

HOMOLOGICAL ILLUSIONS OF PERSISTENCE AND STABILITY

by

Dmitriy Morozov

Department of Computer Science
Duke University



Ph.D. Dissertation
2008

Copyright © 2008 by Dmitriy Morozov
Some rights reserved under Creative Commons Attribution-Noncommercial License.

Abstract

In this thesis we explore and extend the theory of persistent homology, which captures topological features of a function by pairing its critical values. The result is represented by a collection of points in the extended plane called persistence diagram.

We start with the question of ridding the function of topological noise as suggested by its persistence diagram. We give an algorithm for hierarchically finding such ε -simplifications on 2-manifolds as well as answer the question of when it is impossible to simplify a function in higher dimensions.

We continue by examining time-varying functions. The original algorithm computes the persistence pairing from an ordering of the simplices in a triangulation and takes worst-case time cubic in the number of simplices. We describe how to maintain the pairing in linear time per transposition of consecutive simplices. A side effect of the update algorithm is an elementary proof of the stability of persistence diagrams. We introduce a parametrized family of persistence diagrams called persistence vineyards and illustrate the concept with a vineyard describing a folding of a small peptide. We also base a simple algorithm to compute the rank invariant of a collection of functions on the update procedure.

Guided by the desire to reconstruct stratified spaces from noisy samples, we use the vineyard of the distance function restricted to a 1-parameter family of neighborhoods of a point to assess the local homology of a sampled stratified space at that point. We prove the correctness of this assessment under the assumption of a sufficiently dense sample. We also give an algorithm that constructs the vineyard and makes the local assessment in time at most cubic in the size of the Delaunay triangulation of the point sample.

Finally, to refine the measurement of local homology the thesis extends the notion of persistent homology to sequences of kernels, images, and cokernels of maps induced by inclusions in a filtration of pairs of spaces. Specifically, we note that persistence in this context is well defined, we prove that the persistence diagrams are stable, and we explain how to compute them. Additionally, we use image persistence to cope with functions on noisy domains.

Contents

Abstract	iii
List of Figures	vi
List of Tables	viii
Acknowledgments	ix
1 Introduction	1
2 Basics	6
2.1 Complexes and homology	6
2.2 Persistence	10
2.3 Nerves	17
2.4 Distance functions	20
2.5 Piecewise-linear framework	23
3 Simplification	27
3.1 Motivation	27
3.2 Overview	30
3.3 Simplification Details	33
3.4 Lower bound	35
3.5 Higher Dimensions	36
3.6 Discussion	38
4 Persistence Vineyards	39
4.1 Vines and Vineyards	39
4.2 Updating the Pairing	41
4.3 Stability of Persistence Diagrams	47
4.4 Applications	48
4.4.1 Peptide Folding	48
4.4.2 Rank Invariant	52
4.5 Discussion	53

5	Local Homology	54
5.1	Introduction	54
5.2	Local Homology and Distance Functions	57
5.3	The (αr) -Vineyard	60
5.4	Local Homology Inference	62
5.5	Power Cell Algorithm	66
5.6	Correctness	68
5.7	Algorithm Details	70
	5.7.1 Thresholds	70
	5.7.2 Trajectories	75
5.8	Discussion	77
6	Kernels, Images, and Cokernels	79
6.1	Introduction	79
6.2	Algebra	80
6.3	Algorithms	84
6.4	Correctness	87
6.5	Transpositions	91
6.6	Applications	93
	6.6.1 Local Homology	93
	6.6.2 Noisy Domains	96
6.7	Discussion	98
7	Discussion	100
	Bibliography	102
	Index	109
	Biography	111

List of Figures

1.1	Terrain and its persistence diagram	3
2.1	Smith normal form	8
2.2	Maps between nested spaces and pairs of nested spaces	9
2.3	Definition of persistence.	10
2.4	Matrix reduction	13
2.5	Worst case example of a space and its filtration	14
2.6	Inclusions of staggered mapping cylinders.	16
2.7	Convex sets, their nerve, and its barycentric subdivision	17
2.8	Second barycentric subdivision	18
2.9	Persistence diagrams for homology inference	22
2.10	Admissible noise intervals.	23
2.11	Deformation retraction	24
3.1	ε -simplification	28
3.2	Locally-paired barycenters	31
3.3	Cell complex duality	32
3.4	Structure of the filtration during cancellation of a single pair	32
3.5	Reordering of the path	34
3.6	Overlapping pairs in the filtration	34
3.7	Lower bound for persistence-sensitive simplification	35
3.8	Poincaré sphere	37
4.1	Arrangement of piecewise rational curves	41
4.2	Obstacles to transpositions	42
4.3	Evolution of matrix updates	44
4.4	Changes in pairing	45
4.5	Sparse matrix representation	46
4.6	Peptide folding	49
4.7	Arrangement of function values at vertices	49
4.8	Vineyard of pairwise distance function	50
4.9	Pairwise distance function with a cycle	51
4.10	Representation of a bifiltration of a simplicial complex	52

5.1	Pinched torus with a membrane	56
5.2	Sublevel sets of a distance function	57
5.3	Persistence diagrams of a distance function	58
5.4	Inference from persistence diagrams for resolving radius	63
5.5	(αr) -vineyard of the one-dimensional chain of loops	65
5.6	Voronoi decomposition and the power cell	67
5.7	Delaunay triangulations	67
5.8	Spheres and distances	72
5.9	Membership trajectories	75
5.10	Membership trajectories after excision	76
5.11	Same local homology, different local dimension	78
6.1	Square diagram giving rise to kernel, image, and cokernel persistence	80
6.2	The mapping cylinder of $\mathbb{Y} \subseteq \mathbb{X}$	82
6.3	Commuting cube	83
6.4	Matrices with reordered rows	85
6.5	Kernel vineyard for inferring local homology	95
6.6	Diagram implying stability of image persistence on noisy domains	97

List of Tables

5.1	Commuting diagram relating a homology and a cohomology sequences . . .	61
6.1	Five cases in the two function setting relating kernels and images.	81
6.2	Five cases in the two function setting relating cokernels and images.	82
6.3	Six cases in which a simplex changes the kernel, the image, or the cokernel	87

Acknowledgments

I would like to thank my advisor Herbert Edelsbrunner. His guidance, support, and friendship have been invaluable for me. Herbert has introduced me to the wonderful world of topology and geometry as well as to the community of people who create and shape it. By sharing his views on science, philosophy, music, writing, Herbert has shaped mine. I could not have had a more perfect advisor. Working on this dissertation has been one of the greatest experiences of my life, and I am forever indebted to Herbert for making it so.

I am grateful to David Cohen-Steiner and John Harer for devoting so much time to me over the last few years, being patient with my mathematical naïveté, and helping me with every aspect of the research. It is impossible to overestimate their influence on this work. I thank them and Pankaj Agarwal for serving on my PhD committee and encouraging me to remember the big picture.

I would like to thank the funding agencies — NSF, DARPA, and DoE — that through various grants to Herbert Edelsbrunner have supported this work. I particularly would like to thank DARPA Program Manager Benjamin Mann for providing both inspiration and support to pursue many of the research directions explored in this dissertation.

I have been fortunate to get an opportunity to work with many wonderful people over the last few years. I thank Paul Bendich for being always willing to discuss ideas, no matter how silly or crazy, Valerio Pascucci for inspiration, Dominique Attali for her kindness and encouragement.

It has been said that a computer scientist is a device for converting coffee into theorems. The quality of the theorems in this thesis reflects the quality of Shade Tree coffee. I thank Courtney, Tracy, and Greg for creating a productive working environment.

Duke has been an unforgettable experience. Diane Riggs has created a warm, nurturing environment in our department, and I sincerely thank her for that. My fellow graduate students have taught me new ways to appreciate life. Be it good beer, snowboarding, quality food, basketball, a conversation, or a word of advice, Allister, Amit, Bei, Beth, Jaidev, Jeff, Laura, Lee, Madhu, Mason, Monika, Shobana, Varun, Yuriy, and many others have been there for me.

I want to thank Jen for her patience and support.

The path to this degree has been enduring and started long before my graduate studies. I have been lucky to have support of many amazing people in this journey. I thank Millie Cannon for believing in me, and Erich Kaltofen for showing me how much fun research can be.

Finally I would like to thank the people who started it all. I am forever grateful to my parents for instilling in me the love of learning that has served me so well in life. Their sense of humor, willpower, common sense, dedication, faith and love have made me who I am in life. My brother has been not only my first mathematics and computer science teacher, but also my role model. Words cannot express my gratitude to them.

It was my father's dream to see this dissertation. Sadly, he missed it by a year. I dedicate this work to his memory.

Chapter 1

Introduction

We are surrounded by functions. A great number of physical phenomena can be described by real-valued functions. There is a constant drive within the scientific community to make ever more elaborate and precise measurements that produce an abundance of data that is at least as difficult to analyze as it is to obtain.

For some phenomena the function manifests itself. For instance while studying internal combustion engine, one may record *temperature* or *pressure* at sample points in the combustion chamber. Using MRI for medical imaging produces a volume of measurements that represent *density* of the biological material. Measuring the *height* of the terrain at sample points on the planet results in a height function on a sphere.

Perhaps even more interesting than functions that we see directly are those functions that one constructs artificially to help interpret the data. An example that we revisit several times in this thesis is the distance function.

The problem of reconstructing a shape from a nearby point sample has received a lot of attention in the computational geometry community. Suppose that there is a space \mathbb{X} which we do not know, typically a subset of \mathbb{R}^n , and a point sample U that is close to \mathbb{X} . By close we mean that the Hausdorff distance between the space and the sample is small. The reconstruction problem asks for a space K , typically a triangulation with vertices in U , that resembles \mathbb{X} geometrically or topologically. The motivating application in computational geometry has been the need to reconstruct surfaces from point samples collected by ever more accurate and abundant in data laser scanners.

Recently the high-dimensional version of the problem attracted interest of the machine learning community. The key insight has been the recognition that one can interpret data as lying near some low-dimensional space embedded in higher dimensions. If one can recover this hypothetical space, such view can lead to a better understanding of the relationship between the data points if one measures geodesic distances along the low-dimensional space rather than the distances in the ambient space; see [77] for more details. This insight can be used to improve learning algorithms.

Short of reconstructing the space we may ask for its characteristics. One such characteristic is a topological invariant called *homology*. It keeps track of components, tunnels,

voids of the space as well as their high-dimensional counter-parts. A number of techniques have been proposed for learning homology [5, 17, 22, 23, 25, 76, 65]. The best example of practical advantage of knowing homology of the space is the work of Carlsson et al. [17] where the authors discover that at a certain scale the space describing the distribution of patches of natural images resembles a Klein bottle. This realization can be used to develop better image compressions techniques.

While a number of approaches to the reconstruction problem have been explored by the machine learning community, most of the above homology learning methods as well most surface reconstruction techniques in computational geometry [3, 33, 36] reduce to studying the distance function $d_U : \mathbb{R}^n \rightarrow \mathbb{R}$ which is defined on the ambient space of the data by measuring the distance at every point to the closest point in the data set. The resulting function provides a deep view into the shape of the data set, and we discuss it further in the later chapters of this thesis.

Persistence. It is a formidable challenge to analyze the ever-growing data sets. Many techniques have been devised in fields ranging from statistics to scientific visualization. This thesis comes out of the study of the recently introduced [39] theory of persistent homology; see [37] for a survey. By pairing critical points that are responsible for creation and destruction of topological features, persistence represents a function as a collection of paired values. Each pair symbolizes a feature, and the difference between the values represents the significance of the feature.

The idea of persistence may be the easiest to introduce as it applies to terrains¹. Given a real-valued function we consider all those points in the domain whose value is below the given threshold, i.e. the *sublevel set* of the function. As we vary the threshold, the connectivity of the sublevel set changes. For simplicity, in the Introduction we restrict our attention only to the connectivity of components; we delay the formal definition of persistence in all dimensions until the next chapter.

On the left of Figure 1.1 we have a height function $f : \mathbb{R}^2 \rightarrow \mathbb{R}$ defined on the plane with minima at points A, B, D , saddles at points C, E, G , and a maximum at F . As we sweep the threshold from negative infinity to positive infinity, the new components in the sublevel set appear: once we pass the value at D , there are three components in the sublevel set. The components created at B and D merge into one as we pass the value at E . We say that the younger component, created at D , dies, and the older component carries on. As we pass the value at saddle G , the surviving component of B merges with the older component of A , and dies. We pair the values at which a component appears and subsequently dies, and get pairs $(f(D), f(E))$, $(f(B), f(G))$, and $(f(A), \infty)$, with the last one signifying that the component of A lives forever. We record the pairs as points in the extended plane \mathbb{R}^2 with birth on the horizontal and death on the vertical axes. We call the collection *persistence diagram*, shown on the right of Figure 1.1. Since birth happens before death, all the points appear above the diagonal.

The difference between death and birth of a component is its persistence; in the di-

¹ A slightly restricted version of this application, called “topographic prominence”, is the earliest incarnation of persistence known to the author.

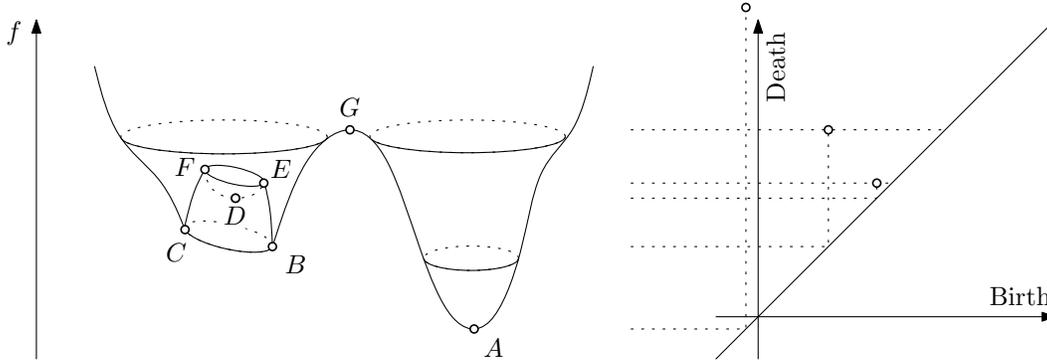


Figure 1.1: Left: terrain whose height is an example of a real-valued function. Right: its 0-dimensional persistence diagram.

agram it is the L_1 distance of the point to the diagonal. Location of the point along the diagonal is its scale, the average of the function values responsible for it. One can view each point as a feature of the function and the persistence of the point as the importance of the feature.

In our example whether the little crater on top of the mountain on the left is a real feature — and the mountain is really a volcano — depends on our application as well as the quality of the data that we have. Some of the points in the persistence diagram may represent noise, some real features. Regardless of which is which persistence diagram does not make a judgment call on what is important. There lies its strength: persistence gives an *omniscalar* description of function’s features. Empowered by the complete picture, the user is able to decide what matters.

When used for the homology learning problem in conjunction with the distance function, the omniscalar nature of persistence translates into techniques that do not require the user to make assumptions about the amount of noise in the data [25]. Instead they present her with a simultaneous view across all scales, which contrasts, for example, the methods of [65] where picking the presumed noise threshold is among the first steps of the algorithm. Moreover, using persistence not only does not require one to inject assumptions about a fixed noise level, it tells what amounts of noise are impossible for a given point sample under very mild assumptions on the conditioning of the space. We explain this observation in detail in Section 2.4.

The concept of persistence can be seen embedded within the theory of spectral sequences [61] but has not been treated as a concept in its own right until [39]. The latter paper also describes a fast algorithm for modulo 2 homology and demonstrates that persistence is relevant to applications, including the study of protein structure. The concept and the algorithm have been extended to homology over fields in [82]. The stability of persistence diagrams has been established in [25], opening the concept up to additional applications, including the inference of homology from point clouds, see also [76], the comparison of shapes, see also [16], and the analysis of discrete curvature measures, see

[24]. Persistence has been applied to the analysis of image patch data [17]. Recent results on the L_p -stability of persistence diagrams of Lipschitz functions [27] have provided insight into periodic gene expression [32]. Independently, the same ideas were developed from a somewhat different angle and restricted to zero-homology by a group of researchers in Italy [15, 44]. Using different terminology, they introduced persistence diagrams and proved stability, albeit only for the evolution of components in the sublevel sets [4]. The ideas of persistence also emerged in the work of Robins [74].

Contributions. In this thesis we explore several applications of persistence, and in doing so extend the theory. Several of the definitions that we introduce along the way are interesting in their own right.

Cohen-Steiner, Edelsbrunner, and Harer [25] express the following view of persistence: “importance [of a feature] can be quantified in terms of the amount of change necessary to eliminate a feature”. While this view is slightly misleading, as Section 3.5 of this thesis shows, it gives a helpful intuition about persistence diagrams. It also raises a natural question of simplifying a persistence diagram by eliminating features with small persistence, which we examine in Chapter 3. We define what an ε -**simplification** of a function is and investigate when it is and is not possible to find such a simplification.

Then we study what happens to persistence diagrams if the underlying function varies. Time-varying data is of great interest in practice (since indeed most physical phenomena are time-dependent processes). We introduce **persistence vineyards** which keep track of how persistence diagrams change as the functions they represent change continuously. A sample vineyard in Chapter 4 captures information about folding of a small peptide.

In subsequent chapters we embrace and explore the omniscalar paradigm of topological data analysis. To guide us we address the homology learning problem by studying the distance function. However, instead of trying to learn the global structure of the sampled space, we investigate the question of recovering its local structure at a point. The reasons for this are manifold including determining whether the space is.

In machine learning, the subfield that studies topological spaces is commonly dubbed *manifold learning*, which discloses an assumption that the spaces under investigation are high-dimensional surfaces. It is not difficult to come up with an example that does match this expectation, e.g., see the pinched torus in Figure 5.1. There is no reason to expect such more complicated spaces to not occur in practice. Moreover, if singularities in the space exist, one may expect them to represent significant behavior in data making their discovery desirable.

The question of learning the local structure is inherently multi-parameter since there is an interplay between what we mean by “local” and the scale of the noise. To preserve our focus on the omniscalar paradigm, we study a local restriction of the distance function as we vary our definition of local. The result is an $(\alpha|r)$ -**vineyard** from which one can infer the local homology of a sampled space in \mathbb{R}^n . In the course of this study we discover the significance of **resolving radius** which we introduce in Section 5.4, and which is interesting in its own right.

Finally, we refine the notion of persistence to cope with a particular setting for pairs

of functions. We use it in Chapter 6 to get a cleaner representation for local homology as well as to cope with the problem of noise in the domain rather than only in the function values.

Chapter 2

Basics

In this thesis we study topological spaces, which are sets given with collections of their open subsets. Often the spaces we consider are subsets of \mathbb{R}^n , in which case their topology is induced by \mathbb{R}^n , i.e. a subset of the space is open if it is an intersection of an open subset of \mathbb{R}^n with the given space. The strongest notion of equivalence that we use is that of *homeomorphism* which is a continuous bijection between two spaces whose inverse is also continuous. However, homeomorphism type of a space is undecidable [58], while we are interested in efficient algorithms to study topological spaces. We therefore restrict our attention to a much weaker, but easily computable topological invariant called *homology*.

2.1 Complexes and homology

We review simplicial complexes, homology groups, and briefly mention concepts from homotopy theory. We refer the reader to Hatcher [52] or Munkres [64] for a thorough study of these subjects.

Complexes. We distinguish between abstract and geometric simplicial complexes. Given a collection of sets, we say that a subset of any $p + 1$ of them is an *abstract p -dimensional simplex*, or abstract p -simplex for short. Any subset of size $q + 1$ of the p -simplex is its q -dimensional face. A *geometric p -simplex* is the convex hull of $p + 1$ affinely independent points. The convex hull of any subset of those points is its *face*. If τ is a face of σ , then σ is a *coface* of τ . A *simplicial complex*, abstract or geometric, is the collection of faces of a finite number of simplices, any two of which are either disjoint or intersect in a common face. If K is a simplicial complex in \mathbb{R}^n , then its *underlying space* is the union of its simplices together with the subspace topology inherited from \mathbb{R}^n . Any abstract simplicial complex with n simplices can be realized in \mathbb{R}^n (for example as a subset of the standard simplex Δ^n , which is the convex hull of points with coordinates $(0, \dots, 0)$, $(1, 0, \dots, 0)$, $(0, 1, 0, \dots, 0)$, \dots , $(0, \dots, 0, 1)$). Therefore we can always talk about the underlying space of a simplicial complex implying the underlying space of some realization if the simplicial complex is abstract.

We consider a subset L of a simplicial complex K , not necessarily a simplicial complex

itself. We define its *star* as the set of simplices in K that have a simplex of L as a face. Its *link* is the set of faces of simplices in the star that do not also belong to the star:

$$\begin{aligned}\text{St } L &= \{\sigma \in K \mid \exists \eta \in L, \eta \subseteq \sigma\}, \\ \text{Lk } L &= \{\tau \in K \mid \tau \subseteq \sigma \in \text{St } L, \tau \notin \text{St } L\}.\end{aligned}$$

We call 0-simplices vertices, 1-simplex edges, 2-simplices triangles, and denote the set of vertices of a simplicial complex K with $\text{Vert } K$.

Given a simplicial complex K , its first *barycentric subdivision*, $\text{Sd } K$, is an abstract simplicial complex with a vertex $\hat{\sigma}$ for each simplex σ in K and a simplex τ for each increasing chain in the face relation of the simplices in K .

The subcomplex of K that consists of all the simplices in K up to dimension p is the *p-dimensional skeleton* of K , or *p-skeleton* for short.

Homology. We consider a topological space \mathbb{X} and a triangulation K of \mathbb{X} , i.e. a simplicial complex whose underlying space is homeomorphic to \mathbb{X} . In simplicial homology, a *p-chain* is a formal sum of p -simplices in K . We use modulo 2 arithmetic implying the coefficients in the formal sum are 0 or 1. We can therefore think of the p -chains as subsets of all p -simplices, namely the ones with coefficient 1. Adding chains modulo 2, we obtain the *group of p-chains*, denoted $C_p(K)$. It is easy to see that $C_p(K)$ is abelian. The *boundary* of a p -simplex is the set of its $(p-1)$ -dimensional faces, and the boundary of a p -chain is the sum of the boundaries of its simplices. Denoting the boundary map by ∂_p , we observe that it is a homomorphism from $C_p(K)$ to $C_{p-1}(K)$. Noting that $\partial_p \partial_{p+1} = 0$, we take the sequence of groups together with the homomorphisms to obtain a *chain complex*,

$$\dots \xrightarrow{\partial_{p+2}} C_{p+1} \xrightarrow{\partial_{p+1}} C_p \xrightarrow{\partial_p} C_{p-1} \xrightarrow{\partial_{p-1}} \dots \xrightarrow{\partial_1} C_0 \xrightarrow{\partial_0} 0.$$

The *group of p-cycles* is the kernel of the p -th boundary homomorphism, $Z_p(K) = \ker \partial_p$, and the *group of p-boundaries* is the image of the $(p+1)$ -st boundary homomorphism, $B_p(K) = \text{im } \partial_{p+1}$. Since $\partial_p \partial_{p+1} = 0$, $B_p(K)$ is a subgroup of $Z_p(K)$. The *p-th homology group* of K is the quotient of the two, $H_p(K) = Z_p(K)/B_p(K)$. The *p-th Betti number* of K is the rank of its p -th homology group, $\beta_p(K) = \text{rank } H_p(K)$. Homology groups and therefore Betti numbers are invariants of the topological space \mathbb{X} , and do not depend on the choice of the triangulation K [52, 64].

Assuming an ordering of the $(p-1)$ -simplices and of the p -simplices, we can interpret the linear operator ∂_p as an *incidence matrix* D_p where $D_p[i, j] = 1$ if the i -th $(p-1)$ -simplex is a face of the j -th p -simplex, and $D_p[i, j] = 0$ otherwise. The boundary of a p -chain can be obtained by multiplication of the corresponding vector with the incidence matrix, $\partial_p(c_p) = D_p c_p$. We observe that the null space of matrix D_p is exactly the kernel of map ∂_p , i.e. the cycle group Z_p , while the column span of D_p is the image of map ∂_p , i.e. the boundary group B_{p-1} . Therefore,

$$\beta_p = \text{rank } Z_p - \text{rank } B_p = \text{rank nullspace}(D_p) - \text{rank colspan}(D_{p+1}).$$

A classical algorithm computes the Betti numbers of K by reducing its incidence matrices to Smith normal form to obtain ranks of their null spaces and column spans. It uses row and column operations to zero out all entries except along an initial portion of the diagonal, as shown in Figure 2.1. In the normal form of D_p , the zero columns form a basis of the p -cycles and the non-zero rows form a basis of the $(p - 1)$ -boundaries. We can thus read the ranks of the cycle and the boundary groups off the normal forms and get β_p for each p .

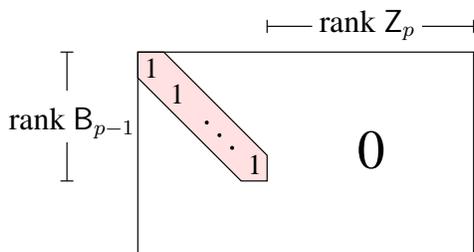


Figure 2.1: Smith normal form of the incidence matrix between $(p - 1)$ -simplices and p -simplices.

It is often advantageous to talk about *reduced Betti numbers* which can be defined by modifying the boundary homomorphism ∂_0 to map into $\mathbb{Z}/2\mathbb{Z}$ rather than 0 with $\partial_0(v) = 1$ for each vertex, and defining a map $\partial_{-1} : \mathbb{Z}/2\mathbb{Z} \rightarrow 0$ to be the zero map. We call the resulting homology groups reduced, denoted $\tilde{H}_p(K)$; their ranks are the reduced Betti numbers, $\tilde{\beta}_p(K) = \text{rank } \tilde{H}_p(K)$. They are equal to the ordinary Betti numbers except for $\tilde{\beta}_0(K) = \beta_0(K) - 1$ if the space K is non-empty, and $\tilde{\beta}_{-1}(K) = 1$ if the space is empty.

It is convenient to collect the homology groups for all dimensions. We write $H(\mathbb{X}) = (\dots, H_p(\mathbb{X}), H_{p+1}(\mathbb{X}), \dots)$ and $\tilde{H}(\mathbb{X}) = (\dots, \tilde{H}_p(\mathbb{X}), \tilde{H}_{p+1}(\mathbb{X}), \dots)$. Similarly, we define a vector of (reduced) Betti numbers, $\beta(\mathbb{X}) = (\dots, \beta_p(\mathbb{X}), \beta_{p+1}(\mathbb{X}), \dots)$ and $\tilde{\beta}(\mathbb{X}) = (\dots, \tilde{\beta}_p(\mathbb{X}), \tilde{\beta}_{p+1}(\mathbb{X}), \dots)$. Of course, only the groups for p between 0 and the dimension of \mathbb{X} are possibly non-trivial. To simplify language, we often ignore the difference between a single homology group and the entire series.

Let \mathbb{X} be a subspace of \mathbb{X}' , another topological space. Inclusion $\mathbb{X} \subseteq \mathbb{X}'$ induces a homomorphism $H_p(\mathbb{X}) \rightarrow H_p(\mathbb{X}')$ between homology groups of the same dimension. As an example consider the space \mathbb{X} in Figure 2.2. It has a single hole marked by the dashed circle that surrounds it. This circle generates a non-trivial class $\gamma \in H_1(\mathbb{X})$. In contrast, the same circle bounds in \mathbb{X}' which implies that the homomorphic image of γ in $H_1(\mathbb{X}')$ is 0. As with groups we combine the homomorphisms to form a series that maps $H(\mathbb{X})$ to $H(\mathbb{X}')$ component-wise.

We also consider pairs of topological spaces $\mathbb{X}_0 \subseteq \mathbb{X}$. We define *relative homology groups*, $H_p(\mathbb{X}, \mathbb{X}_0)$, to consist of classes generated by relative cycles, that is, chains in \mathbb{X} whose boundary is either empty or contained in \mathbb{X}_0 . As before we get one group for each dimension p and we write $H(\mathbb{X}, \mathbb{X}_0)$ for the series. For example in Figure 2.2, the dimension 1 relative homology group of the pair $(\mathbb{X}, \mathbb{X}_0)$ is generated by two classes, the

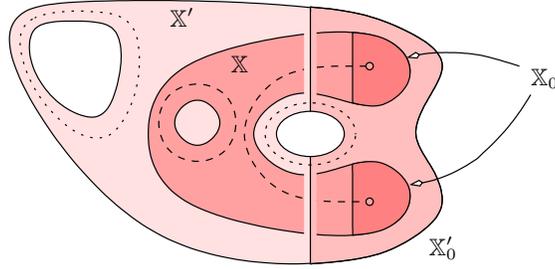


Figure 2.2: The dimension 1 homology groups of \mathbb{X} (dashed closed curve) and of \mathbb{X}' (dotted closed curves) have rank 1 and 2 but the map induced by inclusion has only rank 0. The relative homology groups of $(\mathbb{X}, \mathbb{X}_0)$ (dashed curves) and of $(\mathbb{X}', \mathbb{X}'_0)$ (dotted curves) both have rank 2 but the map induced by inclusion of pairs has only rank 1.

absolute class from before and the relative class generated by the dashed curve whose ends lie in \mathbb{X}_0 . Let $\mathbb{X}'_0 \subseteq \mathbb{X}'$ be another pair of topological spaces. We write $(\mathbb{X}, \mathbb{X}_0) \subseteq (\mathbb{X}', \mathbb{X}'_0)$ if $\mathbb{X} \subseteq \mathbb{X}'$ and $\mathbb{X}_0 \subseteq \mathbb{X}'_0$. In this case, inclusion induces again a homomorphism $H_p(\mathbb{X}, \mathbb{X}_0) \rightarrow H_p(\mathbb{X}', \mathbb{X}'_0)$. In the example in Figure 2.2, the image of the absolute class is zero, as before, and the image of the relative class is another relative class of the pair $(\mathbb{X}', \mathbb{X}'_0)$, namely the one generated by the dotted cycle surrounding the hole on the right in \mathbb{X}' . As before, we simplify notation by considering the series of homomorphisms mapping $H(\mathbb{X}, \mathbb{X}_0)$ to $H(\mathbb{X}', \mathbb{X}'_0)$ component-wise.

Homotopy. While we do not study it in this thesis, we make use of homotopy theory. Even though it is computationally intractable [58], homotopy theory is useful to us since it is stronger than homology. In the rest of this section we recall basic definitions and statements from homotopy theory, and make the former statement precise.

Given two functions $f_0 : \mathbb{X} \rightarrow \mathbb{Y}$ and $f_1 : \mathbb{X} \rightarrow \mathbb{Y}$, we say that they are *homotopic* if there exists a continuous function $F : \mathbb{X} \times [0, 1] \rightarrow \mathbb{Y}$ such that $F(x, 0) = f_0(x)$ and $F(x, 1) = f_1(x)$. We denote this by $f \simeq g$.

Two topological spaces \mathbb{X} and \mathbb{Y} are said to be *homotopy equivalent* if there exist functions $f : \mathbb{X} \rightarrow \mathbb{Y}$ and $g : \mathbb{Y} \rightarrow \mathbb{X}$ such that their compositions are homotopic to the respective identities; $f \circ g \simeq \text{id}_{\mathbb{Y}}$, $g \circ f \simeq \text{id}_{\mathbb{X}}$. In this case we say that the functions f and g are each others' *homotopy inverses*. A homotopy equivalence $f : \mathbb{X} \rightarrow \mathbb{Y}$ treated as a map on chain groups of the spaces induces an isomorphism $f^* : H(\mathbb{X}) \rightarrow H(\mathbb{Y})$ on the homology groups of the spaces. As a result, spaces that are homotopy equivalent have isomorphic homology groups. The converse is not true in general.

A helpful way to show that a subset of a space is homotopy equivalent to it is by constructing the following map. *Deformation retraction* from space \mathbb{X} to its subset $\mathbb{Y} \subseteq \mathbb{X}$ is a homotopy $r : \mathbb{X} \times [0, 1] \rightarrow \mathbb{Y}$ such that $r(x, 0) = x$, $r(\cdot, 1) \subseteq \mathbb{Y}$, and $r(y, \cdot) = y$ for all $y \in \mathbb{Y}$. It follows immediately that $r(x, 1) : \mathbb{X} \rightarrow \mathbb{Y}$ is a homotopy inverse of the inclusion $\text{id}_{\mathbb{Y}} : \mathbb{Y} \rightarrow \mathbb{X}$.

Homotopy theory studies groups formed by maps from spheres to the topological

space, $\gamma : \mathbb{S}^p \rightarrow \mathbb{X}$. The group operation $\gamma_1 + \gamma_2$ is given by collapsing the equator \mathbb{S}^{p-1} of the sphere \mathbb{S}^p to a point and applying γ_1 to the top and γ_2 to the bottom sphere of the resulting dumbbell. In Section 2.3 we make use of homotopy groups by showing that a particular map is a homotopy equivalence since it induces an isomorphism on the homotopy groups of a CW-complex. This result is known as Whitehead's theorem. See [52] for the formal statement and proof of this theorem on page 346, as well as detailed definitions of all of the above concepts.

A topological space is *contractible* if it is homotopy equivalent to a point. In this case, all Betti numbers vanish, except for β_0 which is 1.

2.2 Persistence

A concept that is central to this thesis is persistence which can be defined for any sequence of vector spaces. Suppose $\{V_\alpha\}$ is a sequence of vector spaces with homomorphisms $f_\alpha^\beta : V_\alpha \rightarrow V_\beta$ connecting every pair with $\alpha \leq \beta$, with the property that any homomorphism can be decomposed $f_\alpha^\beta = f_{\beta'}^\beta \circ f_\alpha^{\beta'}$ for any $\alpha \leq \beta' \leq \beta$ and f_α^α is the identity map. We say that a non-zero element $\lambda \in V_\alpha$ is born in V_α if for every $\alpha' < \alpha$ it does not belong to the image of $f_{\alpha'}^\alpha$. Element λ dies entering V_β if $f_\alpha^\beta(\lambda)$ is not in the image of $f_{\alpha'}^\beta$ for all $\alpha' < \alpha < \beta' < \beta$, while $f_\alpha^\beta(\lambda)$ is in the image of $f_{\alpha'}^\beta$ for some $\alpha' < \alpha$. More precisely, an entire coset is born and dies with the birth and death of a single element. We summarize in symbols; see Figure 2.3.

$$V_{\alpha'} \xrightarrow{f_{\alpha'}^\alpha} V_\alpha \xrightarrow{f_\alpha^{\beta'}} V_{\beta'} \xrightarrow{f_{\beta'}^\beta} V_\beta$$

$\lambda \in V_\alpha$ is born in V_α : $\lambda \notin \text{im } f_{\alpha'}^\alpha$ for all $\alpha' < \alpha$,
 λ dies entering V_β : $f_\alpha^{\beta'}(\lambda) \notin \text{im } f_{\alpha'}^{\beta'}$ for all α', β' with $\alpha' < \alpha < \beta' < \beta$,
 $f_\alpha^\beta(\lambda) \in \text{im } f_{\alpha'}^\beta$ for all $\alpha' < \alpha$.

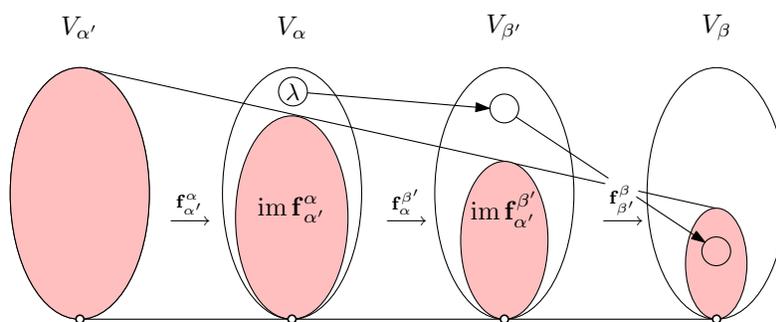


Figure 2.3: Element $\lambda \in V_\alpha$ is born in V_α and dies in V_β .

Persistent homology. The most common setting for persistence, the one that started the study of persistence in its own right is persistent homology [39, 82]. We follow the exposition in [25] in which we have a topological space \mathbb{X} and a continuous function $f : \mathbb{X} \rightarrow \mathbb{R}$.

The *sublevel set* defined by $a \in \mathbb{R}$ consists of all points with function value at most the threshold, $\mathbb{X}_a = f^{-1}(-\infty, a]$. Given $a \leq b$, the inclusion between the sublevel sets, $\mathbb{X}_a \subseteq \mathbb{X}_b$, induces a homomorphism, $\mathbf{f}_a^b : H(\mathbb{X}_a) \rightarrow H(\mathbb{X}_b)$. For $a = b$ this is an isomorphism and for $a < b$ it may or may not be an isomorphism. A value $a \in \mathbb{R}$ is a *homological critical value* of f if there is no sufficiently small $\varepsilon > 0$ for which $\mathbf{f}_{a-\varepsilon}^{a+\varepsilon}$ is an isomorphism. We assume that f is *tame*, that is, it has only finitely many critical values and every sublevel set has only finite rank homology groups.

Let $a_1 < a_2 < \dots < a_m$ be the critical values of f and consider an interleaved sequence $s_{i-1} < a_i < s_i$ for all i . This gives a sequence of spaces, $\mathbb{X}_0 \subseteq \mathbb{X}_1 \subseteq \dots \subseteq \mathbb{X}_m = \mathbb{X}$, where we simplify notation by writing $\mathbb{X}_i = \mathbb{X}_{s_i}$, and a corresponding sequence of homology groups connected by homomorphisms,

$$H(\mathbb{X}_0) \rightarrow H(\mathbb{X}_1) \rightarrow \dots \rightarrow H(\mathbb{X}_m).$$

Using the above definition of persistence for vector spaces, persistent homology concerns itself with the history of individual homology classes within this sequence. Specifically, a class γ in $H(\mathbb{X}_i)$ is *born* at a_i if it is not in the image of $\mathbf{f}_{i-1}^i = \mathbf{f}_{s_{i-1}}^{s_i}$. Furthermore, if γ is born at a_i we say it *dies entering* a_j if $\mathbf{f}_i^{j-1}(\gamma)$ is not contained in the image of \mathbf{f}_{i-1}^{j-1} but $\mathbf{f}_i^j(\gamma)$ is contained in the image of \mathbf{f}_{i-1}^j . The images of the maps \mathbf{f}_i^j are referred to as *persistent homology groups* since they consist of all homology classes born at or before a_i that live beyond a_j .

It is convenient to represent the fact that class γ is born at a_i and dies entering a_j by drawing the point (a_i, a_j) in the two-dimensional plane. By collecting the points for all p -dimensional classes we get the *dimension p persistence diagram* which we denote as $\text{Dgm}_p(f)$. Since birth necessarily happens before death all points lie above the diagonal. It is also possible that a class γ is born at a_i but does not die since it represents a class of $\mathbb{X}_m = \mathbb{X}$. In this case, we draw γ as the point (a_i, ∞) in the diagram. For technical reasons that will become clear later, we consider all points on the diagonal to be part of the persistence diagram. Similar to homology groups we get a diagram for each dimension and we write $\text{Dgm}(f)$ for the infinite series of diagrams; we simplify language by ignoring the difference between a single diagram and an entire series.

Cohen-Steiner, Edelsbrunner, and Harer [26] use relative homology to augment the above sequence of homomorphisms giving rise to persistent homology. Denoting the superlevel set of a function f at threshold a by $\mathbb{X}^a = f[a, \infty)$, equal to the set of points in \mathbb{X} with value at least a , and above simplifying notation $\mathbb{X}^i = \mathbb{X}^{s_i}$, we obtain a sequence.

$$\begin{array}{ccccccc} 0 = & H(\mathbb{X}_0) & \rightarrow & H(\mathbb{X}_1) & \rightarrow & \dots & \rightarrow & H(\mathbb{X}_m) & = & H(\mathbb{X}) \\ & & & & & & & \downarrow & & \\ 0 = & H(\mathbb{X}, \mathbb{X}^0) & \leftarrow & H(\mathbb{X}, \mathbb{X}^1) & \leftarrow & \dots & \leftarrow & H(\mathbb{X}, \emptyset) & = & H(\mathbb{X}) \end{array}$$

Using the definition of persistence for vector spaces, we can distinguish between three types of pairs. There are classes born and dying going up, classes born going up and dying coming down, and classes born and dying coming down. The authors call the resulting persistence pairing *extended persistence* and referred to the three types of pairs as *ordinary*,

extended, and *relative*. It is convenient to record the pairs as points with the value of the sublevel or superlevel set responsible for the homology group as its birth or death value. We get three types of subdiagrams named after the pairs: ordinary, extended, and relative — with the first type always lying above the diagonal, and the last one always below the diagonal.

We use the following theorem throughout this thesis to show equivalence of various sequences of vector spaces, in particular, the ones arising in applications and the ones we can treat algorithmically.

PERSISTENCE EQUIVALENCE THEOREM. Sequences of vector space V_α and V'_α have the same persistence pairing iff for every index α in the sequence there is an isomorphism $j_\alpha : V_\alpha \rightarrow V'_\alpha$, and for every pair of indices α and β the following diagram commutes.

$$\begin{array}{ccc} V_\alpha & \longrightarrow & V_\beta \\ \downarrow j_\alpha & & \downarrow j_\beta \\ V'_\alpha & \longrightarrow & V'_\beta \end{array}$$

The proof of the theorem is trivial, but wordy, so we omit it. This theorem appears as Corollary 3.1 in [82].

