

Towards Stratification Learning through Homology Inference

Paul Bendich*, Bei Wang[†] and Sayan Mukherjee[‡]

August 13, 2010

Abstract

A topological approach to stratification learning is developed for point cloud data drawn from a stratified space. Given such data, our objective is to infer which points belong to the same strata. First we define a multi-scale notion of a stratified space, giving stratifications over a set of resolution levels. We then use methods derived from kernel and cokernel persistent homology to cluster the data points into different strata, and we prove a result which guarantees the correctness of our clustering, given certain topological conditions; some geometric intuition for these topological conditions is also provided. Our correctness result is then given a probabilistic flavor: we give bounds on the minimum number of sample points required to infer with high probability which points belong to the same strata. Finally, we give an explicit algorithm for the clustering, prove its correctness, and apply it to some simulated data.

1 Introduction

Manifold learning is a basic problem in geometry, topology, and statistical inference that has received a great deal of attention. The basic idea is as follows: given a point cloud of data sampled from a manifold in an ambient space \mathbb{R}^k , infer the underlying manifold. A limitation of the problem statement is that it does not apply to sets that are not manifolds. For example, we may consider the more general class of stratified spaces that can be decomposed into strata, which are manifolds of varying dimension, each of which fit together in some uniform way inside the higher dimensional space. In this paper, we study the following problem in stratification learning: given a point cloud sampled from a stratified space, how do we cluster the points so that points in the same cluster are in the same stratum, while points in different clusters are not? We attack the problem with algebro-topological tools; in particular, we define a notion of local equivalence relation via maps between relative homology groups, and we then attempt to infer this relation by using ideas derived from persistent homology [15].

*IST Austria, Klosterneuburg, Austria

[†]Department of Computer Science Duke University, Durham, NC and IST Austria, Klosterneuburg, Austria

[‡]Departments of Statistical Science, Computer Science, Mathematics, and Institute for Genome Sciences and Policy Duke University, Durham, NC

Prior Work Consistency in manifold learning has often been recast as a homology inference statement: as the number of points in a point cloud goes to infinity, the inferred homology converges to the true homology of the underlying space. Results of this nature have been given for manifolds [28, 29] and a large class of compact subsets of Euclidean space [6]. Stronger results in homology inference for closed subsets of a metric space are given in [11].

Geometric approaches to stratification inference have been developed. These include inference of a mixture of linear subspaces [25], mixture models for general stratified spaces [21], and generalized Principal Component Analysis (GPCA) [32] which was developed for dimension reduction for mixtures of manifolds.

The study of stratified spaces has long been a focus of pure mathematics; see, for example, [19, 33]. The problem of inference for the local homology groups of a sampled stratified space in a deterministic setting has been addressed in [3].

Results In this paper we propose an approach to stratification inference based on local homology inference; more specifically, based on inference of the kernels and cokernels of several maps between groups closely related to the multi-scale local homology groups for different pairs of points in the sample. The results in this paper are: (1) a topological definition of two points belonging to the same strata by assessing the multi-scale local structure of the points through kernel and cokernel persistent homology; (2) topological conditions on the point sample under which the topological characterization holds – we call this topological inference; (3) a geometric intuition of these topological conditions based on quantities related to reach and a function of gradients; (4) finite sample bounds for the minimum number of points in the sample required to state with high probability which points belong to the same strata; (5) an algorithm that computes which points belong to the same strata and a proof of correctness for some parts of this algorithm.

Outline We review the needed background in Section 2. In Section 3, we give the topological inference theorem and provide some geometric intuition. The probabilistic statements are provided in Section 4. We describe the clustering algorithm in Section 5; the correctness proof of the algorithm is contained in three Appendices, A through C. The main body of the paper closes with some further discussion in Section 7.

2 Background

We review necessary background on persistent homology and stratified spaces. The treatment here is mostly adapted from [5], although we present the material in slightly simplified form. We first describe general persistence modules, focusing mainly on those that arise from maps between homology groups induced by inclusions of topological spaces. We then discuss stratifications and their connection to the local homology groups of a topological space. We assume basic knowledge of homology itself, referring the reader to [27] or [22], or to [15] for a more computationally oriented treatment.

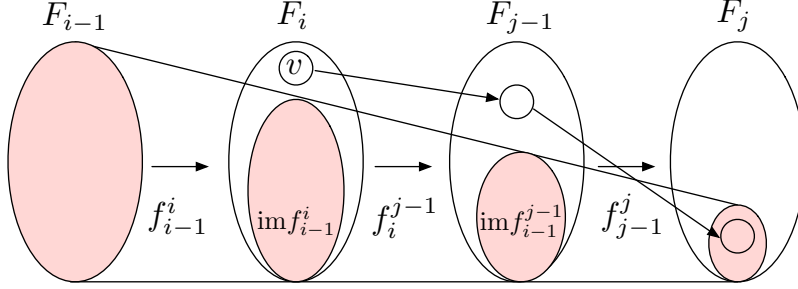


Figure 1: The vector v is born at level i and then it dies at level j .

2.1 Persistence Modules

In [5], the authors define persistence modules over an arbitrary commutative ring R with unity. For simplicity, we restrict immediately to the case $R = \mathbb{Z}/2\mathbb{Z}$. Let A be some subset of \mathbb{R} . Then a *persistence module* \mathcal{F}_A is a collection $\{F_\alpha\}_{\alpha \in A}$ of $\mathbb{Z}/2\mathbb{Z}$ -vector spaces, together with a family $\{f_\alpha^\beta : F_\alpha \rightarrow F_\beta\}_{\alpha \leq \beta \in A}$ of linear maps such that $\alpha \leq \beta \leq \gamma$ implies $f_\alpha^\gamma = f_\beta^\gamma \circ f_\alpha^\beta$. We will assume that the index set A is either \mathbb{R} or $\mathbb{R}_{\geq 0}$ and not explicitly state indices unless necessary.

A real number α is said to be a *regular value* of the persistence module \mathcal{F} if there exists some $\epsilon > 0$ such that, for all $\delta < \epsilon$, the maps $f_{\alpha-\delta}^{\alpha+\delta}$ are all isomorphisms. Otherwise we say that α is a *critical value* of the persistence module; if $A = \mathbb{R}_{\geq 0}$, then $\alpha = 0$ will always be considered to be a critical value. We say that \mathcal{F} is *tame* if it has a finite number of critical values and if all the vector spaces F_α are of finite rank. Any tame $\mathbb{R}_{\geq 0}$ -module \mathcal{F} must have a smallest non-zero critical value $\rho(\mathcal{F})$; we call this number the *feature size* of the persistence module.

Assume \mathcal{F} is tame and so we have a finite ordered list of critical values $0 = c_0 < c_1 < \dots < c_m$. We choose regular values $\{a_i\}_{i=0}^m$ such that $c_{i-1} < a_{i-1} < c_i < a_i$ for all $1 \leq i \leq m$, and we adopt the shorthand notation $F_i \equiv F_{a_i}$ and $f_i^j : F_i \rightarrow F_j$, for $0 \leq i \leq j \leq m$. A vector $v \in F_i$ is said to be *born* at level i if $v \notin \text{im } f_{i-1}^i$, and such a vector *dies* at level j if $f_i^j(v) \in \text{im } f_{i-1}^j$ but $f_i^{j-1}(v) \notin \text{im } f_{i-1}^{j-1}$. We then define $P^{i,j}$ to be the vector space of vectors that are born at level i and then subsequently die at level j , and $\beta^{i,j}$ denotes its rank. This is illustrated in Figure 1.

2.1.1 Persistence Diagrams

The information contained within a tame module \mathcal{F} is often compactly represented by a *persistence diagram*, $\text{Dgm}(\mathcal{F})$. This diagram is a multi-set of points in the extended plane. It contains $\beta^{i,j}$ copies of the points (c_i, c_j) , as well as infinitely many copies of each point along the major diagonal $y = x$. In Figure 3 the persistence diagrams for a curve and a point cloud sampled from it are displayed.

For any two points $u = (x, y)$ and $u' = (x', y')$ in the extended plane, we define $\|u - u'\|_\infty = \max\{|x - x'|, |y - y'|\}$. We define the *bottleneck distance*

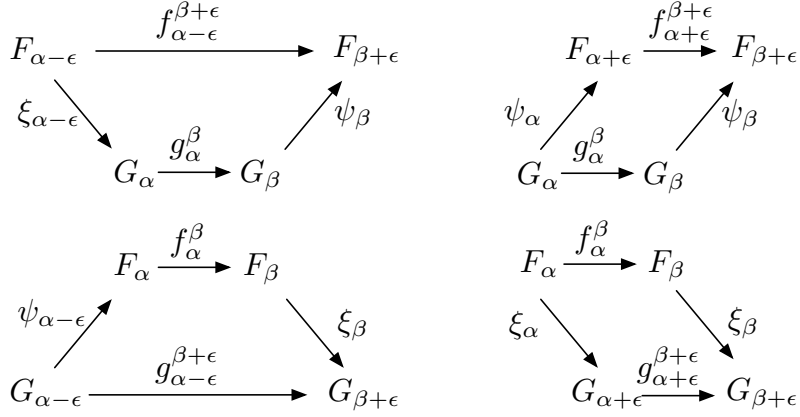


Figure 2: Commuting diagrams for strongly interleaving persistence modules.

between any two persistence diagrams D and D' to be:

$$d_B(D, D') = \inf_{\Gamma: D \rightarrow D'} \sup_{u \in D} \|u - \Gamma(u)\|_{\infty},$$

where Γ ranges over all bijections from D to D' . Under certain conditions which we now describe, persistence diagrams will be stable under the bottleneck distance.

Two persistence modules \mathcal{F} and \mathcal{G} are said to be *strongly ϵ -interleaved* if, for some positive ϵ , there exist two families $\{\xi_{\alpha} : F_{\alpha} \rightarrow G_{\alpha+\epsilon}\}_{\alpha}$ and $\{\psi_{\alpha} : G_{\alpha} \rightarrow F_{\alpha+\epsilon}\}_{\alpha}$ of linear maps which commute with the module maps $\{f_{\alpha}^{\beta}\}$ and $\{g_{\alpha}^{\beta}\}$ in the appropriate manner. More precisely, we require, for all $\alpha \leq \beta$, $f_{\alpha-\epsilon}^{\beta+\epsilon} = \psi_{\beta} \circ g_{\alpha}^{\beta} \circ \xi_{\alpha-\epsilon}$ and $\psi_{\beta} \circ g_{\alpha}^{\beta} = f_{\alpha+\epsilon}^{\beta+\epsilon} \circ \psi_{\alpha}$, as well as the two other equations obtained by exchanging the roles of f and g and ξ and ψ . This is shown in Figure 2.

We can now state the diagram stability result ([5], Theorem 4.4), that we will need later in this paper.

Theorem 2.1 (Diagram Stability Theorem) *Let \mathcal{F} and \mathcal{G} be tame persistence modules and $\epsilon > 0$. If \mathcal{F} and \mathcal{G} are strongly ϵ -interleaved, then*

$$d_B(\text{Dgm}(\mathcal{F}), \text{Dgm}(\mathcal{G})) \leq \epsilon.$$

When we wish to compute the persistence diagram associated to a module \mathcal{F} , it is often convenient to substitute another module \mathcal{G} , usually one defined in terms of simplicial complexes or other computable objects. The following theorem ([15], p.159) gives a condition under which this is possible.

Theorem 2.2 (Persistence Equivalence Theorem) *Given two persistence modules \mathcal{F} and \mathcal{G} , suppose there exist for each α isomorphisms $F_{\alpha} \cong G_{\alpha}$ which commute with the module maps, then $\text{Dgm}(\mathcal{F}) = \text{Dgm}(\mathcal{G})$.*

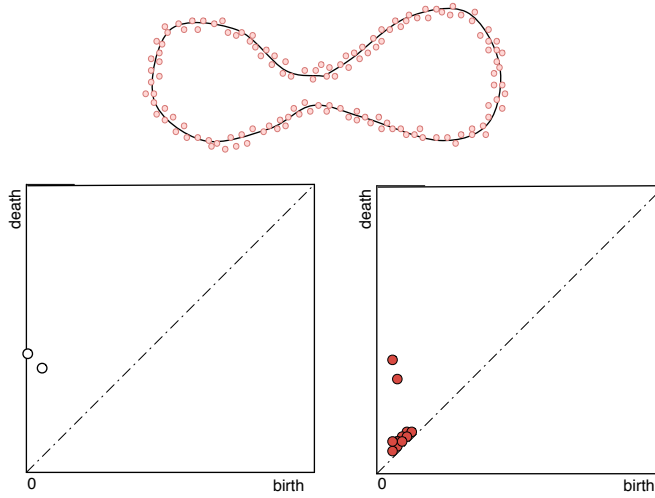


Figure 3: Illustration of a point cloud and its persistence diagram. Top: \mathbb{X} is the curve embedded as shown in the plane and U is the point cloud. Bottom left: the persistence diagram $\text{Dgm}_1(d_{\mathbb{X}})$; Bottom right: the persistence diagram $\text{Dgm}_1(d_U)$.

That is, if all the vertical maps are isomorphisms and all squares commute in the following diagram, then $\text{Dgm}(\mathcal{F}) = \text{Dgm}(\mathcal{G})$.

$$\begin{array}{ccccccc}
 \dots & \rightarrow & F_\alpha & \rightarrow & F_\beta & \rightarrow & \dots \\
 & & \uparrow \cong & & \uparrow \cong & & \\
 \dots & \rightarrow & G_\alpha & \rightarrow & G_\beta & \rightarrow & \dots
 \end{array}$$

2.1.2 (Co)Kernel Modules

Suppose now that we have two persistence modules \mathcal{F} and \mathcal{G} along with a family of maps $\{\phi_\alpha : F_\alpha \rightarrow G_\alpha\}$ which commute with the module maps – for every pair $\alpha \leq \beta$, we have $g_\alpha^\beta \circ \phi_\alpha = \phi_\beta \circ f_\alpha^\beta$. In other words, every square commutes in the diagram below:

$$\begin{array}{ccccccc}
 \dots & \rightarrow & F_\alpha & \xrightarrow{f_\alpha^\beta} & F_\beta & \rightarrow & \dots \\
 & & \downarrow \phi_\alpha & & \downarrow \phi_\beta & & \\
 \dots & \rightarrow & G_\alpha & \xrightarrow{g_\alpha^\beta} & G_\beta & \rightarrow & \dots
 \end{array}$$

Then, for each pair of real numbers $\alpha \leq \beta$, the restriction of f_α^β to $\ker \phi_\alpha$ maps into $\ker \phi_\beta$, giving rise to a new kernel persistence module, with persistence diagram denoted by $\text{Dgm}(\ker \phi)$. Similarly, we obtain a cokernel persistence module, with diagram $\text{Dgm}(\text{cok } \phi)$.

2.2 Homology

Our main examples of persistence modules all come from homology groups, either absolute or relative, and the various maps between them. Homology persistence modules can arise from families of topological spaces $\{\mathbb{X}_\alpha\}$, along with inclusions $\mathbb{X}_\alpha \hookrightarrow \mathbb{X}_\beta$ for all $\alpha \leq \beta$. Whenever we have such a family, the inclusions induce maps $H_j(\mathbb{X}_\alpha) \rightarrow H_j(\mathbb{X}_\beta)$, for each homological dimension $j \geq 0$, and hence we have persistence modules for each j . Defining $H(\mathbb{X}_\alpha) = \bigoplus_j H_j(\mathbb{X}_\alpha)$ and taking direct sums of maps in the obvious way, will also give one large direct-sum persistence module $\{H(\mathbb{X}_\alpha)\}$.

2.2.1 Distance Functions

Here, the families of topological spaces will be produced by the sublevel sets of distance functions. Given a topological space \mathbb{X} embedded in some Euclidean space \mathbb{R}^N , we define $d_{\mathbb{X}}$ as the distance function which maps each point in the ambient space to the distance from its closest point in \mathbb{X} . More formally, for each $y \in \mathbb{R}^N$, $d_{\mathbb{X}}(y) = \inf_{x \in \mathbb{X}} \text{dist}(x, y)$. We let \mathbb{X}_α denote the sublevel set $d_{\mathbb{X}}^{-1}[0, \alpha]$; each sublevel set should be thought of a thickening of \mathbb{X} within the ambient space. Increasing the thickening parameter produces a growing family of sublevel sets, giving rise to the persistence module $\{H(\mathbb{X}_\alpha)\}_{\alpha \in \mathbb{R}_{\geq 0}}$; we denote the persistence diagram of this module by $\text{Dgm}(d_{\mathbb{X}})$ and use $\text{Dgm}_j(d_{\mathbb{X}})$ for the diagrams of the individual modules for each homological dimension j .

In Figure 3, we see an example of such an \mathbb{X} embedded in the plane, along with the persistence diagram $\text{Dgm}_1(d_{\mathbb{X}})$. We also have the persistence diagram $\text{Dgm}_1(d_U)$, where U is a dense point sample of \mathbb{X} . Note that the two diagrams are quite close in bottleneck distance. Indeed, the difference between the two diagrams will always be upper-bounded by the Hausdorff distance between the space and its sample; this follows from Theorem 2.1.

Persistence modules of relative homology groups arise from families of pairs of spaces, as the next example shows. Referring to the left part of Figure 4, we let \mathbb{X} be the space drawn in solid lines and B the closed ball whose boundary is drawn as a dotted circle. By restricting $d_{\mathbb{X}}$ to B and also to ∂B , we produce pairs of sublevel sets $(\mathbb{X}_\alpha \cap B, \mathbb{X}_\alpha \cap \partial B)$. Using the maps induced by the inclusions of pairs, we obtain the persistence module $\{H(\mathbb{X}_\alpha \cap B, \mathbb{X}_\alpha \cap \partial B)\}_{\alpha \in \mathbb{R}_{\geq 0}}$ of relative homology groups. The persistence diagram, for homological dimension 1, appears on the right half of Figure 4.

2.3 Stratified Spaces

We assume that we have a topological space \mathbb{X} embedded in some Euclidean space \mathbb{R}^N . A (purely) d -dimensional stratification of \mathbb{X} is a decreasing sequence of closed subspaces

$$\mathbb{X} = \mathbb{X}_d \supseteq \mathbb{X}_{d-1} \supseteq \dots \supseteq \mathbb{X}_0 \supseteq \mathbb{X}_{-1} = \emptyset,$$

such that for each i , the i -dimensional stratum $\mathbb{S}_i = \mathbb{X}_i - \mathbb{X}_{i-1}$ is a (possibly empty) i -manifold. The connected components of \mathbb{S}_i are called i -dimensional pieces. This is illustrated in Figure 5, where the space \mathbb{X} is a pinched torus with a spanning disc stretched across the hole.

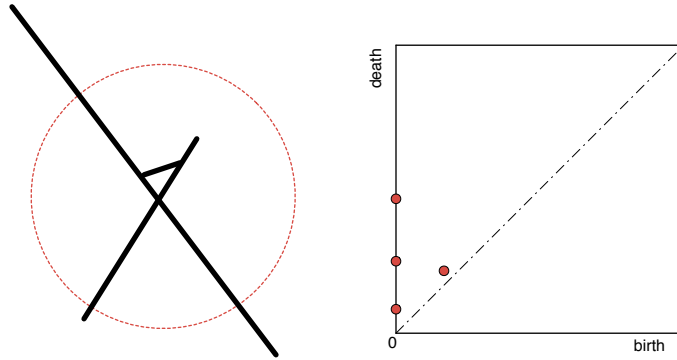


Figure 4: Left: The space \mathbb{X} is in solid line and the closed ball B has dotted boundary. Right: the persistence diagram for the module $\{H_1(\mathbb{X}_\alpha \cap B, \mathbb{X}_\alpha \cap \partial B)\}$.

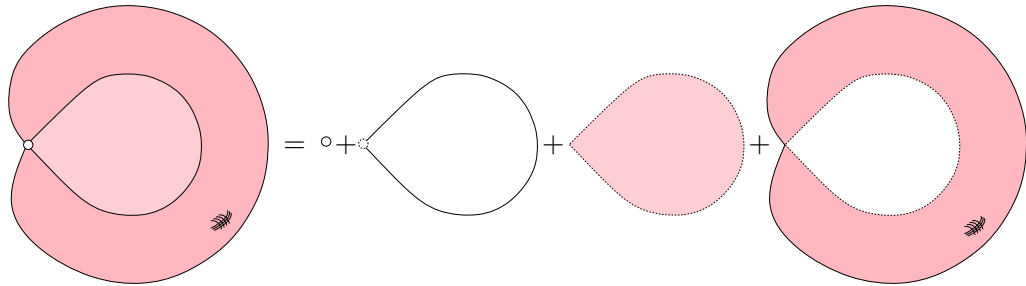


Figure 5: The coarsest stratification of a pinched torus with a spanning disc stretched across the hole.

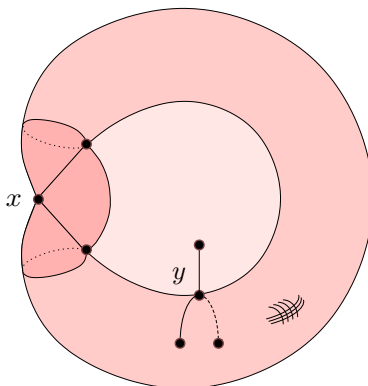


Figure 6: The cones $c(L_x)$ and $c(L_y)$, where x and y are respectively in the 0-stratum and the 1-stratum, are highlighted.

One usually also imposes a requirement to ensure that the various pieces fit together uniformly. There are a number of different ways this can be done (see [23] for an extensive survey). For example, one might assume that for each $x \in \mathbb{S}_i$, there exists a small enough neighborhood $N(x) \subseteq \mathbb{X}$ and a $(d-i-1)$ -dimensional stratified space L_x such that $N(x)$ is stratum-preserving homeomorphic to the product of an i -ball and the cone on L_x ; one can then show that the space L_x depends only on the particular piece containing x . This definition is illustrated in Figure 6.

