

Mathematical
Surveys
and
Monographs

Volume 209



Persistence Theory: From Quiver Representations to Data Analysis

Steve Y. Oudot



American Mathematical Society

Licensed to AMS.

License or copyright restrictions may apply to redistribution; see <http://www.ams.org/publications/ebooks/terms>

Persistence Theory: From Quiver Representations to Data Analysis

**Mathematical
Surveys
and
Monographs**
Volume 209

Persistence Theory: From Quiver Representations to Data Analysis

Steve Y. Oudot



American Mathematical Society
Providence, Rhode Island

EDITORIAL COMMITTEE

Robert Guralnick
Michael A. Singer, Chair
Benjamin Sudakov
Constantin Teleman
Michael I. Weinstein

2010 *Mathematics Subject Classification*. Primary 62-07, 55U10, 16G20, 68U05.

For additional information and updates on this book, visit
www.ams.org/bookpages/surv-209

Library of Congress Cataloging-in-Publication Data

Oudot, Steve Y. 1979–

Persistence theory : from quiver representations to data analysis / Steve Y. Oudot.

pages cm. — (Mathematical surveys and monographs ; volume 209)

Includes bibliographical references and index.

ISBN 978-1-4704-2545-6 (alk. paper)

1. Algebraic topology. 2. Homology theory. I. Title.

QA612.O93 2015

514'.2—dc23

2015027235

Copying and reprinting. Individual readers of this publication, and nonprofit libraries acting for them, are permitted to make fair use of the material, such as to copy select pages for use in teaching or research. Permission is granted to quote brief passages from this publication in reviews, provided the customary acknowledgment of the source is given.

Republication, systematic copying, or multiple reproduction of any material in this publication is permitted only under license from the American Mathematical Society. Permissions to reuse portions of AMS publication content are handled by Copyright Clearance Center's RightsLink® service. For more information, please visit: <http://www.ams.org/rightslink>.

Send requests for translation rights and licensed reprints to reprint-permission@ams.org.

Excluded from these provisions is material for which the author holds copyright. In such cases, requests for permission to reuse or reprint material should be addressed directly to the author(s). Copyright ownership is indicated on the copyright page, or on the lower right-hand corner of the first page of each article within proceedings volumes.

© 2015 by the author. All rights reserved.

Printed in the United States of America.

∞ The paper used in this book is acid-free and falls within the guidelines established to ensure permanence and durability.

Visit the AMS home page at <http://www.ams.org/>

10 9 8 7 6 5 4 3 2 1 20 19 18 17 16 15

Contents

Preface	vii
Introduction	1
Part 1. Theoretical Foundations	11
Chapter 1. Algebraic Persistence	13
1. A quick walk through the theory of quiver representations	14
2. Persistence modules and interval decompositions	19
3. Persistence barcodes and diagrams	21
4. Extension to interval-indecomposable persistence modules	25
5. Discussion	26
Chapter 2. Topological Persistence	29
1. Topological constructions	29
2. Calculations	39
Chapter 3. Stability	49
1. Metrics	50
2. Proof of the stability part of the Isometry Theorem	54
3. Proof of the converse stability part of the Isometry Theorem	60
4. Discussion	61
Part 2. Applications	65
Chapter 4. Topological Inference	67
1. Inference using distance functions	71
2. From offsets to filtrations	78
3. From filtrations to simplicial filtrations	81
Chapter 5. Topological Inference 2.0	85
1. Simple geometric predicates	88
2. Linear size	92
3. Scaling up with the intrinsic dimensionality of the data	95
4. Side-by-side comparison	104
5. Natural images	106
6. Dealing with outliers	110
Chapter 6. Clustering	115
1. Contributions of persistence	117
2. ToMATo	118

3. Theoretical guarantees	122
4. Experimental results	126
5. Higher-dimensional structure	129
Chapter 7. Signatures for Metric Spaces	133
1. Simplicial filtrations for arbitrary metric spaces	138
2. Stability for finite metric spaces	140
3. Stability for totally bounded metric spaces	142
4. Signatures for metric spaces equipped with functions	146
5. Computations	146
Part 3. Perspectives	153
Chapter 8. New Trends in Topological Data Analysis	155
1. Optimized inference pipeline	156
2. Statistical topological data analysis	158
3. Topological data analysis and machine learning	160
Chapter 9. Further prospects on the theory	163
1. Persistence for other types of quivers	163
2. Stability for zigzags	165
3. Simplification and reconstruction	165
Appendix A. Introduction to Quiver Theory with a View Toward Persistence	167
1. Quivers	168
2. The category of quiver representations	168
3. Classification of quiver representations	170
4. Reflections	175
5. Proof of Gabriel's theorem: the general case	186
6. Beyond Gabriel's theorem	191
Bibliography	197
List of Figures	213
Index	217

Preface

It is in the early 2000's that persistence emerged as a new theory in the field of applied and computational topology. This happened mostly under the impulsion of two schools: the one led by H. Edelsbrunner and J. Harer at Duke University, the other led by G. Carlsson at Stanford University. After more than a decade of a steady development, the theory has now reached a somewhat stable state, and the community of researchers and practitioners gathered around it has grown in size from a handful of people to a couple hundred¹. In other words, persistence has become a mature research topic.

The existing books and surveys on the subject [48, 114, 115, 119, 141, 245] are largely built around the topological aspects of the theory, and for particular instances such as the persistent homology of the family of sublevel sets of a Morse function on a compact manifold. While this can be useful for developing intuition, it does create bias in how the subject is understood. A recent monograph [72] tries to correct this bias by focusing almost exclusively on the algebraic aspects of the theory, and in particular on the mathematical properties of persistence modules and of their diagrams.

The goal pursued in the present book is to put the algebraic part back into context², to give a broad view of the theory including also its topological and algorithmic aspects, and to elaborate on its connections to quiver theory on the one hand, to data analysis on the other hand. While the subject cannot be treated with the same level of detail as in [72], the book still describes and motivates the main concepts and ideas, and provides sufficient insights into the proofs so the reader can understand the mechanisms at work.

Throughout the exposition I will be focusing on the currently most stable instance of the theory: 1-dimensional persistence. Other instances, such as multi-dimensional persistence or persistence indexed over general partially ordered sets, are comparatively less well understood and will be mentioned in the last part of the book as directions for future research. The background material on quiver theory provided in Chapter 1 and Appendix A should help the reader understand the challenges associated with them.

Reading guidelines. There are three parts in the book. The first part (Chapters 1 through 3 and Appendix A) focuses on the theoretical foundations of persistence. The second part (Chapters 4 through 7) deals with a selected set of

¹As evidence of this, the Institute for Mathematics and its Applications at the University of Minnesota (<http://www.ima.umn.edu/>) was holding an annual thematic program on *Scientific and Engineering Applications of Algebraic Topology* in the academic year 2013-2014. Their first workshop, devoted to topological data analysis and persistence theory, gathered around 150 people on site, plus 300 simultaneous connections to the live broadcast.

²Let me mention a recent short survey [234] that pursues a similar goal.

applications. The third part (Chapters 8 and 9) talks about future prospects for both the theory and its applications. The document has been designed in the hope that it can provide something to everyone among our community, as well as to newcomers with potentially different backgrounds:

- Readers with a bias towards mathematical foundations and structure theorems will find the current state of knowledge about the decomposability of persistence modules in Chapter 1, and about the stability of their diagrams in Chapter 3. To those who are curious about the connections between persistence and quiver theory, I recommend reading Appendix A.
- Readers with a bias towards algorithms will find a survey of the methods used to compute persistence in Chapter 2, and a thorough treatment of the algorithmic aspects of the applications considered in Part 2.
- Practitioners in applied fields who want to learn about persistence in general will find a comprehensive yet still accessible exposition spanning all aspects of the theory, including its connections to some applications. To those I recommend the following walk through Part 1 of the document:
 - a) The general introduction,
 - b) Sections 1 through 3 of Chapter 1,
 - c) Sections 1.1 and 2.1 of Chapter 2,
 - d) Sections 1, 2.1 and 4 of Chapter 3.

Then, they can safely read Parts 2 and 3.

For the reader's convenience, the introduction of each chapter in Parts 1 and 2 mentions the prerequisites for reading the chapter and provides references to the relevant literature. As a general rule, I would recommend reading [115] or [142] prior to this book, as these references give quite accessible introductions to the field of applied and computational topology.

Acknowledgements. First of all, I want to express my gratitude towards the people who have contributed to shape persistence theory as we know it today. Among them, let me thank my co-authors, with whom I had an exciting time developing some of the ideas presented in this book: Jean-Daniel Boissonnat, Mickaël Buchet, Mathieu Carrière, Frédéric Chazal, David Cohen-Steiner, Vin de Silva, Jie Gao, Marc Glisse, Leonidas Guibas, Benoît Hudson, Clément Maria, Facundo Mémoli, Gary Miller, Maksim Ovsjanikov, Donald Sheehy, Primož Skraba, and Yue Wang.

Second, I want to thank the people who have helped me design the book and improve its content. Among them, my gratitude goes primarily to Michael Lesnick, for his careful reading of early versions of the manuscript and for his insightful comments that greatly helped improve Part 1 and Appendix A. I am also grateful to the anonymous referees, who provided me with valuable feedback on the flow of the book and on its readability. I also want to thank the people who have proof-read excerpts from the manuscript and helped me improve the content and exposition locally: Eddie Aamari, Jean-Daniel Boissonnat, Frédéric Chazal, Jérémy Cochoy, Paweł Dłotko, Marc Glisse, Bertrand Michel. Let me apologize in advance to those whose names I may have forgotten in this list.

Finally, I want to thank Sergei Gelfand, Christine Thivierge, and the American Mathematical Society for their interest in the book and for their support to finalize it.

Palaiseau, June 2015

Introduction

A picture being worth a thousand words, let us introduce our subject by showing a toy example coming from data analysis. Consider the data set of Figure 0.1, which is composed of 176 points sampled along 11 congruent letter-B shapes arranged into a letter A in the plane. When asking about the shape represented by this data set, one usually gets the answer: “It depends”, followed by a list of possible choices, the most common of which being “eleven B’s” and “one A”. To these choices one could arguably add a third obvious possibility: “176 points”. What differentiates these choices is the scale at which each one of them fits the data.

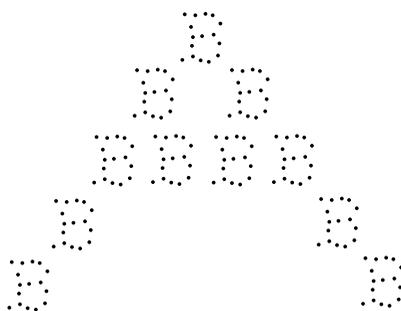


FIGURE 0.1. A planar point set with several underlying geometric structures at different scales.

Finding the ‘right’ scale(s) at which to process a given data set is a common problem faced across the data analysis literature. Most approaches simply ignore it and delegate the choice of scale to the user, who is then reduced to tuning some parameter blindly, usually by trial-and-error. Sometimes the parameter to tune does not even have a direct interpretation as a scale, which makes things even harder. This is where multiscale approaches distinguish themselves: by processing the data at all scales at once, they do not rely on a particular choice of scale. Their feedback gives the user a precise understanding of the relationship between the choice of input parameter and the output to be expected. Eventually, finding the ‘right’ scale to be used to produce the final output is still left to the user, however (s)he can now make an informed choice of parameter.

As an illustration, Figure 0.2 shows the result obtained by hierarchical agglomerative clustering on the aforementioned data set. The hierarchy reveals three relevant scales: at low levels (between 0 and 4), the clustering has one cluster per data point; at intermediate levels (between 8 and 12), the clustering has one cluster per letter B; at the highest level (above 16), there is only one cluster left, which spans the entire letter A.

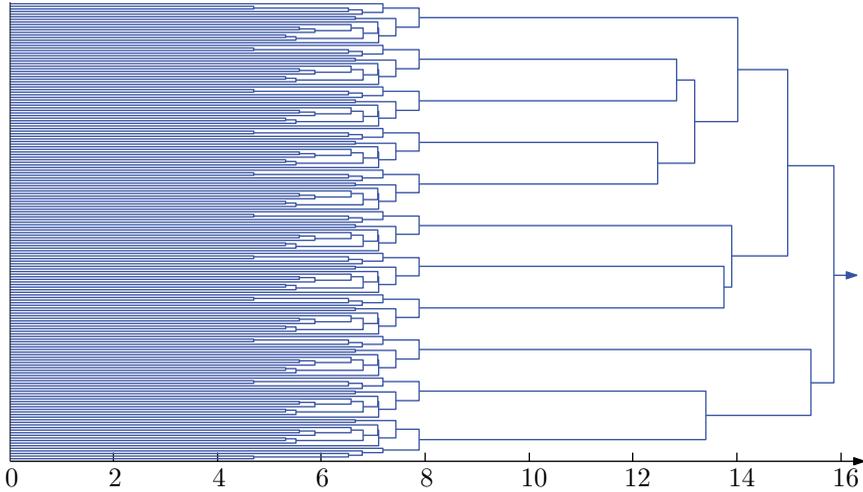


FIGURE 0.2. The hierarchy (also called *dendrogram*) produced by *single-linkage* clustering on the data set of Figure 0.1.

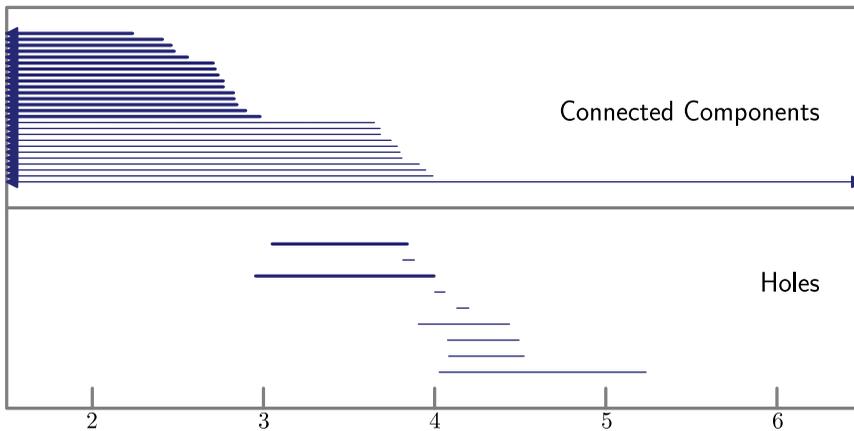


FIGURE 0.3. The barcode produced by persistence on the data set of Figure 0.1. The abscissa line represents the geometric scale with a logarithmic graduation. Left (resp. right) arrows mark left- (resp. right-) infinite intervals, while thin (resp. bold) bars mark intervals with multiplicity 1 (resp. 11).

Persistence produces the same hierarchy but uses a simplified representation for it, shown in the upper half of Figure 0.3. This representation forgets the actual merge pattern between the clusters. When two clusters are merged, they no longer produce a new cluster corresponding to their union in the hierarchy. Instead, one of them ceases to be treated as an independent cluster, to the benefit of the other. The choice of the winner is arbitrary in this case, however in general it is driven by a principle called the *elder rule*, which will be illustrated in the upcoming Figure 0.5. The resulting collection of horizontal bars is called a *persistence barcode*. Each bar is associated to a single data point and represents its *persistence* as an independent

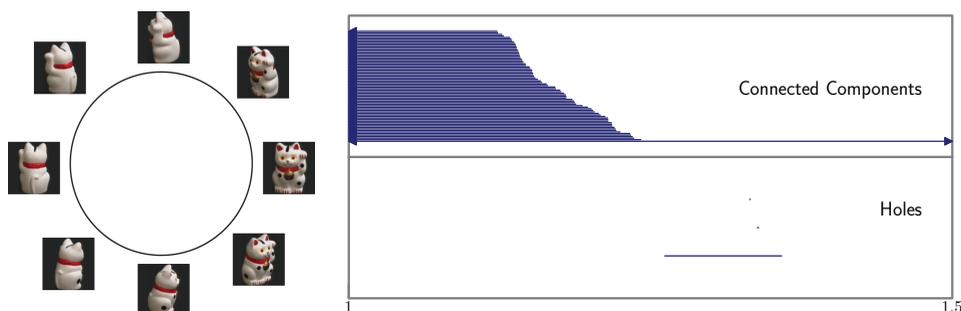


FIGURE 0.4. Left: a collection of 72 color images of size 128×128 pixels coming from the Columbia Object Image Library [203]. The images were obtained by taking pictures of an object while rotating around it. Each image gives a data point in 49 152 dimensions, where by construction the data set lies near some simple closed curve. Right: the corresponding barcode produced by persistence, where the abscissa line represents the geometric scale with a logarithmic graduation, and where left (resp. right) arrows mark left- (resp. right-) infinite intervals.

— Pictures used in the left part of the figure are courtesy of the Computer Vision Laboratory at Columbia University.

cluster. Although weaker than the full hierarchical representation, the barcode is still informative enough to allow for an interpretation. In the present example, the 11 bars with multiplicity 1 come from the 11 B's merging into a single A around scale $2^4 = 16$. Before that, the 15 bars with multiplicity 11 come from each letter B having 16 points that get merged into a single cluster around scale $2^3 = 8$. It takes a bit of time to get used to this kind of representation, in which the actual hierarchy (who is merged with whom) is lost. Nevertheless, this is the price to pay for more stability and generality.

Persistence is indeed able to produce such barcodes for higher-dimensional topological features as well. For instance, the bottom half of Figure 0.3 shows a barcode encoding the lifespans of 'holes' across scales in the data set of Figure 0.1. To understand what is meant by this, imagine each data point being replaced by a ball of radius r at scale r . Persistence detects the holes in the resulting union of balls at every scale, and tracks their persistence across scales. Each bar in the resulting barcode corresponds to a particular hole, and it encodes its lifespan in the growing family of balls. The same can be done for voids in higher dimensions. In the example of Figure 0.3, the 2 bars with multiplicity 11 appearing at lower scales come from each letter B having 2 holes, while the long bar with multiplicity 1 appearing at larger scales comes from the letter A having a single hole. The rest of the bars indicate holes created at intermediate steps in the ball growing process, for instance in places where B's are arranged into a triangle.

Being able to detect the presence of topological features of arbitrary dimensions in data, and to represent these features as a barcode whatever their dimension, is what makes persistence an interesting tool for data visualization and analysis, and a nice complement to more classical techniques such as clustering or dimensionality reduction [183]. Besides, being able to do so in high dimensions and in a robust way,

as illustrated in Figure 0.4, is an asset for applications. It is also an algorithmic challenge, as dealing with high-dimensional data requires to develop a computing machinery that scales up reasonably with the ambient dimension.

Persistence in a nutshell. The theory works at two different levels: topological, and algebraic. At the topological level, it takes as input a sequence of nested topological spaces, called a *filtration*:

$$(0.1) \quad X_1 \subseteq X_2 \subseteq \cdots \subseteq X_n.$$

Such sequences come typically from taking excursion sets (sublevel sets or superlevel sets) of real-valued functions. For instance, in the example of Figure 0.3, the filtration is composed of the sublevel sets of the distance to the point cloud, the r -sublevel set being the same as the union of balls of same radius r about the data points, for every $r \geq 0$. Here already comes a difficulty: in (0.1) we are using a finite sequence, whereas the sublevel sets of a function form a continuous 1-parameter family. While algorithms only work with finite sequences for obvious reasons, the theory is stated for general 1-parameter families. The connection between discrete and continuous families is not obvious in general, and determining the precise conditions to be put on a continuous family so that it behaves ‘like’ a discrete family has been the subject of much investigation, as will be reflected in the following chapters.

Given a sequence like (0.1), we want not only to compute the topological structure of each space X_i separately, but also to understand how topological features persist across the family. The right tool to do this is homology over a field, which turns (0.1) into a sequence of vector spaces (the homology groups $H_*(X_i)$) connected by linear maps (induced by the inclusions $X_i \hookrightarrow X_{i+1}$):

$$(0.2) \quad H_*(X_1) \longrightarrow H_*(X_2) \longrightarrow \cdots \longrightarrow H_*(X_n).$$

Such a sequence is called a *persistence module*. Thus we move from the topological level to the algebraic level, where our initial problem becomes the one of finding bases for the vector spaces $H_*(X_i)$ that are ‘compatible’ with the maps in (0.2). Roughly speaking, being compatible means that for any indices i, j with $1 \leq i \leq j \leq n$, the composition

$$H_*(X_i) \longrightarrow H_*(X_{i+1}) \longrightarrow \cdots \longrightarrow H_*(X_{j-1}) \longrightarrow H_*(X_j)$$

has a (rectangular) diagonal matrix in the bases of $H_*(X_i)$ and $H_*(X_j)$. Then, every basis element can be tracked across the sequence (0.2), and its *birth time* b and *death time* d defined respectively as the first and last indices at which it is part of the current basis. At the topological level, this basis element corresponds to some feature (connected component, hole, void, etc.) appearing in X_b and disappearing in X_{d+1} . Its lifespan is encoded as an interval $[b, d]$ in the persistence barcode³.

The very existence of compatible bases is known from basic linear algebra when $n \leq 2$ and from the structure theorem for finitely generated modules over a principal ideal domain when n is arbitrary (but finite) and the vector spaces $H_*(X_i)$ have finite dimensions. Beyond these simple cases, e.g. when the index set is infinite or when the spaces are infinite-dimensional, the existence of compatible bases is not always assured, and when it is, this is thanks to powerful decomposition theorems from quiver representation theory. Indeed, in its algebraic formulation, persistence

³Some authors rather use the equivalent notation $[b, d + 1)$ for the interval. We will come back to this in Chapter 1.

is closely tied to quiver theory. Their relationship will be stressed in the following chapters, but for now let us say that *quiver* is just another name for (multi-)graph, and that a *representation* is a realization of a quiver as a diagram of vector spaces and linear maps. Thus, (0.2) is a representation of the quiver

$$\bullet \longrightarrow \bullet \longrightarrow \cdots \longrightarrow \bullet$$

1 2 n

Computing a compatible basis is possible when the filtration is *simplicial*, that is, when it is a finite sequence of nested simplicial complexes. It turns out that computing the barcode in this special case is hardly more complicated than computing the homology of the last complex in the sequence, as the standard matrix reduction algorithm for computing homology can be adapted to work with the filtration order. Once again we are back to the question of relating the barcodes of finite (simplicial) filtrations to the ones of more general filtrations. This can be done via the stability properties of these objects.

Stability. The stability of persistence barcodes is stated for an alternate representation called *persistence diagrams*. In this representation, each interval⁴ b, d is viewed as a point (b, d) in the plane, so a barcode becomes a planar multiset. The persistence of a topological feature, as measured by the length $(d - b)$ of the corresponding barcode interval b, d , is now measured by the vertical distance of the corresponding diagram point (b, d) to the diagonal $y = x$. For instance, Figure 0.5 shows the persistence diagrams associated to the filtrations of (the sublevel-sets of) two functions $\mathbb{R} \rightarrow \mathbb{R}$: a smooth function f , and a piecewise linear approximation f' . As can be seen, the proximity between f and f' implies the proximity between their diagrams $\text{dgm}(f)$ and $\text{dgm}(f')$. This empirical observation is formalized in the following inequality, where $\|\cdot\|_\infty$ denotes the supremum norm and d_b denotes the so-called *bottleneck distance* between diagrams:

$$(0.3) \quad d_b(\text{dgm}(f), \text{dgm}(f')) \leq \|f - f'\|_\infty.$$

Roughly speaking, the bottleneck distance provides a one-to-one matching between the diagram points corresponding to highly persistent topological features of f and f' , the topological features with low persistence being regarded as noise and their corresponding diagram points being matched to the nearby diagonal.

Stability, as stated in (0.3) and illustrated in Figure 0.5, is an important property of persistence diagrams for applications, since it guarantees the consistency of the computed results. For instance, it ensures that the persistence diagram of an unknown function can be faithfully approximated from the one of a known approximation. Or, that reliable information about the topology of an unknown geometric object can be retrieved from a noisy sampling under some reasonable noise model.

The proof of the stability result works at the algebraic level directly. For this it introduces a measure of proximity between persistence modules, called the *interleaving distance*, which derives naturally from the proximity between the functions the modules originate from (when such functions exist). In this metric, the stability result becomes in fact an isometry theorem, so that comparing persistence modules is basically the same as comparing their diagrams. From there on, persistence diagrams can be used as signatures for all kinds of objects from which persistence modules are derived, including functions but not only.

⁴We are omitting the brackets to indicate that the interval can be indifferently open, closed, or half-open.

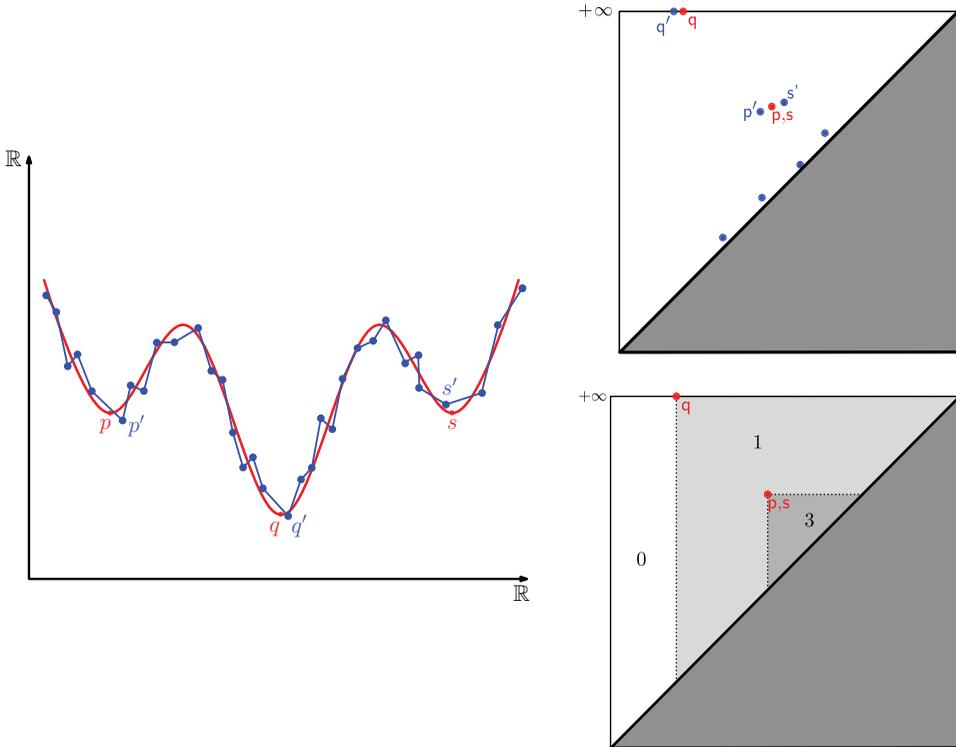


FIGURE 0.5. Left: A smooth function $f : \mathbb{R} \rightarrow \mathbb{R}$ (red) and a piecewise linear approximation f' (blue). Top-right: Superimposition of the persistence diagrams of f (red) and f' (blue). Every red diagram point (b, d) corresponds to some local minimum of f creating an independent connected component in the sublevel set of f at time b , and merging it into the component of some lower minimum at time d , as per the *elder rule*. Idem for blue points and f' . Bottom-right: The size function corresponding to the persistence diagram of f .

The isometry theorem is the cornerstone of the current theory, and its main asset for applications. It is also what makes persistence stand out of classical quiver theory.

Connections to other theories. As mentioned previously, there is a deep connection between the algebraic level of persistence and quiver theory. Meanwhile, the topological level has strong bonds with Morse theory:

- In the special case where the input filtration is given by the sublevel sets of a Morse function f , i.e. a C^∞ -continuous real-valued function with non-degenerate critical points such that all the critical values are distinct, Morse theory describes when and how the topology of the sublevel sets of f changes in the filtration [195, theorems 3.1 and 3.2], thus providing a complete characterization of its persistence diagram. Persistence generalizes this analysis beyond the setting of Morse theory, to cases where the

function f may not be differentiable nor even continuous, and where its domain may not be a smooth manifold nor a manifold at all.

- In a similar way, persistence for simplicial filtrations is related to the discrete version of Morse theory [130]. There are indeed filtered counterparts to the discrete gradient fields and to their associated discrete Morse complexes. These are defined on simplicial filtrations rather than on single simplicial complexes, with the expected property that the persistent homology of the filtered Morse complexes is the same as the one of the filtrations they come from. This connection has been exploited in various ways, for instance to speed up the persistence computation [198].
- Finally, the 0-dimensional aspects of persistence are related to Morse theory in a particular way [55, 213]. Given a Morse function f , the hierarchy on the local minima of f produced by persistence from the family of its sublevel sets is equivalent to the *join tree* of f . Similarly, the hierarchy on the local maxima of f produced from its superlevel sets is equivalent to the *split tree* of f . Once merged together, these two trees form the *contour tree* of f , which is the loop-free version of the *Reeb graph* and is equal to it when the domain of f is both connected and simply connected. There are also some relations between the 1-dimensional persistence of f and the loops of its Reeb graph [86, 108].

As we saw earlier, the connection between the topological and the algebraic levels of persistence happens through the use of homology, which turns sequences of topological spaces into sequences of vector spaces. Using the metaphor of a space changing over time to describe each sequence, we can view persistence as a generalization of classical homology theory to the study of time-evolving spaces. In this metaphor, persistence modules such as (0.2) are the time-dependent analogues of the homology groups, and their barcodes are the time-dependent analogues of the Betti numbers. Although somewhat restrictive, this view of the theory is convenient for interpretation.

Persistence is also a generalization of *size theory* [97, 131], whose concern is with the quantity

$$\text{rank } H_0(X_i) \rightarrow H_0(X_j)$$

defined for all pairs (i, j) such that $1 \leq i \leq j \leq n$, and called the *size function* of the filtration (0.1). The value of the size function at (i, j) measures the number of connected components of X_i that are still disconnected in X_j . The level sets of this function look like staircases in the plane, whose upper-left corners are the points recorded in the 0-dimensional part of the persistence diagram—see Figure 0.5 for an illustration. The stability result (0.3) appeared in size theory prior to the development of persistence, however in a form restricted to 0-dimensional homology.

Finally, the algorithmic aspects of persistence have close connections to *spectral sequences* [81]. Roughly speaking, the spectral sequence algorithm outputs the same barcode as the matrix reduction algorithm, albeit in a different order.

These connections at multiple levels bear witness to the richness of persistence as a theory.

Applications. This richness is also reflected in the diversity of the applications, whose list has been ever growing since the early developments of the theory. The following excerpt⁵ illustrates the variety of the topics addressed:

- analysis of random, modular and non-modular scale-free networks and networks with exponential connectivity distribution [158],
- analysis of social and spatial networks, including neurons, genes, online messages, air passengers, Twitter, face-to-face contact, co-authorship [210],
- coverage and hole detection in wireless sensor fields [98, 136],
- multiple hypothesis tracking on urban vehicular data [23],
- analysis of the statistics of high-contrast image patches [54],
- image segmentation [70, 209],
- 1d signal denoising [212],
- 3d shape classification [58],
- clustering of protein conformations [70],
- measurement of protein compressibility [135],
- classification of hepatic lesions [1],
- identification of breast cancer subtypes [205],
- analysis of activity patterns in the primary visual cortex [224],
- discrimination of electroencephalogram signals recorded before and during epileptic seizures [237],
- analysis of 2d cortical thickness data [82],
- statistical analysis of orthodontic data [134, 155],
- measurement of structural changes during lipid vesicle fusion [169],
- characterization of the frequency and scale of lateral gene transfer in pathogenic bacteria [125],
- pattern detection in gene expression data [105],
- study of plant root systems [115, §IX.4],
- study of the cosmic web and its filamentary structure [226, 227],
- analysis of force networks in granular matter [171],
- analysis of regimes in dynamical systems [25].

In most of these applications, the use of persistence resulted in the definition of new descriptors for the considered data, which revealed previously hidden structural information and allowed the authors to draw original conclusions.

Contents. There are three parts in the book. The first part focuses on the theoretical foundations of persistence. It gives a broad view of the theory, including its algebraic, topological, and algorithmic aspects. It is divided into three chapters and an appendix:

- Chapter 1 introduces the algebraic aspects through the lense of quiver theory. It tries to show both the heritage of quiver theory and the novelty brought in by persistence in its algebraic formulation. It is supplemented with Appendix A, which gives a formal introduction to quiver representation theory and highlights its connections to persistence. Concepts such as persistence module, zigzag module, module homomorphism, interval decomposition, persistence diagram, quiver, representation, are defined in these two chapters.