The persistence homology setting that is crucial for algorithmic purposes is a piecewise constant function $\bar{f} : K \rightarrow \mathbb{R}$ defined on a simplicial complex. We say more about where such functions come from in Sections 2.4 and 2.5, but for now we note that the function value on a face must not exceed that on a coface. Let $\sigma_1, \sigma_2, \dots, \sigma_m$ be a sequence of the simplices in K ordered by the function value with ties broken by dimension, so that faces always precede cofaces. Writing $K_i = \{\sigma_j \mid j \leq i\}$, we call the sequence $\emptyset = K_0 \subset K_1 \subset K_2 \subset \dots \subset K_m = K$ a *filtration* of K if all K_i are complexes or, equivalently, the faces of every simplex precede the simplex in the given sequence. Considering homology groups of each complex in the filtration and the homomorphisms between them induced by inclusion, we are again in the persistent homology setting. Edelsbrunner, Letscher, and Zomorodian [39] give an algorithm for computing the persistence pairing of the sequence of homology groups of a filtration in worst-case time cubic in the number of simplices. We now recall their algorithm using its interpretation in [29].

Computation. To compute persistent homology for the sequence of complexes K_i we let D be the m -by- m incidence matrix. We *reduce* D using left-to-right modulo-2 column additions until the lowest one of every non-zero column is in a unique row. Initializing R and V to the incidence and the identity matrices and letting $lowR[i]$ be the row index of the lowest one in column i of R , or 0 if the entire column is zero, we can formalize the algorithm as follows.

```

R = D; V = I;
for i = 1 to m do
  while  $\exists k < i$  with  $lowR[k] = lowR[i] \neq 0$  do
    add column  $k$  to column  $i$  in  $R$  as well as in  $V$ 

```

Equivalently, the reduced matrix is obtained by multiplying the incidence matrix from the right with an upper-triangular matrix, $R = DV$, such that the map from the non-zero columns of R to the row indices of their lowest ones is injective. As we prove in Chapter 4, R is not unique but the map is. By construction, the rows of R and V correspond to

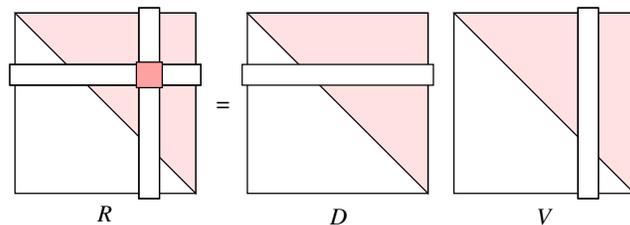


Figure 2.4: The reduced matrix equals the incidence matrix times the chain matrix. All three are upper-triangular.

individual simplices, same as the rows of D , but the columns of R and V correspond to chains. Specifically, column i of R stores the boundary of the chain stored in column i of V . We call σ_i *positive* if its addition to K_{i-1} gives birth to a homology class. Equivalently, column i of V stores a cycle and column i of R is zero. Symmetrically, we call σ_i *negative* if its addition to K_{i-1} gives death to a homology class. Equivalently, column i of V stores a chain that is not a cycle and column i of R is non-zero. The significance of the lowest one in this column of R is that the negative σ_i is paired with the positive σ_ℓ , with $\ell = \text{low}R[i]$, which gives birth to the class that σ_i kills.

PROOF. We prove correctness of the above algorithm by showing that the following two statements are true.

1. $R[i] = 0$ if and only if a class is born in $H(K_i)$;
2. $\text{low}R[j] = i$ if and only if a class born in $H(K_i)$ dies entering $H(K_j)$.

Our proof is inductive: we show that the algorithm is correct for all the complexes in the filtration up to K_l . The claim is trivially true for K_0 .

For inductive step we assume that the claim is true for all $l' < l$. We observe that exactly one class is born or dies in any single complex of the filtration since only one chain group C_p changes by one basis element, and therefore only one of the homology groups H_p or H_{p-1} changes by one basis element. Therefore, once we show that the algorithm reports a birth or a death in complex K_l correctly, the reverse direction in the above statements follows immediately.

If $0 = R[l] = D \cdot V[l]$, then by definition column $V[l]$ represents a cycle. Since matrix V is upper-triangular and invertible, $V[l, l] = 1$, and the cycle $V[l]$ contains simplex σ_l . Therefore, $V[l]$ represents a new class in $H(K_l)$, furthermore the class is non-trivial since K_l contains no cofaces of σ_l . This proves the first statement.

Suppose $\text{low}R[l] = i$, then column $R[l]$ represents a cycle which is born in K_i because of the ordering of the rows. Furthermore, that cycle belongs to a trivial class in $H(K_l)$

because it is by definition the boundary of the chain represented by the column $V[l]$. It remains to show that a class born in $H(K_i)$ does not die before K_l . This follows from inductive hypothesis since if a class born in $H(K_i)$ dies before $H(K_l)$, then there would be a column $l' < l$ with $\text{low}R[l'] = i$. In this case column $R[l]$ could be reduced further, a contradiction. \square

Cubic worst case. We show that in the worst case the above algorithm performs $\Omega(m^3)$ operations. The existence of this example should be contrasted to the experimentally observed sub-quadratic running time for filtrations that arise from applications.

We describe the space as well as the ordering of the simplices. Let $n = \lfloor (m + 29)/7 \rfloor$, $v = \lfloor (n - 1)/2 \rfloor$, and note that both n and v are in $\Omega(m)$. In our filtration, all vertices appear before all edges, and all edges appear before all triangles. The indices that we assign to the simplices are within their respective classes (e.g., edge labeled n appears before the triangle labeled 1); some indices are negative, which is done for simplicity.

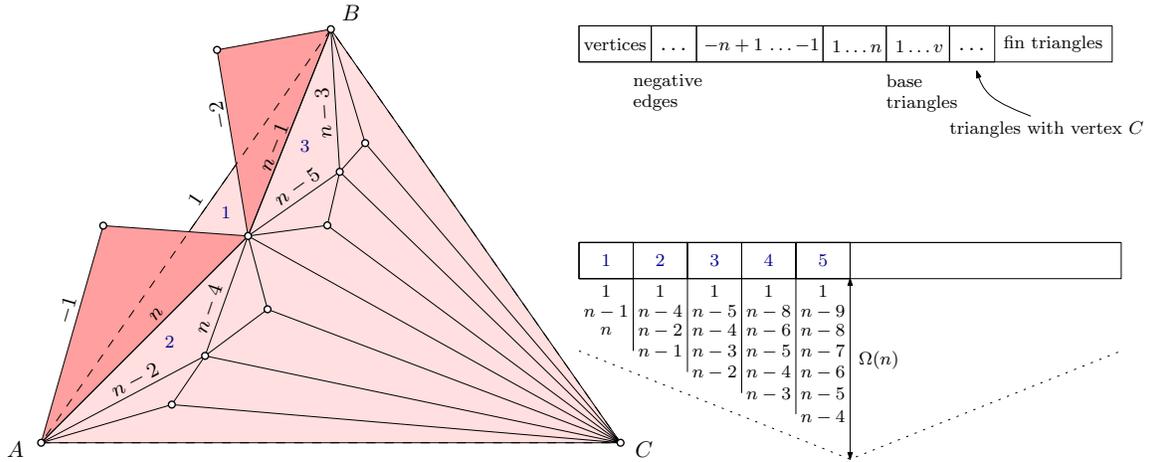


Figure 2.5: Left: space for the worst case performance of the persistence algorithm. Right top: filtration of the space. Right bottom: sparse matrix representation of the portion of the reduced matrix corresponding to the base triangles.

Figure 2.5 illustrates the construction of our space as well as the assignment of indices. Starting with triangle ABC , we add v vertices inside the triangle in the following manner: we place the first vertex V_1 near the middle of edge AB , the second vertex V_2 near the middle of AV_1 , V_3 near the middle of BV_2 , V_4 near AV_2 , V_5 near BV_3 , V_6 near V_1V_2 , V_7 near V_1V_3 , and so on, moving from both ends inwards at each stage. The edges joining C with vertices V_i are the first to appear in the filtration, each one merging its vertex V_i with the component containing C . These edges are not important in our argument, so we do not label them. Edge AB gets index 1, and the remaining edges are assigned indices from the ends inwards similar to the vertices: AV_1 gets n , BV_1 gets $n - 1$, AV_2 gets $n - 2$, BV_3 gets $n - 3$, V_1V_2 gets $n - 4$, V_1V_3 gets $n - 5$, and so on, see Figure 2.5. Similarly, the triangles

are assigned indices from the ends inwards in stages: ABV_1 gets 1, AV_1V_2 gets 2, BV_1V_3 gets 3, AV_2V_4 gets 4, BV_3V_5 gets 5, and so on. We call these triangles the *base triangles*.

In addition, we place $n - 1$ vertices above the plane of triangle ABC , one above each edge AV_i , BV_j , and V_iV_j , and join them to those edges (Figure 2.5 depicts only two of the $n - 1$ such vertices). One of the edges joining the vertex above the edge k to its endpoints merges the component containing the endpoint of the edge k with the vertex above the plane. We do not label this edge. The other edge gets index $k - (n + 1)$ which is negative. The triangle formed gets an index larger than v , so that the triangles not in the plane of ABC appear last in the filtration. We call them *fin triangles*.

Consider what happens when the above algorithm reduces the resulting incidence matrix. There are two interesting parts to its execution. First, the base triangles 1 to v are processed, their columns in matrix R build up $\Omega(n)$ ones each corresponding to edges n to $n - v$, see bottom right of Figure 2.5. Second, when the fin triangles $v + 1$ to $v + n$ are processed, the search for the lowest ones goes through all the columns of the base triangles, adding columns of size $\Omega(n)$. As a result, for $\Omega(n)$ triangles we perform $\Omega(n)$ merges each of which takes time $\Omega(n)$. It follows that the total running time is $\Omega(n^3) = \Omega(m^3)$.

Stability. Cohen-Steiner, Edelsbrunner, and Harer [25] proved a stability result for persistence diagrams that provides a jumping off point for most of this thesis. We restate it here. Given two functions $f, g : \mathbb{X} \rightarrow \mathbb{R}$, defined on some topological space \mathbb{X} , we define the distance between them to be the L_∞ -norm of their difference: $\|f - g\|_\infty = \sup_{x \in \mathbb{X}} |f(x) - g(x)|$. The *bottleneck distance* between the persistence diagrams of f and g is the infimum over all bijections $\gamma : \text{Dgm}_p(f) \rightarrow \text{Dgm}_p(g)$ of the supremum distance between the corresponding points:

$$d_B(\text{Dgm}_p(f), \text{Dgm}_p(g)) = \inf_{\gamma} \sup_{u \in \text{Dgm}_p(f)} \|u - \gamma(u)\|_\infty.$$

For technical reasons the functions are required to be tame, and the space \mathbb{X} triangulable.

STABILITY THEOREM. If $f, g : \mathbb{X} \rightarrow \mathbb{R}$ are two continuous, tame functions then for any $p \geq 0$, the bottleneck distance between their dimension p persistence diagrams is not greater than the distance between the functions: $d_B(\text{Dgm}_p(f), \text{Dgm}_p(g)) \leq \|f - g\|_\infty$.

We give an elementary, combinatorial proof of this theorem in Chapter 4.

Arbitrary maps. While filtrations of sublevel sets of real-valued functions is the most common setting in practice, we do not have to restrict ourselves to homomorphisms between homology groups induced by inclusion. Given a sequence of spaces and arbitrary continuous maps between every pair of them,

$$\mathbb{X}_1 \xrightarrow{h_1} \mathbb{X}_2 \xrightarrow{h_2} \dots \xrightarrow{h_{n-1}} \mathbb{X}_n,$$

we get a sequence of induced homomorphisms between the homology groups of the respective spaces:

$$H(\mathbb{X}_1) \xrightarrow{h_1^*} H(\mathbb{X}_2) \xrightarrow{h_2^*} \dots \xrightarrow{h_{n-1}^*} H(\mathbb{X}_n).$$

To treat this setting algorithmically we reduce it to inclusions of the following spaces. Let $\mathbb{X}'_1 = \mathbb{X}_1$, and $\mathbb{X}'_{i+1} = \mathbb{X}'_i \cup \mathbb{X}_i \times [0, 1] \cup \mathbb{X}_{i+1}$ be the staggered mapping cylinder, where $(x, 0) \in \mathbb{X}_i \times \{0\}$ is identified with $x \in \mathbb{X}_i \subseteq \mathbb{X}'_i$, while $(x, 1) \in \mathbb{X}_i \times \{1\}$ is identified with $h_i(x) \in \mathbb{X}_{i+1}$. Figure 2.6 illustrates this construction.

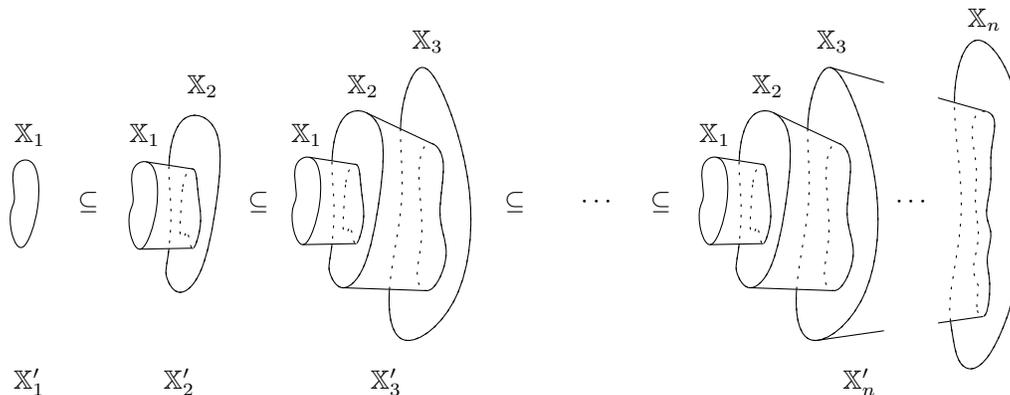


Figure 2.6: Inclusions of staggered mapping cylinders.

It is not difficult to verify that the persistence pairing of the sequence of homology groups $H(\mathbb{X}'_1) \rightarrow \dots \rightarrow H(\mathbb{X}'_n)$ with homomorphisms induced by inclusion is the same as that of the sequence of homomorphisms

$$H(\mathbb{X}_1) \xrightarrow{h_1^*} \dots \xrightarrow{h_{n-1}^*} H(\mathbb{X}_n).$$

Indeed, considering the following diagram

$$\begin{array}{ccc} H(\mathbb{X}_i) & \xrightarrow{h_{j-1}^* \circ \dots \circ h_i^*} & H(\mathbb{X}_j) \\ \downarrow & & \downarrow \\ H(\mathbb{X}'_i) & \longrightarrow & H(\mathbb{X}'_j) \end{array}$$

with all maps except the top horizontal one induced by inclusion, we notice that the vertical maps are isomorphisms since \mathbb{X}'_i deformation retracts onto \mathbb{X}_i . In addition the diagram commutes since for any cycle λ in \mathbb{X}_i , including it into \mathbb{X}'_i and then \mathbb{X}'_j produces a cycle that retracts onto the image $h_{j-1} \circ \dots \circ h_i(\lambda)$ included into \mathbb{X}'_j . Therefore, we can apply Persistence Equivalence Theorem.

Other fields and languages. As Zomorodian and Carlsson point out [82], it is possible to phrase all of the above material (and by extension everything we do in this thesis) using arbitrary fields as coefficients in chain and homology groups with very minor changes: by adding divisions where necessary. We restrict ourselves to $\mathbb{Z}/2\mathbb{Z}$ only to simplify the exposition: there is no mathematical reason for it.

It is also worth noting that all of the material in this thesis can be rephrased in other terms, e.g., modules of [82]. This would be merely a translation between different languages. The language of this thesis is chosen for its simplicity. This judgment, of course, is purely subjective and stems from habit rather than objections to other approaches.

2.3 Nerves

A combinatorial object that is very helpful in bridging the gap between functions that arise in practice and their combinatorial representations as simplicial complexes is the nerve. Given a collection of sets we consider an abstract simplicial complex in which a simplex is present iff all of its sets intersect (it is trivial to verify that it is a simplicial complex). We call such a simplicial complex the *nerve* of the collection of sets. A classic result called the Nerve Lemma states that if the intersection of every subset of the sets in the collection is contractible, the nerve of the collection is homotopy equivalent to its union [57]. We need a stronger statement that the map between the two sets suggested by the Nerve Lemma is a homotopy. In the remainder of this section we make this map precise, and prove the needed theorem.

Realizing nerves. Let C be a finite collection of convex sets in \mathbb{R}^n , write N for the nerve of C , and let $\text{Sd } N$ be the first barycentric subdivision of the nerve. We draw $\text{Sd } N$ in \mathbb{R}^n by mapping each vertex $\hat{\sigma}$ to a point $f(\hat{\sigma})$ in the intersection of the sets that correspond to the vertices of σ , as in Figure 2.7. Extending this map by piecewise linear interpolation to the simplices gives a map $f : |\text{Sd } N| \rightarrow \mathbb{R}^n$. Note that each simplex in $\text{Sd } N$ is contained in a single set in C which implies that the image of f is contained in $\bigcup C$. As suggested

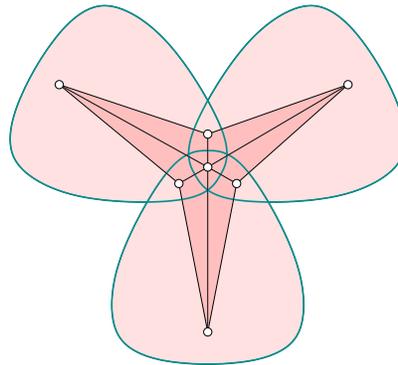


Figure 2.7: A collection of three convex sets in the plane and a piecewise linear drawing of the barycentric subdivision of the nerve in which each vertex maps to a point in the intersection of the corresponding sets.

in Figure 2.7 such a drawing does not generally exist for the nerve, which is the reason we consider its barycentric subdivision.

Let now $\gamma : \mathbb{S}^p \rightarrow \bigcup C$ be a map of the p -dimensional sphere into the union of the convex sets. We are interested in constructing a map $\mu : \mathbb{S}^p \rightarrow |\text{Sd } N|$ whose composition with f is homotopic to γ , $f \circ \mu \simeq \gamma$. For this purpose let L be a triangulation of \mathbb{S}^p that is sufficiently fine such that the image of every simplex is contained in a single but not necessarily unique set in C . In other words, we map each simplex $\tau \in L$ to a set $C_\tau \in C$ such that $\gamma(\tau) \subseteq C_\tau$. Since τ corresponds to a vertex $\hat{\tau}$ in the barycentric subdivision of L and C_τ is also a vertex in the nerve this gives a map from the vertices of $\text{Sd } L$ to the vertices of N . It is easy to see that this is a vertex map, that is, the vertices of a simplex in $\text{Sd } L$ map to the vertices of a simplex in N , a necessary condition to extend the map by piecewise linear interpolation. However, the same is not true if we substitute $\text{Sd } N$ for N . We therefore consider the second barycentric subdivision of L , $\text{Sd}^2 L = \text{Sd}(\text{Sd } L)$, map the new vertices to the corresponding barycenters of simplices in N , and finally extend the thus obtained vertex map by piecewise linear interpolation to $\mu : |\text{Sd}^2 L| \rightarrow |\text{Sd } N|$.

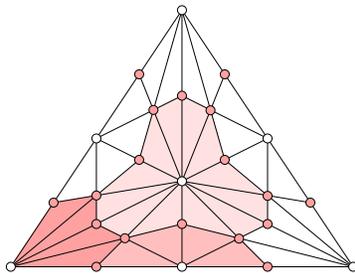


Figure 2.8: The second barycentric subdivision of a triangle in L . The shading indicates the stars of three vertices of the first barycentric subdivision.

We finally construct a homotopy $H : \mathbb{S}^p \times [0, 1] \rightarrow \bigcup C$ that equals $f \circ \mu$ at $t = 0$ and γ at $t = 1$. It deforms the image under $f \circ \mu$ of the star of a vertex $\hat{\tau}$ of $\text{Sd } L$ in $\text{Sd}^2 L$ to the image under γ of the simplex τ of L . Such stars are illustrated in Figure 2.8. To create the deformation we move every new vertex \hat{v} of $\text{Sd}^2 L$ toward the vertex $\hat{\tau}$ of the simplex $v \in \text{Sd } L$ whose corresponding simplex $\tau \in L$ has minimum dimension. To describe the resulting homotopy, we write $g(\hat{v}) = \hat{\tau}$, noting that $g(\hat{v}) = \hat{v}$ if \hat{v} is a vertex that already belongs to the first barycentric subdivision. For the vertices of $\text{Sd}^2 L$ the homotopy is defined by

$$H(\hat{v}, t) = (1 - t)f(\mu(\hat{v})) + t\gamma(g(\hat{v}))$$

and then extended by piecewise linear interpolation for every t . To see that this homotopy is well defined we note that the image of every simplex is contained in the same convex set during the entire homotopy. Let η be a simplex in the star of a vertex $\hat{\tau}$ of $\text{Sd } L$ in $\text{Sd}^2 L$. At the beginning, when $t = 0$, the image of the vertices of η all belong to the set C_τ because they are all equal or adjacent to $\hat{\tau}$ in $\text{Sd}^2 L$ and f maps barycenters to the intersection of the sets that correspond to the vertices of its simplex. By convexity, the image of the entire simplex η is contained in C_τ . At the end, when $t = 1$, the image of η is a subset of $\gamma(\tau)$,

which is contained in C_τ by assumption of L being sufficiently fine. Finally, every point moves linearly so if the image of η is contained in C_τ both at the beginning and at the end then it is contained in C_τ during the entire homotopy.

Applying Whitehead's Theorem. We use a consequence of Whitehead's Theorem [52, p. 346] to construct our second homotopy equivalence, between the pairs $\bigcup C$ and N . To state the theorem, assume the convex sets in the collection C are the closures of the maximal cells in a CW complex. If $\bigcup C$ is connected and the map $f : |\text{Sd } N| \rightarrow \bigcup C$ described above induces an isomorphism on the homotopy groups of $\bigcup C$ and of the nerve N of C for all dimensions p , then the map f is a homotopy equivalence. Constructing a homotopy equivalence thus reduces to proving that the induced map on the homotopy groups is an isomorphism.

NERVE SUBDIVISION LEMMA. Let C be the collection of closures of maximal cells of a CW complex, each a convex set in \mathbb{R}^n , N the nerve of C , and $f : |\text{Sd } N| \rightarrow \bigcup C$ obtained by piecewise linear interpolation of its values at the vertices. If $f(\hat{\sigma})$ 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.

PROOF. First we note that f induces a bijection between the connected components of $\text{Sd } N$ and of $\bigcup C$. We then apply the following argument to each component separately. Equivalently, we assume without loss of generality that $\text{Sd } N$ and $\bigcup C$ are both connected.

We want to show that the induced map on the p -dimensional homotopy group, $f_* : \pi_p(\text{Sd } N) \rightarrow \pi_p(\bigcup C)$, is an isomorphism. We do this in two steps first showing that f_* is surjective and second that it is injective. To show that it is surjective we prove that for each map $\gamma : \mathbb{S}^p \rightarrow \bigcup C$ there is a map $\mu : \mathbb{S}^p \rightarrow |\text{Sd } N|$ such that $f \circ \mu$ is homotopic to γ . But this we already did in the preceding paragraph when we discussed maps of the p -sphere to $\bigcup C$. It remains to show that f_* is injective. For this we consider a map $\gamma : \mathbb{S}^p \rightarrow \bigcup C$ that extends to a map on the ball bounded by the sphere, $\bar{\gamma} : \mathbb{B}^{p+1} \rightarrow \bigcup C$. We use the same construction as before to define a map $\bar{\mu} : \mathbb{B}^{p+1} \rightarrow |\text{Sd } N|$ such that $f \circ \bar{\mu} \simeq \bar{\gamma}$. From the construction it is clear that the restriction μ of $\bar{\mu}$ to the sphere satisfies $f \circ \mu \simeq \gamma$. It follows that the preimage of zero under f_* is zero, in other words, the kernel of f_* is zero. Hence, f_* is injective and therefore bijective. Whitehead's Theorem applies showing that f is indeed a homotopy equivalence, as required. \square

Since Whitehead's theorem works not just for CW complexes but also for the spaces homotopy equivalent to them, it suffices to assume that C is a collection of convex sets in \mathbb{R}^n . We note that it is possible to rephrase the above proof for the case of cells with all intersections contractible. We use convex cells for simplicity because they suffice for the needs of this thesis below and in Chapter 5. Chazal and Oudot have independently obtained a similar result in [23].

2.4 Distance functions

As mentioned in Introduction, a function that captures a great deal of information about an object is a distance function. Suppose Y is a subset of \mathbb{R}^n . We let $d_Y : \mathbb{R}^n \rightarrow \mathbb{R}$ be the *distance function* defined by $d_Y(x) = \inf_{y \in Y} \|x - y\|$. For each real number α the corresponding sublevel set consists of all points at distance at most α , $Y_\alpha = d_Y^{-1}(-\infty, \alpha] = d_Y^{-1}[0, \alpha]$.

In many applications the space of interest is a finite point sample U . In this case we can compute the persistence diagram of the distance function d_U using one of the following two combinatorial objects.

Čech complex. We notice that sublevel set $U_\alpha = d_U^{-1}[0, \alpha]$ is a union of balls of radius α centered at the points in U , $U_\alpha = \bigcup_{u \in U} B_\alpha(u)$. We call the nerve of the collection of balls $\{B_\alpha(u)\}_{u \in U}$ the Čech complex of point set U for parameter α , denoted $\check{C}(\alpha)$. A simplex σ is in the Čech complex $\check{C}(\alpha)$ iff there is a sphere of radius α that encloses the vertices of the simplex. Since the balls are convex, Nerve Lemma tells us the Čech complex has the same homotopy type as the sublevel set of the distance function; $\check{C}(\alpha) \simeq U_\alpha$. It follows that the homology groups of the Čech complex $\check{C}(\alpha)$ and the union of balls U_α are isomorphic. Moreover, each vertex of the first barycentric subdivision of the Čech complex represents an intersection of balls $B_\alpha(u)$. Mapping each vertex into its respective intersection and interpolating linearly on the simplices of the first barycentric subdivision gives us a homotopy equivalence by the Nerve Subdivision Lemma. As a result, the following diagram commutes, and its vertical maps are isomorphisms.

$$\begin{array}{ccc} H(\check{C}(\alpha)) & \longrightarrow & H(\check{C}(\beta)) \\ \downarrow & & \downarrow \\ H(U_\alpha) & \longrightarrow & H(U_\beta) \end{array}$$

Persistence Equivalence Theorem implies that the pairing of births and deaths in the two sequences is the same.

We can define a function $\check{c} : \check{C}(\infty) \rightarrow \mathbb{R}$ which maps each simplex into the radius of its smallest enclosing sphere; $\check{c}(\sigma) = \min\{\alpha \mid \sigma \in \check{C}(\alpha)\}$. By definition, each Čech complex is a sublevel sets of the function, $\check{C}(\alpha) = \check{c}^{-1}(-\infty, \alpha]$. Ordering simplices by the function \check{c} and breaking ties by dimension, we obtain a filtration which can be used with the algorithm described in Section 2.2. Gärtner et al. give an algorithm to compute the value of $\check{c}(\sigma)$ for any simplex of the Čech complex [43, 45].

The size of the Čech complex with parameter infinity is exponential in the number of points, $\text{card } \check{C}(\infty) = 2^{|U|}$. We can minimize the waste by observing that the homology of dimension greater than $(n - 1)$ is trivial, and therefore we are only interested in the n -dimensional skeleton of the Čech complex whose size is $\binom{|U|}{n+1}$.

Alpha shapes. A more efficient combinatorial representation of the distance function is the alpha shape complex, or alpha shapes for short [34]. To introduce it we need to recall

some background from computational geometry. Given a set of points U , the *Voronoi cell* of $u \in U$ is the set of points that are closer to u than to any other point,

$$V(u) = \{x \in \mathbb{R}^n \mid \|x - u\| \leq \|x - v\|, v \in U\}.$$

Each $V(u)$ is the intersection of finitely many closed half-spaces and therefore a convex polyhedron. Collectively, the cells cover the entire space thus forming the *Voronoi decomposition* of \mathbb{R}^n , $\text{Vor}(U|\mathbb{R}^n)$. Its nerve is called the *Delaunay triangulation*, $\text{Del}(U|\mathbb{R}^n)$. Considering the restrictions of the balls around the points to the respective Voronoi cells, $V(u) \cap B_\alpha(u)$, we note that their union is equal to the union of balls. The intersection of convex sets is convex, therefore the nerve of this collection is homotopy equivalent to the union of balls. We call this nerve the *alpha shape* for parameter α , $\text{AS}(\alpha)$. Edelsbrunner showed that inclusion of the alpha shape into the union of balls is a homotopy equivalence [35] by constructing an explicit deformation retraction. The result also follows from the Nerve Subdivision Lemma. Therefore, the following diagram, where all the maps are induced by inclusion,

$$\begin{array}{ccc} \text{H}(\text{AS}(\alpha)) & \longrightarrow & \text{H}(\text{AS}(\beta)) \\ \downarrow & & \downarrow \\ \text{H}(U_\alpha) & \longrightarrow & \text{H}(U_\beta) \end{array}$$

commutes and Persistence Equivalence Theorem tells us that persistence diagrams of the alpha shape and the distance function are the same.

Mapping each simplex to the radius of its smallest empty circumsphere, we get a function $\varrho_0 : \text{Del}(U|\mathbb{R}^n) \rightarrow \mathbb{R}$. Ordering simplices according to this function and breaking ties by dimension we obtain a filtration suitable for the algorithm described in Section 2.2.

Persistence of d_U . The persistence diagrams of the distance function d_U can be obtained from the reduced matrices computed using the algorithm in Section 2.2. Specifically, each lowest one, $l = \text{lowR}[i]$, corresponds to a pair of simplices σ_l and σ_i . We draw the point $(\check{c}(\sigma_l), \check{c}(\sigma_i))$ or $(\varrho_0(\sigma_l), \varrho_0(\sigma_i))$, depending on whether we used Čech or alpha-shape filtration, in the diagram whose dimensions is that of simplex σ_l .

We illustrate some of the strengths of persistence using distance functions. Suppose we want to estimate the homology of some unknown subset \mathbb{X} of \mathbb{R}^n from a point sample U . We call U an ε -*approximation* of \mathbb{X} if the Hausdorff distance between U and \mathbb{X} is at most ε . Equivalently, U is contained in $\mathbb{X}_\varepsilon = d_{\mathbb{X}}^{-1}[0, \varepsilon]$ and, symmetrically, \mathbb{X} is contained in $U_\varepsilon = d_U^{-1}[0, \varepsilon]$. It follows that the maximum difference between the distance functions defined by U and by \mathbb{X} is at most ε . The converse is also true. Therefore U is an ε -approximation of \mathbb{X} iff $\|d_U - d_{\mathbb{X}}\|_\infty \leq \varepsilon$. Additionally we assume that the space \mathbb{X} is well-behaved. Calling the smallest positive homological critical value of the distance function $d_{\mathbb{X}}$ the *homological feature size* of \mathbb{X} , $\text{hfs } \mathbb{X}$, we assume that it exceeds 4ε . In this case the following theorem follows from the Stability Theorem as was pointed out by Cohen-Steiner et al. [25].

HOMOLOGY INFERENCE THEOREM. Given an ε -approximation U of space $\mathbb{X} \subseteq \mathbb{R}^n$ with homological feature size of \mathbb{X} exceeding 4ε , $\text{hfs } \mathbb{X} > 4\varepsilon$, the p -dimensional Betti number of \mathbb{X} is equal to the number of points in the upper-left quadrant of the persistence diagram $\text{Dgm}_p(d_U)$ with lower-right corner at $(\varepsilon, 3\varepsilon)$.

The proof follows from Stability Theorem by examining Figure 2.9. Indeed, because of the

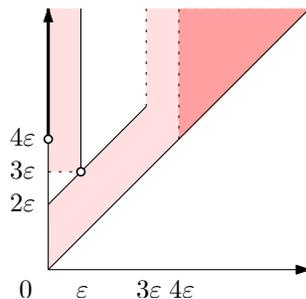


Figure 2.9: Under the conditions of the Homology Inference Theorem, the dark regions contain the persistence diagram of $d_{\mathbb{X}}$. The light regions contain the persistence diagram of d_U .

assumption on the homological feature size of \mathbb{X} it follows that no point in its persistence diagram has a birth or death value in the interval $(0, 4\varepsilon]$, therefore all the points in the persistence diagram of \mathbb{X} lie in the dark shaded regions in the Figure 2.9. I.e. the homology classes in the filtration of the sublevel sets of the distance function $d_{\mathbb{X}}$ are either born at 0 and die after 4ε , or are born after 4ε . Since $\|d_{\mathbb{X}} - d_U\|_{\infty} \leq \varepsilon$, Stability Theorem tells us that the persistence diagram of d_U looks like the lightly shaded portion of Figure 2.9. The number of points in the upper-left quadrant at $(\varepsilon, 3\varepsilon)$ is the same as the number of points with birth equal 0 in the diagram of $d_{\mathbb{X}}$. The claim follows.

It would be misleading to say that we need the power of the Stability Theorem to obtain the Homology Inference Theorem. Indeed, as observed by Chazal and Lieutier [21], its proof follows trivially from the following sequence of homomorphisms between homology groups

$$H(\mathbb{X}) \rightarrow H(U_{\varepsilon}) \rightarrow H(\mathbb{X}_{2\varepsilon}) \rightarrow H(U_{3\varepsilon}) \rightarrow H(\mathbb{X}_{4\varepsilon}).$$

All the homomorphisms are induced by inclusions of respective spaces. The assumption on the homological feature size of space \mathbb{X} tells us that the image of $H(\mathbb{X})$ in $H(\mathbb{X}_{4\varepsilon})$ is isomorphic to $H(\mathbb{X}_{2\varepsilon})$, and therefore the image of $H(U_{\varepsilon})$ in $H(U_{3\varepsilon})$ is isomorphic to $H(\mathbb{X})$ which is exactly the statement of the homology inference theorem.

However, it is equally misleading to say that the Homology Inference Theorem reflects the full power of the Stability Theorem. Examining the structure of the persistence diagrams of function d_U we notice that they must contain two regions with no points, the vertical strip between births of ε and 3ε , and the small triangle adjacent to the death axis.

This suggests that given a point sample U , and persistence diagrams $\text{Dgm}(d_U)$ of its distance function, we can conclude that the Hausdorff distance between U and \mathbb{X} cannot be ε if the regions of the diagram that are supposed to be empty for this ε contain a point. In other words under assumptions that the hypothetical space \mathbb{X} is well-behaved and given sample U is close to \mathbb{X} we can conclude what the Hausdorff distance between U and \mathbb{X} cannot be!

To formalize above discussion we say that ε is an *admissible* noise level, i.e. Hausdorff distance between U and \mathbb{X} , for point (x, y) if the point lies in the shaded region in the diagram of d_U in Figure 2.9. It is inadmissible otherwise. If ε is admissible for every point in the diagram then it is admissible for the entire diagram; otherwise it is not. Considering a single point (x, y) in a persistence diagram of d_U , we get the following three ranges of admissible ε .

$$x > 3\varepsilon \quad \text{i.e.} \quad \varepsilon < x/3 \tag{2.1}$$

$$x \leq \varepsilon \text{ and } y > 3\varepsilon \quad \text{i.e.} \quad x \leq \varepsilon < y/3 \tag{2.2}$$

$$y - x \leq 2\varepsilon \quad \text{i.e.} \quad \varepsilon \geq (y - x)/2 \tag{2.3}$$

The intervals are summarized in Figure 2.10. The first condition accounts for the points

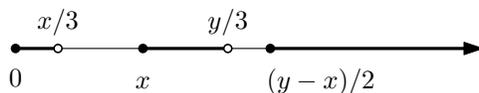


Figure 2.10: Admissible noise intervals.

that represent the features born after 3ε . The second condition accounts for the points in the upper-left quadrant. If the point falls into this region, it counts towards the true homology of the space \mathbb{X} . The third condition represents the expanded diagonal. It is the right-most region in Figure 2.10, and it can be arbitrarily-close to the origin for low persistence points. Therefore, for low persistence points with high birth value all noise levels are admissible.

Taking the union of the inadmissible noise intervals for every point we obtain the inadmissible noise intervals for the entire diagram. The small values of ε are always admissible since we can always interpret a point sample as a space on its own, typically over-fitting our model. Same goes for the large values of ε : zooming out sufficiently far away any space looks like a point.

2.5 Piecewise-linear framework

Another set of functions that is of interest in practice is the real-valued functions $f : |K| \rightarrow \mathbb{R}$ defined on the vertices of a triangulation K and interpolated linearly on the interiors of the simplices. Such functions are common in practice (when the underlying space is sampled at discrete points), and are of interest in scientific visualization. We restrict our

attention to the functions defined on the triangulations of d -manifolds, i.e. $|K| = \mathbb{M}^d$. We assume that f is non-degenerate, that is the function values are different at all the vertices. Using these function values we refine the notions of star and link. Specifically, the *lower star* of a vertex u is the set of simplices in the star for which u has the maximum value of any vertex. The *lower link* of u is the set of faces of simplices in the lower star that do not also belong to the lower star:

$$\begin{aligned} \text{St}_-u &= \{\sigma \in \text{St } u \mid v \in \sigma \Rightarrow f(v) \leq f(u)\}, \\ \text{Lk}_-u &= \{\tau \in \text{Lk } u \mid v \in \tau \Rightarrow f(v) < f(u)\}. \end{aligned}$$

Upper stars and *upper links* are defined symmetrically. We observe that if f is non-degenerate, the lower and upper stars and links of a vertex do not depend on the function values but only on their ordering by the function values.

Considering the evolution of the sublevel sets of such linearly-interpolated function f , we note that their homology can only change when the sublevel set passes a value of a vertex. Indeed, every sublevel set deformation retracts onto a simplicial complex; we construct the retraction as follows. For every simplex whose interior is intersected by the given level set $f^{-1}(a)$ we consider the maximal face not contained in the sublevel set, and the maximal face contained in it. Examining the join of the two faces, which is defined as the disjoint line segments connecting every point of one simplex with every point of the other simplex, we obtain the necessary deformation retraction r_a , see Figure 2.11. Specifically, each point z on a join line segment (x, y) with $f(x) < a$ and $f(y) > a$,

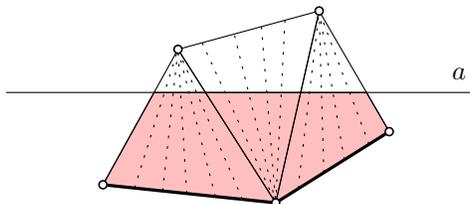


Figure 2.11: Deformation retraction of a sublevel set of a piecewise-linear function. The retraction follows dotted joins of faces above and below the level set threshold. It agrees on the faces of the simplices. The shaded sublevel set deformation retracts on the union of lower stars of the vertices within the sublevel set shown in bold.

follows the retraction $r_a(z, t) = (1 - t)z + tx$. By construction the sublevel set retracts onto the union of lower stars of the vertices in the sublevel set. Denoting the union by $K_a = \bigcup_{v \in \text{Vert } K, f(v) \leq a} \text{St}_-v$, it follows that its inclusion into the sublevel set $f^{-1}(-\infty, a]$ is a homotopy equivalence, with the retraction r_a its homotopy inverse.

Therefore, one may construct simplicial complexes homotopy equivalent to the sublevel sets of the function by adding the lower stars of the vertices in the order of increasing function value. The following diagram, where all maps are induced by inclusion, com-

mates.

$$\begin{array}{ccc}
 \mathrm{H}(f^{-1}(-\infty, a]) & \xrightarrow{f_a^b} & \mathrm{H}(f^{-1}(-\infty, b]) \\
 \uparrow & & \uparrow \\
 \mathrm{H}(K_a) & \xrightarrow{g_a^b} & \mathrm{H}(K_b)
 \end{array}$$

Since the vertical maps are isomorphisms, Persistence Equivalence Theorem tells us that the persistence diagrams of the two sequences of homology groups are the same.

Ensuring that faces precede cofaces within individual lower stars, we obtain a total ordering on the simplices which we call the *lower-star filtration*. Using the algorithm of Section 2.2 we can compute the pairing of the simplices in the filtration. From the above discussion it follows that if a p -dimensional simplex σ in the lower star of vertex s is paired with a $(p + 1)$ -simplex τ in the lower star of vertex t , then we have a point $(f(s), f(t))$ in the p -dimensional persistence diagram of the function, $\mathrm{Dgm}_p(f)$.

If vertices s and t are the same, then point $(f(s), f(t))$ lies on the diagonal and has persistence zero. If all simplices in the lower star of the vertex are paired among themselves, then the homology of the sublevel set does not change as it passes the value of that vertex, and we call the vertex *regular*. It is not difficult to see that this happens iff the reduced Betti numbers of the lower link of the vertex are zero, i.e. its lower link has trivial homology. More generally p -dimensional homology of the sublevel set can change with the addition of the lower star of a vertex only if the $(p - 1)$ - or p -dimensional homology of its lower link is non-trivial. This leads to a convenient characterization of vertices in the lower-star filtration of a manifold. Specifically, we say that if the only non-zero reduced Betti number of its lower link is $\tilde{\beta}_p(\mathrm{Lk}_-v) = 1$, then vertex v is a $(p + 1)$ -saddle, or minimum if $p = -1$, or maximum if $p = (d - 1)$, i.e. entire link is the lower link. Equivalently, we may say that vertex v is critical with *index* $(p + 1)$ if $\tilde{\beta}_p(\mathrm{Lk}_-v) = 1$. If we isolate the persistence pairing of a single lower star within the lower-star filtration, then observation that unpaired simplices reflect the homology of the domain gives us the following lemma.

DIMENSION-INDEX LEMMA. A vertex v is regular iff all simplices in its lower star are paired among themselves. Otherwise, it is critical with index equal to the dimension of the simplex not paired within the lower star.

This classification reflects the classification of critical points in the smooth generic case of Morse functions [59, Theorem 3.2] if the former is considered in terms of the changes to the homology of the sublevel sets produced by each critical point. If the lower link is more complicated than a single non-trivial cycle, we say vertex v is a multi-saddle combining $\tilde{\beta}_p(\mathrm{Lk}_-v)$ $(p + 1)$ -saddles for each p .