Since the topology on \mathbb{X} is that inherited from the ambient space, this neighborhood $N(x)$ will take the form $\mathbb{X} \cap B_r(x)$, where $B_r(x)$ is a small enough ball around x in the ambient space.

We note that the above definition requires all strata to be contained within the closure of the top-dimensional stratum. It is also possible, of course, to have spaces where this is not the case: for example, a two-dimensional plane that has been punctured by a line. In this case, a slight adjustment to the above definitions can be made in order to impose similar notions of uniformity.

2.3.1 Local Homology and Homology Stratifications

Recall ([27]) that the local homology groups of a space \mathbb{X} at a point $x \in \mathbb{X}$ are the groups $H_i(\mathbb{X}, \mathbb{X} - x)$ in each homological dimension i . If \mathbb{X} happens to be a d -manifold, or if x is simply a point in the top-dimensional stratum of a d -dimensional stratification, then these groups are rank one in dimension d and trivial in all other dimensions. On the other hand, the local homology groups for lower-stratum points can be more interesting; for example if x is the crossing point in Figure 7, then $H_1(\mathbb{X}, \mathbb{X} - x)$ has rank three.

If x and y are close enough points in a particular piece of the same stratum, then there is a natural isomorphism between their local homology groups $H(\mathbb{X}, \mathbb{X} - x) \cong H(\mathbb{X}, \mathbb{X} - y)$, which can be understood in the following manner. Taking a small enough radii r and using excision, we see that the two local homology groups in question are in fact just $H(\mathbb{X} \cap B_r(x), \mathbb{X} \cap \partial B_r(x))$ and

$H(\mathbb{X} \cap B_r(y), \mathbb{X} \cap \partial B_r(y))$. Both of these groups will then map, via intersection of chains, isomorphically into the group $H(\mathbb{X} \cap B_r(x) \cap B_r(y), \partial(B_r(x) \cap B_r(y)))$, and the isomorphism above is then derived from these two maps. See the points in Figure 7 for an illustration of this idea.

In [31], the authors define the concept of a homology stratification of a space \mathbb{X} . Briefly, they require a decomposition of \mathbb{X} into pieces such that the locally homology groups are locally constant across each piece; more precisely, that the maps discussed above be isomorphisms for each pair of close enough points in each piece.

3 Topological Inference Theorem

From the discussion above, it is easy to see that any stratification of a topological space will also be a homology stratification. The converse is unfortunately false. However, we can build a useful analytical tool based on the contrapositive: given two points in a point cloud, we can hope to state, based on their local homology groups and the maps between them, that the two points should not be placed in the same piece of any stratification. To do this, we first adapt the definition of these local homology maps into a more multi-scale and robust framework. More specifically, we introduce a radius parameter r and a notion of local equivalence, \sim_r , which allows us to group the points of \mathbb{X} , as well as of the ambient space, into strata at this radius scale. We then give the main result of this section: topological conditions under which the point cloud U can be used to infer the strata at different radius scales.

3.1 Local Equivalence

We assume that we are given some topological space \mathbb{X} embedded in some Euclidean space in \mathbb{R}^N . For each radius $r \geq 0$, and for each pair of points $p, q \in \mathbb{R}^N$, we define the following homology map $\phi^{\mathbb{X}}(p, q, r)$:

$$H(\mathbb{X} \cap B_r(p), \mathbb{X} \cap \partial B_r(p)) \rightarrow H(\mathbb{X} \cap B_r(p) \cap B_r(q), \mathbb{X} \cap \partial(B_r(p) \cap B_r(q))). \quad (1)$$

Intuitively, this map can be understood as taking a chain, throwing away the parts that lie outside the smaller range, and then modding out the new boundary. Alternatively, one may think of it as being induced by a combination of inclusion and excision. A formal definition is given in Appendix A.

Using these maps, we impose an equivalence relation on \mathbb{R}^N .

Definition 3.1 (Local equivalence) *Two points x and y are said to have equivalent local structure at radius r , denoted $x \sim_r y$, iff there exists a chain of points $x = x_0, x_1, \dots, x_m = y$ from \mathbb{X} such that, for each $1 \leq i \leq m$, the maps $\phi^{\mathbb{X}}(x_{i-1}, x_i, r)$ and $\phi^{\mathbb{X}}(x_i, x_{i-1}, r)$ are both isomorphisms.*

In other words, x and y have the same local structure at this radius iff they can be connected by a chain of points which are pairwise close enough and whose local homology groups at radius r map into each other via intersection. Different choices of r will of course lead to different equivalence classes. For example, consider the space \mathbb{X} drawn in the plane as shown in the left half of Figure 7. At the radius drawn, point z is equivalent to the cross point and is not equivalent to

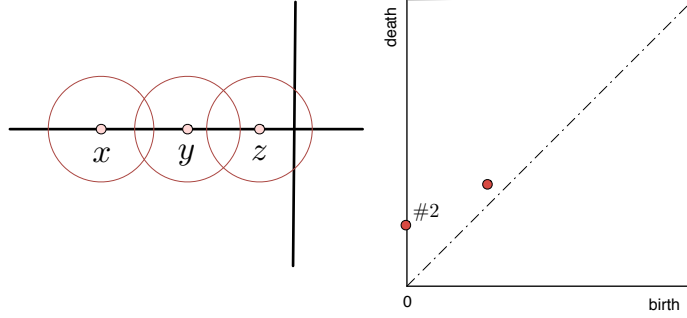


Figure 7: Left: $x \sim_r y, y \sim_r z$. Right: the 1-dim persistence diagram, for the kernel of the map going from the z ball into its intersection with the y ball. A number, i.e., #2, labeling a point in the persistence diagram indicates its multiplicity.

either the point x or y . Note that some points from the ambient space will now be considered equivalent to x and y , and some others will be equivalent to z .

On the other hand, a smaller choice of radius would result in all three of x, y , and z belonging to the same equivalence class.

3.1.1 (Co)Kernel Persistence

In order to relate the point cloud U to the equivalence relation \sim_r , we must first define a multi-scale version of the maps $\phi^{\mathbb{X}}(p, q, r)$; we do so by gradually thickening the space \mathbb{X} . Let $d_{\mathbb{X}} : \mathbb{R}^N \rightarrow \mathbb{R}$ denote the function which maps each point in the ambient space to the distance from its closest point on \mathbb{X} . For each $\alpha \geq 0$, we define $\mathbb{X}_{\alpha} = d_{\mathbb{X}}^{-1}[0, \alpha]$. For each p, q , and r , we will consider the intersection map $\phi_{\alpha}^{\mathbb{X}}(p, q, r)$, which is defined by substituting \mathbb{X}_{α} for \mathbb{X} in (1). Note of course that $\phi^{\mathbb{X}}(p, q, r) = \phi_0^{\mathbb{X}}(p, q, r)$.

For the moment, we fix a choice of p, q , and r , and we use the following shorthand:

$$\begin{aligned} B_p^{\mathbb{X}}(\alpha) &= \mathbb{X}_{\alpha} \cap B_r(p), \\ \partial B_p^{\mathbb{X}}(\alpha) &= \mathbb{X}_{\alpha} \cap \partial B_r(p), \\ B_{pq}^{\mathbb{X}}(\alpha) &= \mathbb{X}_{\alpha} \cap B_r(p) \cap B_r(q), \\ \partial B_{pq}^{\mathbb{X}}(\alpha) &= \mathbb{X}_{\alpha} \cap \partial(B_r(p) \cap B_r(q)). \end{aligned}$$

and we also often write $B_p^{\mathbb{X}} = B_p^{\mathbb{X}}(0)$ and $B_{pq}^{\mathbb{X}} = B_{pq}^{\mathbb{X}}(0)$. By replacing \mathbb{X} with U in this shorthand, we also write $B_p^U(\alpha) = U_{\alpha} \cap B_r(p)$, and so forth.

For any pair of non-negative real values $\alpha \leq \beta$ the inclusion $\mathbb{X}_{\alpha} \hookrightarrow \mathbb{X}_{\beta}$ gives rise to the following commutative diagram:

$$\begin{array}{ccc} \mathrm{H}(B_p^{\mathbb{X}}(\alpha), \partial B_p^{\mathbb{X}}(\alpha)) & \xrightarrow{\phi_{\alpha}^{\mathbb{X}}} & \mathrm{H}(B_{pq}^{\mathbb{X}}(\alpha), \partial B_{pq}^{\mathbb{X}}(\alpha)) \\ \downarrow & & \downarrow \\ \mathrm{H}(B_p^{\mathbb{X}}(\beta), \partial B_p^{\mathbb{X}}(\beta)) & \xrightarrow{\phi_{\beta}^{\mathbb{X}}} & \mathrm{H}(B_{pq}^{\mathbb{X}}(\beta), \partial B_{pq}^{\mathbb{X}}(\beta)) \end{array} \quad (2)$$

Hence there are maps $\ker \phi_\alpha^\mathbb{X} \rightarrow \ker \phi_\beta^\mathbb{X}$ and $\text{cok } \phi_\alpha^\mathbb{X} \rightarrow \text{cok } \phi_\beta^\mathbb{X}$. Allowing α to increase from 0 to ∞ gives rise to two persistence modules, $\{\ker \phi_\alpha^\mathbb{X}\}$ and $\{\text{cok } \phi_\alpha^\mathbb{X}\}$, with diagrams $\text{Dgm}(\ker \phi^\mathbb{X})$ and $\text{Dgm}(\text{cok } \phi^\mathbb{X})$. Recall that a homomorphism is an isomorphism iff its kernel and cokernel are both zero. In our context then, the map $\phi^\mathbb{X}$ is an isomorphism iff neither $\text{Dgm}(\ker \phi^\mathbb{X})$ nor $\text{Dgm}(\text{cok } \phi^\mathbb{X})$ contain any points on the y -axis above 0.

Example. As shown in the left part of Figure 7, x, y and z are points sampled from a cross embedded in the plane. Taking $p = z, q = y$, and r as drawn, we note that the right part of the figure displays $\text{Dgm}_1(\ker \phi^\mathbb{X})$, where $\phi^\mathbb{X} = \phi^\mathbb{X}(p, q, r)$; we now explain this diagram in some detail. The group $H_1(B_p^\mathbb{X}, \partial B_p^\mathbb{X})$ has rank three; as a possible basis we might take the three classes represented by the horizontal line across the ball, the vertical line across the ball, and the two short segments defining the northeast-facing right angle. Under the intersection map $\phi^\mathbb{X} = \phi_0^\mathbb{X}$, the first of these classes maps to the generator of $H_1(B_{pq}^\mathbb{X}, \partial B_{pq}^\mathbb{X})$, while the other two map to zero. Hence $\ker \phi_0^\mathbb{X}$ has rank two. Both classes in this kernel eventually die, one at the α value which fills in the northeast corner of the larger ball, and the other at the α value which fills in the entire right half; these two values are the same here due to symmetry in the picture. At this value, the map $\phi_\alpha^\mathbb{X}$ is an isomorphism and it remains so until the intersection of the two balls fills in completely. This gives birth to a new kernel class which subsequently dies when the larger ball finally fills in. The diagram $\text{Dgm}_1(\ker \phi^\mathbb{X})$ thus contains three points; the leftmost two show that the map $\phi^\mathbb{X}$ is not an isomorphism.

3.2 Inference Theorem

Given a point cloud U sampled from \mathbb{X} consider the following question: for a radius r , how can we infer whether or not any given pair of points in U has the same local structure at this radius? In this subsection, we prove a theorem which describes the circumstances under which we can make the above inference. Naturally, any inference will require that we use U to judge whether or not the maps $\phi^\mathbb{X}(p, q, r)$ are isomorphisms. The basic idea is that if U is a dense enough sample of \mathbb{X} , then the (co)kernel diagrams defined by U will be good enough approximations of the diagrams defined by \mathbb{X} .

3.2.1 (Co)Kernel Stability

Again we fix p, q , and r , and write $\phi^\mathbb{X} = \phi^\mathbb{X}(p, q, r)$. If the y -axes of $\text{Dgm}(\ker \phi^\mathbb{X})$ and $\text{Dgm}(\text{cok } \phi^\mathbb{X})$ are empty above 0, then $\phi^\mathbb{X}$ is an isomorphism. For each $\alpha \geq 0$, we let $U_\alpha = d_U^{-1}[0, \alpha]$. We consider $\phi_\alpha^U = \phi_\alpha^U(p, q, r)$, defined by replacing \mathbb{X} with U_α in (1). Running α from 0 to ∞ , we obtain two more persistence modules, $\{\ker \phi_\alpha^U\}$ and $\{\text{cok } \phi_\alpha^U\}$, with diagrams $\text{Dgm}(\ker \phi^U)$ and $\text{Dgm}(\text{cok } \phi^U)$.

If U is a dense enough sample of \mathbb{X} , then the (co)kernel diagrams defined by U will be good approximations of the diagrams defined by \mathbb{X} . More precisely, we have the following easy consequence of Theorem 2.1:

Theorem 3.1 ((Co)Kernel Diagram Stability) *The bottleneck distances between the (co)kernel diagrams of ϕ^U and $\phi^\mathbb{X}$ are upper-bounded by the Hausdorff dis-*

tance between U and \mathbb{X} :

$$\begin{aligned} d_B(\text{Dgm}(\ker \phi^U), \text{Dgm}(\ker \phi^{\mathbb{X}})) &\leq d_H(U, \mathbb{X}), \\ d_B(\text{Dgm}(\text{cok } \phi^U), \text{Dgm}(\text{cok } \phi^{\mathbb{X}})) &\leq d_H(U, \mathbb{X}). \end{aligned}$$

PROOF. We prove the first inequality; the proof of the second is identical. Put $\epsilon = d_H(U, \mathbb{X})$. Then, for each $\alpha \geq 0$, the inclusions $U_\alpha \hookrightarrow \mathbb{X}_{\alpha+\epsilon}$ and $\mathbb{X}_\alpha \hookrightarrow U_{\alpha+\epsilon}$ induce maps $\ker \phi_\alpha^U \rightarrow \ker \phi_{\alpha+\epsilon}^{\mathbb{X}}$ and $\ker \phi_\alpha^{\mathbb{X}} \rightarrow \ker \phi_{\alpha+\epsilon}^U$. These maps clearly commute with the module maps in the needed way, and hence we have the required ϵ -interleaving and can thus appeal to Theorem 2.1.

3.2.2 Main Inference Result

We now suppose that we have a point sample U of a space \mathbb{X} , where the Hausdorff distance between the two is no more than some ϵ ; in this case, we call U an ϵ -approximation of \mathbb{X} . Given two points $p, q \in U$ and a fixed radius r , we set $\phi^{\mathbb{X}} = \phi^{\mathbb{X}}(p, q, r)$, and we wish to determine whether or not $\phi^{\mathbb{X}}$ is an isomorphism. Since we only have access to the point sample U , we instead compute the diagrams $\text{Dgm}(\ker \phi^U)$ and $\text{Dgm}(\text{cok } \phi^U)$; we provide an algorithm for doing this in Section 5. The main Theorem of this section, Theorem 3.2, gives conditions under which these diagrams enable us to answer the isomorphism question for $\phi^{\mathbb{X}}$. To state the theorem we first need some more definitions.

Given any persistence diagram \mathcal{D} , which we recall is a multi-set of points in the extended plane, and two positive real numbers $a < b$, we let $\mathcal{D}(a, b)$ denote the intersection of \mathcal{D} with the portion of the extended plane which lies above $y = b$ and to the left of $x = a$; note that these points correspond to classes which are born no later than a and die no earlier than b .

For a fixed choice of p, q, r , we consider the following two persistence modules: $\{H(B_p^{\mathbb{X}}(\alpha), \partial B_p^{\mathbb{X}})\}$ and $\{H(B_{pq}^{\mathbb{X}}(\alpha), \partial B_{pq}^{\mathbb{X}})\}$. We let $\sigma(p, r)$ and $\sigma(p, q, r)$ denote their respective feature sizes and then set $\rho(p, q, r)$ to their minimum.

We now state the main theorem of this section, which states that we can use U to decide whether or not $\phi^{\mathbb{X}}(p, q, r)$ is an isomorphism as long as $\rho(p, q, r)$ is large enough relative to the sampling density.

Theorem 3.2 (Topological Inference Theorem) *Suppose that we have an ϵ -sample U from \mathbb{X} . Then for each pair of points $p, q \in \mathbb{R}^N$ such that $\rho = \rho(p, q, r) \geq 3\epsilon$, the map $\phi^{\mathbb{X}} = \phi^{\mathbb{X}}(p, q, r)$ is an isomorphism iff*

$$\text{Dgm}(\ker \phi^U)(\epsilon, 2\epsilon) \cup \text{Dgm}(\text{cok } \phi^U)(\epsilon, 2\epsilon) = \emptyset.$$

PROOF.

To simplify exposition, we will refer to points in $\text{Dgm}(\ker \phi^{\mathbb{X}}) \cup \text{Dgm}(\text{cok } \phi^{\mathbb{X}})$ and $\text{Dgm}(\ker \phi^U) \cup \text{Dgm}(\text{cok } \phi^U)$ as \mathbb{X} -points and U -points, respectively.

Whenever $0 < \alpha < \beta < 3\epsilon < \rho$, the two vertical maps in diagram (2) will by definition both be isomorphisms. Hence the maps $\ker \phi_\alpha^{\mathbb{X}} \rightarrow \ker \phi_\beta^{\mathbb{X}}$ and $\text{cok } \phi_\alpha^{\mathbb{X}} \rightarrow \text{cok } \phi_\beta^{\mathbb{X}}$ must also be isomorphisms. Hence, as α increases from 0 to ∞ , any element of the (co)kernel of $\phi^{\mathbb{X}}$ must live until at least 3ϵ , and any (co)kernel class which is born after 0 must in fact be born after 3ϵ . In other words, any \mathbb{X} -point

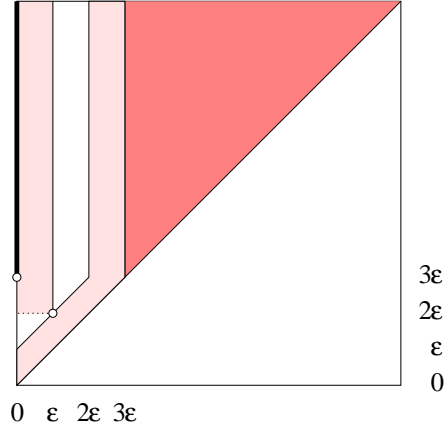


Figure 8: The point in the \mathbb{X} -diagrams lie either along the solid black line or in the darkly shaded region. Adding the lightly shaded regions, we get the region of possible points in the U -diagrams.

must lie either to the right of the line $x = 3\epsilon$ or along the y -axis and above the point $(0, 3\epsilon)$; see Figure 8. Recall that $\phi^{\mathbb{X}}$ is an isomorphism iff $\ker \phi^{\mathbb{X}} = 0 = \text{cok} \phi^{\mathbb{X}}$. Thus $\phi^{\mathbb{X}}$ is an isomorphism iff the black line in Figure 8 contains no \mathbb{X} -points.

On the other hand, Theorem 3.1 requires that every U -point must lie within ϵ of an \mathbb{X} -point. That is, all U -points are contained within the two lightly shaded regions drawn in Figure 8. Since the right such region is more than ϵ away from the thick black line, there will be a U -point in the left region iff there is an \mathbb{X} -point on the thick black line. But the U -points within the left region are exactly the members of $\text{Dgm}(\ker \phi^U)(\epsilon, 2\epsilon) \cup \text{Dgm}(\text{cok} \phi^U)(\epsilon, 2\epsilon)$. \square

Examples. Here we give two examples illustrating the topological inference theorem.

For the first example, suppose we have the space \mathbb{X} in the left half of Figure 9, and we take the labelled points p and q and the radius r as drawn; in this case, one can show that $\rho(p, q, r) = 8.5$, which here is the distance between the line segment and the boundary of the intersection of the two r -balls. First we compute the (co)kernel persistence diagrams for $\phi^{\mathbb{X}}$, showing the kernel diagram in the right half of Figure 9. Since the y -axis of this diagram is free of any points (and the same holds for the un-drawn cokernel diagram), p and q have the same local structure at this radius level.

On the other hand, suppose that we have an ϵ -sample U of \mathbb{X} , with $\epsilon = 2.8 < \rho/3$, as drawn in the left half of Figure 10. We can compute the analogous U -diagrams, with the kernel diagram drawn in the right half of the same figure. Noting that the two rectangles defined by $(\epsilon, 2\epsilon)$ in the two diagrams are indeed empty, and that the same holds for the cokernel diagrams, we can apply Theorem 3.2 to infer that the points have the same local structure at radius level r .

For a second example, suppose \mathbb{X} is the cross on the left half of Figure 11, with p, q, r as drawn. Then p and q are locally different at this radius level, as shown by the presence of two points on the y -axis of the kernel In the left half of Figure 12,

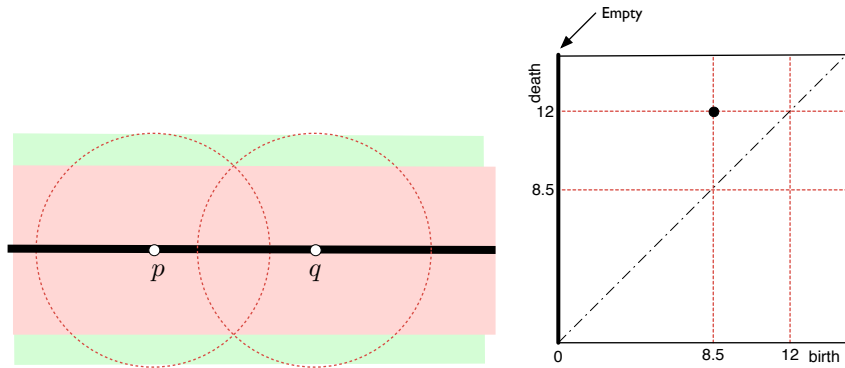


Figure 9: Kernel persistence diagram of two local equivalent points, given X .

