \exp(-na\varepsilon^b/4^b) d\varepsilon \\ &\leq \varepsilon_n + \frac{4n2^b}{b} (na)^{-1/b} \int_{u \geq \log n} u^{1/b-2} \exp(-u) du. \end{aligned}$$

Now, if $b \geq \frac{1}{2}$ then $u^{1/b-2} \leq (\log n)^{1/b-2}$ for any $u \geq \log n$ and then

$$\begin{aligned} \mathbb{E} \left[d_b(\text{dgm}(\text{Filt}(\mathbb{X}_\mu)), \text{dgm}(\text{Filt}(\widehat{\mathbb{X}}_n))) \right] &\leq \varepsilon_n + 4 \frac{2^b}{b} \left(\frac{\log n}{n} \right)^{1/b} (\log n)^{-2} \\ &\leq C_1(a, b) \left(\frac{\log n}{n} \right)^{1/b} \end{aligned} \tag{B.2}$$

where the constant $C_1(a, b)$ only depends on a and b . If $0 < b < \frac{1}{2}$, let $p := \lfloor \frac{1}{b} \rfloor$ and then

$$\begin{aligned} \int_{u \geq u_n := \log n} u^{1/b-2} \exp(-u) du &= u_n^{1/b-2} \exp(u_n) + \left(\frac{1}{b} - 2 \right) u_n^{1/b-3} \exp(u_n) + \dots + \\ &\quad + \prod_{i=2}^p \left(\frac{1}{b} - i \right) u_n^{1/b-p} \exp(u_n) + \int_{u \geq \log n} u^{1/b-p-1} \exp(-u) du \\ &\leq C_2(a, b) \frac{(\log n)^{1/b-2}}{n} \end{aligned}$$

where $C_2(a, b)$ only depends on a and b . Thus (B.2) is also satisfied for $b < \frac{1}{2}$ and the upper bound is proved.

B.2.2 LOWER BOUND

To prove the lower bound, it will be sufficient to consider two Dirac distributions. We take for $P_{0,n} = P_x$ the Dirac distribution on $\mathbb{X}_0 := \{x\}$ and it is clear that $P_0 \in \mathcal{P}(a, b, \mathbb{M})$. Let $P_{1,n}$ be the distribution $\frac{1}{n}\delta_{x_n} + (1 - \frac{1}{n})P_0$. The support of $P_{1,n}$ is denoted $\mathbb{X}_{1,n} := \{x\} \cup \{x_n\}$. Note that for any $n \geq 2$ and any $r \leq \rho(x, x_n)$:

$$P_{1,n}(B(x, r)) = 1 - \frac{1}{n} \geq \frac{1}{2} \geq \frac{1}{2\rho(x, x_n)^b} r^b \geq ar^b$$

and

$$P_{1,n}(B(x_n, r)) = \frac{1}{n} = \frac{1}{n\rho(x, x_n)^b} r^b \geq ar^b.$$

Moreover, for $r > \rho(x, x_n)$, $P_{1,n}(B(0, r)) = P_{1,n}(B(x_n, r)) = 1$. Thus for any $r > 0$ and any $x \in \mathbb{X}_{1,n}$:

$$P_{1,n}(B(x, r)) \geq ar^b \wedge 1$$

and $P_{1,n}$ also belongs to $\mathcal{P}(a, b, \mathbb{M})$.

The probability measure P_0 is absolutely continuous with respect to $P_{1,n}$ and the density of P_0 with respect to $P_{1,n}$ is $p_{0,n} := \frac{n}{n-1} \mathbb{1}_{\{x\}}$. Then

$$\begin{aligned} TV(P_0, P_{1,n}) &= \int_{\mathbb{M}} \left| 1 - \frac{n}{n-1} \mathbb{1}_{\{x\}} \right| dP_{1,n} \\ &= \frac{2}{n}. \end{aligned}$$