In light of the above discussion it is natural to think of function $\bar{f} : K \rightarrow \mathbb{R}$ that assigns to each simplex the maximum value that f attains on it, $\bar{f}(\sigma) = \max_{x \in \sigma} f(x)$. One can think of \bar{f} as a piece-wise constant approximation of function f . If f is generic, i.e. its values on all the vertices of the simplicial complex are unique, then ordering simplices by their values under \bar{f} and breaking the ties by dimension gives us exactly the lower-star filtration of f described above. Therefore, persistence diagrams resulting from \bar{f} are the

same as those of f . Such piece-wise constant functions supply a convenient notation in the next two chapters.

Chapter 3

Simplification

3.1 Motivation

Topological analysis can be used to make sense of real-valued data, to detect interesting features and to observe patterns that cannot be seen in the raw. Regardless of how the data is obtained, whether it is observed in experiments or computed in simulations, data is unfortunately always burdened with noise. While the source of the noise may range from purely physical such as imprecise measurements to purely computational such as the choice of a triangulation, the difficulties it creates always remain. In this chapter we consider the problem of ridding the data of that noise by simplifying the function it defines. Various parts of this chapter are based on the joint work and technical discussions with Nina Amenta, Dominique Attali, Herbert Edelsbrunner, Marc Glisse, Samuel Hornus, Francis Lazarus, and Valerio Pascucci [6, 40].

It is important to note that whether something is noise or a feature is in the eyes of the beholder. We endorse the idea of Cohen-Steiner, Edelsbrunner, and Harer [25] that the importance of a feature can be quantified by the amount of change necessary to eliminate it. We therefore study the question of how one would eliminate a feature in order to both understand what parts of the domain it occupies, and what the function looks like without it.

To state our results, we first introduce the central concept of this chapter. Let \mathbb{X} be a topological space, $f : \mathbb{X} \rightarrow \mathbb{R}$ a continuous function, $\text{Dgm}_p(f)$ its dimension p persistence diagram, and ε a positive constant.

DEFINITION. A *dimension p strong ε -simplification* of f is a non-degenerate function $g : \mathbb{X} \rightarrow \mathbb{R}$ such that $\|f - g\|_\infty \leq \varepsilon$ and all persistence diagrams of g are the same as those of f except for $\text{Dgm}_p(g)$ which is the same as $\text{Dgm}_p(f)$ but with all off-diagonal points at L_1 -distances at most ε from the diagonal removed.

See Figure 3.1 for an illustration. The corresponding notion of a *weak ε -simplification* would allow for the remaining points in the persistence diagram to move by at most ε . In this chapter we concern ourselves only with the strong ε -simplification, and refer to

it as simply an ε -simplification. Naturally, any strong ε -simplification is also a weak ε -simplification.

Once we know that ε -simplifications exist for all dimensions, we can iterate the construction and erase the points close to the diagonal in all persistence diagrams. We refer to the resulting function as an ε -simplification of f . In this chapter, we consider the prob-

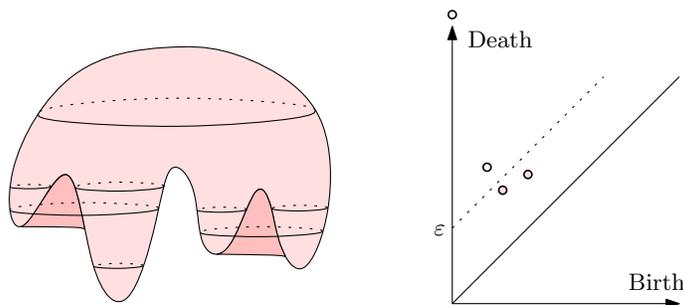


Figure 3.1: Left: two embeddings of a 2-manifold \mathbb{M} in \mathbb{R}^3 . The functions $f, g : \mathbb{M} \rightarrow \mathbb{R}$ are the height functions of the light shaded and the combined light and dark shaded embeddings. Right: the dimension 0 persistence diagrams of f and g . The two points below the threshold distance ε from the diagonal are present in the persistence diagram of f but not in that of its ε -simplification g . The other two points appear in both diagrams.

lem of finding ε -simplifications of a function f , either restricted to a single dimension or iterated across all dimensions. Our main result is a constructive proof that for 2-manifolds such simplifications exist.

SIMPLIFICATION THEOREM FOR 2-MANIFOLDS.

- A. Given a 2-manifold \mathbb{M} , a function $f : \mathbb{M} \rightarrow \mathbb{R}$, a constant $\varepsilon > 0$, and a dimension $p = 0, 1$, there exists a dimension p ε -simplification $g : \mathbb{M} \rightarrow \mathbb{R}$.
- B. For $p = 0, 1$ and all $\varepsilon > \delta > 0$ there exists a 2-manifold \mathbb{M} and a function $f : \mathbb{M} \rightarrow \mathbb{R}$ such that if $g : \mathbb{M} \rightarrow \mathbb{R}$ is a dimension p ε -simplification of f then $\|f - g\|_\infty > \varepsilon - \delta$.

We also demonstrate a negative result in higher dimensions, namely that there are 3-manifolds such that for any function on them, there are values of ε such that ε -simplification does not exist.

The problem of simplifying continuous functions has been studied before, in many different areas and from many different angles. The work related most directly to ours is on the simplification of Morse-Smale complexes initiated in [38]. Such complexes capture information about the gradient vector field by partitioning the domain into regions of uniform flow. While the simplification algorithms given in [14, 38, 49] follow the persistence order, they only simplify the complex and not the function itself. The use of the simplified complex together with the original data may be tolerable for visualization purposes,

but it is not satisfactory when the simplified data is used in the subsequent data analysis stage. An example of a weak ε -simplification technique is the persistence-sensitive flooding described in [1]. It is worth noting that the example used for the proof of part B of the Simplification Theorem shows that the error bound of $\varepsilon/2$ for simplification of a single pair of critical vertices claimed by Bremer et al. [14] is in general unachievable.

In their original paper [39], Edelsbrunner et al. also consider the question of topological simplification. However, there exist significant differences between their work and the results presented in this chapter. The most obvious distinction comes from the problem statement itself. Edelsbrunner et al. propose to move all the points of the persistence diagram towards the diagonal regardless of their persistence and do not bound the magnitude of the involved motion; in this chapter, we require points of persistence higher than ε to remain in place (or move by at most ε in the weak case). In addition, we make explicit guarantees about the distance between the simplified and the original functions.

Forward corollaries. The following results are corollaries of the Pairing Change Theorem from page 45 in the next chapter. We do not create any circular arguments by stating them here since the next chapter does not rely on any results from this one. We choose this order only for the convenience of exposition.

In the next section we replace piece-wise linear function $f : |K| \rightarrow \mathbb{R}$ by a piece-wise constant function $\bar{f} : K \rightarrow \mathbb{R}$ defined on the simplices as described in Section 2.5. Suppose that we continuously change the function values at the simplices. As a result the points in the persistence diagram move, but not more than the amount of change of the values. Even though the motion is therefore continuous, the pairs defining the points in the diagram can *switch* simplices, but only at moments in time when these simplices have the same value.

SWITCH LEMMA. Transposing two consecutive simplices σ_i and σ_{i+1} in the filtration can only affect the persistence pairs containing σ_i and σ_{i+1} .

Chapter 4 gives an algorithm to maintain the pairing if two adjacent simplices are transposed and the new sequence of complexes remains a filtration. In the following sections, we will be transposing adjacent simplices σ_i, σ_{i+1} in the ordering. We get the first constraint on switches between persistence pairs by observing that the indices in each pair are contiguous and increasing.

SAME DIMENSION LEMMA. Transposing simplices of different dimensions preserves the persistence pairing.

A crucial second constraint on how switches between pairs can happen follows from the analysis in Chapter 4. To describe it, we call two pairs of simplices, (σ_i, σ_j) and (σ_k, σ_l) , *nested* if $i < k < l \leq j$ and *disjoint* if $i < j < k < l$. To use these notions for unpaired simplices, we consider them artificially paired with a dummy simplex with subscript equal to infinity and we permit equality when we compare subscripts that are infinite.

NESTED-DISJOINT LEMMA. During a transposition of two consecutive simplices, the pairs can switch these simplices iff the pairs are nested or disjoint both before and after the transposition.

This lemma in particular implies that if before the transposition there exist k and l with $k < i < i + 1 < l$ such that σ_k is paired with σ_{i+1} and σ_i is paired with σ_l , then after σ_i and σ_{i+1} transpose we still have the same two pairs.

3.2 Overview

In this section we present a high-level view of our approach to finding an ε -simplification. We discuss the details of the algorithm in the next section.

Basic strategy. Starting with a function $f : |K| \rightarrow \mathbb{R}$ defined on the vertices and linearly interpolated on the interiors of the simplices, to compute persistence, we construct the lower-star filtration as described in Section 2.5. Our goal is to obtain function $g : |K'| \rightarrow \mathbb{R}$ with $|K'| = |K|$ also defined on the vertices and linearly interpolated, which satisfies the definition of ε -simplification. It is natural to try to manipulate the values at vertices of the function f to obtain g while maintaining the correct pairing in the lower-star filtration. This approach is explored in [40]. However, in the next two sections we pursue a different strategy. We start with a lower-star filtration given by f and assign to each simplex the maximum of the values of its vertices, thus obtaining a piece-wise constant function $\bar{f} : K \rightarrow \mathbb{R}$, introduced in Section 2.5. We manipulate the filtration directly by changing values of individual simplices, i.e. by changing function \bar{f} . Once it is in the form that gives us the right persistence diagrams, we convert it back into the lower-star filtration by defining function values on the vertices of $K' = \text{Sd } K$, the first barycentric subdivision of K . To describe how to do the latter operation, we first consider the following definition.

DEFINITION. In a filtration, simplices σ and τ are said to be *locally paired* if τ comes immediately after σ .

We note that if σ and τ are locally paired, then σ is a codimension 1 face of τ . This observation can be seen, for example, directly from the algorithm that computes the pairing described in Section 2.2. Because of the procedure described in the next paragraph, locally paired simplices can be thought of as having the same value, and thus being a pair of persistence zero.

From filtration to lower-star filtration. Suppose we have a filtration with values assigned to individual simplices. Taking the first barycentric subdivision of the filtered simplicial complex, we can assign the value of each simplex to its barycenter, and linearly interpolate on the interiors. The resulting lower-star filtration gives the same pairing of values as the filtration that we started with. However, if the function is generic, i.e. the value on every vertex is unique, then all vertices in this lower-star filtration are critical in the sense of Section 2.5. Indeed, the lower link of every vertex is the boundary of the simplex that it

subdivided, so on a 2-manifold barycenters of triangles become maxima, barycenters of edges become saddles, while the original vertices become minima.

However, suppose that we have a local pair (σ, τ) in the starting filtration. Let $\hat{\sigma}$ and $\hat{\tau}$ be their respective barycenters. Suppose that we reverse their order, i.e. we exchange the values assigned to the two vertices. We claim that as a result the vertices $\hat{\sigma}$ and $\hat{\tau}$ become regular. Indeed, the lower link of $\hat{\tau}$ is now boundary of τ without σ , which is a ball. While the lower link of $\hat{\sigma}$ becomes $\hat{\tau}$ together with its join to the boundary of σ , which also results in a ball; see Figure 3.2. Performing such a reversal for every local pair we see that



Figure 3.2: First barycentric subdivision of a triangle τ paired with its face σ . Left: lower star of $\hat{\sigma}$ is two vertices that form the boundary of edge σ , and lower star of $\hat{\tau}$ is the 1-sphere that is the boundary of the triangle τ . Right: after reversal the lower stars of $\hat{\sigma}$ and $\hat{\tau}$ are shaded, the edges of their lower links are bold.

starting with a filtration we can obtain a lower-star filtration where the persistence pairing is that of non-locally paired simplices, while barycenters of all locally-paired simplices become regular vertices.

Canceling pairs and duality. Armed with the above reordering procedure, we want to change all pairs below the given threshold into local pairs. We discuss how to simplify a 0-dimensional diagram for a function on a 2-manifold, i.e. vertex-edge pairs in the filtration. The 1-dimensional diagram can be handled symmetrically using the following observations.

If K is a triangulation of a 2-manifold, we denote by K^* the cell complex dual to K , i.e. K^* contains a vertex for each triangle, an edge for each edge, and a 2-dimensional cell for each vertex. The duals of cofaces of simplex σ in triangulation K become faces of its dual σ^* in K^* , see Figure 3.3.

The first barycentric subdivisions of complexes K and K^* are the same. We can assign to each vertex $\hat{\sigma}$ in the barycentric subdivision $\text{Sd } K$ the value equal to the index of σ in the filtration. Interpolating linearly on the simplices, and computing persistence of the resulting function, we get a pairing between indices that is the same as in the original filtration. Cohen-Steiner et al. [26] have shown that for a function f defined on a manifold, the persistence pairing of f and $-f$ is the same, except that we take negative of values, and births and deaths are switched. But the negative of function f defined above is exactly the same as assigning to each vertex $\hat{\sigma}$ of the first barycentric subdivision $\text{Sd } K^* = \text{Sd } K$ the value equal to the index of cell σ^* in the filtration of K^* minus the cardinality of K^* plus one, i.e. $-\bar{f}(\hat{\sigma}) = \text{index}(\sigma^*) - (\text{card } K^* + 1)$. Since the persistence pairing is the

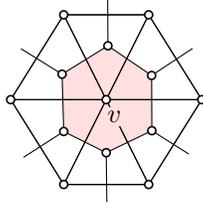


Figure 3.3: Closure of the star of vertex v and its dual (open) cell complex. Shaded cell is dual to vertex v . Its boundary is dual to the cofaces of v .

same, it follows that cells τ^* and σ^* are paired in the filtration of K^* iff σ and τ are paired in the filtration of K .

We perform the cancellations by decreasing the values associated with edges, which interpreted after the above reordering corresponds to lowering saddles. One's initial inclination may be to also increase the values at vertices, however, as example in Section 3.4 shows, it is sometimes impossible to change values on the extrema since they may get stuck as they encounter other critical vertices.

We consider the vertex-edge pairs (v, e) in order of increasing persistence, so that if one pair is nested in another, we process the nested pair first. We lower a contiguous set of simplices T , initially set $T = \{e\}$, and make sure that we do not change the persistence pairing in the process. The structure of the filtration is shown in Figure 3.4. Let ω be the simplex that immediately precedes T in the ordering of the simplices. Lowering T means

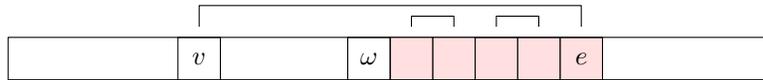


Figure 3.4: Structure of T during cancellation of (v, e) pair. T is contiguous in the filtration and shaded in the figure. It consists of locally paired vertices and edges with edge e being the only non-locally paired simplex. Simplex ω immediately precedes T in the filtration.

either moving ω past T by assigning all simplices in T a value slightly less than $f(\omega)$, or expanding T to include ω by setting the values of all simplices of T equal to $f(\omega)$.

The former approach is preferable. A difficulty arises when ω is locally paired and its paired face is also a face of a simplex in T since moving T below ω would destroy the local pair, and we cannot move T past the shared face. In this case, we expand T which preserves the local pair. If ω is a non-locally paired simplex, its persistence is higher than (v, e) , and we cannot afford to move the corresponding point in the persistence diagram. Therefore, we need to make sure that as we lower T past ω , the pairing of the simplices does not change. This requirement dictates two properties we maintain as invariants, namely that e be the only non-locally paired simplex in T and that T is an open path. We finish when T reaches vertex v , at which point we add v to the set, and reorder the simplices in T , so that they all become locally paired.

3.3 Simplification Details

As we lower the set of simplices T described above we maintain the following three properties as invariants.

- I. The only non-locally paired simplex in T is e ;
- II. T forms an open path;
- III. T is a contiguous subsequence in the filtration.

The invariant is trivially true when $T = \{e\}$.

Case analysis. Let ω be the simplex immediately preceding T in the filtration. Then the following possibilities arise.

- Case 1. ω is a non-locally paired triangle. Since T is a path, all simplices in T are either vertices or edges. From the Same Dimension Lemma it follows that we can transpose T and ω without any changes in pairing.
- Case 2. ω is a non-locally paired edge. The Same Dimension Lemma and the Nested-Disjoint Lemma tell us that we can transpose ω and T . Indeed, when ω precedes a vertex, we can transpose them without pairing changes. Subsequently, when ω precedes an edge, the pairing is neither nested nor disjoint. The same is true when ω precedes edge e because of the order in which we consider the pairs.
- Case 3. $\omega \neq v$ is a non-locally paired vertex. ω cannot be a face of an edge in T . Otherwise it would be paired with e , since the latter would be merging ω with another component. Therefore, we can transpose ω and T without changes in pairing.
- Case 4. $\omega = \tau$ is a triangle paired with edge σ immediately preceding it in the filtration. We can transpose the local pair (σ, τ) and T in the filtration: triangle τ encounters no obstacles by the Same Dimension Lemma, while σ avoids conflicts with vertices by the Same Dimension Lemma, and edges by the Nested-Disjoint Lemma.
- Case 5. $\omega = \tau$ is an edge paired with vertex σ immediately preceding it.
 - Case 5a. If vertex σ is not a face of an edge in T , we can transpose the local pair and the subsequence T . Indeed, considering the transposition of (σ, τ) with a single local pair in T , we see that no change in pairing is possible (the edge in each local pair remains paired with its respective face). Similarly, when (σ, τ) reaches e , transposing τ and e , and then σ and e cannot change the pairing.
 - Case 5b. If vertex σ is a face of an edge in T , then we expand T to include (σ, τ) . Since T is an open path before the expansion, it remains an open path after we add one of its boundary vertices σ and another edge τ that has σ as a face.

It is trivial to check that the invariant is maintained in each of the above cases.

We repeatedly perform transpositions and expansions prescribed by the above case analysis until vertex v immediately precedes subsequence T . Once this occurs, we add v to T , and reorder all the simplices in T , so that they all become local pairs.

Reordering. From the above case analysis and its guiding invariant, it follows that T is a path consisting of locally-paired negative edges and positive vertices with the exception of edge e and vertex v which are not locally paired. Moreover, v is in the boundary of the path T , while its other side is open. We enumerate all the edges and vertices of the path as we traverse it starting from $v = v_1$. This gives us two sequences of vertices and edges: v_1, \dots, v_k , and e_1, \dots, e_k . We replace the sequence T in the filtration by the sequence $v_k, e_k, v_{k-1}, e_{k-1}, \dots, v_1, e_1$. Since each sequence is a traversal of the path T , we see that the resulting sequence satisfies requirements of a filtration, i.e. the faces precede the cofaces. In addition, since v_i is a face of e_i , all the pairs in the resulting filtration are locally paired, which was our original goal. Figure 3.5 illustrates the reordering procedure. In terms of the function value we can think of the reordering as assigning to all

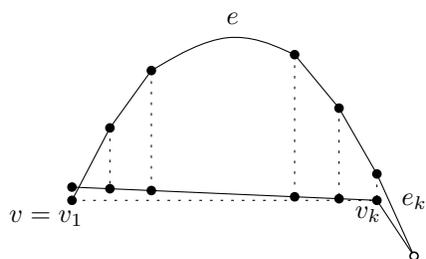


Figure 3.5: Once sequence T reaches vertex v , it represents a half-open path. Reordering the simplices so that the vertices are increasing from the open end of the path to v , and each edge appears immediately after its last face, we get all simplices in sequence T to be locally paired.

the simplices function values in the range $\bar{f}(v)$ and $\bar{f}(v) + \delta$ (for sufficiently small δ), so that all the simplices have distinct values that increase in the order of the above sequence $v_k, e_k, \dots, v_1, e_1$.

Function value change. It remains to verify that after canceling all persistence pairs below the given threshold ε , the change in the function value on any simplex does not exceed ε . The only questionable case arises when we cancel two pairs (v, e) and (v', e') that are neither nested, not disjoint as shown in Figure 3.6. If (v', e') has lower persistence than

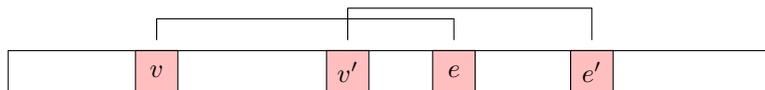


Figure 3.6: Overlapping pairs (v, e) and (v', e') in the filtration.

(v, e) , and therefore is canceled first. It is not immediately obvious what happens to the simplices that are lowered to precede edge e during the simplification of pair (v', e') . If they were to be lowered further during the simplification of pair (v, e) , they would aggregate the total function change. If the sum of the two persistences exceeded ε , so would the total function change. However, we observe that the vertices lowered to precede edge e cannot be faces of any edge that already came before e in the filtration. Therefore, from the above case analysis (specifically, Case 5a), it follows that for intersecting persistence intervals as in Figure 3.6, the function change on every simplex does not exceed ε .

3.4 Lower bound

In this section, we prove part B of the Simplification Theorem for 2-Manifolds stated in Section 3.1: for $p = 0, 1$ and all $\varepsilon > \delta > 0$ there exists a 2-manifold \mathbb{M} and a function $f : \mathbb{M} \rightarrow \mathbb{R}$ such that if $g : \mathbb{M} \rightarrow \mathbb{R}$ is a dimension p ε -simplification of f then $\|f - g\|_\infty > \varepsilon - \delta$. The topology of the 2-manifold is less important for the proof than the details of the function. We thus let \mathbb{M} be the 2-sphere and we choose f as the (vertical) height function of the embedding of \mathbb{M} displayed in Figure 3.7. There are three critical points with similar

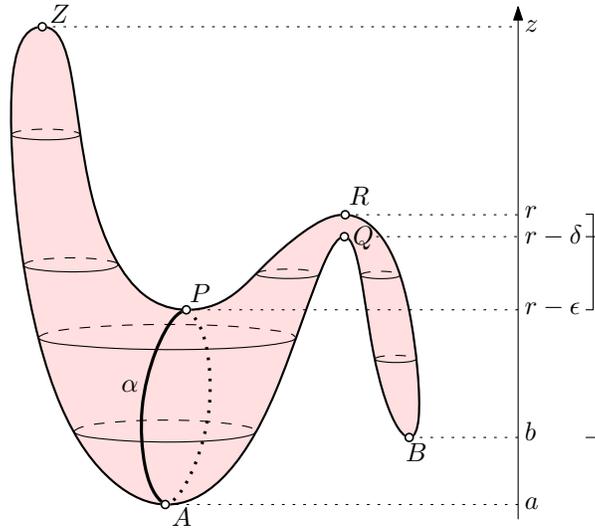


Figure 3.7: Embedding of the 2-sphere \mathbb{M} in \mathbb{R}^3 such that $f : \mathbb{M} \rightarrow \mathbb{R}$ is its height function. There are two minima, A and B , two saddles, P and Q , and two maxima, R and Z . The two ascending paths from A to P decompose \mathbb{M} into a left and a right hemisphere.

heights, $f(P) = r - \varepsilon$, $f(Q) = r - \delta$, $f(R) = r$, where $0 < \delta < \varepsilon$. The two minima have function values $f(A) = a < f(B) = b$ that are both much smaller than r , and the maximum has a function value $f(Z) = z$ that is much larger than r . The critical points are paired as (B, Q) , (P, R) , leaving A and Z unpaired. The off-diagonal points in the

persistence diagrams are therefore

$$\begin{aligned} \text{Dgm}_0(f) &: (a, \infty), (b, r - \delta); \\ \text{Dgm}_1(f) &: (r - \varepsilon, r); \\ \text{Dgm}_2(f) &: (z, \infty). \end{aligned}$$

All points have L_1 -distance larger than ε from the diagonal, except for $(r - \varepsilon, r)$ whose L_1 -distance from the diagonal is ε . To get a dimension 1 ε -simplification, we thus need to cancel P with R and leave the other critical points intact (or replace them by new critical points at the same height). It seems plausible that f does not have a dimension 1 ε -simplification that changes the function by $(\varepsilon - \delta')$ with $\delta < \delta' < \varepsilon$. Indeed, we cannot lower R by more than δ since it gets stuck at Q . Hence we need to raise P by at least $\varepsilon - \delta$. A more formal argument supporting this conclusion will be presented shortly. Since this works for arbitrarily small $\delta > 0$, this implies the claimed lower bound. To prove the same bound for $p = 0$ we use the construction upside-down, that is, we substitute $-f$ for f .

We now give the formal argument for the claim that the difference between f and its ε -simplification g is $\|f - g\|_\infty > \varepsilon - \delta$. To get a contradiction, we assume there is a dimension 1 ε -simplification $g : \mathbb{M} \rightarrow \mathbb{R}$ of f with $\|f - g\|_\infty = \varepsilon - \delta'$ for some $\delta' > \delta$. Let α be the cycle consisting of two monotonically increasing paths from A to P , as drawn in Figure 3.7. It decomposes the 2-sphere into a closed left hemisphere (containing Z) and a closed right hemisphere (containing B, Q, R). Consider the restrictions \bar{f} and \bar{g} of f and g to the right hemisphere. The diagram $\text{Dgm}_0(\bar{f})$ is the same as $\text{Dgm}_0(f)$. By the Stability Theorem, the diagram $\text{Dgm}_0(\bar{g})$ contains a point (b', q') at L_∞ -distance at most ε from $(b, r - \delta)$ in $\text{Dgm}_0(f)$. The value q' is that of a saddle Q' of \bar{g} . By definition of ε -simplification, we have $g(Q') = q' = r - \delta$, which is larger than $g(x) \leq f(x) + (\varepsilon - \delta') < r - \delta$ for any point x on α . This implies that Q' lies in the interior of the right hemisphere and is therefore also a saddle of g . Furthermore, there are no other finite off-diagonal points in the persistence diagrams of g . It follows that g has only one saddle, namely Q' . A similar argument implies that g has only one maximum, Z' , in the left hemisphere and that $g(Z') = z$. Since there is only one maximum and only one saddle, we can draw a path from Z' to Q' that monotonically decreases in g . This path crosses the cycle α . But the points x on α have $g(x) < r - \delta$ which is less than the values of Z' and Q' at the two ends. This contradicts the monotonicity of the path and implies $\|f - g\|_\infty > \varepsilon - \delta$, as required.

3.5 Higher Dimensions

An important question is the extension of the results of this chapter to higher dimensions. We know that there are 3-manifolds for any function on which there are (infinitely many) values of ε such that no ε -simplification of the functions exists for those values.

The proof of this insight can be credited to Henri Poincaré by way of Morse theory. When he originally stated his conjecture [70], Poincaré suggested that a compact 3-manifold with homology of a sphere is a 3-sphere. He discovered a counter-example to

that claim in what we now know as Poincaré homology sphere which has homology of a 3-sphere, but is not simply-connected, famously restating his conjecture to require the manifold to have a trivial fundamental group [71]. The problem remained open for a century until Grigori Perelman recently provided a proof [67, 68, 69]. See [63] for a historical account.

Poincaré homology sphere $\hat{\mathbb{S}}^3$ presents us with an example of a manifold on which not all ε -simplifications exist. Given any function $f : \hat{\mathbb{S}}^3 \rightarrow \mathbb{R}$, its 0- and 3-dimensional persistence diagrams each contain one point with death at infinity. These points reflect the homology of the sphere. The rest of the points have finite persistence, and we let ε_0 be the largest finite persistence of any point in any dimensional diagram. See Figure 3.8.

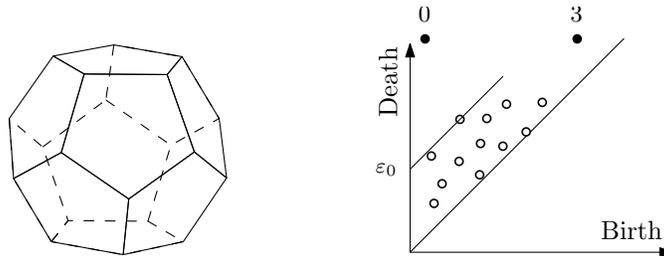


Figure 3.8: Left: Poincaré sphere can be obtained from the dodecahedron by gluing its opposite faces identified with the minimal clockwise twist. Right: persistence diagrams of any function on the Poincaré sphere consist of points at infinity in 0- and 3-dimensional diagrams, shown in black, and finitely persistent points in the diagrams of any dimension.

There does not exist an ε -simplification of f for any $\varepsilon \geq \varepsilon_0$. Indeed if such function existed, its persistence diagrams would consist of two points with infinite deaths, implying that the function had only two critical points. However, from Morse theory we know that if a (generic) function has only two critical points, its domain is a sphere [59, Theorem 3.6], a contradiction.

Similar argument shows that if one can show that an ε -simplification exists for any function on a compact, simply-connected 3-manifold, one would have a proof of the Poincaré conjecture. This speaks of the general difficulty of the question in higher dimensions. Sadly, we do not know how to show the implication in the opposite direction which prevents us from using Perelman’s proof of Poincaré conjecture.

The question that remains open is what happens if the domain is a sphere. In scientific visualization it is common for the data to be given on a 3-dimensional volume embedded in the Euclidean space \mathbb{R}^3 . Compactifying \mathbb{R}^3 we get a 3-sphere, making the open question of great interest in practice.

3.6 Discussion

The main contribution of this chapter is a constructive proof of the existence of ε -simplifications for continuous functions on 2-manifolds. The proof extends to 2-manifolds with boundary since we can convert those into 2-manifolds without boundary by gluing a disk to each boundary cycle. A curious aspect of our proof is that dimension 0 and dimension 1 homology can be simplified independently. Indeed, we can cancel all minimum-saddle pairs of persistence at most ε while leaving all saddle-maximum pairs intact, or vice versa. The only catch one must keep in mind is that not all simplices in a lower star are necessarily locally paired from the start. Therefore, when simplifying, for example, only 0-dimensional diagram, one still has to compute a “0-simplification” of the 1-dimensional diagram to get rid of such non-local pairs; and vice versa. It is also worthwhile to mention that the algorithm is combinatorial and we are free to assign function values that are consistent with the computed ordering of the vertices.

The fact that we can simplify the diagrams in order of increasing persistence allows us to incrementally compute a hierarchy of simplifications without having to start from scratch every time. Attali et al. [6] reexamine the algorithm of this chapter with computation of such hierarchy in mind. They realize that the algorithm affects only negative edges which by definition form a forest. They exploit this realization to get an algorithm which finds an ε -simplification on a 2-manifold in linear time, as opposed to quadratic time algorithm presented in this chapter.

While our motivating problem has been simplification of piece-wise linear functions, the algorithm that we obtain is more general. It works directly on a filtration and reorganizes it in a way that is agnostic to function values — only the order of the simplices matters. Can we use this procedure to reorder an alpha-shape filtration to obtain a simplicial complex with the homology suggested by the Homology Inference Theorem in Section 2.4 that is close to the given point set?

It is important not to underestimate the significance and difficulty of the reordering procedure illustrated in Figure 3.5 that concludes cancellation of a single pair. While trivial in two dimensions, lack of this procedure is the main obstacle to extending results of this chapter to higher dimensions; everything else that we do works in arbitrary dimensions. Its main purpose is to construct a generic function, which cannot be achieved by assigning the same value to all the simplices affected by a cancellation of a pair. It does so by reordering the filtration, so that the canceled pair leaves no combinatorial trace. This in turn allows us to cancel pairs considering each one independently. So far we miss this luxury in higher dimensions.

The simplification of continuous functions is a central problem in visualization. It may be used to clean up Morse-Smale complexes [38] and Reeb graphs [53, 72], which are powerful tools in the study and visualization of continuous data in scientific computing. This strengthens the importance of extending our results to three- and higher-dimensional spheres as suggested in the previous section.

Chapter 4

Persistence Vineyards

4.1 Vines and Vineyards

Motivation and results. If we take a function that changes continuously, from the Stability Theorem we know that its persistence diagrams also change continuously. We can watch them and understand the changes by observing how the points move and rearrange to form fleeting patterns. We may solidify the patterns by stacking up the diagrams, letting each point trace out a curve in space, which we refer to as a vine. This construction is a powerful metaphor aimed at gaining insight into continuous processes and quantifying some of their less tangible aspects. However, constructing the vine turns out to be more difficult than one may expect. Common time-series data is too sparse to compute them by matching the points in contiguous diagrams, and refining the series is expensive and sometimes not sufficiently powerful to remove all ambiguities. Based on the joint work with David Cohen-Steiner and Herbert Edelsbrunner [29], this chapter describes an alternative approach to constructing the vines by maintaining an ordering of the simplices during a homotopy.

The main contributions of this chapter to the theory and practice of persistent homology are:

1. an algorithm that maintains the persistence diagram in time $O(m)$ per transposition, where m is the number of simplices used to represent the topological space and the function;
2. a new and elementary proof of the stability of persistence diagrams;
3. the definition and computation of vineyards (continuous families of persistence diagrams) for time-series of continuous functions;
4. preliminary steps towards the application of vineyards to the study of protein folding trajectories;
5. an algorithm to compute the rank invariant of a bifiltration [18] in quadric time.

Similar to [75], our aim in the application is to learn about protein folding by viewing the process through a quantifiable combinatorial lens. The preliminary results are encouraging and will hopefully lead to a broader and deeper investigation of the subject. To illustrate the utility of the dynamic treatment of a filtration presented in Section 4.2 we give an elementary algorithm to compute the recently introduced rank invariant of a bifiltration [18] in time quadric in the number of simplices.

Stacking up persistence diagrams. We consider a homotopy $F(x, t) : \mathbb{X} \times [0, 1] \rightarrow \mathbb{R}$, and denote its snapshot at a given time-slice by $f_t(x) = F(x, t)$. Assuming every f_t is tame, we have a dimension p persistence diagram for every t and p , and the Stability Theorem relating the various diagrams. We draw $\text{Dgm}_p(f_t)$ in the (extended) plane $x_3 = t$ in $\bar{\mathbb{R}}^3$ thus getting a 1-parameter family of diagrams which we call the *dimension p vineyard* $\text{Vnr}_p(F)$. Each off-diagonal point in $\text{Dgm}_p(f_t)$ moves in time, tracing out a curve we refer to as a *vine*. Each vine is either open (starting and ending on the diagonal plane, $x_1 = x_2$), half-open, or closed (starting at an off-diagonal point in $x_3 = 0$ and ending at an off-diagonal point in $x_3 = 1$). If the homotopy is smooth then so are the vines, except when the pairing of critical values changes. We call such points *knees* and observe that they come in pairs. In practice, homotopies of functions arise from time-series data, given as a sequence of frames which are snapshots of the data at successive moments in time. Naturally, an assumption needs to be made about how the function changes in between the available frames. Ideally, such an assumption reflects the change in the underlying phenomenon described by the function, but in the absence of any such assumption it is convenient to use the straight-line homotopy between the frames.

Construction. We revisit the two types of functions that we discussed in Chapter 2: distance functions to point sets and piecewise linear functions.

A homotopy of the distance functions arises naturally if the underlying point set moves continuously. We recall that in this case the function $\check{c}(\sigma)$ that determines the order of simplices in the Čech filtration is equal to the radius of the smallest sphere enclosing simplex σ . It changes continuously to give us the homotopy $\check{c}_t(\sigma)$. Furthermore, if the points move at constant speeds in straight lines (for example, if we are interpolating between the snapshots of data), then the homotopy $\check{c}_t(\sigma)$ can be described by piecewise rational curves. The degrees of numerator and denominator polynomials are bounded by the dimension of the ambient space plus one, see [34] for explicit formulas for such radii. The continuous, but not necessarily smooth breaks occur between different pieces when the supporting point set of the enclosing sphere of the given simplex adds or loses a point.

A homotopy of piecewise linear functions arises when we interpolate between two snapshots f_0 and f_1 of the function. If vertices follow a straight-line homotopy $f_t(v) = \lambda f_1(v) + (1 - \lambda)f_0(v)$, then the order of the simplices in the lower star filtration is determined by function $f_t(\sigma) = \max_{v \in \sigma} f_t(v)$. The simplices exchange places when the order of their maximal vertices changes.

In both cases the value assigned to any given simplex follows a continuous rational function. The order of the simplices in the filtration changes when two such trajectory

functions cross. To construct the vineyards we first compute the persistence diagram of the initial filtration, which we then update through a sequence of transpositions, as explained in Section 4.2. We generate this sequence by sweeping the arrangement of the piecewise rational functions $P_\sigma : [0, 1] \rightarrow \mathbb{R}$ defined by $P_\sigma(t) = \check{c}_t(\sigma)$ and $P_\sigma(t) = f_t(\sigma)$ in the two cases respectively, as illustrated in Figure 4.1. This sweep can be viewed as performing

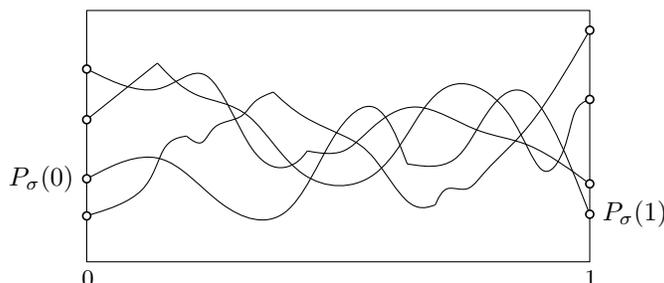


Figure 4.1: Sketch of the arrangement formed by piecewise rational curves representing trajectories of the simplices.

a kinetic sort of the simplices. See [8, 9] for discussion of the framework of kinetic data structures. To maintain persistence diagrams we are interested in updating $R = DV$ decomposition from Section 2.2 which gives us the pairing after two consecutive simplices transpose during the kinetic sort.

4.2 Updating the Pairing

In this section, we present the algorithm that updates the pairing function under a transposition of two simplices in the filtration. We begin with a characterization of the persistence pairing in terms of ranks of submatrices of the incidence matrix.

Uniqueness of pairing function. Let R be a reduced 0-1 matrix as defined in the Section 2.2, and write R_i^j for the lower left minor obtained by deleting the first $i - 1$ rows and the last $m - j$ columns. Any combination of non-zero columns of R_i^j has its last non-zero entry at the same height as the lowest non-zero entry of any of the involved columns. The combination can therefore not be zero implying that the combined non-zero columns are linearly independent. Recall from Section 2.2 that the algorithm in [39] can be interpreted as computing the reduced matrix $R = DV$, where V is invertible and upper-triangular. Since invertible upper-triangular matrices form a group, we can write D as the product RU of the reduced matrix R and the invertible upper-triangular matrix $U = V^{-1}$. When given choice, we prefer $D = RU$ decomposition over $R = DV$ since in applications matrix U is usually much sparser than matrix V . We call such a decomposition an *RU-decomposition* of D . In this decomposition, positive simplices correspond to zero columns and negative simplices to non-zero columns in R . Define

$$r_D(i, j) = \text{rank } D_i^j - \text{rank } D_{i+1}^j + \text{rank } D_{i+1}^{j-1} - \text{rank } D_i^{j-1}.$$

We prove below that the pairing function can be expressed in terms of r_D and is thus independent of the particular RU-decomposition used to define it.

PAIRING UNIQUENESS LEMMA. Letting $D = RU$, we have $\text{low}R[j] = i$ if and only if $r_D(i, j) = 1$. In particular, the pairing function does not depend on the matrix R in the RU-decomposition.

PROOF. Note that adding columns to columns located to their right does not change the rank of lower left minors, so $r_D = r_R$. To prove the claim, it is thus sufficient to show that $\text{low}R[j] = i$ iff $r_R(i, j) = 1$. First assume $\text{low}R[j] = i$. As argued above, the non-zero columns of R_i^j are linearly independent. The last column is non-zero, so $\text{rank} R_i^j - \text{rank} R_i^{j-1} = 1$. Now if we delete the top row from R_i^j then the last column is zero, implying $\text{rank} R_{i+1}^j - \text{rank} R_{i+1}^{j-1} = 0$, as required. Second assume $\text{low}R[j] \neq i$ and consider R_i^j and R_{i+1}^j . If $\text{low}R[j] < i$ the last columns in both matrices are zero and we have $\text{rank} R_i^j = \text{rank} R_i^{j-1}$ as well as $\text{rank} R_{i+1}^j = \text{rank} R_{i+1}^{j-1}$. If $\text{low}R[j] > i$ the last columns in both matrices are non-zero and we have $\text{rank} R_i^j = \text{rank} R_i^{j-1} + 1$ and $\text{rank} R_{i+1}^j = \text{rank} R_{i+1}^{j-1} + 1$. In both cases the claimed result follows. \square

Performing a transposition. To swap the simplices in positions i and $i + 1$, we exchange rows i and $i + 1$ as well as columns i and $i + 1$ in D . The new incidence matrix is therefore PDP , where P is the permutation matrix that swaps i and $i + 1$. To update the pairing function, we just need to repair the RU-decomposition, which we now show how to do in time $O(m)$. In Chapter 6 we will need to maintain $R = DV$ decomposition, so we also state explicitly how to update matrix V . Only one pair R and U , or R and V needs to be maintained in any given application. We have $PDP = PRUP = (PRP)(PUP)$, but this is not necessarily an RU-decomposition. Similarly, $PRP = (PDP)(PVP)$ may violate the conditions on matrices R and V . As illustrated in Figure 4.2, PRP is not reduced iff there are columns k and l with $\text{low}R[k] = i$, $\text{low}R[l] = i + 1$, and $R[i, l] = 1$. We note that since $U = V^{-1}$ and both are upper-triangular, $V[i, i + 1] = U[i, i + 1]$. Therefore, PUP and PVP are not upper-triangular iff $U[i, i + 1] = V[i, i + 1] = 1$. We may assume

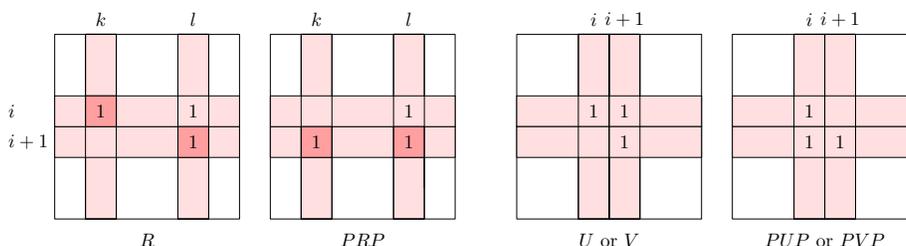


Figure 4.2: The transposition renders this particular R non-reduced and these particular U and V non-upper-triangular.

that the algorithm adds only columns that belong to simplices of the same dimension. The two cases illustrated in Figure 4.2 therefore arise only if the simplices at positions i and

$i + 1$ have the same dimension, which we thus assume. In all other cases, PRP is reduced and PUP and PVP are upper-triangular. In the following case analysis we denote by S_k^l the upper-triangular matrix such that multiplying by it on the right is equivalent to adding column k to column l . In turn multiplying by it on the left implies adding row l to row k . We observe that all such matrices S_k^l are idempotent, i.e. $S_k^l S_k^l = I$.