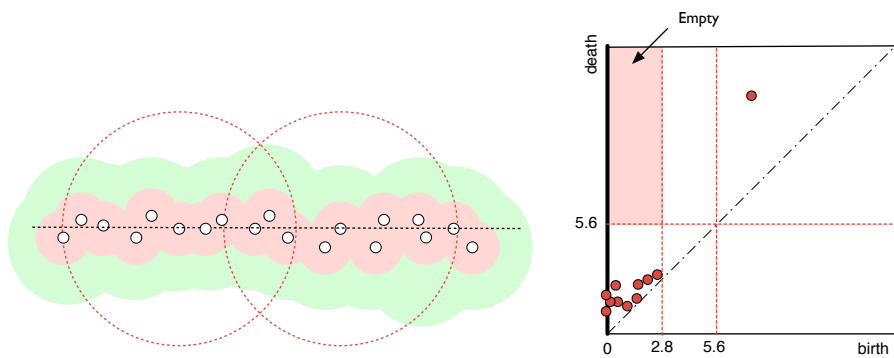


Figure 10: Kernel persistence diagram of two local equivalent points, given U .

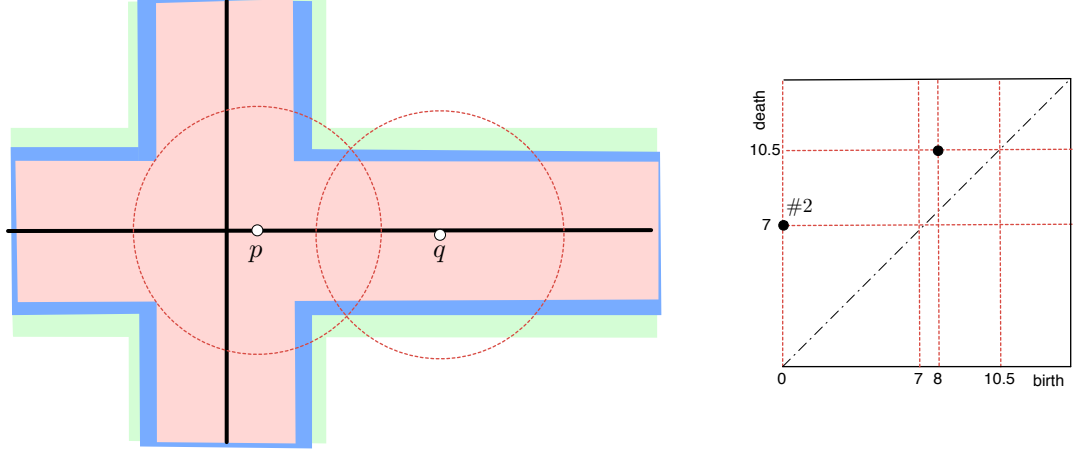


Figure 11: Kernel persistence diagram of two points that are not locally equivalent, given \mathbb{X} . A number, i.e., #2, labeling a point in the persistence diagram indicates its multiplicity.

we show an ϵ -sample U of \mathbb{X} , with $3\epsilon < \rho(p, q, r)$. Note that the kernel diagram for ϕ^U does indeed have two points in the relevant rectangle.

3.3 Geometric Intuition

Theorem 3.2, is stated in terms of a topological parameter, $\rho = \rho(p, q, r)$. We can relate ρ to more geometrically-flavored quantities. Specifically, we will show that ρ is lower bounded by a parameter derived from local variants of *reach*, as well as from parameters related to the gradient of $d_{\mathbb{X}}$. Unfortunately, this lower bound can be quite loose, zero in certain cases, limiting its practical utility. It does provide a geometric intuition to the topological constraints on the point cloud.

Recall that the medial axis \mathcal{M} of an embedded space \mathbb{X} is the subset of the ambient space consisting of all points which have at least two nearest neighbors on \mathbb{X} , and that the reach τ of \mathbb{X} is defined by $\tau = \inf_{x \in \mathbb{X}} \text{dist}(x, \mathcal{M})$. We fix notation for the following four intersections of \mathcal{M} with different subsets of the ambient space: $\mathcal{M}(p, r) = \mathcal{M} \cap B_r(p)$, $\mathcal{M}_0(p, r) = \mathcal{M} \cap \partial B_r(p)$, $\mathcal{M}(p, q, r) = \mathcal{M} \cap B_r(p) \cap B_r(q)$, and $\mathcal{M}_0(p, q, r) = \mathcal{M} \cap \partial(B_r(p) \cap B_r(q))$, and we define a variant of reach for each such space:

$$\begin{aligned} \tau(p, r) &= \inf_{x \in \mathbb{X}} \text{dist}(x, \mathcal{M}(p, r)) \\ \tau_0(p, r) &= \inf_{x \in \mathbb{X}} \text{dist}(x, \mathcal{M}_0(p, r)) \\ \tau(p, q, r) &= \inf_{x \in \mathbb{X}} \text{dist}(x, \mathcal{M}(p, q, r)) \\ \tau_0(p, q, r) &= \inf_{x \in \mathbb{X}} \text{dist}(x, \mathcal{M}_0(p, q, r)). \end{aligned}$$

Note that all four of these quantities are of course upper bounds on τ itself.

Letting $\nabla_{\mathbb{X}}$ be shorthand for the gradient of $d_{\mathbb{X}}$, we define the following subset

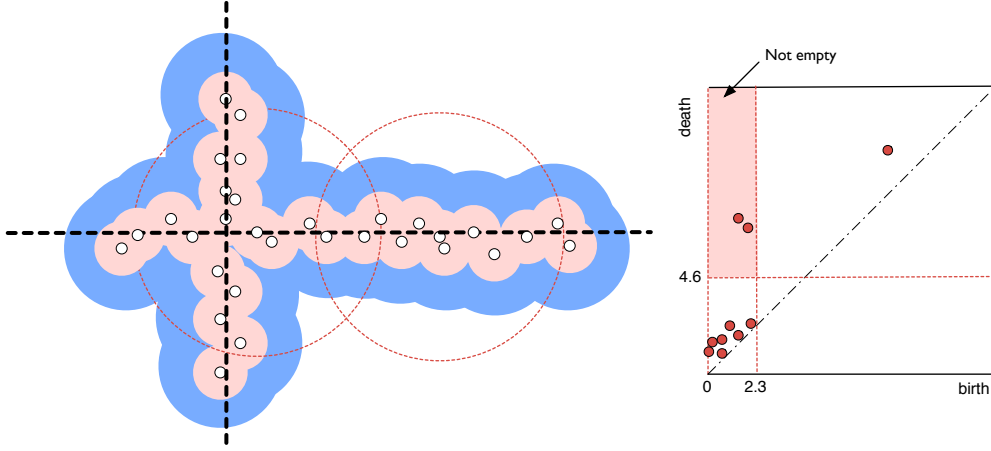


Figure 12: Kernel persistence diagram of two points that are not locally equivalent, given U .

of $\partial B_r(p)$:

$$G(p, r) = \{y \in \partial B_r(p) \mid \nabla_{\mathbb{X}}(y) \perp \partial B_r(p)\},$$

and then set $\eta(p, r) = \inf_{x \in \mathbb{X}} \text{dist}(x, G(p, r))$. We similarly define $G(p, q, r)$ and $\eta(p, q, r)$,

$$G(p, q, r) = \{y \in \partial(B_r(p) \cap B_r(q)) \mid \nabla_{\mathbb{X}}(y) \perp \partial(B_r(p) \cap B_r(q))\},$$

$$\eta(p, q, r) = \inf_{x \in \mathbb{X}} \text{dist}(x, G(p, q, r)).$$

Given the above quantities the following lower bound holds.

Theorem 3.3 (Geometric lower bound) *If we define*

$$\gamma = \gamma(p, q, r) = \min\{\tau(p, r), \tau(p, q, r), \eta(p, r), \eta(p, q, r)\},$$

then $\rho(p, q, r) \geq \gamma(p, q, r)$.

The proof appears in Appendix D.

4 Probabilistic Inference Theorem

The topological inference of Section 3 states conditions under which the point sample U can be used to infer stratification properties of the space \mathbb{X} . The basic condition is that the Hausdorff distance between the two must be small. In this section we describe two probabilistic models for generating the point sample U , and we provide an estimate of how large this point sample should be to infer stratification properties of the space \mathbb{X} with a quantified measure of confidence. More specifically, we provide a local estimate, based on $\rho(p, q, r)$ and $\rho(q, p, r)$, of how many sample points are needed to infer the local relationship at radius level r between two fixed points p and q ; this same theorem can be used to give a global estimate of the number of points needed for inference between any pair of points whose ρ -values are above some fixed low threshold.

4.1 Sampling Strategies

We assume \mathbb{X} to be compact. Since the stratified space \mathbb{X} can contain singularities and maximal strata of varying dimensions, some care is required in the sampling design. Consider for example a sheet of area one, punctured by a line of length one. In this case, sampling from a naively constructed uniform measure on this space would result in no points being drawn from the line. This same issue arose and was dealt with in [29], although in a slightly different approach than we will develop.

The first sampling strategy is to remove the problems of singularities and varying dimension by replacing \mathbb{X} by a slightly thickened version $\mathbb{X} \equiv \mathbb{X}_\delta$. We assume that \mathbb{X} is embedded in \mathbb{R}^k for some k . This new space is a smooth manifold with boundary and our point sample is a set of n points drawn identically and independently from the uniform measure $\mu(\mathbb{X})$ on \mathbb{X} , $U = \{x_1, \dots, x_n\} \stackrel{iid}{\sim} \mu(\mathbb{X})$. This model can be thought of as placing an appropriate measure on the highest dimensional strata to ensure that lower dimensional strata will be sampled from. We call this model M_1 .

The second sampling strategy is to deal with the problem of varying dimensions using a mixture model. In the example of the sheet and line, a uniform measure would be placed on the sheet, while another uniform measure would be placed on the line, and a mixture probability is placed on the two measures; for example, each measure could be drawn with probability $1/2$. We now formalize this approach. Consider each (non-empty) i -dimensional stratum $\mathbb{S}_i = \mathbb{X}_i - \mathbb{X}_{i-1}$ of \mathbb{X} . All strata that are included in the closure of some higher-dimensional strata, in other words all non-maximal strata, are not considered in the model. A uniform measure is assigned to the closure of each maximal stratum, $\mu_i(\mathbb{S}_i)$, this is possible since each such closure is compact. We assume a finite number of maximal strata K and assign to the closure of each such stratum a probability $p_i = 1/K$. This implies the following density

$$f(x) = \frac{1}{K} \sum_{j=1}^K \nu_j(X = x),$$

where ν_j is the density corresponding to measure μ_j . The point sample is generated from the following model: $U = \{x_1, \dots, x_n\} \stackrel{iid}{\sim} f(x)$. We call this model M_2 .

The first model replaces a stratified space with its thickened version, which enables us to place a uniform measure on the thickened space. Although this replacement makes it convenient for sampling, it does not sample directly from the actual space. The second model samples from the actual space, however the sample is not uniform on \mathbb{X} with respect to Lebesgue measure.

4.2 Lower bounds on the sample size of the point cloud

Our first main theorem is the probabilistic analogue of Theorem 3.2. An immediate consequence of this theorem is that, for two points $p, q \in U$, we can infer with probability at least $1 - \xi$ whether p and q are locally equivalent, $p \sim_r q$. The confidence level $1 - \xi$ will be a monotonic function of the size of the point sample.

The theorem involves a parameter $v(\rho)$, for each positive ρ , which is based on the volume of the intersection of ρ -balls with \mathbb{X} . First we note that each maximal stratum of \mathbb{X} comes with its own notion of volume: in the plane punctured by a

line example, we measure volume in the plane and in the line as area and length, respectively. The volume $\text{vol}(\mathbb{Y})$ of any subspace \mathbb{Y} of \mathbb{X} is the sum of the volumes of the intersections of \mathbb{Y} with each maximal stratum. For $\rho > 0$, we define

$$v(\rho) = \inf_{x \in \mathbb{X}} \frac{\text{vol}(B_{\rho/24}(x) \cap \mathbb{X})}{\text{vol}(\mathbb{X})} \quad (3)$$

We can then state:

Theorem 4.1 (Local Probabilistic Sampling Theorem) *Let $U = \{x_1, x_2, \dots, x_n\}$ be drawn from either model M_1 or M_2 . Fix a pair of points $p, q \in \mathbb{R}^N$ and a positive radius r , and put $\rho = \min\{\rho(p, q, r), \rho(q, p, r)\}$. If*

$$n \geq \frac{1}{v(\rho)} \left(\log \frac{1}{v(\rho)} + \log \frac{1}{\xi} \right),$$

then, with probability greater than $1 - \xi$ we can correctly infer whether or not $\phi^{\mathbb{X}}(p, q, r)$ and $\phi^{\mathbb{X}}(q, p, r)$ are both isomorphisms.

PROOF.

A finite collection $U = \{x_1, x_2, \dots, x_n\}$ of points in \mathbb{R}^N is ε -dense with respect to \mathbb{X} if $\mathbb{X} \subseteq U^\varepsilon$; equivalently, U is an ε -cover of \mathbb{X} . Let $C(\varepsilon)$ be the ε -covering number of \mathbb{X} , the minimum number of sets $B_\varepsilon \cap \mathbb{X}$ that cover \mathbb{X} . Let $P(\varepsilon)$ be the ε -packing number of \mathbb{X} , the maximum number of sets $B_\varepsilon \cap \mathbb{X}$ that can be packed into \mathbb{X} without overlap.

We consider a cover of \mathbb{X} with balls of radius $\rho/12$. If there is a sample point in each $\rho/12$ -ball, then U will be an ε -approximation of \mathbb{X} , with $\varepsilon \leq 4(\rho/12) = \rho/3$. This satisfies the condition of the topological inference theorem, and therefore we can infer the local structure between p and q .

The following two results from [28] will be useful in computing the number of sample points n needed to obtain, with confidence, such an ε -approximation.

Lemma 4.1 (Lemma 5.1 in [28]) *Let $\{A_1, A_2, \dots, A_l\}$ be a finite collection of measurable sets with probability measure μ on $\cup_{i=1}^l A_i$, such that for all A_i , $\mu(A_i) > \alpha$. Let $U = \{x_1, x_2, \dots, x_n\}$ be drawn iid according to μ . Then if $n \geq \frac{1}{\alpha} (\log l + \log \frac{1}{\xi})$, with probability $1 - \xi$, $\forall i, U \cap A_i \neq \emptyset$.*

Lemma 4.2 (Lemma 5.2 in [28]) *Let $C(\varepsilon)$ be the covering number of an ε -cover of \mathbb{X} and $P(\varepsilon)$ be the packing number of an ε -packing, then*

$$P(2\varepsilon) \leq C(2\varepsilon) \leq P(\varepsilon).$$

Again, we consider a cover of \mathbb{X} by balls of radius $\rho/12$. Let $\{y_i\}_{i=1}^l \in \mathbb{X}$ be the centers of the balls contained in a minimal sub-cover. Put $A_i = B_{\rho/12}(y_i) \cap \mathbb{X}$. Applying Lemma 4.1, we obtain the estimate

$$n \geq \frac{1}{\alpha} \left(\log l + \log \frac{1}{\xi} \right),$$

where l is the $\rho/12$ -covering number, and $\alpha = \min_i \frac{\text{vol}(B_{\rho/12} \cap \mathbb{X})}{\text{vol}(\mathbb{X})}$.

Applying Lemma 4.2,

$$l = C(\rho/12) \leq P(\rho/24) \leq \frac{\text{vol}(\mathbb{X})}{\text{vol}(B_{\rho/24} \cap \mathbb{X})} \leq \frac{1}{v(\rho)}.$$

On the other hand, $\frac{1}{\alpha} \leq \frac{1}{v(\rho)}$ by definition, and the result follows. \square

To extend the above theorem to a more global result, one can pick a positive ρ and radius r , and consider the set of all pairs of points (p, q) such that $\rho \leq \min\{\rho(p, q, r), \rho(q, p, r)\}$. Applying Theorem 4.1 uniformly to all pairs of points will give the minimum number of sample points needed to settle the isomorphism question for all of the intersection maps between all pairs.

Corollary 4.1 (Global Probabilistic Sampling Theorem) *Let $U = \{x_1, x_2, \dots, x_n\}$ be drawn from either model M_1 or M_2 . For all pairs of points $p, q \in U$ if*

$$n \geq \frac{1}{\alpha_{new}} \left(\log \frac{1}{\alpha_{new}} + \log \frac{1}{\xi} \right),$$

where $\alpha_{new} = \inf_{x \in \mathbb{X}} \frac{\text{vol}(B_{\rho_{min}/24}(x) \cap \mathbb{X})}{\text{vol}(\mathbb{X})}$, with probability greater than $1 - \xi$ we can infer whether $p \sim_r q$.

5 Algorithm

The theorems in the last section give conditions under which a point cloud U , sampled from a stratified space \mathbb{X} , can be used to infer the local equivalences between points on \mathbb{X} and its surrounding ambient space. We now switch gears slightly, and imagine clustering the U -points themselves into strata. The basic strategy is to build a graph on the point set, with edges coming from positive isomorphism judgements. The connected components of this graph will then be our proposed strata. We begin by describing this strategy in Section 5.1. Some of its potential limitations are discussed in Section 5.2, where we also describe a more robust variant based on graph diffusion.

A crucial subroutine in the clustering algorithm is the computation of the diagrams $\text{Dgm}(\ker \phi^U)$ and $\text{Dgm}(\text{cok } \phi^U)$, for $\phi^U = \phi^U(p, q, r)$ between all pairs $(p, q) \in U \times U$. The algorithm for this sub-routine is quite complicated, we describe it in detail in Section 5.3. The correctness proof is even more complicated; we give a proof sketch in Section 5.4, deferring all major details to Appendix C.

We would like to make clear that we consider the algorithm in this paper a first step and several issues both statistical and computational can be improved upon.

5.1 Clustering

We imagine that we are given the following input: a point cloud U sampled from some stratified space \mathbb{X} , and a fixed radius r . We make the assumption that $d_H(U, \mathbb{X}) \leq \epsilon \leq \frac{\rho_{min}}{3}$, where ρ_{min} is the minimum of $\rho(p, q, r)$ for all pairs $(p, q) \in U \times U$. Later we discuss the consequences when this assumption does not hold and a possible solution.

We build a graph where each node in the graph corresponds uniquely to a point from U . Two points $p, q \in U$ (where $\|p - q\| \leq 2r$) are connected by an edge iff both $\phi^{\mathbb{X}}(p, q, r)$ and $\phi^{\mathbb{X}}(q, p, r)$ are isomorphisms, equivalently iff $\text{Dgm}(\ker \phi^U)(\epsilon, 2\epsilon)$ and $\text{Dgm}(\text{cok } \phi^U)(\epsilon, 2\epsilon)$ are empty. The connected components of the resulting graph are our clusters. A more detailed statement of this procedure is giving in pseudo-code, see Algorithm 1. Note that the connectivity of the graph is encoded by a weight matrix, and our clustering strategy is based on a 0/1-weight assignment.

Algorithm 1 Strata-Inference(U, r, ϵ)

```
for all  $p, q \in U$  do  
  if  $\|p - q\| > 2r$  then  
     $W(p, q) = 0$   
  else  
    Compute  $\text{Dgm}(\ker \phi^U(p, q, r))$  and  $\text{Dgm}(\text{cok } \phi^U(p, q, r))$   
    Compute  $\text{Dgm}(\ker \phi^U(q, p, r))$  and  $\text{Dgm}(\text{cok } \phi^U(q, p, r))$   
    if  $\text{Dgm}(\ker \phi^U(p, q, r))(\epsilon, 2\epsilon) \cup \text{Dgm}(\text{cok } \phi^U(p, q, r))(\epsilon, 2\epsilon) \neq \emptyset$  then  
       $W(p, q) = 0$   
    else if  $\text{Dgm}(\ker \phi^U(q, p, r))(\epsilon, 2\epsilon) \cup \text{Dgm}(\text{cok } \phi^U(q, p, r))(\epsilon, 2\epsilon) \neq \emptyset$  then  
       $W(p, q) = 0$   
    else  
       $W(p, q) = 1$   
    end if  
  end if  
end for  
Compute connected components based on  $\mathbf{W}$ .
```

5.2 Robustness of clustering

Two types of errors in the clustering can occur: false positives where the algorithm connects points that should not be connected and false negatives where points that should be connected are not. The current algorithm we state is somewhat brittle with respect to both false positives as well as false negatives. We will suggest a very simple adaptation of our current algorithm that should be more stable with respect to both false positives and false negatives.

The false positives are driven by the condition in Theorem 3.2 that $\rho_{\min} < 3\epsilon$, so if the point cloud is not sampled fine enough we can get incorrect positive isomorphisms and therefore incorrect edges in the graph. If we use transitive closure to define the connected components this can be very damaging in practice since a false edge can collapse disjoint components into one large cluster.

The false negatives occur because our point sample U is not fine enough to capture chains of points that connect pairs in U through isomorphisms, there may be other points in \mathbb{X} which if we had sampled then the chain would be observed. The probability of these events in theory decays exponentially as the sample size increases and the confidence parameter ξ in Theorem 3.2 controls these errors.