⁵Much of the list was provided by F. Chazal, F. Lecci and B. Michel, who recently took an inventory of existing applications of persistence.

- Chapter 2 introduces the topological and algorithmic aspects of persistence theory. It first reviews the topological constructions that are most commonly used in practice to derive persistence modules. It then focuses on the algorithms designed to compute persistence from filtrations: the original algorithm, described in some detail, then a high-level review of its variants and extensions. Concepts such as filtration, zigzag, pyramid, persistent (co-)homology, are defined in this chapter.
- Chapter 3 is entirely devoted to the stability of persistence diagrams, in particular to the statement and proof of the *Isometry Theorem*, which is the central piece of the theory. After introducing and motivating the measures of proximity between persistence modules and between their diagrams which are used in the statement of the theorem, it develops the main ideas behind the proof and discusses the origins and significance of the result. Concepts such as interleaving distance, bottleneck distance and matching, snapping principle, module interpolation, are defined in this chapter.

The second part of the document deals with applications of persistence. Rather than trying to address all the topics covered in the aforementioned list, in a broad and shallow survey, it narrows the focus down to a few selected problems and analyzes in depth the contribution of persistence to the state of the art. Some of these problems have had a lasting influence on the development of the theory. The exposition is divided into four chapters:

- Chapters 4 and 5 introduce the problem of inferring the topology of a geometric object from a finite point sample, which was and continues to be one of the main motivations for the development of the theory. The general approach to the problem is presented in Chapter 4, along with some theoretical guarantees on the quality of the output. Algorithmic aspects are addressed in Chapter 5, which introduces recent techniques to optimize the running time and memory usage, improve the signal-to-noise ratio in the output, and handle a larger variety of input data.
- Chapter 6 focuses more specifically on the 0-dimensional version of topological inference, also known as clustering. It formalizes the connection between persistence and hierarchical clustering, which we saw earlier. It also draws a connection to mode seeking and demonstrates how persistence can be used to stabilize previously unstable hill-climbing methods. Finally, it addresses the question of inferring higher-dimensional structure, to learn about the composition of each individual cluster as well as about their interconnectivity in the ambient space. This part comes with comparatively little effort once the persistence framework has been set up.
- Chapter 7 shifts the focus somewhat and addresses the problem of comparing datasets against one another. After setting up the theoretical framework, in which datasets and their underlying structures are treated as metric spaces, it shows how persistence can be used to define descriptors that are provably stable under very general hypotheses. It also addresses the question of computing these descriptors (or reliable approximations) efficiently. Down the road, this chapter provides material for comparing shapes, images, or more general data sets, with guarantees.

The third part of the document is more prospective and is divided into two short chapters: one is on current trends in topological data analysis (Chapter 8), the other is on further developments of the theory (Chapter 9). This part gathers the many open questions raised within the previous parts, along with some additional comments and references.

Part 1

Theoretical Foundations

Algebraic Persistence

As we saw in the general introduction, the algebraic theory of persistence deals with certain types of diagrams of vector spaces and linear maps, called *persistence modules*. The simplest instances look like this, where the spaces V_1, \dots, V_n and the maps v_1, \dots, v_{n-1} are arbitrary:

$$V_1 \xrightarrow{v_1} V_2 \xrightarrow{v_2} \dots \xrightarrow{v_{n-1}} V_n.$$

Diagrams such as this one are representations of the so-called *linear quiver* L_n . More generally, all persistence modules are representations of certain types of quivers, possibly with relations. Quiver theory provides us not only with a convenient terminology to define persistence modules and describe their properties, but also with a set of powerful structure theorems to decompose them into ‘atomic’ representations called *interval modules*. Thus, in its algebraic formulation, persistence owes a lot to the theory of quiver representations.

Yet, algebraic persistence cannot be quite reduced to a subset of quiver theory. Shifting the focus from quiver representations to signatures derived from their interval decompositions, it provides very general stability theorems for these signatures and efficient algorithms to compute them. Over time, these signatures—known as the *persistence diagrams*—have become its main object of study and its primary tool for applications. This is the story told in this part of the book: first, how persistence develops as an offspring of quiver theory by focusing on certain types of quivers, and second, how it departs from it by shifting the focus from quiver representations to their signatures.

The first part of the chapter introduces persistence modules using the language of quiver theory. It begins naturally with an overview of the required background material from quiver theory (Section 1), followed by a formal introduction to persistence modules and a review of the conditions under which they can be decomposed (Section 2). The emphasis is on the legacy of quiver theory to persistence.

The second part of the chapter introduces persistence diagrams as signatures for decomposable persistence modules (Sections 3). It then shows how these signatures can be generalized to a class of (possibly indecomposable) representations called the *q-tame* modules (Section 4). The stability properties of these signatures are deferred to Chapter 3.

The chapter closes with a general discussion (Section 5).

Prerequisites. No background on quiver theory is required to read the chapter. However, a reasonable background in abstract and commutative algebra, corresponding roughly to Parts I through III of [111], is needed. Also, some basic notions of category theory, corresponding roughly to Chapters I and VIII of [184], can be helpful although they are not strictly required.

1. A quick walk through the theory of quiver representations

This section gives a brief overview of the concepts and results from quiver theory that will be used afterwards. It also sets up the terminology and notations. The progression is from the more classical aspects of quiver theory to the ones more closely related to persistence. It gives a limited view of the theory of representations, which is a far broader subject.

A more thorough treatment is provided in Appendix A for the interested reader. It includes formal definitions for the concepts introduced here, and a proof outline for Gabriel’s theorem, the key result of this section. It also includes further background material, and draws some connections between tools developed independently in persistence and in quiver theory—e.g. *Diamond Principle* of Carlsson and de Silva [49] versus *reflection functors* of Bernstein, Gelfand, and Ponomarev [24].

Quivers. *Quivers* can be thought of as directed graphs, or rather multigraphs, with potentially infinitely many nodes and arrows. Here is a very simple example of quiver:

$$(1.1) \quad \bullet_1 \xrightarrow{a} \bullet_2 \xleftarrow{b} \bullet_3 \xleftarrow{c} \bullet_4 \xrightarrow{d} \bullet_5$$

and here is a more elaborate example:

$$(1.2) \quad \begin{array}{c} \bullet_1 \\ \swarrow a \quad \nwarrow e \\ \bullet_2 \xrightarrow{c} \bullet_3 \\ \searrow b \end{array}$$

The quivers we are most interested in are the so-called A_n -type quivers, which are finite and linear-shaped, with arbitrary arrow orientations, as in (1.1). Here is a formal general description, where a headless arrow means that the actual arrow orientations can be arbitrary:

$$(1.3) \quad \bullet_1 \text{ --- } \bullet_2 \text{ --- } \cdots \text{ --- } \bullet_{n-1} \text{ --- } \bullet_n$$

The special case where all arrows are oriented to the right is called the *linear quiver*, denoted L_n . Not only A_n -type quivers but also their infinite extensions are relevant to persistence theory.

Quiver representations. *Representations* of a quiver Q over a field k are just realizations of Q as a (possibly non-commutative) diagram of k -vector spaces and k -linear maps. For instance, a representation of the quiver (1.1) can be the following:

$$(1.4) \quad k \xrightarrow{\begin{pmatrix} 1 \\ 0 \end{pmatrix}} k^2 \xleftarrow{\begin{pmatrix} 1 \\ 1 \end{pmatrix}} k \xleftarrow{\begin{pmatrix} 0 & 1 \end{pmatrix}} k^2 \xrightarrow{\begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}} k^2$$

or the following:

$$(1.5) \quad k \xrightarrow{0} 0 \xleftarrow{0} k \xleftarrow{1} k \xrightarrow{0} 0$$

Here is an example of a representation of the quiver (1.2)—note that none of the triangles commute:

$$(1.6) \quad \begin{array}{ccc} & \mathbf{k} & \\ \mathbb{1} \curvearrowright & \nearrow \mathbb{1} & \nwarrow (0 \ 1) \\ \mathbf{k} & \xrightarrow{(\begin{smallmatrix} 1 \\ 0 \end{smallmatrix})} & \mathbf{k}^2 \end{array}$$

A *morphism* $\phi : \mathbb{V} \rightarrow \mathbb{W}$ between two representations of \mathbf{Q} is a collection of linear maps $\phi_i : V_i \rightarrow W_i$ at the nodes \bullet_i of \mathbf{Q} , such that the following diagram commutes

for every arrow $\bullet_i \xrightarrow{a} \bullet_j$ of \mathbf{Q} :

$$(1.7) \quad \begin{array}{ccc} V_i & \xrightarrow{v_a} & V_j \\ \phi_i \downarrow & & \downarrow \phi_j \\ W_i & \xrightarrow{w_a} & W_j \end{array}$$

For instance, a morphism ϕ from (1.4) to (1.5) can be the following collection of vertical maps making every quadrangle commute:

$$(1.8) \quad \begin{array}{ccccccc} \mathbf{k} & \xrightarrow{(\begin{smallmatrix} 1 \\ 0 \end{smallmatrix})} & \mathbf{k}^2 & \xleftarrow{(\begin{smallmatrix} 1 \\ 1 \end{smallmatrix})} & \mathbf{k} & \xleftarrow{(0 \ 1)} & \mathbf{k}^2 & \xrightarrow{(\begin{smallmatrix} 1 & 0 \\ 0 & 1 \end{smallmatrix})} & \mathbf{k}^2 \\ \mathbb{1} \downarrow & & \downarrow 0 & & \downarrow -\mathbb{1} & & \downarrow (0 \ -1) & & \downarrow 0 \\ \mathbf{k} & \xrightarrow{0} & 0 & \xleftarrow{0} & \mathbf{k} & \xleftarrow{\mathbb{1}} & \mathbf{k} & \xrightarrow{0} & 0 \end{array}$$

ϕ is called an *isomorphism* between representations when all its linear maps ϕ_i are isomorphisms between vector spaces. The commutativity condition in (1.7) ensures then that \mathbb{V} and \mathbb{W} have the same algebraic structure.

The category of representations. The representations of a given quiver \mathbf{Q} over a fixed base field \mathbf{k} , together with the morphisms connecting them, form a category denoted $\text{Rep}_{\mathbf{k}}(\mathbf{Q})$. This category turns out to be *abelian*, so some of the nice properties of single vector spaces (which can also be viewed as representations of the quiver \bullet having one node and no arrow) carry over to representations of arbitrary quivers. In particular:

- There is a zero object in the category, called the *trivial representation*. It is made up only of zero vector spaces and linear maps. For instance, the trivial representation of the quiver (1.1) is

$$0 \longrightarrow 0 \longleftarrow 0 \longleftarrow 0 \longrightarrow 0$$

- We can form internal and external direct sums of representations. For instance, the external direct sum of (1.4) and (1.5) is the following representation of (1.1), where the spaces and maps are naturally defined as the direct sums of their counterparts in (1.4) and (1.5):

$$\mathbf{k}^2 \xrightarrow{(\begin{smallmatrix} 1 & 0 \\ 0 & 0 \end{smallmatrix})} \mathbf{k}^2 \xleftarrow{(\begin{smallmatrix} 1 & 0 \\ 1 & 0 \end{smallmatrix})} \mathbf{k}^2 \xleftarrow{(\begin{smallmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \end{smallmatrix})} \mathbf{k}^3 \xrightarrow{(\begin{smallmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \end{smallmatrix})} \mathbf{k}^2$$

- We can define the *kernel*, *image* and *cokernel* of any morphism $\phi : \mathbb{V} \rightarrow \mathbb{W}$. These are defined *pointwise*, with respectively $\ker \phi_i$, $\text{im } \phi_i$ and $\text{coker } \phi_i$ attached to each node \bullet_i of \mathbb{Q} , the linear maps between nodes being induced from the ones in \mathbb{V} and \mathbb{W} . For instance, the kernel of the morphism in (1.8) is

$$0 \xrightarrow{0} \mathbf{k}^2 \xleftarrow{0} 0 \xleftarrow{0} \mathbf{k} \xrightarrow{\begin{pmatrix} 1 \\ 0 \end{pmatrix}} \mathbf{k}^2$$

As expected, ϕ is an isomorphism if and only if both $\ker \phi$ and $\text{coker } \phi$ are trivial.

Not all properties of single vector spaces carry over to representations of arbitrary quivers though. Perhaps the most notable exception is semisimplicity, i.e. the fact that a subspace of a vector space always has a complement: there is no such thing for representations of arbitrary quivers. For instance, $\mathbb{W} = 0 \xrightarrow{0} \mathbf{k}$ is a *subrepresentation* of $\mathbb{V} = \mathbf{k} \xrightarrow{1} \mathbf{k}$, i.e. its spaces are subspaces of the ones in \mathbb{V} and its maps are the restrictions of the ones in \mathbb{V} , yet \mathbb{W} is not a *summand* of \mathbb{V} , i.e. there is no subrepresentation \mathbb{U} such that $\mathbb{V} = \mathbb{U} \oplus \mathbb{W}$. This obstruction is what makes the classification of quiver representations an essentially more challenging problem than for single vector spaces.

Classification of quiver representations. Given a fixed quiver \mathbb{Q} and a fixed base field \mathbf{k} , what are the isomorphism classes of representations of \mathbb{Q} over \mathbf{k} ? This central problem in quiver theory has a decisive impact on persistence, as it provides decomposition theorems for persistence modules. Although solving it in full generality is an essentially impossible task, under some restrictions it becomes remarkably simple. For instance, assuming the quiver \mathbb{Q} is finite and every representation of \mathbb{Q} under consideration has finite *dimension* (defined as the sum of the dimensions of its constituent vector spaces), we benefit from the Krull-Remak-Schmidt principle, that is:

THEOREM 1.1 (Krull, Remak, Schmidt). *Let \mathbb{Q} be a finite quiver, and let \mathbf{k} be a field. Then, every finite-dimensional representation \mathbb{V} of \mathbb{Q} over \mathbf{k} decomposes as a direct sum*

$$(1.9) \quad \mathbb{V} = \mathbb{V}^1 \oplus \dots \oplus \mathbb{V}^r$$

where each \mathbb{V}^i is itself indecomposable i.e. cannot be further decomposed into a direct sum of at least two nonzero representations. Moreover, the decomposition (1.9) is unique up to isomorphism and permutation of the terms in the direct sum.

The proof of existence is an easy induction on the dimension of \mathbb{V} , while the proof of uniqueness can be viewed as a simple application of Azumaya's theorem [14]. This result turns the classification problem into the one of identifying the isomorphism classes of indecomposable representations. Gabriel [133] settled the question for a small subset of the finite quivers called the *Dynkin* quivers. His result asserts that Dynkin quivers only have finitely many isomorphism classes of indecomposable representations, and it provides a simple way to identify these classes. It also introduces a dichotomy on the finite connected quivers, between the ones that are Dynkin, for which the classification problem is easy, and the rest, for which

the problem is significantly harder if at all possible¹. Luckily for us, A_n -type quivers are Dynkin, so Gabriel’s theorem applies and takes the following special form:

THEOREM 1.2 (Gabriel for A_n -type quivers). *Let \mathbb{Q} be an A_n -type quiver, and let \mathbf{k} be a field. Then, every indecomposable finite-dimensional representation of \mathbb{Q} over \mathbf{k} is isomorphic to some interval representation $\mathbb{I}_{\mathbb{Q}}[b, d]$, described as follows:*

$$\underbrace{0 \xrightarrow{0} \dots \xrightarrow{0} 0}_{[1, b-1]} \xrightarrow{0} \underbrace{\mathbf{k} \xrightarrow{1} \dots \xrightarrow{1} \mathbf{k}}_{[b, d]} \xrightarrow{0} \underbrace{0 \xrightarrow{0} \dots \xrightarrow{0} 0}_{[d+1, n]}$$

Combined with Theorem 1.1, this result asserts that every finite-dimensional representation of an A_n -type quiver \mathbb{Q} decomposes uniquely (up to isomorphism and permutation of the terms) as a direct sum of interval representations. This not only gives an exhaustive classification of the finite-dimensional representations of \mathbb{Q} , but it also provides complete descriptors for their isomorphism classes, as any such class is fully described by the collection of intervals $[b, d]$ involved in its decomposition. This collection of intervals is at the origin of our persistence barcodes.

Unfortunately, Theorems 1.1 and 1.2 are limited in two ways for our purposes. First, by only considering quivers indexed over finite sets, whereas we would like to consider arbitrary subsets of \mathbb{R} . Second, by restricting the focus to finite-dimensional representations, whereas in our case we may have to deal with representations including infinitely many nontrivial spaces, or spaces of infinite dimension. The rest of Section 1 is a review of several extensions of the theorems that address these limitations.

Infinite-dimensional representations. Theorems 1.1 and 1.2 turn out to be still valid if infinite-dimensional representations are considered as well. This is thanks to a powerful result of Auslander [12] and Ringel and Tachikawa [217] dealing with left modules over Artin algebras². The connection with quiver representations is done through the construction of the so-called *path algebra* of a quiver \mathbb{Q} , denoted $\mathbf{k}\mathbb{Q}$, which in short is the \mathbf{k} -algebra generated by the finite oriented paths in \mathbb{Q} , with the product operator induced by concatenations of paths. The path algebra is an Artin algebra whenever \mathbb{Q} is finite with no oriented cycle, which happens e.g. when \mathbb{Q} is an A_n -type quiver. There is then an equivalence of categories between $\text{Rep}_{\mathbf{k}}(\mathbb{Q})$ and the left modules over $\mathbf{k}\mathbb{Q}$. The result of Auslander [12] and Ringel and Tachikawa [217], combined with Gabriel’s and Azumaya’s theorems, gives the following decomposition theorem for finite- or infinite-dimensional representations of A_n -type quivers:

THEOREM 1.3 (Auslander, Ringel, Tachikawa, Gabriel, Azumaya). *Let \mathbb{Q} be an A_n -type quiver, and let \mathbf{k} be a field. Then, every indecomposable representation of \mathbb{Q} over \mathbf{k} is isomorphic to an interval representation $\mathbb{I}_{\mathbb{Q}}[b, d]$, and every representation of \mathbb{Q} , whether finite- or infinite-dimensional, is isomorphic to a (possibly infinite) direct sum of interval representations. Moreover, this decomposition is unique up to isomorphism and permutation of the terms.*

¹There is in fact a trichotomy: beside the Dynkin quivers are the so-called *tame* quivers, for which the classification problem is still feasible (albeit harder); the rest of the quivers are called *wild* because for them the classification problem is an essentially impossible task.

²An Artin algebra is a finitely generated algebra over an Artinian ring, i.e. a commutative ring that satisfies the descending chain condition on ideals: every nested sequence of ideals $I_1 \supseteq I_2 \supseteq \dots$ stabilizes eventually.

Representations of arbitrary subposets of (\mathbb{R}, \leq) . We now consider extensions of the index set to arbitrary subsets T of \mathbb{R} . For this we work with the poset (T, \leq) directly, rather than with some associated quiver. Regarding (T, \leq) as a category in the natural way⁴, we let a *representation of (T, \leq)* be a functor to the category of vector spaces. What this means concretely is that the representation defines vector spaces $(V_i)_{i \in T}$ and linear maps $(v_i^j : V_i \rightarrow V_j)_{i \leq j \in T}$ satisfying the following constraints induced by functoriality:

$$(1.10) \quad \begin{aligned} v_i^i &= \mathbb{1}_{V_i} && \text{for every } i \in T, \text{ and} \\ v_i^k &= v_j^k \circ v_i^j && \text{for every } i \leq j \leq k \in T. \end{aligned}$$

More generally, one can define representations for any given poset as functors from that poset to the category of vector spaces.

Crawley-Boevey [93] has extended Theorems 1.1 and 1.2 to representations of arbitrary subposets of (\mathbb{R}, \leq) . Pointwise finite-dimensionality is understood as in Theorem 1.4. The proof uses a specialized version of the *functorial filtration* method of Ringel [216]. The uniqueness of the decomposition once again follows from Azumaya's theorem.

THEOREM 1.6 (Crawley-Boevey). *Let \mathbf{k} be a field and let $T \subseteq \mathbb{R}$. Then, any pointwise finite-dimensional representation of (T, \leq) over \mathbf{k} is a direct sum of interval representations.*

Note that there is a larger variety of interval representations for the posets (\mathbb{Z}, \leq) and (\mathbb{R}, \leq) than for the A_n -type quivers. Indeed, some intervals may be left-infinite, or right-infinite, or both. Moreover, since \mathbb{R} has limit points, some intervals for (\mathbb{R}, \leq) may be open or half-open. We will elaborate on this point in the next section.

REMARK. The connection to quiver theory is somewhat more subtle in this general setting than in the previous ones. First of all, any quiver can be viewed as a category, with one object per node and one morphism per finite oriented path. Its representations are then interpreted as functors to the category of vector spaces. In the case of the quivers \mathbb{N} and \mathbb{Z} , the corresponding categories are equivalent to the posets (\mathbb{N}, \leq) and (\mathbb{Z}, \leq) respectively. However, not every quiver is equivalent (as a category) to a poset, and conversely, not every poset is equivalent to a quiver. The reason for the latter limitation is that paths sharing the same source and the same target are not considered equal in a quiver (recall (1.2) and (1.6)), whereas they are in a poset by transitivity. The workaround is to equip the quivers with *relations* that identify the paths sharing the same source and the same target. The representations of the resulting *quivers with relations* have to reflect these identifications. This way, any poset can be made equivalent (as a category) to some quiver with relations, and its representations can be viewed themselves as quiver representations.

2. Persistence modules and interval decompositions

The background material on quiver theory given in Section 1 provides us with a convenient terminology to introduce persistence modules—see also Section 5 for a historical account. From now on and until the end of the chapter, the field \mathbf{k} over which representations are taken is fixed.

⁴i.e. with one object per element $i \in T$ and a single morphism per couple $i \leq j$.

DEFINITION 1.7. Given $T \subseteq \mathbb{R}$, a *persistence module* over T is a representation of the poset (T, \leq) .

This definition follows (1.10) and includes the representations of the quivers L_n , N and Z as special cases. However, it does not include the representations of general A_n -type quivers, which are gathered into a different concept called *zigzag module*.

DEFINITION 1.8. Given $n \geq 1$, a *zigzag module* of length n is a representation of an A_n -type quiver.

The term ‘zigzag’ is justified by the following special situation motivated by applications, where every other arrow is oriented backwards:

$$(1.11) \quad V_1 \xrightarrow{v_1} V_2 \xleftarrow{v_2} V_3 \xrightarrow{v_3} \cdots \xleftarrow{v_{n-3}} V_{n-2} \xrightarrow{v_{n-2}} V_{n-1} \xleftarrow{v_{n-1}} V_n$$

Zigzag modules can also be thought of as poset representations. As a directed acyclic graph, an A_n -type quiver Q is the Hasse diagram⁵ of some partial order relation \preceq on the set $\{1, \dots, n\}$. Since Q has at most one oriented path between any pair of nodes, it is equivalent (as a category) to the poset $(\{1, \dots, n\}, \preceq)$, and its representations are also representations of $(\{1, \dots, n\}, \preceq)$. Thus, we can rewrite Definition 1.8 as follows, which emphasizes its connection to Definition 1.7:

DEFINITION 1.8 (rephrased). Given $n \geq 1$, a *zigzag module* of length n is a representation of the poset $(\{1, \dots, n\}, \preceq)$, where \preceq is any partial order relation whose Hasse diagram is of type A_n .

Section 1 provides us with powerful structure theorems to decompose these objects. The basic building blocks are the interval representations, called *interval modules* in the persistence literature. Given an arbitrary index set $T \subseteq \mathbb{R}$, an *interval* of T is a subset $S \subseteq T$ such that for any elements $i \leq j \leq k$ of T , $i, k \in S$ implies $j \in S$. The associated interval module has the field \mathbf{k} at every index $i \in S$ and the zero space elsewhere, the maps between copies of \mathbf{k} being identities and all other maps being zero. This definition is oblivious to the actual map orientations, which depend on the order relation \preceq that equips the index set T . When this relation is obvious from the context, we simply write $\mathbb{I}S$ for the interval module associated with S ; otherwise we write $\mathbb{I}_{\preceq}S$, or even $\mathbb{I}_Q S$ when \preceq is specified through its Hasse diagram Q .

The following theorem summarizes the structural results from Section 1 (Theorems 1.1, 1.2, 1.3 and 1.6) and can be thought of as our main heritage from quiver theory. The conditions under which it guarantees the existence of an interval decomposition are sufficient for our purposes.

THEOREM 1.9 (Interval Decomposition). *Given an index set $T \subseteq \mathbb{R}$ and a partial order relation \preceq on T , a representation \mathbb{V} of the poset (T, \preceq) can be decomposed as a direct sum of interval modules in each of the following situations:*

- (i) *T is finite and the Hasse diagram of \preceq is of type A_n (which happens in particular when \preceq is the natural order \leq on T , whose Hasse diagram is L_n),*
- (ii) *T is arbitrary, \preceq is the natural order \leq , and \mathbb{V} is pointwise finite-dimensional. Moreover, the decomposition, when it exists, is unique up to isomorphism and permutation of the terms in the direct sum, and each term is indecomposable.*

⁵Defined as the graph having $\{1, \dots, n\}$ as vertex set, and one edge $i \rightarrow j$ per couple $i \prec j$ such that there is no k with $i \prec k \prec j$.

So, concretely:

- (i) any zigzag module decomposes uniquely as a direct sum of interval modules,
- (i)-(ii) any persistence module whose index set is finite or whose vector spaces are finite-dimensional decomposes uniquely as a direct sum of interval modules.

A persistence or zigzag module \mathbb{V} that decomposes as a direct sum of interval modules is called *interval-decomposable*. The converse (called *interval-indecomposable*) means that either \mathbb{V} decomposes into indecomposable representations that are not interval modules, or \mathbb{V} does not decompose at all as a direct sum of indecomposable representations. In principle, interval-decomposability is a stronger concept than the classical notion of decomposability into indecomposables from quiver theory. Nevertheless, the tools introduced by Webb [238] can be used to prove both concepts equivalent for modules over (subsets of) \mathbb{Z} [192], while to our knowledge the question is still not settled for modules over \mathbb{R} .

3. Persistence barcodes and diagrams

In order to simplify the description of interval modules over arbitrary subsets of \mathbb{R} , including subsets with limit points, we need to decide on a simple and unified writing convention for intervals. For this purpose we will use *decorated real numbers*, which are written as ordinary real numbers with an additional superscript $+$ (plus) or $-$ (minus). Whenever the decoration of a number is unknown or irrelevant, we will use the superscript \pm . The order on decorated numbers is the obvious one: $b^\pm < d^\pm$ if $b < d$, or if $b = d$ and $b^\pm = b^-$ and $d^\pm = b^+$. The corresponding dictionary for finite intervals of \mathbb{R} is the following one, where $b^\pm < d^\pm$:

$$\begin{aligned} [b^-, d^-] &\text{ stands for } [b, d), \\ [b^-, d^+] &\text{ stands for } [b, d], \\ [b^+, d^-] &\text{ stands for } (b, d), \\ [b^+, d^+] &\text{ stands for } (b, d]. \end{aligned}$$

We will also use the symbols $-\infty$ and $+\infty$ for infinite endpoints. Since intervals are always open at infinity, these implicitly carry the superscripts $-\infty^+$ and $+\infty^-$, which we will generally omit in the notations, so for instance $[-\infty, d^-]$ stands for the open interval $(-\infty, d)$.

Given an arbitrary index set $T \subseteq \mathbb{R}$, we can now rewrite each interval S of T as $[b^\pm, d^\pm] \cap T$, for some interval $[b^\pm, d^\pm]$ of \mathbb{R} . Note that the choice of $[b^\pm, d^\pm]$ may not be unique when T is a strict subset of \mathbb{R} . For instance, letting $S = \{2, 3\}$ and $T = \{1, 2, 3, 4\}$, we can write S indifferently as $[2, 3] \cap T$, or $(1, 3] \cap T$, or $(1, 4) \cap T$, or $[2, 4) \cap T$, or more generally $[b^\pm, d^\pm] \cap T$ for any $[2, 3] \subseteq [b^\pm, d^\pm] \subseteq (1, 4)$. To remove ambiguities, unless otherwise stated we will always pick the interval of \mathbb{R} that is smallest with respect to inclusion, as per the following rule⁶:

RULE 1.10. *Given an interval S of an index set T , among the intervals of \mathbb{R} whose intersection with T is S , pick the one that is smallest with respect to inclusion, e.g. $[2, 3]$ in the previous example.*

⁶This rule has been applied implicitly until now. It is arbitrary and not motivated by mathematical considerations. Other rules can be applied as well, leading to different writings of the same interval S of T .

$$\begin{array}{cc}
 [b_j, d_j] = [b_j^-, d_j^+] & [b_j^+, d_j^+] = (b_j, d_j) \\
 \times & \\
 [b_j, d_j] = [b_j^-, d_j^-] & [b_j^+, d_j^-] = (b_j, d_j)
 \end{array}$$

FIGURE 1.1. The four decorated points corresponding to intervals $[b_j^\pm, d_j^\pm]$.

For simplicity we will also omit the index set T in the notation when it is irrelevant or obvious from the context. Then, any interval-decomposable module \mathbb{V} can be written uniquely (up to permutation of the terms) as

$$(1.12) \quad \mathbb{V} \cong \bigoplus_{j \in J} \mathbb{I}[b_j^\pm, d_j^\pm].$$

The set of intervals $[b_j^\pm, d_j^\pm]$, ordered by the lexicographical order on the decorated coordinates, is called the *persistence barcode* of \mathbb{V} . Technically it is a multiset, as an interval may occur more than once. Another representation of the persistence barcode is as a multiset of *decorated points* in the extended plane $\bar{\mathbb{R}}^2 = [-\infty, +\infty]^2$, where each interval $[b_j^\pm, d_j^\pm]$ is identified with the point of coordinates (b_j, d_j) decorated with a diagonal tick according to the convention of Figure 1.1. This multiset of decorated points is called the *decorated persistence diagram* of \mathbb{V} , noted $\text{Dgm}(\mathbb{V})$. From (1.12),

$$(1.13) \quad \text{Dgm}(\mathbb{V}) = \{(b_j^\pm, d_j^\pm) \mid j \in J\}.$$

The *undecorated persistence diagram* of \mathbb{V} , noted $\text{dgm}(\mathbb{V})$, is the same multiset without the decorations:

$$(1.14) \quad \text{dgm}(\mathbb{V}) = \{(b_j, d_j) \mid j \in J\}.$$

Let us give a concrete example taken from our traditional zoo—the background details will be given in Chapter 2.

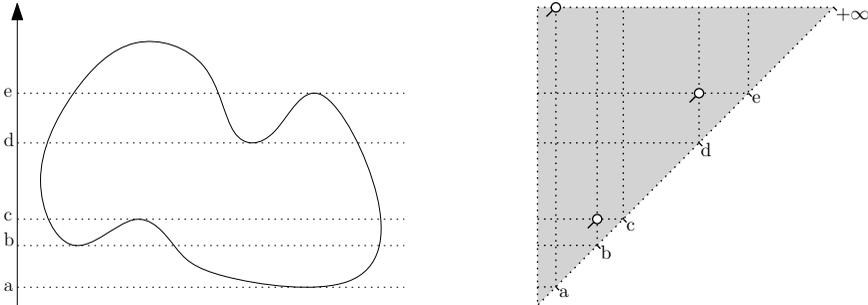


FIGURE 1.2. A classical example in persistence theory. Left: a smooth planar curve X and its y -coordinate or ‘height’ function $f : X \rightarrow \mathbb{R}$. Right: the decorated persistence diagram of $H_0(\mathcal{F})$.

— From Chazal et al. [72].

EXAMPLE 1.11. Consider the curve X in \mathbb{R}^2 shown in Figure 1.2, filtered by the height function f . Take the family \mathcal{F} of sublevelsets $F_y = f^{-1}((-\infty, y])$ of f , where parameter y ranges over \mathbb{R} . This family is nested, that is, $F_y \subseteq F_{y'}$ whenever $y \leq y'$.