Case 1 Both i and $i + 1$ are positions of positive simplices. Since column i in R is zero we may set $U[i, i + 1] = V[i, i + 1] = 0$, if this is not the case. It follows that PUP and PVP are upper-triangular and we only need to consider PRP .

Case 1.1 There are columns k and l with $lowR[k] = i$, $lowR[l] = i + 1$, and $R[i, l] = 1$.

Case 1.1.1 $k < l$, as in Figure 4.2, left. To reduce PRP , we add column k to column l . Matrix S_k^l performs this operation, and we have $PDP = (PRPS_k^l)(S_k^l PUP)$. Similarly, $PRPS_k^l = (PDP)(PVP S_k^l)$. By construction, $PRPS_k^l$ is reduced. Furthermore, adding row l to row k preserves PUP as an upper-triangular matrix much like adding column k to column l preserves PVP as upper-triangular. It follows that this is an RU-decomposition of the new incidence matrix.

Case 1.1.2 $l < k$. To reduce PRP , we add column l to column k on its right. We have $PDP = (PRPS_l^k)(S_l^k PUP)$ as an RU-decomposition, same as Case 1.1.1, and, similarly, $PRPS_l^k = (PDP)(PVP S_l^k)$.

Case 1.2 There are no columns k and l as in Case 1.1. Then after transposition $PDP = (PRP)(PUP)$ is an RU-decomposition, and PVP also requires no update.

Case 2 Both i and $i + 1$ are positions of negative simplices. In this case, rows i and $i + 1$ cannot contain the lowest 1s of any columns. It follows that PRP is reduced and we only need to consider PUP and PVP .

Case 2.1 $U[i, i + 1] = V[i, i + 1] = 1$, as in Figure 4.2, right. To make the second matrix in the decomposition upper-triangular, we add row $i + 1$ of U to row i , equivalently column i of V to column $i + 1$. Matrix S_i^{i+1} performs this operation, we have $PDP = (PRS_i^{i+1}P)(PS_i^{i+1}UP)$, and $PRS_i^{i+1}P = (PDP)(PVS_i^{i+1}P)$. The effect of S_i^{i+1} on R is it adds column i to column $i + 1$.

Case 2.1.1 $lowR[i] < lowR[i + 1]$. Then RS_i^{i+1} as well as $PRS_i^{i+1}P$ are reduced and we have an RU-decomposition.

Case 2.1.2 $lowR[i + 1] < lowR[i]$. Then RS_i^{i+1} is not reduced, but we can reduce it by adding column $i + 1$ to column i . After the transposition, this is adding column i of $RS_i^{i+1}P$ to column $i + 1$ and we get

$$PDP = (PRS_i^{i+1}PS_i^{i+1})(S_i^{i+1}PS_i^{i+1}UP)$$

$$PRS_i^{i+1}PS_i^{i+1} = (PDP)(PVS_i^{i+1}PS_i^{i+1}).$$

The updated matrices U and V are upper-triangular and the updated matrix R is reduced, as illustrated in Figure 4.3, top row.

Case 2.2 $U[i, i + 1] = V[i, i + 1] = 0$. Then $PDP = (PRP)(PUP)$ is an RU-decomposition, and PVP also requires no update.

Case 3 i is the position of a negative simplex and $i + 1$ is the position of a positive simplex.

Case 3.1 $U[i, i + 1] = V[i, i + 1] = 1$, as in Figure 4.2, right. Just like in Case 2.1, we add row $i + 1$ of U to row i and get $PDP = (PRS_i^{i+1}P)(PS_i^{i+1}UP)$ and $PRS_i^{i+1}P = (PDP)(PVS_i^{i+1}P)$. The updated matrices U and V are upper-triangular. However, the updated matrix R is not reduced, and we reduce it by adding column i of $RS_i^{i+1}P$ to column $i + 1$, giving

$$PDP = (PRS_i^{i+1}PS_i^{i+1})(S_i^{i+1}PS_i^{i+1}UP)$$

$$PRS_i^{i+1}PS_i^{i+1} = (PDP)(PVS_i^{i+1}PS_i^{i+1})$$

as the final decompositions; see Figure 4.3, bottom row.

Case 3.2 $U[i, i + 1] = V[i, i + 1] = 0$. Then $PDP = (PRP)(PUP)$ is an RU-decomposition, and PVP also requires no update.

Case 4 i is the position of a positive simplex and $i + 1$ is the position of a negative simplex.

This is the reverse of Case 3.2. Indeed we set $U[i, i + 1] = V[i, i + 1] = 0$ if this is not the case and get the RU-decomposition $PDP = (PRP)(PUP)$.

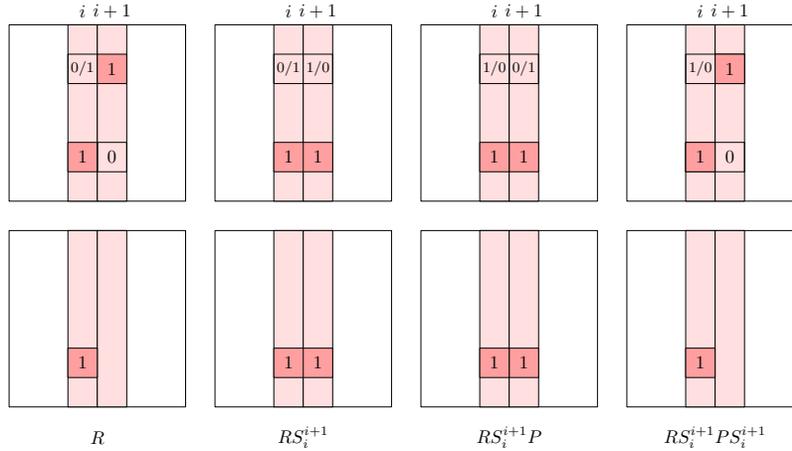


Figure 4.3: Evolutions of R to $RS_i^{i+1}PS_i^{i+1}$, in Case 2.1.2 on top and in Case 3.1 on bottom.

Changes in pairing. There are four cases in which the pairing function changes, all illustrated in Figure 4.4. We summarize them in the following theorem.

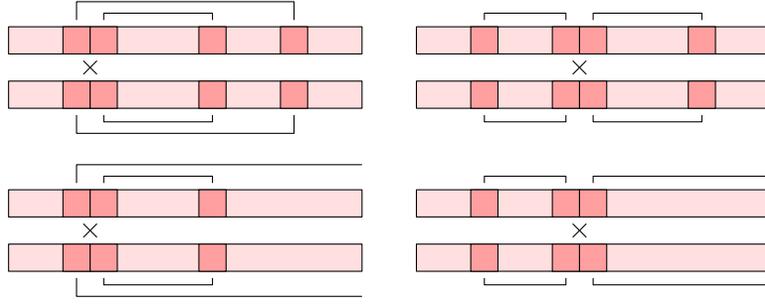


Figure 4.4: On the top left we see Case 1.1.2 and, if read backwards, Case 2.1.2. On the bottom left we see a particular setting in Case 1.2. On the right we see Case 3.1.

PAIRING CHANGE THEOREM. After the transposition of two consecutive simplices with indices i and $i + 1$, the pairing function may change only if the dimension of the simplices is the same. The pairing changes if and only if:

- Case 1a** $i = \text{low}R[k] < i + 1 = \text{low}R[l] < l < k, R[i, l] = 1$, in which two nested intervals swap their left endpoints to remain nested. This is Case 1.1.2 in the above analysis.
- Case 1b**¹ $i + 1 = \text{low}R[l], R[i, l] = 1$, and there is no k with $i = \text{low}R[k]$, in which an infinite and a nested finite intervals switch their left endpoints to remain nested. This is a possibility in Case 1.2 above.
- Case 2** $\text{low}R[i + 1] < \text{low}R[i] < i < i + 1, U[i, i + 1] = V[i, i + 1] = 1$, in which two nested intervals swap their right endpoints to remain nested. This is Case 2.1.2 above.
- Case 3** $R[i] \neq R[i + 1] = 0, U[i, i + 1] = V[i, i + 1] = 1$, in which two disjoint intervals swap their near endpoints to remain disjoint. This is Case 3.1 above.

In Cases 1a and 1b of the theorem the simplices responsible for the births of two vines in the same vineyard switch their roles giving us *Type 1* switch and two knees in the same vineyard. In Case 2 the simplices responsible for the deaths of two vines in the same vineyard switch their roles giving us *Type 2* switch and again two knees in the same vineyard. However, in Case 3 one of the switching simplices is responsible for the death of a vine, while the other one for the birth of a vine in the vineyards of contiguous dimension. We refer to this switch as *Type 3*, and observe that the two knees belong to the vineyards of contiguous dimension.

A trivial corollary of the above theorem, which we used in Chapter 3, is the following.

NESTED-DISJOINT COROLLARY. After transposition of two consecutive simplices of the same dimension, their persistence pairing may switch only if their persistence intervals are either nested or disjoint.

¹ The author would like to thank Samuel Hornus for pointing out that [29] forgets to highlight this case.

See Figure 4.4 for an illustration.

Running time. In every case, a transposition triggers at most one row exchange, one column exchange, and two column additions in the matrix R . Symmetrically, there are at most one column exchange, one row exchange, and two row additions in U or two column additions in V . This takes time at most linear in the number of simplices.

In all applications we have encountered, R , U and V are sparse and we can save time and storage using a sparse matrix implementation. We explain such a data structure for R consisting of two linear arrays, one for the set of columns and one for the set of rows, and a singly linked list for each column, as sketched in Figure 4.5. The j -th element of the column array points to the linked list of 1s in the column. The i -th element in the row array represents the i -th row in the original row sequence and stores its index in the current row sequence. We also store the reverse link, from the current row back to its corresponding original row. Each node in a linked list stores its index in the original row sequence, which we interpret as a pointer into the row array. To exchange two columns,

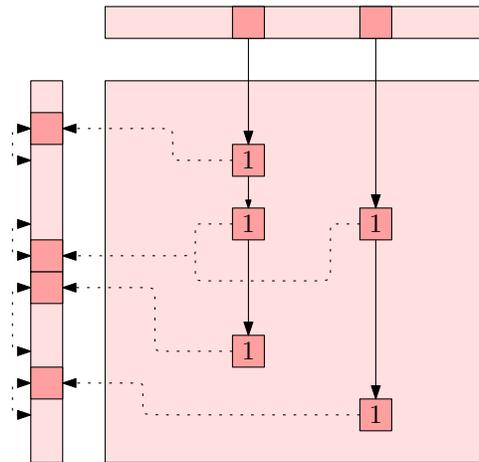


Figure 4.5: The sparse matrix representation of R sketched by showing the row and column arrays and the linked lists of two columns.

we swap their pointers (lists), which takes only constant time. To add a column to another, we merge the two lists, deleting nodes that come in duplicates, and substitute the result for the second column. This takes time proportional to the number of 1s in the two columns as long as the lists are consistently sorted. We achieve this by protecting the lists from row exchanges, keeping them sorted with respect to the original row indices. A row exchange thus only updates the correspondence between the original and the current orderings of the rows, which takes only constant time.

We use the same data structure for storing matrix V , and a symmetric sparse matrix data structure preferring rows over columns for U . The result is an implementation that takes storage proportional to m plus the number of 1s in R and in V or U . The amortized time for each operation is at most proportional to the number of 1s in the affected rows

and columns. The worst-case time is $O(m)$ per update, as before, but in our experiments the average update time is about constant.

4.3 Stability of Persistence Diagrams

In this section we give an elementary proof of the Stability Theorem.

Let $f, g : \mathbb{X} \rightarrow \mathbb{R}$ be continuous functions, K a simplicial complex, and Φ a homeomorphism from the underlying space of K to \mathbb{X} . The function $\bar{f} : K \rightarrow \mathbb{R}$ that maps each simplex $\sigma \in K$ to $\bar{f}(\sigma) = \max_{x \in \sigma} f(\Phi(x))$ is a piecewise constant approximation of f that allows us to construct a filtration on K suitable for the persistence algorithm in Section 2.2. Similarly, $\bar{g}(\sigma) = \max_{x \in \sigma} g(\Phi(x))$ is a piecewise constant approximation of g .

COMBINATORIAL STABILITY THEOREM. For functions $\bar{f}, \bar{g} : K \rightarrow \mathbb{R}$ and any dimension p , the bottleneck distance between the two dimension p persistence diagrams satisfies $d_B(\text{Dgm}_p(\bar{f}), \text{Dgm}_p(\bar{g})) \leq \|\bar{f} - \bar{g}\|_\infty$.

PROOF. Consider the straight-line homotopy $\bar{f}_t(\sigma) = (1 - t)\bar{f}(\sigma) + t\bar{g}(\sigma)$, and note that \bar{f}_t allows to construct a filtration for each $0 \leq t \leq 1$. Let t_1 to t_k be the values at which the ordering of the simplices changes by one or several transpositions, and set $t_0 = 0$ and $t_{k+1} = 1$. Let $t_i \leq r < s < t_{i+1}$, for any $0 \leq i \leq k$, and consider a pair of simplices, (σ, τ) , defined for the ordering that exists during the open time interval. Then $u_r = (\bar{f}_r(\sigma), \bar{f}_r(\tau))$ is a point in $\text{Dgm}_p(\bar{f}_r)$ and $u_s = (\bar{f}_s(\sigma), \bar{f}_s(\tau))$ is a point in $\text{Dgm}_p(\bar{f}_s)$. The Manhattan distance between the two points is the larger of the two coordinate differences, which implies

$$\begin{aligned} d_B(\text{Dgm}_p(\bar{f}_r), \text{Dgm}_p(\bar{f}_s)) &\leq \|\bar{f}_r - \bar{f}_s\|_\infty \\ &= (s - r) \cdot \|\bar{f} - \bar{g}\|_\infty. \end{aligned}$$

A transposition changes the pairing but it does not affect the persistence diagram. Hence,

$$\begin{aligned} d_B(\text{Dgm}_p(\bar{f}), \text{Dgm}_p(\bar{g})) &\leq \sum_{i=0}^k d_B(\text{Dgm}_p(\bar{f}_{t_i}), \text{Dgm}_p(\bar{f}_{t_{i+1}})) \\ &\leq \|\bar{f} - \bar{g}\|_\infty \sum_{i=0}^k (t_{i+1} - t_i). \end{aligned}$$

The latter sum is 1, which implies the claimed inequality. \square

As mentioned in Section 2.5, the persistence diagrams of the piecewise constant maps \bar{f}_t are the same as those of the piecewise linear maps defined by the same values at the vertices. The theorem thus implies the stability of persistence diagrams for the class of piecewise linear functions.

The original proof of Stability Theorem in [25] extends the proof of stability for piecewise linear functions to continuous functions defined on triangulable spaces. The extension works by taking arbitrarily fine triangulations of the domain to approximate the given function; it is described in the “Finale” paragraph of Section 3.3 in [25]. The argument applies directly to the Combinatorial Stability Theorem extending it to the usual Stability Theorem as stated in Section 2.2.

4.4 Applications

We illustrate the concepts introduced in this chapter by first computing a sequence of vineyards for the folding trajectory of a protein. They are generated by a piecewise linear function on a fixed triangulation, which is a common situation we encounter in applications. Then we show how to compute the rank invariant of a bifiltration in time $O(m^4)$.

4.4.1 Peptide Folding

Folding trajectories. The question of how proteins fold is a grand challenge in molecular biology and only modest progress has been reported in the last decades. It appears that the scientific community has not yet succeeded in simulating the folding process computationally. Exceptions are very short sequences or simulations over very short time intervals. We feel that vineyards can be useful in understanding the few folding trajectories that have been computed. One such trajectory describes the simulated folding motion of BBA5, a short peptide of $N = 23$ amino acids [73]. The trajectory is given as $n + 1 = 201$ frames covering a total of 40 picoseconds at regular intervals of 200 femtoseconds. For each $0 \leq i \leq n$, the i -th snapshot is a configuration of this backbone represented by a sequence S_i of N points in \mathbb{R}^3 , each the center of an alpha carbon along the backbone; see Figure 4.6. We turn the folding trajectory into vineyards using a 1-parameter family of functions described below.

Pairwise distance. Given a curve $b : [0, 1] \rightarrow \mathbb{R}^3$ in space, the pairwise distance function $[0, 1]^2 \rightarrow \mathbb{R}$ is defined by mapping (r, s) to $\|b(r) - b(s)\|$. Each function we consider is a piecewise linear approximation of such a pairwise distance function defined by the corresponding backbone configuration. We need some notation. Recall that S_i is the sequence of points describing the i -th backbone and let $c_{i,j}$ be the j -th point in S_i , for $1 \leq j \leq N$. Let K be the triangulation of $[1, N]^2$ obtained by connecting contiguous integer points along common horizontal, vertical, and 45-degree lines. It consists of N^2 vertices, $(3N - 1)(N - 1)$ edges, $2(N - 1)^2$ triangles, and therefore of $m < 6N^2$ simplices in total. For each $0 \leq i \leq n$, we construct $f_i : [1, N]^2 \rightarrow \mathbb{R}$ by defining $f_i(j, k) = \|c_{i,j} - c_{i,k}\|$ and extending the values at the vertices by linear interpolation over the edges and triangles. To form a homotopy from f_0 to f_m that passes through all intermediate functions, we finally define $f_{i+\lambda} = (1 - \lambda)f_i + \lambda f_{i+1}$, for all $0 \leq i < n$ and all $0 \leq \lambda \leq 1$.

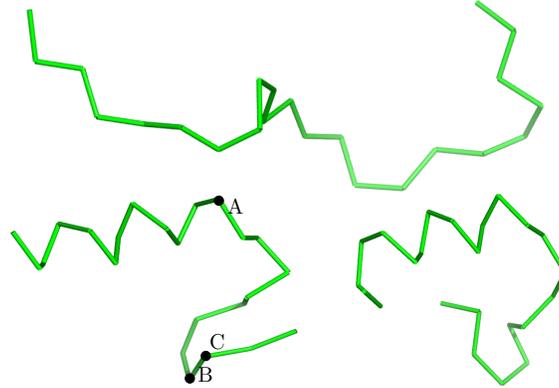


Figure 4.6: Top: snapshot 0, the initial backbone. Lower left: snapshot 68, the alpha helix is complete. Lower right: snapshot 200, the final backbone.

To construct the vineyards, we first compute the persistence diagrams of f_0 , which we then update through a sequence of transpositions, as explained in Section 4.2. We generate this sequence by sweeping the arrangement of polylines $P_{jk} : [0, n] \rightarrow \mathbb{R}$ defined by $P_{jk}(t) = f_t(j, k)$, as illustrated in Figure 4.7. We have N^2 polylines with at most n

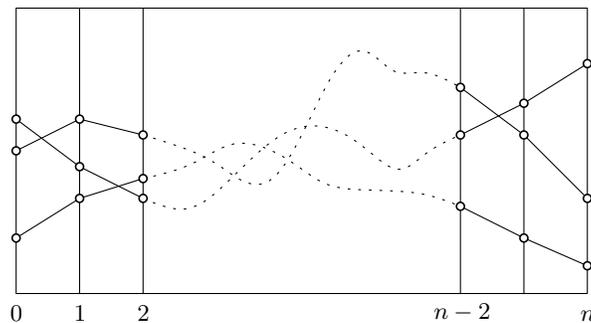


Figure 4.7: Sketch of the arrangement formed by the N^2 polylines representing the variation of function value at the vertices of K .

crossings between each pair. Each crossing corresponds to a transposition of two vertices. Using a standard plane-sweep algorithm, we can compute the ordered sequence in time $O(\log m)$ per crossing. The resulting algorithm takes worst-case time $O(nm^3)$ to construct the vineyards. In practice, the algorithm runs significantly faster, first because nm^2 is a gross over-estimate of the usual number of crossings, and second because our sparse-matrix implementation takes only about constant time per update.

Discussion of the vineyards. The results are illustrated in Figure 4.8, which shows the dimension 0 and 1 vineyards of the pairwise distance function. Each vine is drawn twice, as viewed from the front (normal to the diagonal direction) and from the side (along the diagonal direction). To interpret Figure 4.8, we fix a value $t \in [0, m]$ and consider a horizontal cross-section at height t . We note that two points (r, s) and (v, w) belong to a

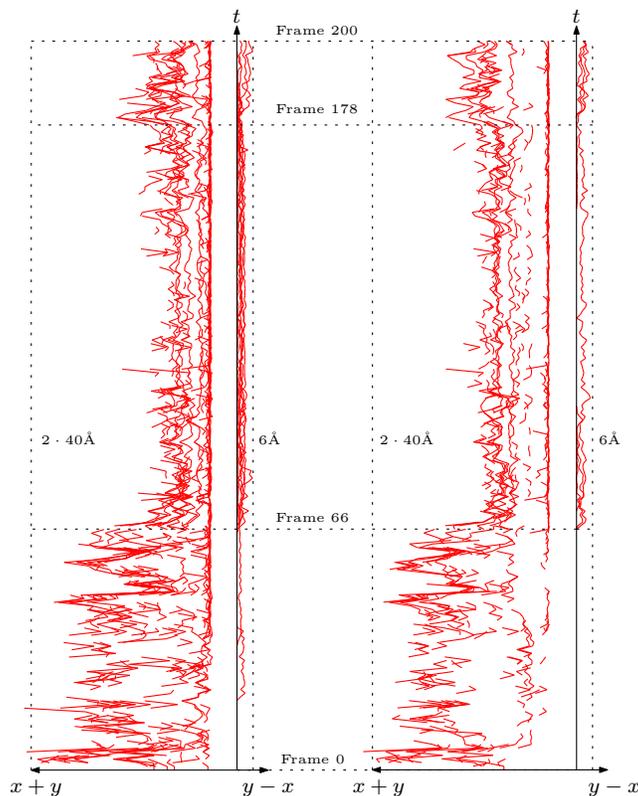


Figure 4.8: The front view $(x+y, t)$ and the side view $(y-x, t)$ of the dimension 0 vineyard on the left and the dimension 1 vineyard on the right. The side views are simplified by removing vines with lifetime less than 20 frames.

common component of the sublevel set $f_t^{-1}[0, \alpha]$ iff the component contains a path from the first point to the second. In other words, we can continuously move point $b(r)$ to $b(v)$ and simultaneously $b(s)$ to $b(w)$, both along the backbone b , such that the distance is less than or equal to α at all times. For $\alpha = 0$, the sublevel set consists of a single component, the diagonal of the domain. As we increase α , we see new components start at off-diagonal minima and components merge at saddles of f_t . The first critical points with non-zero value appear at α between 5 and 6 Å, causing the characteristic gap of about twice 5.5 Å to the time axis in the front views of the vineyards. The gap becomes particularly well defined when the alpha-helix is formed, suggesting the gap measures the distance between two alpha carbons separated by a single turn of the helix.

The dimension 1 vineyard is somewhat more difficult to interpret. It helps to break the folding process into three stages, the first from Frame 0 to 68, the second from 68 to 170, and the third from 170 to 200. The first stage is characterized by large and seemingly chaotic motions of the backbone that precipitate in vines across a relatively wide range of scales visible in the front view, both for the dimension 0 and 1 vineyards. At the end of the first stage, an alpha helix forms and the backbone assumes a rough S-shape, which remains until the end of the second stage. Covering almost the same time interval, we

see a dimension 1 vine emerging from the diagonal at Frame 66 and surviving until the Frame 178 when it disappears into the diagonal. There are 59 knees on this vine, and its maximum persistence (distance from the diagonal which is visible in the side view) is less than 6\AA and at times drops well below 1\AA . Nevertheless, the vine is very long-lived which suggests that even subtle configurations can stay around for a while. Let us take a closer look at this long vine representing a cycle created at a saddle and destroyed at a maximum. While the atom pairs responsible for the saddle and the maximum (AC and AB in Figure 4.6 for Frame 68) change as the vine evolves, they always span the tail of the backbone which remains intact during the second stage. Figure 4.9 shows the graph of the function together with a cycle in the homology class of the feature. During the third stage, we see the tail of the S-shape turn around and point back to the alpha helix. In the dimension 1 vineyard, we see three vines of persistence up to 5\AA emerge from the diagonal shortly after Frame 178 and survive until the end, at Frame 200.

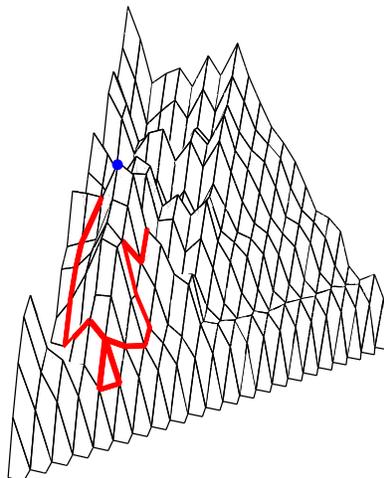


Figure 4.9: Pairwise distance function for Frame 68 of the BBA5 folding trajectory. The highlighted cycle is destroyed by the marked maximum and belongs to the homology class responsible for the long-lived vine.

In conclusion, we note that the folding process is very complex and it seems difficult to agree on when exactly events begin and end. This is in sharp contrast to a vine, which is unambiguously associated to a feature and has a precisely defined beginning and end. Furthermore, at any moment in time, the scale and the persistence of that feature are quantitatively expressed by the coordinates of the corresponding point on the vine. We thus believe that vines can be used to objectively and quantitatively encapsulate events in the process described by a homotopy.

4.4.2 Rank Invariant

Carlsson and Zomorodian [18] consider the invariants that arise for a multi-filtration of simplices. They propose to investigate the rank invariant which generalizes persistent homology; we briefly review it below. The main goal of this section is to give an elementary algorithm that uses the linear time update described in Section 4.2 to compute the rank invariant in quadric time.

Given a simplicial complex K , suppose each simplex in the complex is assigned injectively two indices $f(\sigma) = (i, j)$, each in range $[1, m]$, where m is the cardinality of K . The indices may for example come from ordering the simplices with respect to two different functions, and considering their sublevel sets. Furthermore, suppose the indices assigned to any face are pairwise smaller than the indices assigned to its coface; in other words, the simplices form a filtration with respect to each index. We define the sublevel set $K_{i,j} = \{\sigma \mid f(\sigma) = (i', j') \text{ and } i' \leq i, j' \leq j\}$. Considering the homology groups of the sublevel sets, we call the collection of ranks of their images induced by inclusion the *rank invariant* of the bifiltration. In symbols, $\text{RI}(f, g) = \{\text{rank } H(K_{i',j'}) \rightarrow H(K_{i,j}) \text{ for all } i' \leq i, j' \leq j\}$.

The collection of sublevel sets $K_{i,j}$ can be represented as an m -by- m matrix. To see how, we record each simplex σ in row i , column j of the matrix if $f(\sigma) = (i, j)$. Each sublevel set is the union of the simplices in the upper-left submatrix with the lower right corner (i, j) , see Figure 4.10. Ordering simplices according to any monotonically increasing path

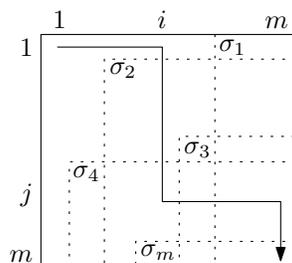


Figure 4.10: Representation of a bifiltration of a simplicial complex. Sublevel sets that contain a simplex are marked with dotted lines and have that simplex in the upper-left corner. (i, j) -path through the filtration is shown with a solid line.

from cell $(1, 1)$ through the matrix to the cell (m, m) gives us a filtration of simplices. Computing persistence of this filtration provides the ranks of maps induced by inclusion of any sublevel set into any other sublevel set on the path.

We consider all the paths starting at cell $(1, 1)$ that move i steps to the right, j steps down, then $(m - i)$ steps to the right, and $(m - j)$ steps down. We call such a path an (i, j) -path. We observe that from the (i, j) -path one can read ranks of all maps of the form $H(K_{i,*}) \rightarrow H(K_{*,j})$. We also note that the filtrations defined by paths (i, j) and $(i + 1, j)$, or symmetrically $(i - 1, j)$, $(i, j + 1)$, or $(i, j - 1)$, differ by a single simplex changing its position in the order. We can model such a move by at most m transpositions of

consecutive simplices, which in turn takes quadratic time to update the filtration. Since there are m^2 (i, j) -paths, we can iterate through them by changing the value of i or j by at most one per step, and the rank of the map between any pair of sublevel sets can be read off from the persistence pairing of such a path, we can compute the rank invariant in quadratic time.

The above construction naturally extends to a multi-filtration of k functions. Using the same procedure we can compute its rank invariant in time $O(m^{k+2})$.

4.5 Discussion

We conclude with a small number of questions aimed at improving and extending the results presented in this chapter.

- Are there variants of the update algorithm that are more efficient than the one described in Section 4.2 or that are simpler and just as efficient? For example, can we update the pairing by only maintaining the reduced matrix, R , and not the matrix U that does the reducing? Is there an advantage in treating rows and columns differently or is the symmetric version more efficient?
- In many applications, the points in the persistence diagram further away from the diagonal are more important than the points close to the diagonal. Can we use or adopt the update algorithm to compute the points with persistence beyond some threshold without spending time on the others? We desire an algorithm whose running time depends only on the size of the output it produces and not on the size of the entire diagram.
- Vineyards trace critical values and do not require any notion of critical points. However, when critical points are available, such as for smooth and for piecewise linear functions on manifolds [7, 62], we can use the update algorithm to maintain their association with the points in the persistence diagram. Can we exploit this ability to gain a better understanding of the stability or instability of critical points? In particular, can this ability be developed into a global alignment algorithm for shapes that is more general and more reliable than what is currently available [46, 55]?

Finally, we would like to suggest that vineyards should not be limited to homotopies but rather considered an analysis and visualization tool for parametrized families of functions. A point in case is the elevation function [2] whose maxima have been useful in coarse protein docking [79]. For a surface \mathbb{M} in space, this function is based on the sphere of height functions, which provide a homotopy $F : \mathbb{M} \times \mathbb{S}^2 \rightarrow \mathbb{R}$. The elevation function can be constructed from the \mathbb{S}^2 -parametrized vineyard of this homotopy. For a piece-wise linear surface with m simplices, elevation function can be computed using the update procedure of Section 4.2 in time $O(m^5)$.

In the next chapter we explore vineyards of a parametrized family of functions on a ball to infer local homology of a sampled space.

Chapter 5

Local Homology

5.1 Introduction

Much of the progress in today's experimental sciences is predicated on the ability to collect larger sets of more accurate measurements faster. Each data element is a tuple interpreted as a point in a space of the appropriate dimension. The resulting point set is often referred to as a *point cloud* so we are reminded that there is a lot of accumulated data. The main task is to detect patterns and to infer properties of the measured process from the structure of the point cloud.

Motivation. A common phenomenon in experimental measurements is that the data appears to describe a space whose intrinsic dimension is significantly smaller than that of the ambient space \mathbb{R}^n . This statement needs some clarification since every finite point set is, by definition, zero-dimensional. What we mean is that there is a relatively simple subset $X \subseteq \mathbb{R}^n$ of dimension $m \ll n$ such that all data points lie on or near this subset. The reason for this phenomenon is perhaps self-inflicted by our inability to make sense of processes that depend on a large number of independent parameters. The problem of reconstructing this subset (or one such subset from the class of possibilities) is often referred to as *manifold learning* [13]. The name betrays the tacit assumption that the subset is thought to have the topological structure of a manifold. In other words, it is locally homeomorphic to \mathbb{R}^m , or possibly to the m -dimensional half-space if we allow the manifold to have boundary. In case this assumption is grossly false it is suggested that these violations are artifacts of the mapping into the ambient space.

Similar to [51], we take a deliberate departure from the manifold assumption. While their methods are statistical in nature, we use local homology to recognize locations where the assumption is violated. Specifically, based on the joint work with Paul Bendich, David Cohen-Steiner, Herbert Edelsbrunner, and John Harer [11], we consider samplings of spaces that are partitioned into strata, each a manifold of dimension m or less, and we focus on the characterization of how these strata connect to each other. Stratified spaces can be described relatively compactly while significantly generalizing the class of spaces beyond manifolds. They include smooth images of manifolds into \mathbb{R}^n and permit different

local dimensions at different locations, but they do not include sets of fractal dimension. Since our goal is to deal with scientific data, the inability to describe fractal behavior might be considered a serious drawback. We argue otherwise, delegating the expression of fractal or chaotic behavior to the multi-scale description of the data. While this is not yet addressed in this thesis, we lay the foundations by parametrizing all our results in terms of scale parameters.

Results and prior work. Stratified spaces have been studied extensively in the mathematical literature [48, 81]. Particularly relevant to the line of work presented in this chapter is the development of intersection homology [47] and of persistence for intersection homology [12]. There is a striking paucity in computational studies of the reconstruction of stratified spaces from point samples which testifies for the technical difficulties caused by the presence of singularities. In general, the reconstruction of spaces from point samples is an important topic in a number of fields, each putting its own emphasis on the subject. Computer graphics and visualization stresses fast algorithms inspired by work in numerical analysis and image processing and focuses on data that describes surfaces in \mathbb{R}^3 [54, 80]. Computational geometry favors combinatorial algorithms based on Delaunay triangulations [33] and provides proofs of correctness under assumptions of dense sampling [3, 20]. Machine learning uses statistical methods to study high-dimensional data that describes relatively low-dimensional manifolds [10, 65]. Finally, topological data analysis relies on algebraic methods to reveal the topological structure of high-dimensional data [76]. In this chapter we combine characteristics found in computational geometry and in topological data analysis:

- we turn the algebraic concept of local homology into a multi-scale notion by constructing extended series of homology groups;
- we describe the $(\alpha|r)$ -vineyard, which we introduce as a practical tool for studying the local homology of a sampled space;
- we explain how the $(\alpha|r)$ -vineyard expresses the local homology of the sampled space at a point and prove the relation under the assumption of a dense sampling;
- we give an algorithm that computes the $(\alpha|r)$ -vineyard of a point cloud in time cubic in the number of simplices in the Delaunay triangulation.

Stratified Spaces. While the methods in this chapter apply to any compact subset of the Euclidean space, we are particularly interested in stratified spaces, which we introduce in the rest of this section. This interest mirrors the belief that this class of spaces is both computationally tractable and provides greater insight into data than its manifold interpretation alone.

Recall that a topological m -manifold is a space \mathbb{M} such that every point $z \in \mathbb{M}$ has a neighborhood homeomorphic to \mathbb{R}^m . If a space fails to be a manifold, it is because of the existence of singular points where no such neighborhood exists. For example, the figure-8 is not a 1-manifold; the singular crossing point has no neighborhood homeomorphic to \mathbb{R} .

On the other hand, every other point has such a neighborhood; in other words the figure-8 minus the crossing point is a 1-manifold and the crossing point itself is a 0-manifold. In general, a *stratification* of a topological space \mathbb{X} is a filtration by closed subsets,

$$\emptyset = \mathbb{X}_{-1} \subseteq \mathbb{X}_0 \subseteq \dots \subseteq \mathbb{X}_{m-1} \subseteq \mathbb{X}_m = \mathbb{X},$$

such that $\mathbb{X}_i - \mathbb{X}_{i-1}$ is a (possibly empty) i -manifold for each i . The set $\mathbb{X}_i - \mathbb{X}_{i-1}$ is called the i -*stratum* and its components are the *dimension i pieces* of \mathbb{X} .

Local structure. The above definition does not require that the points on a piece have similar neighborhoods outside the piece. Such requirements are usually added by extra conditions. Although there are many different conditions that might be added, each with its own subtleties [56], the following will do for our purpose. A stratified space $\mathbb{X} \subseteq \mathbb{R}^n$ with the above stratification is a *cs-space* if every point $x \in \mathbb{X}_i - \mathbb{X}_{i-1}$ has a neighborhood in \mathbb{X} homeomorphic to the product of an open i -ball in $\mathbb{X}_i - \mathbb{X}_{i-1}$ and the open cone on a compact topological space. The homeomorphism is assumed to take the product of the open i -ball and the cone point to the intersection of the neighborhood with \mathbb{X}_i . This is illustrated in Figure 5.1 where \mathbb{X} is a torus with one of the meridian circles pinched to a point and a disk stretched taut across its tunnel. If we remove the boundary circle of the

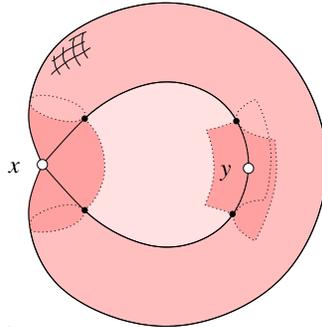


Figure 5.1: A 2-dimensional cs-space. The neighborhood of the pinch point x and of another point y along the boundary circle of the disk are highlighted. The two neighborhoods are not homeomorphic but they have the same homotopy type and therefore isomorphic homology groups.

disk, we disconnect \mathbb{X} and obtain a 2-manifold. The removed circle itself is a 1-manifold. However, the local structure is not uniform throughout the circle as the pinch point has a different neighborhood than any other point on the circle. Clearly, the pinch point is distinguishable from any other point on the circle.

For a cs-space, the cone in the definition of the neighborhood at a point x depends only on the piece that contains x . Since this piece is itself a manifold, the open balls are also homeomorphic. Hence the condition on the neighborhoods enforces exactly the requirement that each point in a piece has the same local structure in \mathbb{X} .

Connection with local homology. We note that the filtration in the definition of the stratified space is not unique. However, there is a natural coarsest filtration [50] which consists of the components in the partition of \mathbb{X} defined by calling points x and y equivalent if there exist neighborhoods of x and y and a homeomorphism between these neighborhoods that maps x to y . Any cs-space meets this condition already; for the coarsest filtration, we just impose the converse.

Now if two points x and y have such neighborhoods, then their local homology groups are also the same. It is the contrapositive of this statement that we hope to use in finding the best stratification of point cloud data. In this chapter we take the first step towards finding such stratification by considering how one could estimate the local homology of a sampled compact space. We pursue the omniscalar paradigm discussed in Chapter 1 by not forcing the user to choose the underlying parameters.

5.2 Local Homology and Distance Functions

In the rest of this chapter we use distance functions, introduced in Section 2.4, so we briefly review them as well as the classical notion of local homology.

Distance, filtrations, and diagrams. Let $d_Y : \mathbb{R}^n \rightarrow \mathbb{R}$ be the *distance function* defined by $d_Y(x) = \inf_{y \in Y} \|x - y\|$. For each real number α the corresponding *sublevel set* consists of all points at distance at most α , $Y_\alpha = d_Y^{-1}[0, \alpha]$; see Figure 5.2. We call

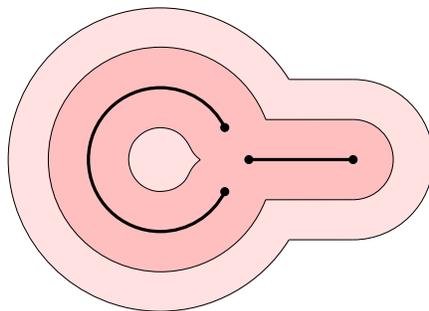


Figure 5.2: The space Y consists of the brace on the left and the stick to its right, both shown in bold. It is also the first sublevel set in the picture, $Y = d_Y^{-1}[0, 0]$. The second sublevel set merges the two components and creates a hole. The third sublevel set fills the hole.

α an *absolute homological regular value* if there is a sufficiently small $\varepsilon > 0$ such that the maps between homology groups induced by the inclusion $Y_{\alpha-\delta} \subseteq Y_{\alpha+\delta}$ form a series of isomorphisms for every $0 < \delta < \varepsilon$. Otherwise, α is an *absolute homological critical value*. We also define *superlevel sets* $Y^\alpha = d_Y^{-1}[\alpha, \infty)$, using them to define pairs (\mathbb{R}^n, Y^α) . We call α a *relative homological regular value* if there is a sufficiently small $\varepsilon > 0$ such that the maps between relative homology groups induced by the inclusion $(\mathbb{R}^n, Y^{\alpha+\delta}) \subseteq (\mathbb{R}^n, Y^{\alpha-\delta})$ form a series of isomorphisms for every $0 < \delta < \varepsilon$. Otherwise,

α is a *relative homological critical value*. As in Section 2.2 we call a function *tame* if it has finitely many (absolute and relative) homological critical values and its sublevel and superlevel sets have finite rank (absolute and relative) homology groups.

We consider the extended sequence of absolute and relative homology groups as introduced in [26],

$$\begin{aligned} 0 \rightarrow H(Y_\alpha) \rightarrow \dots \rightarrow H(\mathbb{R}^n) \\ \rightarrow H(\mathbb{R}^n, Y^\alpha) \rightarrow \dots \rightarrow 0, \end{aligned} \quad (5.1)$$

where α increases from 0 to infinity going up during the first half and then decreases from infinity to 0 coming down during the second half of the sequence. We denote the series of resulting persistence diagrams, $\text{Dgm}(d_Y)$. Following the notion of extended persistence [26], see Section 2.2, we distinguish between classes born and dying going up, classes born going up and dying coming down, and classes born and dying coming down. The corresponding three types are referred to as *ordinary*, *extended*, and *relative points* and they make up the *ordinary*, *extended*, and *relative subdiagrams* of the persistence diagram.