We now state a simple adaptation of the algorithm that will make it more robust. It is natural to think of the 0/1-weight assignment on pairs of points $p, q \in U$ as an association matrix \mathbf{W} . A classic approach for robust partitioning is via spectral graph theory [26, 24, 9]. This approach is based an eigen-decomposition of the the graph Laplacian, $\mathbf{L} = \mathbf{D} - \mathbf{W}$ with the diagonal matrix $\mathbf{D}_{ii} = \sum_j \mathbf{W}_{ij}$. The smallest nontrivial eigenvalue λ_1 of \mathbf{W} is called the Fiedler constant and estimates of how well the vertex set can be partitioned [17]. The corresponding eigenvector v_1 is used to partition the vertex set. There are strong connections between spectral clustering and diffusions or random walks on graphs [9].

The problems of spectral clustering and lower dimensional embeddings have been examined from a manifold learning perspective [1, 2, 18]. The idea central

to these analyses is given a point sample from a manifold construct an appropriate graph Laplacian and use its eigenvectors to embed the point cloud in a lower dimensional space. A theoretical analysis of this idea involves proving convergence of the graph Laplacian to the Laplace-Beltrami operator on the manifold and the convergence of the eigenvectors of the graph Laplacian to the eigenvalues of the Laplace-Beltrami operator. A key quantity in this analysis is the Cheeger constant which is the first nontrivial eigenvalue of the Laplace-Beltrami operator [8]. An intriguing question is whether the association matrix we construct from the point cloud can be related to the Laplacian on high forms.

5.3 Diagram Computation

We now describe the computation of the diagrams $\text{Dgm}(\ker \phi^U)$ and $\text{Dgm}(\text{cok } \phi^U)$. To do this, we need for each $\alpha \geq 0$ a simplicial analogue of the map

$$\phi_\alpha^U : \mathbf{H}(B_p^U(\alpha), \partial B_p^U(\alpha)) \rightarrow \mathbf{H}(B_{pq}^U(\alpha), \partial B_{pq}^U(\alpha)).$$

To produce this, we first define, for each $\alpha \geq 0$, two pairs of simplicial complexes $L_0(\alpha) \subseteq L(\alpha)$ and $K_0(\alpha) \subseteq K(\alpha)$, and relative homology map

$$\psi_\alpha : \mathbf{H}(L(\alpha), L_0(\alpha)) \rightarrow \mathbf{H}(K(\alpha), K_0(\alpha))$$

between them. We will later show that

$$\text{Dgm}(\ker \phi^U) = \text{Dgm}(\ker \psi) \text{ and } \text{Dgm}(\text{cok } \phi^U) = \text{Dgm}(\text{cok } \psi).$$

To compute the diagrams involving ψ , we reduce various boundary matrices; since we follow very closely the (co)kernel persistence algorithm described in [12], we omit any further details here.

5.3.1 Complexes

To construct the simplicial complexes in our algorithm, we take the nerves of several collections of sets which are derived from a variety of Voronoi diagrams of different spaces. Here we briefly review these concepts.

Nerves. The *nerve* $N(\mathcal{C})$ of a finite collection of sets \mathcal{C} is defined to be the abstract simplicial complex with vertices corresponding to the sets in \mathcal{C} and with simplices corresponding to all non-empty intersections among these sets; that is, $N(\mathcal{C}) = \{S \subseteq \mathcal{C} \mid \bigcap S \neq \emptyset\}$. Every abstract simplicial complex can be geometrically realized, and therefore the concept of homotopy type makes sense. Under certain conditions, for example whenever the sets in \mathcal{C} are all closed and convex subsets of Euclidean space ([15], p.59), the nerve of \mathcal{C} has the same homotopy type, and thus the same homology groups, as the union of sets in \mathcal{C} .

Voronoi diagram. If U is a finite collection of points in \mathbb{R}^k and $u_i \in U$, then the *Voronoi cell* of u_i is defined to be:

$$V_i = V(u_i) = \{x \in \mathbb{R}^k \mid \|x - u_i\| \leq \|x - u_j\|, \forall u_j \in U\}.$$

The set of cells V_i cover the entire space and form the *Voronoi diagram* of \mathbb{R}^k , denoted as $\text{Voi}(U|\mathbb{R}^k)$. Each V_i restricted to a subset $X \subseteq \mathbb{R}^k$, $V_i \cap X$, cover the

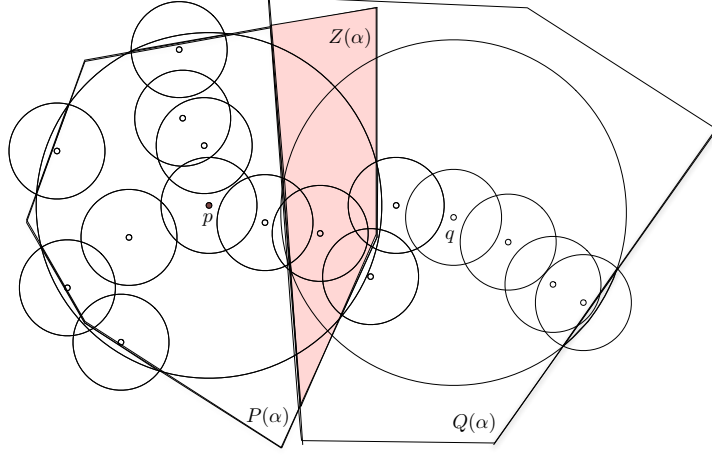


Figure 13: Illustration of intersection power cell $Z(\alpha)$, as the shaded region. The unshaded convex regions are $P(\alpha)$ and $Q(\alpha)$ respectively.

space X , and form the *restricted Voronoi diagram*, denoted as $\text{Voi}(U|X)$. For a simplex $\sigma \in U$, we set $V_\sigma = \cap_{u_i \in \sigma} V_i$.

The nerve of the restricted Voronoi diagram $\text{Voi}(U|X)$ is called the *restricted Delaunay triangulation*, denoted as $\text{Del}(U|X)$. It contains the set of simplices σ for which $V_\sigma \cap X \neq \emptyset$.

Power cells. An important task in our algorithm is the computation of the relative homology groups $H(B_p^U(\alpha), \partial B_p^U(\alpha))$ and $H(B_{pq}^U(\alpha), \partial B_{pq}^U(\alpha))$. Now to compute $H(U_\alpha)$, the absolute homology of the thickened point cloud, we would need only to compute the nerve of the collection of sets $V_i \cap U_\alpha$. This is because each such set is convex and their union obviously equals the space U_α . Such a direct construction will not work in our context, for the simple reason that the sets $V_i \cap \partial B_p^U(\alpha)$ and $V_i \cap \partial B_{pq}^U(\alpha)$ need not be convex.

To get around this problem, we first define $P(\alpha)$, the *power cell* with respect to $B_r(p)$, to be:

$$P(\alpha) = \{x \in \mathbb{R}^k \mid \|x - p\|^2 - r^2 \leq \|x - u\|^2 - \alpha^2, \forall u \in U\}, \quad (4)$$

and we set $P_0(\alpha) = B_r(p) - \text{int } P(\alpha)$. To define $Q(\alpha)$, the power cell with respect to $B_r(q)$, we replace p with q in (4). Finally, we set $Z(\alpha) = P(\alpha) \cap Q(\alpha)$, and $Z_0(\alpha) = (B_r(p) \cap B_r(q)) - \text{int } Z(\alpha)$. This is illustrated in Figure 13. We note that $P_0(\alpha)$ and $Z_0(\alpha)$ are both contained in U_α , as can be seen by playing around with the inequalities in their definitions.

Replacing $\partial B_p^U(\alpha)$ with $P_0(\alpha)$ and $\partial B_{pq}^U(\alpha)$ with $Z_0(\alpha)$ has no effect on the relative homology groups in question, as is implied by the following two lemmas. The first lemma was proven in [3]; a proof of the second appears in Appendix C.

Lemma 5.1 (Power Cell Lemma) *Assume $B_r(p) - P_0(\alpha) \neq \emptyset$. The identity on $B_p^U(\alpha)$ is a homotopy equivalence of $(B_p^U(\alpha), \partial B_p^U(\alpha))$ and $(B_p^U(\alpha), P_0(\alpha))$, as a map of pairs.*

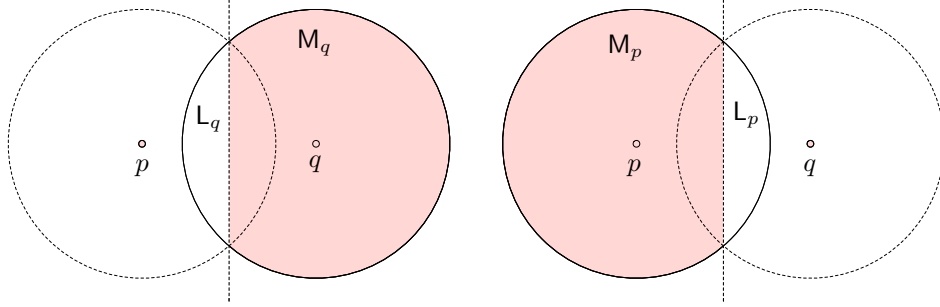


Figure 14: Illustration of the lune and the moon. The shaded regions are the respective moons. The white regions within solid circles are the respective lunes.

Lemma 5.2 (Intersection Power Cell Lemma) *Assume $B_r(p) \cap B_r(q) - Z_0(\alpha) \neq \emptyset$. The identity on $B_{pq}^U(\alpha)$ is a homotopy equivalence of $(B_{pq}^U(\alpha), \partial B_{pq}^U(\alpha))$ and $(B_{pq}^U(\alpha), Z_0(\alpha))$, as a map of pairs.*

Lune and moon. It can be shown ([3]) that the sets $V_i \cap P_0(\alpha)$ are convex. Unfortunately, it is still possible for some set $V_i \cap Z_0(\alpha)$ to be non-convex. To fix this, we must further divide the Voronoi cells in a manner we now describe.

We consider the hyperplane P of points in \mathbb{R}^k which are equidistant to p and q . This will divide \mathbb{R}^k into two half-spaces; let P_p and P_q denote the half-spaces containing p and q , respectively. We also define the p -lune, $L_p = P_q \cap B_r(p)$, and the p -moon, $M_p = P_p \cap B_r(p)$, as illustrated in Figure 14. The lune and the moon divide each Voronoi cell into two parts, $V_i^L = V_i \cap L_p$ and $V_i^M = V_i \cap M_p$. These sets are convex, assuming they are non-empty, since they are each the intersection of two convex sets. Furthermore, we have the following lemma whose simple but technical proof we omit:

Lemma 5.3 (Convexity Lemma) *The sets $V_i^L \cap Z_0(\alpha)$ and $V_i^M \cap Z_0(\alpha)$ are all convex, assuming they are non-empty.*

Of course the nonempty sets among $V_i^L \cap P_0(\alpha)$ and $V_i^M \cap P_0(\alpha)$ will also be convex.

To construct the simplicial complexes needed in our algorithm, we define \mathcal{A} to be the collection of the nonempty sets among $V_i^L \cap B_p^U(\alpha)$ and $V_i^M \cap B_p^U(\alpha)$, and we define \mathcal{A}_0 to be the collection of the nonempty sets among $V_i^L \cap P_0(\alpha)$ and $V_i^M \cap P_0(\alpha)$. Note that $\cup \mathcal{A} = B_p^U(\alpha)$ and $\cup \mathcal{A}_0 = P_0(\alpha)$. Taking the nerve of both collections, we define the simplicial complexes $L(\alpha) = N(\mathcal{A})$ and $L_0(\alpha) = N(\mathcal{A}_0)$.

Similarly, we define \mathcal{C} and \mathcal{C}_0 to be the collections of the nonempty sets among, respectively, $V_i^L \cap B_{pq}^U(\alpha)$ and $V_i^M \cap B_{pq}^U(\alpha)$, and $V_i^L \cap Z_0(\alpha)$ and $V_i^M \cap Z_0(\alpha)$. And we define $K(\alpha) = N(\mathcal{C})$ and $K_0(\alpha) = N(\mathcal{C}_0)$.

An example of the simplicial complexes constructed in \mathbb{R}^2 for a given U are illustrated in Figure 15. A direct approach to construct these simplicial complexes runs into difficulties as the corners of the convex sets created by the bisector can be shared by many sets, we defer the technicalities to Appendix B.

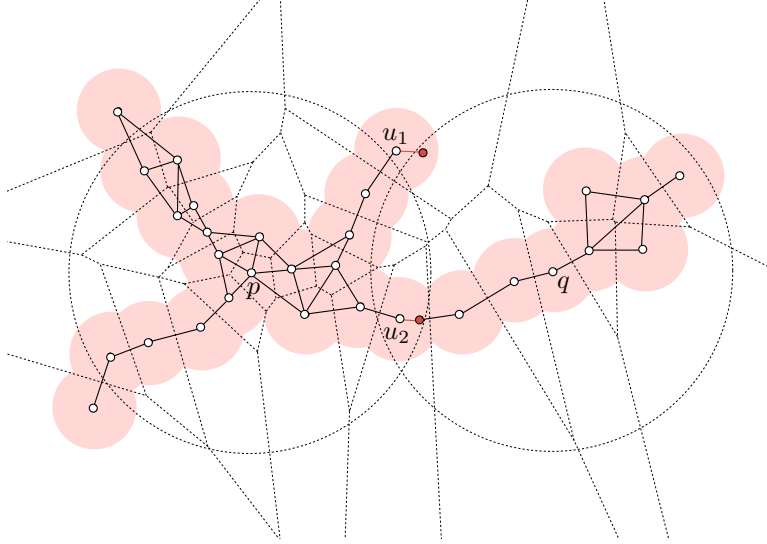


Figure 15: Illustration of the simplicial complexes constructed around two points p and q . The underlying Voronoi decomposition of the space is shown in thin dotted lines. u_1 and u_2 in U are the points whose restricted Voronoi regions intersect with the lune at non-convex regions.

5.3.2 Maps

We now construct simplicial analogues

$$\psi_\alpha : \mathbf{H}(L(\alpha), L_0(\alpha)) \rightarrow \mathbf{H}(K(\alpha), K_0(\alpha)).$$

of the maps ϕ_α^U .

The containments $L_0(\alpha) \subseteq L(\alpha)$ and $K_0(\alpha) \subseteq K(\alpha)$ are obvious. In order to define ψ_α , we first need the following technical lemma:

Lemma 5.4 (Containment Lemma) *Assume that a simplex σ is in $L_0(\alpha)$. If σ is also in $K(\alpha)$, then σ is in $K_0(\alpha)$, as well.*

PROOF. . By definition, $\sigma \in L_0(\alpha)$ iff there exists some point $x \in V^\sigma \cap P_0(\alpha)$. We must show that the set $V^\sigma \cap Z_0(\alpha)$ is non-empty. Note that $x \in P_0(\alpha)$ implies that $x \in B_r(p)$, while $x \notin \text{int } P(\alpha)$ implies that $x \notin \text{int } Z(\alpha)$. If $x \in B_r(q)$, then we are done, since $Z_0(\alpha) = B_r(p) \cap B_r(q) - \text{int } Z(\alpha)$.

Otherwise, choose some point $y \in V^\sigma \cap U_\alpha \cap B_r(p) \cap B_r(q)$, which is possible since $\sigma \in K(\alpha)$. Since both x and y belong to the same convex set $V^\sigma \cap U_\alpha \cap B_r(p)$, there exists a directed line segment γ from x to y within this set connecting them. We imagine moving along γ and first we suppose that γ intersects $B_r(q)$ before it intersects $\text{int } Q(\alpha)$. Let z be the first point of intersection. Then $z \in B_r(p) \cap B_r(q)$, $z \notin \text{int } Q(\alpha)$. Therefore $z \in V^\sigma \cap Z_0(\alpha)$. On the other hand, we may prove by contradiction that it is impossible for γ to intersect $Q(\alpha)$ before it intersects $B_r(q)$. Let z' be the first point of such an intersection. Since $z' \in Q(\alpha)$, by definition $\|z' - q\|^2 - r^2 \leq \|z' - u_i\|^2 - \alpha^2, \forall u_i \in U$. Since $z' \in U_\alpha, \forall u_i \in \sigma$, $\|z' - u_i\|^2 - \alpha^2 \leq 0$. Therefore $\|z' - q\|^2 - r^2 \leq \|z' - u_i\|^2 - \alpha^2 \leq 0, \forall u_i \in \sigma$. Since z' is outside $B_r(q)$, $\|z' - q\|^2 - r^2 > 0$. This is a contradiction. \square

To define ψ_α , we first construct a chain map $g = g_\alpha : \mathbb{C}(L(\alpha)) \rightarrow \mathbb{C}(K(\alpha))$ as follows. Given a simplex $\sigma \in L(\alpha)$, we define $g(\sigma) = \sigma$ if $\sigma \in K(\alpha)$, and $g(\sigma) = 0$ otherwise; we then extend g to a chain map by linearity. Using the Containment Lemma, we see that $g(\mathbb{C}(L_0(\alpha))) \subseteq \mathbb{C}(K_0(\alpha))$, and thus g descends to a relative chain map $f = f_\alpha : \mathbb{C}(L(\alpha), L_0(\alpha)) \rightarrow \mathbb{C}(K(\alpha), K_0(\alpha))$. Since f clearly commutes with all boundary operators, it induces a map on relative homology, and this is our $\psi = \psi_\alpha$.

5.4 Correctness

We show that our algorithm is correct by proving the following theorem. A sketch of the proof is given here, with the details deferred to Appendix C.

Theorem 5.1 (Correctness Theorem) *The persistence diagrams involving simplicial complexes are equal to the persistence diagrams involving the point cloud, that is, $\text{Dgm}(\ker \phi^U) = \text{Dgm}(\ker \psi)$ and $\text{Dgm}(\text{cok } \phi^U) = \text{Dgm}(\text{cok } \psi)$.*

Proof sketch. To prove Theorem 5.1, we will prove, for each $\alpha \leq \beta$, that the following diagram (as well as a similar diagram involving cokernels) commutes, with the vertical maps being isomorphisms.

$$\begin{array}{ccccccc} \dots & \rightarrow & \ker \phi_\alpha^U & \rightarrow & \ker \phi_\beta^U & \rightarrow & \dots \\ & & \uparrow \cong & & \uparrow \cong & & \\ \dots & \rightarrow & \ker \psi_\alpha & \rightarrow & \ker \psi_\beta & \rightarrow & \dots \end{array} \quad (5)$$

Applying Theorem 2.2 then finishes the claim. $\text{Dgm}(\ker \phi^U) = \text{Dgm}(\ker \psi)$ and $\text{Dgm}(\text{cok } \phi^U) = \text{Dgm}(\text{cok } \psi)$.

6 Simulations

We use a simulation on simple synthetic data with points sampled from grids to illustrate how the algorithm performs. In these simulations we assume we know ε , and we run our algorithm for $0 \leq \alpha \leq 2\varepsilon$. The data sets are shown in Figure 16.

We use the following results to demonstrate that the inference on local structure, at least for these very simple examples, is correct. As shown in Figure 17 top, if two points are locally equivalent, their corresponding ker/cok persistence diagrams contain the empty quadrant prescribed by our theorems, while in Figure 17 bottom, the diagrams associated to two non-equivalent points do not contain such empty quadrants. Similar results are shown for other data sets in Figure 18 and Figure 19.

7 Discussion

As the title of the paper suggests we have presented a first step towards learning stratified spaces. In the discussion we state some future problems and extensions of interest.

Algorithmic efficiency: The algorithm to compute the (co)kernel diagrams from the thickened point cloud is based on an adaption of Delaunay triangulation and

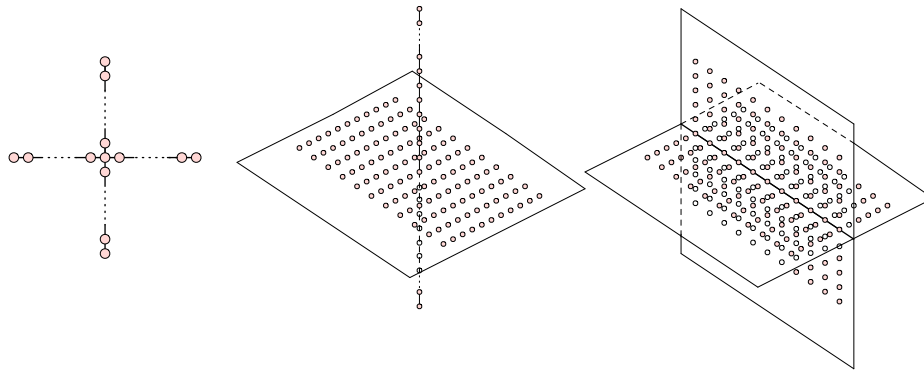


Figure 16: From left to right: points sampled from a cross; points sampled from a plane intersecting a line; points sampled from two intersecting planes. All points are located on the grid.