Next, $[1 - TV(P_0, P_{1,n})]^{2n} = (1 - \frac{2}{n})^{2n} \rightarrow e^{-4}$ as n tends to infinity. It remains to compute $d_b(\text{dgm}(\text{Filt}(\mathbb{X}_0)), \text{dgm}(\text{Filt}(\mathbb{X}_{1,n})))$. We only consider here the Rips case, the other filtrations can be treated in a similar way. The barcode of $\text{Filt}(\mathbb{X}_0)$ is composed of only one segment $(0, +\infty)$ for the 0-cycles. The barcode of $\text{Filt}(\mathbb{X}_{1,n})$ is composed of the segment of $\text{Filt}(\mathbb{X}_0)$ and one more 0-cycle : $(0, \rho(x, x_n))$. Thus we have:

$$\begin{aligned} d_b(\text{dgm}(\text{Filt}(\mathbb{X}_0)), \text{dgm}(\text{Filt}(\mathbb{X}_{1,n}))) &= d_\infty(\Delta, (0, \rho(x, x_n))) \\ &= \frac{\rho(x, x_n)}{2}. \end{aligned}$$

The proof is then complete using Lecam’s Lemma (Lemma 9).

B.3 Proof of Theorem 5

Let A be the interval $[0, 1]$ and c a positive constant to be chosen further. We consider k “holes” H_i of length $c \frac{\log n}{n}$ each, distant enough from each other that we remain $(\frac{1}{2}, 1)$ -standard when we remove any number of H_i from A , which is possible as long as $k c \frac{\log n}{n} < \frac{1}{2}$. We denote $A_i = A \setminus H_i$. For $I \subset \{1, \dots, k\}$, $A_I = \bigcap_{i \in I} A_i$, $B = A^n$, $B_I = A_I^n$. Denoting the uniform measure on $[0, 1]$ by λ , we have $\lambda^{\otimes n}(B_I) = (1 - |I|c \frac{\log n}{n})^n \sim n^{-|I|c}$.

The main idea is that when sampling n points from A , most likely (at least) one of the H_i contains no points. Without points in H_i , the estimator cannot distinguish A from A_i , but since those two have diagrams at distance $c \frac{\log n}{n}$, this gives a bound on the quality of the estimator. The technical difficulty is that several H_i can be empty at the same time.

For a given n , let $\widehat{\text{dgm}}_n$ be an estimator of persistence diagram of the sampling distribution support. Assume for the moment that $\widehat{\text{dgm}}_n$ satisfies

$$\sup_{\mu \in \mathcal{P}(\frac{1}{2}, 1, [0, 1])} \mathbb{E} \left[d_b \left(\text{dgm}(\text{Filt}(\mathbb{X}_\mu)), \widehat{\text{dgm}}_n \right) \right] \leq \frac{1}{2} \frac{\log n}{n}. \tag{B.3}$$

Under this assumption, our goal is to lower bound $\mathbb{E} \left[d_b \left(\text{dgm}(\text{Filt}(\mathbb{X}_\mu)), \widehat{\text{dgm}}_n \right) \right]$ for μ equal to λ the uniform distribution on $[0, 1]$. The estimator $\widehat{\text{dgm}}_n$ can also be written as $\widehat{\text{dgm}}_n = g(X_1, \dots, X_n)$ where g is a measurable application from B into the set of persistence diagrams. First we note that when the observations are sampled according to λ :

$$\begin{aligned} \mathbb{E} \left[d_b \left(\text{dgm}(\text{Filt}(\mathbb{X}_\lambda)), \widehat{\text{dgm}}_n \right) \right] &= \int_B d_b(\text{dgm}(A), g(x)) d\lambda^{\otimes n}(x) \\ &\geq \int_{\bigcup_{1 \leq i \leq k} B_i} d_b(\text{dgm}(A), g(x)) d\lambda^{\otimes n}(x) =: R_n \end{aligned}$$

so it will be sufficient to bound this last integral. Applying Inequality B.3 to the uniform distribution μ_I on the set A_I , we find that

$$\frac{1}{\lambda^{\otimes n}(B_I)} \int_{B_I} d_B(\text{dgm}(A_I), g(x)) d\lambda^{\otimes n}(x) \leq \frac{1}{2} \frac{\log n}{n}.$$