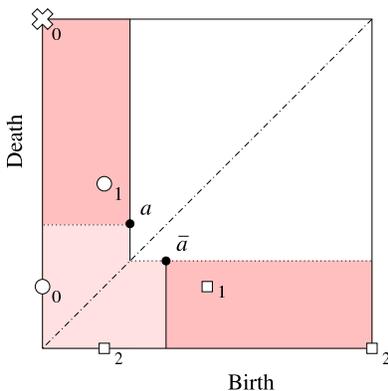


Figure 5.3: The persistence diagrams of the distance function in Figure 5.2. We draw time of birth from left to right and time of death from bottom to top, each ranging from 0 to ∞ . The white dots, crosses, and squares represent the ordinary, extended, and relative points in the diagrams of dimension given by the subscripts. The shading shows quadrants and half-planes defined by the points $a = (\xi, \zeta)$ and $\bar{a} = (\zeta, \xi)$.

To recover information from the diagrams we count points in subregions. Given $a = (\xi, \zeta)$ we count the classes that are alive during the entire interval, on the way up or the way down. Assuming $\xi \leq \zeta$ this number is

$$\begin{aligned} \#_p^a(d_Y) &= \text{rank}(\text{im}(H_p(Y_\xi) \rightarrow H_p(Y_\zeta))) \\ &+ \text{rank}(\text{im}(H_p(\mathbb{R}^n, Y^\zeta) \rightarrow H_p(\mathbb{R}^n, Y^\xi))) \\ &- \text{rank}(\text{im}(H_p(Y_\xi) \rightarrow H_p(\mathbb{R}^n, Y^\xi))). \end{aligned}$$

Of course we get a number for each dimension p and we write $\#^a(d_Y)$ for the series. To cover the other case, when $\zeta \leq \xi$, we set $\#^a(d_Y) = \#\bar{a}(d_Y)$ with $\bar{a} = (\zeta, \xi)$. As illustrated

in Figure 5.3, $\#^a$ counts the ordinary points in the upper left quadrant, $[0, \xi] \times (\zeta, \infty)$, the relative points in the lower right quadrant, $[\zeta, \infty) \times [0, \xi]$, as well as the extended points in the union of the left and lower half-planes, $[0, \xi] \times \mathbb{R} \cup \mathbb{R} \times [0, \xi]$. The resulting series for the example illustrated in Figures 5.2 and 5.3 is then $\#^a = (\dots, 1, 2, 1, \dots)$, where we only show the numbers for dimensions $p = 0, 1, 2$.

Local homology. Let $z \in \mathbb{R}^n$ be a point and $d_z : \mathbb{R}^n \rightarrow \mathbb{R}$ the distance function defined by $d_z(x) = \|x - z\|$. We write $B_r = d_z^{-1}[0, r]$ and $B^r = d_z^{-1}[r, \infty)$ for the sublevel and superlevel sets defined by r . Fix $\alpha \geq 0$ and consider Y_α , the space of points at distance at most α from Y . The sublevel sets and superlevel sets restricted to Y_α are $Y_\alpha \cap B_r$ and $Y_\alpha \cap B^r$. Traditionally, the *local homology groups* at a point z in a space Y_α are defined to be the relative homology groups $H(Y_\alpha, Y_\alpha - z)$ [64]. In words, a *local cycle* γ at z is a chain whose boundary misses z . Additionally, the boundary of γ must also miss some small open set $Y_\alpha \cap \text{int } B_r$ containing z , that is, γ belongs to $H(Y_\alpha, Y_\alpha - (Y_\alpha \cap \text{int } B_r)) = H(Y_\alpha, Y_\alpha \cap B^r)$. Now for any $s < r$, we have a map $H(Y_\alpha, Y_\alpha \cap B^r) \rightarrow H(Y_\alpha, Y_\alpha \cap B^s)$ induced by inclusion and excision. Our local cycle γ must lie in the image of this map for all possible choices of s . As a consequence, we see that the above definition of local homology at a point z is equivalent to the *direct limit*, $\lim_{r \rightarrow 0} H(Y_\alpha, Y_\alpha \cap B^r)$. To make this a multi-scale concept we consider again the extended sequence of homology groups,

$$\begin{aligned} 0 \rightarrow H(Y_\alpha \cap B_r) \rightarrow \dots \rightarrow H(Y_\alpha) \\ \rightarrow H(Y_\alpha, Y_\alpha \cap B^r) \rightarrow \dots \rightarrow 0, \end{aligned} \tag{5.2}$$

where r first increases from 0 to infinity and then decreases from infinity back to 0. As before, we record the evolution of homology classes using the thus defined series of persistence diagrams, $\text{Dgm}(d_z|Y_\alpha)$. The relative subdiagrams contain the information most directly relevant to estimating the local homology at z .

Discontinuity in α . The extended sequence of homology groups (5.2) provides a feasible approach to assessing local homology if the space, Y_α , is fixed. In the context of this chapter, we assume that the space has not been reconstructed, and we examine it at various scales by varying α . A drawback of the above construction is that the diagrams are not continuous in α . To see this let $0 < a' < a''$ be the distance thresholds of the three sublevel sets shown in Figure 5.2. Let z be the right endpoint of the stick in Y . For $\alpha = a'$ we have a one-dimensional homology class, γ , that is born going up and dies coming down with r . The class is represented by an off-diagonal point in the extended subdiagram of $\text{Dgm}_1(d_z|Y_{a'})$. In contrast, there are no one-dimensional classes for $\alpha = 0$ and for $\alpha = a''$. The class γ first appears when α reaches half the distance between the ends of the brace and the left endpoint of the stick. Right from the start, the representing point in the diagram is some distance away from the diagonal. Later, γ disappears when α reaches the radius of the brace and the representing point in the diagram merges into the diagonal.

5.3 The $(\alpha|r)$ -Vineyard

In this section, we introduce the main algebraic tool we use to study sampled stratified spaces.

Two filtrations. To obtain a continuous expression of the 2-parameter variation, we swap the order and vary α to construct the diagrams. Let $z \in \mathbb{R}^n$, as before, but now fix $r > 0$. Recall that $d_Y : \mathbb{R}^n \rightarrow \mathbb{R}$ is the distance function defined by $Y \subseteq \mathbb{R}^n$. The sublevel and superlevel sets of its restrictions to the ball of radius r around z are $Y_\alpha \cap B_r$ and $Y^\alpha \cap B_r$. Going first up with α from 0 to ∞ and then down from ∞ to 0 we get

$$\begin{aligned} 0 \rightarrow \mathrm{H}(Y_\alpha \cap B_r) &\rightarrow \dots \rightarrow \mathrm{H}(B_r) \\ &\rightarrow \mathrm{H}(B_r, Y^\alpha \cap B_r) \rightarrow \dots \rightarrow 0, \end{aligned} \quad (5.3)$$

and we write $\mathrm{Dgm}(d_Y|B_r)$ for the series of persistence diagrams that records the evolution of the homology classes in the sequence. The notion of local homology suggests we modify the filtration (5.3) and take the homology of Y_α within B_r relative to Y_α within the sphere ∂B_r . Constructing the extended sequence by first going up with α from 0 to ∞ and then down from ∞ to 0, we get

$$\begin{aligned} 0 \rightarrow \mathrm{H}(Y_\alpha \cap B_r, Y_\alpha \cap \partial B_r) &\rightarrow \dots \rightarrow \mathrm{H}(B_r, \partial B_r) \\ &\rightarrow \mathrm{H}(B_r, \partial B_r \cup (Y^\alpha \cap B_r)) \rightarrow \dots \rightarrow 0. \end{aligned} \quad (5.4)$$

The evolution of the homology classes is again recorded in the series of persistence diagrams, which we denote as $\mathrm{Dgm}(d_Y|(B_r, \partial B_r))$. In the rest of the chapter we need a mild assumption on Y , namely that the restrictions of its distance function to balls and to ball-sphere pairs are tame.

Equivalence of diagrams. We now show that the two diagrams contain the same information. Specifically, we establish isomorphisms between the homology groups in (5.3) and (5.4) and show that the corresponding pairings are dual and thus give the same diagrams. To shorten the notation and clarify the relations we set $X = Y_\alpha \cap B_r$, decompose its boundary $\partial X = F \cup G$ where $F = Y_\alpha \cap Y^\alpha \cap B_r$ and $G = \partial B_r \cap Y_\alpha$, and set $A = F \cap G$. Generically, X is an n -manifold with boundary, F and G are $(n-1)$ -manifolds with boundary, and A is an $(n-2)$ -manifold without boundary. Assuming tameness of Y , we use excision to rewrite (5.3) and (5.4), running them anti-parallel against each other:

$$\begin{array}{ccccc} \rightarrow & \mathrm{H}_{n-p}(X) & \rightarrow & \mathrm{H}_{n-p}(X, F) & \rightarrow \\ & \otimes & & \otimes & \\ \leftarrow & \mathrm{H}_p(X, \partial X) & \leftarrow & \mathrm{H}_p(X, G) & \leftarrow \\ & \downarrow & & \downarrow & \\ & \mathbb{Z}/2\mathbb{Z} & & \mathbb{Z}/2\mathbb{Z} & \end{array}$$

By Lefschetz duality, the first vertical pairing is perfect. This means that the paired groups are isomorphic and the persistence pairs in the first half of the two sequences are the same [26]. The other vertical pairing is also perfect, but it takes a little more effort to prove this.

ISOMORPHISM LEMMA. For every dimension p , the intersection pairing on X induces a perfect pairing

$$H_{n-p}(X, F) \otimes H_p(X, G) \rightarrow \mathbb{Z}/2\mathbb{Z}.$$

$$\begin{array}{ccccccccc} H_{p+1}(X, \partial X) & \rightarrow & H_p(F, A) & \rightarrow & H_p(X, G) & \rightarrow & H_p(X, \partial X) & \rightarrow & H_{p-1}(F, A) \\ \downarrow & & \downarrow & & \downarrow & & \downarrow & & \downarrow \\ H^{n-p-1}(X) & \rightarrow & H^{n-p-1}(F) & \rightarrow & H^{n-p}(X, F) & \rightarrow & H^{n-p}(X) & \rightarrow & H^{n-p}(F) \end{array}$$

Table 5.1: Commuting diagram with isomorphisms between the terms in the exact homology sequence of the triple $(X, \partial X, G)$ at the top and the exact cohomology sequence of the pair (X, F) on the bottom.

PROOF. First we notice that by excision, the relative homology groups $H_p(\partial X, G)$ and $H_p(F, A)$ are isomorphic. Next we consider the exact cohomology sequence of the pair (X, F) , shown in the bottom row in Table 5.1, and the exact homology sequence of the triple $(X, \partial X, G)$,

$$\begin{aligned} \rightarrow H_{p+1}(X, \partial X) &\rightarrow H_p(\partial X, G) \rightarrow H_p(X, G) \\ &\rightarrow H_p(X, \partial X) \rightarrow H_{p-1}(\partial X, G) \rightarrow \end{aligned}$$

Replacing $H_p(\partial X, G)$ by $H_p(F, A)$ we get the diagram in Table 5.1. Here each vertical arrow is the Poincaré-Lefschetz duality map defined by $\gamma \rightarrow f_\gamma$ where $f_\gamma(\delta) = \gamma \cdot \delta$, the intersection number between the two classes. It is not difficult to check that this diagram commutes. The two vertical maps on the left and the two vertical maps on the right are isomorphisms by Poincaré-Lefschetz duality. The Steenrod Five-Lemma then tells us that the center vertical map is also an isomorphism [64]. Finally we note that the Poincaré duality map $H_p(X, G) \rightarrow H^{n-p}(X, F)$ being an isomorphism implies that the intersection pairing is perfect. Indeed, every non-zero $\gamma \in H_p(X, G)$ has at least one $\delta \in H_p(X, G)$ such that $f_\gamma(\delta) = \gamma \cdot \delta \neq 0$. The claim follows. \square

We see that the pairings between the groups in (5.3) and (5.4) are perfect and all diagrams commute. It follows that if we use the superscript T to denote reflection across the diagonal we have $\text{Dgm}_p(d_Y|B_r) = \text{Dgm}_{n-p}^T(d_Y|(B_r, \partial B_r))$ for all dimensions p and all radii r .

Vineyard. An important result is that the diagrams of the filtrations in (5.3) and (5.4) vary continuously with r . It is convenient to show this for (5.3) again using the assumption of tameness. We let d_B be the supremum bottleneck distance between corresponding persistence diagrams in the series.

STABILITY LEMMA. Let $Y \subseteq \mathbb{R}^n$ and $z \in \mathbb{R}^n$ such that the restriction of $d_Y : \mathbb{R}^n \rightarrow \mathbb{R}$ to any ball and any ball-sphere pair centered at z is tame. Then the bottleneck distance between the series of persistence diagrams for two radii $r \leq r'$ is

$$d_B(\text{Dgm}(d_Y|B_r), \text{Dgm}(d_Y|B_{r'})) \leq r' - r.$$

PROOF. Letting $f, g : \mathbb{R}^n \rightarrow \mathbb{R}$ be defined by $f(x) = d_Y(z + rx)$ and $g(x) = d_Y(z + r'x)$, the restrictions of d_Y correspond to the restrictions of f and g to the unit ball. Changing coordinates does not affect the diagrams. Since d_Y is a distance function we have

$$\begin{aligned} |f(x) - g(x)| &= |d_Y(z + rx) - d_Y(z + r'x)| \\ &\leq \|x\| (r' - r). \end{aligned}$$

Since $\|x\| \leq 1$ the difference between the two functions is $\|f - g\|_\infty \leq r' - r$. The extension of the Stability Theorem in [25] to extended persistence as described in [26] implies the claim. \square

The stability of the persistence diagram suggests we vary r within $[0, \infty)$ and describe the homology in the neighborhood of $z \in \mathbb{R}^n$ by the resulting 1-parameter family of persistence diagrams. Stacking up the diagrams in \mathbb{R}^3 , as done in the previous chapter, using r as the third coordinate, each point sweeps out a curve which we refer to as a *vine*. Together the vines form a collection of curves which we refer to as the *vineyard* of the two distance functions; see Chapter 4. Specifically, we denote the vineyard obtained by stacking up the dimension p persistence diagrams by $V_{\text{nr},p}(d_Y|d_z)$ and the series of vineyards by $V_{\text{nr}}(d_Y|d_z)$. On occasion we call this the series of $(\alpha|r)$ -vineyards thus emphasizing that the diagrams are obtained by varying the threshold α for the distance to Y while fixing the threshold r for the distance to z , and the vines are obtained by varying r . This series of vineyards is the main technical ingredient in our approach to understanding sampled stratified spaces.

5.4 Local Homology Inference

In this section, we follow the techniques of Cohen-Steiner, Edelsbrunner, and Harer [25], described in Section 2.4, to prove that even with rather mild assumptions on the sampling of a space it is possible to infer its local homology. Perhaps more important than the guaranteed recognition is the interpretation of our result as describing the set of spaces that can possibly give rise to the sample.

Sample. The data we consider is a finite set of points, $U \subseteq \mathbb{R}^n$. It will be convenient to index the points in this set as u_i . We assume that U is sampled from or near a compact space $\mathbb{X} \subseteq \mathbb{R}^n$. For example, \mathbb{X} may be a compact stratified space but the existence of a stratification will play no role in what we prove in this section. It will, however, be important that the diagram of the restricted distance functions of \mathbb{X} be stable. We therefore assume that $d_{\mathbb{X}}|B_r$ is tame for every $z \in \mathbb{R}^n$ and every ball B_r centered at z . As mentioned earlier, this is a rather mild assumption whose violation requires infinitely many oscillations, like in the topologist's sine curve [64, p. 168], or similar phenomena. Recall that the space \mathbb{X} is unknown and the main question we ask is how much we can find out about \mathbb{X} under what assumptions relating U with \mathbb{X} .

Throughout this chapter we use a constant $\varepsilon > 0$ that quantifies the relation between \mathbb{X} and its sample. Suppose U is an ε -approximation of \mathbb{X} . We recall from Section 2.4 that by definition that Hausdorff distance between U and \mathbb{X} is at most ε , i.e. $\|d_U - d_{\mathbb{X}}\|_{\infty} \leq \varepsilon$.

Resolution. When we refer to the local homology at a point z , we consider the family of balls B_r centered at z and for each r we study the sequence of homology groups

$$\begin{aligned} 0 \rightarrow H(Y_{\alpha} \cap B_r) \rightarrow \dots \rightarrow H(B_r) \\ \rightarrow H(B_r, Y^{\alpha} \cap B_r) \rightarrow \dots \rightarrow 0, \end{aligned} \quad (5.5)$$

where Y is either U or \mathbb{X} ; see sequence (5.3) in Section 5.3. For each radius $r > 0$ we thus consider the series of persistence diagrams $\text{Dgm}(d_Y|_{B_r})$. The only non-trivial homology group of B_r is H_0 which has rank one. There is therefore only one extended point in this series tracking the first component that appears in the filtration. To determine the local homology of \mathbb{X} at a point z from the sample U it is necessary that the points sample all relevant features of the space fine enough to be recognized. To make this precise, we consider the homological critical values of the distance function of \mathbb{X} restricted to the ball.

DEFINITION. A radius r resolves \mathbb{X} at z to ε if the smallest positive absolute and relative homological critical values of $d_{\mathbb{X}}$ restricted to B_r exceed 4ε .

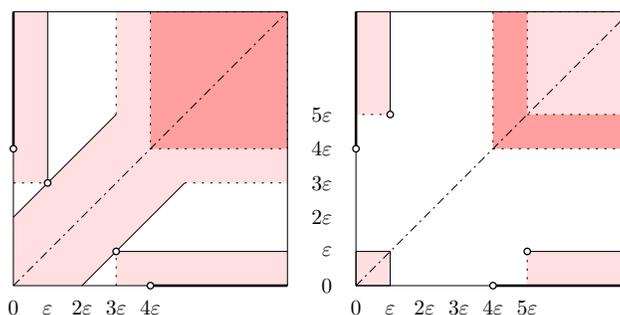


Figure 5.4: Left: the dark regions contain the persistence diagram of $d_{\mathbb{X}}$ for every radius $r \in R_{\mathbb{X}}(\varepsilon)$ and the light regions expand them to contain the persistence diagram of d_U for every radius $r \in R'_U(\varepsilon)$. Right: the light regions contain the persistence diagram of d_U for every radius $r \in R''_U(\varepsilon)$ and the dark regions contain the persistence diagrams of the distance function of U_{ε} .

For a radius r that resolves \mathbb{X} to ε there are no births and no deaths in the interval $(0, 4\varepsilon]$. In other words, the corridors separating the two boldface segments from the dark square in Figure 5.4, left, are empty. It follows that everything born at $\alpha = 0$ lives for a while and if it dies on the way up, as α increases, then it dies strictly after 4ε . Symmetrically, everything that dies at $\alpha = 0$ must have lived for a while and if it was born on the way down, as α decreases, then it was born strictly before 4ε . Radii that have this property are of special interest, so we define $R_{\mathbb{X}}(\varepsilon)$ as the set of radii r for which the points in

$\text{Dgm}(d_{\mathbb{X}}|B_r)$ all lie in the dark portion of Figure 5.4, left, which includes the vertical segment with lower endpoint $(0, 4\varepsilon)$, the horizontal segments with left endpoint $(4\varepsilon, 0)$, and the quadrant $(4\varepsilon, \infty)^2$.

Inference. By slightly extending the notation introduced earlier, we write $\#^a(d_{\mathbb{X}}|B_r)$ for the series of point counts of the corresponding diagrams in the region illustrated in Figure 5.3. For example, if a is the origin, then $\#^0(d_{\mathbb{X}}|B_r)$ counts the points on the horizontal Birth-axis and the vertical Death-axis.

LOCAL HOMOLOGY INFERENCE THEOREM. Let $\varepsilon > 0$, \mathbb{X} a compact space, U an ε -approximation of \mathbb{X} , and z a point in \mathbb{R}^n . Then $\#^0(d_{\mathbb{X}}|B_r) = \#^{(\varepsilon, 3\varepsilon)}(d_U|B_r)$ for every radius $r \in R_{\mathbb{X}}(\varepsilon)$.

PROOF. We will prove $R_{\mathbb{X}}(\varepsilon) \subseteq R'_U(\varepsilon)$, where the latter set consists of all radii r for which the points in $\text{Dgm}(d_U|B_r)$ all lie in the shaded portion of Figure 5.4, left, which expands the dark regions and the diagonal by ε in the vertical as well as the horizontal direction. We will see that this containment of sets implies the claimed equality.

Since $r \in R_{\mathbb{X}}(\varepsilon)$ we have $\#^0(d_{\mathbb{X}}|B_r) = \#^a(d_{\mathbb{X}}|B_r)$ for every $a \in [0, 4\varepsilon]^2$. Since $\|d_U - d_{\mathbb{X}}\|_{\infty} \leq \varepsilon$, the Stability Theorem of extended persistence implies a bijection such that each point in $\text{Dgm}(d_{\mathbb{X}}|B_r)$ lies within L_{∞} -distance ε from its corresponding point in $\text{Dgm}(d_U|B_r)$. This implies that all points of $\text{Dgm}(d_U|B_r)$ lie inside the ε -expanded region depicted in Figure 5.4, left. This region consists of three disjoint subregions, one expanding the vertical segment, one expanding the horizontal segment, and the third expanding the quadrant that contains the remaining points of $\text{Dgm}(d_{\mathbb{X}}|B_r)$ as well as the diagonal. By disjointness of the three subregions, the points of $\text{Dgm}(d_{\mathbb{X}}|B_r)$ in the two segments cannot map to any points other than the ones in the subregions that expand them. The points of $\text{Dgm}(d_U|B_r)$ in these two subregions are counted by $\#^a(d_U|B_r)$ with $a = (\varepsilon, 3\varepsilon)$. This implies the claimed equality. \square

Inverse. Recall that U is known but \mathbb{X} is not. The way we hope to use the Local Homology Inference Theorem is that we identify radii r for which the white corridors in Figure 5.4, left, are empty. For each such r there is a chance that it belongs to $R_{\mathbb{X}}(\varepsilon)$ and if it does we know the local homology of \mathbb{X} for this radius r . The trouble is that we can generally not be sure that r really belongs to $R_{\mathbb{X}}(\varepsilon)$. However, we can further restrict the regions that contain the points of $\text{Dgm}(d_U|B_r)$ so that they imply the existence of a space \mathbb{X} for which U is an ε -approximation and r is in $R_{\mathbb{X}}(\varepsilon)$. Let $R''_U(\varepsilon)$ be the set of radii r for which the points in $\text{Dgm}(d_U|B_r)$ are contained in the light shaded region in Figure 5.4, right.

INVERSE LHI THEOREM. Let $\varepsilon > 0$, U a subset of \mathbb{R}^n , and z a point in \mathbb{R}^n . Then there exists a compact space $\mathbb{X} \subseteq \mathbb{R}^n$ for which U is an ε -approximation and $R''_U(\varepsilon) \subseteq R_{\mathbb{X}}(\varepsilon)$.

PROOF. Set $\mathbb{X} = U_{\varepsilon}$ and note that U is an ε -approximation of \mathbb{X} . The distance function defined by \mathbb{X} is $d_{\mathbb{X}}(x) = \max\{0, d_U(x) - \varepsilon\}$. It follows that each birth and each death happens at 0 or ε earlier than before. The corresponding transformation of persistence

diagrams is a shift by ε down and a shift by ε to the left, except that a movement stops before the point enters the negative regions of birth or of death. If $r \in R_U''(\varepsilon)$ then all points in the diagrams of \mathbb{X} lie on the two segments and the quadrant that define $R_{\mathbb{X}}(\varepsilon)$. \square

Multi-scale example. Observe that the Local Homology Inference Theorem describes the relationship between the persistence diagrams of \mathbb{X} and of U for a fixed radius r . It is difficult to know ahead of time which value of r is most appropriate and in many situations it is not even desirable to make a choice. We cope with this difficulty by examining the persistent behavior across all radii. We use the example in Figure 5.5 to illustrate what we have in mind. Here \mathbb{X} is a one-dimensional space embedded in \mathbb{R}^2 . It consists of a string

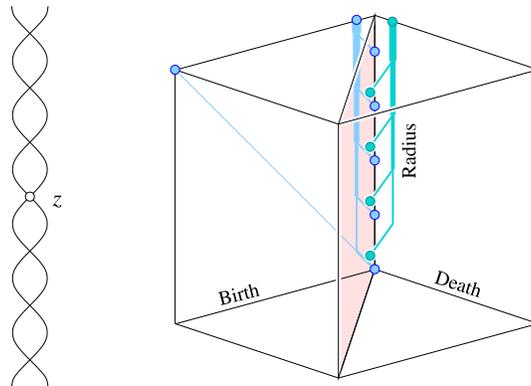


Figure 5.5: Left: the one-dimensional chain of loops, \mathbb{X} . Right: the dimension 1 $(\alpha|r)$ -vineyard of \mathbb{X} at z . The only significantly persistent vine runs roughly diagonally in the Birth-Radius plane and tracks a relative 1-cycle. All other vines run near the Radius-axis and track classes caused by the loops in the chain.

of loops, each connected to the loop before and the loop after. Its dimension 1 vineyard at the point z contains a prominent vine that has high persistence across all values of r . This vine tracks a dimension 1 relative homology class and corresponds to the chain itself which, from a distance, may be seen as a single curve. It can be detected even for rather sparse samples. Furthermore, the vineyard contains two small vines per loop, one tracking a relative and the other an absolute homology class. The relative class emerges at the moment the ball B_r first intersects the loop. It attains its largest persistence when B_r reaches the maximum near the center of the loop after which time the corresponding point in the diagram stops moving and sweeps out a vertical vine. At the same moment the absolute class emerges and attains its largest persistence when B_r reaches the other end of the loop after which time the corresponding point stops moving and sweeps out a vertical vine, as before.

For the study of local homology we are primarily interested in small values of r , that is, the lower portion of the vineyard. Of course, what small means is in the eye of a beholder.

On the other hand, the Local Homology Inference Theorem and its inverse can be used to make informed guesses. If the space \mathbb{X} in Figure 5.5 is sampled sufficiently densely, then small values of r resolve it, and we are able to detect the three dimension 1 cycles in the local homology of z . Specifically, there are three vines emerging from the origin, each tracking a relative homology class. If the sampling is not sufficiently dense then we cannot distinguish \mathbb{X} from a 1-manifold. Indeed, an arc passing through the vertices joining the loops could conceivably produce the same sample.

5.5 Power Cell Algorithm

In this section, we describe an algorithm for constructing the series of $(\alpha|r)$ -vineyards of a finite set of points as seen from a fixed point $z \in \mathbb{R}^n$. The algorithm is based on comparing various subcomplexes of the Delaunay triangulation of the finite set.

Voronoi decompositions. In Section 2.4 we introduced Voronoi decompositions and Delaunay triangulations. In this section we review and extend them to the case of weighted distance. Letting $u \in \mathbb{R}^n$ be a point with weight $w \in \mathbb{R}$, the *weighted square distance* of $x \in \mathbb{R}^n$ from u is $\pi_u(x) = \|x - u\|^2 - w$. For the common case in which the weight vanishes the weighted square distance is the squared Euclidean distance. Given a set of weighted points U , the (*weighted*) *Voronoi cell* of $u \in U$ is

$$V(u) = \{x \in \mathbb{R}^n \mid \pi_u(x) \leq \pi_v(x), v \in U\}.$$

For the time being we are interested in the case in which U is finite and all weights are zero. We index the points and use the shorter notation $V_i = V(u_i)$ for their Voronoi cells. Each V_i is the intersection of finitely many closed half-spaces and therefore a convex polyhedron. Collectively, the V_i cover the entire space thus forming the *Voronoi decomposition* of \mathbb{R}^n , which we denote as $\text{Vor}(U|\mathbb{R}^n)$; see Figure 5.6. We are also interested in the Voronoi decompositions of the sublevel sets inside the ball and on the sphere, which we denote as $\text{Vor}(U|U_\alpha \cap B_r)$ and $\text{Vor}(U|U_\alpha \cap \partial B_r)$. The former consists of cells $V_i \cap U_\alpha \cap B_r$, which are convex and generically either empty or n -dimensional. The latter consists of cells $V_i \cap U_\alpha \cap \partial B_r$, which are intersections of spherical caps and generically either empty or $(n - 1)$ -dimensional but not necessarily topologically simple. For example, in Figure 5.6 we see a Voronoi edge that intersects ∂B_r twice so that one of the two incident Voronoi cells intersects ∂B_r in two components. To cope with the resulting difficulties, we introduce the set $Z(\alpha)$ of points $x \in \mathbb{R}^n$ that satisfy $\|x - z\|^2 - r^2 \leq \|x - u_i\|^2 - \alpha^2$ for all $u_i \in U$. This is the Voronoi cell of z in $\text{Vor}(U \cup \{z\}|\mathbb{R}^n)$ in which every point has weight α^2 except for z which has weight r^2 . To distinguish it from the other Voronoi cells we refer to $Z(\alpha)$ as the *power cell* of z . More important than $Z(\alpha)$ itself is the complement of its interior, $Z_0(\alpha) = B_r - \text{int } Z(\alpha)$. We will see in Section 5.6 that $Z_0(\alpha)$ behaves topologically like $U_\alpha \cap \partial B_r$. A distinct advantage of the former set is that its intersections with the Voronoi cells are convex. Furthermore, it is not difficult to show that every point in $Z_0(\alpha)$ also belongs to U_α , that is, $Z_0(\alpha) \subseteq U_\alpha \cap B_r$.

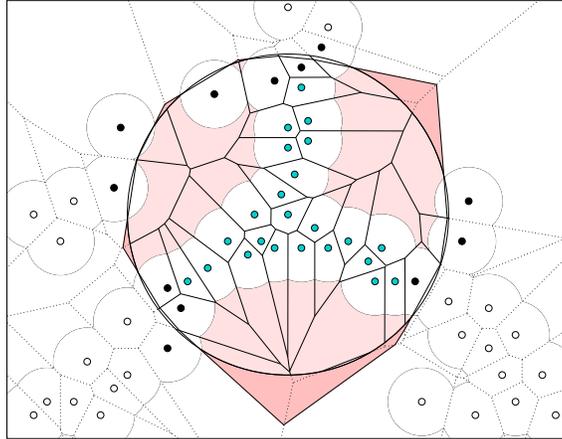


Figure 5.6: Voronoi decompositions of space, of the sublevel set of points at distance at most α from U , and of restricted versions of the same sublevel set. The Voronoi cells of the black dots contribute to the decompositions of $U_\alpha \cap \partial B_r$ and of $Z_0(\alpha)$, those of the gray and black dots contribute to the decomposition of $U_\alpha \cap B_r$, and the Voronoi cells of all dots contribute to the decompositions of U_α and of \mathbb{R}^n .

Delaunay triangulations. Computationally more convenient than the Voronoi decompositions are their dual *Delaunay triangulations*. For a subset $X \subseteq \mathbb{R}^n$ this is the set $\text{Del}(U|X)$ of simplices $\sigma \subseteq U$ for which X and $V_\sigma = \bigcap_{u_i \in \sigma} V_i$ have a non-empty intersection. In other words, $\text{Del}(U|X)$ is the nerve of the collection of sets $X \cap V_i$. For $X = \mathbb{R}^n$ we get the usual notion of Delaunay triangulation and for $X \subset \mathbb{R}^n$ we get the restricted Delaunay triangulation as defined in [42]. Generically, $\text{Del}(U|\mathbb{R}^n)$ is a simplicial complex geometrically realized in \mathbb{R}^n ; see Figure 5.7. We are also interested in the restric-

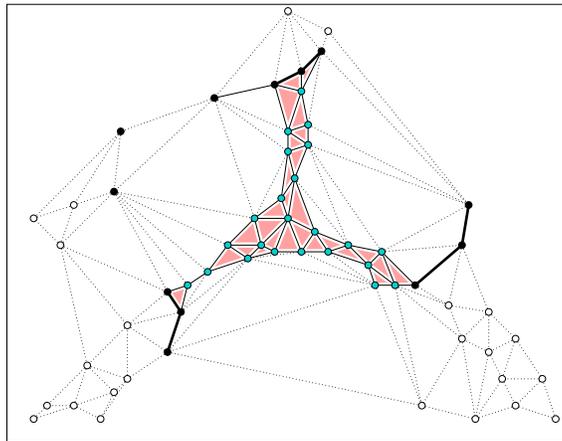


Figure 5.7: Delaunay triangulations dual to the Voronoi decompositions of the rectangular window in Figure 5.6 as well as of $U_\alpha \cap B_r$ and of $Z_0(\alpha)$. The drawing style identifies which vertices, edges, and triangles belong to which Delaunay triangulations.

tions to $U_\alpha \cap B_r$ and to $Z_0(\alpha)$. In these cases the Delaunay triangulations depend on α

and we write $K(\alpha) = \text{Del}(U|U_\alpha \cap B_r)$ and $K_0(\alpha) = \text{Del}(U|Z_0(\alpha))$. Since the restricting domains are subsets of each other, the three Delaunay triangulations are subcomplexes of each other, namely $K_0(\alpha) \subseteq K(\alpha) \subseteq \text{Del}(U|\mathbb{R}^n)$; see Figure 5.7.

Computing persistence. We now discuss the construction of the series of persistence diagrams of the distance function d_U restricted to the ball of fixed radius r around z . Specifically, we compute the diagram that describes the evolution of the homology classes in the sequence (5.3). Alternatively, we could compute the diagram of the sequence (5.4), which by the Isomorphism Lemma contains the same information as (5.3) but read backwards. We do neither and instead compute the diagrams from the respective first halves since this avoids the need to subdivide the Delaunay triangulation and leads to a simpler and more efficient implementation. Indeed, we substitute the sequence of complexes $K(\alpha)$, for $0 \leq \alpha < \infty$, and the homomorphisms induced by inclusion for the first half of (5.3). A formal proof that this substitution does not affect the persistence diagrams will be given in Section 5.6. To do the actual computation, we construct a filtration of the simplices in $K(\infty) = \text{Del}(U|B_r)$, and compute its pairing. Details of this construction are described in Section 5.7.

Constructing vineyards. Recall that the series of vineyards, $\text{Vnrd}(d_U|d_z)$, may be identified with the 1-parameter family of persistence diagrams, $\text{Dgm}(d_U|B_r)$, for $0 \leq r < \infty$. This is also how we construct it, by maintaining the persistence diagrams while growing r from zero to infinity. Indeed, we just need to maintain the two orderings of the simplices and update the persistence pairing whenever these orderings change. For each simplex $\sigma \in \text{Del}(U|\mathbb{R}^n)$ there are functions in r^2 that characterize when σ belongs to $K(\alpha)$ and when to $K_0(\alpha)$. Importantly, these functions are mostly continuous so that the maintenance of the orderings reduces to transpositions of contiguous simplices; see Section 5.7. Furthermore, each function is piecewise algebraic, where the number of pieces and the degrees are bounded from above by some constant. It follows that the graphs of any two functions cross at most some constant number of times. The total number of transpositions is therefore in $O(m^2)$. Finally, the persistence pairing can be maintained in time $O(m)$ per transposition as described in the previous chapter. It follows that the entire algorithm takes time at most cubic in the number of simplices.

$(\alpha|r)$ -VINEYARD THEOREM. Given the Delaunay triangulation of U in \mathbb{R}^n and a point $z \in \mathbb{R}^n$, the series of $(\alpha|r)$ -vineyards of the two distance functions, $\text{Vnrd}(d_U|d_z)$, can be constructed in time $O(m^3)$, where m is the number of simplices in $\text{Del}(U|\mathbb{R}^n)$.

5.6 Correctness

In this section, we prove that substituting the complexes $K_0(\alpha) \subseteq K(\alpha)$ for the spaces $U_\alpha \cap \partial B_r \subseteq U_\alpha \cap B_r$ does not affect the persistence diagrams. We do this in two steps, first constructing homotopy equivalences between spaces and second embedding

the induced isomorphisms between the corresponding homology groups in commuting diagrams.

Homotopy equivalence of pairs. Beyond homotopy equivalences between spaces we need them between pairs of spaces. Specifically, (X, X_0) is *homotopy equivalent* to (Y, Y_0) , denoted as $(X, X_0) \simeq (Y, Y_0)$, if there exist maps of pairs in both directions whose compositions are homotopic to the respective identities [60, p. 27]. We begin by establishing a homotopy equivalence between the pairs $(U_\alpha \cap B_r, U_\alpha \cap \partial B_r)$ and $(U_\alpha \cap B_r, Z_0(\alpha))$.

POWER CELL LEMMA. Let U, α, z, r be such that $B_r - Z_0(\alpha) \neq \emptyset$. Then the identity on $U_\alpha \cap B_r$ is a homotopy equivalence of $(U_\alpha \cap B_r, U_\alpha \cap \partial B_r)$ and $(U_\alpha \cap B_r, Z_0(\alpha))$ as a map of pairs.

PROOF. It suffices to show that the restriction of the identity, $i : U_\alpha \cap \partial B_r \rightarrow Z_0(\alpha)$, is a homotopy equivalence. Let y be a point in $B_r - Z_0(\alpha)$. Every point x in $Z_0(\alpha)$ belongs to B_r but not to the interior of $Z(\alpha)$. The weighted square distance of x from z is therefore non-positive and not smaller than the smallest weighted square distance to a point in U . Hence $\|x - u_i\|^2 - \alpha^2 \leq 0$ for at least one $u_i \in U$ which implies $Z_0(\alpha) \subseteq U_\alpha$. Now draw the ray that starts at y and passes through x and let x' be the point where it crosses ∂B_r . We map x to x' and thus define a retraction $j : Z_0(\alpha) \rightarrow U_\alpha \cap \partial B_r$. The composition $i \circ j$ is the identity on $U_\alpha \cap \partial B_r$. The other composition, $j \circ i$, is homotopic to the identity of $Z_0(\alpha)$, as established by the straight-line homotopy $\lambda : Z_0(\alpha) \times [0, 1] \rightarrow Z_0(\alpha)$ defined by $\lambda(x, t) = (1 - t)x + tx'$. This implies that the identity is a homotopy equivalence as a map of pairs, as claimed. \square

We note that when $B_r \cap Z(\alpha) = \emptyset$ then there is no homotopy equivalence between the pairs. Indeed, we then have $B_r \subseteq U_\alpha$ so that $(U_\alpha \cap B_r) - (U_\alpha \cap \partial B_r) = B_r - \partial B_r$ is an open ball while $(U_\alpha \cap B_r) - Z_0(\alpha) = \emptyset$.

We use the Nerve Subdivision Lemma from Section 2.3 in our context, and let C be the collection of cells $V_i \cap U_\alpha \cap B_r$ in the Voronoi decomposition of the restricted sublevel set. Recall that $K(\alpha)$ is the nerve of this collection of sets. Next we construct a map $h_\alpha : |\text{Sd } K(\alpha)| \rightarrow U_\alpha \cap B_r$ by specifying it at the vertices and extending it by piecewise linear interpolation. Recall that V_σ is the intersection of the Voronoi cells of all vertices of σ . To define the map we set

$$h_\alpha(\hat{\sigma}) = \arg \min_{x \in V_\sigma \cap U_\alpha \cap B_r} d_U^2(x) - d_z^2(x).$$

By construction, $h_\alpha(\hat{\sigma})$ belongs to the intersection of the cells that correspond to the vertices of σ . We can therefore apply the Nerve Subdivision Lemma and conclude that h_α is a homotopy equivalence. We are also interested in the restriction of h_α to the barycentric subdivision of $K_0(\alpha)$. Recall that $\sigma \in K(\alpha)$ belongs to $K_0(\alpha)$ iff $V_\sigma \cap Z_0(\alpha)$ is non-empty. By construction, the point $h_\alpha(\hat{\sigma})$ then lies in this intersection. The restriction of h_α is therefore a map $h'_\alpha : |\text{Sd } K_0(\alpha)| \rightarrow Z_0(\alpha)$ that again satisfies the assumptions of the Nerve Subdivision Lemma. Hence, h_α is a homotopy equivalence as a map of pairs.

Commuting diagrams. The Power Cell and Nerve Subdivision Lemmas imply that the series $H(U_\alpha \cap B_r, U_\alpha \cap \partial B_r)$, $H(U_\alpha \cap B_r, Z_0(\alpha))$, and $H(K(\alpha), K_0(\alpha))$ are isomorphic. However, to use Persistence Equivalence Theorem from Section 2.2 and conclude that the corresponding sequences of homology groups give rise to the same persistence diagrams we need more, namely that the groups form a commuting diagram whose vertical maps are isomorphisms. We draw the diagram of spaces and maps between them from which the commuting diagram can be obtained by application of the homology functor:

$$\begin{array}{ccc}
(U_\alpha \cap B_r, U_\alpha \cap \partial B_r) & \xrightarrow{j_\alpha^{\alpha'}} & (U_{\alpha'} \cap B_r, U_{\alpha'} \cap \partial B_r) \\
\downarrow i_\alpha & & \downarrow i_{\alpha'} \\
(U_\alpha \cap B_r, Z_0(\alpha)) & \xrightarrow{j_\alpha^{\alpha'}} & (U_{\alpha'} \cap B_r, Z_0(\alpha')) \\
\uparrow h_\alpha & & \uparrow h_{\alpha'} \\
(K(\alpha), K_0(\alpha)) & \xrightarrow{g_\alpha^{\alpha'}} & (K(\alpha'), K_0(\alpha')),
\end{array}$$

where $\alpha \leq \alpha'$. By the Power Cell and Nerve Subdivision Lemmas, the vertical maps induce isomorphisms between the homology groups of the spaces. The maps i and j are inclusions which implies that the upper square of the corresponding diagram of homology groups commutes. To prove the same for the lower square we consider the maps $e = j_\alpha^{\alpha'} \circ h_\alpha$ and $e' = h_{\alpha'} \circ g_\alpha^{\alpha'}$. Consider the map $H : |K(\alpha)| \times [0, 1] \rightarrow U_\alpha \cap B_r$ defined by $H(x, t) = h_{\alpha_t} \circ g_\alpha^{\alpha_t}(x)$, where $\alpha_t = (1-t)\alpha + t\alpha'$. 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, $e_* = e'_*$.