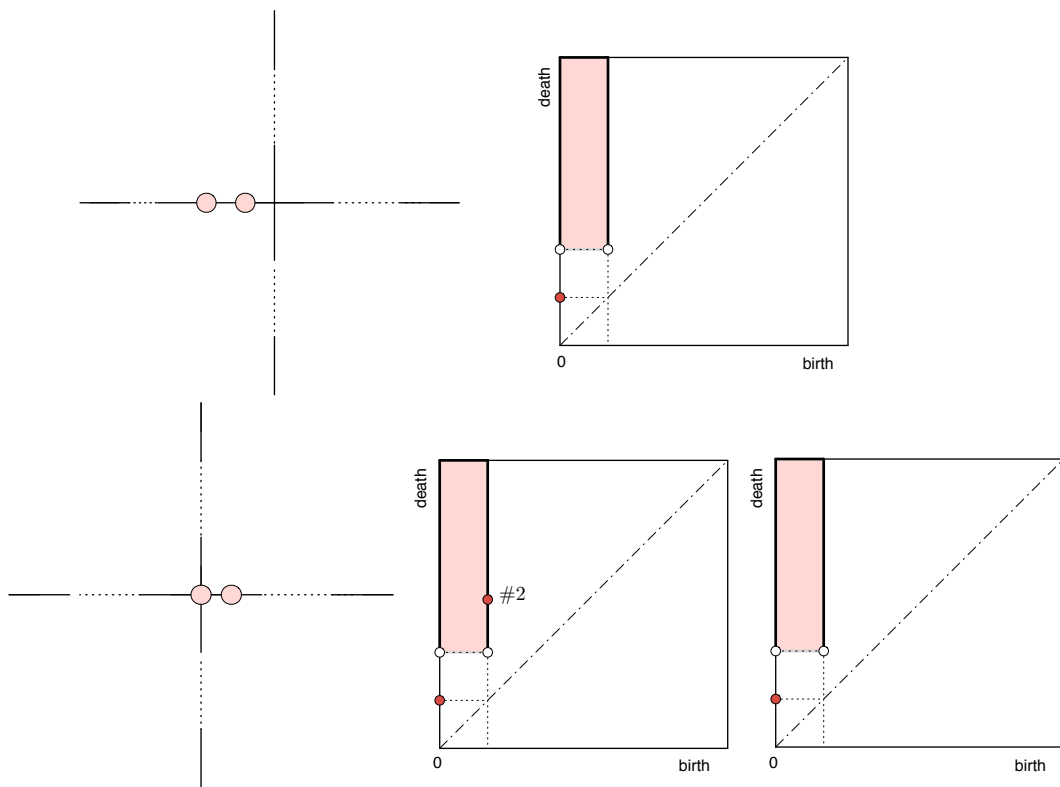


Figure 17: Top: both points are from 1-strata. Bottom: one point from 0-strata, one point from 1-strata. Left part shows the locations of the points. Right part shows the ker/cok persistence diagram of two points respectively, if the diagrams are the same, only one is shown. A number labeling a point in the persistence diagram indicates its multiplicity.

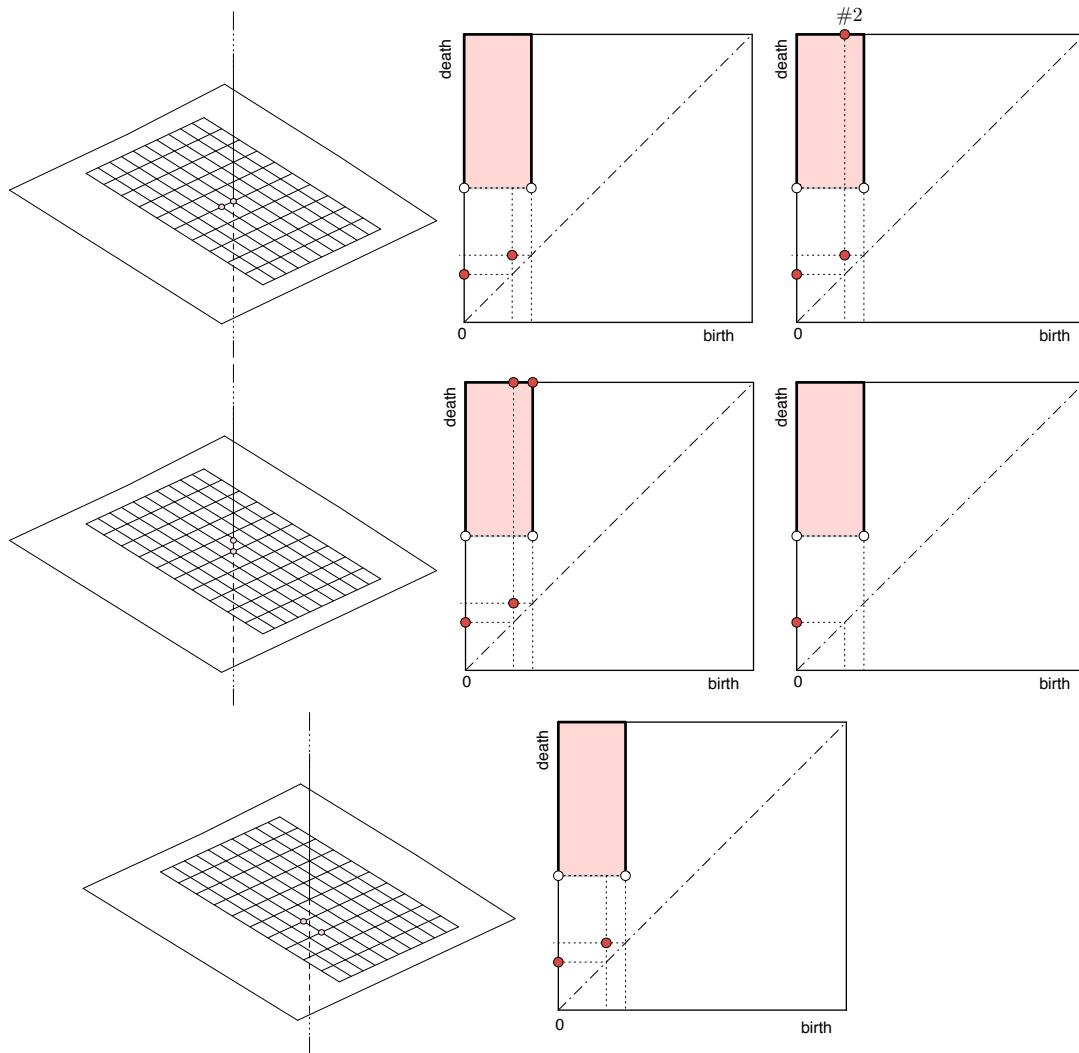


Figure 18: Top: one point from 0-strata, one point from 2-strata. Middle: one point from 0-strata, one from 1-strata. Bottom: both points are from 2-strata. A number labeling a point in the persistence diagram indicates its multiplicity.

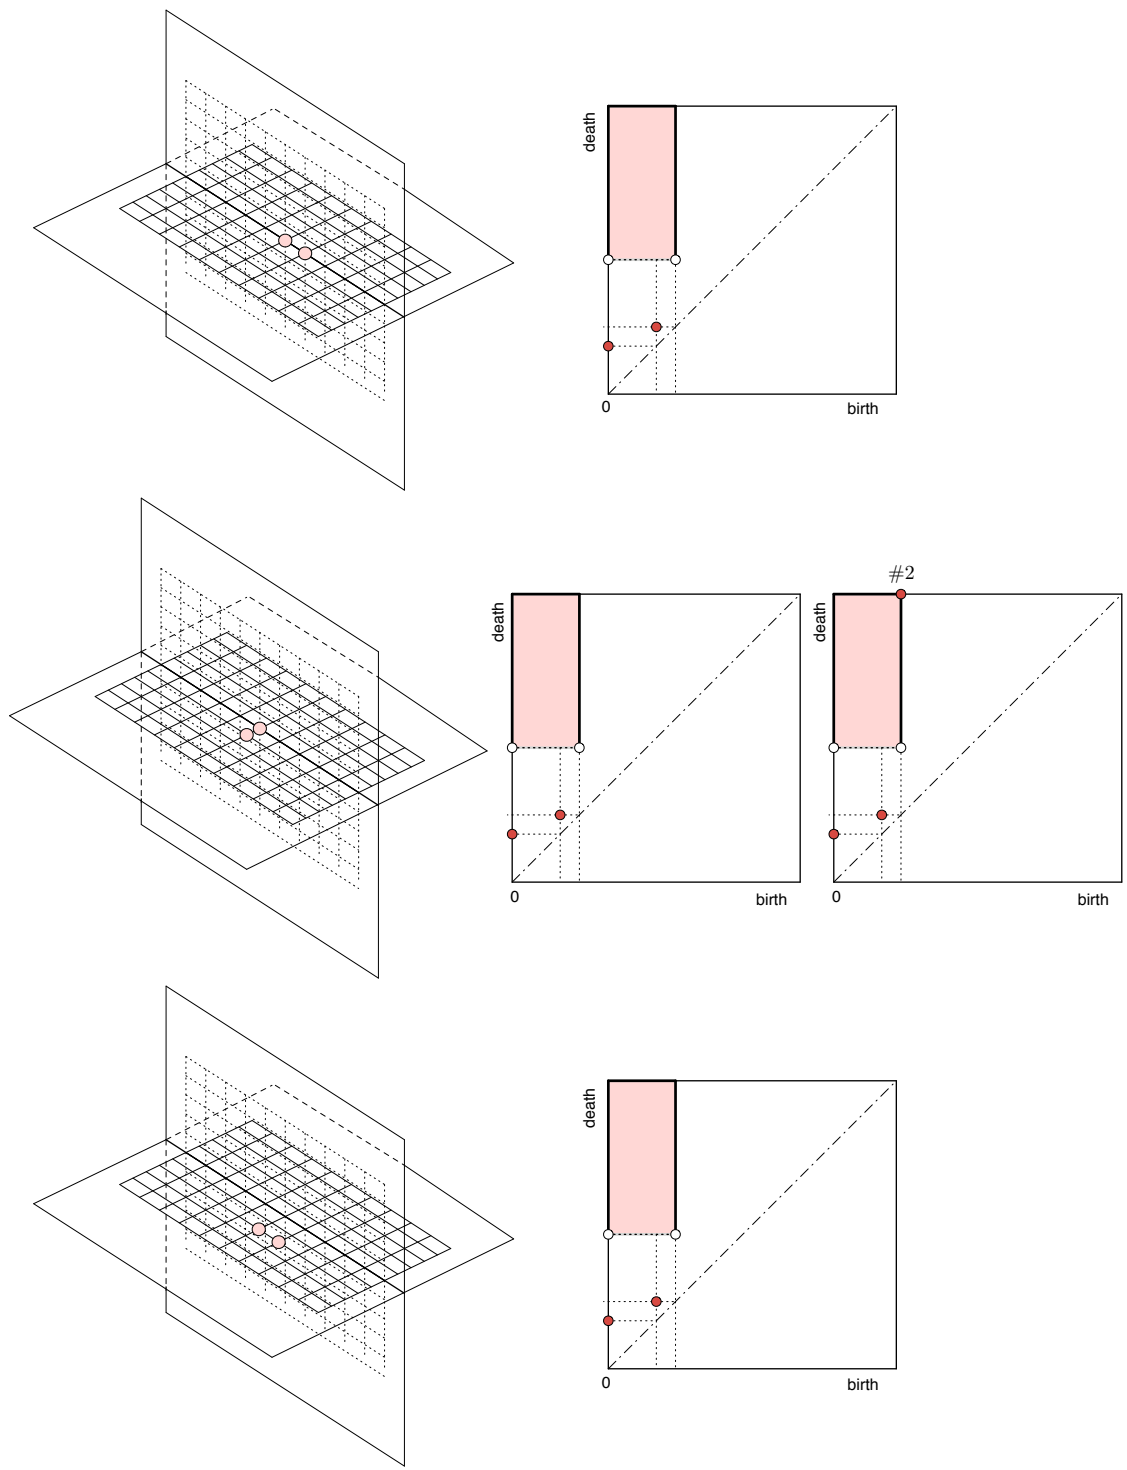


Figure 19: Top: both points from 1-strata. Middle: one point from 1-strata, one from 2-strata. Bottom: both points are from 2-strata. A number labeling a point in the persistence diagram indicates its multiplicity.

the power-cell construction. This algorithm should be quite slow when the dimensionality of the ambient space is high due to the runtime complexity of Delaunay triangulation. One idea to address this bottleneck is to use Rips or Witness complexes [13]; at the moment, we are not sure how to approach a proof of correctness, due to problems presented by the boundaries of the r -balls. Another approach is to use dimension reduction techniques such as principal components analysis (PCA) or random projection that approximately preserve distance [10] as a preprocessing step. Another idea that may work if the ambient dimension is not too high is using faster algorithms to construct Delaunay triangulations [4].

Weighting local equivalence: Currently we use a graph with 0/1 weights based on the local equivalence between two points. Extending this idea to assign fractional weights between points is appealing as it suggests a more continuous metric for local equivalence. This may also allow for greater robustness when using spectral methods to assign points to strata.

Curvature moderated tubes: Markus J. Pflaum [30] introduced a concept called *curvature moderation* that regulates the behavior of the tangent spaces of a stratum near the boundary. In other words, a stratum is curvature moderate if it curves near the boundary in a controlled way, this includes the higher derivatives of the curvature. This is yet another way to describe how strata and their tubular neighborhood are “glued together nicely”. It would be interesting to connect this concept to our idea of “local reach”.

Noisy data: Our sampling models draw points from the underlying topological space. A more general model would sample points that are concentrated on the topological space. A version of this type of sampling model is discussed in [28]. It would be of interest to study how well our approach is suited to such a model.

Adaptive sampling conditions: Throughout this paper we use ε -approximation to characterize the similarity of the point sample to the topological space. There are other approximation criteria that may be interesting to study and may provide better sampling estimates. One such criteria is the ε -sample [14] which is adaptive in that it is proportional to the local feature size. Another criteria of possible interest is the weak feature size [7].

Appendices

A Defining the Map ϕ

We give a more precise description of the map

$$\phi = \phi_\alpha^U : \mathbb{H}(B_p^U(\alpha), \partial B_p^U(\alpha)) \rightarrow \mathbb{H}(B_{pq}^U(\alpha), \partial B_{pq}^U(\alpha)).$$

The definition will be made on the chain level and will be given in terms of singular chains.

A.1 Background

We give here some necessary background as well as some material from algebraic topology and homological algebra which will be needed in Appendix C. Most of the descriptions are adapted from [22] and [27].

Chain homotopies. For our purposes, a *chain complex* \mathcal{C} is a sequence of $\mathbb{Z}/2\mathbb{Z}$ -vector spaces C_p , one for each integer p , connected by boundary homomorphisms $\partial_p^{\mathcal{C}} : C_p \rightarrow C_{p-1}$ such that $\partial_{p-1} \circ \partial_p = 0$ for all p . The p -th homology group of such a chain complex is defined by $H_p = \ker \partial_p / \text{im } \partial_{p+1}$.

A *chain map* $\eta : \mathcal{C} \rightarrow \mathcal{D}$ between two chain complexes is a sequence of homomorphisms $\eta_p : C_p \rightarrow D_p$ which commute with the boundary homomorphisms: $\partial_p^{\mathcal{D}} \circ \eta_p = \eta_{p-1} \circ \partial_p^{\mathcal{C}}$. Every chain map induces a map η_* between the homology groups of the two complexes.

A *chain homotopy* F between two chain maps $\eta, \eta' : \mathcal{C} \rightarrow \mathcal{D}$ is a sequence of homomorphisms $F_p : C_p \rightarrow D_{p+1}$ which satisfy the following formula for each p : $\eta - \eta' = \partial_{p+1}^{\mathcal{D}} \circ F_p - F_{p-1} \circ \partial_p^{\mathcal{C}}$. We say that η and η' are *chain homotopic* and note that they must then induce the same maps on homology: $\eta_* = \eta'_*$. Finally, η is called a *chain homotopy equivalence* if there exists a chain map $\rho : \mathcal{D} \rightarrow \mathcal{C}$ such that $\eta \circ \rho$ and $\rho \circ \eta$ are both chain homotopic to the identity. In this case η and ρ will both induce homology isomorphisms.

Singular homology. The *standard p -simplex* is the subset of \mathbb{R}^{p+1} given by

$$\Delta_p = \{(t_0, \dots, t_p) \in \mathbb{R}^{p+1} \mid \sum_{i=0}^p t_i = 1, \forall i, t_i \geq 0\}.$$

The $p + 1$ vertices of Δ_p are points $\{e_i\} \subset \mathbb{R}^{p+1}$ where

$$\begin{aligned} e_0 &= (1, 0, 0, \dots, 0), \\ e_1 &= (0, 1, 0, \dots, 0), \\ &\dots \\ e_p &= (0, 0, 0, \dots, 1). \end{aligned}$$

A *singular p -simplex* of a topological space X is a continuous map $\delta : \Delta_p \rightarrow X$. By taking formal sums of singular simplices (using binary coefficients for our purposes) one forms $C_p(X)$, the *singular chain group* of X in dimension p . Given points a_0, \dots, a_p in some Euclidean space, which need not be independent, there is a unique affine map l of Δ_p that maps the vertices e_i of Δ_p to a_i . This map defines the *linear singular simplex* determined by a_0, \dots, a_p , denoted as $l(a_0, \dots, a_p)$. One then defines a boundary homomorphism $\partial_p : C_p(X) \rightarrow C_{p-1}(X)$ by:

$$\partial_p(\delta) = \sum_{i=0}^p (\delta \circ l(\varepsilon_0, \dots, \hat{\varepsilon}_i, \dots, \varepsilon_p)),$$

and defines the singular homology groups $H_p(X)$ as above. A continuous map f from X to another topological space Y induces a chain map $f_{\#} : C_p(X) \rightarrow C_p(Y)$ given by the formula $f_{\#}(\delta) = f \circ \delta$, and thus also a homology map $f_* : H_p(X) \rightarrow H_p(Y)$.

The *minimal carrier* of a singular simplex δ is its image $\delta(\Delta_p)$, and the minimal carrier of a singular p -chain $\sum \delta_i$ is the union of the minimal carriers of the δ_i .

Isomorphism between simplicial and singular homology. The (simplicial) homology groups of a simplicial complex K and the singular homology groups of its realization $|K|$ are isomorphic. To show an explicit isomorphism ([27]), we first define a chain map

$$\eta : C(K) \rightarrow C(|K|)$$

as follows [27]: choose a partial ordering of the vertices of K that induces a linear ordering on the vertices of each simplex of K . Orient the simplices of K by using this ordering, and define

$$\eta([v_0, \dots, v_p]) = l(v_0, \dots, v_p),$$

where $v_0 < \dots < v_p$ in the given ordering. We refer to the linear singular simplex $l(v_0, \dots, v_p)$ as a *simplicial linear singular simplex* and it is important in the subsequent sections. The chain map η is in fact a chain equivalence as it has a *chain-homotopy inverse*, for which a specific formula can be found in [16]. Hence the induced homology map η_* provides an isomorphism of simplicial with singular homology.

A.2 Intersection Map Details

We now give the full and formal definition of the homology map $\phi = \phi_\alpha^U$, starting on the chain level. For compactness, we will use the following shorthand:

$$\begin{aligned} X &= B_p^U(\alpha) = U_\alpha \cap B_r(p), \\ B &= \partial B_p^U(\alpha) = U_\alpha \cap \partial B_r(p), \\ S &= B_{pq}^U(\alpha) = U_\alpha \cap B_r(p) \cap B_r(q), \\ A &= U_\alpha \cap B_r(p) - \text{int}(S), \\ U &= U_\alpha \cap B_r(p) - S. \end{aligned}$$

Note that $X - U = S = B_{pq}^U(\alpha)$ and $A - U = \partial S = \partial B_{pq}^U(\alpha)$. So to define ϕ we need only define a chain map $j : C(X, B) \rightarrow C(X - U, A - U)$ and then take ϕ as the map induced on homology. The map j is defined as the composition $j = k \circ i$. The chain map $i : C(X, B) \rightarrow C(X, A)$ is induced by inclusion on the second factor, while the chain map $k : C(X, A) \rightarrow C(X - U, A - U)$ is an excision, although this latter statement requires further elaboration.

Excisions. The inclusion map of pairs $(X - U, A - U) \rightarrow (X, A)$ is called an excision if it induces a homology isomorphism; in this case one says that U can be excised. We will make use of the following two results about excision (see, e.g., [20]):

Theorem A.1 (*Excision Theorem*) *If the closure of U is contained in the interior of A , then U can be excised.*

Theorem A.2 (*Excision Extension*) *Suppose $V \subset U \subset A$ and*

- (i) *V can be excised.*
- (ii) *$(X - U, A - U)$ is a deformation retract of $(X - V, A - V)$.*

Then U can be excised.

In our context, the closure of U need not be contained in the interior of A , and so we must define a suitable $V \subset U$. Although there are many ways to do this, one direct way is to choose some small enough positive δ .

$$\begin{aligned} I &= U_\alpha \cap \partial(B_r(p) \cap B_r(q)) \cap \text{cl}(U), \\ I_\delta &= \{x \in \text{cl}(U) \mid d_I(x) \leq \delta\}, \\ V &= U - I_\delta, \end{aligned}$$

where $d_I(x)$ is the distance from x to the set I . It is straightforward to verify that $V \subset U \subset A$ satisfies the hypotheses of Theorem A.2. In other words, the inclusion of pairs $(X - V, A - V) \rightarrow (X, A)$ is an excision; its induced chain map has a chain-homotopy inverse, which we denote as $s : C(X, A) \rightarrow C(X - V, A - V)$. Finally, we define $k = r_\# \circ s$, where $r_\#$ is the chain map induced by the retraction $r : (X - V, A - V) \rightarrow (X - U, A - U)$.

Subdivision. In order to fully carry out the analysis in Appendix B, we must first decompose the maps i and k through subdivision. Given a topological space X and a collection \mathcal{A} of subsets of X whose interiors form an open cover of X , a singular simplex of X is said to be \mathcal{A} -small if its image set is entirely contained in a single element of \mathcal{A} . For each dimension p , the chain group $C_p^{\mathcal{A}}(X)$ is the subgroup of $C_p(X)$ spanned by the \mathcal{A} -small singular p -simplices. These groups form a chain complex, with homology $H^{\mathcal{A}}(X)$. Of course, any singular simplex on X can be subdivided into a sum of \mathcal{A} -small simplices, so it is plausible, and in fact true ([22]), that the inclusion $C^{\mathcal{A}}(X) \rightarrow C(X)$ is a chain homotopy equivalence.