Let $M_I := \int_{B_I} d_b(\text{dgm}(A), g(x)) d\lambda^{\otimes n}(x)$. Knowing that $d_b(\text{dgm}(A), \text{dgm}(A_I)) = c \frac{\log n}{n}$ and using the triangular inequality, we find that

$$\left| \frac{M_I}{\lambda^{\otimes n}(B_I)} - c \frac{\log n}{n} \right| \leq \frac{1}{2} \frac{\log n}{n}. \tag{B.4}$$

By applying the inclusion-exclusion principle for the union of the B_i 's, we find that $R_n \geq R_{1,n} - R_{2,n}$ where $R_{1,n} = \sum_i M_i$ and $R_{2,n} = \sum_{i < j} M_{\{i,j\}}$. According to (B.4) we have

$$R_{1,n} \geq k \left(c - \frac{1}{2} \right) \frac{\log n}{n} \left(1 - c \frac{\log n}{n} \right)^n$$

and

$$R_{2,n} \leq \frac{k(k-1)}{2} \left(c + \frac{1}{2} \right) \frac{\log n}{n} \left(1 - 2c \frac{\log n}{n} \right)^n.$$

We take $c = \frac{3}{4}$. Then the lower bound of $R_{1,n}$ is equivalent to $\frac{k}{4} \frac{\log n}{n} n^{-3/4}$ and

$$\frac{R_{2,n}}{R_{1,n}} \leq 5 \frac{k-1}{2} \frac{\left(1 - \frac{3}{2} \frac{\log n}{n} \right)^n}{\left(1 - \frac{3}{4} \frac{\log n}{n} \right)^n} \sim_{n \rightarrow \infty} 5 \frac{k-1}{2} n^{-\frac{3}{4}}.$$

We take $k = k_n := \left\lceil \frac{n^{3/4}}{5} \right\rceil$ in order to have $\frac{R_{2,n}}{R_{1,n}}$ tending to $\frac{1}{2}$ as n tends to infinity. We thus have

$$\liminf_n \frac{n}{\log n} R_n \geq \frac{1}{40}.$$

Moreover, note that

$$k_n c \frac{\log n}{n} \sim \frac{3}{20} n^{3/4} \frac{\log n}{n}$$

which is smaller than $\frac{1}{2}$ for n large enough so the $(\frac{1}{2}, 1)$ -standard assumption is verified for n large enough.

To summarize, for any n and any estimator $\widehat{\text{dgm}}_n$: either $\widehat{\text{dgm}}_n$ satisfies (B.3) and then

$$\sup_{\mu \in \mathcal{P}(\frac{1}{2}, 1, [0, 1])} \mathbb{E} \left[d_b \left(\text{dgm}(\text{Filt}(\mathbb{X}_\mu)), \widehat{\text{dgm}}_n \right) \right] \geq R_n,$$

or

$$\sup_{\mu \in \mathcal{P}(\frac{1}{2}, 1, [0, 1])} \mathbb{E} \left[d_b \left(\text{dgm}(\text{Filt}(\mathbb{X}_\mu)), \widehat{\text{dgm}}_n \right) \right] \geq \frac{1}{2} \frac{\log n}{n}.$$

Finally we have that for any estimator $\widehat{\text{dgm}}_n$:

$$\liminf_n \frac{n}{\log n} \sup_{\mu \in \mathcal{P}(\frac{1}{2}, 1, [0, 1])} \mathbb{E} \left[d_b \left(\text{dgm}(\text{Filt}(\mathbb{X}_\mu)), \widehat{\text{dgm}}_n \right) \right] \geq \min \left(\frac{1}{2}, \liminf_n \frac{n}{\log n} R_n \right) \geq \frac{1}{40}$$

and the theorem is proved.

B.4 Proofs for Section 4.1

Lemma 11 1. Under assumption [B], we have $d_H(G_0, \mathbb{X}_f) = 0$.

2. Under Assumptions [A] and [B], μ satisfies a standard assumption with $b = \alpha + k$ and with a depending on $\mathcal{F}(\alpha)$.