Let us apply the 0-homology functor H_0 to \mathcal{F} and study the resulting persistence module $H_0(\mathcal{F})$, which encodes the evolution of the connectivity of the sublevel sets F_y as parameter y ranges from $-\infty$ to $+\infty$. This module decomposes as follows:

$$H_0(\mathcal{F}) \cong \mathbb{I}[a, +\infty) \oplus \mathbb{I}[b, c) \oplus \mathbb{I}[d, e) = \mathbb{I}[a^-, +\infty] \oplus \mathbb{I}[b^-, c^-] \oplus \mathbb{I}[d^-, e^-],$$

which intuitively means that 3 different independent connected components appear in the sublevel set F_y during the process: the first one at $y = a$, the second one at $y = b$, the third one at $y = d$; while the first one remains until the end, the other two eventually get merged into it, at times $y = c$ and $y = e$ respectively. A pictorial description of this decomposition is provided by the decorated persistence diagram in Figure 1.2.

The persistence measure. Let \mathbb{V} be an interval-decomposable module. From its decorated persistence diagram $\text{Dgm}(\mathbb{V})$ we derive the following measure on rectangles $R = [p, q] \times [r, s]$ in the extended plane with $-\infty \leq p < q \leq r < s \leq +\infty$:

$$(1.15) \quad \mu_{\mathbb{V}}(R) = \text{card}(\text{Dgm}(\mathbb{V})|_R),$$

where the membership relation for a decorated point $(b^\pm, d^\pm) \in \text{Dgm}(\mathbb{V})$ is defined by:

$$(1.16) \quad (b^\pm, d^\pm) \in R \iff [q, r] \subseteq [b^\pm, d^\pm] \subseteq [p, s].$$

The pictorial view of (1.16) is that point (b, d) and its decoration tick belong to the closed rectangle R , as illustrated in Figure 1.3. Then, (1.15) merely defines $\mu_{\mathbb{V}}$ as the counting measure over the restrictions of $\text{Dgm}(\mathbb{V})$ to rectangles, which takes values in $\{0, 1, 2, \dots, +\infty\}$ (we do not distinguish between infinite values). The

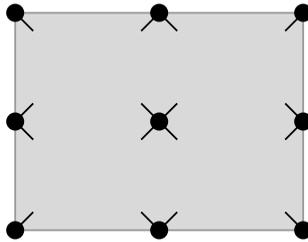


FIGURE 1.3. A decorated point (b^\pm, d^\pm) belongs to a rectangle R if (b, d) belongs to the interior of R , or if it belongs to the boundary of R with its tick pointing towards the interior of R .

— From Chazal et al. [72].

term *measure* is motivated by the fact that $\mu_{\mathbb{V}}$ is additive with respect to splitting a rectangle into two rectangles, either vertically or horizontally, with the convention that $x + \infty = +\infty + x = +\infty$:

$$(1.17) \quad \forall p < x < q \leq r < y < s, \quad \begin{cases} \mu_{\mathbb{V}}([p, q] \times [r, s]) = \mu_{\mathbb{V}}([p, x] \times [r, s]) + \mu_{\mathbb{V}}([x, q] \times [r, s]) \\ \mu_{\mathbb{V}}([p, q] \times [r, s]) = \mu_{\mathbb{V}}([p, q] \times [r, y]) + \mu_{\mathbb{V}}([p, q] \times [y, s]) \end{cases}$$

This additivity property is illustrated in Figure 1.4, where the claim is that $\mu_{\mathbb{V}}(R) = \mu_{\mathbb{V}}(A) + \mu_{\mathbb{V}}(B) = \mu_{\mathbb{V}}(C) + \mu_{\mathbb{V}}(D)$. Notice how the decoration of a given point on the border between two subrectangles assigns this point uniquely to one of them.

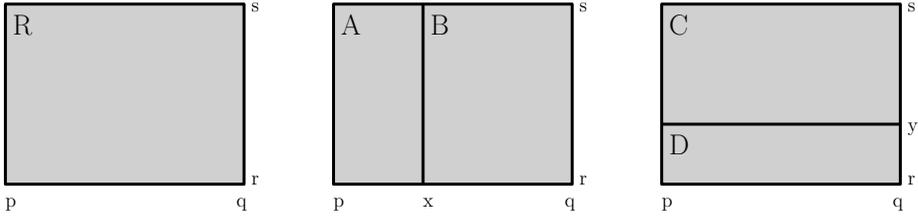


FIGURE 1.4. Additivity of $\mu_{\mathbb{V}}$ under vertical / horizontal splitting.
 — From Chazal et al. [72].

When $\mathbb{V} = (V_i, v_i^j)$ is a persistence module over \mathbb{R} , we can also relate its persistence measure $\mu_{\mathbb{V}}$ more directly to \mathbb{V} through the following well-known inclusion-exclusion formulas, which hold provided that all the ranks are finite or, less stringently, that all but $\text{rank } v_q^r$ are finite—in which case the values of the alternating sums are $+\infty$:

$$\begin{aligned}
 &\forall p < q \leq r < s \in \mathbb{R}, \\
 (1.18) \quad &\begin{cases} \mu_{\mathbb{V}}([-\infty, q] \times [r, +\infty]) &= \text{rank } v_q^r, \\ \mu_{\mathbb{V}}([-\infty, q] \times [r, s]) &= \text{rank } v_q^r - \text{rank } v_q^s, \\ \mu_{\mathbb{V}}([p, q] \times [r, +\infty]) &= \text{rank } v_q^r - \text{rank } v_p^r, \\ \mu_{\mathbb{V}}([p, q] \times [r, s]) &= \text{rank } v_q^r - \text{rank } v_q^s + \text{rank } v_p^s - \text{rank } v_p^r. \end{cases}
 \end{aligned}$$

The first formula counts the number of decorated points of $\text{Dgm}(\mathbb{V})$ that lie inside the quadrant $Q = [-\infty, q] \times [r, +\infty]$, the second formula inside the horizontal strip $H = [-\infty, q] \times [r, s]$, the third formula inside the vertical strip $V = [p, q] \times [r, +\infty]$, the fourth formula inside the rectangle $R = [p, q] \times [r, s]$. These patterns are illustrated in Figure 1.5. Notice how the second, third and fourth formulas are obtained from the first one by counting points inside quadrants and by removing multiple counts (hence the term ‘inclusion-exclusion formulas’). These formulas are useful to localize points in the persistence diagram from the sole knowledge of the ranks of the linear maps in \mathbb{V} . As we will see next, they can be used to generalize the definition of persistence diagram to a certain class of interval-indecomposable persistence modules.

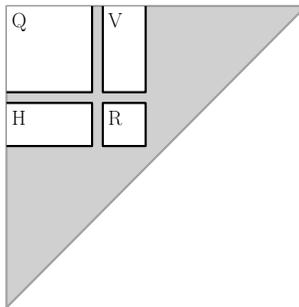


FIGURE 1.5. A quadrant, horizontal strip, vertical strip, and finite rectangle in the half plane above the diagonal.
 — From Chazal et al. [72].

4. Extension to interval-indecomposable persistence modules

So far we have restricted our attention to modules that are decomposable into interval summands. For these modules we have defined persistence diagrams that encode their algebraic structure in a unique and complete way. These diagrams, together with their derived persistence measures, serve as signatures for the modules, and as we will see in Chapter 3 they can be compared against one another in a natural and theoretically sound way.

In order to extend the definition of persistence diagram to persistence modules \mathbb{V} that are not interval-decomposable, we proceed backwards compared to Section 3: first, we use the formulas of (1.18) as axioms to define the persistence measure $\mu_{\mathbb{V}}$ of \mathbb{V} ; then, we show the existence and uniqueness of a multiset of points in the extended plane whose counting measure restricted to rectangles coincides with $\mu_{\mathbb{V}}$. In order to use (1.18), we need to assume that the ranks in the alternating sums are finite, except rank v_q^r , which brings up the following notion of tameness for \mathbb{V} :

DEFINITION 1.12. A persistence module $\mathbb{V} = (V_i, v_i^j)$ over \mathbb{R} is *quadrant-tame*, or *q-tame* for short, if $\text{rank } v_i^j < +\infty$ for all $i < j$.

The reason for this name is obvious from the first formula of (1.18): when \mathbb{V} is interval-decomposable, $\text{rank } v_i^j$ represents the total multiplicity of the diagram inside the quadrant $[-\infty, i] \times [j, +\infty]$. The q-tameness property allows this multiplicity to be infinite when $i = j$ (i.e. when the lower-right corner of the quadrant touches the diagonal line $y = x$) but not when $i < j$ (i.e. when the quadrant lies strictly above the diagonal line). Note that forcing the multiplicity to be finite even when $i = j$ would bring us back to the concept of pointwise finite-dimensional module.

For a q-tame persistence module \mathbb{V} , we define the persistence measure $\mu_{\mathbb{V}}$ by the formulas of (1.18), which as we saw are well-founded in this case. This measure on rectangles does take values in $\{0, 1, 2, \dots, +\infty\}$ as before, although the fact that it cannot be negative is not obvious at first sight: it follows from the observation that the formulas of (1.18) actually count the multiplicity of the summand $\mathbb{I}[q, r]$ in the interval-decomposition of the restriction of the module \mathbb{V} to some finite index set J , which is $J = \{q, r\}$ for the first formula, $J = \{q, r, s\}$ for the second, $J = \{p, q, r\}$ for the third, and $J = \{p, q, r, s\}$ for the fourth. In each case, the restriction of \mathbb{V} to this finite index set J is interval-decomposable by Theorem 1.9 (i), so the multiplicity of the summand $\mathbb{I}[q, r]$ is well-defined and non-negative—for further details on this specific point, please refer to [72, §2.1]. In addition to being non-negative, $\mu_{\mathbb{V}}$ is also additive under vertical and horizontal splittings, a straight consequence of its definition (the duplicated terms in the alternating sums cancel out).

THEOREM 1.13 ([72, theorem 2.8 and corollary 2.15]). *Given a q-tame module \mathbb{V} , there is a uniquely defined locally finite decorated multiset $\text{Dgm}(\mathbb{V})$ in the open extended upper half-plane $\{(x, y) \in \overline{\mathbb{R}^2} \mid x < y\}$ such that for any rectangle $R = [p, q] \times [r, s]$ with $-\infty \leq p < q < r < s \leq +\infty$,*

$$\mu_{\mathbb{V}}(R) = \text{card}(\text{Dgm}(\mathbb{V})|_R).$$

The proof of this result works by subdividing the rectangle R recursively into subrectangles, and by a limit process it eventually charges the local mass of $\mu_{\mathbb{V}}$ to a finite set of decorated points, whose uniqueness is obtained as an easy consequence of the construction. A preliminary version of this argument appeared in [71] and was

based on measures defined on specific families of rectangles. The final version in [72] is more cleanly formalized and considers general rectangle measures, therefore we recommend it to the interested reader.

Theorem 1.13 defines the persistence diagram of a \mathfrak{q} -tame module \mathbb{V} uniquely as a multiset $\mathbf{Dgm}(\mathbb{V})$ of decorated points lying above the diagonal line $\Delta = \{(x, x) \mid x \in \mathbb{R}\}$. Alternately, one can use a persistence barcode representation, in which every diagram point (b^\pm, d^\pm) becomes the interval $\lceil b^\pm, d^\pm \rceil$. The undecorated version of $\mathbf{Dgm}(\mathbb{V})$, denoted $\mathbf{dgm}(\mathbb{V})$, is obtained as in (1.14) by forgetting the decorations.

These definitions agree with the ones from Section 3 in the sense that, if a persistence module \mathbb{V} over \mathbb{R} is both \mathfrak{q} -tame and interval-decomposable, then the multiset $\mathbf{Dgm}(\mathbb{V})$ derived from Theorem 1.13 agrees with the one from (1.13) everywhere above the diagonal Δ —see [72, proposition 2.18]. The two multisets may differ along Δ though, since the interval-decomposition of \mathbb{V} may contain summands of type $\mathbb{I}[b, b] = \mathbb{I}[b^-, b^+]$, which are not captured by the rectangles not touching Δ . It turns out that either definition of $\mathbf{Dgm}(\mathbb{V})$ can be used in practice, as the natural measures of proximity between persistence modules and between their diagrams, which will be presented in Chapter 3, are oblivious to the restrictions of the diagrams to the diagonal Δ . We will therefore be using both definitions indifferently in the following.

5. Discussion

To conclude the chapter, let us discuss some of its concepts further and put them into perspective.

Persistence modules: a historical account. Several definitions of a persistence module coexist in the persistence literature, following the steady development of the theory towards greater generality and abstraction.

The term *persistence module* was coined originally by Zomorodian and Carlsson [243], but the concept appeared already in [118], where it referred to a finite sequence of finite-dimensional vector spaces connected by linear maps as follows:

$$(1.19) \quad V_1 \xrightarrow{v_1} V_2 \xrightarrow{v_2} \cdots \xrightarrow{v_{n-1}} V_n$$

In other words, a persistence module as per Edelsbrunner, Letscher, and Zomorodian [118] is a finite-dimensional representation of the linear quiver L_n . Zomorodian and Carlsson [243] extended the concept to diagrams indexed over the natural numbers, that is, to representations of the quiver N . Cohen-Steiner, Edelsbrunner, and Harer [87] further extended the concept to work with diagrams indexed over the real line. They defined a persistence module as an indexed family of finite-dimensional vector spaces $\{V_i\}_{i \in \mathbb{R}}$ together with a doubly indexed family of linear maps $\{v_i^j : V_i \rightarrow V_j\}_{i \leq j}$ that satisfy the following identity and composition rules:

$$(1.20) \quad \begin{aligned} \forall i, \quad v_i^i &= \mathbf{1}_{V_i}, \\ \forall i \leq j \leq k, \quad v_i^k &= v_j^k \circ v_i^j. \end{aligned}$$

Such objects are pointwise finite-dimensional representations of the poset (\mathbb{R}, \leq) , the identity and composition rules (1.20) following from functoriality as in (1.10).

Chazal et al. [71] dropped the finite-dimensionality condition on the vector spaces, and then Chazal et al. [72] replaced \mathbb{R} by any subset $T \subseteq \mathbb{R}$ equipped with the same order relation \leq . Hence Definition 1.7.

Carlsson and de Silva [49] extended the concept of persistence module in a different way, by choosing arbitrary orientations for the linear maps connecting the finite-dimensional vector spaces in (1.19). This gave rise to the concept of zigzag module, presented in Definition 1.8 without the finite-dimensionality condition.

More recently, Bubenik and Scott [41] then Bubenik, de Silva, and Scott [40] proposed to generalize the concept of persistence module to representations of arbitrary posets. This generalization reaches far beyond the setting of 1-dimensional persistence, losing some of its fundamental properties along the way, such as the ability to define complete discrete invariants like persistence barcodes in a systematic way. Nevertheless, it still guarantees some ‘soft’ form of stability, as will be discussed at the end of Chapter 3 and then in Chapter 9.

Interval decompositions in the persistence literature. Although Theorem 1.9 is presented as a byproduct of representation theory in the chapter, it actually took our community some time to realize this connection.

Historically, Zomorodian and Carlsson [243] were the first ones to describe persistence modules in terms of representations. They pointed out the connection between the persistence modules over \mathbb{N} and the graded modules over the polynomial ring $\mathbf{k}[t]$ (mentioned after Theorem 1.4), and they used the structure theorem for finitely generated modules over a principal ideal domain as decomposition theorem for finite-dimensional persistence modules over \mathbb{N} .

Some time later, Carlsson and de Silva [49] introduced zigzag modules and connected them to finite-dimensional representations of A_n -type quivers. This connection induces a decomposition theorem for finite-dimensional zigzag modules via Gabriel’s theorem.

More recently, Chazal et al. [72] pointed out the connection between persistence or zigzag modules over finite sets and representations of finitely generated algebras, and they referred to the work of Auslander [12] and Ringel and Tachikawa [217] to decompose arbitrary persistence or zigzag modules (including infinite-dimensional ones) over finite index sets—Theorem 1.9 (i).

In the meantime, Lesnick [179] introduced our community to the work of Webb [238], which generalizes the decomposition theorem used by Zomorodian and Carlsson [243] to pointwise finite-dimensional modules over \mathbb{N} . Crawley-Boevey [93] further extended it to a decomposition theorem for persistence modules over arbitrary subsets of \mathbb{R} under the pointwise finite-dimensionality condition—Theorem 1.9 (ii). His proof turns out to hold under a somewhat weaker (albeit technical) condition [95], which gives hope for tackling the interval decomposability question in greater generality, as will be discussed next.

These results have contributed to shape the theory as we know it today. Among them, let us point out [49] as a key contribution, for bringing the existence of quiver theory and its connection to persistence to the attention of our community, and conversely, for creating an opportunity to advertise persistence among the representation theory community and stimulate interactions. Besides, Carlsson and de Silva [49] proposed a genuinely new constructive proof of Gabriel’s theorem in the special case of A_n -type quivers, which is self-contained, requires no prior knowledge of quiver theory, and has eventually led to a practical algorithm for computing decompositions of zigzag modules [50]. For the interested reader, we analyze this proof and establish connections to the so-called *reflection functors* of Bernstein, Gelfand, and Ponomarev [24] in Section 4.4 of Appendix A.

q-tameness versus interval-decomposability. The concepts of interval-decomposable and q -tame persistence modules over \mathbb{R} are closely related but not identical, and neither of them is a direct generalization of the other. For instance, the module $\bigoplus_{j \in J} \mathbb{I}[b_j^\pm, d_j^\pm]$, where the decorated pairs form a dense subset of the half-plane above Δ , is interval-decomposable but not q -tame—in fact its persistence measure is infinite on every rectangle. By contrast, the module $\prod_{n \geq 1} \mathbb{I}[0, \frac{1}{n}]$ is q -tame but not interval-decomposable.

Both concepts are related though. As we will see in Chapter 3, q -tame modules are the limits of pointwise finite-dimensional modules in some metric called the *interleaving distance*. Recall that pointwise finite-dimensional modules are themselves q -tame by definition and interval-decomposable by Theorem 1.9 (ii), so q -tame modules are limits of (a subset of) the interval-decomposable modules. Furthermore, Crawley-Boevey [95] proved that any q -tame persistence module \mathbb{V} admits interval-decomposable submodules \mathbb{W} whose interleaving distance to \mathbb{V} is zero, so q -tame modules are in fact indistinguishable from interval-decomposable modules in that metric. This result led Bauer and Lesnick [20] to define the undecorated persistence diagram of a q -tame module \mathbb{V} directly as the diagram of any interval-decomposable submodule \mathbb{W} of \mathbb{V} lying at interleaving distance zero from \mathbb{V} . The upcoming Isometry Theorem (Theorem 3.1) implies that this is a sound definition, all such submodules \mathbb{W} having in fact the same undecorated diagram.

One of these submodules stands out: the so-called *radical* $\text{rad}(\mathbb{V})$, defined as follows (where $\mathbb{V} = (V_i, v_i^j)$):

$$(1.21) \quad \forall j \in \mathbb{R}, \text{rad}(\mathbb{V})_j = \sum_{i < j} \text{im } v_i^j.$$

Although not always pointwise finite-dimensional, $\text{rad}(\mathbb{V})$ is interval-decomposable [95], furthermore it makes the quotient module $\mathbb{U} = \mathbb{V}/\text{rad}(\mathbb{V})$ *ephemeral*⁷, that is:

$$(1.22) \quad \forall i < j \in \mathbb{R}, \text{rank } u_i^j = 0.$$

Thus, every q -tame module is interval-decomposable ‘modulo’ some ephemeral module. Chazal, Crawley-Boevey, and de Silva [63] formalized this idea by introducing the so-called *observable category* of persistence modules, defined as the quotient category of q -tame modules ‘modulo’ ephemeral modules in the sense of Serre’s theory of localization. In this new category, q -tame modules become interval-decomposable, and (undecorated) persistence diagrams are a complete invariant for them.

The conclusion of this discussion is that q -tame modules appear as a natural extension of (a subset of) the interval-decomposable modules. In addition, experience shows that q -tame modules occur rather widely in applications. For instance, as we will see in Chapter 7, the Vietoris-Rips and Čech complexes of a compact metric space have q -tame persistent homology, whereas they can be very badly behaved when viewed non-persistently. Such examples support the claim that the q -tame modules are also a natural class of persistence modules to work with in practice.

⁷In fact, $\text{rad}(\mathbb{V})$ is the smallest submodule of \mathbb{V} such that the quotient module is ephemeral.

Topological Persistence

Persistence modules arise naturally as algebraic invariants of families of topological spaces. Such families, called *filtrations*, can be viewed as sequences of topological spaces linked by continuous maps (usually taken to be inclusions). The connection between the topological level and the algebraic level happens through some functor, typically the homology or cohomology functor, which turns the filtrations into persistence modules, taking sequences of topological spaces and continuous maps to sequences of vector spaces and linear maps. Then, persistence diagrams can be defined for filtrations and used as signatures for functions over topological spaces. The magic is that the connection between the topological and algebraic levels remains implicit, so users do not have to manipulate persistence modules directly in practice, even when computing the persistence diagrams.

This chapter reviews the most common topological constructions from which persistence modules are derived (Section 1). It also describes how the persistence of these constructions is computed without an explicit representation of their corresponding persistence modules (Section 2). Here we do not pretend to be exhaustive, but rather to give the reader a flavor of the variety of the contributions in this area.

Throughout the chapter, unless otherwise mentioned, we will be using singular homology or cohomology with coefficients in a fixed field \mathbf{k} —omitted in the notations for simplicity.

Prerequisites. We will be assuming familiarity with simplicial complexes as well as simplicial (resp. singular) homology and cohomology. A good introduction to these topics can be found in Chapters 1, 4 and 5 of [202]. We will also use some basic notions of Morse theory, which can be found in Part I of [195].

1. Topological constructions

This section reviews the most common topological constructions in persistence theory: filtrations and excursion sets, relative and extended filtrations, zigzags and level sets, and finally kernels, images and cokernels. It is a good opportunity to introduce some terminology and to set up the notations.

1.1. Filtrations and excursion sets. Filtrations are the simplest kind of topological constructions and also the most widely used in practice.

DEFINITION 2.1. Given a subset $T \subseteq \mathbb{R}$, a *filtration* \mathcal{X} over T is a family of topological spaces X_i , parametrized by $i \in T$, such that $X_i \subseteq X_j$ whenever $i \leq j \in T$.

Thus, \mathcal{X} is a special type of representation of the poset (T, \leq) in the category of topological spaces. \mathcal{X} is called (*finitely*) *simplicial* if the spaces X_i are (finite) simplicial complexes and X_i is a subcomplex of X_j whenever $i \leq j$. When the X_i

are subsets of a common topological space X , with $\bigcap_{i \in T} X_i = \emptyset$ and $\bigcup_{i \in T} X_i = X$, \mathcal{X} is called a *filtration of X* , and the pair (X, \mathcal{X}) is called a *filtered space*.

Applying the homology functor to the sequence of spaces and maps in \mathcal{X} gives one persistence module per homology dimension p , denoted $H_p(\mathcal{X})$. The collection of these modules for p ranging over all dimensions is called the *persistent homology of \mathcal{X}* and denoted $H_*(\mathcal{X})$. Each module $H_p(\mathcal{X})$ has the homology group $H_p(X_i)$ as vector space at every index $i \in T$, and the homomorphism $H_p(X_i) \rightarrow H_p(X_j)$ induced by the inclusion map $X_i \hookrightarrow X_j$ at any pair of indices $i \leq j \in T$. The image $\text{im } H_p(X_i) \rightarrow H_p(X_j)$ is sometimes called a *p -th persistent homology group* of \mathcal{X} , by a natural extension of the concept of p -th homology group to filtrations.

The *persistent cohomology* $H^*(\mathcal{X})$ is defined similarly, but the maps are oriented backwards, so the induced persistence modules are indexed over the ‘backwards copy’ of the real line \mathbb{R} , denoted \mathbb{R}^{op} and defined as follows:

$$(2.1) \quad \mathbb{R}^{\text{op}} = \{\underline{i} \mid i \in \mathbb{R}\} \text{ ordered by } \underline{i} \leq \underline{j} \Leftrightarrow i \geq j.$$

Persistence diagrams. Suppose the modules in $H_*(\mathcal{X})$ have well-defined persistence diagrams. Then, the superimposition of these diagrams in the extended plane, with different indices for different homology dimensions, is called the *persistence diagram* of the filtration \mathcal{X} . It has a decorated version, denoted $\text{Dgm}(\mathcal{X})$, and an undecorated version, denoted $\text{dgm}(\mathcal{X})$. An illustration is provided in Example 2.4 below.

From Chapter 1 (Theorems 1.9 and 1.13) we know that the persistence modules in $H_*(\mathcal{X})$ have well-defined diagrams at least in the following situations:

- (i) when the index set T is finite and the modules $H_*(\mathcal{X})$ are arbitrary,
- (ii) when T is arbitrary and the modules $H_*(\mathcal{X})$ are pointwise finite-dimensional,
- (iii) when $T = \mathbb{R}$ and the modules $H_*(\mathcal{X})$ are \mathfrak{q} -tame.

In case (iii), we say that \mathcal{X} itself is \mathfrak{q} -tame. The kind of filtrations considered originally were finitely simplicial filtrations over finite index sets [118]. These satisfy both (i) and (ii). They are still widely used nowadays because they are the ones handled by algorithms, as we will see in Section 2.

Interpretation. Persistence uses a convenient terminology to read into the persistence diagram of a filtration \mathcal{X} . For this it treats \mathcal{X} not as a family of topological spaces but as a single space ‘transforming’ over time. $\text{Dgm}(\mathcal{X})$ encodes then the evolution of the topological structure of that space as the filtration parameter ranges over the index set T , from low values to high values. Assuming the persistent homology of \mathcal{X} is interval-decomposable, to each decorated point (b^\pm, d^\pm) of $\text{Dgm}(\mathcal{X})$ corresponds an interval in the decomposition of some module $H_p(\mathcal{X})$, called a *p -dimensional feature* of \mathcal{X} . At the topological level, this feature materializes as an independent connected component ($p = 0$), or hole ($p = 1$), or void ($p = 2$), etc. Its *lifespan* is encoded in (b^\pm, d^\pm) , meaning that the feature is present (*alive*) in the current space X_i at every index $i \in [b^\pm, d^\pm] \cap T$ and nowhere else.

EXAMPLE 2.2. Take the filtration \mathcal{X} of the octahedron indexed over $\{1, 2, 3, 4\}$ shown in Figure 2.1. Its persistent homology decomposes as follows:

$$\begin{aligned} H_0(\mathcal{X}) &\cong \mathbb{I}[1, 4] \oplus \mathbb{I}[2, 2] \\ H_1(\mathcal{X}) &\cong \mathbb{I}[3, 3] \\ H_2(\mathcal{X}) &\cong \mathbb{I}[4, 4]. \end{aligned}$$

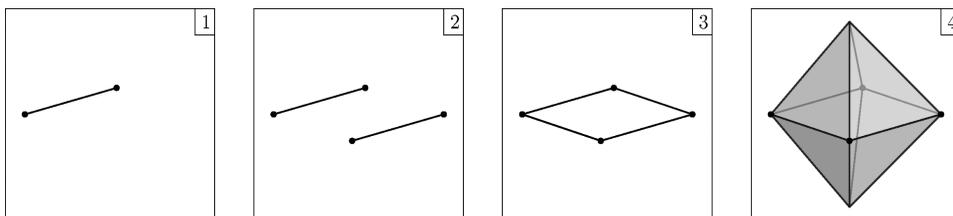


FIGURE 2.1. A simplicial filtration of the octahedron, indexed over the set $\{1, 2, 3, 4\}$.

Using the terminology just introduced, we interpret the decomposition as follows: there is a 0-dimensional feature (connected component) living throughout the filtration, while another 0-dimensional feature lives only at time 2, a 1-dimensional feature (hole) lives only at time 3, and a 2-dimensional feature (void) lives only at time 4.

By extension, we use the same terminology to read into $\mathbf{Dgm}(\mathcal{X})$ when \mathcal{X} is \mathfrak{q} -tame, however it is important to bear in mind that the resulting interpretation will not fully reflect the algebraic structure of $H_*(\mathcal{X})$ if the latter is not interval-decomposable.

Let us take this opportunity to emphasize that persistence provides finer information than just pointwise homology. In the above example, persistence distinguishes between the two connected components that appear in \mathcal{X} and detects that one lives throughout the filtration whereas the other is ephemeral. By contrast, computing the homology of each space X_i in the filtration separately would not make possible to relate the topological features across indices.

Let us also mention that some authors use the concepts of *birth* and *death* of a feature to interpret interval decompositions. For instance, in the above example, one would say that the first connected component is born at time 1 and lives on forever, while the second component is born at time 2 and dies at time 3. This terminology is well-suited for persistence modules over finite index sets, however it is unsafe for persistence modules over general subsets of \mathbb{R} because it assumes implicitly that the intervals in the decomposition can be written in the form $[b, d)$, which is possible in the above example but not always over \mathbb{R} —recall Rule 1.10 and the corresponding discussion in Chapter 1. It must therefore be used with care. As a side note, let us mention that the death of the component born last when two components get merged, as occurs at time 3 in Example 2.2, defines the *elder rule* mentioned in the general introduction.

Excursion sets. Given a topological space X , a *filter* of X is a function $f : X \rightarrow \mathbb{R}$. Each value $i \in \mathbb{R}$ gives rise to two excursion sets: the sublevel set $F_i = f^{-1}((-\infty, i])$, and the superlevel set $F^i = f^{-1}([i, +\infty))$. These excursion sets are usually taken to be closed, and we are following the general trend here.

Consider now the family \mathcal{F}_{\leq} of sublevel sets F_i for i ranging over \mathbb{R} . This family is nested, that is, we have $F_i \subseteq F_j$ whenever $i \leq j$. Moreover, we have $\bigcap_{i \in \mathbb{R}} F_i = \emptyset$ and $\bigcup_{i \in \mathbb{R}} F_i = X$, so \mathcal{F}_{\leq} is a filtration of X , called the *sublevel-sets filtration* of f . One can also consider the *superlevel-sets filtration* \mathcal{F}_{\geq} of f , composed of the superlevel sets F^i and indexed over the ‘backwards copy’ of \mathbb{R} defined in (2.1). By default, when mentioning the ‘persistence of a function’, people

refer to the persistence of its sublevel-sets filtration \mathcal{F}_\leq . The same will be true in the following unless otherwise stated, and we will be using the following vocabulary and notations:

- the decorated and undecorated persistence diagrams of f , denoted respectively $\text{Dgm}(f)$ and $\text{dgm}(f)$, will refer to the ones of \mathcal{F}_\leq ,
- the persistent homology and cohomology of f , denoted $\text{H}_*(f)$ and $\text{H}^*(f)$ respectively, will refer to the ones of \mathcal{F}_\leq ,
- f will be said \mathfrak{q} -tame whenever \mathcal{F}_\leq is.

There are important practical cases in which the function f is naturally \mathfrak{q} -tame. For instance,

PROPOSITION 2.3 (Chazal et al. [72]). *The function $f : X \rightarrow \mathbb{R}$ is \mathfrak{q} -tame in the following cases:*

- (i) *when X is finitely triangulable, i.e. homeomorphic to the underlying space of a finite simplicial complex, and f is continuous,*
- (ii) *when X is locally finitely triangulable, i.e. homeomorphic to the underlying space of a locally finite simplicial complex, and f is continuous, bounded from below, and proper in the sense that the preimage of any compact interval is compact.*

In both cases the proof that $\text{rank } \text{H}_*(F_i) \rightarrow \text{H}_*(F_j) < +\infty$ for $i < j$ relies on inserting some space Y with finitely generated homology between F_i and F_j . More precisely, Y is composed of the simplices of a sufficiently fine subdivision of X that intersect F_i .

The conditions on X and f are tight in that we can build examples where f is no longer \mathfrak{q} -tame when a single one of them is not met. For instance, $X = \mathbb{Z}$ is locally finitely triangulable and $f : n \mapsto n$ is continuous and proper but not bounded from below, and we have $\text{rank } \text{H}_0(F_i) \rightarrow \text{H}_0(F_j) = +\infty$ for any $i \leq j$. This does not mean that the conditions are always necessary though, and in fact there are many cases where they are not met yet f is still \mathfrak{q} -tame. For instance, let $X \subset \mathbb{R}^2$ be the so-called *Hawaiian earring*, i.e. the union of the circles of center $(\frac{1}{n}, 0)$ and radius $\frac{1}{n}$ for n ranging over $\mathbb{N} \setminus \{0\}$, and let $f : X \rightarrow \mathbb{R}$ be the map

$$(x, y) \mapsto \begin{cases} 0 & \text{if } x = 0 \\ \frac{1}{x} & \text{if } x > 0 \end{cases}$$

Then, f is \mathfrak{q} -tame (in fact all its sublevel sets have finite-dimensional 0-homology and trivial higher-dimensional homology) whereas X is not locally triangulable and f is not continuous at the origin.