Returning to our context, we set $\mathcal{A} = \{X - V, A\}$ and denote by l the chain inclusion $C^{\mathcal{A}}(X, A) \rightarrow C(X, A)$. We also let $\rho : C(X, B) \rightarrow C^{\mathcal{A}}(X, B)$ be the chain homotopy inverse of the chain inclusion $C^{\mathcal{A}}(X, B) \rightarrow C(X, B)$, and let $t : C^{\mathcal{A}}(X, B) \rightarrow C^{\mathcal{A}}(X, A)$ be the chain map induced by inclusion on the second factor. Finally we note that $i = l \circ t \circ \rho$.

We also decompose k as $k = r_\# \circ \eta \circ \rho$, where η is the chain homotopy inverse of the chain map $C(X - V, A - V) \rightarrow C(X, A) \rightarrow C^{\mathcal{A}}(X, A)$.

Summary. To summarize, our map $\phi = j_*$, where j is the chain map defined by the following sequence of chain maps

$$j = k \circ i = (r_\# \circ \eta \circ \rho) \circ (l \circ t \circ \rho).$$

Following the same framework as above, we also define a chain map j' and its induced homology map $\phi' = j'_* : H(B_p^U(\alpha), P_0(\alpha)) \rightarrow H(B_{pq}^U(\alpha), Z_0(\alpha))$ simply by adopting the notation:

$$\begin{aligned} X &= B_p^U(\alpha) = U_\alpha \cap B_r(p), \\ B' &= P_0(\alpha), \\ S &= B_{pq}^U(\alpha) = U_\alpha \cap B_r(p) \cap B_r(q), \\ A' &= U_\alpha \cap B_r(p) - S + Z_0(\alpha), \\ U &= U_\alpha \cap B_r(p) - S, \\ I &= U_\alpha \cap \partial(B_r(p) \cap B_r(q)) \cap \text{cl}(U), \\ I_\delta &= \{x \in \text{cl}(U) \mid d_I(x) \leq \delta\}, \\ V &= U - I_\delta, \end{aligned}$$

defining our open cover to be $\mathcal{A}' = \{X - V, A'\}$, and otherwise proceeding exactly as before.

Similarly, we create a chain map f' which induces $\psi' = f'_* : H(|\text{Sd } L|, |\text{Sd } L_0|) \rightarrow H(|\text{Sd } K|, |\text{Sd } K_0|)$, using the notation

$$\begin{aligned} X' &= |\text{Sd } L|, \\ B' &= |\text{Sd } L_0|, \\ A' &= (|\text{Sd } L| - |\text{Sd } K|) \cup |\text{Sd } K_0|, \\ U' &= |\text{Sd } L| - \text{int } |\text{Sd } K|, \\ I &= |\text{Sd } K| \cap \text{cl } (U'), \\ I_\delta &= \{x \in \text{cl } (U') \mid d_I(x) \leq \delta\}, \\ V' &= U' - I_\delta, \end{aligned}$$

with $\mathcal{A}'' = \{A'', X'' - V''\}$.

B Algorithm Details

We describe the details in constructing the simplicial complexes described in our algorithm. The various simplicial complexes, L , L_0 , K and K_0 , are the nerves of collections of convex sets. Here we go through the construction of L ; construction of the others is similar.

Implicit Perturbations. A direct approach to constructing L , the nerve of the collection B , runs into difficulties as the corners of the convex sets created by the bisector P can be shared by many sets. To cope with this difficulty, we perturb these convex sets ever so slightly so that they meet in general position. Note that this is not done by perturbing the bisector; rather, it is done by decomposing the bisector into pieces.

We are interested in the restricted Voronoi diagram of the sublevel sets inside the ball $B_r(p)$, which we denote as $\mathcal{V} = \text{Voi}(U|U_\alpha \cap B_r(p))$. The *restricted Voronoi cell* of u_i is defined as $V(u_i|U_\alpha \cap B_r(p)) = V(u_i) \cap B_r(p)$.

Given \mathcal{V} , we create three sets of points. Let \mathcal{T}_{pq} be the set of points $u_i \in U$ whose restricted Voronoi cells have non-trivial intersection with the bisector: $V(u_i|U_\alpha \cap B_r(p)) \cap P \neq \emptyset$. We impose an ordering of points in \mathcal{T}_{pq} , w.l.o.g., let the ordered set be $\mathcal{T}_{pq} = \{x_1, x_2, \dots, x_m\}$. \mathcal{T}_p is the set of points $u_i \in U$ which are not in \mathcal{T}_{pq} and are closer to p than they are to q . Similarly, \mathcal{T}_q is the set of points $u_i \in U$, which are not in \mathcal{T}_{pq} and are closer to q .

By construction, the bisector P intersects the restricted Voronoi cells of the points in \mathcal{T}_{pq} . We denote these corresponding intersections as $\{P_1, P_2, \dots, P_m\}$. We perturb each P_i slightly such that no two pieces are colinear. Note that P_i is perpendicular to the direction $q - p$. One possible choice of perturbation would move each P_i within the restricted Voronoi cell along the direction $q - p$ for $i\varepsilon$ distance, where ε is sufficiently small. An example in \mathbb{R}^2 is shown in Figure 20.

Given such a perturbation, we let $\tilde{\mathcal{A}}$ be the resulting collection of perturbed convex sets, and we compute $\tilde{L} = \text{Nerve}(\tilde{\mathcal{A}})$ instead of $L = \text{Nerve}(\mathcal{A})$. By the properties of nerve construction, $\text{Nerve}(\tilde{\mathcal{A}}) \simeq \bigcup \tilde{\mathcal{A}}$, $\text{Nerve}(\mathcal{A}) \simeq \bigcup \mathcal{A}$. Since $\bigcup \tilde{\mathcal{A}} = \bigcup \mathcal{A}$, then we have $\tilde{L} \simeq L$. We now describe how we construct \tilde{L} .

Case Analysis. Let L' be the restricted Delaunay triangulation, $L' = \text{Del}(U|_{U_\alpha} \cap B_r(p))$. We read the simplicies from $\tilde{\mathcal{A}}$ without explicit perturbations. Specifically, we follow a set of rules, described below, to construct \tilde{L} from L' .

The bisector divides the restricted Voronoi cell of a point $x \in \mathcal{T}_{pq}$ into two convex sets. Let x^p represent the perturbed convex set closer to p in the nerve construction, and let x^q represent the other set. Let σ be a simplex in L' with k vertices, that is, $\sigma = [y_1, y_2, \dots, y_k]$. There are seven cases regarding the membership of the points $\{y_1, y_2, \dots, y_k\}$.

1. All $y_i \in \sigma$ belong to \mathcal{T}_p . We add the simplex $[y_1, y_2, \dots, y_k]$ to \tilde{L} .
2. All $y_i \in \sigma$ belong to \mathcal{T}_q . Same as case 1. We add the simplex $[y_1, y_2, \dots, y_k]$ to \tilde{L} .
3. All $y_i \in \sigma$ belong to \mathcal{T}_{pq} . Suppose $\{y_1, y_2, \dots, y_k\}$ are sorted according to the ordering in \mathcal{T}_{pq} . We add the following simplicies and their faces to \tilde{L} :

$$\begin{aligned} & [y_1^p, \dots, y_m^p, y_1^q] \\ & [y_2^p, \dots, y_m^p, y_1^q, y_2^q] \\ & [y_3^p, \dots, y_m^p, y_1^q, y_2^q, y_3^q] \\ & \dots \\ & [y_m^p, y_1^q, y_2^q, \dots, y_m^q] \end{aligned}$$

4. Some y_i are in \mathcal{T}_p , the rest are in \mathcal{T}_{pq} . Suppose $\{y_{i_1}, \dots, y_{i_n}\} \subseteq \mathcal{T}_p$ and $\{y_{j_1}, \dots, y_{j_l}\} \subseteq \mathcal{T}_{pq}$. We add $[y_{i_1}, \dots, y_{i_n}, y_{j_1}^p, \dots, y_{j_l}^p]$ to \tilde{L} .
5. Some y_i are in \mathcal{T}_q , the rest are in \mathcal{T}_{pq} . Similar to case 4, suppose $\{y_{i_1}, \dots, y_{i_n}\} \subseteq \mathcal{T}_q$ and $\{y_{j_1}, \dots, y_{j_l}\} \subseteq \mathcal{T}_{pq}$. We add $[y_{i_1}, \dots, y_{i_n}, y_{j_1}^q, \dots, y_{j_l}^q]$ to \tilde{L} .
6. Some y_i are in \mathcal{T}_p , the rest are in \mathcal{T}_q . We show that Case 6 is impossible. Choose $y_i \in \mathcal{T}_p$ and $y_j \in \mathcal{T}_q$ such that y_i and y_j are connected by an edge. Since y_i and y_j are on the opposite sides of P , this edge must intersect P at some point z . Then their corresponding restricted Voronoi cells, $V(y_i|_{U_\alpha} \cap B_r(p))$ and $V(y_j|_{U_\alpha} \cap B_r(p))$, must meet at a Voronoi face, which contains the point z . Suppose that the Voronoi face is in general position, that is, it is not parallel to P . Then P intersects the Voronoi face, by definition, y_i and y_j must belong to \mathcal{T}_{pq} . This is a contradiction.
7. Some y_i are in \mathcal{T}_p , some are in \mathcal{T}_q , and the rest are in \mathcal{T}_{pq} . We show that case 7 is impossible using the same proof in case 6.

A simple example is shown in Figure 20. Given $[y_1, y_2, y_3] \in L'$, simplex $[y_1, y_2^p, y_3^p]$ is added to \tilde{L} according to case 4. Given $[y_2, y_3] \in L$, simplicies $[y_2^p, y_3^p, y_2^q]$, $[y_3^p, y_2^q, y_3^q]$ and their faces are added to \tilde{L} according to case 3.

In summary, we construct \tilde{L} by iterating through all simplicies σ in L' , adding new simplicies to \tilde{L} constructed from σ following the above cases.

C Algorithmic Correctness

We prove the correctness of the algorithm described in Section 5.3 by proving Theorem 5.1. More precisely, we will prove that diagram 5 commutes, with the vertical arrows being isomorphisms, for some arbitrary but fixed choice of $\alpha < \beta$; we will omit the very similar argument about cokernels. The proof is unfortunately

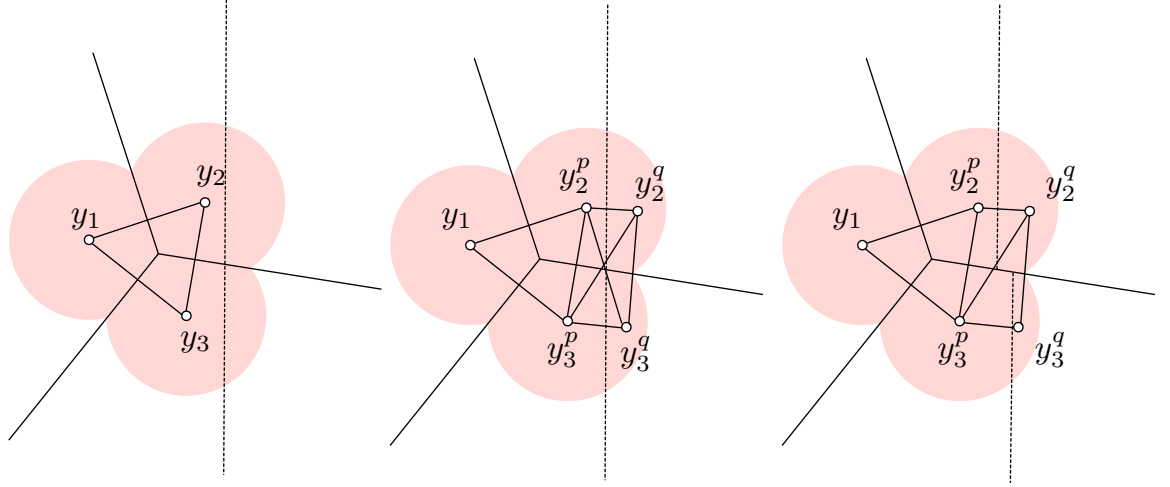


Figure 20: An example of the implicit perturbation. Dotted lines are the bisectors. A simplex $[y_1, y_2, y_3] \in L'$ is shown in the left. The simplices in L and in \tilde{L} are shown in the middle and right, respectively.

lengthy, and at times a bit technical, for in order to prove our statements about diagram 5, we must also prove similar statements about several other interlocking diagrams. For sanity and clarity of presentation, we first exhibit all the diagrams at once in the form of the following double-cube (Figure 21).

C.1 Bottom Face

The bottom face of the double-cube has been detached and drawn in diagram 6. The horizontal maps in the upper square are induced by inclusion of pairs, and so the upper square certainly commutes.

$$\begin{array}{ccc}
 \mathrm{H}(B_p^U(\alpha), \partial B_p^U(\alpha)) & \xrightarrow{j_\alpha^\beta} & \mathrm{H}(B_p^U(\beta), \partial B_p^U(\beta)) \\
 \downarrow i_\alpha & & \downarrow i_\beta \\
 \mathrm{H}(B_p^U(\alpha), P_0(\alpha)) & \xrightarrow{j_\alpha^\beta} & \mathrm{H}(B_p^U(\beta), P_0(\beta)) \\
 \uparrow h_\alpha & & \uparrow h_\beta \\
 \mathrm{H}(L(\alpha), L_0(\alpha)) & \xrightarrow{g_\alpha^\beta} & \mathrm{H}(L(\beta), L_0(\beta)). \tag{6}
 \end{array}$$

We have already shown that the two vertical maps in the upper square are isomorphisms; this was the content of the Power Cell Lemma in Section 5.3. To show that the vertical maps in the lower square are isomorphisms requires a bit more work. We make use of the following lemma, proven in [3].

Lemma C.1 (General Nerve Subdivision Lemma (GNSL)) *Let \mathcal{C} be the collection of maximal cells of a CW complex, each a convex set in \mathbb{R}^k . Define $f : |\mathrm{Sd} N| \rightarrow \cup \mathcal{C}$ by piecewise linear interpolation of its values at the vertices. If*

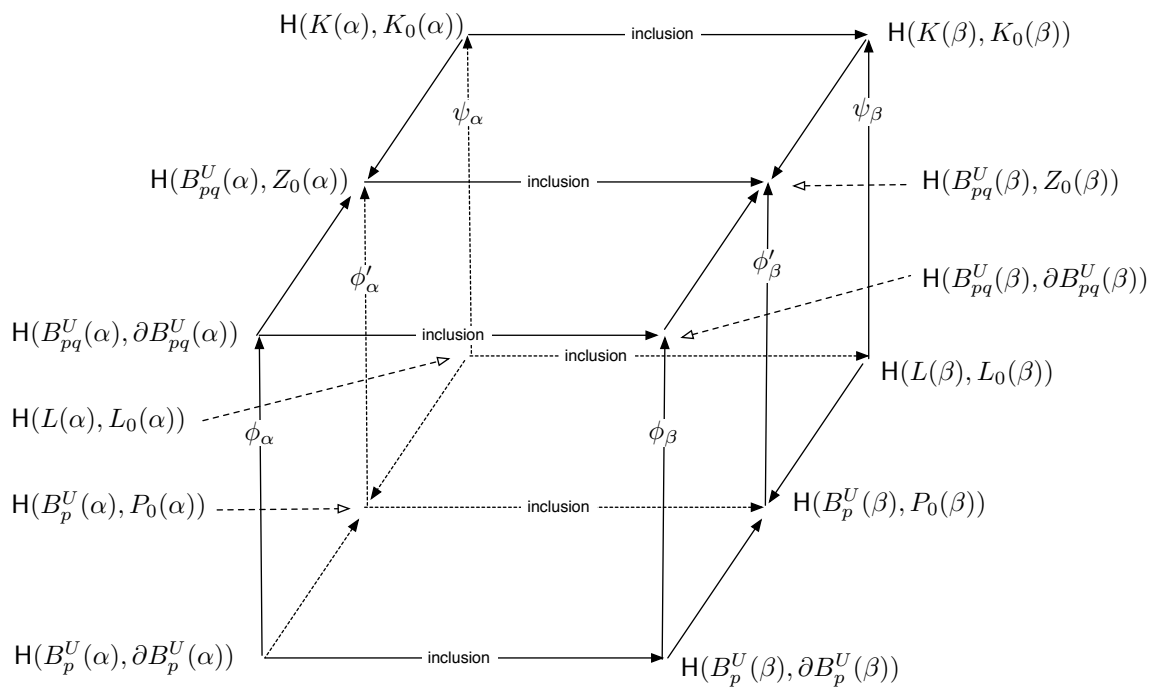


Figure 21: Two adjacent commuting cubes.

$f(\hat{\sigma})$ is contained in the intersection of the cells that correspond to the vertices of σ , for each simplex $\sigma \in N$, then f is a homotopy equivalence.

The vertical isomorphisms in the bottom square then follow from this next lemma, where we may of course replace α with β if we wish.

Lemma C.2 (Nerve Subdivision Lemma) Define $h = h_\alpha : |\text{Sd } L(\alpha)| \rightarrow B_p^U(\alpha)$ on the vertices $\hat{\sigma}$ of $\text{Sd } L(\alpha)$ by the formula

$$h_\alpha(\hat{\sigma}) = \arg \min_{x \in V^\sigma \cap U_\alpha \cap B_r(p)} d_U^2(x) - d_p^2(x),$$

and extend it by linear interpolation. Then h_α is a homotopy equivalence of pairs from $(|\text{Sd } L(\alpha)|, |\text{Sd } L_0(\alpha)|)$ to $(B_p^U(\alpha), P_0(\alpha))$.

PROOF. By construction, $h(\hat{\sigma})$ is contained in the intersection of the cells that correspond to the vertices of σ . By the GNSL then, h is a homotopy equivalence.

Now we need to prove that the restriction of h to $\text{Sd } L_0(\alpha)$ is also a homotopy equivalence. Let $\sigma \in L_0(\alpha)$ and put $h(\hat{\sigma}) = z$. For purposes of contradiction, suppose $z \notin P_0(\alpha)$. This means that $z \in \text{int } P(\alpha)$, by definition, and hence $d_U(z)^2 - d_p(z)^2 > \alpha^2 - r^2$.

Now choose some $z' \in V^\sigma \cap P_0$, which must exist since $\sigma \in L_0(\alpha)$. Then by definition we have $d_p(z')^2 - r^2 \geq d_U(z')^2 - \alpha^2$, or $d_U(z')^2 - d_p(z')^2 \leq \alpha^2 - r^2$. Combining the above inequalities, we have $d_U(z')^2 - d_p(z')^2 \leq \alpha^2 - r^2 < d_U(z)^2 - d_p(z)^2$, which contradicts the assumption that $h(\hat{\sigma}) = z$. We conclude that $z \in V_\sigma \cap P_0(\alpha)$. Applying the GNSL once more finishes the proof. \square

To show that the lower square commutes, we put $e = j_\alpha^\beta \circ h_\alpha$ and $e' = h_\beta \circ g_\alpha^\beta$, and we consider the map $H : |L(\alpha)| \times [0, 1] \rightarrow U_\alpha \cap B_r(p)$ defined by $H(x, t) = h_{\alpha_t} \circ g_{\alpha_t}^\alpha(x)$, where $\alpha_t = (1-t)\alpha + t\beta$. Since the maps g and j are inclusions and the maps h vary continuously with α , H is a homotopy between e and e' . This implies that the induced homomorphisms between the corresponding homology groups are the same, $e_* = e'_*$.

C.2 Top Face

We detach the top face of Figure 21, drawing it in diagram 7. As before, we prove that all vertical maps are isomorphisms. The commutativity of the two smaller squares follows from nearly identical arguments to the ones used for the bottom face.

$$\begin{array}{ccc} \text{H}(B_{pq}^U(\alpha), \partial B_{pq}^U(\alpha)) & \rightarrow & \text{H}(B_{pq}^U(\beta), \partial B_{pq}^U(\beta)) \\ \downarrow i'_\alpha & & \downarrow i'_\beta \\ \text{H}(B_{pq}^U(\alpha), Z_0(\alpha)) & \rightarrow & \text{H}(B_{pq}^U(\beta), Z_0(\beta)) \\ \uparrow h'_\alpha & & \uparrow h'_\beta \\ \text{H}(K(\alpha), K_0(\alpha)) & \rightarrow & \text{H}(K(\beta), K_0(\beta)) \end{array} \quad (7)$$

The Intersection Power Cell Lemma tells us that the vertical maps in the top square are isomorphisms. As promised, we give the proof of this lemma here, repeating the statement for completeness.

Lemma C.3 (Intersection Power Cell Lemma) Assume $B_r(p) \cap B_r(q) - Z_0(\alpha) \neq \emptyset$. The identity i' on $B_{pq}^U(\alpha)$ is a homotopy equivalence of pairs between $(B_{pq}^U(\alpha), \partial B_{pq}^U(\alpha))$ and $(B_{pq}^U(\alpha), Z_0(\alpha))$.

PROOF. It suffices to show that the restriction of the identity, $i' = i'_\alpha : \partial_{pq}^U(\alpha) \rightarrow Z_0(\alpha)$, is a homotopy equivalence. To do this, we first define a retraction $j : Z_0(\alpha) \rightarrow \partial_{pq}^U(\alpha)$ as follows. Fix a point $y \in \text{int } Z(\alpha)$, recalling that this set is nonempty by assumption. For each point $x \in Z_0(\alpha)$, we consider the unique ray starting at y and passing through x , and we let $x' = j(x)$ denote its intersection with $\partial(B_r(p) \cap B_r(q))$. Note that $x' \in Z(\alpha) \subseteq U(\alpha)$, and so j is certainly well-defined. That j is a retraction, meaning $j \circ i'$ is the identity on $\partial B_{pq}^U(\alpha)$, is obvious. On the other hand, the map

$$\lambda : Z_0(\alpha) \times [0, 1] \times Z_0(\alpha)$$

defined by $\lambda(x, t) = (1-t)x + tx'$ is a homotopy between $i' \circ j$ and the identity map on $Z_0(\alpha)$, and so the claim follows. \square

To prove that the vertical maps in the lower square are isomorphisms, we again make use of the GNSL.