To summarize, we have isomorphisms connecting the groups in the columns of a commuting diagram. It follows that each of the three rows gives rise to the same series of persistence diagrams. In other words, our algorithm which computes persistence diagrams using the complexes $K(\alpha)$ and $K_0(\alpha)$ is correct.

5.7 Algorithm Details

In this section we go through the details of constructing and maintaining the filtrations which we use to compute the persistence diagrams and vineyards described above.

5.7.1 Thresholds

Given a point $z \in \mathbb{R}^n$ and a radius $r \geq 0$, we use the restriction of the Delaunay triangulation to the ball B_r and the pair $(B_r, \partial B_r)$, both centered at z , to assess the local homology of the data at z . Specifically, we consider the complexes

$$\begin{aligned}
K(\alpha) &= \text{Del}(U|U_\alpha \cap B_r); \\
K_0(\alpha) &= \text{Del}(U|Z_0(\alpha)),
\end{aligned}$$

where we recall that $Z_0(\alpha) = B_r - \text{int } Z(\alpha)$ is contained in U_α . In this section, we study under what conditions these complexes contain a simplex in $\text{Del}(U|\mathbb{R}^n)$.

Threshold for A . A necessary condition for a simplex $\sigma \in \text{Del}(U|\mathbb{R}^n)$ to belong to $K(\alpha)$ is it belongs to the Delaunay triangulation of the sublevel set,

$$A(\alpha) = \text{Del}(U|U_\alpha).$$

This complex is also known as the alpha complex of U [41]. To characterize when σ belongs to $A(\alpha)$ we consider the smallest $(n-1)$ -sphere that passes through the points of $\sigma \subseteq U$ and encloses none of the points of U . We call this the *smallest empty circumsphere* of σ and let $c_0 \in \mathbb{R}^n$ be its center and ϱ_0 its radius. If $\sigma = \{u_i\}$ is a vertex then the $(n-1)$ -sphere degenerates to the point $c_0 = u_i$ and we have $\varrho_0 = 0$. Generically, c_0 and ϱ_0 are unique and we set $\alpha_\sigma = \varrho_0$.

CONDITION A. Let σ be a simplex in $\text{Del}(U|\mathbb{R}^n)$. Then

$$\sigma \in A(\alpha) \quad \text{iff} \quad \alpha_\sigma^2 \leq \alpha^2.$$

There is an alternative geometric interpretation of this condition in terms of the smallest but possibly non-empty circumsphere of σ . It passes through the points of σ but may enclose other points of U . Unless $\dim \sigma = n$, its center c and radius ϱ are not necessarily the same as c_0 and ϱ_0 . The point c is also the center of the common intersection of the balls of radius α centered at the vertices of σ and the plane of dimension $n - \dim \sigma$ whose points are equidistant from these vertices. If this common intersection is non-empty then it is either a point or a ball of dimension $n - \dim \sigma$ and square radius $\alpha^2 - \varrho^2$. The interpretation of Condition A is now that α_σ^2 is the smallest value of α^2 for which this ball has a non-empty intersection with $V_\sigma = \bigcap_{u_i \in \sigma} V_i$. In other words, $\sigma \in A(\alpha)$ iff $U_\alpha \cap V_\sigma \neq \emptyset$ which is consistent with the definition of $A(\alpha)$.

Thresholds for E and E_0 . We take an indirect approach to the restrictions of the Delaunay triangulation to $U_\alpha \cap B_r$ and to $Z_0(\alpha)$. We begin by considering the restrictions to the power cell and its boundary,

$$\begin{aligned} E(\alpha) &= \text{Del}(U|Z(\alpha)); \\ E_0(\alpha) &= \text{Del}(U|\partial Z(\alpha)). \end{aligned}$$

We note that $E(\alpha)$ subdivides the underlying space of the closed star of z in $W(\alpha)$, the Delaunay triangulation of $U \cup \{z\}$ in which z has weight r^2 and all other points have weight α^2 . Furthermore $E_0(\alpha)$ is the link of z in $W(\alpha)$. To characterize when σ belongs to $E(\alpha)$ and $E_0(\alpha)$ we set $\eta_\sigma^2 = \varphi_\sigma^2 = r^2 + \varrho_0^2 - \|z - c_0\|^2$ if σ is an n -simplex. Otherwise, we set η_σ^2 equal to the maximum value η_τ^2 assigned to any n -simplex τ that has σ as a face. Symmetrically, we set φ_σ^2 equal to the minimum value φ_τ^2 assigned to any n -simplex τ that has σ as a face.

CONDITION E. Let σ be a simplex in $\text{Del}(U|\mathbb{R}^n)$. Then

$$\begin{aligned}\sigma \in E(\alpha) & \text{ iff } \alpha^2 \leq \eta_\sigma^2; \\ \sigma \in E_0(\alpha) & \text{ iff } \varphi_\sigma^2 \leq \alpha^2 \leq \eta_\sigma^2.\end{aligned}$$

There is again an alternative geometric interpretation of this condition. For a point $x \in V_\sigma$ let $\varrho(x)$ be the distance to the points $u_i \in \sigma$. Then x belongs to $Z(\alpha)$ iff $\|x - z\|^2 - r^2 \leq \varrho(x)^2 - \alpha^2$. The interpretation is now that $\varphi_\sigma^2 \leq \alpha^2 \leq \eta_\sigma^2$ iff V_σ has a non-empty intersection with $\partial Z(\alpha)$. For $\alpha^2 = \varphi_\sigma^2$ the rest of V_σ lies inside $Z(\alpha)$ while for $\alpha^2 = \eta_\sigma^2$ the rest of V_σ lies outside $Z(\alpha)$.

Thresholds for F and F_0 . Next we consider the restrictions of the sublevel set to the power cell and its boundary,

$$\begin{aligned}F(\alpha) & = \text{Del}(U|U_\alpha \cap Z(\alpha)); \\ F_0(\alpha) & = \text{Del}(U|U_\alpha \cap \partial Z(\alpha)).\end{aligned}$$

By definition, $F(\alpha)$ is a subset of $A(\alpha) \cap E(\alpha)$ and, similarly, $F_0(\alpha)$ is a subset of $A(\alpha) \cap E_0(\alpha)$, but generally we do not have equality. By definition of the power cell we have $U_\alpha \cap V_\sigma \subseteq Z(\alpha)$ if and only if $U_\alpha \cap V_\sigma \subseteq B_r$ as well as $U_\alpha \cap V_\sigma \cap Z(\alpha) = \emptyset$ if and only if $U_\alpha \cap V_\sigma \cap B_r = \emptyset$. We use this to formulate a test for deciding when a simplex in $E(\alpha)$ belongs to $F(\alpha)$ and the same for $E_0(\alpha)$ and $F_0(\alpha)$. Specifically, we give a condition when $U_\alpha \cap V_\sigma$ intersects ∂B_r . Consider the affine hull of the simplex, $D = \text{aff } \sigma$, and the affine hull of the corresponding intersection of Voronoi cells, $V = \text{aff } V_\sigma$. By construction, D is a plane of dimension $\dim \sigma$, V is a plane of dimension $n - \dim \sigma$, and the two intersect orthogonally in the point c defined earlier; see Figure 5.8. Assuming $\dim \sigma < n$, the balls of radius α centered at the vertices of σ intersect V in a ball of dimension $n - \dim \sigma$ whose center is c and whose square radius is $\alpha^2 - \varrho^2$. Let d be the distance

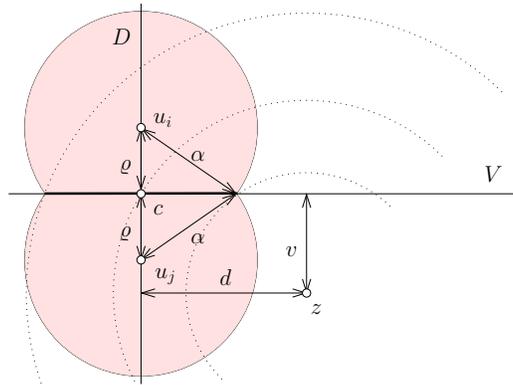


Figure 5.8: The point c is the center of the smallest circumsphere of the edge $\sigma = \{u_i, u_j\}$. The distance of z from c is $(d^2 + v^2)^{1/2}$.

of z from D and v the distance from V . When $|d - (r^2 - v^2)^{1/2}| = (\alpha^2 - \varrho^2)^{1/2}$ then the

$(n - \dim \sigma)$ -ball centered at c touches the sphere of radius r around z , and depending on the sign of $d - (r^2 - v^2)^{1/2}$ it does this either from the inside or from the outside. Taking the square and solving for α^2 we get $w(r^2) = r^2 + \varrho^2 + d^2 - v^2 - 2d(r^2 - v^2)^{1/2}$. This function has a minimum at $r^2 = d^2 + v^2$ for which the value is $w(r^2) = \varrho^2$. We now set

$$\begin{aligned}\kappa_\sigma^2 &= \begin{cases} w(r^2) & \text{if } r^2 \leq d^2 + v^2; \\ \varrho^2 & \text{if } d^2 + v^2 \leq r^2; \end{cases} \\ \lambda_\sigma^2 &= \begin{cases} \varrho^2 & \text{if } r^2 \leq d^2 + v^2; \\ w(r^2) & \text{if } d^2 + v^2 \leq r^2. \end{cases}\end{aligned}$$

They are both undefined for $r^2 < v^2$, the case when $w(r^2)$ is undefined. The above derivation does not make sense for an n -simplex σ but the following condition does, provided we set $\kappa_\sigma^2 = \lambda_\sigma^2 = \varrho^2$.

CONDITION F. Let σ be a simplex in $\text{Del}(U|\mathbb{R}^n)$. Then

$$\begin{aligned}\sigma \in F(\alpha) &\text{ iff } \max\{\alpha_\sigma^2, \kappa_\sigma^2\} \leq \alpha^2 \leq \eta_\sigma^2; \\ \sigma \in F_0(\alpha) &\text{ iff } \max\{\alpha_\sigma^2, \varphi_\sigma^2, \kappa_\sigma^2, \lambda_\sigma^2\} \leq \alpha^2 \leq \eta_\sigma^2.\end{aligned}$$

The geometric interpretation of Condition F should be clear. The only way for a simplex σ to belong to $A(\alpha)$ and $E(\alpha)$ but not to $F(\alpha)$ is that its circumcenter c lies outside the sphere around z , $r^2 < d^2 + v^2$, and the $(n - \dim \sigma)$ -ball does not touch this sphere, $\alpha^2 < \kappa_\sigma^2$. Similarly, the only way for σ to belong to $A(\alpha)$, $E_0(\alpha)$, and $F(\alpha)$ but not to $F_0(\alpha)$ is that c lies inside the ball around z , $d^2 + v^2 < r^2$, and the $(n - \dim \sigma)$ -ball does not touch the sphere, $\alpha^2 < \lambda_\sigma^2$.

Thresholds for K and K_0 . We now derive conditions for the restrictions to $U_\alpha \cap B_r$ and $Z_0(\alpha)$. Specifically, we use the fact that both these two spaces are swept out by the boundary of the sublevel set restricted to the power cell.

SWEEP LEMMA. Let $\alpha \geq 0$. Then

$$\begin{aligned}U_\alpha \cap B_r &= \bigcup_{0 \leq s \leq \alpha} \partial U_s \cap Z(s); \\ Z_0(\alpha) &= \bigcup_{0 \leq s \leq \alpha} U_s \cap \partial Z(s).\end{aligned}$$

PROOF. To establish the first equation we show that $x \in \partial U_s$ implies $x \in B_r$ iff $x \in Z(s)$. Indeed, $x \in \partial U_s$ implies $\|x - u_i\|^2 - s^2 = 0$, where $u_i \in U$ minimizes the distance to x . Thus $x \in Z(s)$ iff $\|x - z\|^2 - r^2 \leq 0$, as required. This implies that $\partial U_s \cap Z(s)$ sweeps out $U_\alpha \cap B_r$ as we increase s from 0 to α .

To prove the second equation we show that $x \in \partial Z(s)$ implies $x \in B_r$ iff $x \in U_\alpha$. Indeed, if $x \in \partial Z(s)$ then $\|x - z\|^2 - r^2 = \|x - u_i\|^2 - s^2$, where $u_i \in U$ again minimizes the distance to x . Both sides are non-positive at the same time, as required. This implies that $U_s \cap \partial Z(s)$ sweeps out $Z_0(\alpha) = B_r - \text{int } Z(\alpha)$ as we increase s from 0 to α . \square

The first equation in the Sweep Lemma implies that $U_\alpha \cap B_r$ is the union of the $U_s \cap Z(s)$, even without taking the boundary of U_s . Hence, $K(\alpha)$ is the union of the $F(s)$, for $0 \leq s \leq \alpha$. Similarly, the second equation implies that $K_0(\alpha)$ is the union of the $F_0(s)$, for $0 \leq s \leq \alpha$. We observe that if η_σ^2 is less than α_σ^2 or κ_σ^2 then $F(s)$ is empty, for all s , and so are $K(\alpha)$ and $K_0(\alpha)$. Otherwise, we get the conditions by dropping the upper bounds in Condition F.

CONDITION K. Let σ be a simplex in $\text{Del}(U|\mathbb{R}^n)$. If $\eta_\sigma^2 < \max\{\alpha_\sigma^2, \kappa_\sigma^2\}$ then $K(\alpha) = K_0(\alpha) = \emptyset$. Otherwise,

$$\begin{aligned} \sigma \in K(\alpha) & \text{ iff } \max\{\alpha_\sigma^2, \kappa_\sigma^2\} \leq \alpha^2; \\ \sigma \in K_0(\alpha) & \text{ iff } \max\{\alpha_\sigma^2, \varphi_\sigma^2, \kappa_\sigma^2, \lambda_\sigma^2\} \leq \alpha^2. \end{aligned}$$

Indeed, the upper bound just guarantees that $U_\alpha \cap V_\sigma$ has a non-empty intersection with $Z(\alpha)$. Since we increase s from 0 to α there is a non-empty intersection between $U_s \cap V_\sigma$ and $Z(s)$, for some s , iff there is a non-empty intersection between $U_\alpha \cap V_\sigma$ and B_r , which is captured by the remaining inequality.

Thresholds for A and A_0 . We can further simplify the condition by restricting the Delaunay triangulation to the sublevel set outside the power cell,

$$A_0(\alpha) = \text{Del}(U|U_\alpha - \text{int } Z(\alpha)).$$

Equivalently, $A_0(\alpha) = (A(\alpha) - K(\alpha)) \cup K_0(\alpha)$. By construction, $A_0(\alpha)$ is a subcomplex of $A(\alpha)$ and the difference is $A(\alpha) - A_0(\alpha) = K(\alpha) - K_0(\alpha)$. We also consider the diagram of relative homology groups with induced homomorphisms,

$$\begin{array}{ccc} (K(\alpha), K_0(\alpha)) & \xrightarrow{g_\alpha^{\alpha'}} & (K(\alpha'), K_0(\alpha')), \\ \downarrow i_\alpha & & \downarrow i_{\alpha'} \\ (A(\alpha), A_0(\alpha)) & \xrightarrow{j_\alpha^{\alpha'}} & (A(\alpha'), A_0(\alpha')), \end{array}$$

Since all maps are inclusions the diagram commutes and so does the corresponding diagram of relative homology groups. Finally, excision implies that the vertical maps induce isomorphisms. It follows that the persistence diagrams we get from the $(A(\alpha), A_0(\alpha))$ are indeed the same as the one of the $(K(\alpha), K_0(\alpha))$. The reason for making this final substitution is computational convenience.

CONDITION A'. Let σ be a simplex in $\text{Del}(U|\mathbb{R}^n)$. Then

$$\begin{aligned} \sigma \in A(\alpha) & \text{ iff } \alpha_\sigma^2 \leq \alpha^2; \\ \sigma \in A_0(\alpha) & \text{ iff } \max\{\alpha_\sigma^2, \varphi_\sigma^2, \lambda_\sigma^2\} \leq \alpha^2. \end{aligned}$$

5.7.2 Trajectories

In this subsection, we extend the results of the last subsection to the situation in which the radius of the restricting ball varies. It is convenient to add r to the notation, writing $K(\alpha, r) = K(\alpha)$ and similar for other complexes.

Absolute homology. Recall that the first half of the sequence (5.3) gives the same persistence diagrams as the sequence of the $K(\alpha, r)$, where r is fixed. Each simplex σ in the Delaunay triangulation undergoes the same kind of evolution as α goes from 0 to ∞ :

Step 1. σ becomes a member of $K(\alpha, r)$;

a step that may also be skipped. The thresholds that determine membership in this complex vary continuously with the radius. Looking at the squares and parametrizing by r^2 , we get three functions, $\alpha_\sigma^2, \kappa_\sigma^2, \eta_\sigma^2 : [0, \infty) \rightarrow [0, \infty)$. Rewriting the first half of Condition K we find how these functions control membership in $K(\alpha, r)$.

ABSOLUTE HOMOLOGY EVOLUTION LEMMA. Let σ be a simplex in $\text{Del}(U|\mathbb{R}^n)$ and $r \geq 0$. If $\eta_\sigma^2(r^2) < \max\{\alpha_\sigma^2(r^2), \kappa_\sigma^2(r^2)\}$ then $K(\alpha, r) = \emptyset$. Otherwise,

$$\sigma \in K(\alpha, r) \quad \text{iff} \quad \max\{\alpha_\sigma^2(r^2), \kappa_\sigma^2(r^2)\} \leq \alpha^2.$$

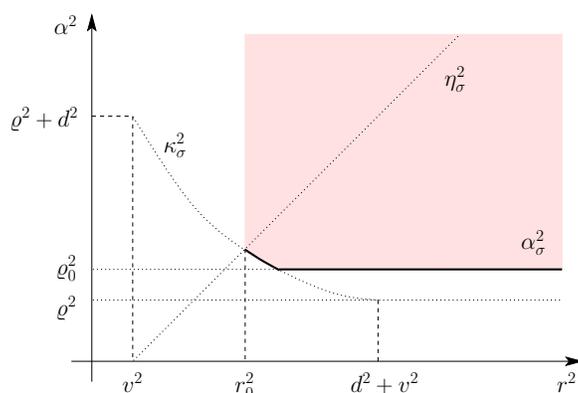


Figure 5.9: The simplex σ belongs to $K(\alpha, r)$ iff the point (α^2, r^2) lies in the shaded region above the function that discriminates between membership and non-membership.

The inequality is illustrated in Figure 5.9. The smallest value of r^2 at which σ belongs to any of the $K(\alpha, r)$ is when $\max\{\alpha_\sigma^2(r^2), \kappa_\sigma^2(r^2)\} = \eta_\sigma^2(r^2)$. Letting this value be r_0^2 , we have a continuous function from $[r_0^2, \infty)$ to $[0, \infty)$ that discriminates membership of σ in $K(\alpha, r)$ from non-membership. In general, this function consists of two portions, the one on the left contributed by κ_σ^2 and the one on the right contributed by α_σ^2 . If $c = c_0$ then $r_0^2 \leq d^2 + v^2$ which guarantees that the left portion is non-empty. Otherwise, it is possible that the function is constant over the entire interval it is defined.

Relative homology. Recall that the first half of the sequence (5.4) gives the same persistence diagrams as the sequence of the $(A(\alpha, r), A_0(\alpha, r))$, where r is again fixed. Each simplex $\sigma \in \text{Del}(U|\mathbb{R}^n)$ undergoes the same kind of evolution as α goes from 0 to ∞ :

Step 1. σ becomes a member of $A(\alpha, r)$;

Step 2. σ becomes a member of $A_0(\alpha, r)$,

two steps that may also occur simultaneously. Similar to the absolute homology case we get three continuous functions, $\alpha_\sigma^2, \varphi_\sigma^2, \lambda_\sigma^2 : [0, \infty) \rightarrow [0, \infty)$. Rewriting Condition A' we find how these functions control membership in $(A(\alpha, r), A_0(\alpha, r))$.

RELATIVE HOMOLOGY EVOLUTION LEMMA. Let σ be a simplex in $\text{Del}(U|\mathbb{R}^n)$ and $r \geq 0$. Then

$$\begin{aligned} \sigma \in A(\alpha, r) & \text{ iff } \alpha_\sigma^2(r^2) \leq \alpha^2; \\ \sigma \in A_0(\alpha, r) & \text{ iff } \max\{\alpha_\sigma^2(r^2), \varphi_\sigma^2(r^2), \lambda_\sigma^2(r^2)\} \leq \alpha^2. \end{aligned}$$

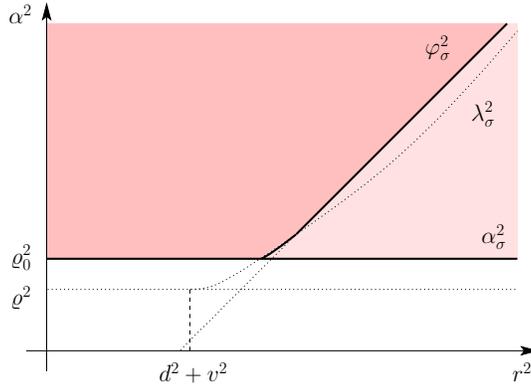


Figure 5.10: The simplex σ belongs to $A(\alpha, r)$ and $A_0(\alpha, r)$ iff the point (α^2, r^2) lies in the shaded region above both membership discriminating functions, and σ belongs to $A(\alpha, r)$ but not to $A_0(\alpha, r)$ iff the point lies in the light shaded region between the two functions.

Figure 5.10 illustrates the inequalities for a simplex σ whose smallest circumsphere is different from its smallest empty circumsphere: $c \neq c_0$ and $\varrho < \varrho_0$. Often the picture is even simpler. For example, if $\dim \sigma = n$ then $\varrho = \varrho_0$ and the function that controls the membership of σ in $A_0(\alpha, r)$ simplifies to $r^2 + \varrho^2 - (d^2 + v^2)$. It follows that the region in which σ belongs to $A(\alpha, r)$ but not to $A_0(\alpha, r)$ is a wedge to the right of the corner point at $r^2 = d^2 + v^2$ and $\alpha^2 = \varrho^2$.

Crossings and transpositions. The two Evolution Lemmas introduce three *membership functions* per simplex,

$$\begin{aligned} f_{\sigma, K}(r^2) &= \max\{\alpha_\sigma(r^2), \kappa_\sigma^2(r^2)\}; \\ f_{\sigma, A}(r^2) &= \alpha_\sigma(r^2); \\ f_{\sigma, A_0}(r^2) &= \max\{\alpha_\sigma(r^2), \varphi_\sigma^2(r^2), \lambda_\sigma^2(r^2)\}. \end{aligned}$$

All three are continuous except for $f_{\sigma,K}$ which is undefined for $r^2 < r_0^2$. We set $f_{\sigma,K}(r^2) = \infty$ for $0 \leq r^2 < r_0^2$, and viewing infinity as just another value we thus get a function with a single discontinuity. We will see shortly that having just one discontinuity per function does not substantially affect the results we reap. Each membership function consists of pieces contributed by α_σ^2 , φ_σ^2 , κ_σ^2 , and λ_σ^2 . The α_σ^2 are constant functions and the φ_σ^2 are functions of degree one in the variable r^2 . The κ_σ^2 and λ_σ^2 are also of degree one in r^2 except that they have an additional square root term. Each membership function consists of a constant number of pieces. The only case in which this is perhaps not entirely obvious is for φ_σ^2 which is a point-wise minimum of functions φ_τ^2 over all n -simplices τ in $\text{Del}(U|\mathbb{R}^n)$ that contain σ as a face. There can be an arbitrary number of such n -simplices but since the corresponding functions are of the form $\varphi_\tau^2(r^2) = r^2 + \text{const}$, only one provides all the minima. Two functions $f, g : [0, \infty) \rightarrow [0, \infty)$ cross at r_1 if

$$[f(r_1 - \varepsilon) - g(r_1 - \varepsilon)][f(r_1 + \varepsilon) - g(r_1 + \varepsilon)] < 0$$

for all sufficiently small $\varepsilon > 0$. For continuous functions this happens at a point at which they agree, $f(r_1) = g(r_1)$. Similarly, we define when f and g cross at an interval along which they agree. Any two pieces of the membership functions have constant complexity and cross at most a constant number of times. Since each membership function consists only of a constant number of such pieces, this implies that any two membership functions cross at most some constant number of times. Drawing the graphs of the membership functions of all m simplices in $\text{Del}(U|\mathbb{R}^n)$ we therefore get an arrangement of $3m$ curves with at most some constant times m^2 vertices, edges, and regions. When we sweep the arrangement with a vertical line from left to right we do at most some constant times m^2 transpositions at the vertices and another constant times m^2 transpositions because of the m discontinuities. The $(\alpha|r)$ -Vineyard Theorem follows.

5.8 Discussion

The main contribution of this chapter is the development of topological data analysis methods for the algorithmic study of sampled stratified spaces. Specifically, we show how to assess the local homology at a point and prove that a sufficiently dense sample implies the correctness of our assessment. While non-trivial, the described algorithm is readily implementable and runs in time at most cubic in the number of simplices in the Delaunay triangulation.

We expect that in practice the rate-limiting step of our algorithm will be the construction of the Delaunay triangulation. In most cases, only a small subcomplex of the Delaunay triangulation is relevant for the assessment. More generally, we may restrict the construction to simplices that connect points at distance at most some threshold r from each other. Examples of complexes that limit themselves to such simplices are the alpha, Čech, and Vietoris-Rips complexes; see [41, 76]. It would be interesting to develop fast output-sensitive algorithms for these complexes and to substitute them for the Delaunay triangulation of \mathbb{R}^n on which the methods in this chapter are currently based.

The ability to assess the local homology of a stratified space at a point from a finite sample is an important step in a more ambitious program. The larger goal is the construction of the stratified space or a description of the class of stratified spaces that possibly give rise to the observed sample.

The logical next step is determining the local dimension of the space at a point, i.e. the dimensions of the strata close to the given point. Knowing local homology alone is not enough as the following example illustrates. However, it allows us to group points into strata whose dimensions are readily definable. Consider the following two examples in Figure 5.11. The first is a tetrahedron with its barycenter joined to each one of its edges.

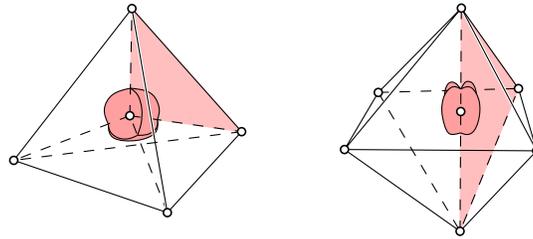


Figure 5.11: Left: tetrahedron with its barycenter joined to its edges. Right: octahedron with an edge connecting opposite vertices joined to the remaining vertices. Local neighborhoods of the barycenter on the left, and a point on the additional edge on the right are highlighted. The local homology is the same in both cases, but the local dimensions differ.

The local homology groups at the barycenter consist of three independent 2-cycles, and the barycenter belongs to the 0-stratum. The second example is an octahedron with an edge added between a pair of opposite vertices. The edge is joined with each one of the four remaining vertices resulting in four triangles. The local homology groups of every point on the edge consist of three independent 2-cycles with the edge forming a piece of the 1-stratum.

We note that the notion of local dimension that we desire to compute reflects the stratification of the space and as such is different than the dimension typically considered in computational geometry and machine learning literature. A point with enough rays coming out of it can span a space of arbitrary dimension, while still belonging to a 0-dimensional stratum.

Chapter 6

Kernels, Images, and Cokernels

6.1 Introduction

Up to this point we have considered nested sequences of spaces, $\mathbb{X}_0 \subseteq \mathbb{X}_1 \subseteq \dots \subseteq \mathbb{X}_m$, and persistent homology has arisen from considering the corresponding sequences of homology groups, $H(\mathbb{X}_0) \rightarrow H(\mathbb{X}_1) \rightarrow \dots \rightarrow H(\mathbb{X}_m)$, connected from left to right by homomorphic maps induced by inclusion. Persistence tracks when a homology class is born and dies in this sequence. As defined in Section 2.2 this can also be done for an arbitrary sequence of vector spaces connected by homomorphic maps. Motivation for studying such more general sequences is derived from the following two applications. First, we describe how sequences of kernels can be used to refine the multi-scale assessment of local homology from the previous chapter. Second, we use sequences of images to introduce a notion of persistence that filters out noise induced by imprecise specifications of domains. This contrasts standard persistence which can handle imprecise function values but not imprecise domains. As an application, we approximate the persistence diagram of a function knowing only its values at a finite set of points. This chapter is based on the joint work with David Cohen-Steiner, Herbert Edelsbrunner, and John Harer [28]; its main contributions are two-fold:

- an algorithm that computes the persistence diagrams of sequences of kernels, images, and cokernels in time at most cubic in the size of the simplicial complexes representing the data is given in Section 6.3, and its correctness is shown in Section 6.4;
- applications of the algebraic and algorithmic results to measuring local homology and to approximating persistence diagrams of noisy functions on noisy domains are described in Section 6.6.

6.2 Algebra

This section extends the concept of persistent homology to sequences of kernels, images, and cokernels. It also proves that the persistence diagrams of these extensions are stable.

Kernels, images, and cokernels. For the extension of persistence to kernels, images, and cokernels we consider two functions, $f : \mathbb{X} \rightarrow \mathbb{R}$ and a majorizing function $g : \mathbb{Y} \rightarrow \mathbb{R}$ defined on a subspace $\mathbb{Y} \subseteq \mathbb{X}$, that is, $f(y) \leq g(y)$ for all $y \in \mathbb{Y} \subseteq \mathbb{X}$. Assuming both functions are tame, we order the collection of critical values of f and g and interleave them with a sequence of real values s_i as in Section 2.2. The corresponding sequences of sublevel sets give rise to two parallel sequences of homology groups,

$$\begin{array}{ccccccc} \mathrm{H}(\mathbb{X}_0) & \rightarrow & \mathrm{H}(\mathbb{X}_1) & \rightarrow & \dots & \rightarrow & \mathrm{H}(\mathbb{X}_m) \\ & & \uparrow j_0 & & \uparrow j_1 & & \dots & & \uparrow j_m \\ & & \mathrm{H}(\mathbb{Y}_0) & \rightarrow & \mathrm{H}(\mathbb{Y}_1) & \rightarrow & \dots & \rightarrow & \mathrm{H}(\mathbb{Y}_m), \end{array}$$

where $\mathbb{X}_i = f^{-1}(-\infty, s_i]$ and $\mathbb{Y}_i = g^{-1}(-\infty, s_i]$. The two sequences are connected by homomorphisms $j_i : \mathrm{H}(\mathbb{Y}_i) \rightarrow \mathrm{H}(\mathbb{X}_i)$ induced by the inclusions $\mathbb{Y}_i \subseteq \mathbb{X}_i$. We call this the *two function setting*, in contrast to the more special *one function setting* in which g is the restriction of f to \mathbb{Y} . More about the relationship between the two settings later. We are

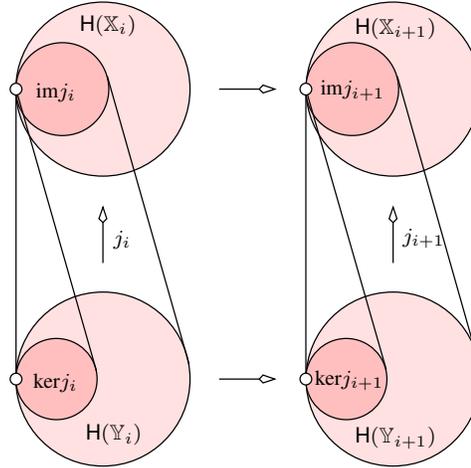


Figure 6.1: A square of four homology groups and the maps between them. The square commutes because all four maps are induced by inclusions.

interested in the kernels, images, and cokernels of the connecting homomorphisms,

$$\begin{aligned} \ker j_i &= \{\gamma \in \mathrm{H}(\mathbb{Y}_i) \mid j_i(\gamma) = 0 \in \mathrm{H}(\mathbb{X}_i)\}; \\ \mathrm{im} j_i &= \{j_i(\gamma) \in \mathrm{H}(\mathbb{X}_i) \mid \gamma \in \mathrm{H}(\mathbb{Y}_i)\}; \\ \mathrm{cok} j_i &= \mathrm{H}(\mathbb{X}_i) / \mathrm{im} j_i. \end{aligned}$$

We note that the “coimage” of the map j_i , in symbols $H(\mathbb{Y}_i)/\ker j_i$, is isomorphic to the image of this map, and therefore does not deserve any special attention. Figure 6.1 illustrates this construction for two contiguous spaces in both sequences. The square defined by the four homology groups commutes. It follows that the inclusion $\mathbb{Y}_i \subseteq \mathbb{Y}_{i+1}$ induces a homomorphism $\ker j_i \rightarrow \ker j_{i+1}$. Similarly, the inclusion $\mathbb{X}_i \subseteq \mathbb{X}_{i+1}$ induces a homomorphism $\operatorname{im} j_i \rightarrow \operatorname{im} j_{i+1}$ and another homomorphism $\operatorname{cok} j_i \rightarrow \operatorname{cok} j_{i+1}$. We thus get sequences of kernels, images, and cokernels,

$$\begin{aligned} \operatorname{Ker}(g \rightarrow f) &: \ker j_0 \rightarrow \ker j_1 \rightarrow \dots \rightarrow \ker j_m; \\ \operatorname{Im}(g \rightarrow f) &: \operatorname{im} j_0 \rightarrow \operatorname{im} j_1 \rightarrow \dots \rightarrow \operatorname{im} j_m; \\ \operatorname{Cok}(g \rightarrow f) &: \operatorname{cok} j_0 \rightarrow \operatorname{cok} j_1 \rightarrow \dots \rightarrow \operatorname{cok} j_m, \end{aligned}$$

all connected from left to right by homomorphisms. Homology classes are born and die in these sequences same as in the sequences of homology groups. Using the definition of persistence for vector spaces we can therefore define *persistent kernels*, *persistent images*, and *persistent cokernels* as well as construct the corresponding persistence diagrams, which we denote as $\operatorname{Dgm}(\ker g \rightarrow f)$, $\operatorname{Dgm}(\operatorname{im} g \rightarrow f)$, and $\operatorname{Dgm}(\operatorname{cok} g \rightarrow f)$.

Birth-death combinations. We consider the generic case in which changes happen one at a time. An event thus corresponds to a birth, a death, or no change in the kernel, in the image, and in the cokernel, giving rise to 27 different combinations. But the ranks of these groups are not independent, that is,

$$\begin{aligned} \operatorname{rank} \ker j_i + \operatorname{rank} \operatorname{im} j_i &= \operatorname{rank} H(\mathbb{Y}_i); \\ \operatorname{rank} \operatorname{im} j_i + \operatorname{rank} \operatorname{cok} j_i &= \operatorname{rank} H(\mathbb{X}_i), \end{aligned}$$

for all i . We can therefore relate the births and deaths in the three sequences using the births and deaths in the sequences of homology groups of the \mathbb{Y}_i and of the \mathbb{X}_i . The first equation eliminates two of the nine combinations for kernels and images. Another combination is eliminated by $\ker j_i$ being a subgroup of $H(\mathbb{Y}_i)$, hence a death in the kernel implies a death in the homology group. Table 6.1 lists the remaining six cases. Case A

Case	$\ker j_i$	$\operatorname{im} j_i$	$H(\mathbb{Y}_i)$
A	birth	death	—
B	—	birth	birth
C, D	—	—	—
E	—	death	death
F	death	—	death
P	birth	—	birth

Table 6.1: The five cases in the two function setting relating kernels and images. Except for Case P they also occur in the one function setting.

occurs for example when $\mathbb{Y}_{i-1} = \mathbb{X}_{i-1} = \mathbb{Y}_i$ is a circle and \mathbb{X}_i is obtained by adding a spanning disk. Case B occurs when $\mathbb{Y}_{i-1} = \mathbb{X}_{i-1}$ is a point and $\mathbb{Y}_i = \mathbb{X}_i$ is obtained by

adding an arc that completes the point to a circle. Case C occurs when $\mathbb{Y}_{i-1} = \mathbb{X}_{i-1} = \mathbb{Y}_i$ is a point and \mathbb{X}_i is again obtained by adding an arc that forms a circle. We also retain ranks in Case D which occurs when \mathbb{X}_{i-1} is a circle, $\mathbb{Y}_{i-1} = \mathbb{Y}_i$ is a point on this circle, and \mathbb{X}_i is obtained by adding a spanning disk. Case E occurs when $\mathbb{Y}_{i-1} = \mathbb{X}_{i-1}$ is a circle and $\mathbb{Y}_i = \mathbb{X}_i$ is obtained by adding a spanning disk. Case F occurs when \mathbb{X}_{i-1} is a disk, \mathbb{Y}_{i-1} is its boundary circle, and we get \mathbb{Y}_i and \mathbb{X}_i by adding another spanning disk to both spaces. Finally, Case P occurs when $\mathbb{X}_{i-1} = \mathbb{X}_i$ is a disk, \mathbb{Y}_{i-1} is a point of its boundary circle, and \mathbb{Y}_i is obtained by adding the rest of the circle. This last case happens in the two function setting but not in the one function setting because it requires points that are added to the sublevel set of g strictly after they are added to the sublevel set of f .

Similarly, the second equation eliminates two of the nine combinations for images and cokernels. Another combination is eliminated by $\text{im } j_i$ being a subgroup of $H(\mathbb{X}_i)$, hence a death in the image implies a death in the homology group. Table 6.2 lists the remaining six cases. Cases A to F have been described above and the example for Case P also works

Case	$\text{im } j_i$	$\text{cok } j_i$	$H(\mathbb{X}_i)$
A, E	death	—	death
B	birth	—	birth
C, F	—	birth	birth
D	—	death	death
Q	—	—	—
R	birth	death	—

Table 6.2: The five cases in the two function setting relating cokernels and images. Except for Cases Q and R they also occur in the one function setting.

for Case Q. Case R occurs when $\mathbb{X}_{i-1} = \mathbb{Y}_i = \mathbb{X}_i$ is a circle and \mathbb{Y}_{i-1} is a point on that circle. Cases Q and R do not happen in the one function setting in which every change has a non-zero effect on the rank of the homology group of \mathbb{X}_i .

Mapping cylinder. We reduce the two function setting to the one function setting using a construction that will be exploited by the algorithm described in Section 6.3. Specifically, the *mapping cylinder* of the pair $\mathbb{Y} \subseteq \mathbb{X}$ is the space $\mathbb{X}' = \mathbb{X} \cup (\mathbb{Y} \times [0, 1])$ obtained by gluing $\mathbb{Y} \subseteq \mathbb{X}$ to $\mathbb{Y} \times \{0\} \subseteq \mathbb{Y} \times [0, 1]$. This is illustrated in Figure 6.2. The function

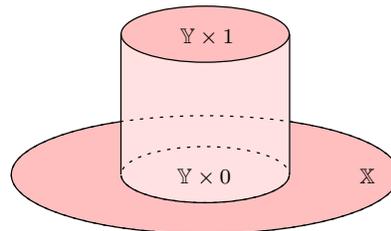


Figure 6.2: The mapping cylinder of $\mathbb{Y} \subseteq \mathbb{X}$.

$f' : \mathbb{X}' \rightarrow \mathbb{R}$ agrees with f on \mathbb{X} and with g on $\mathbb{Y}' = \mathbb{Y} \times \{1\}$, linearly interpolating in between, that is, $f'(x) = f(x)$ for every $x \in \mathbb{X}$ and $f'(y, t) = (1 - t)f(y) + tg(y)$ for every $y \in \mathbb{Y}$ and every $t \in [0, 1]$. The pair of functions f' and $g' = f'|_{\mathbb{Y}'}$ induces homomorphisms $j'_i : H(\mathbb{Y}'_i) \rightarrow H(\mathbb{X}'_i)$; see Figure 6.3. The corresponding sequences of kernels, images, and cokernels are

$$\begin{aligned} \text{Ker}(g' \rightarrow f') & : \ker j'_0 \rightarrow \ker j'_1 \rightarrow \dots \rightarrow \ker j'_m; \\ \text{Im}(g' \rightarrow f') & : \text{im } j'_0 \rightarrow \text{im } j'_1 \rightarrow \dots \rightarrow \text{im } j'_m; \\ \text{Cok}(g' \rightarrow f') & : \text{cok } j'_0 \rightarrow \text{cok } j'_1 \rightarrow \dots \rightarrow \text{cok } j'_m. \end{aligned}$$

We claim that they contain the same information as the sequences $\text{Ker}(g \rightarrow f)$, $\text{Im}(g \rightarrow f)$, and $\text{Cok}(g \rightarrow f)$.

$$\begin{array}{ccccc} & & H(\mathbb{X}'_i) & \longrightarrow & H(\mathbb{X}'_{i+1}) \\ & \nearrow & \uparrow & & \nearrow \\ H(\mathbb{X}_i) & \longrightarrow & H(\mathbb{X}_{i+1}) & \xrightarrow{j'_{i+1}} & H(\mathbb{Y}'_{i+1}) \\ & \searrow & \downarrow j'_i & & \downarrow \\ & & H(\mathbb{Y}'_i) & \xrightarrow{j_{i+1}} & H(\mathbb{Y}'_{i+1}) \\ & \nearrow & \uparrow j_i & & \nearrow \\ H(\mathbb{Y}_i) & \longrightarrow & H(\mathbb{Y}_{i+1}) & & \end{array}$$

Figure 6.3: The diagram of homology groups of two contiguous sublevel sets of f, g, f', g' . The diagram commutes because ten of the twelve maps are induced by inclusions, and the left and right maps of the bottom square are induced by the inverses of the mapping cylinder retractions restricted to \mathbb{Y}'_i and \mathbb{Y}'_{i+1} respectively.