Proof First, note that we always have

$$\overset{\circ}{G}_0 \subset \mathbb{X}_f \subset \overline{G}_0. \tag{B.5}$$

Indeed, if $\overset{\circ}{G}_0 \cap (\chi \setminus \mathbb{X}_f)$ is non empty, let x be in the intersection. Then there exists $\varepsilon > 0$ such that $B(x, \varepsilon) \subset G_0$ and $B(x, \varepsilon) \subset (\chi \setminus \mathbb{X}_f)$ since \mathbb{X}_f is assumed to be closed. The first inclusion then gives that $\mu(B(x, \varepsilon)) > 0$ whereas the second inclusion gives that $\mu(B(x, \varepsilon)) = 0$. Thus $\overset{\circ}{G}_0 \cap (\chi \setminus \mathbb{X}_f)$ is empty, the second inclusion in (B.5) is obvious since \mathbb{X}_f is assumed to be closed.

Then,

$$\begin{aligned} d_H(\mathbb{X}_f, G_0) &= \max(\sup_{x \in \mathbb{X}_f} d(x, G_0), \sup_{x \in G_0} d(x, \mathbb{X}_f)) \\ &= \max(\sup_{x \in \mathbb{X}_f} d(x, \overline{G}_0), \sup_{x \in \overline{G}_0} d(x, \mathbb{X}_f)) \\ &= \sup_{x \in \overline{G}_0} d(x, \mathbb{X}_f) \\ &= \sup_{x \in \partial G_0} d(x, \mathbb{X}_f) \end{aligned} \tag{B.6}$$

where we use the continuity of the distance function for the second equality and (B.5) for the two last ones. It follows from assumption [B] that for any $x \in \partial G_0$, $d(x, \overset{\circ}{G}_0) = 0$. Thus $d(x, \mathbb{X}_f) = 0$ according to (B.5) and we have proved that (B.6) is equal to zero.

We now prove the second point of the Lemma. Let $x \in \overset{\circ}{G}_0$ and let $r > 0$ such that

$$\frac{r}{2} \left(1 \wedge \frac{1}{C_b} \right) < \varepsilon_0 \wedge \left(\frac{\delta_a}{C_a} \right)^{1/\alpha}. \tag{B.7}$$

According to Assumption [B], for $\varepsilon = \frac{r}{2} \left(1 \wedge \frac{1}{C_b} \right)$, there exists $y \in I_\varepsilon(G_0)$ such that $d(x, y) \leq C_b \varepsilon \leq \frac{r}{2}$. Then, there exists $z \in I_\varepsilon$ such that $y \in B(z, \varepsilon) \subset I_\varepsilon$. Since $\varepsilon \leq \frac{r}{2}$ we find that $B(z, \varepsilon) \subset B(x, r) \cap G_0$. Thus,

$$\begin{aligned} \mu(B(x, r)) &\geq \int_{B(z, \varepsilon)} f(u) \, d\lambda(u) \\ &\geq \int_{B(z, \varepsilon)} \delta_a \wedge C_a d(u, \partial G_0)^\alpha \, d\lambda(u) \\ &\geq C_a \int_{B(z, \varepsilon)} (\varepsilon - \|u - z\|)^\alpha \, d\lambda(u) \\ &\geq C_a s_{k-1} \int_0^\varepsilon (\varepsilon - r)^\alpha r^{k-1} \, dr \end{aligned}$$

where s_{k-1} denotes the surface area of the unit $k - 1$ -sphere of \mathbb{R}^k , and where we have used Assumption [A] for the second inequality and the fact $C_a \varepsilon^\alpha \leq \delta_a$ for the third one. Finally we find that for any r satisfying (B.7):

$$\begin{aligned} \mu(B(x, r)) &\geq \frac{C_a s_{k-1} (k-1)!}{(\alpha+1) \dots (\alpha+k)} \varepsilon^{\alpha+k} \\ &\geq \frac{C_a s_{k-1} (k-1)! (1 \wedge \frac{1}{C_b})^{\alpha+k}}{2^{\alpha+k} (\alpha+1) \dots (\alpha+k)} r^{\alpha+k} \end{aligned}$$

and we obtain that μ satisfies that standard assumption with $b = \alpha + k$. ■

B.4.1 PROOF OF PROPOSITION 6

The first point of the proposition is an immediate consequence of the first point of Theorem 4 together with Lemma 11. We now prove the lower bound by adapting some ideas from the proof of Proposition 3 in Singh et al. (2009) about the Hausdorff lower bound. At the price of loosing a logarithm term in the lower bound, we propose here a proof based on a two-alternative analysis.