Lemma C.4 (Intersection Nerve Subdivision Lemma (INSL)) *Define $h' = h'_\alpha : |\text{Sd } K(\alpha)| \rightarrow B_{pq}^U(\alpha)$ by setting*

$$h'_\alpha(\hat{\sigma}) = \arg \min_{x \in V_\sigma \cap U_\alpha \cap B_r(p) \cap B_r(q)} \min\{d_U^2(x) - d_p^2(x), d_U^2(x) - d_q^2(x)\},$$

where $\hat{\sigma}$ is the barycentre of $\sigma \in K\alpha$, and then extending by linear interpolation. Then h' is a homotopy equivalence of pairs between $(|\text{Sd } K(\alpha)|, |\text{Sd } K_0(\alpha)|)$ and $(B_{pq}^U(\alpha), Z_0(\alpha))$.

PROOF. The proof is quite similar to that of the NSL. By construction, $h'(\hat{\sigma})$ is contained in the intersection of the cells that correspond to the vertices of σ , and so we need only prove that the restriction of h' to the barycentric subdivision of $K_0(\alpha)$ is also a homotopy equivalence. Let $\sigma \in K_0(\alpha)$ and put $h'(\hat{\sigma}) = z$.

Suppose $z \notin Z_0(\alpha)$, and thus $z \in \text{int } Z(\alpha)$. By definition then, $d_p(z)^2 - r^2 < d_U(z)^2 - \alpha^2$ and $d_q(z)^2 - r^2 < d_U(z)^2 - \alpha^2$. In other words, $\min\{d_U^2(x) - d_p^2(x), d_U^2(x) - d_q^2(x)\} > \alpha^2 - r^2$.

Choose some point $z' \in V_\sigma \cap Z_0(\alpha)$. Then one of the following inequalities must hold: $d_p(z')^2 - r^2 \geq d_U(z')^2 - \alpha^2$, or $d_q(z')^2 - r^2 \geq d_U(z')^2 - \alpha^2$. That is, $\min\{d_U^2(z') - d_p^2(z'), d_U^2(z') - d_q^2(z')\} \leq \alpha^2 - r^2$.

Therefore, combining both inequalities, $\min\{d_U^2(z') - d_p^2(z'), d_U^2(z') - d_q^2(z')\} \leq \alpha^2 - r^2 < \min\{d_U^2(z) - d_p^2(z), d_U^2(z) - d_q^2(z)\}$, which contradicts the definition of z . \square

C.3 Left and Right Faces

We now come to the final and most complicated part of the correctness proof, involving the left face (diagram 8) of the double-cube; of course, everything we prove here will also hold for the right face. We have already established that all

vertical maps are isomorphisms, and now must show that both squares commute.

$$\begin{array}{ccc}
\mathbf{H}(B_p^U, \partial B_p^U) & \xrightarrow{\phi} & \mathbf{H}(B_{pq}^U, \partial B_{pq}^U) \\
\downarrow i_* & & \downarrow i'_* \\
\mathbf{H}(B_p^U, P_0) & \xrightarrow{\phi'} & \mathbf{H}(B_{pq}^U, Z_0) \\
\uparrow h_* & & \uparrow h'_* \\
\mathbf{H}(L, L_0) & \xrightarrow{\psi} & \mathbf{H}(K, K_0).
\end{array} \tag{8}$$

The top square will in fact commute even on the chain level. The bottom square is a little more complicated, and we start by addressing this first.

In diagram 9, this bottom square has been expanded into two smaller squares of chain groups connected by chain maps. We show that the lower of these squares commutes on the chain level, and that the two choices of path across the upper square are connected by a chain homotopy.

$$\begin{array}{ccc}
C(B_p^U, P_0) & \xrightarrow{j'} & C(B_{pq}^U, Z_0) \\
\uparrow h_{\#} & & \uparrow h'_{\#} \\
C(|\text{Sd } L|, |\text{Sd } L_0|) & \xrightarrow{f'} & C(|\text{Sd } K|, |\text{Sd } K_0|). \\
\uparrow \eta & & \uparrow \eta \\
C(\text{Sd } L, \text{Sd } L_0) & \xrightarrow{f} & C(\text{Sd } K, \text{Sd } K_0).
\end{array} \tag{9}$$

C.3.1 Map Details

First we need to discuss two of the horizontal chain maps from diagram 9 in more explicit detail.

Upper map. We analyze the effect of j' on an arbitrary linear singular simplex $\omega : \Delta_p \rightarrow B_p^U$, where $\omega = l(a_0, \dots, a_p)$ for some points a_i in Euclidean space. The analysis can be broken up into three main cases:

(A.1) $\omega(\Delta_p) \subseteq B_q^U$: Then j' maps ω through unchanged, meaning:

$$[\omega : \Delta_p \rightarrow B_p^U] \xrightarrow{j'} [\omega : \Delta_p \rightarrow B_{pq}^U].$$

From now on we simplify notation by omitting the domain and range of the singular simplex, writing instead: $\omega \xrightarrow{j'} \omega$.

(A.2) $\omega(\Delta_p) \cap B_q^U = \emptyset$: Then $\omega \xrightarrow{j'} 0$.

(A.3) $\omega(\Delta_p) \not\subseteq B_q^U$ and $\omega(\Delta_p) \cap B_q^U \neq \emptyset$: here we have two sub-cases:

(A.3.a) ω is \mathcal{A}' -small: This implies that $\omega(\Delta_p) \subseteq X - V$. Map j' can be interpreted as a retraction. That is, letting $T = \omega(\Delta_p)$, $S = \omega(\Delta_p) \cap B_q^U$ and $R = \omega(\Delta_p) \cap \partial B_q^U$, we define $r : T \rightarrow S$ by: for $x \in S$, $r(x) = x$; for $x \in T - S$, $r(x) = x'$, where $x' \in R$, as shown in the left of Figure 22. Then $\omega \xrightarrow{j'} \tau$, where $\tau : \Delta_p \rightarrow B_{pq}^U$ is defined by: for $\varepsilon \in \Delta_p$ where $\omega(\varepsilon) \in S$, $\tau(\varepsilon) = \omega(\varepsilon)$; otherwise for $\varepsilon \in \Delta_p$ where $\omega(\varepsilon) \notin S$, $\tau(\varepsilon) = r \circ \omega(\varepsilon)$.

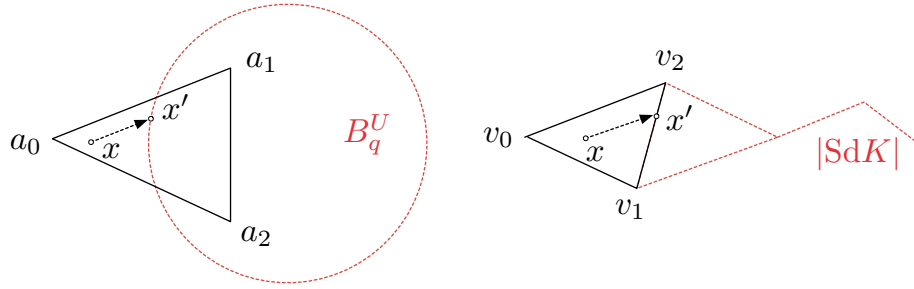


Figure 22: Left: map j' for a linear singular simplex $l(a_0, a_1, a_2)$. Right: map f' for a simplicial linear singular simplex $l(v_0, v_1, v_2)$.

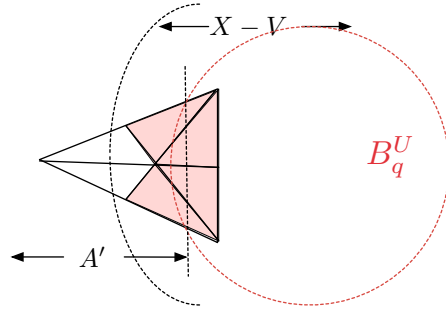


Figure 23: Map j' for a linear singular simplex that requires barycentric subdivision. In this illustrated example, all four shaded regions are the images of the four singular simplices in the first barycentric subdivision which are \mathcal{A}' -small and have their images in $X - V$. Their formal sum gives a singular chain in $X - V$. Their retraction result in a singular chain in B_{pq}^U .

(A.3.b) ω is not \mathcal{A}' -small: We barycentrically subdivide ω enough times m until $sd^m \omega$ is a \mathcal{A}' -small singular chain. Then each \mathcal{A}' -small singular simplex in $sd^m \omega$ that has its image in $X - V$ follows the pattern of (A.3.a), resulting in a singular simplex $\tau_i : \Delta_p \rightarrow B_{pq}^U$. In the end we have, $\omega \xrightarrow{j'} c_\tau$, where c_τ is the singular chain, $c_\tau = \sum \tau_i$. This is shown in Figure 23.

Middle map. We now describe the action of f' on an arbitrary simplicial linear singular simplex. Let $\delta : \Delta_p \rightarrow |\text{Sd } L|$ be such a simplex with $\delta = \eta(\sigma) = l(v_0, \dots, v_p)$, for some simplex $\sigma = [v_0, \dots, v_p] \in \text{Sd } K$. As above, we have three cases to consider::

- (B.1) $\delta(\Delta_p) \subseteq |\text{Sd } K|$: then $\delta \xrightarrow{f'} \delta$.
- (B.2) $\delta(\Delta_p) \cap |\text{Sd } K| = \emptyset$: $\delta \xrightarrow{f'} 0$.
- (B.3) $\delta(\Delta_p) \not\subseteq |\text{Sd } K|$ and $\delta(\Delta_p) \cap |\text{Sd } K| \neq \emptyset$: From Lemma C.5 below, we know that $(\delta(\Delta_p) \cap |\text{Sd } K|) \subseteq |\text{Sd } K_0|$, and so $\delta \xrightarrow{f'} 0$.

Lemma C.5 *Given a simplex $\sigma \in L$, if $\sigma \notin K$ and there exists $\tau < \sigma$ such that $\tau \in K$, then $\tau \in K_0$.*

PROOF. Suppose there exists $\omega < \tau$ such that $\omega \in K - K_0$. This implies that V^ω is completely contained in B_{pq}^U . Since V^σ is the intersection of V^ω with the partial Voronoi cells of vertices in σ that are not in ω , then V^σ should be completely contained in B_{pq}^U . This means that σ is in K , which leads to a contradiction. \square

C.3.2 Lower Square

As promised, we now show that the lower square in diagram 9 commutes. Choose an arbitrary $\sigma = [v_0, \dots, v_p] \in \text{Sd } L$, where each v_i is a barycenter of some simplex σ' in L ; as always, we assume that the vertices are ordered by increasing dimension of their defining simplices. We have two cases:

- (C.1) $\sigma \in \text{Sd } K$: by definition, $\eta(\sigma) = l(v_0, \dots, v_p)$ has its image in $|\text{Sd } K|$, and f is the identify map, that is, $\sigma \xrightarrow{f} \sigma \xrightarrow{\eta} \eta(\sigma)$. Meanwhile, by case (B.1), $\sigma \xrightarrow{\eta} \eta(\sigma) \xrightarrow{f'} \eta(\sigma)$. Therefore $(\eta \circ f)(\sigma) = f' \circ \eta(\sigma)$.
- (C.2) $\sigma \notin \text{Sd } K$: then $\sigma \xrightarrow{f} 0 \xrightarrow{\eta} 0$. On the other had, since $\sigma \notin \text{Sd } K$, we know that the image of $\delta = \eta(\sigma) = l(v_0, \dots, v_p)$ cannot be entirely contained within $|\text{Sd } K|$. There are then two sub-cases to consider:

(C.2.a) $\delta(\Delta_p) \cap |\text{Sd } K| = \emptyset$: this is case (B.2). We have $\sigma \xrightarrow{\eta} \delta \xrightarrow{f'} 0$.

(C.2.b) $\delta(\Delta_p) \cap |\text{Sd } K| \subseteq |\text{Sd } K_0|$: this is case (B.3). We have $\sigma \xrightarrow{\eta} \delta \xrightarrow{f'} 0$.

C.3.3 Upper Square

Finally, we show that the upper square in diagram 9 commutes up to chain homotopy; that is, we will construct a chain homotopy D between the two chain maps $e = j' \circ h_\#$ and $e' = h'_\# \circ f'$. This will of course imply that $e_* = e'_*$; in other words, that the induced homology diagram commutes. For clarity, we zoom in on diagram 9 and draw the relevant portion below as diagram 10.

$$\begin{array}{ccc} C(B_p^U, P_0) & \xrightarrow{j'} & C(B_{pq}^U, Z_0) \\ \uparrow h_\# & & \uparrow h'_\# \\ C(|\text{Sd } L|, |\text{Sd } L_0|) & \xrightarrow{f'} & C(|\text{Sd } K|, |\text{Sd } K_0|). \end{array} \quad (10)$$

For notational ease, we set $X = |\text{Sd } L|$ and $Y = B_{pq}^U$. To construct D , we will define for each p a chain map $F_p : C_p(X \times I) \rightarrow C_p(Y)$, and then we will set $D_p = F_{p+1} \circ G_p$, where $G_p : C_p(X \times I) \rightarrow C_{p+1}(X \times I)$ is given by Lemma C.6 below.

Construction of F. Let $\pi : X \times I \rightarrow X$ be projection on the first factor, and fix an arbitrary simplicial linear singular simplex $\kappa : \Delta_p \rightarrow X \times I$. Then $\pi_\#(\kappa) = \delta = l(\hat{\sigma}_0, \dots, \hat{\sigma}_p)$, for some simplex $\sigma = [\hat{\sigma}_0, \dots, \hat{\sigma}_p]$ in $\text{Sd } L$. We define F in stages, based on properties of δ , as follows.

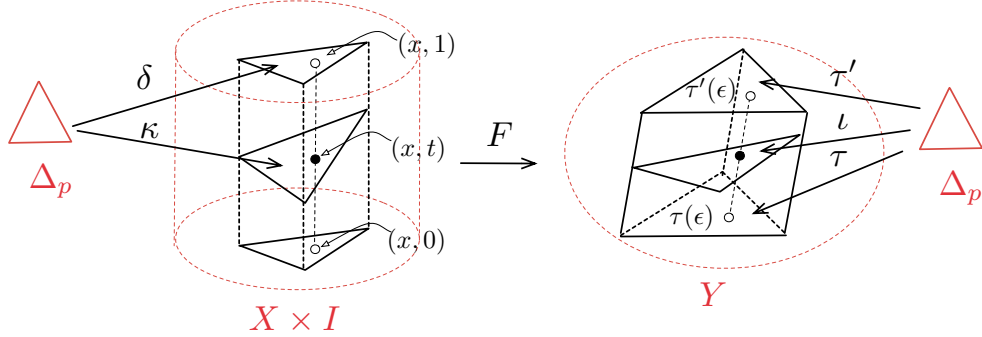


Figure 24: Case (D.1.a): illustration of F .

(D.1) $\delta(\Delta_p) \subseteq |\text{Sd } K|$: following the e' -path and case (B.1), we have $\delta \xrightarrow{f'} \omega$
 $\delta \xrightarrow{h_{\#}'} \tau'$. On the other hand, following the e -path results in $\delta \xrightarrow{h_{\#}} \omega$. We
now have three sub-cases, based on varying properties of ω :

(D.1.a) $\omega(\Delta_p) \subseteq B_q^U$: following the e -path and case (A.1). we have, $\delta \xrightarrow{h_{\#}} \omega$
 $\omega \xrightarrow{j'} \tau$, where $\tau = \omega$ except for differing range. We then can define
 $F(\kappa) = \iota$, where $\iota : \Delta_p \rightarrow Y$ is given by: for every $\epsilon \in \Delta_p$, where
 $\kappa(\epsilon) = (x, t) \in X \times I$, $\iota(\epsilon) = (1-t)\tau(\epsilon) + t\tau'(\epsilon)$. This formula is
illustrated in Figure 24.

(D.1.b) $\omega(\Delta_p) \cap B_q^U = \emptyset$: This is case (A.2). We branch further as follows:

- (i) $\delta(\Delta_p) \subseteq |\text{Sd } K_0|$: Following the e' -path, $\delta \xrightarrow{f'} 0 \xrightarrow{h_{\#}'} 0$. Similarly,
following the e -path, $\delta \xrightarrow{h_{\#}} \omega \xrightarrow{j'} 0$. We define $F(\kappa) = 0$.
- (ii) $\delta(\Delta_p) \not\subseteq |\text{Sd } K_0|$: this is not possible. Suppose it were. This
implies that there exists at least one vertex $\hat{\sigma}_i$ of σ such that $V^{\sigma_i} \cap B_{pq}^U \neq \emptyset$
and $V^{\sigma_i} \cap Z_0 = \emptyset$. This means that V^{σ_i} is completely contained in $B_r(q)$.
Therefore $h(\hat{\sigma}_i)$ is contained in $B_r(q)$, which contradicts our assumption.

(D.1.c) $\omega(\Delta_p) \not\subseteq B_q^U$ and $\omega(\Delta_p) \cap B_q^U \neq \emptyset$: we must consider two further
sub-cases.

(i) ω is \mathcal{A}' -small: this is case (A.3.a), and we define $F(\kappa)$ similarly
to case (D.1.a).

(ii) ω is not \mathcal{A}' -small: this is case (A.3.b). Then $\delta \xrightarrow{h_{\#}} \omega \xrightarrow{j'} c_{\tau}$,
where $c_{\tau} = \sum \tau_i$ for some collection of $\tau_i : \Delta_p \rightarrow B_{pq}^U$. We now
define $F(\kappa) = c_{\iota}$, where $c_{\iota} = \sum \iota_i$ and each singular simplex
 $\iota_i : \Delta_p \rightarrow Y$ is defined as follows. Let m be the smallest integer
such that $sd^m \omega$ is \mathcal{A}' -small. For each singular simplex τ_i in c_{τ} ,
there exists a singular simplex ω_i in $sd^m \omega$ such that $j'(\omega_i) = \tau_i$.
For each such ω_i , there exists a singular simplex δ_i in $sd^m \delta$ such
that $h_{\#}(\delta_i) = \omega_i$. In other words, for each τ_i in c_{τ} , there exist
 δ_i in $sd^m \delta$, such that following the e -path, $\delta_i \xrightarrow{h_{\#}} \omega_i \xrightarrow{j'} \tau_i$.

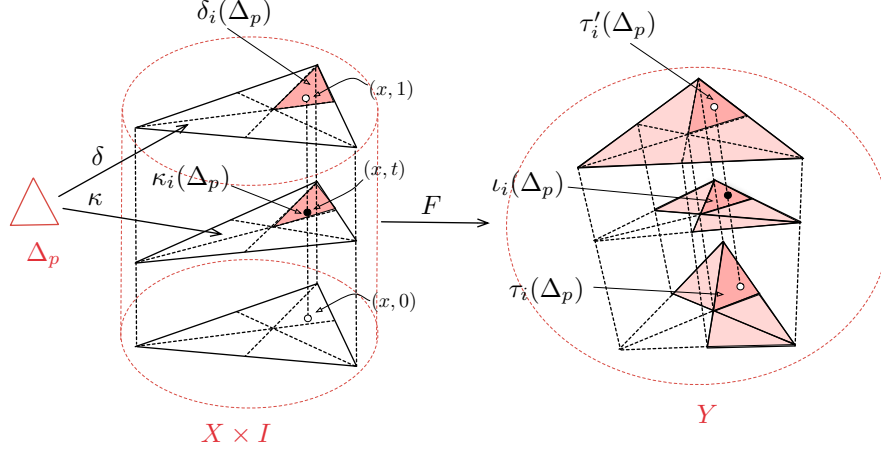


Figure 25: Case (D.1.c): illustration of F . Left: the dark shaded region is the minimal carrier of δ_i and κ_i . Right: the shaded regions from top to bottom are the minimal carriers of τ_i' , c_i and c_τ respectively; the dark shaded regions from top to bottom are the minimal carriers of τ_i' , l_i and τ_i , respectively. For simplicity, we illustrate the minimal carrier of the singular chain c_τ as the union of the minimal carriers of its simplexes before their retraction.

Meanwhile, for each such δ_i , following the e' -path gives $\delta_i \xrightarrow{f'} 0$
 $\delta_i \xrightarrow{h'_\#} \tau_i'$.

On the other hand, for each such δ_i , there exists a corresponding κ_i in $sd^m \kappa$, such that $\delta_i = \pi(\kappa_i)$. We now define l_i for each such κ_i . For all $\varepsilon \in \Delta_p$ where $\kappa_i(\varepsilon) = (x, t) \in X \times I$, $l_i(\varepsilon) = (1-t)\tau_i(\varepsilon) + t\tau_i'(\varepsilon)$. This is illustrated in Figure 25.

(D.2) $\delta(\Delta_p) \not\subseteq |\text{Sd } K|$: we again have two sub-cases:

(D.2.a) $\delta(\Delta_p) \cap |\text{Sd } K| = \emptyset$: following the e' -path and case (B.2), $\delta \xrightarrow{f'} 0 \xrightarrow{h'_\#} 0$. Since $\delta(\Delta_p) \cap |\text{Sd } K| = \emptyset$, this implies that its corresponding $\sigma \notin \text{Sd } K$. That is, for all $\hat{\sigma}_i$, $V^{\sigma_i} \cap B_{pq}^U = \emptyset$, therefore all $h(\hat{\sigma}_i)$ lie outside of $B_r(q)$. Let $\omega = h_\#(\delta) = h \circ \delta$. This means ω has its image outside of B_q^U . Then following the e -path, $\delta \xrightarrow{h_\#} \omega \xrightarrow{j'} 0$. We define $F(\kappa) = 0$.