MAPPING CYLINDER LEMMA. The pairs of functions f, g and f', g' define the same persistence diagrams for kernels, images, and cokernels:

$$\text{Dgm}(\text{grp } g \rightarrow f) = \text{Dgm}(\text{grp } g' \rightarrow f')$$

for $\text{grp} \in \{\ker, \text{im}, \text{cok}\}$.

PROOF. We note that \mathbb{X}_i is a deformation retract of \mathbb{X}'_i . This implies that the map from $H(\mathbb{X}_i)$ to $H(\mathbb{X}'_i)$ induced by the inclusion $\mathbb{X}_i \subseteq \mathbb{X}'_i$ is an isomorphism. Similarly, the map from $H(\mathbb{Y}_i)$ to $H(\mathbb{Y}'_i)$ implied by the inverse of the retraction is an isomorphism. The maps $j_i : H(\mathbb{Y}_i) \rightarrow H(\mathbb{X}_i)$ and $j'_i : H(\mathbb{Y}'_i) \rightarrow H(\mathbb{X}'_i)$ are also induced by inclusions which implies that the left square in the diagram of Figure 6.3 commutes. It follows that the pairs of kernels, images, and cokernels are isomorphic, $\ker j_i \simeq \ker j'_i$, $\text{im } j_i \simeq \text{im } j'_i$, and $\text{cok } j_i \simeq \text{cok } j'_i$. Similarly, the right square commutes and the kernels, images, and cokernels of j_{i+1} and j'_{i+1} are isomorphic. To prove that the two sequences of kernels define the same persistence diagram we still need to consider the bottom square in Figure 6.3. The left and right maps are isomorphisms and the square commutes by construction. Hence $\text{Dgm}(\ker g \rightarrow f) = \text{Dgm}(\ker g' \rightarrow f')$. Similarly, we get

$\text{Dgm}(\text{im } g \rightarrow f) = \text{Dgm}(\text{im } g' \rightarrow f')$ and $\text{Dgm}(\text{cok } g \rightarrow f) = \text{Dgm}(\text{cok } g' \rightarrow f')$ by considering the top square in Figure 6.3. \square

Stability. An important property of the persistence diagrams is their stability originally proved in [25]. As stated in Section 2.2, the bottleneck distance between the diagrams of two functions f and f'' is bounded from above by the difference between the functions,

$$d_B(\text{Dgm}(f), \text{Dgm}(f'')) \leq \|f - f''\|_\infty.$$

Here d_B is the maximum of $\|u - \gamma(u)\|_\infty$, where u is a point in the diagram of f and γ is a dimension-preserving bijection between the diagrams of f and of f'' . Recall that the points on the diagonal belong to the diagrams and can therefore be used in the effort to find a matching γ that minimizes the length of the longest edge. The proof of stability given in [25] can be adapted to the setting in this chapter. Specifically, we consider the maps $j_a : H(\mathbb{Y}_a) \rightarrow H(\mathbb{X}_a)$ and $j''_{a+\varepsilon} : H(\mathbb{Y}''_{a+\varepsilon}) \rightarrow H(\mathbb{X}''_{a+\varepsilon})$, where ε is the larger of the two differences between functions, $\|f - f''\|_\infty$ and $\|g - g''\|_\infty$, and $\mathbb{Y}_a, \mathbb{X}_a, \mathbb{Y}''_{a+\varepsilon}, \mathbb{X}''_{a+\varepsilon}$ are the sublevel sets of g, f, g'', f'' for thresholds a and $a + \varepsilon$. To adapt the proof we need that the maps induced by the inclusions $\mathbb{Y}_a \subseteq \mathbb{Y}''_{a+\varepsilon}$ and $\mathbb{X}_a \subseteq \mathbb{X}''_{a+\varepsilon}$ send the kernel of j_a into the kernel of $j''_{a+\varepsilon}$. But this follows from the commutativity of the diagram

$$\begin{array}{ccc} H(\mathbb{X}_a) & \longrightarrow & H(\mathbb{X}''_{a+\varepsilon}) \\ \uparrow j_a & & \uparrow j''_{a+\varepsilon} \\ H(\mathbb{Y}_a) & \longrightarrow & H(\mathbb{Y}''_{a+\varepsilon}). \end{array}$$

Similarly, we need that the inclusions $\mathbb{Y}''_a \subseteq \mathbb{Y}_{a+\varepsilon}$ and $\mathbb{X}''_a \subseteq \mathbb{X}_{a+\varepsilon}$ send $\ker j''_a$ into $\ker j_{a+\varepsilon}$ which follows by symmetry. With this property the original proof of stability goes through and we refer to [25] for details. The arguments for the images and the cokernels are the same and we state the results.

STABILITY THEOREM. Let $f, f'' : \mathbb{X} \rightarrow \mathbb{R}$ and $g, g'' : \mathbb{Y} \rightarrow \mathbb{R}$ with $f(y) \leq g(y)$ and $f''(y) \leq g''(y)$ for every $y \in \mathbb{Y} \subseteq \mathbb{X}$ and $\varepsilon = \max\{\|f - f''\|_\infty, \|g - g''\|_\infty\}$. Then the bottleneck distance between the persistence diagrams is bounded from above by the difference between the functions:

$$d_B(\text{Dgm}(\text{grp } g \rightarrow f), \text{Dgm}(\text{grp } g'' \rightarrow f'')) \leq \varepsilon,$$

for $\text{grp} \in \{\ker, \text{im}, \text{cok}\}$, provided f, g, f'' , and g'' are continuous and tame and there is a triangulation of \mathbb{X} in which \mathbb{Y} arises as a subcomplex.

6.3 Algorithms

In this section, we describe the algorithms for computing the persistence diagrams of the sequences of kernels, images, and cokernels. At their core is the reduction of a matrix as described in Section 2.2.

Partial and reordered matrices. We prepare the computation of the persistence diagrams by reducing five matrices. By the Mapping Cylinder Lemma, we can restrict ourselves to the one function setting. We therefore assume two simplicial complexes, $L \subseteq K$, let $f : K \rightarrow \mathbb{R}$ be an injective function whose sublevel sets are subcomplexes of K , and let g be the restriction of f to L . We write D_f for the incidence matrix of K whose rows and columns are ordered by f . Similarly, we write D_g for the incidence matrix of L whose rows and columns are ordered by g .

Step 1 Reduce the two incidence matrices to get $R_f = D_f V_f$ and $R_g = D_g V_g$.

Step 2 Reorder the rows of D_f leaving the columns untouched to get a new matrix D_{im} . Specifically, its rows correspond to the simplices in L , ordered by g , followed by the simplices in $K - L$, ordered by f . Reduce the new matrix to get $R_{\text{im}} = D_{\text{im}} V_{\text{im}}$; see Figure 6.4.

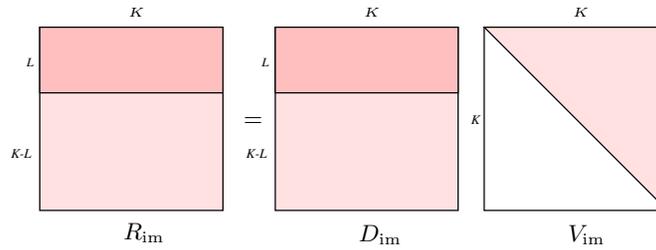


Figure 6.4: Matrices computed in the reduction of the incidence matrix of K with reordered rows. The matrix V_{im} is upper-triangular with all ones in the diagonal.

Step 3 Delete some of the columns from V_{im} and reorder the rows to get a new matrix D_{ker} . Specifically, keep the columns that represent cycles and remove all others. Furthermore, reorder the rows so they correspond to the simplices in L , ordered by g , followed by the simplices in $K - L$, ordered by f . Finally, reduce the new matrix to get $R_{\text{ker}} = D_{\text{ker}} V_{\text{ker}}$.

Step 4 Starting again with D_f , replace some of the columns to get a new matrix D_{cok} . Specifically, substitute columns in V_g that represent cycles for the corresponding columns in D_f , adding zeros to compensate for the simplices in $K - L$, which are missing in V_g . Reduce the new matrix to get $R_{\text{cok}} = D_{\text{cok}} V_{\text{cok}}$.

We note that reducing D_f is redundant because the type information it furnishes is also available from R_{im} . We still use R_f because this clarifies which information is used where.

Births, deaths, and pairs. We use the reduced matrices to compute the persistence diagrams of the sequences of kernels, images, and cokernels. Specifically, R_f and R_g (and in one case R_{im}) decide which simplices give birth and which give death and R_{ker} , R_{im} , R_{cok} determine how the births match up with the deaths. We begin with the sequence of kernels and recall the relevant Cases A and F in Table 6.1.

Algorithm for kernels:

Birth. A simplex σ gives birth in $\text{Ker}(g \rightarrow f)$ iff $\sigma \in K - L$, σ is negative in R_f , and the lowest one in its column in R_{im} corresponds to a simplex in L .

Death. A simplex τ gives death in $\text{Ker}(g \rightarrow f)$ iff $\tau \in L$, τ is negative in R_g , and τ is positive in R_f . In this case, the lowest one in the column of τ in R_{ker} corresponds to a simplex $\sigma \in K - L$ that gives birth in $\text{Ker}(g \rightarrow f)$. Then (σ, τ) is a pair.

A dimension p homology class is given birth to in the kernel by a $(p + 1)$ -simplex and it dies at the hand of another $(p + 1)$ -simplex. The dimension p persistence diagram thus consists of all points $(f(\sigma), f(\tau))$ encoding pairs of $(p + 1)$ -simplices identified in the Death case as well as all points $(f(\sigma), \infty)$ encoding unpaired $(p + 1)$ -simplices identified in the Birth case. We continue with the sequence of images and recall the relevant Cases A, B, E in Tables 6.1 and 6.2.

Algorithm for images:

Birth. A simplex σ gives birth in $\text{Im}(g \rightarrow f)$ iff $\sigma \in L$ and σ is positive in R_g .

Death. A simplex τ gives death in $\text{Im}(g \rightarrow f)$ iff τ is negative in R_f and the lowest one in its column in R_{im} corresponds to a simplex $\sigma \in L$. Then (σ, τ) is a pair.

Note that the Death case splits into Case A with $\tau \in K - L$ and Case E with $\tau \in L$. The dimension p persistence diagram consists of all points $(f(\sigma), f(\tau))$ encoding pairs of p - and $(p + 1)$ -simplices identified in the Death case as well as points $(f(\sigma), \infty)$ encoding unpaired p -simplices identified in the Birth case. We continue with the sequence of cokernels and recall the relevant Cases C, F and D in Table 6.2.

Algorithm for cokernels:

Birth. A simplex σ gives birth in $\text{Cok}(g \rightarrow f)$ iff σ is positive in R_f and it is either in $K - L$ or negative in R_g .

Death. A simplex τ gives death in $\text{Cok}(g \rightarrow f)$ iff τ is negative in R_f and the lowest one in its column in R_{im} corresponds to a simplex in $K - L$. In this case, the lowest one in the column of τ in R_{cok} corresponds to a simplex σ that gives birth in $\text{Cok}(g \rightarrow f)$. Then (σ, τ) is a pair.

The dimension p persistence diagram consists of all points $(f(\sigma), f(\tau))$ encoding pairs of p - and $(p + 1)$ -simplices identified in the Death case as well as points $(f(\sigma), \infty)$ encoding unpaired p -simplices identified in the Birth case. The running time of the three algorithms is $O(m^3)$, same as the original persistence algorithm given in [39]. Furthermore, it is possible to extend the algorithm in Chapter 4 so that it maintains the reduced matrices in time $O(m)$ per transposition of contiguous simplices in the ordered sequences; details can be found in Section 6.5. This is the method of choice for computing the vineyard of a pair of 1-parameter families of functions f and g , as they arise in applications considered in Section 6.6.

6.4 Correctness

We prove the correctness of the algorithms inductively, by considering one simplex at a time. For each index i , we consider the *actual* births, deaths, and pairs that occur in the sequences up to j_i , and the *computed* births, deaths, and pairs reported by the algorithm working on the simplices up to σ_i . To prove that the corresponding sets are the same at the end, for $i = m$, we show they are the same throughout, for all i . We do this in two steps, first proving that the algorithms are necessary and second that they are sufficient. In other words, we first prove that the computed information is correct and second that it is complete.

Preparation. We begin with a few preliminary observations. Recall that Tables 6.1 and 6.2 list the possible combinations of births and deaths under the simplifying assumption that each group has at most one change happening at any one time. This is indeed the situation if we add individual simplices to a growing complex; see Appendix B. We can therefore use the two tables in the correctness proof, but since we only consider the one function setting, we can further simplify and combine them into Table 6.3. We get K_i by

Case	$\ker j_i$	$\text{im } j_i$	$\text{cok } j_i$	$H(L_i)$	$H(K_i)$	σ_i
A	birth	death	—	—	death	$K - L$
B	—	birth	—	birth	birth	L
C	—	—	birth	—	birth	$K - L$
D	—	—	death	—	death	$K - L$
E	—	death	—	death	death	L
F	death	—	birth	death	birth	L

Table 6.3: In the one function setting there are six cases in which the addition of σ_i changes the kernel, the image, or the cokernel.

adding σ_i to K_{i-1} . If $\sigma_i \in L$ then $L_i = L_{i-1} \cup \{\sigma_i\}$ else $L_i = L_{i-1}$. In Cases B, E, F, the addition of σ_i changes the homology of L_{i-1} , which can only happen if $\sigma_i \in L$. In the remaining three cases, the addition of σ_i changes the homology of K_{i-1} but not that of L_{i-1} , which can only happen if $\sigma_i \in K - L$. Note also that the change in the homology of K_{i-1} is unambiguous in all cases, that is, σ_i is positive in Cases B, C, F, and negative in Cases A, D, E. We note that each death is paired with a unique birth but some births remain unpaired until the very end. It is convenient to rephrase the pairing condition in a form that is most directly useful in the argument below. We say a cycle *appears* in $\text{grp } j_l$ if the class it represents is born at that group, where $\text{grp} \in \{\ker, \text{im}, \text{cok}\}$ as usual. We note that the cycle might exist in the complex before it appears in the group.

DEATH LEMMA. Let $l < i$ be indices and z a cycle that appears first in $\text{grp } j_l$ and is zero in $\text{grp } j_i$. If there is no index $i' < i$ for which there is a cycle that first appears in $\text{grp } j_l$ and is zero in $\text{grp } j_{i'}$ then the class represented by z is born at $\text{grp } j_l$ and dies entering $\text{grp } j_i$.

Next we consider the incidence matrices used to compute the persistence diagrams of the sequences of kernels, images, and cokernels. To simplify language, we let $M[i]$ be the column of matrix M that corresponds to σ_i or, alternatively, the set of simplices whose corresponding rows have a one in this column. We will refer to it as *column i* of M and note that in some cases it is not the i -th column from the left, for example when $M = R_{\ker}$.

OBSERVATION. Recall that $R_f, R_g, R_{\ker}, R_{\text{im}}, R_{\text{cok}}$ are the reduced incidence matrices computed by the algorithms in Section 6.3.

- (i) If $R_g[i] = 0$ then $R_f[i] = 0$.
- (ii) $R_f[i] = 0$ iff $R_{\text{im}}[i] = 0$.
- (iii) If $\sigma_i \in L$ and $R_f[i] \neq 0$ then the lowest one in $R_{\text{im}}[i]$ corresponds to a simplex in L .
- (iv) The columns of R_{\ker} are all non-zero.
- (v) If $\sigma_i \in K - L$ then the lowest one in $R_{\ker}[i]$ corresponds to σ_i .
- (vi) If $\sigma_i \in L$, $R_g[i] \neq 0$, and $R_f[i] = 0$ then the lowest one in $R_{\ker}[i]$ corresponds to a simplex in $K - L$.

PROOF. Except for added zeros the columns of a simplex in L are the same in D_g and in D_f . This implies (i). The reordering of rows does not change the rank of the matrix. This implies (ii). By Observation (i), $R_g[i] \neq 0$ follows from $\sigma_i \in L$ and $R_f[i] \neq 0$. Since the columns of σ_i in D_g and D_{im} are the same, except for added zeros as before, the lowest one in $R_{\text{im}}[i]$ cannot be lower than that in $R_g[i]$. This implies (iii). The matrix V_{im} is upper triangular with a diagonal of ones. It thus has full rank and so does D_{\ker} which consists of a subset of the columns in V_{im} . This implies (iv). We get D_{\ker} from this subset of columns by reordering the rows, moving simplices in L up and simplices in $K - L$ down. The reordering maintains the relative order of the simplices in $K - L$. This implies (v).

Finally, we prove (vi) by contradiction, assuming the lowest one in $R_{\ker}[i]$ corresponds to a simplex $\sigma_l \in L$. Since $R_f[i] = 0$ we have $R_{\text{im}} = 0$. It follows that column i of V_{im} is part of D_{\ker} , after reordering the rows. Because of the upper triangular structure of V_{im} , the diagonal ones may be moved by the reordering but they are not cancelled in the reduction. By assumption, the lowest one in $R_{\ker}[i]$ corresponds to a simplex in L which can therefore only be σ_i , that is, $l = i$. But then $R_{\ker}[i]$ stores a cycle in L , a contradiction to $R_g[i] \neq 0$. \square

Inductive step. We assume inductively that the actual and the computed sets of births, deaths, and pairs are the same up to index $i - 1$. This is clearly true for $i - 1 = 0$, when all sets are empty.

Necessity. Using this as the induction hypothesis, we first show that the computed sets of births, deaths, and pairs are subsets of the corresponding actual sets up to index i .

Images. The algorithm for the images reports a birth for the new simplex iff $\sigma_i \in L$ and $R_g[i] = 0$. In this case, $V_g[i]$ stores a cycle that represents a new class in the homology

group of L_i as well as in $\text{im } j_i$. Hence, σ_i gives rise to an actual birth in the sequence of images.

The algorithm reports a death iff $R_f[i] \neq 0$ and the lowest one in $R_{\text{im}}[i]$ corresponds to a simplex $\sigma_l \in L$. The reduced column is a sum of boundaries in K_i ,

$$R_{\text{im}}[i] = \sum_{\sigma_k \in V_{\text{im}}[i]} D_{\text{im}}[k].$$

This cycle first appears in $\text{im } j_l$ and it is zero in K_i and therefore also in $\text{im } j_i$. Furthermore, there is no $i' < i$ for which there is a chain that first appears in $\text{im } j_l$ and is zero in $\text{im } j_{i'}$. Otherwise, $(\sigma_l, \sigma_{i'})$ would be a pair and the lowest one in $R_{\text{im}}[i']$ would correspond to σ_l , by inductive hypothesis. But then R_{im} could be reduced further, a contradiction. By the Death Lemma, σ_i gives rise to an actual death and (σ_l, σ_i) is an actual pair in the sequence of images.

Cokernels. The algorithm for the cokernels reports a birth iff $R_f[i] = 0$ and either $\sigma_i \in K - L$ or else $\sigma_i \in L$ and $R_g[i] \neq 0$. In this case, we have indeed a new class in the cokernel, namely the one represented by $V_f[i]$.

The algorithm reports a death iff $R_f[i] \neq 0$ and the lowest one in $R_{\text{im}}[i]$ corresponds to a simplex in $K - L$. By Observation (iii), this implies $\sigma_i \in K - L$. Letting σ_l correspond to the lowest one in $R_{\text{cok}}[i]$, the algorithm reports (σ_l, σ_i) as a pair. In this case, the reduced column is a sum of boundaries in K_i and cycles in L_i ,

$$R_{\text{cok}}[i] = \sum_{\sigma_k \in V_{\text{cok}}[i]} D_{\text{cok}}[k].$$

This cycle appears first in $\text{cok } j_l$ and it is zero in K_i and therefore also in $\text{cok } j_i$. Furthermore, there is no index $i' < i$ for which there is a chain that first appears in $\text{cok } j_l$ and is zero in $\text{cok } j_{i'}$. As before, we use induction and the fact that R_{cok} is reduced to prove this claim. By the Death Lemma, σ_i gives rise to an actual death and (σ_l, σ_i) is an actual pair in the sequence of cokernels.

Kernels. The algorithm for the kernels reports a birth iff $\sigma_i \in K - L$, $R_f[i] \neq 0$, and the lowest one in $R_{\text{im}}[i]$ corresponds to a simplex in L . In this case, $R_{\text{im}}[i]$ is a cycle in L and $V_{\text{im}}[i]$ is a chain whose boundary is this cycle. Furthermore, i is the smallest index for which such a chain exists, else we could use induction to show that R_{im} can be further reduced. Since σ_i belongs to $V_{\text{im}}[i]$, this chain does not belong to L . Hence, $R_{\text{im}}[i]$ represents a class in the kernel and σ_i gives rise to an actual birth.

The algorithm reports a death iff $\sigma_i \in L$, $R_g[i] \neq 0$, and $R_f[i] = 0$. By Observation (vi), the lowest one in $R_{\text{ker}}[i]$ corresponds to a simplex $\sigma_l \in K - L$. The algorithm reports (σ_l, σ_i) as a pair. To prove that there is an actual death, we recall that the columns in D_{ker} are cycles in K . We write each cycle as a sum of two chains, one in L and the other in its complement, $D_{\text{ker}}[k] = \sum \lambda_\ell + \sum \kappa_\ell$, where the λ_ℓ belong to L and the κ_ℓ belong to $K - L$. The two chains share their boundary, which we denote as $z_k = \partial \sum \lambda_\ell = \partial \sum \kappa_\ell$. Clearly, z_k is a cycle in L . Because it bounds the sum of the κ_ℓ , the cycle belongs to the

kernel, and because it bounds the sum of the λ_ℓ , the cycle is zero in L and therefore also zero in the kernel. Consider the cycle

$$z = \sum_{\sigma_k \in V_{\ker}[i]} z_k.$$

We claim that the class it represents in the kernel is born at $\ker j_l$. Indeed, if it were born earlier there would be $l' < l$ and a chain $c \in K_{l'}$ whose boundary is z . But then σ_l would be positive and by Observation (v) it would be the lowest one of its own column and not that of column i . Now, $V_{\ker}[i]$ provides the sum we need to finish the proof using the Death Lemma. As before, we use the induction hypothesis and the fact that R_{\ker} is reduced to conclude that there is no index $i' < i$ for which a cycle appears in $\ker j_l$ and is zero in $\ker j_{i'}$. Therefore, σ_i gives rise to an actual death and (σ_l, σ_i) is an actual pair in the sequence of kernels.

Sufficiency. We second show that the algorithms are complete, that is, the actual births, deaths, and pairs are subsets of the corresponding computed sets. Since there is a bijection between the deaths and the pairs, it suffices to prove the containments for the births and the deaths. We use Table 6.3 to do this by exhaustive case analysis.

Case 1 $\sigma_i \in K - L$.

Case 1.1 $R_f[i] = 0$. This is Case C in Table 6.3. The only change is a birth in the cokernel and this is correctly reported by the algorithms.

Case 1.2 $R_f[i] \neq 0$. Let σ_l correspond to the lowest one in $R_{\text{im}}[i]$.

Case 1.2.1 $\sigma_l \in K - L$. From the above analysis we know that this corresponds to a death in the cokernel. This is Case D in Table 6.3. There are no other changes and this is correctly reported by the algorithms.

Case 1.2.2 $\sigma_l \in L$. From the above analysis we know that this corresponds to a birth in the kernel and a death in the image. This is Case A in Table 6.3. There are no other changes and this is correctly reported by the algorithms.

Case 2 $\sigma_i \in L$.

Case 2.1 $R_f[i] = 0$.

Case 2.1.1 $R_g[i] \neq 0$. This is Case F in Table 6.3. There is a death in the kernel, a birth in the cokernel, and no change in the image, and this is correctly reported by the algorithms.

Case 2.1.2 $R_g[i] = 0$. This is Case B in Table 6.3. The only change is a birth in the image and this is correctly reported by the algorithms.

Case 2.2 $R_f[i] \neq 0$. By Observation (i), this implies $R_g[i] \neq 0$. This is Case E in Table 6.3. The only change is a death in the image which is correctly reported by the algorithms.

We conclude that the actual births and deaths are subsets of the computed births and deaths, and similarly that the actual pairs are a subset of the computed pairs.

We now have the containment relations in both directions which implies that corresponding sets of computed and actual births, deaths, and pairs are in fact the same. This concludes the proof that the algorithms in Section 6.3 correctly compute the persistence diagrams of the sequences of kernels, images, and cokernels.

6.5 Transpositions

We consider the actions necessary to maintain the matrix decompositions under the transposition of two contiguous simplices σ_i and σ_{i+1} . The motivation to study this operation is same as in previous chapters: construction of vineyards. We discuss a kernel persistence vineyard of interest in applications in the next section.

The maintenance of the $R = DU$ decomposition under such transpositions has been considered in Section 4.2 following [29] and that algorithm applies directly to the maintenance of $R_f = D_f V_f$ and $R_g = D_g V_g$. However, maintaining the other three decompositions is more difficult because D_{\ker} , D_{im} , and D_{cok} are not ordinary incidence matrices. The algorithm in Section 4.2 expresses an update in terms of pre- and post-multiplications by idempotent matrices. We observe that multiplying V by the same matrices as R maintains the equality $R = DV$ as well as $VU = I$.

Step 2, the image. Consider first the maintenance of $R_{\text{im}} = D_{\text{im}} V_{\text{im}}$ which is made difficult by ordering the rows and columns differently. We distinguish the case in which σ_i and σ_{i+1} both belong to L or to $K - L$ from the case in which one belongs to L and the other to $K - L$. In the former case, we let P' be the transposition matrix of the rows σ_i and σ_{i+1} , and in the latter case we set $P' = I$. Now we can follow the case analysis in Section 4.2 replacing the pre-multiplication of D_{im} and R_{im} by P with P' . The only other adjustment that we must make is in Case 1 in Section 4.2 in which both simplices σ_i and σ_{i+1} are positive. Namely, we do not need to check whether a collision is introduced in R_{im} by the transposition of rows (Case 1.1 in Section 4.2) if σ_i and σ_{i+1} do not belong to the same group, since in this case $P' = I$ and no rows transpose.

Step 3, the kernel. Consider second the maintenance of $R_{\ker} = D_{\ker} V_{\ker}$. Recall that the matrix D_{\ker} is obtained from V_{im} by keeping only the columns of positive simplices and reordering the rows. If both transposing simplices are positive in R_{\ker} then their columns in V_{\ker} do not change. Therefore, their columns in D_{\ker} and in R_{\ker} transpose as well as both their rows and columns in V_{\ker} . If as a result V_{\ker} ceases to be upper-triangular then we can perform an update similar to Case 2.1 in Section 4.2 by adding the column of σ_i to the column of σ_{i+1} before the transposition. Using the matrix S_i^{i+1} for this operation, we get

$$P' R_{\ker} S_i^{i+1} P = (P' D_{\ker} P)(P V_{\ker} S_i^{i+1} P).$$

Observe that this operation may render R_{\ker} non-reduced which can also happen as a result of a row transposition if σ_i and σ_{i+1} both belong to L or to $K - L$. If R_{\ker} becomes non-reduced we can fix it by multiplying again by S_i^{i+1} which gives

$$P' R_{\ker} S_i^{i+1} P S_i^{i+1} = (P' D_{\ker} P)(P V_{\ker} S_i^{i+1} P S_i^{i+1}).$$

This update may result in a Type 2 switch in the pairing, that is both transposing simplices are responsible for deaths of homology classes. If the two transposing simplices are both negative then there are no column exchanges in D_{\ker} , R_{\ker} , and V_{\ker} . However, the rows of D_{\ker} and R_{\ker} exchange if both σ_i and σ_{i+1} are in L or in $K - L$. The former case is easy because the transposing simplices do not contain the lowest one in their rows. However, in the latter case, similarly to Case 1.1 of Section 4.2, the matrix R_{\ker} may become non-reduced. Denote by σ_l and σ_k (both in L) the simplices in R_{\ker} that have in their columns the lowest ones in the positions of simplices σ_i and σ_{i+1} , respectively. We can reduce R_{\ker} by adding the preceding column to the succeeding column. We thus obtain

$$\begin{aligned} P' R_{\ker} S_l^k &= (P' D_{\ker})(V_{\ker} S_l^k) \text{ or} \\ P' R_{\ker} S_k^l &= (P' D_{\ker})(V_{\ker} S_k^l). \end{aligned}$$

If this update is necessary and k precedes l then we have a Type 1 switch in the pairing that is, both transposing simplices are responsible for births of homology classes. If σ_i is negative and σ_{i+1} is positive in R_f , that is, we are in Case 3 of Section 4.2, then regardless of whether the pairing switches or not, the columns of D_{\ker} remain the same. This observation follows from the update rule in case the pairing switches. This leaves the column of the negative simplex to be the sum of the transposing columns while leaving the column of the positive simplex intact. The rows of D_{\ker} transpose only if σ_i and σ_{i+1} are in L or in $K - L$. However, this transposition has no effect since $D_{\ker}[i, i+1] = V_{\ker}[i, i+1] = 1$ if and only if there is a switch in the pairing in R_{\ker} . As a result, we have

$$P' R_{\ker} = (P' D_{\ker}) V_{\ker}.$$

If $\sigma_i \in K - L$, $\sigma_{i+1} \in L$ and the pairing switches then the pair disappears from the kernel entirely: the column of D_{\ker} that used to represent $\sigma_{i+1} \in L$ now represents the simplex $\sigma_i \in K - L$. If σ_i is positive and σ_{i+1} is negative then no switches in the pairing may occur in R_{\ker} , so the only concern is the row transposition in D_{\ker} and R_{\ker} if σ_i and σ_{i+1} are both in L or in $K - L$. In the latter case, σ_i remains paired with itself and is unaffected by the row transposition. In the former case, σ_{i+1} can never be the lowest one in a column of R_{\ker} while σ_i can only be the lowest one in its own column, which cannot contain σ_{i+1} . Therefore, no changes in the pairing are possible, and

$$P' R_{\ker} = (P' D_{\ker}) V_{\ker}$$

is a proper decomposition. We observe that a Type 3 switch in the pairing never arises in the analysis of changes to the decomposition in Step 3 of the algorithm. This is to be

expected since the pairs in the kernel are always between simplices of the same dimension, and Step 3 affects neither image, nor cokernel pairs.

Step 4, the cokernel. Consider finally the maintenance of $R_{\text{cok}} = D_{\text{cok}} V_{\text{cok}}$. For reasons that will become apparent shortly, we switch to maintaining the decomposition $D_{\text{cok}} = R_{\text{cok}} U_{\text{cok}}$, where $U_{\text{cok}} = V_{\text{cok}}^{-1}$. We might as well since the matrix V_{cok} plays no role in the construction of the persistence diagrams. It is a curious property of the algorithm in Section 4.2 maintaining $D = RU$ (or equivalently $R = DV$) that the only time the columns of V corresponding to positive simplices can change is in the preprocessing step of Case 1, namely when σ_i and σ_{i+1} are both positive in R and $V[i, i+1] = 1$. Applied to the decomposition of D_f , we add column $V_g[i]$ to $V_g[i+1]$ whenever this happens. To offset the resulting change in D_{cok} , we add the row of σ_{i+1} to the row of σ_i in U_{cok} , noting that the two rows are not necessarily adjacent. We can update U_{cok} in linear time while the equivalent fix to V_{cok} would require a quadratic number of operations. The only remaining changes to D_{cok} are transpositions of rows and columns which we handle directly using the algorithm in Section 4.2. It follows that the maintenance of the decomposition $D_{\text{cok}} = R_{\text{cok}} U_{\text{cok}}$ takes linear time per operation.

6.6 Applications

In this section, we use kernels to measure local homology and images to approximate persistence diagrams of noisy functions specified on noisy domains. There are additional applications that are sufficiently straightforward that we can leave the details to the interested reader. One is the denoising of alpha-beta witness complexes as introduced in [5]; see also [76]. By considering maps from one complex to another, more tolerantly constructed complex we can preserve persistent features without accidentally introducing new ones. By varying the scale parameter, α , we thus get a persistence diagram that is less noisy than the diagram of the sequence of complexes for fixed tolerance parameter β . Another application is the computation of rank invariants for a doubly-filtered space $\mathbb{X}_{i,k}$ as considered in [18] and described in Section 4.4.2. Here we get a simple algorithm by encoding the ranks of the images of $H_p(\mathbb{X}_{i,k}) \rightarrow H_p(\mathbb{X}_{i',k'})$, for fixed $i < i'$ and variable $k < k'$, in a single persistence diagram of the images of the maps $H_p(\mathbb{X}_{i,k}) \rightarrow H_p(\mathbb{X}_{i',k})$. We can compute all such diagrams in time quartic in the number of simplices in the triangulation of \mathbb{X} .

6.6.1 Local Homology

We begin with the application of kernels to measuring the local homology of a space in \mathbb{R}^n at a point not necessarily in the space. Following earlier work, we assume that the space is not known other than indirectly through a finite set of points sampled in or near the space.

Measuring local homology. In the previous chapter we studied the reconstruction of a stratified space \mathbb{X} from a finite point sample U in \mathbb{R}^n ; we refer to that space as \mathbb{S} in this

section to avoid clashes of notation. Specifically, we used persistence to define a multi-scale version of the local homology of \mathbb{S} at a point $z \in \mathbb{R}^n$. Let \mathbb{S}_α be the sets of points at Euclidean distance at most α from \mathbb{S} , \mathbb{S}^α the set of points at distance at least α from \mathbb{S} , and B_r the closed ball of radius r centered at z . We expressed the homology within a fixed distance r of z in terms of the persistence diagram of the sequence

$$\begin{aligned} 0 \rightarrow \mathrm{H}(\mathbb{S}_\alpha \cap B_r) &\rightarrow \dots \rightarrow \mathrm{H}(B_r) \\ &\rightarrow \mathrm{H}(B_r, \mathbb{S}^\alpha \cap B_r) \rightarrow \dots \rightarrow 0, \end{aligned} \quad (6.1)$$

where α first goes up, from 0 to ∞ , and then down, from ∞ to 0. The first half of the sequence captures the development of the absolute homology of \mathbb{S}_α within B_r , and the second half captures the development of the relative homology of \mathbb{S}_α within the pair $(B_r, \partial B_r)$. Cycles that lie entirely inside the ball are captured twice, once in each half. Finally, we varied r from 0 to ∞ and this way obtain a vineyard that expresses the local homology of \mathbb{S} at z under the 2-parameter variation of α and r . We also proved relationships between this vineyard and the similarly defined vineyard of a finite set of points $U \subseteq \mathbb{R}^n$ sampled near \mathbb{S} .

In this chapter, we substitute a sequence of kernels for (6.1). Specifically, let $\mathbb{X} = B_r$, $\mathbb{Y} = \partial B_r$, and let $f : \mathbb{X} \rightarrow \mathbb{R}$, $g : \mathbb{Y} \rightarrow \mathbb{R}$ map each point to its Euclidean distance from \mathbb{S} . For each value α we write $\mathbb{X}_\alpha = f^{-1}(-\infty, \alpha]$, $\mathbb{Y}_\alpha = g^{-1}(-\infty, \alpha]$ and let $j_\alpha : \mathrm{H}(\mathbb{Y}_\alpha) \rightarrow \mathrm{H}(\mathbb{X}_\alpha)$ be the map induced by the inclusion $\mathbb{Y}_\alpha \subseteq \mathbb{X}_\alpha$. Assuming f and g are both tame we have a finite set of critical values and thus a finite sequence of kernels,

$$\mathrm{Ker}(g \rightarrow f) \quad : \quad \ker j_0 \rightarrow \ker j_1 \rightarrow \dots \rightarrow \ker j_m, \quad (6.2)$$

which traces the evolution of the relative homology classes in (6.1) that have a non-zero boundary in $\mathbb{Y} = \partial B_r$. The Stability Theorem in Section 6.2 implies that varying r from 0 to ∞ gives a vineyard. It tracks a homology class as long as the boundary of the ball with radius r intersects all its representatives. The interval of radii expresses relevant size information, namely how far away from z the class starts and ends. It is thus no longer necessary to include absolute homology classes in the measurement and the vineyard simplifies into a form that more closely reflects the shape of the space in the neighborhood of the point z ; see Figure 6.5. Relationships between the vineyard of \mathbb{S} and that of a finite point sample $U \subseteq \mathbb{R}^n$ similar to the Local Homology Inference and the Inverse LHI Theorems in Section 5.4 can be proved using the same methods. Details are omitted.

Computing local homology. Following the previous chapter, we use the Delaunay triangulations of U restricted to the ball B_r and to the ball without the interior of the power cell of the point z , $Z_0(r) = B_r - \mathrm{int} Z(r)$; see Section 5.5 for details. Let $K_\alpha = \mathrm{Del}(U|_{U_\alpha \cap B_r})$ and $L_\alpha = \mathrm{Del}(U|_{U_\alpha \cap Z_0(r)})$ be the Delaunay triangulations that are further restricted to the set of points at distance at most α from U . Let furthermore

$$\begin{aligned} i_\alpha & : \quad \mathrm{H}(U_\alpha \cap Z_0(r)) \rightarrow \mathrm{H}(U_\alpha \cap B_r); \\ l_\alpha & : \quad \mathrm{H}(L_\alpha) \rightarrow \mathrm{H}(K_\alpha) \end{aligned}$$

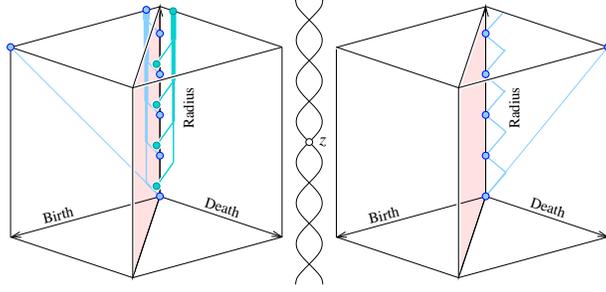


Figure 6.5: The vineyard of the chain of loops defined using the sequence (6.1) on the left and the sequence (6.2) of kernels on the right. To compare the two diagrams we see that the kernel sequence reflects the left half of the left vineyard across the diagonal plane and corrects for the removed absolute homology groups.

be the maps between the homology groups induced by inclusion. To justify the use of the restricted Delaunay triangulations, we need to show that the persistence diagrams of the kernels of these maps are the same. This requires that the following diagram is well-defined, commutative, and its vertical maps are isomorphisms whenever $\alpha \leq \alpha'$:

$$\begin{array}{ccc}
 \ker j_\alpha & \rightarrow & \ker j_{\alpha'} \\
 \downarrow i_\alpha & & \downarrow i_{\alpha'} \\
 \ker i_\alpha & \rightarrow & \ker i_{\alpha'} \\
 \uparrow h_\alpha & & \uparrow h_{\alpha'} \\
 \ker l_\alpha & \rightarrow & \ker l_{\alpha'}.
 \end{array}$$

Consider the following diagram whose maps are all induced by inclusions except for the lower vertical maps which are motivated by the Nerve Subdivision Lemma, see Sections 2.3 and 5.6.

$$\begin{array}{ccc}
 \mathrm{H}(U_\alpha \cap \partial B_r) & \xrightarrow{j_\alpha} & \mathrm{H}(U_\alpha \cap B_r) \\
 \downarrow & & \downarrow \\
 \mathrm{H}(U_\alpha \cap Z_0(r)) & \xrightarrow{i_\alpha} & \mathrm{H}(U_\alpha \cap B_r) \\
 \uparrow & & \uparrow \\
 \mathrm{H}(L_\alpha) & \xrightarrow{l_\alpha} & \mathrm{H}(K_\alpha).
 \end{array}$$

The top square commutes. The bottom square also commutes since the horizontal maps are induced by inclusion and the vertical map for L_α is the restriction of the one for K_α .

Therefore the kernel diagram is well-defined. Since the kernels are subgroups of the domains of their defining maps, the kernel diagram is a restriction of a diagram considered in Section 5.6. The analysis there implies that it commutes and its vertical maps are isomorphisms, as required.

The Stability Theorem of Section 6.2 implies that the kernel persistence diagrams change continuously with the radius of the restricted ball B_r . We construct the implied vineyard by maintaining the ordering of the simplices and the reduced matrices. However, unlike with relative homology in the previous chapter we cannot use excision to maintain different orderings of a static complex. We need to be able to handle insertion of simplices into the ordering when the power cell of the point z expands with r to include new simplices. Fortunately, the new simplices are paired amongst each other so that these updates can be done in linear time per insertion. Once a simplex is inserted, its position in the ordering can be described by a continuous function; see Sections 5.7.1 and 5.7.2. It therefore suffices to maintain the decompositions in the four steps of the algorithm under transpositions of contiguous simplices. Details on how to perform these operations are given in Section 6.5.

6.6.2 Noisy Domains

Persistent homology has proven to be well-suited for dealing with noisy functions. Indeed, the stability of persistence diagrams implies that the topological features of an unknown ideal function $f : \mathbb{X} \rightarrow \mathbb{R}$ can be approximately recovered knowing only a noisy approximation \tilde{f} of f . We claim that the persistence for images can be used to furthermore filter out the topological noise induced by the domain itself.