The function f_0 is defined on χ as follows for $r_0 > 0$ small enough:

$$f_0 = \begin{cases} C_a \|x\|^\alpha & \text{if } \|x\| \leq r_0 \\ C_0 & \text{if } r_0 \leq \|x\| \leq 2r_0 \\ C_a (3r_0 - \|x\|)^\alpha & \text{if } 2r_0 \leq \|x\| \leq 3r_0 \\ 0 & \text{elsewhere} \end{cases}$$

where

$$C_0 = \frac{1 - C_a s_{k-1} r_0^{k+\alpha} (\frac{1}{k+\alpha} + I_\alpha)}{s_{k-1} r_0^k (2^k - 1)/k} \quad \text{with } I_\alpha = \int_2^3 k^{k-1} (3-u)^\alpha du.$$

For $n \geq 1$ let $\varepsilon_n := n^{-1/(k+\alpha)}$, the function $f_{1,n}$ is defined on χ by

$$f_{1,n} = \begin{cases} \|x\|^\alpha & \text{if } \varepsilon_n \leq \|x\| \leq r_0 \\ C_{1,n} & \text{if } r_0 \leq \|x\| \leq 2r_0 \\ C_a (3r_0 - \|x\|)^\alpha & \text{if } 2r_0 \leq \|x\| \leq 3r_0 \\ 0 & \text{elsewhere} \end{cases}$$

where

$$\begin{aligned} C_{1,n} &= \frac{1 - C_a s_{k-1} \left\{ r_0^{k+\alpha} (\frac{1}{k+\alpha} + I_\alpha) - \frac{\varepsilon_n^{k+\alpha}}{k+\alpha} \right\}}{s_{k-1} r_0^k (2^k - 1)/k} \\ &= C_0 + \frac{k C_a \varepsilon_n^{k+\alpha}}{(k + \alpha) r_0^k (2^k - 1)}. \end{aligned}$$

We assume that δ_a is small enough so that we can choose r_0 such that $\delta_a \leq C_0$ for n large enough. Then f_0 and $f_{1,n}$ are both densities and they both belong to $\mathcal{F}(\alpha)$ for n large enough. The support of $f_0 d\lambda$ is equal to $\mathbb{X}_0 := \bar{B}(0, 3r_0)$ whereas the support of $f_{1,n} d\lambda$ is equal to $\mathbb{X}_{1,n} = \bar{B}(0, 3r_0) \setminus \bar{B}(0, \varepsilon_n)$. Next,

$$\begin{aligned} \text{TV}(f_0 d\lambda, f_{1,n} d\lambda) &= \int_\chi |f_0 - f_{1,n}| dx \\ &= s_{k-1} C_a \int_0^{\varepsilon_n} r^{\alpha+k-1} dr + s_{k-1} \int_{r_0}^{2r_0} (C_{1,n} - C_0) r^{k-1} dr \\ &= \frac{2s_{k-1} C_a}{k + \alpha} \varepsilon_n^{k+\alpha} \end{aligned}$$

Note that $(1 - \text{TV}(f_0 d\lambda, f_{1,n} d\lambda))^{2n} \rightarrow \exp(-\frac{4s_{k-1} C_a}{k+\alpha})$ as n tends to infinity. It remains to compute $d_b(\text{dgm}(\text{Filt}(\mathbb{X}_0)), \text{dgm}(\text{Filt}(\mathbb{X}_{1,n})))$. We only consider here the Rips case, the other filtrations can be treated in a similar way. The barcode of $\text{Filt}(\mathbb{X}_0)$ is composed of only one segment $(0, +\infty)$ for the 0-cycles. The barcode of $\text{Filt}(\mathbb{X}_{1,n})$ is composed of the segment of $\text{Filt}(\mathbb{X}_0)$ and one more 1-cycle : $(0, 2\varepsilon_n)$. Thus we have:

$$\begin{aligned} d_b(\text{dgm}(\text{Filt}(\mathbb{X}_0)), \text{dgm}(\text{Filt}(\mathbb{X}_{1,n}))) &= d_\infty(\Delta, (0, \varepsilon)) \\ &= \varepsilon_n. \end{aligned}$$

We then finish the proof using Lecam’s Lemma.