(D.2.b) $(\delta(\Delta_p) \cap |\text{Sd } K|) \subseteq |\text{Sd } K_0|$: following the e' -path and case (B.3), we have, $\sigma \xrightarrow{f'} 0 \xrightarrow{h'_\#} 0$. On the other hand, let $\omega = h_\#(\delta) = h \circ \delta$. Then $\omega(\Delta_p) \subseteq P_0$ and so following the e -path give $\sigma \xrightarrow{h_\#} \omega \xrightarrow{j'} 0$. We define $F(\kappa) = 0$.

Construction of D. To define our chain homotopy D , we first need the following lemma:

Lemma C.6 ([27], page 171) *There exists for each space X and each non-negative integer p , a homomorphism $G_p : C_p(X) \rightarrow C_{p+1}(X \times I)$, having the following*

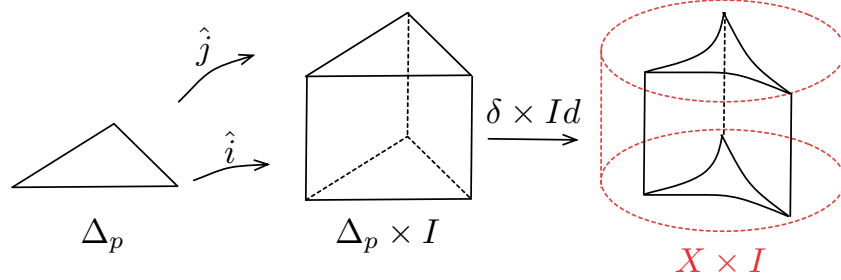


Figure 26: Illustration of G .

property: if $\delta : \Delta_p \rightarrow X$ is a singular simplex, then $\partial G\delta + G\partial\delta = j_{\#}(\delta) + i_{\#}(\delta)$, where the map $i : X \rightarrow X \times I$ carries x to $(x, 0)$, and the map $j : X \rightarrow X \times I$ carries x to $(x, 1)$.

This homomorphism is illustrated intuitively in Figure 26, where $\delta \times Id$ carries a singular $p + 1$ chain that fills up the entire prism $\Delta_p \times I$ to a singular chain on $X \times I$, and the maps $\hat{i}, \hat{j} : \Delta_p \rightarrow \Delta_p \times I$ carry each x to $(x, 0)$ and $(x, 1)$ respectively. Then, as promised, we set $D_p = F_{p+1} \circ G_p$. To show that D is a chain homotopy between e and e' , we calculate

$$\begin{aligned}
 \partial D &= \partial(FG) \\
 &= F\partial G \\
 &= F(j_{\#} + i_{\#} + G\partial) \\
 &= Fj_{\#} + Fi_{\#} + FG\partial \\
 &= Fj_{\#} + Fi_{\#} + D\partial
 \end{aligned}$$

Hence we need only show that $Fj_{\#} = e'$ and $Fi_{\#} = e$ to complete the argument. In the case when $F(\kappa)$ is defined to be 0, the corresponding $e(\delta)$ and $e'(\delta)$ are also 0, so this is no problem. In the case when $F(\kappa)$ is not defined to be 0, as shown in Figure 24 and Figure 25, $Fj_{\#}(\delta) = e'(\delta)$, and $Fi_{\#}(\delta) = e(\delta)$. This concludes the proof that the upper square in diagram 9 commutes up to chain homotopy, and thus that the bottom square of diagram 8 commutes.

C.3.4 Top Square of Diagram 8

As promised above, we now prove that the top square of diagram 8 commutes, which will complete the proof that the left face of Figure 21 commutes. In fact, the top square commutes on the chain level, which we draw directly below.

$$\begin{array}{ccc}
 C(B_p^U, \partial B_p^U) & \xrightarrow{j} & C(B_{pq}^U, \partial B_{pq}^U) \\
 \downarrow i_{\#} & & \downarrow i'_{\#} \\
 C(B_p^U, P_0) & \xrightarrow{j'} & C(B_{pq}^U, Z_0)
 \end{array}$$

Setting $e = j' \circ i_{\#}$ and $e' = i'_{\#} \circ j$, we show, once again via an exhaustive case analysis, that $e = e'$.

First we need to understand the map j for a linear singular simplex. The interpretation of j is almost the same as that of j' (case (A)). More specifically, we let $\omega : \Delta_p \rightarrow B_p^U$ be an arbitrary linear singular simplex. There are three cases:

(E.1) $\omega(\Delta_p) \subseteq B_q^U$: then $\omega \xrightarrow{j} \omega$.

(E.2) $\omega(\Delta_p) \cap B_q^U = \emptyset$: then $\omega \xrightarrow{j} 0$.

(E.3) $\omega(\Delta_p) \not\subseteq B_q^U$ and $\omega(\Delta_p) \cap B_q^U \neq \emptyset$: We have two sub-cases:

(E.3.a) ω is \mathcal{A} -small: then $\omega \xrightarrow{j} \gamma$, where $\gamma : \Delta_p \rightarrow B_{pq}^U$ is defined via the retraction-type arguments above.

(E.3.b) ω is not \mathcal{A} -small: then $\omega \xrightarrow{j'} c_\gamma$, where $c_\gamma = \sum \gamma_i$, with each $\gamma_i : \Delta_p \rightarrow B_{pq}^U$ described by the subdivision and retraction arguments we have already given.

To complete the proof, we fix an arbitrary singular simplex $\delta : \Delta_p \rightarrow B_p^U$, and again argue by cases.

(F.1) $\delta(\Delta_p) \subseteq B_q^U$: exploiting the analysis above, we note that following the e' -path results in $\delta \xrightarrow{j} \delta \xrightarrow{i'_{\#}} \delta$, while following the e -path gives $\delta \xrightarrow{i_{\#}} \delta \xrightarrow{j'} \delta$, as needed.

(F.2) $\delta(\Delta_p) \cap B_q^U = \emptyset$: here both paths result in 0.

(F.3) $\delta(\Delta_p) \not\subseteq B_q^U$ and $\delta(\Delta_p) \cap B_q^U \neq \emptyset$: here we must analyze two sub-cases:

(F.3.a) δ is \mathcal{A} -small: this implies that $\delta(\Delta_p) \subseteq X - V$. Following the e -path gives. $\delta \xrightarrow{j} \gamma \xrightarrow{i'_{\#}} \gamma$. On the other hand, δ is also \mathcal{A}' -small, since \mathcal{A} and \mathcal{A}' share the element $X - V$, and hence the e' path

$$\delta \xrightarrow{j} \delta \xrightarrow{i'_{\#}} \tau.$$

But really the fact that $X - V$ is part of \mathcal{A} and \mathcal{A}' means that τ and γ follow the same retraction, and thus $\gamma = \tau$.

(F.3.b) δ is not \mathcal{A} -small: the analysis here is the same as the last case, with some words about subdivision added.

C.4 Finale

We are now ready to finish the proof of Theorem 5.1, which boils down to verifying that diagram 5 commutes, with the vertical maps being isomorphisms. That is,

$$\begin{array}{ccccc} \dots & \rightarrow & \ker \phi_\alpha^U & \rightarrow & \ker \phi_\beta^U & \rightarrow & \dots \\ & & \uparrow \cong & & \uparrow \cong & & \\ \dots & \rightarrow & \ker \psi_\alpha & \rightarrow & \ker \psi_\beta & \rightarrow & \dots \end{array}$$

But this is now just easy diagram-chasing. Commutativity of diagram 5 follows directly from the commutativity of the bottom face of the double-cube in Figure 21, and the leftmost (rightmost) vertical isomorphism derives from our statements about the left (right) face of the double-cube. The commutativity of the top face implies that the cokernel analogue to diagram 5 commutes, after a little more algebra which we omit.

D Proof of Theorem 3.3

In this Appendix, we give a proof for Theorem 3.3. First we need a technical lemma involving some simple algebraic topology.

D.1 Absolute Homology Modules

Recall from before that $\sigma(p, r)$ is the feature size of the relative homology persistence module $\{H(B_p^{\times}, \partial B_p^{\times})\}$. On the other hand, the same thickening process also defines two absolute homology persistence modules, $\{H(B_p^{\times})\}$ and $\{H(\partial B_p^{\times})\}$. We let $\sigma_i(p, r)$ and $\sigma_b(p, r)$ denote the feature sizes of these modules. Similarly, we define $\sigma_i(p, q, r)$ and $\sigma_b(p, q, r)$, respectively, to be the feature sizes of the absolute homology persistence modules $\{H(B_{pq}^{\times})\}$ and $\{H(\partial B_{pq}^{\times})\}$.

Theorem D.1 (Relative/Absolute Lemma) *The feature size of each relative module is at least the minimum of those of its two associated absolute modules:*

$$\begin{aligned}\sigma(p, r) &\geq \min\{\sigma_i(p, r), \sigma_b(p, r)\}, \\ \sigma(p, q, r) &\geq \min\{\sigma_i(p, q, r), \sigma_b(p, q, r)\}.\end{aligned}$$

PROOF. We prove the first equality; the second can then be proven with only minor notational adjustment. For any two non-negative reals $\alpha < \beta$, and for each homological dimension $i \geq 0$, consider the following commutative diagram:

$$\begin{array}{ccccccccc} H_i(\partial B_p^{\times}(\alpha)) & \rightarrow & H_i(B_p^{\times}(\alpha)) & \rightarrow & H_i(B_p^{\times}(\alpha), \partial B_p^{\times}(\alpha)) & \rightarrow & H_{i-1}(\partial B_p^{\times}(\alpha)) & \rightarrow & H_{i-1}(B_p^{\times}(\alpha)) \\ & & \downarrow & & \downarrow & & \downarrow & & \downarrow \\ H_i(\partial B_p^{\times}(\beta)) & \rightarrow & H_i(B_p^{\times}(\beta)) & \rightarrow & H_i(B_p^{\times}(\beta), \partial B_p^{\times}(\beta)) & \rightarrow & H_{i-1}(\partial B_p^{\times}(\beta)) & \rightarrow & H_{i-1}(B_p^{\times}(\beta)) \end{array} \quad (11)$$

where the vertical maps are induced by the inclusion $X_{\alpha} \hookrightarrow X_{\beta}$ and the two rows come from the long exact sequences of the pairs $(B_p^{\times}(\alpha), \partial B_p^{\times}(\alpha))$ and $(B_p^{\times}(\beta), \partial B_p^{\times}(\beta))$ ([27]).

Suppose that the middle vertical map fails to be an isomorphism. Then the Five-Lemma ([27], p.140) tells us that at least one of the other four vertical maps will also fail to be an isomorphism. In other words, any change within the relative module must be accompanied by a simultaneous change in at least one of the two absolute modules. The inequality follows. \square

This theorem together with the definition of $\rho(p, q, r)$ implies the following inequality

$$\rho(p, q, r) \geq \min\{\sigma_i(p, r), \sigma_b(p, r), \sigma_i(p, q, r), \sigma_b(p, q, r)\}. \quad (12)$$

D.2 Proof

To prove Theorem 3.3, we will further lower bound the σ -parameters above. This is accomplished via one more lemma.

Lemma D.1 (Deformation Lemmas) *The following four claims all hold for every small enough $\delta > 0$. In each of the claims, the homotopy equivalence is given by a deformation retraction:*

$$\begin{aligned} \forall \alpha < \min\{\tau(p, r), \eta(p, r)\}, (\mathbb{X}_\alpha \cap B_r(p)) &\simeq (\mathbb{X}_\delta \cap B_r(p)), \\ \forall \alpha < \min\{\tau_0(p, q, r), \eta(p, r)\}, (\mathbb{X}_\alpha \cap \partial B_r(p)) &\simeq (\mathbb{X}_\delta \cap \partial B_r(p)), \\ \forall \alpha < \min\{\tau(p, q, r), \eta(p, q, r)\}, (\mathbb{X}_\alpha \cap B_r(p) \cap B_r(q)) &\simeq (\mathbb{X}_\delta \cap B_r(p) \cap B_r(q)), \\ \forall \alpha < \min\{\tau_0(p, q, r), \eta(p, q, r)\}, (\mathbb{X}_\alpha \cap \partial(B_r(p) \cap B_r(q))) &\simeq (\mathbb{X}_\delta \cap \partial(B_r(p) \cap B_r(q))). \end{aligned}$$

PROOF. All four claims follow from Stratified Morse Theory [19]. We prove only the first claim; the other three can be proven with only slight modifications. Consider the stratification of $B_r(p)$ with singular set $\Sigma = \mathcal{M}(p, r) \cup \partial B_r(p)$ and whatever further decomposition of Σ is needed. Setting $d = d_{\mathbb{X}}|_{B_r(p)} : B_r(p) \rightarrow \mathbb{R}$, we note that the sets $X_\alpha \cap B_r(p)$ are simply the sublevel sets of d for various parameters α . Generically, d will be a Stratified Morse function on $B_r(p)$ with its above stratification. Consider the set H of all critical points of d which have positive d -value.

We claim $H \subset (\mathcal{M}(p, r) \cup G(p, r))$: to see this, we suppose $y \in H$ and we assume first that y is in the interior of $B_r(p)$. Then y is also a critical point of the globally defined function $d_{\mathbb{X}}$, and since $d(x) = d_{\mathbb{X}}(x) > 0$, we know that $y \in \mathcal{M}$. Since y is also in $B_r(p)$ by assumption, we know in fact that $y \in \mathcal{M}(p, r)$. On the other hand, suppose that $y \in \partial B_r(p)$; we can also assume that $y \notin \mathcal{M}(p, r)$ or we are already done. Then by definition y is a critical point of the restriction of $d_{\mathbb{X}}$ to $\partial B_r(p)$. Since the gradient of this latter function is simply the projection of $\nabla d_{\mathbb{X}}$ onto $\partial B_r(p)$, we can conclude $y \in G(p, r)$.

In other words, if $\alpha < \{\tau(p, r), \eta(p, r)\}$, then $(\mathbb{X}_\alpha \cap B_r(p)) \cap H = \emptyset$, and hence the interval $[\delta, \alpha]$ contains no critical values of d . The claim then follows from the first fundamental theorem of Stratified Morse Theory [19]. \square

Finally, we finish the proof of Theorem 3.3, which we restate here for convenience.

Theorem D.2 (Geometric lower bound) *If we define*

$$\gamma = \gamma(p, q, r) = \min\{\tau(p, r), \tau(p, q, r), \eta(p, r), \eta(p, q, r)\},$$

then $\rho(p, q, r) \geq \gamma(p, q, r)$.

PROOF.

Note that $\tau(p, r) \leq \tau_0(p, r)$ and $\tau(p, q, r) \leq \tau_0(p, q, r)$ so we need not consider $\tau_0(p, r)$ and $\tau_0(p, q, r)$.

Recall $\sigma_i(p, r)$ and $\sigma_b(p, r)$ were defined to be the feature sizes of the persistence modules $\{H(B_p^{\mathbb{X}}(\alpha))\}$ and $\{H(\partial B_p^{\mathbb{X}}(\alpha))\}$, respectively.

By the first and second of the Deformation Lemmas the following holds

$$\sigma_i(p, r), \sigma_b(p, r) \geq \min\{\tau(p, r), \eta(p, r)\}.$$

For the same reason

$$\sigma_i(p, q, r), \sigma_b(p, q, r) \geq \min\{\tau(p, q, r), \eta(p, q, r)\}.$$

These inequalities, together with (12)

$$\rho(p, q, r) \geq \min\{\sigma_i(p, r), \sigma_b(p, r), \sigma_i(p, q, r), \sigma_b(p, q, r)\},$$

prove the theorem, $\rho(p, q, r) \geq \gamma(p, q, r)$. \square

Acknowledgments

All the authors would like to thank Herbert Edelsbrunner and John Harer for useful discussions and suggestions. PB would like to thank David Cohen-Steiner and Dmitriy Morozov for helpful discussions, and SM would like to thank Shmuel Weinberger for useful comments. SM would like to acknowledge the support of NSF Grant DMS-07-32260, NIH Grant R01 CA123175-01A1 and NIH Grant P50 GM 081883.

References

- [1] Mikhail Belkin and Partha Niyogi. Laplacian eigenmaps for dimensionality reduction and data representation. *Neural Computation*, 15:1373–1396, 2002.
- [2] Mikhail Belkin and Partha Niyogi. Towards a theoretical foundation for laplacian-based manifold methods. *Journal of Computer and System Sciences*, 74(8):1289–1308, 2008.
- [3] Paul Bendich, David Cohen-Steiner, Herbert Edelsbrunner, John Harer, and Dmitriy Morozov. Inferring local homology from sampled stratified spaces. In *Proceedings 48th Annual IEEE Symposium on Foundations of Computer Science*, pages 536–546, 2007.
- [4] Jean-Daniel Boissonnat, Olivier Devillers, and Samuel Hornus. Incremental construction of the delaunay triangulation and the delaunay graph in medium dimension. *Proceedings 25th Annual Symposium on Computational Geometry*, pages 208–216, 2009.
- [5] Frédéric Chazal, David Cohen-Steiner, Marc Glisse, Leonidas J. Guibas, and Steve Y. Oudot. Proximity of persistence modules and their diagrams. In *Proceedings 25th Annual Symposium on Computational Geometry*, pages 237–246, 2009.
- [6] Frédéric Chazal, David Cohen-Steiner, and André Lieutier. A sampling theory for compact sets in euclidean space. *Discrete and Computational Geometry*, 41:461–479, 2009.
- [7] Frédéric Chazal and André Lieutier. Weak feature size and persistent homology: computing homology of solids in r^n from noisy data samples. *Proceedings 21st Annual Symposium on Computational Geometry*, pages 255–262, 2005.
- [8] Jeffrey Cheeger. A lower bound for the smallest eigenvalue of the Laplacian. In *Problems in Analysis*, pages 195–199, Princeton, NJ, USA, 1970. Princeton University Press.

- [9] Fan R.K. Chung. *Spectral Graph Theory (CBMS Regional Conference Series in Mathematics, No. 92)*. American Mathematical Society, 1997.
- [10] Kenneth L. Clarkson. Tighter bounds for random projections of manifolds. *Proceedings 24th Annual Symposium on Computational Geometry*, pages 39–48, 2008.
- [11] David Cohen-Steiner, Herbert Edelsbrunner, and John Harer. Stability of persistence diagrams. *Discrete and Computational Geometry*, **37**:103–120, 2007.
- [12] David Cohen-Steiner, Herbert Edelsbrunner, John Harer, and Dmitriy Morozov. Persistence homology for kernels, images and cokernels. *Proceedings 20th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1011–1020, 2009.
- [13] V. de Silva and G. Carlsson. Topological estimation using witness complexes. *Symposium on Point-Based Graphics*, pages 157–166, 2004.
- [14] Tamal K. Dey. *Curve and Surface Reconstruction*. Cambridge University Press, 2007.
- [15] Herbert Edelsbrunner and John Harer. *Computational Topology: An Introduction*. American Mathematical Society, 2010.
- [16] Samuel Eilenberg and Norman Steenrod. *Foundations of Algebraic Topology*. Princeton University Press, 1952.
- [17] Miroslav Fiedler. A property of eigenvectors of nonnegative symmetric matrices and its application to graph theory. *Czechoslovak Mathematical Journal*, 25(4):619–633, 1975.
- [18] Evarist Giné and Vladimir Koltchinskii. Empirical Graph Laplacian Approximation of Laplace-Beltrami Operators: Large Sample Results. In *Lecture Notes-Monograph Series, Vol. 51, High Dimensional Probability*, pages 238–259. Institute of Mathematical Statistics, 2006.
- [19] Mark Goresky and Robert MacPherson. *Stratified Morse Theory*. Springer-Verlage, 1988.
- [20] Marvin J. Greenberg and John R. Harper. *Algebraic Topology A First Course*. Addison-Wesley, 1981.
- [21] Gloria Haro, Gregory Randall, and Guillermo Sapiro. Stratification learning: Detecting mixed density and dimensionality in high dimensional point clouds. *Advances in NIPS*, 17, 2005.
- [22] Allen Hatcher. *Algebraic Topology*. Cambridge University Press, 2002.
- [23] Bruce Hughes and Shmuel Weinberger. Surgery and stratified spaces. *Surveys on Surgery Theory*, pages 311–342, 2000.
- [24] R. Kannan, S. Vempala, and A. Veta. On clusterings-good, bad and spectral. In *FOCS '00: Proceedings of the 41st Annual Symposium on Foundations of Computer Science*, page 367, Washington, DC, USA, 2000. IEEE Computer Society.
- [25] Gilad Lerman and Teng Zhang. Probabilistic recovery of multiple subspaces in point clouds by geometric lp minimization, 2010.

- [26] Marina Meilă and Jianbo Shi. Learning segmentation by random walks. In *In Advances in Neural Information Processing*, pages 470–477. MIT Press, 2000.
- [27] James R. Munkres. *Elements of algebraic topology*. Addison-Wesley, Redwood City, California, 1984.
- [28] Partha Niyogi, Stephen Smale, and Shmuel Weinberger. Finding the homology of submanifolds with high confidence from random samples. *Discrete Computational Geometry*, 39:419–441, 2008.
- [29] Partha Niyogi, Stephen Smale, and Shmuel Weinberger. A topological view of unsupervised learning from noisy data. Manuscript, 2008.
- [30] Markus J. Pflaum. *Analytic and Geometric Study of Stratified Spaces*. Springer, 2001.
- [31] Colin Rourke and Brian Sanderson. Homology stratifications and intersection homology. *Geometry and Topology Monographs*, 2:455–472, 1999.
- [32] R. Vidal, Y. Ma, and S. Sastry. Generalized principal component analysis (GPCA). *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 27:1945 – 1959, 2005.
- [33] Shmuel Weinberger. *Chicago Lectures in Mathematics*, chapter The topological classification of stratified spaces. University of Chicago Press, Chicago, IL, 1994.