By contrast, putting stronger conditions on X and f results in finer properties for the persistent homology of f . For instance, if X is a compact manifold and f is a Morse function, or if X is a finitely triangulable space and f is a piecewise-linear function, then there is a finite set of values $c_1 < c_2 < \dots < c_{n-1} < c_n$ such that the inclusion $F_i \hookrightarrow F_j$ induces isomorphisms of finite rank in homology whenever (i, j) does not meet any of these values¹. The persistent homology $\text{H}_*(f)$ is then constant over each of the intervals $(-\infty, c_1)$, $[c_1, c_2)$, \dots , $[c_{n-1}, c_n)$, $[c_n, +\infty)$, and it decomposes into finitely many interval summands of type $\mathbb{I}[c_k, c_l)$ or $\mathbb{I}[c_n, +\infty)$.

¹This is what Cohen-Steiner, Edelsbrunner, and Harer [87] call ‘tameness’ for a function f . It is a more restrictive condition than \mathfrak{q} -tameness, however it guarantees that $\text{H}_*(f)$ is interval-decomposable, with half-open intervals.

The impact on the persistence diagram is that $\text{Dgm}(f)$ is finite, with all decorations pointing to the bottom-left.

Let us give a concrete example, which completes Example 1.11 from Chapter 1.

EXAMPLE 2.4. Consider the planar curve X and its height function $f : X \rightarrow \mathbb{R}$ depicted in Figure 1.2. This is a Morse function over a compact manifold, therefore not only is it \mathfrak{q} -tame, but its persistent homology is interval-decomposable, with the following finite decomposition:

$$\begin{aligned} H_0(f) &\cong \mathbb{I}[a, +\infty) \oplus \mathbb{I}[b, c) \oplus \mathbb{I}[d, e) \\ H_1(f) &\cong \mathbb{I}[m, +\infty), \end{aligned}$$

where $m = \max_{x \in X} f(x)$. The last interval is explained by the fact that a hole appears in the sublevel-set F_i when parameter i reaches the level m , hole that does not disappear subsequently because $F_i = X = F_m$ for all $i \geq m$. Figure 2.2 shows the completed decorated persistence diagram of f sketched in Figure 1.2. Notice the point decorations, which are all pointing to the bottom-left as explained previously.

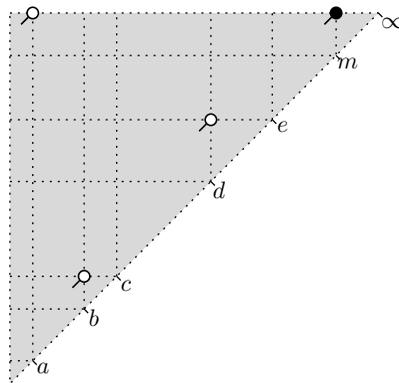


FIGURE 2.2. The full persistence diagram of the height function from Example 1.11. The dimension of each topological feature is coded in the color of the corresponding diagram point: white for 0, black for 1.

— From Chazal et al. [72].

A natural question to ask is whether all \mathbb{R} -indexed filtrations of a topological space X are sublevel-sets filtrations of functions $f : X \rightarrow \mathbb{R}$. This is not exactly the case, as the following simple example shows.

EXAMPLE 2.5. Take for X the discrete space $\{a, b, c\}$, and let \mathcal{X} be the following filtration of X indexed over \mathbb{R} :

$$X_i = \begin{cases} \{a\} & \text{if } i \leq 0 \\ \{a, b\} & \text{if } i \in (0, 1) \\ \{a, b, c\} & \text{if } i \geq 1. \end{cases}$$

Suppose there is a map $f : X \rightarrow \mathbb{R}$ whose sublevel-sets filtration \mathcal{F}_{\leq} equals \mathcal{X} . Then,

$$\begin{aligned} b \in F_i \ \forall i > 0 &\Rightarrow f(b) \leq 0, \\ b \notin F_i \ \forall i < 0 &\Rightarrow f(b) \geq 0, \end{aligned}$$

and so $f(b) = 0$. But then $F_0 \ni b \notin X_0$, which contradicts the assumption that $\mathcal{F}_{\leq} = \mathcal{X}$. The mismatch between the two filtrations at time $i = 0$ can be resolved by taking open sublevel sets instead of closed sublevel sets, but then another mismatch appears at time $i = 1$, with $X_1 \ni c \notin F_1$.

Nevertheless, the persistence diagrams of sublevel-sets filtrations form a dense subset of the persistence diagrams of \mathbb{R} -indexed filtrations, in the following sense²:

PROPOSITION 2.6. *Let \mathcal{X} be a \mathfrak{q} -tame filtration of X . Then, there is a \mathfrak{q} -tame function $f : X \rightarrow \mathbb{R}$ whose undecorated persistence diagram coincides with the one of \mathcal{X} . The decorated diagrams $\text{Dgm}(f)$ and $\text{Dgm}(\mathcal{X})$ may be different though, as in Example 2.5. For f one can take for instance the time of appearance function defined by $f(x) = \inf\{i \in \mathbb{R} \mid x \in X_i\}$.*

Typical applications. Filtrations, in particular sublevel-sets filtrations, are being widely used in topological data analysis. Perhaps their most emblematic use is in inferring the homology of an unknown object from a finite sampling of it, as will be discussed in Chapters 4 and 5.

1.2. Relative filtrations and extended persistence. Suppose we are given a topological space X and a filtration \mathcal{X} of that space, indexed over $T \subseteq \mathbb{R}$. Beside the filtration \mathcal{X} itself, another object of interest is the family of pairs of spaces (X, X_i) for $i \in T$, connected by the componentwise inclusion maps $(X, X_i) \hookrightarrow (X, X_j)$ for $i \leq j \in T$. We will call such a family a *relative filtration*, denoted (X, \mathcal{X}) . The homology functor turns it into a persistence module over T for each homology dimension. The collection of these modules is called the *persistent homology* of the relative filtration (X, \mathcal{X}) , denoted $\text{H}_*(X, \mathcal{X})$. The rest of the terminology introduced for filtrations can be used verbatim for relative filtrations.

Extended persistence. Cohen-Steiner, Edelsbrunner, and Harer [86] introduced extended persistence (EP) as a mean to capture a greater part of the homological information carried by a pair (X, f) . Indeed, some but not all of this information is captured by the persistent homology of the sublevel-sets filtration \mathcal{F}_{\leq} of f . The idea of extended persistence is to grow the space X from the bottom up through the sublevel-sets filtration \mathcal{F}_{\leq} , then to relativize X from the top down with the superlevel-sets filtration \mathcal{F}_{\geq} . The resulting persistence module \mathbb{V}^p at the p -th homology level is indexed over the set

$$\mathbb{R}_{\text{EP}} = \mathbb{R} \cup \{+\infty\} \cup \mathbb{R}^{\text{op}},$$

where \mathbb{R}^{op} is the ‘backwards copy’ of \mathbb{R} defined in (2.1). The order on \mathbb{R}_{EP} is completed by $i < +\infty < \underline{j}$ for all $i, j \in \mathbb{R}$. As a poset, \mathbb{R}_{EP} is isomorphic to (\mathbb{R}, \leq) ,

²This result is a straight consequence of the stability theorem that will be presented in Chapter 3.

but we will keep it as it is for clarity. The module \mathbb{V}^p is defined as follows:

$$(2.2) \quad \begin{aligned} V_i^p &= H_p(F_i) && \text{for } i \in \mathbb{R} \\ V_{+\infty}^p &= H_p(X) \cong H_p(X, \emptyset) \\ V_{\underline{i}}^p &= H_p(X, F^i) && \text{for } \underline{i} \in \mathbb{R}^{\text{op}}, \end{aligned}$$

the morphisms between these spaces being induced by the inclusion maps $\emptyset \hookrightarrow F_i \hookrightarrow F_j \hookrightarrow X$ and $\emptyset \hookrightarrow F^j \hookrightarrow F^i \hookrightarrow X$ for $i \leq j \in \mathbb{R}$. When this module is \mathfrak{q} -tame, e.g. when X is finitely triangulable and f is continuous, its decorated persistence diagram $\text{Dgm}(\mathbb{V}^p)$ is well-defined everywhere in the extended plane except on the diagonal Δ , by a straightforward adaptation of Theorem 1.13. The superimposition of these diagrams forms the so-called *extended persistence diagram of f* , which contains three types of points:

- points with coordinates $i \leq j \in \mathbb{R}$, which belong also to the diagram of the sublevel-sets filtration \mathcal{F}_{\leq} and lie above Δ ,
- points with coordinates $\underline{i} \leq \underline{j} \in \mathbb{R}^{\text{op}}$, which belong also to the diagram of the relative filtration (X, \mathcal{F}_{\geq}) and lie below Δ ,
- points with coordinates $i \in \mathbb{R}, \underline{j} \in \mathbb{R}^{\text{op}}$, which lie on either side of Δ .

These three types of points form respectively the *ordinary*, the *relative*, and the *extended* parts of the extended persistence diagram of f .

EXAMPLE 2.7. Consider once again the planar curve X and its height function $f : X \rightarrow \mathbb{R}$ depicted in Figure 1.2. This is a Morse function over a compact manifold, therefore the extended persistence modules \mathbb{V}^p are interval-decomposable, with the following finite decompositions:

$$\begin{aligned} \mathbb{V}^0 &\cong \mathbb{I}[a, \underline{m}] \oplus \mathbb{I}[b, c] \oplus \mathbb{I}[d, e] \\ \mathbb{V}^1 &\cong \mathbb{I}[m, \underline{a}] \oplus \mathbb{I}[\underline{c}, \underline{b}] \oplus \mathbb{I}[\underline{e}, \underline{d}], \end{aligned}$$

where $m = \max_{x \in X} f(x)$. The corresponding diagram is shown in Figure 2.3.

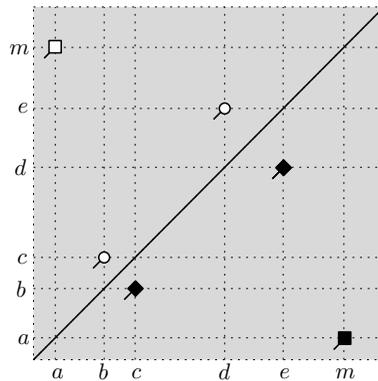


FIGURE 2.3. The extended persistence diagram of the height function from Example 1.11. The dimension of each topological feature is encoded in the color of the corresponding diagram point: white for 0, black for 1. The shape of each diagram point indicates to which part of the diagram it belongs: the ordinary part for a round shape, the relative part for a diamond shape, the extended part for a square shape.

Notice the symmetry in the diagram of Figure 2.3. It is a general property of Morse functions over compact manifolds, and a consequence of the Lefschetz duality between absolute and relative homology groups of complementary dimensions:

THEOREM 2.8 (EP Duality [86]). *Let f be a Morse function over a compact d -manifold. Then, the extended persistence diagram of f has the following reflection symmetries across the diagonal Δ : for all $0 \leq p \leq d$,*

- *the ordinary part of $\mathrm{Dgm}(\mathbb{V}^p)$ and the relative part of $\mathrm{Dgm}(\mathbb{V}^{d-p})$ are mirror of each other,*
- *the extended parts of $\mathrm{Dgm}(\mathbb{V}^p)$ and $\mathrm{Dgm}(\mathbb{V}^{d-p})$ are mirror of each other,*
where the \mathbb{V}^p are the extended persistence modules derived from f according to (2.2).

From this result follows another one relating the extended persistence diagrams of f and $-f$. This time the reflections are performed across the minor diagonal $\{(x, -x) \mid x \in \mathbb{R}\}$:

THEOREM 2.9 (EP Symmetry [86]). *Let f be a Morse function over a compact d -manifold. Then, the extended persistence diagrams of f and $-f$ are mirror of each other across the minor diagonal $\{(x, -x) \mid x \in \mathbb{R}\}$ in the following sense:*

- *the ordinary parts of $\mathrm{Dgm}(\mathbb{V}^p)$ and $\mathrm{Dgm}(\mathbb{W}^{d-p-1})$ are mirror of each other,*
- *the relative parts of $\mathrm{Dgm}(\mathbb{V}^p)$ and $\mathrm{Dgm}(\mathbb{W}^{d-p-1})$ are mirror of each other,*
- *the extended parts of $\mathrm{Dgm}(\mathbb{V}^p)$ and $\mathrm{Dgm}(\mathbb{W}^{d-p})$ are mirror of each other,*
where the \mathbb{V}^p and \mathbb{W}^p are the extended persistence modules derived respectively from f and $-f$.

Edelsbrunner and Kerber [117] make even finer claims in the special case where $X = \mathbb{S}^d$. Given a Morse function $f : \mathbb{S}^d \rightarrow \mathbb{R}$ and two subsets $L, W \subset \mathbb{S}^d$ —the *land* and the *water*—such that $L \cup W = \mathbb{S}^d$ and $L \cap W$ is an $(n-1)$ -submanifold of \mathbb{S}^d , their *land and water theorem* shows (under some conditions) that a symmetry exists between the persistence diagrams of $f|_L$ and $f|_W$, which follows from Alexander duality.

These symmetries are the gem of extended persistence theory, generalizing classical duality theorems from manifolds to real-valued functions over manifolds in a way that is both natural and elegant. As we will see in the next section (Theorem 2.11), they can be further generalized to certain classes of functions over non-manifold spaces.

Typical applications. Extended persistence has been used by Bendich et al. [22] to define multiscale variants of the concept of local homology group at a point in a topological space. These are useful e.g. to infer the local structure and dimension of an unknown stratified space from a finite sampling of it, which was the main motivation in the paper.

1.3. Zigzags and level sets. A *zigzag* \mathcal{Z} is a diagram of topological spaces of the following form, where the maps are inclusions and can be oriented arbitrarily:

$$(2.3) \quad Z_1 \text{ --- } Z_2 \text{ --- } \cdots \text{ --- } Z_{n-1} \text{ --- } Z_n.$$

Thus, zigzags are a special type of representations of A_n -type quivers in the category of topological spaces. Applying the homology functor to a zigzag \mathcal{Z} gives one zigzag module per homology dimension p , denoted $H_p(\mathcal{Z})$. The collection of these modules is called the *persistent homology* of \mathcal{Z} , denoted $H_*(\mathcal{Z})$. One can also consider its

persistent cohomology $H^*(\mathcal{Z})$, obtained by applying the cohomology functor. According to Theorem 1.9 (i), the modules in $H_*(\mathcal{Z})$ are always interval-decomposable, so their persistence diagrams are well-defined. Their superimposition in the plane, with different indices for different homology dimensions, is called the *persistence diagram* of the zigzag \mathcal{Z} . Thus, most of the terminology introduced for filtrations carries over to zigzags.

Level-sets zigzag. Given a topological space X and a function $f : X \rightarrow \mathbb{R}$, each interval $[i, j] \subset \mathbb{R}$ gives rise to a *slice* $F_i^j = f^{-1}([i, j])$. The slice F_i^i is called a *level set*. Given a sequence of real values $i_0 < i_1 < \dots < i_n$, consider the following diagram where all maps are inclusions between slices:

$$(2.4) \quad F_{i_0}^{i_0} \rightarrow F_{i_0}^{i_1} \leftarrow F_{i_1}^{i_1} \rightarrow F_{i_1}^{i_2} \leftarrow \dots \rightarrow F_{i_{n-1}}^{i_n} \leftarrow F_{i_n}^{i_n}.$$

This diagram is called a *level-sets zigzag* of f , denoted \mathcal{Z} generically. There are as many such diagrams as there are finite increasing sequences of real values. However, in some cases there is a class of such zigzags that is ‘canonical’. For instance, when X is a compact manifold and f is a Morse function, there is a finite set of values $c_1 < \dots < c_n$ such that the preimage of each open interval $S = (-\infty, c_1), (c_1, c_2), \dots, (c_{n-1}, c_n), (c_n, +\infty)$ is homeomorphic to a product of the form $Y \times S$, where Y has finite-dimensional homology and f is the projection onto the factor S . Then, choosing $i_0 < \dots < i_n$ that are interleaved with the c_i ’s as follows:

$$(2.5) \quad i_0 < c_1 < i_1 < c_2 < \dots < i_{n-1} < c_n < i_n,$$

we have that, up to isomorphism, the persistent homology of \mathcal{Z} is independent of the choice of indices i_0, \dots, i_n . In such a case, \mathcal{Z} is called *the* level-sets zigzag of f : even though it is not unique, its persistent homology is (up to isomorphism). Note that the assumptions X compact and f Morse are not exploited in full here. One obtains the same characterization for \mathcal{Z} from any function that is of *Morse type* in the following sense:

DEFINITION 2.10 (Carlsson, de Silva, and Morozov [50]). Given a topological space X , a function $f : X \rightarrow \mathbb{R}$ is of *Morse type* if there is a finite set of values $c_1 < \dots < c_n$ such that the preimage of each open interval $S = (-\infty, c_1), (c_1, c_2), \dots, (c_{n-1}, c_n), (c_n, +\infty)$ is homeomorphic to a product of the form $Y \times S$, where Y has finite-dimensional homology and f is the projection onto the factor S . Moreover, each homeomorphism $Y \times S \rightarrow f^{-1}(S)$ should extend to a continuous function $Y \times \bar{S} \rightarrow F_{c_i}^{c_{i+1}}$.

The pyramid. Assume $f : X \rightarrow \mathbb{R}$ is of Morse type as in Definition 2.10, and let $i_0 < \dots < i_n$ be chosen as in (2.5). Then, one can build a gigantic lattice of (relative) homology groups of slices, where the maps are induced by inclusions. With some imagination, this lattice looks like a pyramid of apex $H_*(F_{i_0}^{i_n})$ seen from the top—hence its name. It is shown for $n = 3$ in Figure 2.4.

This pyramid was introduced by Carlsson, de Silva, and Morozov [50] as a tool to relate the level-sets zigzag \mathcal{Z} of f to its extended persistence. Indeed, the persistent homology of \mathcal{Z} appears in the bottom row of the lattice, while the extended persistent homology of f appears³ in the diagonal $H_*(F_{i_0}^{i_0}) \rightarrow \dots \rightarrow H_*(F_{i_0}^{i_n}) \rightarrow H_*(F_{i_0}^{i_n}, F_{i_n}^{i_n}) \rightarrow \dots \rightarrow H_*(F_{i_0}^{i_n}, F_{i_0}^{i_n}) = 0$, and the extended persistent

³Thanks to the fact that f is of Morse type, its extended persistent homology and the diagonal of the pyramid have the same interval decomposition up to some shifts in the interval endpoints.

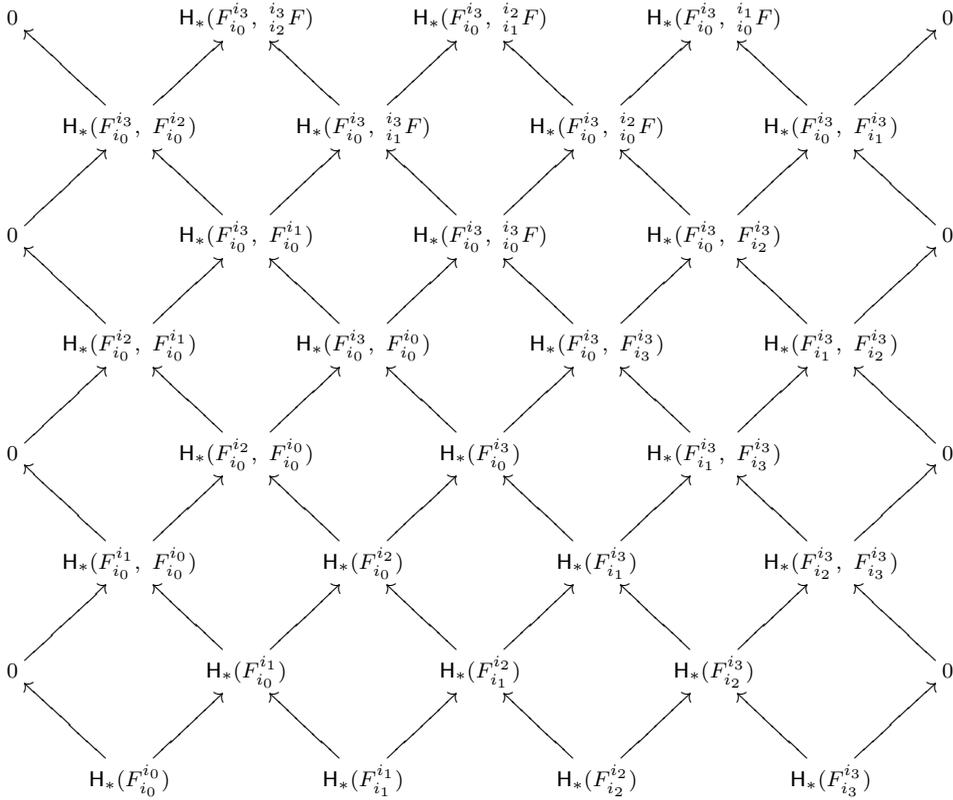
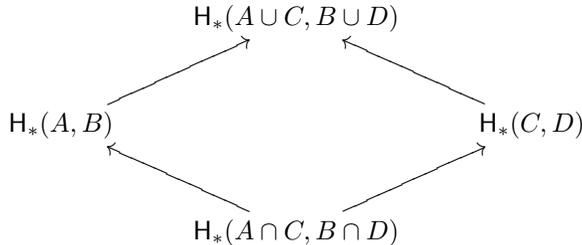


FIGURE 2.4. The pyramid for $n = 3$. Each homology group $H_*(F_{i_r}^{i_s})$ is identified with the relative homology group $H_*(F_{i_r}^{i_s}, \emptyset)$. Moreover, the $H_*(F_{i_r}^{i_s}, F_{i_r}^{i_s})$ are identified with 0, and ${}^{i_s}F$ stands for the union $F_{i_0}^{i_r} \cup F_{i_s}^{i_n}$.
 — Based on Carlsson, de Silva, and Morozov [50].

homology of $-f$ in the minor diagonal $0 = H_*(F_{i_0}^{i_n}, F_{i_0}^{i_n}) \leftarrow \dots \leftarrow H_*(F_{i_0}^{i_n}, F_{i_0}^{i_0}) \leftarrow H_*(F_{i_0}^{i_n}, F_{i_0}^{i_n}) \leftarrow \dots \leftarrow H_*(F_{i_n}^{i_n})$.

The remarkable property of the pyramid is that all the diamonds are of the following type, called *Mayer-Vietoris*:



There is then a bijection between the interval decompositions of any two x -monotone paths in the pyramid that differ by one diamond. The bijection may slightly shift

the intervals' endpoints or their associated homology dimensions, however it does match the two decompositions completely. This is the so-called *Mayer-Vietoris Diamond Principle*⁴ of Carlsson and de Silva [49].

Thus, by 'travelling down' the pyramid, from either the diagonal or the minor diagonal, to the bottom zigzag, and by applying the Mayer-Vietoris Diamond Principle at every diamond encountered, one can derive bijections between the decompositions of the extended persistent homologies of f and $-f$ and the decomposition of $H_*(\mathcal{Z})$. The details of these bijections are not important, what matters is that they generalize the EP Symmetry Theorem 2.9 from Morse functions over manifolds to Morse-type functions over arbitrary topological spaces:

THEOREM 2.11 (Pyramid [50]). *Let $f : X \rightarrow \mathbb{R}$ be of Morse type. Then, up to bijections, the extended persistent homologies of f and $-f$ have the same interval decomposition as the persistent homology of the level-sets zigzag of f .*

Typical applications. Beside making possible to generalize the EP Symmetry Theorem, zigzags have been used as a theoretical tool in the analysis of various homology inference techniques [107, 221], and as a smaller-sized data structure to improve their efficiency [208]. We will elaborate on this in Chapters 5 and 7.

1.4. Kernels, images, and cokernels. Suppose we are given two filtrations \mathcal{X}, \mathcal{Y} , together with the promise that \mathcal{X} is 'dominated' by \mathcal{Y} in the sense that $X_i \subseteq Y_i$ for all $i \in T$. Then, the family of inclusion maps $X_i \hookrightarrow Y_i$ induces a morphism between persistence modules $H_p(\mathcal{X}) \rightarrow H_p(\mathcal{Y})$ at the p -th homology level, for any homology dimension p . The kernel, image, and cokernel of this morphism are persistence modules themselves, and they have an interest of their own. For instance, they are \mathfrak{q} -tame when both \mathcal{X} and \mathcal{Y} are, but not only: they can also sometimes be \mathfrak{q} -tame when only \mathcal{X} , or only \mathcal{Y} , or neither, is.

Typical applications. Persistence for kernels, images and cokernels has been introduced by Cohen-Steiner et al. [89], who used it to refine the method of Bendich et al. [22] for measuring the local homology of a sampled topological space. As we will see in Chapter 6, it can also be used effectively to approximate the persistence diagram of a real-valued function from a finite sampling of its domain.

2. Calculations

Algorithms that compute persistence take in a (finitely) simplicial filtration \mathcal{K} indexed over a finite index set $T \subset \mathbb{R}$. Let $t_1 < \dots < t_n$ be the elements of T , and let $K = K_{t_n}$, so the filtration \mathcal{K} is written as follows:

$$(2.6) \quad \emptyset \subseteq K_{t_1} \subseteq K_{t_2} \subseteq \dots \subseteq K_{t_n} = K.$$

Define the *time of appearance* of a simplex $\sigma \in K$ to be

$$t(\sigma) = \min\{t_i \mid \sigma \in K_{t_i}\}.$$

The level sets of t are the $K_{t_i} \setminus K_{t_{i-1}}$, where by convention we let $K_{t_0} = \emptyset$. The order induced by t on the simplices of K is only partial because it is unspecified within each level set. However, it is compatible with the incidence relations in K ,

⁴This is the topological counterpart of the Diamond Principle presented in Section 4.4 of Appendix A. It provides a complete matching between the interval decompositions of the two x -monotone paths considered. Its proof combines the Diamond Principle with the Mayer-Vietoris exact sequence associated to the central diamond—see [49, §5.3] for the details.

that is, a simplex never appears in \mathcal{K} before its faces. Pick then any total order on the simplices of K that is compatible both with t and with the incidence relations within its level sets, and label the simplices of K accordingly:

$$(2.7) \quad K = \{\sigma_1, \sigma_2, \dots, \sigma_m\}, \text{ where } i < j \text{ whenever } t(\sigma_i) < t(\sigma_j) \\ \text{or } \sigma_i \text{ is a proper face of } \sigma_j.$$

The sequence of simplices $\sigma_1, \sigma_2, \dots, \sigma_m$ is the actual input taken in by persistence algorithms. It defines the following simplicial filtration of K , denoted \mathcal{K}^σ :

$$\emptyset \subseteq \{\sigma_1\} \subseteq \{\sigma_1, \sigma_2\} \subseteq \dots \subseteq \{\sigma_1, \dots, \sigma_m\} = K.$$

\mathcal{K}^σ is a refinement of the original filtration \mathcal{K} . It is not strictly equivalent to \mathcal{K} in terms of persistence, however it is closely related to it. The relationship is best described in terms of half-open intervals⁵:

LEMMA 2.12 (folklore). *There is a partial matching between the interval decompositions of $H_*(\mathcal{K})$ and $H_*(\mathcal{K}^\sigma)$ such that:*

- *summands $\mathbb{I}[i, +\infty)$ of $H_p(\mathcal{K}^\sigma)$ are matched with summands $\mathbb{I}[t(\sigma_i), +\infty)$ of $H_p(\mathcal{K})$, and vice-versa,*
- *summands $\mathbb{I}[i, j)$ of $H_p(\mathcal{K}^\sigma)$ where $t(\sigma_i) < t(\sigma_j)$ are matched with summands $\mathbb{I}[t(\sigma_i), t(\sigma_j))$ of $H_p(\mathcal{K})$, and vice-versa,*
- *all other summands of $H_p(\mathcal{K}^\sigma)$, i.e. summands $\mathbb{I}[i, j)$ where $t(\sigma_i) = t(\sigma_j)$, are unmatched.*

This result can be viewed as an instance of the *snapping principle* that will be presented in Chapter 3. It asserts that $H_*(\mathcal{K}^\sigma)$ carries a superset of the persistence information carried by $H_*(\mathcal{K})$, and that there is a simple rule to recover the interval decomposition of $H_*(\mathcal{K})$ from the one of $H_*(\mathcal{K}^\sigma)$. Hence, from now on we will work with \mathcal{K}^σ instead of \mathcal{K} .

Since every simplex of K has its own index in \mathcal{K}^σ , each interval summand $\mathbb{I}[i, j)$ in the decomposition of $H_*(\mathcal{K}^\sigma)$ defines a *pairing* between simplices σ_i and σ_j of K . In persistence terms, the topological feature represented by the interval summand is created by σ_i and then destroyed by σ_j in \mathcal{K}^σ . Hence, σ_i is said to be *creating*, or *positive*, while σ_j is said to be *destroying*, or *negative*. Meanwhile, each summand $\mathbb{I}[i, +\infty)$ in the decomposition represents a topological feature created by the (thus positive) unpaired simplex σ_i , and living throughout the rest of the filtration \mathcal{K}^σ . This feature is called *essential* because it corresponds to a non-trivial element in $H_*(K)$.

It is easy to see that every simplex must be either positive or negative—this is in fact a consequence of the matrix reduction algorithm that will be presented next. The partial pairing of the simplices of K describes the persistence of \mathcal{K}^σ completely and is the information sought for by algorithms.

In Section 2.1 below we present the basic matrix reduction algorithm, then in Section 2.2 we review its improvements and extensions.

⁵For this we are implicitly extending \mathcal{K} to a filtration over \mathbb{R} by putting the space K_{t_i} at every real index $x \in [t_i, t_{i+1})$, the space K at every index $x \geq t_n$, and the empty space at every index $x < t_1$. The summands in the interval decomposition of $H_*(\mathcal{K})$ are then transformed into interval modules over \mathbb{R} by the same process, and by construction they still decompose $H_*(\mathcal{K})$. We apply the same implicit extension to \mathcal{K}^σ .

2.1. Matrix reduction. At a high level, the approach consists in performing Gaussian elimination on the matrix of the boundary operator of K while preserving the column and row orders prescribed by the input filtration.

Details of the algorithm. Given a labelling of the simplices of K as in (2.7), and the corresponding filtration \mathcal{K}^σ of K , the algorithm builds a square $m \times m$ matrix M representing the boundary operator of K . For ease of exposition we assume that the field of coefficients is \mathbb{Z}_2 , so M is a binary matrix. Specifically, M has one row and one column per simplex of K , with $M_{ij} = 1$ if σ_i is a face of codimension 1 of σ_j and $M_{ij} = 0$ otherwise. Moreover, for the needs of the algorithm, the rows and columns of M are ordered as the simplices in the sequence of (2.7). Since the sequence is compatible with the incidence relations in K , the matrix M is upper-triangular.

Once M is built, the algorithm processes its columns once from left to right. When processing the j -th column, it keeps adding columns from the left until the following loop invariant is satisfied: the submatrix spanned by columns 1 through j has the property that the lowest nonzero entry of every nonzero column lies in a unique row. The pseudo-code is given in Algorithm 2.1, where $\text{low}(j, M)$ denotes the row index of the lowest nonzero entry in column j of the matrix M — $\text{low}(j, M) = 0$ if column j is entirely zero.

Algorithm 2.1: Matrix reduction

Input: $m \times m$ binary matrix M

```

1 for  $j = 1$  to  $m$  do
2   while there exists  $k < j$  with  $\text{low}(k, M) = \text{low}(j, M) \neq 0$  do
3     | add (modulo 2) column  $k$  to column  $j$  in  $M$ ;
4   | end
5 end

```

Output: the reduced matrix M

Upon termination, the matrix M has the property that the lowest nonzero entry of every nonzero column lies in a unique row. Its structure is then interpreted as follows:

- every zero (resp. nonzero) column j corresponds to a positive (resp. negative) simplex σ_j ,
- every nonzero column j is paired with the column $i = \text{low}(j, M)$ and gives rise to a summand $\mathbb{I}[i, j)$ in the interval decomposition of $\mathbf{H}_*(\mathcal{K}^\sigma)$,
- every remaining unpaired zero column j gives rise to a summand $\mathbb{I}[j, +\infty)$ in the decomposition.