B.5 Proof of Proposition 8

We only need to prove the lower bound since the upper bound is a direct corollary of Theorem 3 in Genovese et al. (2012b). To prove the lower bound, we may use the particular manifolds defined in Genovese et al. (2012a) and also used by the same authors for the proof of Theorem 2 in Genovese et al. (2012b). Without loss of generality, we assume that

$\chi = [-L, L]^D$ and that $\kappa < L/2$. For $\ell \leq L$, let M and M' be the two manifolds of χ defined by

$$M = [-\ell, \ell]^D \cap \{x \in \chi \mid x_{k+1} = \dots = x_D = 0\} \quad \text{and} \quad M' = 2\kappa e_{k+1} + M$$

where e_{k+1} is the $k + 1$ -th vector of the canonical basis in \mathbb{R}^D . We assume that ℓ is chosen so that $b < 2(2\ell)^{-k} < B$. Let μ_0 be the uniform measure on $\mathbb{X}_0 := M \cup M'$ and then $\mu_0 \in \mathcal{H}$.

According to Genovese et al. (2012a, Theorem 6), for $0 < \gamma < \kappa$, we can define a manifold M_γ which can be seen as a perturbation of M such that:

- $\Delta(M_\gamma) = \kappa$
- $d_H(M_\gamma, M) = \gamma$ and $d_H(M_\gamma, M') = 2\kappa - \gamma$
- If $A = \{x \in M_\gamma \mid x \notin M\}$ then $\mu_1(A) \leq C\gamma^{k/2}$ where $C > 0$ and where μ_1 is the uniform measure on $\mathbb{X}_1 := M_\gamma \cup M'$.

For small enough γ we see that μ_1 satisfies $[H_2]$ and thus $\mu_1 \in \mathcal{H}$.

As before, we only consider here filtrations of Rips complexes. The persistence diagrams of $\text{Filt}(\mathbb{X}_0)$ and $\text{Filt}(\mathbb{X}_1)$ are exactly the same except for the diagram of 0-cycles : the first filtration has a barcode with a segment $(0, 2\kappa)$ whereas the corresponding barcode for $\text{Filt}(\mathbb{X}_1)$ is $(0, 2\kappa - \gamma)$. Thus, $d_b(\text{Filt}(\mathbb{X}_0), \text{Filt}(\mathbb{X}_1)) = \gamma$. Moreover, $\text{TV}(\mu_0, \mu_1) \leq |\mu_0(A) - \mu_1(A)| \leq C\gamma^{k/2}$. Finally, we choose $\gamma = (1/n)^{k/2}$ as in the proof of Genovese et al. (2012b, Theorem 2) and we conclude using Lecam's Lemma.

References

- Dragomir Anguelov, Praveen Srinivasan, Daphne Koller, Sebastian Thrun, Jim Rodgers, and James Davis. SCAPE: shape completion and animation of people. In *Proceedings of SIGGRAPH*, pages 408–416, 2005.
- Sivaraman Balakrishnan, Alessandro Rinaldo, Don Sheehy, Aarti Singh, and Larry A. Wasserman. Minimax rates for homology inference. *Journal of Machine Learning Research - Proceedings Track*, 22:64–72, 2012.
- Ulrich Bauer, Axel Munk, Hannes Sieling, and Max Wardetzky. Persistent homology meets statistical inference - a case study: Detecting modes of one-dimensional signals. *ArXiv:1404.1214*, 2014.
- G erard Biau, Beno t Cadre, David M. Mason, and Bruno Pelletier. Asymptotic normality in density support estimation. *Electronic Journal of Probability*, 14:2617–2635, 2009.
- Peter Bubenik. Statistical topology using persistence landscapes. *ArXiv:*, July 2012.
- Peter Bubenik and Peter T Kim. A statistical approach to persistent homology. *Homology, Homotopy and Applications*, 9(2):337–362, 2007.
- Dmitri Burago, Yuri Burago, and Sergei Ivanov. *A Course in Metric Geometry*, volume 33. American Mathematical Society Providence, 2001.