Stability. We assume an unknown ideal domain given as the zero sublevel set of the unknown ideal function $\tilde{h} : \mathbb{R}^n \rightarrow \mathbb{R}$, that is, $\tilde{\mathbb{X}} = \tilde{h}^{-1}(-\infty, 0]$. On this domain we consider another unknown ideal function but because we will vary the domain we assume it is defined on the entire ambient space, $\tilde{f} : \mathbb{R}^n \rightarrow \mathbb{R}$. Can we estimate the persistence diagram of the restriction $\tilde{f}|_{\tilde{\mathbb{X}}} : \tilde{\mathbb{X}} \rightarrow \mathbb{R}$ knowing only noisy approximations h, f of \tilde{h}, \tilde{f} ? We give an affirmative answer under mild requirements on the functions. To describe these requirements we use superscripts for sublevel sets of h and \tilde{h} and subscripts for sublevel sets of f and \tilde{f} , that is, $\mathbb{X}^u = h^{-1}(-\infty, u]$, $\mathbb{X}_a = f^{-1}(-\infty, a]$, $\mathbb{X}_a^u = \mathbb{X}_a \cap \mathbb{X}^u$ and similarly for $\tilde{\mathbb{X}}^u, \tilde{\mathbb{X}}_a, \tilde{\mathbb{X}}_a^u$. Writing $\varepsilon = \|h - \tilde{h}\|_\infty$ we require that \tilde{h} is smooth and the norm of its gradient is bounded away from 0 where this is relevant, that is, $\|\nabla \tilde{h}\| \geq \mu > 0$ on $\tilde{\mathbb{X}}^{2\varepsilon} - \tilde{\mathbb{X}}^{-2\varepsilon}$. Furthermore, we write $\delta = \|f - \tilde{f}\|_\infty$ and require that f is Lipschitz with constant κ , that is, $|f(x) - f(y)| \leq \kappa\|x - y\|$ for all $x, y \in \mathbb{R}^n$.

We note that the requirement on \tilde{h} implies a homotopy between the identity on $\tilde{\mathbb{X}}^{2\varepsilon}$ and a retraction ϱ from $\tilde{\mathbb{X}}^{2\varepsilon}$ to $\tilde{\mathbb{X}}^0$. To construct the homotopy we consider the integral lines of the vector field $-\nabla \tilde{h}$ starting at points $x \in \tilde{\mathbb{X}}^{2\varepsilon}$. Let $\varrho(x)$ be the first point on the curve starting at x that satisfies $\tilde{h}(\varrho(x)) = \tilde{h}(x) - 2\varepsilon$ or $\tilde{h}(\varrho(x)) \leq -2\varepsilon$. In words, ϱ moves the fringe outside the boundary of $\tilde{\mathbb{X}}$ to the fringe inside that boundary and it moves the fringe inside that boundary to the boundary of $\tilde{\mathbb{X}}^{-2\varepsilon}$. Since the gradient has norm no

smaller than $\mu > 0$ along the integral line, the retraction $\varrho : \tilde{\mathbb{X}}^{2\varepsilon} \rightarrow \tilde{\mathbb{X}}^{2\varepsilon}$ is well defined. By construction, there is a homotopy between the identity on $\tilde{\mathbb{X}}^{2\varepsilon}$ and ϱ that moves points by at most $2\varepsilon/\mu$. We are now ready to state our result.

STABILITY THEOREM FOR NOISY DOMAINS. Let functions $h, \tilde{h} : \mathbb{R}^n \rightarrow \mathbb{R}$ with $\varepsilon = \|h - \tilde{h}\|_\infty$, \tilde{h} smooth, and the norm of the gradient satisfying $\|\nabla \tilde{h}\| \geq \mu > 0$ on $\tilde{\mathbb{X}}^{2\varepsilon} - \tilde{\mathbb{X}}^{-2\varepsilon}$. Furthermore, let $f, \tilde{f} : \mathbb{R}^n \rightarrow \mathbb{R}$ with $\delta = \|f - \tilde{f}\|_\infty$ and f Lipschitz with constant κ . Then

$$d_B(\text{Dgm}(\tilde{f}|_{\tilde{\mathbb{X}}}), \text{Dgm}(\text{im } f|_{\mathbb{X}^{-\varepsilon}} \rightarrow f|_{\mathbb{X}^\varepsilon})) \leq 2\kappa\varepsilon/\mu + \delta$$

provided the restrictions of \tilde{f} to $\tilde{\mathbb{X}}$ and of f to $\tilde{\mathbb{X}}, \mathbb{X}^\varepsilon, \mathbb{X}^{-\varepsilon}$ are continuous and tame and there exists a triangulation of \mathbb{X}^ε in which $\mathbb{X}^{-\varepsilon}$ and $\tilde{\mathbb{X}}$ arise as subcomplexes.

PROOF. By the Stability Theorem for ordinary persistence we have bottleneck distance $d_B(\text{Dgm}(\tilde{f}|_{\tilde{\mathbb{X}}}), \text{Dgm}(f|_{\tilde{\mathbb{X}}})) \leq \delta$. It remains to show that the bottleneck distance between $\text{Dgm}(f|_{\tilde{\mathbb{X}}})$ and $\text{Dgm}(\text{im } f|_{\mathbb{X}^{-\varepsilon}} \rightarrow f|_{\mathbb{X}^\varepsilon})$ is bounded from above by $c = 2\kappa\varepsilon/\mu$.

Writing F_a for $H(\tilde{\mathbb{X}}_a^0)$, the diagram $\text{Dgm}(f|_{\tilde{\mathbb{X}}})$ is obtained from the sequence formed by the maps $F_a \rightarrow F_b$ induced by the inclusion $\tilde{\mathbb{X}}_a^0 \subseteq \tilde{\mathbb{X}}_b^0$ for all $a \leq b$. Similarly, writing J_a for the image of the map $H(\mathbb{X}_a^{-\varepsilon}) \rightarrow H(\mathbb{X}_a^\varepsilon)$, the diagram $\text{Dgm}(\text{im } f|_{\mathbb{X}^{-\varepsilon}} \rightarrow f|_{\mathbb{X}^\varepsilon})$ is obtained from the sequence of maps $J_a \rightarrow J_b$ again induced by inclusion and for all $a \leq b$. To adapt the proof of stability given in [25], we need to connect these two sequences by maps $\phi_a : F_{a-c} \rightarrow J_a$ and $\psi_a : J_{a-c} \rightarrow F_a$, for all $a \in \mathbb{R}$, in such a way that the diagram formed by the two sequences together with the new maps commutes. We construct the

$$\begin{array}{ccccccc}
 & & & & H(\tilde{\mathbb{X}}_{a-c}^0) & & \\
 & & & & \downarrow & & \\
 & & & & \swarrow & & \\
 H(\tilde{\mathbb{X}}_{a-2\varepsilon}^0) & \rightarrow & H(\mathbb{X}_a^{-\varepsilon}) & \rightarrow & H(\tilde{\mathbb{X}}_a^0) & \rightarrow & H(\mathbb{X}_a^\varepsilon) \rightarrow H(\tilde{\mathbb{X}}_a^{2\varepsilon}) \\
 & & & & \downarrow & & \swarrow \\
 & & & & H(\tilde{\mathbb{X}}_{a+c}^0) & &
 \end{array}$$

Figure 6.6: The diagram used to define the maps ϕ_a and ψ_a . The sequence $F_{a-c} \rightarrow F_a \rightarrow F_{a+c}$ is drawn vertically from top to bottom and J_a can be seen as the image of the composition of two horizontal maps. All maps except for the diagonal ones are induced by inclusion.

new maps using the homotopy between the identity and the retraction $\varrho : \tilde{\mathbb{X}}^{2\varepsilon} \rightarrow \tilde{\mathbb{X}}^{2\varepsilon}$. As mentioned earlier, the homotopy moves a point by at most $2\varepsilon/\mu$ and since f is Lipschitz with constant κ , ϱ maps $\tilde{\mathbb{X}}_{a-c}^0$ to $\tilde{\mathbb{X}}_a^{-2\varepsilon}$ which is included in $\mathbb{X}_a^{-\varepsilon}$. The induced homomorphism from $H(\tilde{\mathbb{X}}_{a-c}^0)$ to $H(\mathbb{X}_a^{-\varepsilon})$ composed with the induced homomorphism from $H(\mathbb{X}_a^{-\varepsilon})$ to $H(\mathbb{X}_a^\varepsilon)$ gives the map ϕ_a . As shown in Figure 6.6, ϕ_a connects the two sequences with a shift of $c = 2\kappa\varepsilon/\mu$. By a similar process, we construct the map ψ_a connecting the two

sequences of vector spaces in the other direction and again with a shift of c . Because ϱ is homotopic to the identity, the diagram formed by the two sequences and the maps ϕ_a and ψ_a commutes. The remainder of the proof can be adapted directly from [25]. \square

Diagram approximation. An interesting case of the above theorem arises when we consider a finite set of points, U , sampling an unknown shape, $\mathbb{S} \subseteq \mathbb{R}^n$. Let $\tilde{h} : \mathbb{R}^n \rightarrow \mathbb{R}$ be the distance function of \mathbb{S} , that is, $\tilde{h}(x) = \inf_{y \in \mathbb{S}} \|x - y\|$. Similarly, let $h : \mathbb{R}^n \rightarrow \mathbb{R}$ be the distance function of U and set ε to the Hausdorff distance between \mathbb{S} and U . For technical reasons we may have to replace \tilde{h} by a smooth approximation, for example obtained by convolution with an infinitesimally narrow Gaussian so that the assumptions in the theorem are satisfied. In this setting, the requirement that the norm of the gradient of \tilde{h} is bounded from below by μ is equivalent to the μ -reach of \mathbb{S} exceeding 4ε . Here we recall that the μ -reach as recently introduced in [20] is a notion of feature size that permits the treatment of non-smooth objects. Under this assumption, the Stability Theorem for Noisy Domains implies that it is possible to estimate the persistence diagram of a function \tilde{f} restricted to the 2ε -offset of \mathbb{S} knowing only a Lipschitz function f that approximates \tilde{f} and the point set U that samples \mathbb{S} .

To approximate the persistence diagram of \tilde{f} restricted to \mathbb{S} itself, we exploit the existence of an isotopy ι from $\mathbb{S}^{2\varepsilon}$ to an arbitrarily small offset \mathbb{S}^η of the shape such that the points move by less than $2\varepsilon/\mu$ during the deformation. The construction of ι is similar to the construction of the homotopy between the identity and ϱ described above. The isotopy implies that $\text{Dgm}(f|_{\mathbb{S}^{2\varepsilon}})$ equals $\text{Dgm}(f \circ \iota^{-1}|_{\mathbb{S}^\eta})$ which in turn is c -close to $\text{Dgm}f|_{\mathbb{S}^\eta}$ by stability. Hence, the latter diagram can also be estimated from f and U with an accuracy of $2c$. Perhaps surprisingly, $\text{Dgm}(f|_{\mathbb{S}^\eta})$ may not converge to $\text{Dgm}(f|_{\mathbb{S}})$ as η goes to 0; see [22] for examples of shapes \mathbb{S} that lack this convergence property. However, convergence holds for sufficiently regular spaces \mathbb{S} , such as smooth submanifolds or geometrically realized simplicial complexes. In these cases we can estimate $\text{Dgm}(f|_{\mathbb{S}})$ with precision $2c$.

In the more realistic case in which f is only known at a finite set of points, U , a valid approach replaces f by the function \bar{f} that is constant on the Voronoi cells of the points and coincides with f on U . While \bar{f} is not continuous, it is almost everywhere continuous in a way that does not disrupt the proof of stability. Furthermore, f and \bar{f} differ by at most $4\kappa\varepsilon$ on $\tilde{h}^{-1}(-\infty, 4\varepsilon]$. By the Stability Theorem for Noisy Domains, the persistence diagrams of the images defined by f and \bar{f} are close. The diagram for \bar{f} can be computed using the alpha shape filtration of U . We thus get a practical algorithm for estimating the persistence diagram of functions given only at a finite set of points.

6.7 Discussion

In this chapter, we consider persistent homology for sequences of kernels, images, and cokernels defined by a pair of topological spaces, $\mathbb{Y} \subseteq \mathbb{X}$, and two functions, $f : \mathbb{X} \rightarrow \mathbb{R}$ and $g : \mathbb{Y} \rightarrow \mathbb{R}$, with $f(y) \leq g(y)$ for every $y \in \mathbb{Y}$. Since g majorizes the restriction of f to \mathbb{Y} , its sublevel sets are contained in those of f , $\mathbb{Y}_a = g^{-1}(-\infty, a] \subseteq \mathbb{X}_a = f^{-1}(-\infty, a]$,

and we have homomorphisms $j_a : H(\mathbb{Y}_a) \rightarrow H(\mathbb{X}_a)$ induced by the inclusions. To see that persistent homology is well defined we just need to note that the diagrams

$$\begin{array}{ccc} H(\mathbb{X}_a) & \longrightarrow & H(\mathbb{X}_b) \\ \uparrow j_a & & \uparrow j_b \\ H(\mathbb{Y}_a) & \longrightarrow & H(\mathbb{Y}_b) \end{array}$$

commute for any $a \leq b$ and thus induce homomorphic maps $\ker j_a \rightarrow \ker j_b$, $\text{im } j_a \rightarrow \text{im } j_b$, and $\text{cok } j_a \rightarrow \text{cok } j_b$.

It is worth noting that the mapping cylinder construction described in Section 6.2 can be used to extend the framework from inclusion $\mathbb{Y} \subseteq \mathbb{X}$ to an arbitrary continuous map $j : \mathbb{Y} \rightarrow \mathbb{X}$. If we have two functions $f : \mathbb{X} \rightarrow \mathbb{R}$, and $g : \mathbb{Y} \rightarrow \mathbb{R}$, such that $j(\mathbb{Y}_a) \subseteq \mathbb{X}_a$ with the sublevel set of each space taken with respect to its own map, then the maps $j_a : \mathbb{Y}_a \rightarrow \mathbb{X}_a$ between the sublevel sets induce homomorphisms on homology groups just the same. Constructing a mapping cylinder $\mathbb{X}' = \mathbb{X} \cup \mathbb{Y} \times [0, 1]$ by identifying $(y, 0) \in \mathbb{Y} \times \{0\}$ with $j(y) \in \mathbb{X}$, it is easy to verify that the sequences of kernels, images, and cokernels induced by inclusion $\mathbb{Y}_a = \mathbb{Y}_a \times \{1\} \subseteq \mathbb{X}'_a$ give the same persistence pairing as the three sequences induced by the continuous maps j_a defined above.

One can also extend the framework described in this chapter from spaces $\mathbb{Y}_a \subseteq \mathbb{X}_a$ to pairs of spaces $(\mathbb{Y}_a, \mathbb{Y}_{a_0}) \subseteq (\mathbb{X}_a, \mathbb{X}_{a_0})$ using the cone construction exploited for the computation of extended persistence [26]. Indeed, observing that the relative homology groups $H(\mathbb{Y}_a, \mathbb{Y}_{a_0})$ are isomorphic to the homology groups of \mathbb{Y}_a with a cone on \mathbb{Y}_{a_0} rel the cone point, i.e. $H(\mathbb{Y}_a, \mathbb{Y}_{a_0}) \simeq H(\mathbb{Y}_a \cup C\mathbb{Y}_{a_0}, \omega)$, we can compute the persistence of kernels, images, and cokernels induced by the inclusions of pairs of spaces by using the algorithms of this chapter on corresponding cones.

The mapping cylinder construction and the cone construction can be combined to cope with arbitrary continuous maps between pairs of spaces rather than only inclusions.

The algorithms for computing the persistence diagrams of the sequences of kernels, images, and cokernels are variants of the classic Smith normal form algorithm; see [64, Chapter 1.11]. More directly, we build on the algorithm in [29] which reduces the entire incidence matrix at once, paying careful attention to the orderings of the rows and the columns which are the same and consistent with the ordering of the sublevel sets. The main difficulty in the new setting is that we have two orderings, one for K triangulating \mathbb{X} and the other for $L \subseteq K$ triangulating \mathbb{Y} . We cope using the mapping cylinder construction and matrices in which the rows are reordered so that L precedes $K - L$. The resulting algorithms run in time at most cubic in the size of K , same as the reduction algorithm in [29] as well as the classic Smith normal form algorithm for modulo-2 arithmetic. The fact that a simple reordering of the rows does the trick suggests that there may be other interesting pieces of information that can be extracted from reduced reordered matrices. Does the pairing defined by the lowest ones in a reduced incidence matrix in which columns and rows are ordered independently and arbitrarily have an intuitive interpretation that carries topological meaning?

Chapter 7

Discussion

The theory of persistent homology has emerged as an effective tool for handling the omniscalar nature of data. Its solid theoretical foundation provides for elegant and concise solutions to problems that arise in areas ranging from machine learning to scientific visualization to computational biology. Perhaps more importantly persistence provides a philosophical perspective and an example of how science could approach measuring and interpreting nature, how continuity of scales in physical phenomena and their interaction can be a blessing, not a burden.

This thesis branches out in the direction of dynamic treatment of persistence. Its original motivation is algorithmic analysis of time-varying data. The results obtained in the process turn out to be instrumental for the study of sampled stratified spaces, and in particular for applying the omniscalar data analysis paradigm introduced by persistence.

This dissertation begins with Chapter 2 recollecting the necessary theory. The Persistence Equivalence Theorem emphasizes the requirements for clear proofs of correctness of persistence pairing as illustrated in [11, 23]. The cubic worst-case example sheds light on the behavior of the persistence algorithm. The Nerve Subdivision Theorem provides a powerful tool for proving correctness of persistence pairing of filtered collections of convex sets and their nerves, as Chapter 5 demonstrates. The view of persistence computation as an RU -decomposition provides an insight whose utility we see in Chapters 4 and 6.

The thesis continues with the question of persistence-sensitive simplification presenting an algorithm for functions on 2-manifolds as well as limitations in higher dimensions. The extension of simplification to higher-dimensional spheres is an important open question.

Subsequently we introduce persistence vineyards by stacking up persistence diagrams within a homotopy of functions. The resulting curves allow one to keep track of homological features in time-varying data. Interesting questions that remain open are more efficient update of persistence pairing in particular the kind that relies only on the reduced matrix, and an algorithm for top-down computation of persistence diagrams.

In the following two chapters we are motivated by learning sampled stratified spaces. We initiate this study with the question of determining local homology of such a space at a point. However, this merely scratches the surface of the rich field that lies beneath. Ulti-

mately, in the spirit of omniscalar data analysis, one may desire to reconstruct a hierarchy of stratified spaces with accompanying information that allows the user to make decisions about meaningful scales within the hierarchy. This goal is ambitious and presents a longer-term research agenda. A more immediate and realistic question in this line of research is determining the local dimensions of the space at a point as suggested in Section 5.8.

The inference theorems in Chapters 2, 5, and 6 exploit the fact that distance functions of a space and its ε -approximation differ by at most ε . Unfortunately, while convenient for analysis, when dealing with high dimensional data, distance functions present more of a liability than advantage. Delaunay triangulation is exponential in the ambient dimension of the point set making it desirable to use some approximation of the distance function rather than computing it explicitly. The recent work using Vietoris-Rips and Witness complexes [5, 23, 76] can be seen as approximating distance functions for homology inference, especially as distilled in [19, 23]. Adopting the techniques for assessing local homology presented in this thesis to Vietoris-Rips [78] and Witness complexes is a practically important and promising course of study.

Finally, while theoretically simple and easy to analyze, the Hausdorff model of noise allowed by the notion of ε -approximation is often unrealistic. It would be far more desirable to be able to work with a point set drawn from a distribution that is uniformly random on a topological space and Gaussian in normal directions. This line of research was initiated by Dasgupta [30, 31] for the restricted case of 0-manifolds, i.e. the mixture of Gaussians. Carlsson et al. use various input preprocessing heuristics aimed to deal with outlier sensitivity of distance functions [17]. Most recently Niyogi, Smale, and Weinberger [66] propose a clean-up algorithm for making the homology inference problem in the Gaussian noise setting manageable with the distance function approach. Development of persistence-based omniscalar methods for dealing with Gaussian noise is an important research direction left open by this thesis.

Bibliography

- [1] PANKAJ K. AGARWAL, LARS ARGE, AND KE YI. I/O-Efficient Batched Union-Find and Its Applications to Terrain Analysis. *Proceedings of the Annual Symposium on Computational Geometry*, pages 167–176, New York, New York, June 2006.
- [2] PANKAJ K. AGARWAL, HERBERT EDELSBRUNNER, JOHN HARER, AND YUSU WANG. Extreme Elevation on a 2-Manifold. *Discrete and Computational Geometry*, **36**:553–572, 2006.
- [3] NINA AMENTA AND MARSHAL BERN. Surface Reconstruction by Voronoi Filtering. *Discrete and Computational Geometry*, **22**:481–504, 1999.
- [4] MICHELE D’AMICO, PATRIZIO FROSINI, AND CLAUDIA LANDI. Optimal Matching between Reduced Size Functions. Università degli studi di Modena e Reggio Emilia, Technical Report 35, 2003.
- [5] DOMINIQUE ATTALI, HERBERT EDELSBRUNNER, JOHN HARER, AND YURIY MILEYKO. Alpha-Beta Witness Complexes. *Proceedings of the Workshop on Algorithms and Data Structures*, pages 386–397, Halifax, Canada, Springer-Verlag, Lecture Notes in Computer Science **4619**, August 2007.
- [6] DOMINIQUE ATTALI, MARC GLISSE, SAMUEL HORNUS, FRANCIS LAZARUS, AND DMITRIY MOROZOV. Persistence-Sensitive Simplification of Functions on Surfaces in Linear Time. Manuscript, INRIA, 2008.
- [7] THOMAS F. BANCHOFF. Critical Points and Curvature for Embedded Polyhedra. *Journal of Differential Geometry*, **1**:245–256, 1967.
- [8] JULIEN BASCH. Kinetic Data Structures. Ph.D. Thesis, Stanford University, 1999.
- [9] JULIEN BASCH, LEONIDAS J. GUIBAS, AND JOHN HERSHBERGER. Data Structures for Mobile Data. *Journal of Algorithms*, **31**(1):1–28, 1999.
- [10] MIKHAIL BELKIN AND PARTHA NIYOGI. Laplacian Eigenmaps for Dimensionality Reduction and Data Representation. *Neural Computation*, **15**(6):1373–1396, 2003.

- [11] PAUL BENDICH, DAVID COHEN-STEINER, HERBERT EDELSBRUNNER, JOHN HARER, AND DMITRIY MOROZOV. Inferring Local Homology from Sampled Stratified Spaces. *Proceedings of the Symposium on Foundations of Computer Science*, pages 536–546, Providence, Rhode Island, October 2007.
- [12] PAUL BENDICH, JOHN HARER, AND HENRY KING. Persistent Intersection Homology for Stratified Spaces. Manuscript, Department of Mathematics, Duke University, Durham, North Carolina, 2007.
- [13] CHRISTOPH BREGLER AND STEPHEN M. OMOHUNDRO. Nonlinear Manifold Learning for Visual Speech Recognition. *Proceedings of the IEEE International Conference on Computer Vision*, pages 494–499, Cambridge, Massachusetts, June 1995.
- [14] PEER-TIMO BREMER, HERBERT EDELSBRUNNER, BERND HAMANN, AND VALERIO PASCUCCI. Topological Hierarchy for Functions on Triangulated Surfaces. *IEEE Transactions on Visualization and Computer Graphics*, **10**:385–396, 2004.
- [15] FRANCESCA CAGLIARI, MASSIMO FERRI, AND PAOLA POZZI. Size Functions from the Categorical Viewpoint. *Acta Applicandae Mathematicae*, **67**(3):225–235, 2001.
- [16] GUNNAR CARLSSON, ANNE COLLINS, LEONIDAS J. GUIBAS, AND AFRA ZOMORODIAN. Persistence Barcodes for Shapes. *International Journal of Shape Modeling*, **11**(2):149–187, 2005.
- [17] GUNNAR CARLSSON, TIGRAN ISHKHANOV, VIN DE SILVA, AND AFRA ZOMORODIAN. On the Local Behavior of Spaces of Natural Images. *International Journal of Computer Vision*, **76**:1–12, 2008.
- [18] GUNNAR CARLSSON AND AFRA ZOMORODIAN. The Theory of Multidimensional Persistence. *Proceedings of the Annual Symposium on Computational Geometry*, pages 184–193, Sedona, Arizona, June 2007.
- [19] FRÉDÉRIC CHAZAL, DAVID COHEN-STEINER, LEONIDAS J. GUIBAS, AND STEVE Y. OUDOT. The Stability of Persistence Diagrams Revisited. INRIA Saclay – Île-de-France, Technical Report 6568, 2008.
- [20] FRÉDÉRIC CHAZAL, DAVID COHEN-STEINER, AND ANDRÉ LIEUTIER. A Sampling Theory for Compact Sets in Euclidean Space. *Proceedings of the Annual Symposium on Computational Geometry*, pages 319–326, New York, New York, June 2006.
- [21] FRÉDÉRIC CHAZAL AND ANDRÉ LIEUTIER. Weak Feature Size and Persistent Homology: Computing Homology of Solids in \mathbb{R}^n from Noisy Data Samples. *Proceedings of the Annual Symposium on Computational Geometry*, pages 255–262, Pisa, Italy, June 2005.

- [22] FRÉDÉRIC CHAZAL AND ANDRÉ LIEUTIER. Stability and Computation of Topological Invariants of Solids in \mathbb{R}^n . *Discrete and Computational Geometry*, **37**:601–617, 2007.
- [23] FRÉDÉRIC CHAZAL AND STEVE Y. OUDOT. Towards Persistence-Based Reconstruction in Euclidean Spaces. *Proceedings of the Annual Symposium on Computational Geometry*, pages 232–241, College Park, Maryland, June 2008.
- [24] DAVID COHEN-STEINER AND HERBERT EDELSBRUNNER. Inequalities for the Curvature of Curves and Surfaces. *Proceedings of the Annual Symposium on Computational Geometry*, pages 272–277, Pisa, Italy, June 2005.
- [25] DAVID COHEN-STEINER, HERBERT EDELSBRUNNER, AND JOHN HARER. Stability of Persistence Diagrams. *Discrete and Computational Geometry*, **37**:103–120, 2007.
- [26] DAVID COHEN-STEINER, HERBERT EDELSBRUNNER, AND JOHN HARER. Extending Persistence Using Poincaré and Lefschetz Duality. *Foundations of Computational Mathematics*, to appear, 2008.
- [27] DAVID COHEN-STEINER, HERBERT EDELSBRUNNER, JOHN HARER, AND YURIY MILEYKO. Lipschitz Functions Have L_p -Stable Persistence. Manuscript, Department of Computer Science, Duke University, Durham, North Carolina, 2008.
- [28] DAVID COHEN-STEINER, HERBERT EDELSBRUNNER, JOHN HARER, AND DMITRIY MOROZOV. Persistent Homology for Kernels, Images, and Cokernels. Manuscript, Department of Computer Science, Duke University, Durham, North Carolina, 2008.
- [29] DAVID COHEN-STEINER, HERBERT EDELSBRUNNER, AND DMITRIY MOROZOV. Vines and Vineyards by Updating Persistence in Linear Time. *Proceedings of the Annual Symposium on Computational Geometry*, pages 119–126, New York, New York, June 2006.
- [30] SANJOY DASGUPTA. Learning Mixtures of Gaussians. *Proceedings of the Symposium on Foundations of Computer Science*, pages 634–644, New York, New York, October 1999.
- [31] SANJOY DASGUPTA. Learning Probability Distributions. Ph.D. Thesis, University of California at Berkeley, 2000.
- [32] MARY-LEE DEQUEANT, SEBASTIAN AHNERT, HERBERT EDELSBRUNNER, THOMAS M. A. FINK, EARL F. GLYNN, GAYE HATTEM, ANDRZEJ KUDLICKI, YURIY MILEYKO, JASON MORTON, ARCADY R. MUSHEGIAN, LIOR PACHTER, MAGA ROWICKA, ANNE J. SHIU, BERND STURMFELS, AND OLIVIER POURQUIE. Comparison of Pattern Detection Methods Applied to the Temporal

- Transcription Program of the Mouse Segmentation Clock. *PLoS One*, to appear, 2008.
- [33] TAMAL K. DEY. *Curve and Surface Reconstruction*. Cambridge University Press, England, 2006.
- [34] HERBERT EDELSBRUNNER. Weighted Alpha Shapes. Department of Computer Science, University of Illinois, Urbana, Illinois, Technical Report UIUCDCS-R-92-1760, 1992.
- [35] HERBERT EDELSBRUNNER. Union of Balls and Its Dual Shape. *Discrete and Computational Geometry*, **13**:415–440, 1995.
- [36] HERBERT EDELSBRUNNER. Surface Reconstruction by Wrapping Finite Point Sets in Space. *Discrete and Computational Geometry*, **32**:231–244, 2003.
- [37] HERBERT EDELSBRUNNER AND JOHN HARER. Persistent Homology — A Survey. *Twenty Years After*, J. E. Goodman, J. Pach, and R. Pollack, eds. AMS, to appear, 2008.
- [38] HERBERT EDELSBRUNNER, JOHN HARER, AND AFRA ZOMORODIAN. Hierarchical Morse-Smale Complexes for Piecewise Linear 2-Manifolds. *Discrete and Computational Geometry*, **30**:87–107, 2003.
- [39] HERBERT EDELSBRUNNER, DAVID LETSCHER, AND AFRA ZOMORODIAN. Topological Persistence and Simplification. *Discrete and Computational Geometry*, **28**:511–533, 2002.
- [40] HERBERT EDELSBRUNNER, DMITRIY MOROZOV, AND VALERIO PASCUCCI. Persistence-Sensitive Simplification of Functions on 2-Manifolds. *Proceedings of the Annual Symposium on Computational Geometry*, pages 127–134, New York, New York, June 2006.
- [41] HERBERT EDELSBRUNNER AND ERNST P. MÜCKE. Three-Dimensional Alpha Shapes. *ACM Transactions on Graphics*, **13**:43–72, 1994.
- [42] HERBERT EDELSBRUNNER AND NIMISH R. SHAH. Triangulating Topological Spaces. *International Journal of Computational Geometry and Applications*, **7**:365–378, 1997.
- [43] KASPAR FISCHER, BERND GÄRTNER, AND MARTIN KUTZ. Fast Smallest-Enclosing-Ball Computation in High Dimensions. *Proceedings of the European Symposium on Algorithms*, pages 630–641, Budapest, Hungary, September 2003.
- [44] PATRIZIO FROSINI. Discrete Computation of Size Functions. *Journal of Combinatorics, Information and System Sciences*, **17**(3–4):232–250, 1992.

- [45] BERND GÄRTNER. Fast and Robust Smallest Enclosing Balls. *Proceedings of the European Symposium on Algorithms*, pages 325–338, Prague, Czech Republic, July 1999.
- [46] NATASHA GELFAND, NILOY J. MITRA, LEONIDAS J. GUIBAS, AND HELMUT POTTMANN. Robust Global Alignment. *Proceedings of the Symposium on Geometric Processing*, pages 197–206, Vienna, Austria, July 2005.
- [47] MARK GORESKEY AND ROBERT D. MACPHERSON. Intersection Homology. *Topology*, **19**:135–162, 1982.
- [48] MARK GORESKEY AND ROBERT D. MACPHERSON. *Stratified Morse Theory*. Springer-Verlag, Heidelberg, Germany, 1988.
- [49] ATTILA GYULASSY, VIJAY NATARAJAN, VALERIO PASCUCCI, PEER-TIMO BREMER, AND BERND HAMANN. Topology-Based Simplification for Feature Extraction from 3D Scalar Fields. *Proceedings of the IEEE Visualization Conference*, pages 275–280, Minneapolis, Minnesota, October 2005.
- [50] NATHAN HABEGGER AND LESLIE D. SAPER. Intersection Cohomology of cs -Spaces and Zeeman’s Filtration. *Inventiones Mathematicae*, **105**:247–272, 1991.
- [51] GLORIA HARO, GREGORY RANDALL, AND GUILLERMO SAPIRO. Stratification Learning: Detecting Mixed Density and Dimensionality in High Dimensional Point Clouds. *Proceedings of the Advances in Neural Information Processing Systems*, pages 553–560, B. Schölkopf, J. Platt, and T. Hoffman, eds. Vancouver, Canada, 2007.
- [52] ALLEN HATCHER. *Algebraic Topology*. Cambridge University Press, England, 2001.
- [53] MASAKI HILAGA, YOSHIHISA SHINAGAWA, TAKU KOHMURA, AND TOSIYASU L. KUNII. Topology Matching for Fully Automatic Similarity Estimation of 3D Shapes. *Proceedings of the Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH)*, pages 203–212, Los Angeles, California, 2001.
- [54] HUGUES HOPPE, TONY DE ROSE, TOM DUCHAMP, JOHN McDONALD, AND WERNER STÜTZLE. Surface Reconstruction from Unorganized Points. *Proceedings of the Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH)*, pages 71–78, Chicago, Illinois, 1992.
- [55] DANIEL HUBER AND MARTIAL HEBERT. Fully Automatic Registration of Multiple 3D Data Sets. *Image and Vision Computing*, **21**:637–650, 2003.
- [56] BRUCE HUGHES AND SHMUEL WEINBERGER. Surgery and Stratified Spaces. *Surveys on Surgery Theory*, pages 311–342, S. Cappell, A. Ranicki, and J. Rosenberg, eds. Princeton University Press, 2000.

- [57] JEAN LERAY. Sur la Forme des Espaces Topologiques et sur les Points Fixes des Représentations. *Journal de Mathématiques Pures et Appliquées*, **24**:95–167, 1945.
- [58] ANDREY ANDREYEVICH MARKOV. Insolubility of the Problem of Homeomorphy. *Proceedings of the International Congress of Mathematicians*, pages 300–306, Edinburgh, Scotland, August 1958.
- [59] YUKIO MATSUMOTO. *An Introduction to Morse Theory*. American Mathematical Society, Providence, Rhode Island, 2001.
- [60] C. R. F. MAUNDER. *Algebraic Topology*. Dover, New York, New York, 1996.
- [61] JOHN MCCLEARY. *User’s Guide to Spectral Sequences*. Publish or Perish, 1985.
- [62] JOHN MILNOR. *Morse Theory*. Princeton University Press, Princeton, New Jersey, 1963.
- [63] JOHN MILNOR. Towards the Poincaré Conjecture and the Classification of 3-Manifolds. *Notices of the AMS*, **50**(10):1226–1233, 2003.
- [64] JAMES R. MUNKRES. *Elements of Algebraic Topology*. Addison Wesley Publishing Company, 1984.
- [65] PARTHA NIYOGI, STEPHEN SMALE, AND SHMUEL WEINBERGER. Finding the Homology of Submanifolds with High Confidence from Random Samples. *Discrete and Computational Geometry*, **39**:419–441, 2008.
- [66] PARTHA NIYOGI, STEPHEN SMALE, AND SHMUEL WEINBERGER. A Topological View of Unsupervised Learning from Noisy Data. Manuscript, 2008.
- [67] GRIGORI PERELMAN. The Entropy Formula for the Ricci Flow and Its Geometric Applications. arXiv:math.DG/0211159, 2002.
- [68] GRIGORI PERELMAN. Ricci Flow with Surgery on Three-Manifolds. arXiv:math.DG/0303109, 2003.
- [69] GRIGORI PERELMAN. Finite Extinction Time for the Solutions to the Ricci Flow on Certain Three-Manifolds. arXiv:math.DG/0307245, 2003.
- [70] HENRI POINCARÉ. Second Complément à L’Analysis Situs. *Proceedings of the London Mathematical Society*, **32**:277–308, 1900.
- [71] HENRI POINCARÉ. Cinquième Complément à L’Analysis Situs. *Rendiconti del Circolo Matematico di Palermo*, **18**:45–110, 1904.
- [72] GEORGES REEB. Sur les Points Singuliers d’une Forme de Pfaff Complément Intégrable ou d’une Fonction Numérique. *Comptes Rendus de L’Académie des Séances*, **222**:847–849, 1946.

- [73] YOUNG MIN RHEE, ERIC J. SORIN, GUHA JAYACHANDRAN, ERIK LINDAHL, AND VIJAY S. PANDE. Simulations of the Role of Water in the Protein-Folding Mechanism. *Proceedings of the National Academy of Sciences of the United States of America*, **101**:6456–6461, 2004.
- [74] VANESSA ROBINS. Toward Computing Homology from Finite Approximations. *Topology Proceedings*, **24**:503–532, 1999.
- [75] DANIEL RUSSEL AND LEONIDAS J. GUIBAS. Exploring Protein Folding Trajectories Using Geometric Spanners. *Proceedings of the Pacific Symposium on Biocomputing*, pages 40–51, World Scientific Publishing Co, Big Island, Hawaii, January 2005.
- [76] VIN DE SILVA AND GUNNAR CARLSSON. Topological Estimation Using Witness Complexes. *Proceedings of the Symposium on Point-Based Graphics*, pages 157–166, Morgan Kaufmann Publishers Inc, San Francisco, California, June 2004.
- [77] JOSHUA B. TENENBAUM, VIN DE SILVA, AND JOHN C. LANGFORD. A Global Geometric Framework for Nonlinear Dimensionality Reduction. *Science*, **290**(5500):2319–2323, 2000.
- [78] LEOPOLD VIETORIS. Über den Höheren Zusammenhang Kompakter Räume und eine Klasse von Zusammenhangstreuen Abbildungen. *Mathematische Annalen*, **97**:454–472, 1927.
- [79] YUSU WANG, PANKAJ K. AGARWAL, PAUL BROWN, HERBERT EDELSBRUNNER, AND JOHANNES RUDOLPH. Coarse and Reliable Geometric Alignment for Protein Docking. *Proceedings of the Pacific Symposium on Biocomputing*, pages 65–75, World Scientific Publishing Co, Big Island, Hawaii, January 2005.
- [80] JIANNIG WANG, MANUEL M. OLIVEIRA, AND ARIE E. KAUFMAN. Reconstructing Manifold and Non-Manifold Surfaces from Point Clouds. *Proceedings of the IEEE Visualization Conference*, page 53, Minneapolis, Minnesota, October 2005.
- [81] SHMUEL WEINBERGER. *The Topological Classification of Stratified Spaces*. University of Chicago Press, Chicago, Illinois, 1994.
- [82] AFRA ZOMORODIAN AND GUNNAR CARLSSON. Computing Persistent Homology. *Discrete and Computational Geometry*, **33**(2):249–274, 2005.

Index

- $(\alpha|r)$ -vineyard, 62
 - Theorem, 68
- ε -approximation, 21
- ε -simplification, 27
- Čech complex, 20

- alpha shape, 21

- barycentric subdivision, 7
- Betti number, 7
 - reduced, 8

- chain, 7
- contractible space, 10
- cs-space, 56

- deformation retraction, 9
- Delaunay triangulation, 21, 67
- Dimension-Index Lemma, 25
- distance function, 20

- extended persistence, 11

- filtration, 12

- homological critical value, 11
- homological feature size, 21
- homology group, 7
 - relative, 8
- Homology Inference Theorem, 21
- homotopic functions, 9
- homotopy equivalence, 9
- homotopy inverse, 9

- incidence matrix, 7
- index of critical vertex, 25
- Inverse LHI Theorem, 64

- knee, 40, 45

- local homology groups, 59
- Local Homology Inference Theorem, 64
- lower link, 24
- lower star, 24
- lower-star filtration, 25

- Mapping Cylinder Lemma, 83

- nerve, 17
 - realization, 19
- Nerve Subdivision Lemma, 19
- Nested-Disjoint Corollary, 29, 45

- Pairing Change Theorem, 45
- Pairing Uniqueness Lemma, 42
- persistence, 10
 - persistent homology, 10
 - persistent homology group, 11
- persistence diagram, 2
- Persistence Equivalence Theorem, 12
- piece, 56
- Power Cell Lemma, 69

- rank invariant, 52
- resolving radius, 63
- RU-decomposition, 41

- Same Dimension Lemma, 29
- simplex, 6
 - coface, 6
 - face, 6
 - link, 7
 - negative, 13
 - positive, 13
 - star, 7
- simplicial complex, 6
- Simplification Theorem for 2-Manifolds, 28

- skeleton, 7
- Stability Theorem, 15
 - Combinatorial Stability Theorem, 47
 - for Kernels, Images, Cokernels, 84
 - for Noisy Domains, 97
- stratification, 56
- stratum, 56
- sublevel set, 2
- Switch Lemma, 29

- tame function, 11
- triangulation, *see* simplicial complex

- underlying space, 6
- upper link, 24
- upper star, 24

- vine, 40
- vineyard, 40
- Voronoi cell, 21
 - weighted, 66
- Voronoi decomposition, 21
 - weighted, 66

Biography

Dmitriy Morozov was born in Snezhinsk, Russia on December 18, 1982. He received the Bachelor of Science degrees summa cum laude in Computer Science and Applied Mathematics from North Carolina State University in 2003. Later that year he started graduate studies in Computer Science at Duke University.

At NCSU Dmitriy received an Undergraduate Research Award in 2002. At Duke he was awarded a Graduate Fellowship in 2003. Dmitriy's research interests lie in Computational Topology and Geometry.

5. Paul Bendich, David Cohen-Steiner, Herbert Edelsbrunner, John Harer and Dmitriy Morozov. Inferring Local Homology from Sampled Stratified Spaces. In *Proceedings of the Forty-Eighth Annual Symposium on Foundations of Computer Science*, pages 536–546, 2007.
4. Herbert Edelsbrunner, Dmitriy Morozov and Valerio Pascucci. Persistence-Sensitive Simplification of Functions on 2-Manifolds. In *SCG '06: Proceedings of the Twenty-Second Annual Symposium on Computational Geometry*, pages 127–134, New York, NY, USA, 2006. ACM Press.
3. David Cohen-Steiner, Herbert Edelsbrunner and Dmitriy Morozov. Vines and Vineyards by Updating Persistence in Linear Time. In *SCG '06: Proceedings of the Twenty-Second Annual Symposium on Computational Geometry*, pages 119–126, New York, NY, USA, 2006. ACM Press.
2. Dmitriy Morozov. Persistence Algorithm Takes Cubic Time in the Worst Case. *Bio-Geometry News*, Department of Computer Science, Duke University, Durham, NC, 2005.
1. Erich Kaltofen, Dmitriy Morozov and George Yuhasz. Generic matrix multiplication and memory management in LinBox. In *ISSAC '05: Proceedings of the 2005 International Symposium on Symbolic and Algebraic Computation*, pages 216–223, New York, NY, USA, 2005. ACM Press.