EXAMPLE 2.13. Take the simplicial filtration \mathcal{K}^σ shown in Figure 2.5. From the reduced boundary matrix we can read off the interval decomposition of the persistent homology of \mathcal{K}^σ :

$$\begin{aligned} \mathbf{H}_0(\mathcal{K}^\sigma) &\cong \mathbb{I}[1, +\infty) \oplus \mathbb{I}[2, 4) \oplus \mathbb{I}[3, 5) \\ \mathbf{H}_1(\mathcal{K}^\sigma) &\cong \mathbb{I}[6, 7) \end{aligned}$$

Observe that the essential feature $\mathbb{I}[1, +\infty)$ gives the homology of the solid triangle as expected.

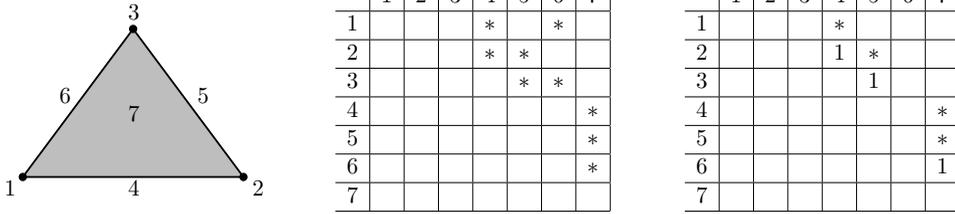


FIGURE 2.5. Left: filtration of a solid triangle (each simplex is marked with its index in the sequence). Center: the corresponding boundary matrix M , where simplices are identified with their index in the sequence, and where stars mark the nonzero entries. Right: the reduced matrix, where lowest nonzero entries are marked with 1's while the other nonzero entries are marked with stars.

— Based on Edelsbrunner and Harer [115].

Correctness. The crux of the matter is to prove the previous interpretation correct. Formally:

THEOREM 2.14 (de Silva, Morozov, and Vejdemo-Johansson [99]). *Upon termination, the simplicial chains $\hat{\sigma}_1, \dots, \hat{\sigma}_n$ represented by the columns of M yield a partition $F \sqcup G \sqcup H$ of the index set $\{1, \dots, m\}$, and the lowest nonzero entries in the columns induce a bijective pairing $G \leftrightarrow H$ such that:*

- (i) *For every index i , the chains $\hat{\sigma}_1, \dots, \hat{\sigma}_i$ form a basis of the chain group of the complex $\{\sigma_1, \dots, \sigma_i\}$,*
- (ii) *For every index $f \in F$, $\partial \hat{\sigma}_f$ is a cycle, i.e. $\partial \hat{\sigma}_f = 0$,*
- (iii) *For every pair of indices (g, h) given by the pairing $G \leftrightarrow H$, $\partial \hat{\sigma}_h = \hat{\sigma}_g$ and so $\partial \hat{\sigma}_g = 0$.*

Item (i) is equivalent to the assertion that the leading term of each simplicial chain $\hat{\sigma}_i$ is σ_i . Item (ii) asserts that the set F identifies the positive simplices that do not get paired. Item (iii) asserts that the set G identifies the positive simplices that do get paired, while the set H identifies the corresponding negative simplices.

The proof of the theorem is a sequence of simple lemmas. Although there is no record in the official literature, it can be found in some course notes such as [29]. Here we prefer to give the following high-level interpretation.

High-level interpretation. As pointed out after Theorem 1.4, the persistent homology $H_*(\mathcal{K}^\sigma)$ can be viewed⁶ as a \mathbb{Z} -graded module over the polynomial ring $\mathbf{k}[t]$, so the classical structure theorem for finitely generated graded modules over graded principal ideal domains guarantees the existence and uniqueness of a decomposition into interval summands. Furthermore, the proof of the structure theorem for a given module \mathbb{V} shows that the decomposition is given by the Smith normal form of the matrix of some module homomorphism going from the free module spanned by the generators of \mathbb{V} onto \mathbb{V} . In the context of standard, non-persistent homology over the integers, the homomorphism is induced by the boundary operator of the simplicial complex under consideration, whose coefficients are in \mathbb{Z} —see e.g. [202, §1.11] for the details. In the context of persistent homology over a field \mathbf{k} , the

⁶Indeed, $H_*(\mathcal{K}^\sigma)$ can be extended to a representation of the poset \mathbb{Z} by adding zero vector spaces and maps outside the range $[1, m]$, so the remark following Theorem 1.4 applies.

homomorphism is induced by the boundary operator of the simplicial filtration \mathcal{K}^σ , whose coefficients are in $\mathbf{k}[t]$. This provides us with an algorithm for computing persistence: build the matrix of the boundary operator of \mathcal{K}^σ , then compute its Smith normal form over $\mathbf{k}[t]$ using Gaussian elimination on the rows and columns, then read off the interval decomposition. The algorithm is the same as for non-persistent homology, except the ring of coefficients is $\mathbf{k}[t]$ instead of \mathbb{Z} .

This approach for computing persistence was first suggested by Zomorodian and Carlsson [243], who also observed that it is in fact sufficient to reduce the boundary matrix to a column-echelon form over the field \mathbf{k} itself, which has two benefits: first, Gaussian elimination only has to be applied on the columns of the matrix; second, and more importantly, there is no more need to compute greatest common divisors between polynomials. The catch is that now the insertion order of the simplices must be preserved, so no swaps between columns are allowed, which means that the reduced matrix is in column-echelon form up to a permutation of the columns only. When $\mathbf{k} = \mathbb{Z}_2$, this simplified approach coincides with the original persistence algorithm designed by Edelsbrunner, Letscher, and Zomorodian [118] for subcomplexes of the sphere \mathbb{S}^3 . It is the version presented here.

Complexity. The time complexity of the algorithm is at most cubic in the number m of simplices of K . To see this, observe that the j -th iteration of the `for` loop modifies only the j -th column of the matrix, therefore every column $k < j$ is considered at most once by the inner `while` loop at that iteration. Moreover, each column k can store the row number of its lowest entry once and for all after the k -th iteration of the `for` loop, so finding k such that $\text{low}(k, M) = \text{low}(j, M)$ at a subsequent iteration j can be done in time $O(j)$. Thus, the running time of the algorithm is indeed $O(m^3)$. A more careful analysis—see e.g. [115, §VII.2]—using a sparse matrix representation for M gives a tighter running-time bound⁷ in $O(\sum_{j=1}^m p_j^2)$, where $p_j = j$ if simplex σ_j is unpaired and $p_j = j - i$ if σ_j is paired with σ_i . The quantity $\sum_{j=1}^m p_j^2$ is of the order of m^3 in the worst case, and Morozov [199] has worked out a worst-case example on which the total running time is indeed cubic. However, $\sum_{j=1}^m p_j^2$ is typically much smaller than that in practice, where worst-case scenarios are unlikely to occur and a near-linear running time is usually observed.

Implementations. To our knowledge, all existing implementations of the persistence algorithm work with finite fields of coefficients (some only with \mathbb{Z}_2) to avoid numerical issues. There currently are two implementations that follow the guidelines of [115, §VII.2] using a sparse matrix representation: the first one, by A. Zomorodian, was destined to be integrated into the C++ library CGAL (<http://www.cgal.org>) but was never released; the second one, by D. Morozov, was released as part of the C++ library DIONYSUS (<http://www.mrv.org/software/dionysus/>). The other implementations use various tricks to sparsify the data structures and to optimize for speed, as we will see next.

2.2. Further developments. A lot of progress has been made since the introduction of the original matrix reduction algorithm. Every new contribution has aimed at a specific type of improvement, whether it be extending the approach to other topological constructions, or increasing its practical efficiency in terms of running time and memory usage, or reducing the theoretical complexity of the problem.

⁷This bound assumes the dimension of each simplex to be constant.

2.2.1. *Extensions of the matrix reduction algorithm.* Cohen-Steiner, Edelsbrunner, and Harer [86] adapted the matrix reduction algorithm so it computes extended persistence, while Cohen-Steiner et al. [89] adapted it to compute persistence for kernels, images and cokernels. These extensions are fairly straightforward, casting the considered problems into that of the standard matrix reduction. Other extensions include:

Persistence over several coefficient fields. Boissonnat and Maria [32] adapted the method so that it can compute persistence over several finite fields of prime order at once in a single matrix reduction. Their approach is based on a clever use of the Chinese Remainder Theorem. The gain in practice is not only in terms of running time, but also in terms of richness of the output. Indeed, combined with the Universal Coefficient Theorem, their approach provides information about the prime divisors of the torsion coefficients of the integral homology groups, which calculations over a single field typically do not provide.

Zigzag persistence. Carlsson, de Silva, and Morozov [50] considered the problem of computing zigzag persistence from an input simplicial zigzag. Adapting the matrix reduction to this context is subtle, partly because the column order is not known in advance. The adaptation proposed by Carlsson, de Silva, and Morozov [50] follows the proof of Gabriel's theorem devised by Carlsson and de Silva [49], which we outline for the interested reader in Section 4.4 of Appendix A. In practice it comes down to scanning through the input zigzag once from left to right, loading a single complex in main memory at a time, and performing sequential simplex insertions and deletions while maintaining the so-called *right filtration* and *birth index* of the zigzag.

Vineyards. Cohen-Steiner, Edelsbrunner, and Morozov [85] considered the problem of updating the interval decomposition of the persistent homology of \mathcal{K}^σ as the order in the sequence of simplices changes. They showed how the reduced matrix M can be updated in time $O(m)$ after a *simple transposition*, i.e. a transposition between consecutive simplices in the sequence. From there, any permutation of the simplex order can be handled by decomposing it into a sequence of simple transpositions. In the worst case, the number of simple transpositions can be up to quadratic in m and therefore induce a cubic total update time, which is as bad as recomputing persistence from scratch. However, in practical situations the number of simple transpositions can be much smaller. This approach is particularly well suited for tracking the evolution of the persistence diagram of a function $f : K \rightarrow \mathbb{R}$ that changes continuously over time. Indeed, continuous changes in f induce a series of simple transpositions, whose effect on the diagram can be computed efficiently. By stacking up all the instances of the diagram over time, one obtains a series of z -monotone trajectories⁸, one for each diagram point. Each trajectory is seen as a *vine*, and the entire collection is called the *vineyard* of f . Further details together with examples of practical applications can be found in [199].

Persistent cohomology. Instead of applying the homology functor to the filtration \mathcal{K}^σ , one can apply the cohomology functor and get a reversed sequence of vector spaces (the cohomology groups) and linear maps induced by inclusions.

⁸These trajectories are continuous, as guaranteed by the Stability Theorem that will be presented in Chapter 3.

de Silva, Morozov, and Vejdemo-Johansson [100] adapted the matrix reduction technique so it computes the interval decomposition of such a persistence module, using the same ingredients as the zigzag persistence algorithm of Carlsson, de Silva, and Morozov [50]. The beautiful thing about this approach is that homology groups and cohomology groups over a field are dual to each other, and that the universal coefficient theorem [152] implies that the persistence diagram obtained by the cohomology algorithm is in fact the same as the one obtained by the standard persistence algorithm. As both algorithms have the same cubic worst-case running time, the comparison of their respective performances is done on practical data. de Silva, Morozov, and Vejdemo-Johansson [99] have reported the cohomology algorithm to outperform the standard homology algorithm in this respect, although this advantage seems to fade away when comparing optimized versions. In particular, it seems to be highly dependent on the data considered.

Simplex annotations. An elegant description of the cohomology algorithm, using the notion of *simplex annotation*, has been given by Dey, Fan, and Wang [107]. The idea is to compute the decomposition of the persistent homology of \mathcal{K}^σ by maintaining cohomology bases instead of homology bases. Each basis element is encoded by the values it takes on the simplices of K , and the collection of these values for a simplex forms its annotation. Beside its simplicity and efficiency, this approach applies to more general sequences of simplicial complexes and maps, such as for instance sequences of inclusions and edge contractions. Although this new setting reduces easily to the standard setting described in (2.6) by turning each map into an inclusion in its mapping cylinder [199], the cost of building the cylinders of all the maps and to glue them together can be high. Therefore, working directly with the maps themselves is relevant, all the more as the overall complexity of the approach is proved to remain controllable.

2.2.2. Efficient implementations. As we will see in forthcoming chapters, the total number m of simplices in the filtered simplicial complex K can be huge in practice, so it is highly recommended to use time-wise and memory-wise efficient implementations to compute persistence. Among the few currently available implementations, there are two main competitors in terms of practical efficiency:

Simplex tree and compressed annotation matrix. Boissonnat and Maria [33] introduced a lightweight data structure to represent simplicial complexes, called the *simplex tree*. Roughly speaking, this is a spanning tree of the Hasse diagram of the simplicial complex K , derived from the lexicographical order on the sorted vertex sequences representing the simplices of K . It has size $O(m)$ while the full Hasse diagram has size $O(dm)$, where d is the dimension of K and m is its number of simplices. Combined with a compressed version of the annotation matrix of Dey, Fan, and Wang [107] developed by Boissonnat, Dey, and Maria [30], this data structure gives the currently most lightweight sequential algorithm to compute persistence. Its time performance is on par with the state of the art, outperforming the standard persistence algorithm by one or two orders of magnitude. The implementation is available in the C++ library GUDHI (<http://pages.saclay.inria.fr/vincent.rouvreau/gudhi/>).

Clear and compress. Bauer, Kerber, and Reininghaus [17] presented a highly parallelizable version of the matrix reduction algorithm. Their version differs from the standard version by first computing persistence pairs within local chunks, then simplifying the unpaired columns, and finally applying standard reduction on the

simplified matrix. Their implementation is available in the C++ library PHAT (<http://phat.googlecode.com/>). The sequential code is reported to outperform the standard reduction algorithm and to be competitive to other variants with a good practical behavior. The parallelized version yields reasonable speed-up factors and demonstrates the usefulness of parallelization for this problem.

2.2.3. Other approaches and theoretical optimizations. As shown by Morozov [199], the cubic upper bound for the worst-case time complexity of the matrix reduction algorithm is tight. However, the optimal bound for the persistence computation problem itself still remains unknown. The following contributions have shed new light on this question. The first item shows that the complexity of the problem in full generality is subcubic. The second item shows that it can be expressed in a more output-sensitive way. The third item (which has been known since [118]) shows that the complexity of the problem can be drastically reduced under some restrictive assumptions.

Fast matrix multiplication. As mentioned in Section 2.1, the matrix reduction algorithm is essentially Gaussian elimination with known column order. As such, it can be optimized by using fast matrix multiplication techniques. Specifically, the PLU factorization algorithm of Bunch and Hopcroft [43] can be applied with minor modification to reduce the complexity of the persistence algorithm from $O(m^3)$ down to $O(m^\omega)$, where $\omega \in [2, 2.373)$ is the best known exponent for multiplying two $m \times m$ matrices. The details of the adaptation can be found in [196]. The approach has been recently extended by Milosavljevic, Morozov, and Skraba [197] to handle zigzags, a trickier setting where the order of the columns is not known in advance. Let us point out that these contributions are mostly theoretical, considering the technicality of the approach and the size of the constant factors involved in the complexity bounds.

Rank computation. Chen and Kerber [75] proposed another approach that leads to an output-sensitive algorithm to compute persistence. Their approach consists in using the formulas of (1.18) to localize and count the multiplicity of the points of $\text{dgm}(\mathcal{K}^\sigma)$. This requires performing rank computations on submatrices of the boundary matrix. Given a user-defined threshold $\delta > 0$, the algorithm returns only those topological features whose lifespan in the filtration is at least δ . The total running time is $O(n_{\delta(1-\alpha)}R(m) \log m)$, where $n_{\delta(1-\alpha)}$ is the number of topological features of lifespan at least $\delta(1-\alpha)$ for arbitrarily small $\alpha > 0$, m is the total number of simplices in the filtration, and $R(m)$ is the time required to compute the rank of a binary $m \times m$ matrix. Depending on the choice of the rank computation algorithm, one gets either a deterministic $O(n_{\delta(1-\alpha)}m^\omega \log m)$, or a Las-Vegas $O(n_{\delta(1-\alpha)}m^{2.28})$, or a Monte-Carlo $O(n_{\delta(1-\alpha)}m^{2+\epsilon})$ time algorithm. The algorithm performs no better asymptotically than the standard matrix reduction when $n_{\delta(1-\alpha)} = \Omega(m)$, however it does perform better when this quantity is reduced. The price to pay is that only a subset of the topological features are detected, nevertheless these are the most persistent and therefore also the most relevant ones in applications, as we will see in the second part of the book.

Union-find. When one is only interested in 0-dimensional homology, computing persistence boils down to tracking the connected components as they appear and then are merged in the filtration \mathcal{K}^σ . Only the 1-skeleton graph of the simplicial complex K is involved in the process, with each vertex insertion creating a new connected component (so each vertex is positive), while an edge insertion either

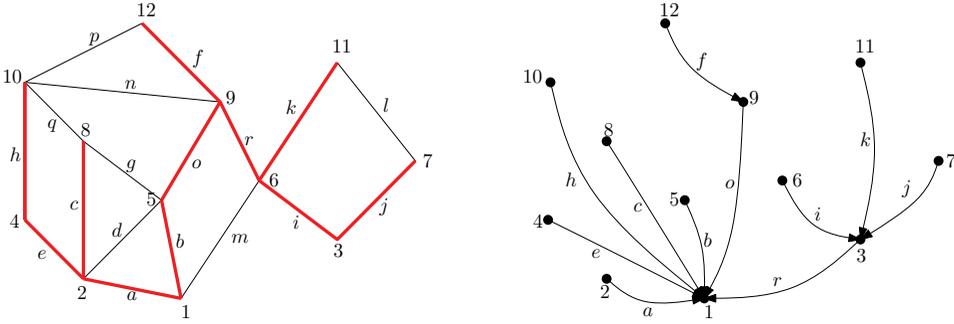


FIGURE 2.6. Left: a planar graph filtered by the order $1, \dots, 12, a, \dots, r$. The negative edges (in bold red) form a spanning tree of the graph. Right: every negative edge connects two previously independent connected components, thus forming a hierarchy on the vertices. The persistence pairs are given by the arrows and their tail in the hierarchy.

merges two components if the edge is negative, or creates a cycle if the edge is positive—the latter case being ignored. Tracking the component creations and merges can be done in $O(m\alpha(m))$ time, where m is the size of the 1-skeleton graph of K and α denotes the inverse Ackermann function, by means of an efficient union-find data structure such as a disjoint-sets forest [92]. The total running time of the persistence algorithm is then $O(m \log m)$ because vertices and edges need first to be sorted according to their time of appearance in \mathcal{K}^σ . Thus, computing persistence is done in near-linear time in this special case. Let us also point out that the output persistence pairs have several additional features, such as the fact that they form a spanning tree of the input graph (as in Kruskal’s minimum spanning tree algorithm), or that they define a hierarchy on the vertices of K , as illustrated in Figure 2.6. These extra features are instrumental in a number of applications, including the one presented in Chapter 6.

2.2.4. *Preprocessing for faster computations.* In addition to the previous approaches, some techniques have been developed to preprocess the input in order to speed up the persistence computation. The bottom line is to reduce the size of the input filtration as much as possible before computing its persistence, in order to lower the impact of the potentially cubic complexity bound on the total running time.

Morse theory for filtrations. Given a fixed cell complex K , discrete Morse theory [130] introduces the concept of a discrete vector field as a partial matching between the cells of K such that the dimensions of the simplices in a matched pair differ by 1. A discrete flow line is then the concatenation of a sequence of matched pairs such that two consecutive pairs (σ_i, τ_i) and $(\sigma_{i+1}, \tau_{i+1})$ satisfy the property that σ_{i+1} is a face of τ_i of codimension 1. The theory highlights those vector fields whose flow lines have no nontrivial cycles, which it calls *gradient vector fields*. These are the analogue, in the discrete setting, of the continuous gradient vector fields of Morse functions defined over smooth manifolds. By analogy, the unpaired cells are called *critical*, and together with the incidence relations induced by flow lines starting and ending at them, they form a cell complex called the *Morse complex*

of the gradient vector field. A central result in the theory is that this complex has the same homology as K . Mischaikow and Nanda [198] have extended these concepts and results to the filtered setting, introducing the notions of a *filtered gradient field* and its associated *filtered Morse complex*, which they prove to have the same persistent homology as the input filtration. Thus, the filtered Morse complex can be used instead of the input filtration in the persistence calculation. This is interesting because the new input size cannot be larger than the previous one, and in some cases it can be much smaller. The challenge is to find filtered gradient fields that minimize the number of critical cells, a problem known to be NP-hard in general [167] but for which heuristics exist. The one proposed by Mischaikow and Nanda [198] is based on the coreduction homology algorithm of Mrozek and Batko [200], which they report to give good speed-ups on practical data. The approach can also be iterated, each time by computing a new filtered gradient field over the previously obtained filtered Morse complex, in order to further reduce the size of the filtration, at the price of an increased preprocessing time. An implementation is available as part of the C++ library PERSEUS (<http://www.sas.upenn.edu/~vnanda/perseus/index.html>).

Complexes from graphs. Specific optimizations can be made for filtrations that are composed of *clique complexes*. A clique complex is itself composed of the cliques of its 1-skeleton graph, and as such it is fully determined by this graph. Rips complexes (defined in Chapter 5) are examples of such complexes. The challenge is to be able to compute their homology directly from the graph, without performing the full clique expansion, and with a minimal memory overhead. A seemingly more accessible goal is to be able to simplify the combinatorial structure of the complexes while preserving their homology, as cliques are known to have trivial reduced homology regardless of their size. Several methods have been proposed to reduce the size of a clique complex K while preserving its homology [8, 242]. The bottom line is to perform a greedy sequence of simplex contractions and collapses in K , under some local sufficient conditions ensuring that the homology is preserved by these operations. The simplification process takes place at the 1-skeleton graph level, with some controlled extra bookkeeping, instead of at the full clique expansion level. It ends when a minimal simplicial complex (or simplicial set) has been reached, so no more collapses or contractions satisfying the sufficient local conditions can be applied. The size of the output is not guaranteed to be minimal in any way, as finding an optimal sequence of collapses or contractions is a hard problem. However, practical experiments show huge improvements over computing the full expansion. Unfortunately, as of now these methods only work with single complexes. Extending them to filtrations is a challenging open problem.

Stability

As we saw in the previous chapters, an important contribution of persistence theory has been to introduce persistence diagrams and to promote them as signatures for persistence and zigzag modules, and by extension, for topological filtrations and functions. Of particular interest is the fact that these signatures do not require the modules under consideration to be decomposable in the sense of quiver theory, which gives them more flexibility and a larger potential for applications.

Two fundamental questions remain: are these signatures stable? are they informative? In mathematical terms, we want to know whether (dis-)similar persistence modules have (dis-)similar persistence diagrams, for suitable choices of (dis-)similarity measures. This is where the Isometry Theorem comes into play and acts as the cornerstone of the theory:

THEOREM 3.1 (Isometry). *Let \mathbb{V}, \mathbb{W} be \mathfrak{q} -tame persistence modules over \mathbb{R} . Then,*

$$d_b(\mathbf{dgm}(\mathbb{V}), \mathbf{dgm}(\mathbb{W})) = d_i(\mathbb{V}, \mathbb{W}).$$

The statement of the theorem falls into two parts: on the one hand, the ‘stability’ part ($d_b(\mathbf{dgm}(\mathbb{V}), \mathbf{dgm}(\mathbb{W})) \leq d_i(\mathbb{V}, \mathbb{W})$), which guarantees that persistence diagrams are stable signatures for persistence modules; on the other hand, the ‘converse stability’ part ($d_i(\mathbb{V}, \mathbb{W}) \leq d_b(\mathbf{dgm}(\mathbb{V}), \mathbf{dgm}(\mathbb{W}))$), which guarantees that they are also informative signatures. The choice of distances is an essential aspect of the theorem. Roughly speaking, the *interleaving distance* d_i between persistence modules measures how far they are from being isomorphic, while the *bottleneck distance* d_b between locally finite multisets of points in the plane measures the cost of the best transport plan between them¹—see Section 1 for the formal definitions. In addition to being natural, these metrics give the tightest bounds in the theorem.

The stability part of the Isometry Theorem is reputed to be difficult to prove, whereas the converse stability part is easy once the stability part is given. We will review their proofs in Sections 2 and 3 respectively, following the work of Chazal et al. [72]. Several other proofs exist in the literature, both for the stability part [20, 87] and for the converse stability part [41, 179]. These proofs use somewhat different languages and various sets of hypotheses. We will give a brief historical account and put them into perspective in Section 4, where we will also discuss other choices of metrics.

Prerequisites. This chapter requires no extra background compared to the previous chapters.

¹This metric is oblivious to the decorations in the persistence diagrams, so throughout the chapter and unless otherwise stated the diagrams under consideration will be undecorated.

1. Metrics

We introduce the bottleneck distance first because the intuition behind it is easier to grasp. We then proceed with the interleaving distance.

1.1. Bottleneck distance. We treat undecorated persistence diagrams as plain multisets of points in the extended plane $\bar{\mathbb{R}}^2$. Given two such multisets P, Q , a *partial matching* between P, Q is understood as in graph theory, that is, it is a subset M of $P \times Q$ that satisfies the following constraints:

- every point $p \in P$ is matched with at most one point of Q , i.e. there is at most one point $q \in Q$ such that $(p, q) \in M$,
- every point $q \in Q$ is matched with at most one point of P , i.e. there is at most one point $p \in P$ such that $(p, q) \in M$.

We will use the notation $M : P \leftrightarrow Q$ to indicate that M is a partial matching between P and Q . By convention, the *cost* of a matched pair $(p, q) \in M$ is $\|p - q\|_\infty$, the ℓ^∞ -distance between p and q , while the cost of an unmatched point $s \in P \sqcup Q$ is $\frac{|s_y - s_x|}{2}$, the ℓ^∞ -distance between s and its closest point on the diagonal Δ . Note that these quantities can be infinite when the points considered lie at infinity. The chosen cost function for partial matchings $M : P \leftrightarrow Q$ is the *bottleneck cost* $c(M)$:

$$(3.1) \quad c(M) = \max \left\{ \sup_{(p,q) \in M} \|p - q\|_\infty, \sup_{s \in P \sqcup Q \text{ unmatched}} \frac{|s_y - s_x|}{2} \right\}.$$

The *bottleneck distance* between the two multisets P, Q is the smallest possible bottleneck cost achieved by partial matchings between them:

$$(3.2) \quad d_b(P, Q) = \inf_{M: P \leftrightarrow Q} c(M).$$

Let us give an illustrative example to help the reader grasp the intuition behind the choice of costs in this definition.

EXAMPLE 3.2. Take the function $f : \mathbb{R} \rightarrow \mathbb{R}$ and its noisy approximation f' depicted in Figure 3.1(a) (reproduced from Figure 0.5 in the general introduction), and consider the persistent homology of their sublevel-sets filtrations, described in the diagrams of Figure 3.1(b).

The three distinguished minima p', q', s' of f' can be viewed as the counterparts of the minima p, q, s of f after addition of the noise. Their corresponding points $\mathbf{p}', \mathbf{q}', \mathbf{s}'$ in $\mathbf{dgm}(f')$ are therefore naturally matched with the points $\mathbf{p}, \mathbf{q}, \mathbf{s}$ corresponding to p, q, s in $\mathbf{dgm}(f)$. The actual matching between $\{\mathbf{p}, \mathbf{s}\}$ and $\{\mathbf{p}', \mathbf{s}'\}$ does not really matter since the two possible choices give the same cost, points $\mathbf{p}, \mathbf{s} \in \mathbf{dgm}(f)$ having the same location. What matters though is that $\mathbf{dgm}(f')$ has two points in the vicinity of \mathbf{p}, \mathbf{s} , a constraint that is well-captured by the bottleneck distance². The other minima of f' are inconsequential and therefore treated as topological noise, which the bottleneck distance captures by leaving their corresponding points in $\mathbf{dgm}(f')$ unmatched, or equivalently, by matching them with the diagonal Δ .

²As opposed to other distances between multisets of points, such as the Hausdorff distance.

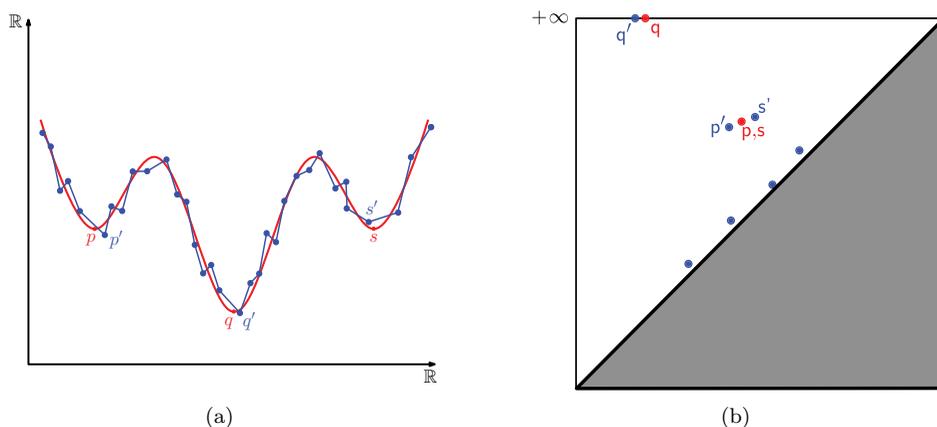


FIGURE 3.1. (a) A smooth function $f : \mathbb{R} \rightarrow \mathbb{R}$ (red) and a piecewise linear approximation f' (blue). (b) Superimposition of the persistence diagrams of f (red) and f' (blue).

Let us now focus on the persistent homologies of f and f' , which decompose as follows:

$$\begin{aligned} H_*(f) &\cong \mathbb{I}[p_x, p_y] \oplus \mathbb{I}[q_x, q_y] \oplus \mathbb{I}[s_x, s_y], \\ H_*(f') &\cong \mathbb{I}[p'_x, p'_y] \oplus \mathbb{I}[q'_x, q'_y] \oplus \mathbb{I}[s'_x, s'_y] \oplus \bigoplus_{t' \in \text{dgm}(f') \setminus \{p', q', s'\}} \mathbb{I}[t'_x, t'_y]. \end{aligned}$$

The cost of a matched pair—say (q, q') —measures the difference between their corresponding intervals— $[q_x, q_y]$ and $[q'_x, q'_y]$ —in the decompositions of $H_*(f)$ and $H_*(f')$. The cost of an unmatched point—say $t' \in \text{dgm}(f') \setminus \{p', q', s'\}$ —measures the smallest possible difference between its associated interval $[t'_x, t'_y]$ and a length-zero interval. The actual choice of norm in \mathbb{R}^2 does not matter fundamentally since all norms are equivalent, however the tightest bounds in the isometry theorem are achieved using the ℓ^∞ -norm.

The bottleneck distance is an extended pseudometric on the multisets of points in the extended plane. Indeed, it is symmetric, it takes values in $[0, +\infty]$, with $d_b(P, Q) = 0$ whenever $P = Q$, moreover it satisfies the triangle inequality

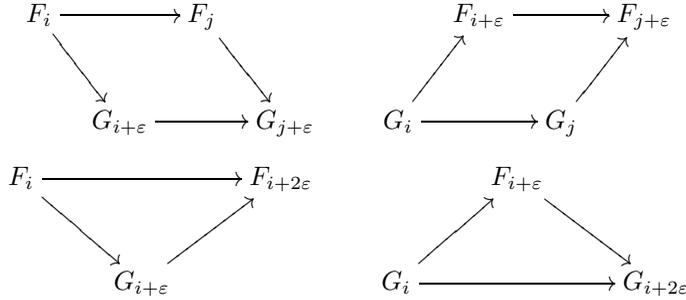
$$d_b(P, S) \leq d_b(P, Q) + d_b(Q, S)$$

as matchings $P \leftrightarrow Q$ and $Q \leftrightarrow S$ can be composed in a natural way to get a matching $P \leftrightarrow S$. It is not a true distance though, even on the multisets not touching the diagonal Δ , as for instance the multisets $\mathbb{Q}^2 \setminus \Delta$ and $(\mathbb{Q} + \sqrt{2})^2 \setminus \Delta$ are distinct but at bottleneck distance zero—the infimum in (3.2) is zero but not attained in this case. In [72, theorem 4.10] a compactness argument is used to show that the infimum in (3.2) is always attained when P, Q are locally finite multisets in $\bar{\mathbb{R}}^2 \setminus \Delta$. This implies that d_b is a true distance for such multisets, which by Theorem 1.13 include in particular all undecorated persistence diagrams of \mathfrak{q} -tame modules.