- Claire Caillerie, Frédéric Chazal, Jérôme Dedecker, and Bertrand Michel. Deconvolution for the Wasserstein metric and geometric inference. *Electronic Journal of Statistics*, 5: 1394–1423, 2011.
- Gunnar Carlsson. Topology and data. *AMS Bulletin*, 46(2):255–308, 2009.
- Frédéric Chazal, David Cohen-Steiner, Marc Glisse, Leonidas J. Guibas, and Steve Y. Oudot. Proximity of persistence modules and their diagrams. In *SCG*, pages 237–246, 2009a.
- Frédéric Chazal, David Cohen-Steiner, Leonidas J. Guibas, Facundo Mémoli, and Steve Y. Oudot. Gromov-hausdorff stable signatures for shapes using persistence. *Computer Graphics Forum (proc. SGP 2009)*, pages 1393–1403, 2009b.
- Frédéric Chazal, David Cohen-Steiner, and Quentin Mérigot. Geometric inference for probability measures. *Foundations of Computational Mathematics*, 11(6):733–751, 2011.
- Frédéric Chazal, Vin de Silva, Marc Glisse, and Steve Oudot. The structure and stability of persistence modules. *ArXiv:1207.3674*, 2012a.
- Frédéric Chazal, Vin de Silva, and Steve Oudot. Persistence stability for geometric complexes. *ArXiv*, july 2012b.
- Frédéric Chazal, Leonidas J. Guibas, Steve Y. Oudot, and Primoz Skraba. Persistence-based clustering in Riemannian manifolds. *J. ACM*, 60(6):41:1–41:38, November 2013.
- Frédéric Chazal, Brittany Terese Fasy, Fabrizio Lecci, Bertrand Michel, Alessandro Rinaldo, and Larry Wasserman. Robust topological inference: Distance to a measure and kernel distance. *ArXiv:1412.7197*, 2014a.
- Frédéric Chazal, Brittany Terese Fasy, Fabrizio Lecci, Bertrand Michel, Alessandro Rinaldo, and Larry Wasserman. Subsampling methods for persistent homology. *ArXiv:1406.1901*, 2014b.
- David Cohen-Steiner, Herbert Edelsbrunner, and John Harer. Stability of persistence diagrams. *Discrete & Computational Geometry*, 37(1):103–120, 2007.
- Antonio Cuevas. Set estimation: another bridge between statistics and geometry. *Boletín de Estadística e Investigación Operativa*, 25(2):71–85, 2009.
- Antonio Cuevas and Ricardo Fraiman. A plug-in approach to support estimation. *The Annals of Statistics*, 25(6):2300–2312, 1997.
- Antonio Cuevas and Alberto Rodríguez-Casal. On boundary estimation. *Advances in Applied Probability*, 36(2):340–354, 2004.
- Antonio Cuevas, Ricardo Fraiman, and Beatriz Pateiro-López. On statistical properties of sets fulfilling rolling-type conditions. *Advances in Applied Probability*, 44(2):311–329, 2012.

- Vin De Silva and Robert Ghrist. Homological sensor networks. *Notices of the American Mathematical Society*, 54(1), 2007.
- Ernesto De Vito, Lorenzo Rosasco, and Alessandro Toigo. Learning sets with separating kernels. *Applied and Computational Harmonic Analysis*, 37(2):185 – 217, 2014.
- Luc Devroye and Gary L. Wise. Detection of abnormal behavior via nonparametric estimation of the support. *SIAM Journal on Applied Mathematics*, 38(3):480–488, 1980.
- Lutz Dümbgen and Günther Walther. Rates of convergence for random approximations of convex sets. *Advances in Applied Probability*, 28(2):384–393, 1996.
- Herbert Edelsbrunner. The union of balls and its dual shape. *Discrete & Computational Geometry*, 13(1):415–440, 1995.
- Herbert Edelsbrunner and John L Harer. *Computational Topology: an Introduction*. American Mathematical Soc., 2010.
- Herbert Edelsbrunner, David Letscher, and Afra Zomorodian. Topological persistence and simplification. *Discrete & Computational Geometry*, 28:511–533, 2002.
- Brittany Terese Fasy, Fabrizio Lecci, Alessandro Rinaldo, Larry Wasserman, Sivaraman Balakrishnan, Aarti Singh, et al. Confidence sets for persistence diagrams. *The Annals of Statistics*, 42(6):2301–2339, 2014.
- Herbert Federer. Curvature measures. *Transactions of the American Mathematical Society*, 93:418–491, 1959.
- Christopher R. Genovese, Marco Perone-Pacifico, Isabella Verdinelli, and Larry Wasserman. Minimax manifold estimation. *Journal of Machine Learning Research*, 13:1263–1291, july 2012a.
- Christopher R. Genovese, Marco Perone-Pacifico, Isabella Verdinelli, and Larry Wasserman. Manifold estimation and singular deconvolution under hausdorff loss. *The Annals of Statistics*, 40:941–963, 2012b.
- Allen Hatcher. *Algebraic Topology*. Cambridge Univ. Press, 2001.
- Peter M Kasson, Afra Zomorodian, Sanghyun Park, Nina Singhal, Leonidas J Guibas, and Vijay S Pande. Persistent voids: a new structural metric for membrane fusion. *Bioinformatics*, 23(14):1753–1759, 2007.
- Aleksandr Petrovič Korostelëv and Alexandre B. Tsybakov. *Minimax Theory of Image Reconstruction*, volume 82 of *Lecture Notes in Statistics*. Springer-Verlag, New York, 1993.
- Aleksandr Petrovič Korostelëv, Leopold Simar, and Alexandre B. Tsybakov. Efficient estimation of monotone boundaries. *Ann. Statist.*, 23(2):476–489, 1995.

- Pascal Massart. *Concentration Inequalities and Model Selection*. Springer, Berlin, 2007. Lectures from the 33rd Summer School on Probability Theory held in Saint-Flour, July 6–23, 2003.
- Alexander Meister. *Deconvolution Problems in Nonparametric Statistics*, volume 193. Springer, 2009.
- Yuriy Mileyko, Sayan Mukherjee, and John Harer. Probability measures on the space of persistence diagrams. *Inverse Problems*, 27(12), 2011.
- Partha Niyogi, Stephen Smale, and Shmuel Weinberger. Finding the homology of submanifolds with high confidence from random samples. *Discrete & Computational Geometry*, 39(1-3):419–441, March 2008.
- Partha Niyogi, Stephen Smale, and Shmuel Weinberger. A topological view of unsupervised learning from noisy data. *SIAM Journal on Computing*, 40(3):646–663, 2011.
- Alberto Rodríguez-Casal. Set estimation under convexity type assumptions. *Annales de l'Institut Henri Poincaré (B) Probability and Statistics*, 43(6):763 – 774, 2007.
- Aarti Singh, Clayton Scott, and Robert Nowak. Adaptive Hausdorff estimation of density level sets. *The Annals of Statistics*, 37(5B):2760–2782, 2009.
- Gurjeet Singh, Facundo Memoli, Tigran Ishkhanov, Guillermo Sapiro, Gunnar Carlsson, and Dario L Ringach. Topological analysis of population activity in visual cortex. *Journal of Vision*, 8(8), 2008.
- Alexandre B. Tsybakov. On nonparametric estimation of density level sets. *The Annals of Statistics*, 25(3):948–969, 1997.
- Alexandre B Tsybakov and Vladimir Zaiats. *Introduction to Nonparametric Estimation*, volume 11. Springer, 2009.
- Jianzhong Wang. *Geometric Structure of High-Dimensional Data and Dimensionality Reduction*. Springer, 2012.
- Shmuel Weinberger. The complexity of some topological inference problems. *Foundations of Computational Mathematics*, 14(6):1277–1285, 2014.
- Bin Yu. Assouad, Fano, and Le Cam. In *Festschrift for Lucien Le Cam*, pages 423–435. Springer, New York, 1997.
- Afra Zomorodian and Gunnar Carlsson. Computing persistent homology. *Discrete & Computational Geometry*, 33(2):249–274, 2005.