1.2. Interleaving distance. Let us begin with a special case that carries the intuition behind the concept of interleaving. Consider two maps $f, g : X \rightarrow \mathbb{R}$ such that $\|f - g\|_\infty \leq \varepsilon$. Their sublevel-sets filtrations \mathcal{F} and \mathcal{G} are interleaved as follows:

$$(3.3) \quad \forall i \in \mathbb{R}, F_i \subseteq G_{i+\varepsilon} \text{ and } G_i \subseteq F_{i+\varepsilon}.$$

More precisely, for all indices $i \leq j \in \mathbb{R}$ we have the following commutative diagram of sublevel sets and inclusion maps:



After applying the homology functor H_* , we get a commutative diagram of vector spaces and linear maps, which involves the persistence modules $H_*(\mathcal{F})$ and $H_*(\mathcal{G})$ together with the cross maps $\phi_i : H_*(F_i) \rightarrow H_*(G_{i+\varepsilon})$ and $\psi_i : H_*(G_i) \rightarrow H_*(F_{i+\varepsilon})$ induced at the homology level by $F_i \hookrightarrow G_{i+\varepsilon}$ and $G_i \hookrightarrow F_{i+\varepsilon}$. For simplicity of notations we have renamed $H_*(\mathcal{F}) = \mathbb{V}$ and $H_*(\mathcal{G}) = \mathbb{W}$:

$$(3.4) \quad \begin{array}{ccc} V_i & \longrightarrow & V_j \\ \phi_i \searrow & & \searrow \phi_j \\ W_{i+\varepsilon} & \longrightarrow & W_{j+\varepsilon} \end{array} \quad \begin{array}{ccc} V_{i+\varepsilon} & \longrightarrow & V_{j+\varepsilon} \\ \psi_i \nearrow & & \nearrow \psi_j \\ W_i & \longrightarrow & W_j \end{array}$$

$$(3.5) \quad \begin{array}{ccc} V_i & \longrightarrow & V_{i+2\varepsilon} \\ \phi_i \searrow & & \searrow \psi_{i+\varepsilon} \\ W_{i+\varepsilon} & & \end{array} \quad \begin{array}{ccc} & & V_{i+\varepsilon} \\ \psi_i \nearrow & & \searrow \phi_{i+\varepsilon} \\ W_i & \longrightarrow & W_{i+2\varepsilon} \end{array}$$

This is what we call an ε -interleaving between persistence modules. As one can see, it is the direct translation, at the algebraic level, of the interleaving between filtrations, although it does not actually need filtrations to start with in order to be stated.

DEFINITION 3.3. Let \mathbb{V}, \mathbb{W} be two persistence modules over \mathbb{R} , and let $\varepsilon \geq 0$. An ε -interleaving between \mathbb{V}, \mathbb{W} is given by two families of linear maps $(\phi_i : V_i \rightarrow W_{i+\varepsilon})_{i \in \mathbb{R}}$ and $(\psi_i : W_i \rightarrow V_{i+\varepsilon})_{i \in \mathbb{R}}$ such that the diagrams of (3.4) and (3.5) commute for all $i \leq j \in \mathbb{R}$. The *interleaving distance* between \mathbb{V} and \mathbb{W} is

$$(3.6) \quad d_i(\mathbb{V}, \mathbb{W}) = \inf \{ \varepsilon \geq 0 \mid \text{there is an } \varepsilon\text{-interleaving between } \mathbb{V} \text{ and } \mathbb{W} \}.$$

Note that there are no conditions on the persistence modules \mathbb{V}, \mathbb{W} , which can be arbitrary as long as they are defined over the same ground field \mathbf{k} . When there is no ε -interleaving between \mathbb{V} and \mathbb{W} for any $\varepsilon \geq 0$, we let $d_i(\mathbb{V}, \mathbb{W}) = +\infty$.

Let us now rephrase the concept of interleaving in categorical terms, which will give a nice and compact definition. A family $\phi = (\phi_i : V_i \rightarrow W_i)_{i \in \mathbb{R}}$ of linear maps

such that the leftmost diagram of (3.4) commutes for all $i \leq j \in \mathbb{R}$ is called a *morphism of degree ε* from \mathbb{V} to \mathbb{W} . This notion extends the concept of morphism between persistence modules by adding a shift by ε in the indices. In particular, a morphism in the classical sense is a morphism of degree 0 in the present context. We write

$$\begin{aligned}\mathrm{Hom}^\varepsilon(\mathbb{V}, \mathbb{W}) &= \{\text{morphisms } \mathbb{V} \rightarrow \mathbb{W} \text{ of degree } \varepsilon\}, \\ \mathrm{End}^\varepsilon(\mathbb{V}) &= \{\text{morphisms } \mathbb{V} \rightarrow \mathbb{V} \text{ of degree } \varepsilon\}.\end{aligned}$$

Composition gives a map

$$\mathrm{Hom}^{\varepsilon'}(\mathbb{V}, \mathbb{W}) \times \mathrm{Hom}^\varepsilon(\mathbb{U}, \mathbb{V}) \rightarrow \mathrm{Hom}^{\varepsilon+\varepsilon'}(\mathbb{U}, \mathbb{W}).$$

When $\mathbb{U} = \mathbb{W}$, the shift by $\varepsilon + \varepsilon'$ prevents the composition from being equal to the identity $\mathbb{1}_{\mathbb{U}}$. However, there is a natural counterpart to $\mathbb{1}_{\mathbb{U}}$ here: the so-called $(\varepsilon + \varepsilon')$ -*shift* morphism, noted $\mathbb{1}_{\mathbb{U}}^{\varepsilon+\varepsilon'} \in \mathrm{End}^{\varepsilon+\varepsilon'}(\mathbb{U})$, which is nothing but the collection of maps $(u_i^{\varepsilon+\varepsilon'})_{i \in \mathbb{R}}$ taken from the persistence module structure on \mathbb{U} . Intuitively, this morphism ‘travels’ along \mathbb{U} upwards by $\varepsilon + \varepsilon'$. Then, Definition 3.3 can be rephrased as follows:

DEFINITION 3.3 (rephrased). Two persistence modules \mathbb{V}, \mathbb{W} over \mathbb{R} are ε -*interleaved* if there exist $\phi \in \mathrm{Hom}^\varepsilon(\mathbb{V}, \mathbb{W})$ and $\psi \in \mathrm{Hom}^\varepsilon(\mathbb{W}, \mathbb{V})$ such that $\psi \circ \phi = \mathbb{1}_{\mathbb{V}}^{2\varepsilon}$ and $\phi \circ \psi = \mathbb{1}_{\mathbb{W}}^{2\varepsilon}$. The *interleaving distance* between \mathbb{V} and \mathbb{W} is

$$(3.6) \quad d_i(\mathbb{V}, \mathbb{W}) = \inf \{\varepsilon \geq 0 \mid \mathbb{V} \text{ and } \mathbb{W} \text{ are } \varepsilon\text{-interleaved}\}.$$

Saying that $\phi \in \mathrm{Hom}^\varepsilon(\mathbb{V}, \mathbb{W})$ and $\psi \in \mathrm{Hom}^\varepsilon(\mathbb{W}, \mathbb{V})$ is equivalent to saying that the diagrams of (3.4) commute for all $i \leq j \in \mathbb{R}$. Meanwhile, having $\psi \circ \phi = \mathbb{1}_{\mathbb{V}}^{2\varepsilon}$ and $\phi \circ \psi = \mathbb{1}_{\mathbb{W}}^{2\varepsilon}$ is equivalent to having the diagrams of (3.5) commute for all $i \in \mathbb{R}$. Thus, the two versions of Definition 3.3 are equivalent to each other.

The interleaving distance is an extended pseudometric between the isomorphism classes of persistence modules over \mathbb{R} . Indeed, d_i is symmetric, it takes values in $[0, +\infty]$, with $d_i(\mathbb{V}, \mathbb{W}) = 0$ whenever $\mathbb{V} \cong \mathbb{W}$, and moreover $d_i(\mathbb{U}, \mathbb{W}) \leq d_i(\mathbb{U}, \mathbb{V}) + d_i(\mathbb{V}, \mathbb{W})$ as ε -interleavings between \mathbb{U}, \mathbb{V} can be composed with ε' -interleavings between \mathbb{V}, \mathbb{W} to obtain $(\varepsilon + \varepsilon')$ -interleavings between \mathbb{U}, \mathbb{W} . However, d_i is not a metric since $d_i(\mathbb{V}, \mathbb{W}) = 0$ does not imply $\mathbb{V} \cong \mathbb{W}$. For instance $\mathbb{V} = \mathbb{I}[0, 1]$ and $\mathbb{W} = \mathbb{I}(0, 1)$ are at interleaving distance zero yet nonisomorphic. In fact, they are ε -interleaved for any $\varepsilon > 0$ but not 0-interleaved, so the infimum in (3.6) is zero but not attained. If they were 0-interleaved, then by definition they would be isomorphic.

Beside allowing the isometry theorem to work, the interleaving distance enjoys the following fundamental properties:

- It is stable in the sense that for any topological space X and functions $f, g : X \rightarrow \mathbb{R}$ we have $d_i(\mathbf{H}_*(\mathcal{F}), \mathbf{H}_*(\mathcal{G})) \leq \|f - g\|_\infty$, where \mathcal{F}, \mathcal{G} denote the sublevel-sets filtrations of f, g .
- It is universal in the sense that any other stable pseudometric d between persistence modules satisfies $d \leq d_i$.

These properties strongly suggest using persistence modules as algebraic signatures for topological spaces and their functions, moreover they make the interleaving distance a natural choice for comparing these signatures. Lesnick [179] proved these properties under some restrictions on either the persistence modules (which

should be interval-decomposable) or on their ground field (which should be prime), and he conjectured that they should in fact hold without restrictions.

Lesnick's proof is interesting in its own right, as it shows that the interleaving distance can be 'lifted' from the algebraic level back to the topological level, thus reversing the effect of the homology functor. More precisely, if two persistence modules \mathbb{V}, \mathbb{W} are ε -interleaved, then for any homology dimension $p > 0$ there are a CW-complex X and functions $f, g : X \rightarrow \mathbb{R}$ such that $\mathbb{V} \cong H_p(\mathcal{F})$, $\mathbb{W} \cong H_p(\mathcal{G})$, and $\|f - g\|_\infty = \varepsilon$. Once again, there currently are some restrictions on either the modules or their ground field for this lifting to work, but the bottomline is that one can move from the topological level to the algebraic level and back.

2. Proof of the stability part of the Isometry Theorem

We will give two proofs for this part. The first proof yields a loose upper bound on the bottleneck distance, namely $d_b(\mathbf{dgm}(\mathbb{V}), \mathbf{dgm}(\mathbb{W})) \leq 3d_i(\mathbb{V}, \mathbb{W})$, but it is very simple and geometrically flavored, while the bound obtained, although not tight in full generality, is nonetheless optimal in some sense. Follows then the proof of the tight upper bound $d_b(\mathbf{dgm}(\mathbb{V}), \mathbf{dgm}(\mathbb{W})) \leq d_i(\mathbb{V}, \mathbb{W})$, which requires more work with the algebra and in particular the introduction of a novel ingredient.

2.1. A loose bound. The entire proof revolves around some kind of snapping principle, which we will introduce first. We refer the reader to Figure 3.2 for an illustration.

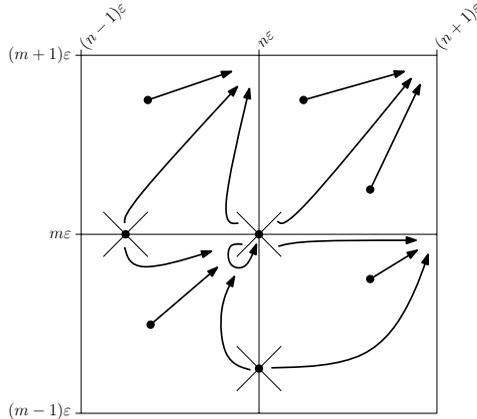


FIGURE 3.2. The snapping rule: each point moves to the upper-right corner of the grid cell containing its decoration.

Snapping principle. Given a persistence module $\mathbb{V} = (V_i, v_i^j)$ over \mathbb{R} and a parameter $\varepsilon > 0$, the ε -discretization of \mathbb{V} , denoted $\mathbb{V}_{\varepsilon\mathbb{Z}}$, is the restriction of \mathbb{V} to the index set $\varepsilon\mathbb{Z}$. Its spaces are the $V_{n\varepsilon}$ for $n \in \mathbb{Z}$, and its maps are the $v_{n\varepsilon}^{m\varepsilon}$ for $n \leq m \in \mathbb{Z}$. Sometimes it is convenient to extend $\mathbb{V}_{\varepsilon\mathbb{Z}}$ to a persistence module over \mathbb{R} by putting the space $V_{n\varepsilon}$ at every index $i \in [n\varepsilon, (n+1)\varepsilon)$. This extension is performed implicitly in the following description.

The effect of the discretization on the interval decomposition of \mathbb{V} (when it exists) is intuitively clear. The birth of a feature occurring strictly between times $(n-1)\varepsilon$ and $n\varepsilon$ is detected only at time $n\varepsilon$ in $\mathbb{V}_{\varepsilon\mathbb{Z}}$. Similarly, the death of a

feature occurring strictly between times $(m - 1)\varepsilon$ and $m\varepsilon$ is detected only at time $m\varepsilon$ in $\mathbb{V}_{\varepsilon\mathbb{Z}}$. Thus, an interval summand $\mathbb{I}[b^\pm, d^\pm]$ with $(n - 1)\varepsilon < b < n\varepsilon$ and $(m - 1)\varepsilon < d < m\varepsilon$ turns into $\mathbb{I}[n\varepsilon, m\varepsilon]$. In the persistence diagram representation, the corresponding point (b, d) gets ‘snapped’ to the upper-right corner of the cell of the grid $\varepsilon\mathbb{Z} \times \varepsilon\mathbb{Z}$ it belongs to. In the limit cases where $b = n\varepsilon$ or $d = m\varepsilon$ or both (i.e. point (b, d) belongs to a cell boundary), it is the decoration of (b, d) that tells where the point must be snapped, using the same rule as before, that is: (b, d) gets snapped to the upper-right corner of the grid cell its decorated version (b^\pm, d^\pm) belongs to. This principle is illustrated in Figure 3.2 and formalized for \mathfrak{q} -tame modules in the following lemma from [71]—see also [72, §2.9]:

LEMMA 3.4 (Snapping). *If \mathbb{V} is \mathfrak{q} -tame, then $\mathbf{dgm}(\mathbb{V}_{\varepsilon\mathbb{Z}})$ is well-defined and $\mathbf{dgm}(\mathbb{V}_{\varepsilon\mathbb{Z}}) \setminus \Delta$ is localized at the grid vertices $\{(n\varepsilon, m\varepsilon)\}_{n < m \in \mathbb{Z}}$ where $\mathbb{Z} = \mathbb{Z} \cup \{-\infty, +\infty\}$. The multiplicity of each grid vertex is given by the persistence measure $\mu_{\mathbb{V}}$ of its lower-left cell, according to the following rules³:*

$$(3.7) \quad \begin{aligned} \text{mult}(n\varepsilon, m\varepsilon) &= \mu_{\mathbb{V}}([(n - 1)\varepsilon, n\varepsilon] \times [(m - 1)\varepsilon, m\varepsilon]), \\ \text{mult}(n\varepsilon, +\infty) &= \mu_{\mathbb{V}}([(n - 1)\varepsilon, n\varepsilon] \times \{+\infty\}), \\ \text{mult}(-\infty, n\varepsilon) &= \mu_{\mathbb{V}}(\{-\infty\} \times [(n - 1)\varepsilon, n\varepsilon]), \\ \text{mult}(-\infty, +\infty) &= \mu_{\mathbb{V}}(\{-\infty\} \times \{+\infty\}). \end{aligned}$$

These rules define a partial matching of bottleneck cost at most ε between $\mathbf{dgm}(\mathbb{V})$ and $\mathbf{dgm}(\mathbb{V}_{\varepsilon\mathbb{Z}})$.

The proof of this lemma follows the intuition given above for interval-decomposable modules. The technical details can be found in [71, 72], where it is also proved that discretizations can be performed over arbitrary index sets that do not have limit points in \mathbb{R} . For instance, we can discretize \mathbb{V} over the index set $\varepsilon\mathbb{Z} + \frac{\varepsilon}{2}$ and still benefit from the Snapping Lemma 3.4 (with adapted indices in the formulas), in particular $d_b(\mathbf{dgm}(\mathbb{V}), \mathbf{dgm}(\mathbb{V}_{\varepsilon\mathbb{Z} + \frac{\varepsilon}{2}})) \leq \varepsilon$. This observation will be instrumental in the upcoming proof of the loose bound.

The proof. Let $\mathbb{V} = (V_i, v_i^j)$ and $\mathbb{W} = (W_i, w_i^j)$ be two \mathfrak{q} -tame persistence modules. The goal is to show that $d_b(\mathbf{dgm}(\mathbb{V}), \mathbf{dgm}(\mathbb{W})) \leq 3\varepsilon$ for all $\varepsilon > d_i(\mathbb{V}, \mathbb{W})$. The case $d_i(\mathbb{V}, \mathbb{W}) = +\infty$ is trivial, so let us assume that $d_i(\mathbb{V}, \mathbb{W})$ is finite and, given $\varepsilon > d_i(\mathbb{V}, \mathbb{W})$, let us take two degree- ε morphisms $\phi \in \text{Hom}^\varepsilon(\mathbb{V}, \mathbb{W})$ and $\psi \in \text{Hom}^\varepsilon(\mathbb{W}, \mathbb{V})$ that form an ε -interleaving of \mathbb{V}, \mathbb{W} as per Definition 3.3. We then have (in particular) the following commutative diagram, where parameter n ranges over \mathbb{Z} :

$$(3.8) \quad \begin{array}{ccccccc} \cdots & \longrightarrow & V_{(2n-2)\varepsilon} & \xrightarrow{v_{(2n-2)\varepsilon}^{2n\varepsilon}} & V_{2n\varepsilon} & \xrightarrow{v_{2n\varepsilon}^{(2n+2)\varepsilon}} & V_{(2n+2)\varepsilon} & \longrightarrow & \cdots \\ & & \searrow & \nearrow \psi_{(2n-1)\varepsilon} & & \searrow \phi_{2n\varepsilon} & \nearrow & & \\ & & \phi_{(2n-2)\varepsilon} & & & & & & \\ & & \cdots & \longrightarrow & W_{(2n-1)\varepsilon} & \xrightarrow{w_{(2n-1)\varepsilon}^{(2n+1)\varepsilon}} & W_{(2n+1)\varepsilon} & \longrightarrow & \cdots \\ & & & & \nearrow & \searrow \psi_{(2n+1)\varepsilon} & & & \end{array}$$

³Here we are using an extended version of the persistence measure of Chapter 1, which allows singular rectangles at infinity. For this we are following the convention of [72, §2.6], where $\mu_{\mathbb{V}}([p, q] \times \{+\infty\})$ is defined as the limit $\lim_{r \rightarrow +\infty} \mu_{\mathbb{V}}([p, q] \times [r, +\infty]) = \min_r \mu_{\mathbb{V}}([p, q] \times [r, +\infty])$, so every point (b, d) with $p < b < q$ and $d = +\infty$ is counted as being inside the rectangle $[p, q] \times \{+\infty\}$ even though its decoration is $(b^\pm, +\infty^-)$. The persistence measures of rectangles $\{-\infty\} \times [r, s]$ and $\{-\infty\} \times \{+\infty\}$ are defined similarly.

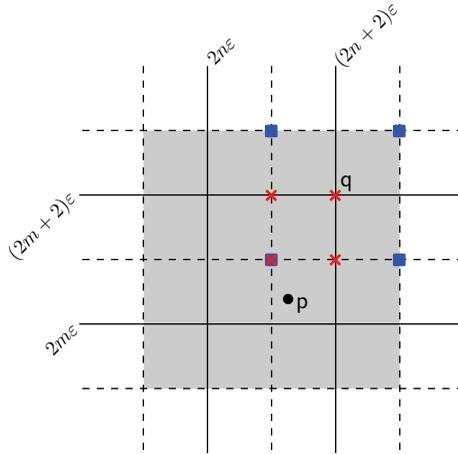


FIGURE 3.3. Tracking down point $p \in \text{dgm}(\mathbb{V})$ through its successive matchings.

— Based on Chazal et al. [71].

This diagram involves three distinguished persistence modules, namely:

- the discretization of \mathbb{V} over $2\varepsilon\mathbb{Z}$, obtained by following the top row:

$$\mathbb{V}_{2\varepsilon\mathbb{Z}} = \cdots \longrightarrow V_{(2n-2)\varepsilon} \longrightarrow V_{2n\varepsilon} \longrightarrow V_{(2n+2)\varepsilon} \longrightarrow \cdots$$

- the discretization of \mathbb{W} over $2\varepsilon\mathbb{Z} + \varepsilon$, obtained by following the bottom row:

$$\mathbb{W}_{2\varepsilon\mathbb{Z} + \varepsilon} = \cdots \longrightarrow W_{(2n-1)\varepsilon} \longrightarrow W_{(2n+1)\varepsilon} \longrightarrow \cdots$$

- the following mixed module, obtained by following the diagonal maps:

$$\mathbb{U} = \cdots \longrightarrow V_{(2n-2)\varepsilon} \longrightarrow W_{(2n-1)\varepsilon} \longrightarrow V_{2n\varepsilon} \longrightarrow W_{(2n+1)\varepsilon} \longrightarrow V_{(2n+2)\varepsilon} \longrightarrow \cdots$$

The key observation is that, by commutativity, the modules $\mathbb{V}_{2\varepsilon\mathbb{Z}}$ and $\mathbb{W}_{2\varepsilon\mathbb{Z} + \varepsilon}$ are also 2ε -discretizations of \mathbb{U} . We can then apply the Snapping Lemma to the four pairs of modules independently: $(\mathbb{V}, \mathbb{V}_{2\varepsilon\mathbb{Z}})$, $(\mathbb{U}, \mathbb{V}_{2\varepsilon\mathbb{Z}})$, $(\mathbb{U}, \mathbb{W}_{2\varepsilon\mathbb{Z} + \varepsilon})$, and finally $(\mathbb{W}, \mathbb{W}_{2\varepsilon\mathbb{Z} + \varepsilon})$. Each time, Lemma 3.4 gives an upper bound of 2ε on the bottleneck distance between the persistence diagrams of the two modules considered, so the triangle inequality gives

$$\begin{aligned} d_b(\text{dgm}(\mathbb{V}), \text{dgm}(\mathbb{W})) &\leq d_b(\text{dgm}(\mathbb{V}), \text{dgm}(\mathbb{V}_{2\varepsilon\mathbb{Z}})) + d_b(\text{dgm}(\mathbb{V}_{2\varepsilon\mathbb{Z}}), \text{dgm}(\mathbb{U})) \\ &\quad + d_b(\text{dgm}(\mathbb{U}), \text{dgm}(\mathbb{W}_{2\varepsilon\mathbb{Z} + \varepsilon})) + d_b(\text{dgm}(\mathbb{W}_{2\varepsilon\mathbb{Z} + \varepsilon}), \text{dgm}(\mathbb{W})) \leq 8\varepsilon. \end{aligned}$$

This crude upper bound on the bottleneck distance can be reduced to 3ε by carefully following the movements of each point of $\text{dgm}(\mathbb{V})$ as it is matched successively with the diagrams of $\mathbb{V}_{2\varepsilon\mathbb{Z}}$, \mathbb{U} , $\mathbb{W}_{2\varepsilon\mathbb{Z} + \varepsilon}$, and \mathbb{W} . The typical situation is depicted in Figure 3.3: point $p \in \text{dgm}(\mathbb{V})$ is first matched with $q \in \text{dgm}(\mathbb{V}_{2\varepsilon\mathbb{Z}})$, then with one of the four red cross-shaped points in $\text{dgm}(\mathbb{U})$, then with one of the blue square-shaped points in $\text{dgm}(\mathbb{W}_{2\varepsilon\mathbb{Z} + \varepsilon})$, and finally with some point within the grey area in $\text{dgm}(\mathbb{W})$. All in all, point p has moved by at most 3ε in the ℓ^∞ -distance. Since this is true for any point of $\text{dgm}(\mathbb{V})$ and any $\varepsilon > 3d_1(\mathbb{V}, \mathbb{W})$, we conclude that $d_b(\text{dgm}(\mathbb{V}), \text{dgm}(\mathbb{W})) \leq 3d_1(\mathbb{V}, \mathbb{W})$.

REMARK. As we know, $3d_i(\mathbb{V}, \mathbb{W})$ is not a tight upper bound. However, it becomes so if one replaces the concept of interleaving as of Definition 3.3 by the weaker version given in (3.8). Indeed, it is easy to build examples of persistence modules \mathbb{V}, \mathbb{W} that are thus weakly ε -interleaved but whose persistence diagrams lie 3ε apart of each other in the bottleneck distance [71].

2.2. The tight bound. The previous proof does not exploit the full power of interleavings, restricting itself to the weaker version (3.8). The key property ignored so far is that the set $\varepsilon\mathbb{Z}$ over which discretizations are taken can be shifted arbitrarily. The novel ingredient that captures this property is commonly referred to as the *Interpolation Lemma*, and it is certainly the most subtle part in the proof of the Isometry Theorem. We will present it first, before giving the proof of the tight bound.

Interpolation lemma. The idea is to interpolate between pairs of interleaved persistence modules while maintaining the interleaving property. In other words, one seeks to find some kind of ‘paths’ of bounded length between persistence modules in the interleaving distance.

LEMMA 3.5 (Interpolation). *Suppose \mathbb{V}, \mathbb{W} are ε -interleaved. Then there is a 1-parameter family $(\mathbb{U}_x)_{x \in [0, \varepsilon]}$ of persistence modules such that $\mathbb{U}_0 = \mathbb{V}$, $\mathbb{U}_\varepsilon = \mathbb{W}$, and $\mathbb{U}_x, \mathbb{U}_y$ are $|y - x|$ -interleaved for all $x, y \in [0, \varepsilon]$. This family is not unique in general.*

PROOF OUTLINE. There is a nice pictorial representation of the interpolating family, as an interleaving between persistence modules can itself be thought of as a representation of a certain poset in the plane. More precisely, consider the standard partial order on the plane:

$$(x, y) \preceq (x', y') \iff x \leq x' \text{ and } y \leq y'.$$

For any real number r , define the shifted diagonal

$$\Delta_r = \{(x, y) \mid y - x = 2r\} \subset \mathbb{R}^2.$$

As a poset, this is isomorphic to the real line, for instance by identifying $t \in \mathbb{R}$ with $(t - r, t + r) \in \Delta_r$. Through this, we get a canonical identification between the persistence modules over \mathbb{R} and the representations of the poset Δ_r . Moreover, it is not hard to see that a $|y - x|$ -interleaving between two persistence modules induces a representation of the poset $\Delta_x \cup \Delta_y$ and vice-versa, which we formalize as follows:

$$(3.9) \quad \begin{aligned} &\mathbb{V}, \mathbb{W} \text{ are } |y - x| \text{-interleaved} \\ &\iff \\ &\exists \mathbb{U} \in \text{Rep}_{\mathbf{k}}(\Delta_x \cup \Delta_y) \text{ such that } \mathbb{U}|_{\Delta_x} = \mathbb{V} \text{ and } \mathbb{U}|_{\Delta_y} = \mathbb{W}. \end{aligned}$$

To be more specific, the persistence module structures on \mathbb{V}, \mathbb{W} are identified respectively with the maps along the shifted diagonals Δ_x, Δ_y , while the morphisms $\phi \in \text{Hom}^\varepsilon(\mathbb{V}, \mathbb{W})$ and $\psi \in \text{Hom}^\varepsilon(\mathbb{W}, \mathbb{V})$ that form the interleaving between \mathbb{V}, \mathbb{W} are identified respectively with the families of vertical and horizontal morphisms between Δ_x, Δ_y , the rest of the maps being obtained by composition. See Figure 3.4 for an illustration.

Then, the proof of Lemma 3.5 amounts to showing that, if there exists a representation \mathbb{U} of the pair of lines $\Delta_0 \cup \Delta_\varepsilon$ such that $\mathbb{U}|_{\Delta_0} = \mathbb{V}$ and $\mathbb{U}|_{\Delta_\varepsilon} = \mathbb{W}$, then

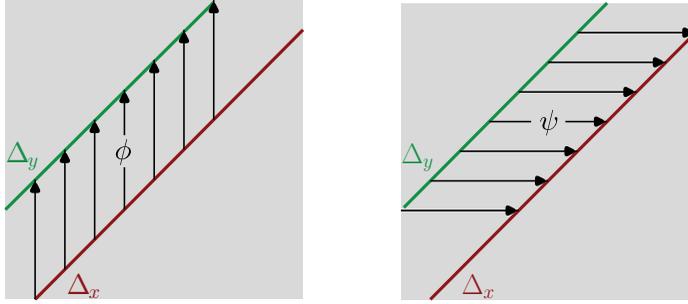


FIGURE 3.4. Identifying ϕ, ψ with the families of vertical and horizontal morphisms in the representation of $\Delta_x \cup \Delta_y$.
 — From Chazal et al. [72].

there exists also a representation $\bar{\mathbb{U}}$ of the diagonal strip

$$\Delta_{[0,\varepsilon]} = \{(x, y) \mid 0 \leq y - x \leq 2\varepsilon\} \subset \mathbb{R}^2$$

such that $\bar{\mathbb{U}}|_{\Delta_0} = \mathbb{V}$ and $\bar{\mathbb{U}}|_{\Delta_\varepsilon} = \mathbb{W}$. Indeed, in that case the 1-parameter family of persistence modules interpolating between \mathbb{V} and \mathbb{W} is given by the restrictions of $\bar{\mathbb{U}}$ to the shifted diagonals Δ_x for $0 \leq x \leq \varepsilon$. Let us give the formal construction of the representation $\bar{\mathbb{U}}$ without proof, referring the reader to [72] for the technical fact checking.

First of all, from \mathbb{V}, \mathbb{W} we construct the following representations of \mathbb{R}^2 :

- \mathbb{A} defined by $A_{(x,y)} = V_x$ and $a_{(x,y)}^{(z,t)} = v_x^z$
- \mathbb{B} defined by $B_{(x,y)} = W_{y-\varepsilon}$ and $b_{(x,y)}^{(z,t)} = w_{y-\varepsilon}^{t-\varepsilon}$
- \mathbb{C} defined by $C_{(x,y)} = V_y$ and $c_{(x,y)}^{(z,t)} = v_y^t$
- \mathbb{D} defined by $D_{(x,y)} = W_{x+\varepsilon}$ and $d_{(x,y)}^{(z,t)} = w_{x+\varepsilon}^{z+\varepsilon}$

Next, we construct the four morphisms:

- $\mathbb{1}_{\mathbb{V}} : \mathbb{A} \rightarrow \mathbb{C}$ defined at (x, y) to be $v_x^y : V_x \rightarrow V_y$
- $\Phi : \mathbb{A} \rightarrow \mathbb{D}$ defined at (x, y) to be $\phi_x : V_x \rightarrow W_{x+\varepsilon}$
- $\Psi : \mathbb{B} \rightarrow \mathbb{C}$ defined at (x, y) to be $\psi_{y-\varepsilon} : W_{y-\varepsilon} \rightarrow V_y$
- $\mathbb{1}_{\mathbb{W}} : \mathbb{B} \rightarrow \mathbb{D}$ defined at (x, y) to be $w_{y-\varepsilon}^{x+\varepsilon} : W_{y-\varepsilon} \rightarrow W_{x+\varepsilon}$

These four morphisms are defined over the region where $0 \leq y - x \leq 2\varepsilon$, i.e. precisely over the diagonal strip $\Delta_{[0,\varepsilon]}$. We now define a morphism $\mathbb{A} \oplus \mathbb{B} \rightarrow \mathbb{C} \oplus \mathbb{D}$ by the 2×2 matrix:

$$\begin{pmatrix} \mathbb{1}_{\mathbb{V}} & \Psi \\ \Phi & \mathbb{1}_{\mathbb{W}} \end{pmatrix}$$

Its image has the desired properties to be our representation $\bar{\mathbb{U}}$. □

REMARK. Our construction of $\bar{\mathbb{U}}$ shows in fact the stronger statement that any representation of $\Delta_0 \cup \Delta_\varepsilon$ can be extended to a representation of $\Delta_{[0,\varepsilon]}$. The

extension is not unique, and in our outline we only give one possibility. Other possibilities include taking the kernel or cokernel of the morphism $\begin{pmatrix} \mathbf{1}_{\mathbb{V}} & -\Psi \\ -\Phi & \mathbf{1}_{\mathbb{W}} \end{pmatrix}$. The effect on the interpolation differs from one choice to the next, as will be discussed at the end of the section.

Proof of the tight bound. Let $\mathbb{V} = (V_i, v_i^j)$ and $\mathbb{W} = (W_i, w_i^j)$ be ε -interleaved \mathfrak{q} -tame persistence modules, and let $(\mathbb{U}_x)_{x \in [0, \varepsilon]}$ be the interpolating 1-parameter family of modules provided by the Interpolation Lemma 3.5. We will assume for simplicity that $\mathbf{dgm}(\mathbb{U}_x)$ has finitely many points for every $x \in [0, \varepsilon]$. Otherwise, an additional compactness argument is needed to finish off the proof, see [72, §4.8]. The proof divides into the following steps:

- (1) For any $0 \leq x \leq y \leq \varepsilon$, it shows that $d_H(\mathbf{dgm}(\mathbb{U}_x), \mathbf{dgm}(\mathbb{U}_y))$, the Hausdorff distance between $\mathbf{dgm}(\mathbb{U}_x)$ and $\mathbf{dgm}(\mathbb{U}_y)$, is at most $|y - x|$.
- (2) It shows that $d_b(\mathbf{dgm}(\mathbb{U}_x), \mathbf{dgm}(\mathbb{U}_y)) = d_H(\mathbf{dgm}(\mathbb{U}_x), \mathbf{dgm}(\mathbb{U}_y))$ when $|y - x|$ is small enough.
- (3) Using the previous steps, it tracks the changes in $\mathbf{dgm}(\mathbb{U}_x)$ as parameter x ranges from 0 to ε , to derive the desired upper bound on $d_b(\mathbf{dgm}(\mathbb{V}), \mathbf{dgm}(\mathbb{W}))$.

Steps 1 and 2 rely on the following box inequalities, known under the name of *Box Lemma* [72, lemma 4.22], where δ denotes the quantity $|y - x|$, where $R = [p, q] \times [r, s] \subset \mathbb{R}^2$ is any rectangle such that $r > q + 2\delta$, and where $R^\delta = [p - \delta, q + \delta] \times [r - \delta, s + \delta]$ is its δ -thickening:

$$(3.10) \quad \mu_{\mathbb{U}_x}(R) \leq \mu_{\mathbb{U}_y}(R^\delta) \text{ and } \mu_{\mathbb{U}_y}(R) \leq \mu_{\mathbb{U}_x}(R^\delta).$$

These inequalities relate the persistence measures of $\mathbb{U}_x, \mathbb{U}_y$ locally, and they are proven using the machinery introduced in Section 2.1: first the persistence modules $\mathbb{U}_x, \mathbb{U}_y$ are discretized over the finite index set $\{p - \delta, p, q, q + \delta, r - \delta, r, s, s + \delta\}$, second the inequalities are derived from the same point tracking strategy using the Snapping Lemma 3.4.

Step 1 follows then from (3.10) by letting R converge to a single point of $\mathbf{dgm}(\mathbb{U}_x)$, which implies that $\mu_{\mathbb{U}_y}$ must be positive within a δ -thickening of that point. Step 2 follows along the way, once it has been observed that when δ is below some threshold δ_x (typically a fraction of the minimum inter-point distance in $\mathbf{dgm}(\mathbb{U}_x)$), we eventually get $\mu_{\mathbb{V}}(R) = \mu_{\mathbb{W}}(R^\delta) = \mu_{\mathbb{V}}(R^{2\delta})$ as R converges to a single point in $\mathbf{dgm}(\mathbb{U}_x)$.

For step 3, consider the family of open intervals $(x - \delta_x, x + \delta_x)$. This family forms an open cover of the compact interval $[0, \varepsilon]$, from which a finite subcover can be extracted, whose elements are centered at the values $0 = x_0 < x_1 < \dots < x_k = \varepsilon$. For any consecutive values x_i, x_{i+1} , we have from steps 1-2 that

$$d_b(\mathbf{dgm}(\mathbb{U}_{x_i}), \mathbf{dgm}(\mathbb{U}_{x_{i+1}})) = d_H(\mathbf{dgm}(\mathbb{U}_{x_i}), \mathbf{dgm}(\mathbb{U}_{x_{i+1}})) \leq |x_{i+1} - x_i|,$$

which gives $d_b(\mathbf{dgm}(\mathbb{V}), \mathbf{dgm}(\mathbb{W})) \leq \varepsilon$ as desired, by the triangle inequality.

REMARK. To illustrate the proof, Figure 3.5 shows how the interpolation between persistence modules (here $\mathbb{I}[0, 4]$ and $\mathbb{I}[1, 6]$) translates into an interpolation between their persistence diagrams. As said earlier, the modules interpolation is not unique, and the figure shows the effects of various choices of parameters, including the interleaving parameter ε , on the resulting diagrams interpolation—see [72, figures 9-10] for further details. It is interesting to note that the interpolated points in the plane do not always follow shortest paths in the ℓ^∞ -distance.

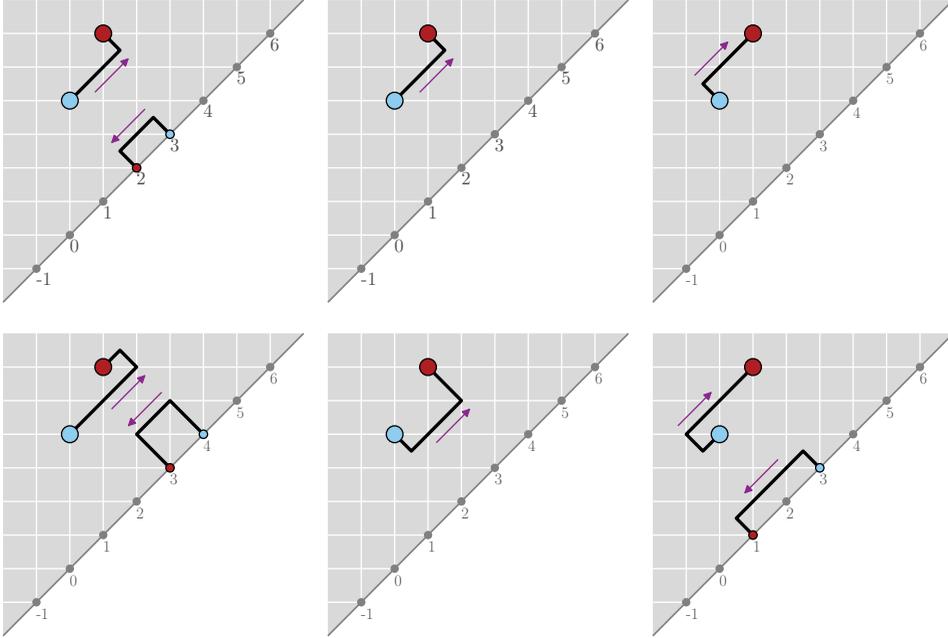


FIGURE 3.5. Effects of various interpolation parameters on the persistence diagrams for the modules $\mathbb{I}[0, 4]$ and $\mathbb{I}[1, 6]$. Top row: $\varepsilon = 2$, bottom row: $\varepsilon = 3$. Left column: using cokernels in the extension part of the proof of Lemma 3.5, center column: using images, right column: using kernels.

— From Chazal et al. [72].

3. Proof of the converse stability part of the Isometry Theorem

Let \mathbb{V}, \mathbb{W} be two \mathfrak{q} -tame persistence modules. The proof that $d_i(\mathbb{V}, \mathbb{W}) \leq d_b(\text{dgm}(\mathbb{V}), \text{dgm}(\mathbb{W}))$ proceeds in two steps:

- (1) It proves the result in the special case where \mathbb{V}, \mathbb{W} are interval-decomposable, which is fairly straightforward.
- (2) It extends the result to arbitrary \mathfrak{q} -tame modules by showing that every such module is the limit, in the interleaving distance, of a sequence of pointwise finite-dimensional (hence interval-decomposable) modules.

Interval-decomposable case. Let $\varepsilon > d_b(\text{dgm}(\mathbb{V}), \text{dgm}(\mathbb{W}))$, and take an arbitrary matching of bottleneck cost at most ε between the two undecorated diagrams. This defines a partial matching between the interval summands of \mathbb{V} and the ones of \mathbb{W} , so the two persistence modules can be decomposed as follows:

$$\mathbb{V} \cong \bigoplus_{j \in J} \mathbb{V}_j, \quad \mathbb{W} \cong \bigoplus_{j \in J} \mathbb{W}_j,$$

where each pair $(\mathbb{V}_j, \mathbb{W}_j)$ is one of the following:

- a pair of matched interval summands, i.e. $\mathbb{V}_j = \mathbb{W}_j = \mathbb{I}[b_j^\pm, d_j^\pm]$,
- \mathbb{V}_j is an unmatched interval summand and $\mathbb{W}_j = 0$,
- $\mathbb{V}_j = 0$ and \mathbb{W}_j is an unmatched interval summand.

It is then easy to exhibit an ε -interleaving between the two decompositions, by proceeding independently for each pair $(\mathbb{V}_j, \mathbb{W}_j)$ and by taking the direct sums of the obtained interleaving maps.

General case. Consider the space \mathcal{P}_q of q -tame persistence modules, equipped with the pseudometric d_i . Let $\mathcal{P}_{\text{pfd}} \subset \mathcal{P}_q$ be the subspace of pointwise finite-dimensional persistence modules. Take the map

$$f : \mathcal{P}_q \times \mathcal{P}_q \longrightarrow [-\infty, +\infty], \quad (\mathbb{V}, \mathbb{W}) \longmapsto d_i(\mathbb{V}, \mathbb{W}) - d_b(\text{dgm}(\mathbb{V}), \text{dgm}(\mathbb{W})),$$

where by convention we let $f = 0$ when both d_i and d_b are infinite. The stability part of Theorem 3.1 implies that $\mathbb{V} \mapsto \text{dgm}(\mathbb{V})$ is a 1-Lipschitz map, so f is non-negative and continuous. Meanwhile, the converse stability part of Theorem 3.1, specialized to interval-decomposable modules, implies that $f = 0$ over $\mathcal{P}_{\text{pfd}} \times \mathcal{P}_{\text{pfd}}$. If we can prove that \mathcal{P}_{pfd} is dense in \mathcal{P}_q , then, by a continuity argument, we will have $f = 0$ over the whole space $\mathcal{P}_q \times \mathcal{P}_q$, which will conclude the proof of the converse stability part of Theorem 3.1 in the general q -tame case.

Thus, our task is to prove that every q -tame persistence module $\mathbb{V} = (V_i, v_i^j)$ lies in the topological closure of \mathcal{P}_{pfd} . Given $\varepsilon > 0$, let \mathbb{V}^ε be the module whose spaces are $V_i^\varepsilon = \text{im } v_{i-\varepsilon}^{i+\varepsilon}$ for $i \in \mathbb{R}$, and whose maps are induced by the maps in \mathbb{V} . This module is called the ε -smoothing of \mathbb{V} , not to be confused with the $\varepsilon\mathbb{Z}$ -discretization $\mathbb{V}_{\varepsilon\mathbb{Z}}$ introduced in Section 2.1. It is easily seen from the definition that \mathbb{V}^ε is pointwise finite-dimensional, and that $\mathbb{V}, \mathbb{V}^\varepsilon$ are ε -interleaved. Hence, \mathbb{V} is a limit of the sequence $(\mathbb{V}^{1/n})_{n \in \mathbb{N}}$ and therefore lies in the topological closure of \mathcal{P}_{pfd} .

REMARK. It is interesting to note that the continuity argument presented here actually extends the full Isometry Theorem to q -tame modules, not just the converse stability part. Therefore, it is enough to prove the Isometry Theorem for pointwise finite-dimensional modules. This is the approach adopted e.g. in [20].

4. Discussion

Origin of the Isometry Theorem. The original proof of the stability part of Theorem 3.1 was given by Cohen-Steiner, Edelsbrunner, and Harer [87]. It was a seminal contribution as it already contained many of the ideas of the modern proofs, including the Box Lemma and the interpolation argument used in Section 2.2, albeit in more restrictive forms. The result was stated for the persistent homology of real-valued functions and had the following flavor⁴:

COROLLARY 3.6. *Let $f, g : X \rightarrow \mathbb{R}$ be q -tame functions. Then,*

$$d_b(\text{dgm}(f), \text{dgm}(g)) \leq \|f - g\|_\infty.$$

In fact, the version stated in [87] used a more restrictive notion of tameness, briefly mentioned in Chapter 2 (see Footnote 1 therein), and it added the extra conditions that X is finitely triangulable and that f, g are continuous. The reason was that the authors did not have the concept of interleaving between persistence modules and the corresponding algebraic Interpolation Lemma 3.5 at their disposal.

⁴This result is a direct consequence of the Isometry Theorem. Indeed, as we saw in Section 1.2, f, g have $\|f - g\|_\infty$ -interleaved sublevel-sets filtrations in the sense of (3.3), so their persistent homologies are $\|f - g\|_\infty$ -interleaved in the sense of Definition 3.3, and Theorem 3.1 applies.

To overcome this lack, they interpolated between the functions f, g at the topological level, and took the persistent homology of the interpolating function as the interpolating persistence module. The problem was that this module could become wild during the process, which led them to add these extra conditions on X, f, g in order to force the interpolating module to remain tame. While this may look a small technicality in the analysis, it does have an impact on applications, where it is common to deal with spaces or maps that do not comply with the requirements, especially in the presence of noise in the data. Such examples can be found in the following chapters.

This was the state of affairs until Chazal and Oudot [66] introduced interleavings between filtrations as in (3.3), which eventually led to the concept of interleaving between persistence modules as in Definition 3.3. Together with D. Cohen-Steiner and L. Guibas, they reworked the proof of Cohen-Steiner, Edelsbrunner, and Harer [87] within this new context, but were unable to derive a tight upper bound on the bottleneck distance without the interpolating argument, nor to turn this argument into a purely algebraic one. Nevertheless, they were able to prove the loose upper bound of Section 2.1 without it, and this was the first purely algebraic proof of stability for persistence diagrams.

At that point, the algebraic Interpolation Lemma was still the key missing ingredient in the picture. It is M. Glisse who brought it to them and thus made it possible to complete the algebraic proof of the stability part of the Isometry Theorem with a tight upper bound [71]. Then, V. de Silva brought them the measure-theoretic view on the definition of persistence diagrams and the categorical view on the proof of the algebraic Interpolation Lemma [72].

Recently, Bauer and Lesnick [20] gave a more direct proof of the stability part of the Isometry Theorem, building the matching between the persistence diagrams directly from one of the two degree- ε morphisms involved in the ε -interleaving between the persistence modules. This allowed them to avoid a blind interpolation argument, and thus to get a better control over the interpolation—recall Figure 3.5. Meanwhile, Bubenik and Scott [41] then Bubenik, de Silva, and Scott [40] rephrased the concept of interleaving in categorical terms, generalizing it to representations of arbitrary posets, and they derived ‘soft’ stability results that bound the interleaving distance in terms of the distance between the topological objects the poset representations are derived from originally. These are promising new developments on the subject, which down the road may help tackle the still unresolved question of defining and proving stability for persistence diagrams of zigzag modules.

The converse stability part of the Isometry Theorem has a more recent history. To our knowledge, it is Lesnick [179] who gave the first proof, in the context of pointwise finite-dimensional modules. Around the same time, Chazal et al. [72] proposed a similar proof and completed it with the continuity argument of Section 3, which allowed them to extend the result to \mathfrak{q} -tame modules. Once again, Bubenik and Scott [41] rephrased the approach in categorical terms and suggested the name ‘Isometry Theorem’ for the combination of the stability and converse stability results.

Balancing the Isometry Theorem. The careful reader may have noticed that the presentation of the Isometry Theorem given in these pages is somewhat unbalanced. Indeed, as we saw in Section 1, the interleaving distance d_i between \mathfrak{q} -tame persistence modules is only a pseudometric between their isomorphism classes, whereas

the bottleneck distance d_b between their undecorated persistence diagrams is a true metric. This asymmetry is overcome in the so-called *observable category*, introduced originally by Chazal, Crawley-Boevey, and de Silva [63] and mentioned briefly at the end of our Chapter 1. In this quotient category, d_i becomes a true metric between the isomorphism classes of \mathfrak{q} -tame modules, so we have the following clean and balanced theory thanks to the Isometry Theorem:

THEOREM 3.7. *For any \mathfrak{q} -tame modules \mathbb{V} and \mathbb{W} , the following are equivalent:*

- \mathbb{V} and \mathbb{W} are isomorphic in the observable category,
- the interleaving distance between \mathbb{V} and \mathbb{W} is zero,
- the bottleneck distance between $\text{dgm}(\mathbb{V})$ and $\text{dgm}(\mathbb{W})$ is zero,
- $\text{dgm}(\mathbb{V})$ and $\text{dgm}(\mathbb{W})$ are equal.

Wasserstein distances. Partial matchings between multisets in the plane can be viewed as transport plans between measures—see [235] for a formal introduction to optimal transportation theory. In this viewpoint, the bottleneck distance d_b is closely related to the Wasserstein distance W_∞ . Other Wasserstein distances W_p with $p < +\infty$ can be considered as well, and new stability results can be derived from the Isometry Theorem, such as:

THEOREM 3.8 (Adapted from Cohen-Steiner et al. [88]). *Let \mathbb{V}, \mathbb{W} be \mathfrak{q} -tame persistence modules over \mathbb{R} . Then,*

$$W_p(\text{dgm}(\mathbb{V}), \text{dgm}(\mathbb{W})) \leq (\text{Pers}(\mathbb{V}) + \text{Pers}(\mathbb{W}))^{\frac{1}{p}} d_i(\mathbb{V}, \mathbb{W})^{1-\frac{1}{p}},$$

where $\text{Pers}(\mathbb{V}) = \sum_{p \in \text{dgm}(\mathbb{V})} (p_y - p_x)$ is the total persistence of \mathbb{V} .

The problem with this kind of upper bound is that it depends on the total persistences of the modules, which are unstable quantities. This dependence might just be an artefact of the proof, which invokes the Isometry Theorem as a black-box and then manipulates the formulas to make the p -th power appear. Yet, as of now this is the only known way to bound degree- p Wasserstein distances between persistence diagrams, and unfortunately the bounds do not seem to be tight as in the case of the Isometry Theorem.

Signatures for topological spaces and functions. In Section 1 we saw two ways of building signatures for spaces, functions, or other topological objects. After building filtrations from these objects, the first approach computes their persistent homologies and compares them in the interleaving distance, whereas the second approach computes their persistence diagrams and compares them in the bottleneck distance. While the Isometry Theorem guarantees that these two approaches are equally powerful in terms of stability and discrimination power, persistence diagrams are consistently preferred over persistence modules in practice, for several reasons. First, they are easy to compute from filtrations, as we saw in Section 2 of Chapter 2. Second, they have a canonical representation in the plane, which is easy to visualize and readily interpretable. Third, and last, the bottleneck distance between them is comparatively easy to compute as it reduces to a bottleneck matching problem [115], whereas computing the interleaving distance is polynomially equivalent to deciding the solvability of some systems of multivariate quadratic equations [179]. These features make the persistence diagrams a practical tool to work with in applications, as we will see in the following chapters.

Part 2

Applications

Topological Inference

Discovering the structure of an unknown geometric object from a finite collection of data samples is becoming ubiquitous in the sciences. Indeed, the recent explosion in the amount and variety of available data (correlated with the increase of the world's technological storage capacity, going e.g. from 2.6 optimally compressed exabytes in 1986 to 15.8 in 1993, over 54.5 in 2000, and to 295 optimally compressed exabytes in 2007 [157]) calls for effective methods to organize them.

What kind of data are we referring to? Typically, clouds of points equipped with a notion of distance or (dis-)similarity between the points. A point in such a cloud can represent for instance a patch in an image, or an image in a collection, or a 3d shape in a database, or a user in a social network, etc. Examples include the MNIST handwritten digits database [176], the Columbia Object Image Library [203], the Princeton Shape Benchmark [223], the Social Computing Data Repository [241], or the Stanford Large Network Dataset Collection [178]. The exploration of such data sets serves two purposes: first, to interpret the data and identify intrinsic phenomena; second, to summarize the properties of the data for further processing or comparison tasks. In each case, uncovering the geometric structure(s) underlying the data is a key step, for which the following challenges arise:

- (1) The interpretation of the data is tied to the scale at which they are considered. We gave an illustrative example in the general introduction of the book (Figure 0.1).
- (2) The data can be corrupted by noise and outliers that hinder the identification of potential underlying structures. An illustration is given in Figure 4.1.
- (3) The data can live in high dimensions, which causes undesired phenomena (concentration of distances, exponential growth of the metric entropy, etc.) falling under the hood of the so-called *curse of dimensionality*. For instance, the COIL data set [203] has near 50,000 dimensions.
- (4) Finally, the data can be massive, which implies that they can no longer be stored and processed locally. For instance, some of the network data sets [178] contain hundreds of millions of edges.

These challenges call for the development of data analysis methods that are multi-scale, robust to noise, and highly parallelizable. But clearly, capturing the structure of arbitrarily large-dimensional geometric objects is out of reach due to the size of the required sampling. A common assumption across data analysis is that the objects underlying the input data have small intrinsic dimension m , regardless of how large the ambient dimension d may be. For instance, although the COIL data set [203] has dozens of thousands of dimensions, it lies close to some 1-dimensional structure as we saw in the general introduction of the book (Figure 0.4). Under

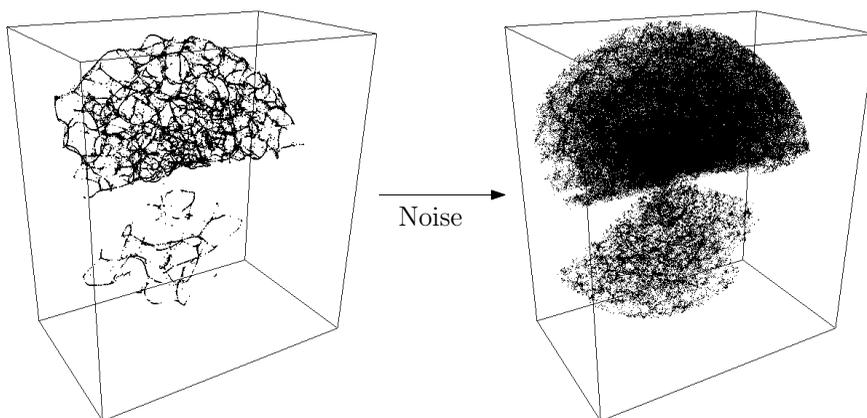


FIGURE 4.1. Matter is believed to be spread unevenly in the Universe, due to the influence of Gravity. Large-scale simulations suggest that matter gathers around filamentary structures, as shown in the 3-dimensional data set on the left—where each data point gives the position of the center of mass of an entire galaxy. However, physical measurements such as produced by telescopes are corrupted with noise and outliers hiding the filamentary structures, as shown on the right.

— The data presented here were produced from the database of the Sloan Digital Sky Survey [123].

the ‘large d - small m ’ assumption, a popular approach to overcome the curse of dimensionality is to map the data points to some lower-dimensional space, ideally of dimension proportional to m , a process known as *dimensionality reduction*. Not only does it help detect the intrinsic parameters of the data and remove the bad effects of high dimensionality, but it also reduces the algorithmic complexity of the problem and makes the use of classical techniques meant for small d applicable.

The recent years have seen the emergence of a new challenge in data analysis: topology. Indeed, dimensionality reduction assumes implicitly that the topological structure of the object underlying the data is simple, by assuming for instance linear or developable manifold structures [27, 230]. By contrast, it happens that modern datasets carry some nontrivial topology. Examples include the space of image patches [177], the layout of a wireless sensor field [98], or the energy landscape of a protein [187]—the first example will be described in detail in Section 5 of Chapter 5. Uncovering the topological structure carried by such data is paramount to performing their analysis. This is the subject of topological inference.

Topological inference. The usual setting is described as follows. The ambient space is \mathbb{R}^d , equipped with the Euclidean norm, denoted $\|\cdot\|$. In this space lives an object K , say a compact set, that remains unknown or hidden from us. Instead, we are given a finite set of points $P \subset \mathbb{R}^d$, called a *point cloud*, with the promise that the points of P live on or close to K in the ambient space. This is captured

by the *Hausdorff distance* d_H between P and K being small, say ε :

$$(4.1) \quad d_H(K, P) := \max \left\{ \sup_{x \in K} \inf_{p \in P} \|x - p\|, \sup_{p \in P} \inf_{x \in K} \|p - x\| \right\} = \varepsilon,$$

where the suprema and infima are in fact maxima and minima since both P and K are compact. Our task is then to uncover the topology of the unknown object K from the given finite sampling P . As we saw in the general introduction, this problem is known to be ill-posed, as several objects with different topological types can be sampled by P in the sense of (4.1). Further evidence of this phenomenon is given in Figure 4.2, where the object underlying the data can be either a simple closed curve or a torus. Both objects are perfectly valid underlying structures, however

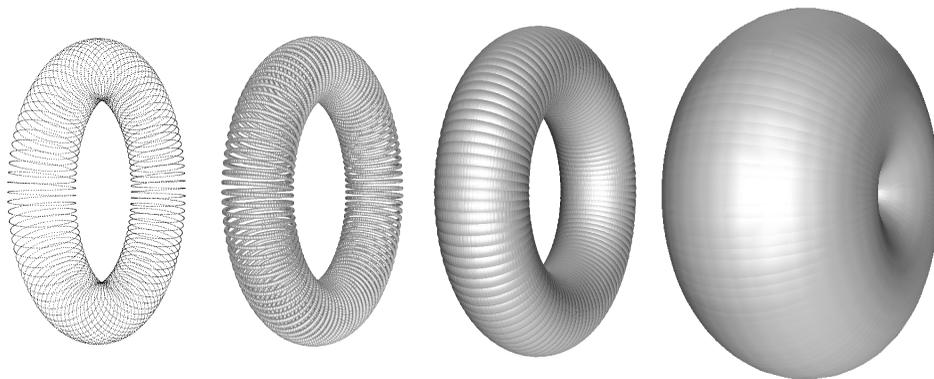


FIGURE 4.2. 10,000 points distributed uniformly along a curve winding around a torus in \mathbb{R}^3 . From left to right: i -offsets of the input points for increasing values of i , starting at $i = 0$. The second offset carries the homotopy type of the underlying curve, while the third offset carries the homotopy type of the underlying torus.

the difference between them resides in the scale at which the data are considered. To give an analogy, looking at the point cloud from far away reveals the torus to the observer, while looking at it from very close reveals the curve. The problem of choosing a relevant set of scales is ubiquitous in the literature on data analysis. This is where persistence theory comes into play: after deriving suitable filtrations from the input point cloud, such as for instance its *offsets filtration* defined below, one can use the persistence algorithm to quantify the importance of each topological feature by measuring its persistence across scales. Indeed, topological inference was and continues to be the most emblematic application of persistence theory, and one of the main motivations for its development [48].

The approach is backed up by a solid sampling theory that provides sufficient conditions under which the *offsets* of a compact set K can be approximated by the ones of a finite sampling P . To be more specific, the distance to K is defined over \mathbb{R}^d by

$$(4.2) \quad d_K(x) = \min_{y \in K} \|x - y\|,$$

and the i -*offset* of K , denoted K_i , is defined for any level $i \in \mathbb{R}$ as the sublevel set $d_K^{-1}((-\infty, i])$. This set is empty when $i < 0$, equal to K when $i = 0$, and one has

$\bigcup_{i \in \mathbb{R}} K_i = \mathbb{R}^d$. Thus, the family of offsets of K for i ranging over \mathbb{R} forms a filtration of \mathbb{R}^d , called the *offsets filtration* of K and denoted \mathcal{K} . Similarly, one can define the distance function d_P of P , its i -offsets for all $i \in \mathbb{R}$, and its offsets filtration \mathcal{P} . The same goes for any compact set in \mathbb{R}^d . Assuming Hausdorff proximity between K and P as in (4.1) is equivalent to assuming sup-norm proximity between their distance functions:

$$(4.3) \quad \|d_K - d_P\|_\infty = \varepsilon.$$

This implies in particular that the offsets filtrations of K and P are ε -interleaved in the sense of (3.3), that is:

$$(4.4) \quad \forall i \in \mathbb{R}, K_i \subseteq P_{i+\varepsilon} \text{ and } P_i \subseteq K_{i+\varepsilon}.$$

Therefore the persistence diagrams $\text{dgm}(\mathcal{K})$ and $\text{dgm}(\mathcal{P})$ are ε -close in the bottleneck distance, as guaranteed by the Isometry Theorem 3.1. The question now is whether the *topological signal* carried by $\text{dgm}(\mathcal{K})$ can be distinguished from the extra *topological noise* present in $\text{dgm}(\mathcal{P})$, despite the fact that the bottleneck matching between the two diagrams is unknown. The aim of the sampling theory is precisely to quantify the *signal-to-noise ratio* in $\text{dgm}(\mathcal{P})$ with respect to the sampling density—measured by ε in (4.1)—on the one hand, with respect to the ‘regularity’ of the geometry of K —measured by some quantity to be defined—on the other hand. The rationale is that when the signal-to-noise ratio is large enough, the user can read off the topology of K (more precisely, its homology or cohomology) from the persistence diagram of \mathcal{P} .

We briefly introduce the theory of distance functions in Section 1, and explain how it acts as a sampling theory for topological inference in Section 2. The filtrations involved in this theory are made of offsets of compact sets, which are continuous objects and therefore not naturally amenable to manipulation on a computer. For practical purposes it is therefore necessary to replace them by purely combinatorial filtrations that carry the same topological information. The mechanism is described in Section 3, where we introduce the two main simplicial filtrations considered for this purpose: the so-called *Čech filtration* and *α -complex filtration*—the latter being renamed *Delaunay filtration* in the following for clarity.

The main interest of our community, as reflected in this chapter, is to lay down the mathematical foundations of a working pipeline for doing topological inference on a computer using persistence. The pipeline goes as follows: it takes a point cloud P as input, builds a simplicial filtration (Čech, Delaunay) on top of P , computes its persistence barcode using the persistence algorithm or one of its variants from Chapter 2, and passes the result on to the user for interpretation. Proving this result correct under various sampling conditions is the topic of this chapter and was the original aim of our community. Algorithmic questions, including the use of more lightweight simplicial filtrations to scale up nicely with the size and dimensionality of the data, started to be addressed later and are therefore deferred to the next chapter.

Prerequisites. For more background on data analysis we recommend reading [151], especially Chapter 14. For a survey of linear and non-linear dimensionality reduction techniques, see [236]. Much of the current chapter is devoted to introducing distance functions and their inference properties. The literature on the subject is vast and bears many connections to data analysis. Here we assume no prior exposure to the subject and we restrict ourselves to a few selected results—Theorems 4.1,

4.4 and 4.7—that are most relevant to us, referring the interested reader to [59] and the references therein for a more comprehensive survey. The rest of the material used in the chapter comes from Part 1, with the notable exception of the Nerve Lemma whose statement is recalled for the reader’s convenience—background on homotopy equivalences and deformation retractions can be found e.g. in Chapter 0 of [152].

1. Inference using distance functions

As mentioned in the introduction of the chapter, the ambient space is \mathbb{R}^d equipped with the Euclidean norm $\|\cdot\|$. This is in fact a simplifying assumption, as the concepts and techniques presented in this section, in particular the generalized gradient vector field and its associated flow, extend naturally to Riemannian manifolds, and beyond that, to Alexandrov spaces [211], so the theory can be rewritten in this more general context.

1.1. Reach. The set of points $y \in K$ that realize the minimum in (4.2) is called the *projection set* of x , denoted $\Pi_K(x)$. This set is never empty, and when it is a singleton $\{y\}$ we let $\pi_K(x) = y$ and call y the *projection* of x . The map $x \mapsto \pi_K(x)$ is defined everywhere in \mathbb{R}^d except on the *medial axis* of K , denoted $\mathfrak{M}(K)$, which is the set of points x such that $|\Pi_K(x)| \geq 2$. See Figure 4.3 for an example showing that the medial axis may be neither open nor closed.

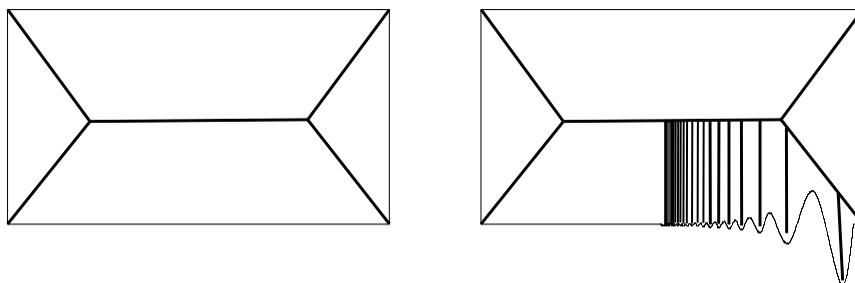


FIGURE 4.3. Left: the medial axis of the boundary of a rectangle, marked in bold lines, is clearly not open. Right: the medial axis of the same rectangle whose bottom edge has been replaced by the C^1 -continuous function $x \mapsto x^3 \sin \frac{1}{x}$ for $x \geq 0$ and $x \mapsto 0$ for $x < 0$, does not contain the limit vertical segment above the C^2 -discontinuity at $x = 0$ and is therefore not closed either.

— Reprinted from Lieutier [180]: “Any open bounded subset of \mathbb{R}^n has the same homotopy type as its medial axis”, *Computer-Aided Design*, 36(11):1029–1046, ©2004, with permission from Elsevier.

The following properties are well-known and easy to prove [128]. They involve the topological closure of the medial axis, denoted $\overline{\mathfrak{M}}(K)$:

- d_K^2 is continuously differentiable over $\mathbb{R}^d \setminus \overline{\mathfrak{M}}(K)$, where its gradient is $2(x - \pi_K(x))$. Note that the gradient vanishes on K .

- d_K is 1-Lipschitz over \mathbb{R}^d and continuously differentiable over $\mathbb{R}^d \setminus (K \cup \overline{\mathfrak{M}}(K))$, where its gradient, denoted ∇_K , is

$$(4.5) \quad \nabla_K(x) = \frac{x - \pi_K(x)}{\|x - \pi_K(x)\|} = \frac{x - \pi_K(x)}{d_K(x)}.$$

Note that $\|\nabla_K\| = 1$.

- The projection map π_K is continuous, so the opposite gradient vector field $-\nabla_K$ can be integrated into a continuous flow:

$$(4.6) \quad \begin{aligned} \mathbb{R}^+ \times \mathbb{R}^d \setminus \overline{\mathfrak{M}}(K) &\longrightarrow \mathbb{R}^d \setminus \overline{\mathfrak{M}}(K), \\ (t, x) &\longmapsto \begin{cases} x - t\nabla_K(x) & \text{if } t < d_K(x), \\ \pi_K(x) & \text{otherwise.} \end{cases} \end{aligned}$$

Note however that the gradient field ∇_K itself does not always yield a continuous flow. An illustration is given in Figure 4.3 (right), where the induced flow has a discontinuity across the limit vertical segment above the C^2 -discontinuity of K .

In this context, the *reach* of K is defined as the minimum distance between K and its medial axis. Specifically, the Euclidean distance of a point $x \in K$ to $\overline{\mathfrak{M}}(K)$ is called the *reach of K at x* :

$$(4.7) \quad \text{rch}(K, x) = d_{\overline{\mathfrak{M}}(K)}(x),$$

and the infimum (in fact minimum) of this quantity over K is the *reach of K* :

$$(4.8) \quad \text{rch}(K) = \min_{x \in K} \text{rch}(K, x).$$

The concept of reach was introduced by Federer [128] to define curvature measures on sets that are neither convex nor C^2 -continuous—see Figure 4.4 for an example. The reach at a point $x \in K$ is also sometimes called the *local feature size* at x in the literature [3].

When a compact set K has positive reach, i.e. when it does not come close to its medial axis, the flow of (4.6) can be applied in its neighborhood. Applying it to the points of an offset K_j with $j < \text{rch}(K)$ results in a deformation retraction of that offset to K . Along the way, it also results in weak deformation retractions of K_j to the offsets K_i with $0 < i < j$. Hence, the inclusion maps $K \hookrightarrow K_i \hookrightarrow K_j$ are homotopy equivalences and therefore induce isomorphisms at the homotopy and homology levels.

Niyogi, Smale, and Weinberger [206] proposed to reproduce this argument with the offsets of K replaced by the ones of a sufficiently close point cloud P . The main bottleneck was to prove that the intersections of the flow lines of K with the offsets of P are contractible, which they were able to do under the extra condition that K is a submanifold of \mathbb{R}^d . This was the first inference result using distance functions.

THEOREM 4.1 (Niyogi, Smale, and Weinberger [206]). *Let K be a compact submanifold of \mathbb{R}^d with positive reach. Let P be a point cloud in \mathbb{R}^d such that $d_H(K, P) = \varepsilon < \sqrt{\frac{3}{20}} \text{rch}(K)$. Then, for any level $i \in (2\varepsilon, \sqrt{\frac{3}{5}} \text{rch}(K))$, the offset P_i deformation retracts onto K , so the inclusion map $K \hookrightarrow P_i$ is a homotopy equivalence.*

It turns out that both the manifold and the positive reach conditions are superfluous, as we will see next.

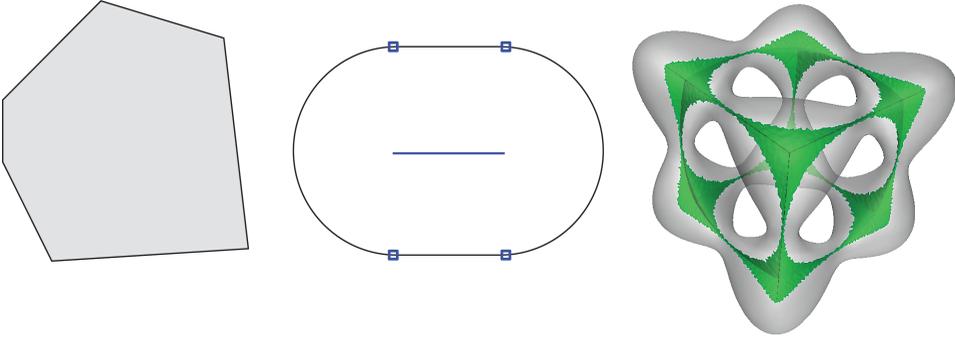


FIGURE 4.4. Some compact sets with positive reach. Left: the solid polygon is convex, therefore its medial axis is empty and its reach is infinite. Center: the curve made of two half-circles connected by two line segments has C^2 -discontinuities at the gluing points (marked by squares). Its medial axis is the central horizontal line segment. Right: the quartic surface called ‘tanglecube’ is C^2 -continuous. It is shown here in transparency, superimposed with an approximation of its inner medial axis. Note that the sets in Figure 4.3 have zero reach.

1.2. Weak feature size. Although the distance function d_K is not differentiable on the medial axis of K , it is possible to extend its gradient to a *generalized gradient* function $\nabla_K : \mathbb{R}^d \setminus K \rightarrow \mathbb{R}^d$ as follows:

$$(4.9) \quad \nabla_K(x) = \frac{x - c(\Pi_K(x))}{d_K(x)},$$

where $c(\Pi_K(x))$ denotes the center of the minimum enclosing ball of the projection set $\Pi_K(x)$. This center plays the role of the projection $\pi_K(x)$ in the original gradient, and it is in fact equal to it outside $\mathfrak{M}(K)$, so (4.9) extends (4.5) in a natural way. We let $\nabla_K = 0$ over K by convention, so ∇_K is defined over the entire space \mathbb{R}^d . For $x \notin K$ we have

$$\|\nabla_K(x)\|^2 = 1 - \frac{r(\Pi_K(x))^2}{d_K(x)^2},$$

where $r(\Pi_K(x))$ denotes the radius of the smallest enclosing ball of $\Pi_K(x)$. Equivalently, $\|\nabla_K\|$ is the cosine of the half-angle of the smallest cone of apex x that contains $\Pi_K(x)$. See Figure 4.5 for an illustration.

Although the generalized gradient vector field ∇_K is not continuous, it is locally semi-Lipschitz in the sense that

$$(\nabla_K(x) - \nabla_K(y)) \cdot (x - y) \leq \frac{1}{i}(x - y)^2$$

for any $i > 0$ and any points $x, y \notin K_i$. Moreover, the map $x \mapsto \|\nabla_K(x)\|$ is lower-semicontinuous, meaning that

$$\liminf_{y \rightarrow x} \|\nabla_K(y)\| \geq \|\nabla_K(x)\|$$

for any $x \in \mathbb{R}^d$. These are important properties since they allow to integrate the generalized gradient vector field ∇_K using Euler schemes. As the integration step

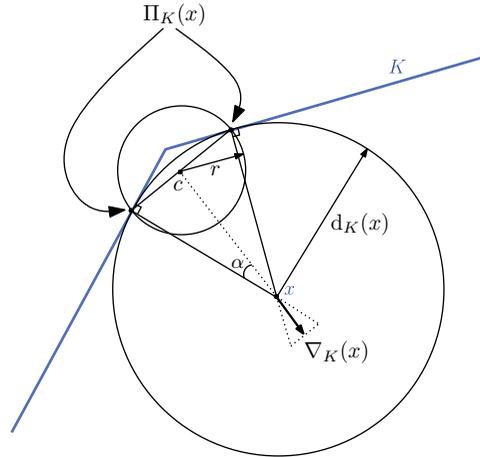


FIGURE 4.5. The generalized gradient of the distance to K , where c and r are shorthands for $c(\Pi_K(x))$ and $r(\Pi_K(x))$ respectively. The norm of the gradient is given by $\|\nabla_K(x)\| = \cos \alpha$.

— Reprinted from Chazal, Cohen-Steiner, and Lieutier [60]: “A Sampling Theory for Compact Sets in Euclidean Space”, *Discrete & Computational Geometry*, 41(3):461–479, ©2009, with kind permission from Springer Science and Business Media.

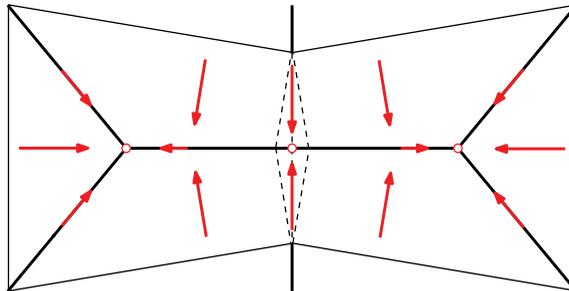


FIGURE 4.6. Generalized gradient of d_K and its associated flow. Bold lines mark the medial axis, arrows mark the gradient and also the direction of the flow, circles mark the critical points of d_K located outside of K .

— Based on Lieutier [180].

decreases, the Euler schemes converge uniformly to a continuous flow $\mathbb{R}^+ \times \mathbb{R}^d \rightarrow \mathbb{R}^d$ that is right-differentiable and whose right derivative coincides with ∇_K [180].

An illustration of the gradient and its flow is given in Figure 4.6. Note that the flow line l starting at a point $x \in \mathbb{R}^d$ can be parametrized by arc length, so it is possible to integrate $\|\nabla_K\|$ along l to obtain the value of $d_K(y)$ at any downstream point y . In this respect, ∇_K acts as a gradient for d_K over \mathbb{R}^d . With the new gradient flow at hand, one can construct the same kind of deformations retractions as in Section 1.1. However, extending the construction to the entire space \mathbb{R}^d requires to develop a generalized critical point theory.

DEFINITION 4.2. A point $x \in \mathbb{R}^d$ is *critical* if its generalized gradient $\nabla_K(x)$ is zero, or equivalently, if x lies in the convex hull of its projection set $\Pi_K(x)$. A value $v \in \mathbb{R}$ is *critical* if $v = d_K(x)$ for some critical point x . In particular, all points of K are critical, therefore 0 is a critical value.

This definition agrees with the ones from nonsmooth analysis [83] and Riemannian geometry [74, 145]. In particular, we can rely on the following result by Grove [145], proved independently by Chazal and Lieutier [57] using the generalized gradient flow. This result relates the offsets of K in the same spirit as in Section 1.1, but for all offset values $i > 0$. It extends a classical result of Morse theory [195, theorem 3.1] to the current nonsmooth setting:

LEMMA 4.3. *If $0 < i < j$ are such that there is no critical value of d_K within the closed interval $[i, j]$, then K_j deformation retracts onto K_i , so the inclusion map $K_i \hookrightarrow K_j$ is a homotopy equivalence.*

In this context, the *weak feature size* of K is defined as follows:

$$(4.10) \quad \text{wfs}(K) = \inf\{i > 0 \mid i \text{ is a critical value of } d_K\}.$$

It is clear that $\text{rch}(K) \leq \text{wfs}(K)$ since the critical points of d_K lie either on K or on its medial axis. The inequality can be strict, as in Figure 4.6, so the class of compact sets with positive weak feature size is larger than the one of compact sets with positive reach. It is in fact much larger, containing for instance all polyhedra or piecewise analytic sets [57].

When $\text{wfs}(K)$ is nonzero, Lemma 4.3 guarantees that all the offsets K_i for $i \in (0, \text{wfs}(K))$ are homotopy equivalent. Thus, the weak feature size can be viewed as ‘the minimum size of the topological features’ of the compact set K . Note that the homotopy equivalence does not extend to the 0-offset K in general, as 0 is a critical value of d_K by definition, and furthermore there are examples of compact sets K that do not have the same homotopy type as their small offsets. A well-known such example is given in Figure 4.7.

Chazal and Lieutier [65] introduced the weak feature size and proposed the following generalization of Theorem 4.1—see also [87] for a similar result:

THEOREM 4.4. *Let K be a compact set in \mathbb{R}^d with positive weak feature size. Let P be a point cloud in \mathbb{R}^d such that $d_H(K, P) = \varepsilon < \frac{1}{4}\text{wfs}(K)$. Then, for any levels i, j such that $\varepsilon < i < i + 2\varepsilon \leq j < \text{wfs}(K) - \varepsilon$, the persistent homology group $\text{im } H_*(P_i) \rightarrow H_*(P_j)$ induced by the inclusion map $P_i \hookrightarrow P_j$ is isomorphic to $H_*(K_r)$ for any $r \in (0, \text{wfs}(K))$.*

The result follows from basic rank arguments¹ once it has been observed that the following sequence of inclusions is implied by the hypothesis and (4.4):

$$K_{i-\varepsilon} \hookrightarrow P_i \hookrightarrow K_{i+\varepsilon} \hookrightarrow P_j \hookrightarrow K_{j+\varepsilon},$$

where compositions $K_{i-\varepsilon} \hookrightarrow K_{i+\varepsilon} \hookrightarrow K_{j+\varepsilon}$ induce isomorphisms at the homology level according to Lemma 4.3.

It is important to note that using persistent homology groups instead of homology groups of offsets of P is necessary to capture the homology of the small offsets K_r of K in Theorem 4.4. Indeed, there are cases where there is not a single

¹Chazal and Lieutier [65] also proved the result for homotopy groups, for which the proof is somewhat more elaborate but keeps the same flavor.

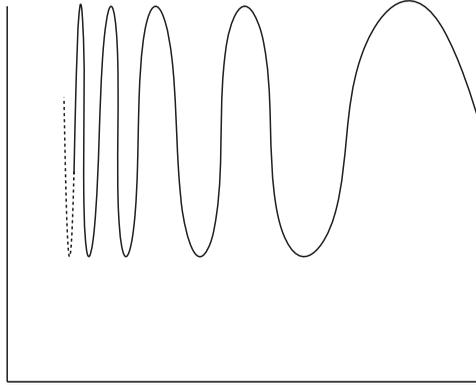


FIGURE 4.7. The compact set $K \subset \mathbb{R}^2$ is the union of the four sets $\{(x, y) \mid x = 0, -2 \leq y \leq 1\}$, $\{(x, y) \mid 0 \leq x \leq 1, y = -2\}$, $\{(x, y) \mid x = 1, -2 \leq y \leq 0\}$, and $\{(x, y) \mid 0 < x \leq 1, y = \sin \frac{2\pi}{x}\}$. This set K is simply connected with positive weak feature size, however its (small enough) offsets are homeomorphic to annuli and therefore not homotopy equivalent to K .

— Based on Spanier [228], §2.4.8.

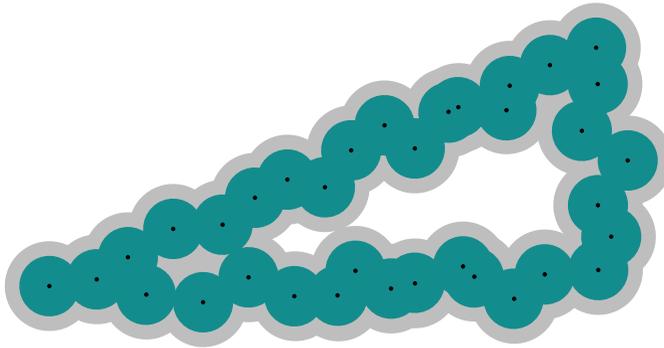


FIGURE 4.8. Sampling P of the boundary K of a triangle in the plane. Because of the small angle in the triangle and of the noise in the sampling, there is not a single level i such that the i -offset of P captures the homology of K . Only much denser samplings P have this property, and reducing the angle arbitrarily makes the threshold in density tend to infinity while keeping the weak feature size of K constant. Yet, for a fixed angle there is a finite threshold in density beyond which single offsets of P capture the homology of K , which suggests that there may be a meaningful intermediate quantity between the weak feature size and the reach.

— Based on Chazal and Cohen-Steiner [59].

level i such that P_i has the same homological type as K_r . Figure 4.8 gives such an example, and suggests that there may be an intermediate class of compact sets K between those with positive weak feature size and those with positive reach, which

realizes a trade-off between weakening the sampling conditions on P and still capturing the homology of the small offsets K_r from single offsets of P . This is the ‘raison d’être’ of the μ -reach, presented in the next section.

1.3. μ -reach. A point $x \in \mathbb{R}^d$ is called μ -critical if $\|\nabla_K(x)\| < \mu$, and the set of μ -critical points located outside K is called the μ -medial axis² of K , denoted $\mathfrak{M}_\mu(K)$. Note that a point of $\mathbb{R}^d \setminus K$ is 1-critical if and only if it belongs to the medial axis of K , and μ -critical for all $\mu > 0$ if and only if it is critical in the sense of Definition 4.2. Thus, outside K , the nested family of μ -medial axes interpolates between the medial axis from Section 1.1 and the set of critical points from Section 1.2. In this context, the μ -reach of K , denoted $r_\mu(K)$, was introduced by Chazal, Cohen-Steiner, and Lieutier [60] as an interpolating quantity between the reach and the weak feature size of K , defined as follows—notice the similarity with (4.8):

$$(4.11) \quad \forall \mu \in (0, 1], \quad r_\mu(K) = \min_{x \in K} d_{\overline{\mathfrak{M}_\mu(K)}}(x).$$

As μ decreases, $\mathfrak{M}_\mu(K)$ shrinks therefore $r_\mu(K)$ increases. Moreover, we have $r_1(K) = \text{rch}(K)$ on the one hand, and $\lim_{\mu \rightarrow 0^+} r_\mu(K) = \text{wfs}(K)$ provided that $\lim_{\mu \rightarrow 0^+} r_\mu(K) > 0$ on the other hand³. Thus, $r_\mu(K)$ indeed interpolates between $\text{rch}(K)$ and $\text{wfs}(K)$.

The main property of the μ -critical points of d_K is to be stable under small Hausdorff perturbations of K :

LEMMA 4.5 (Chazal, Cohen-Steiner, and Lieutier [60]). *Let K, K' be compact sets in \mathbb{R}^d , and let $\varepsilon = d_H(K, K')$. Then, for any μ -critical point x of $d_{K'}$, there is a $(2\sqrt{\varepsilon/d_{K'}(x)} + \mu)$ -critical point of d_K at distance at most $2\sqrt{\varepsilon d_{K'}(x)}$ from x .*

This stability property implies a separation between the critical values of $d_{K'}$ when K has positive μ -reach:

LEMMA 4.6 (Chazal, Cohen-Steiner, and Lieutier [60]). *Let K, K' be compact sets in \mathbb{R}^d , and let $\varepsilon = d_H(K, K')$. Then, for any $\mu \in (0, 1]$, there is no critical value of $d_{K'}$ within the range $(4\varepsilon/\mu^2, r_\mu(K) - 3\varepsilon)$.*

Taking $\mu = 1$ in the above statement gives that the critical points of $d_{K'}$ are concentrated around K and its medial axis when K has positive reach, a well-known separation result (at least in the case of smoothly embedded surfaces) [109].

Applying Lemma 4.6 with $K' = P$ in the context of the proof of Theorem 4.4 allows to relate the homotopy type of small offsets of K directly to that of single offsets of P :

THEOREM 4.7 (Chazal, Cohen-Steiner, and Lieutier [60]). *Let K be a compact set in \mathbb{R}^d with positive μ -reach for some $\mu \in (0, 1]$. Let P be a point cloud in \mathbb{R}^d such that $d_H(K, P) = \varepsilon < \frac{\mu^2}{5\mu^2 + 12} r_\mu(K)$. Then, for any level i such that $\frac{4\varepsilon}{\mu^2} \leq i < r_\mu(K) - 3\varepsilon$, the i -offset of P is homotopy equivalent to K_r for any $r \in (0, \text{wfs}(K))$.*

²Not to be confused with the λ -medial axis of Chazal and Lieutier [57], which is closely related but not equal.

³One may have $\lim_{\mu \rightarrow 0^+} r_\mu(K) < \text{wfs}(K)$ when $\lim_{\mu \rightarrow 0^+} r_\mu(K) = 0$. For instance, when K is the union of two tangent disks in the plane, $\mathfrak{M}_\mu(K)$ converges to the emptyset as μ tends to 0, so $\text{wfs}(K) = +\infty$ whereas $r_\mu(K) = 0$ for all $\mu > 0$.

REMARK. The statement can be further strengthened into an isotopic reconstruction theorem involving the level sets of d_K and d_P . The idea is to use the stability properties of the generalized gradient vector field to turn it into a C^∞ -continuous vector field that is ‘transverse’ to the level sets of d_K and d_P . Isotopies between level-sets of K and of P are then derived from the flow induced by this modified vector field. The details can be found in [61].

1.4. Distance-like functions. As pointed out by Chazal, Cohen-Steiner, and M erigot [62], the generalized gradient vector field and its induced flow can be built not only for the distance to a compact set K , but also for any *distance-like* function, i.e. any non-negative function $d : \mathbb{R}^d \rightarrow \mathbb{R}$ that satisfies the following axioms:

- A1 $d(x)$ tends to infinity whenever x does,
- A2 d is 1-Lipschitz, and
- A3 d^2 is 1-semiconcave, i.e. the map $x \mapsto \|x\|^2 - d^2(x)$ is convex.

The Lipschitz continuity implies that d is differentiable almost everywhere. In particular, the *medial axis*, defined as the set of points of \mathbb{R}^d where d is not differentiable, has zero d -volume. The semiconcavity imposes further regularity on d , in particular twice differentiability almost everywhere. It plays a central role in the proof of existence of the flow induced by the generalized gradient of d .

Suppose now we are working with a given class of objects living in Euclidean space \mathbb{R}^d , such as for instance its compact subsets. Suppose we can derive a distance-like function f_K from each object K , in such a way that the map $K \mapsto f_K$ is Lipschitz continuous with respect to some metric between objects and to the supremum distance between distance-like functions. For instance, we saw in (4.3) that the map $K \mapsto f_K$ is 1-Lipschitz in the class of compact subsets of \mathbb{R}^d equipped with the Hausdorff distance. Then, the analysis of Sections 1.2 and 1.3 can be reproduced and inference results similar to Theorems 4.4 and 4.7 (up to constant factors) can be stated.

This observation allows to generalize the topological inference theory beyond the mere compact subsets of \mathbb{R}^d , to other classes of objects such as probability measures, as we will see in Section 6 of the next chapter.

2. From offsets to filtrations

The results of Section 1 guarantee the existence of offsets (or pairs of offsets) of the input point cloud P that carry the topology of its underlying object K , under some sampling conditions. However, the quantities involved in the bounds, such as the reach or the weak feature size of K , remain unknown, so in practice the user is left with the difficult question of choosing the ‘right’ offset parameter(s), which amounts to choosing the ‘right’ scale(s).

This is where persistence comes in. By looking at all scales at once, and by encoding the evolution of the topology of the offsets across scales, it gives valuable feedback to the user for choosing relevant scale(s). The choice is guided by the results of Section 1, which translate into statements about the structure of the persistence barcode or diagram of the offsets filtration \mathcal{P} of P , describing where and under what form the homology of K appears. The barcode representation is usually preferred over diagram representation in this context because it is more readily interpretable, as Figures 4.9 and 4.10 illustrate. We will therefore use it in the statements.

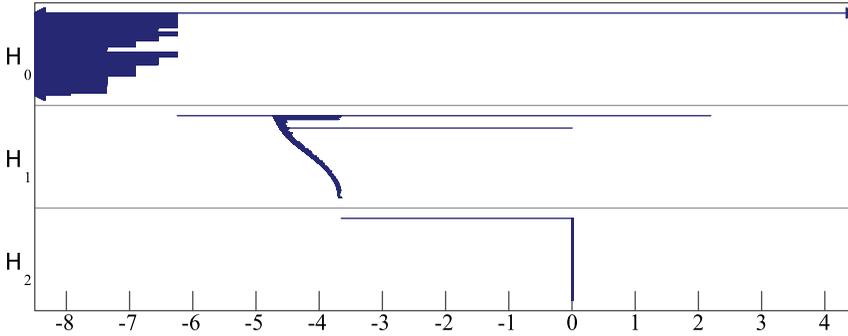


FIGURE 4.9. Barcode of the offsets filtration of the point cloud from Figure 4.2, drawn on a logarithmic scale. Left/Right arrow heads indicate left/right-infinite intervals. The sweet range for the helicoidal curve is roughly $(-6.25, -4.8)$, the one for the torus is approximately $(-3.7, 0)$. There is also a sweet range $(0, 2)$ for the solid torus, and a sweet range $(2, +\infty)$ for the entire ambient space \mathbb{R}^3 .

REMARK. Before stating the results we must check that \mathcal{P} indeed has a well-defined persistence diagram. It turns out that all distances to compact sets in \mathbb{R}^d have well-defined diagrams. This is a consequence of their being q -tame by Proposition 2.3 (ii), and it holds regardless of the regularity of the compact sets.

Compact sets with positive weak feature size. If we apply Theorem 4.4 with $i \rightarrow \varepsilon^+$ and $j \rightarrow (\text{wfs}(K) - \varepsilon)^-$ on the one hand, with $j = i + 2\varepsilon$ on the other hand, then we obtain the following guarantee on the existence of a *sweet range* over which the homology of the underlying compact set K can be read off:

COROLLARY 4.8. *Let K be a compact set in \mathbb{R}^d with positive weak feature size. Let P be a point cloud in \mathbb{R}^d such that $d_{\text{H}}(K, P) = \varepsilon < \frac{1}{4}\text{wfs}(K)$. Then, there is a sweet range $T = (\varepsilon, \text{wfs}(K) - \varepsilon)$ whose intersection with the barcode of the offsets filtration \mathcal{P} has the following properties:*

- *The intervals that span T encode the homology of K , in the sense that their number for each homology dimension p is equal to the dimension of $H_p(K_r)$, for any $r \in (0, \text{wfs}(K))$.*
- *The remaining intervals have length at most 2ε .*

This result partitions the restriction of the barcode of \mathcal{P} to the sweet range T into two classes: on the one hand, the intervals spanning T compose what is called the *topological signal*; on the other hand, intervals not spanning T compose what is called the *topological noise*⁴. The ratio of the minimum length of a signal interval (which is also the width of the sweet range) to the maximum length of a noise interval in T is called the *signal-to-noise ratio*. Under the hypotheses of the theorem, it is guaranteed to be at least $\frac{\text{wfs}(K) - 2\varepsilon}{2\varepsilon}$.

⁴Note that there may be long intervals in the full barcode whose restrictions to T are treated as topological noise because they do not span T . The theorem guarantees that such intersections are short.

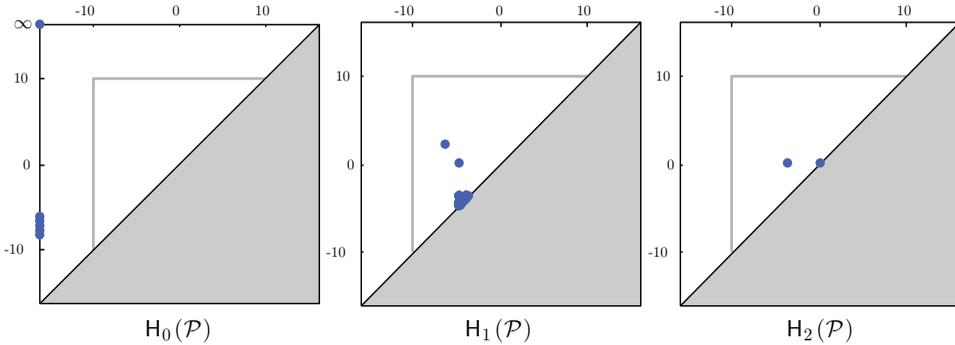


FIGURE 4.10. Persistence diagrams of the offsets filtration of the point cloud P from Figure 4.2 for homology dimensions 0, 1, 2 (from left to right). The diagrams are plotted on a log-log scale.

An illustration is given in Figure 4.9, showing the barcode of the offsets filtration of the point cloud from Figure 4.2. The barcode is represented on a logarithmic scale to better reveal the phenomena occurring at smaller scales. As expected, sweet ranges for the helicoidal curve and for the torus appear, but not only. The barcode also reveals two other underlying structures: the solid torus, and the entire space \mathbb{R}^d at the largest scales. At the other end of the barcode, i.e. at the lowest scales, the point cloud itself can be viewed as its own underlying space. These structures are also revealed by the persistence diagram, shown in Figure 4.10, however in a less readable form⁵, which justifies the use of the barcode representation in practice.

It is worth pointing out that in this example the sweet ranges contain no topological noise and therefore have infinite signal-to-noise ratio. This is a consequence of the underlying structures having positive μ -reach, as we will see next.

Compact sets with positive μ -reach. Applying Theorem 4.7 gives the stronger guarantee that there exists a *sweeter range* over which the homology of the compact set K is encoded with no noise at all (hence an infinite signal-to-noise ratio):

COROLLARY 4.9. *Let K be a compact set in \mathbb{R}^d with positive μ -reach for some $\mu \in (0, 1]$. Let P be a point cloud in \mathbb{R}^d such that $d_H(K, P) = \varepsilon < \frac{\mu^2}{5\mu^2 + 12} r_\mu(K)$. Then, there is a sweeter range $T = \left[\frac{4\varepsilon}{\mu^2}, r_\mu(K) - 3\varepsilon \right]$ whose intersection with the barcode of the offsets filtration \mathcal{P} has the following properties:*

- *The intervals that span T encode the homology of K , in the sense that their number for each homology dimension p is equal to the dimension of $H_p(K_r)$, for any $r \in (0, \text{wfs}(K))$.*
- *There are no other intervals.*

Since compact sets with positive μ -reach also have positive weak feature size, their sweeter ranges are in fact included in larger sweet ranges. Needless to say that the bounds on the sweeter range given in Corollary 4.9 indeed lie between the bounds on the sweet range from Corollary 4.8. As μ goes down while the Hausdorff

⁵For instance, distinguishing the helicoidal curve from the torus requires some imagination, although the difference does exist in the abscissae of the diagram points.

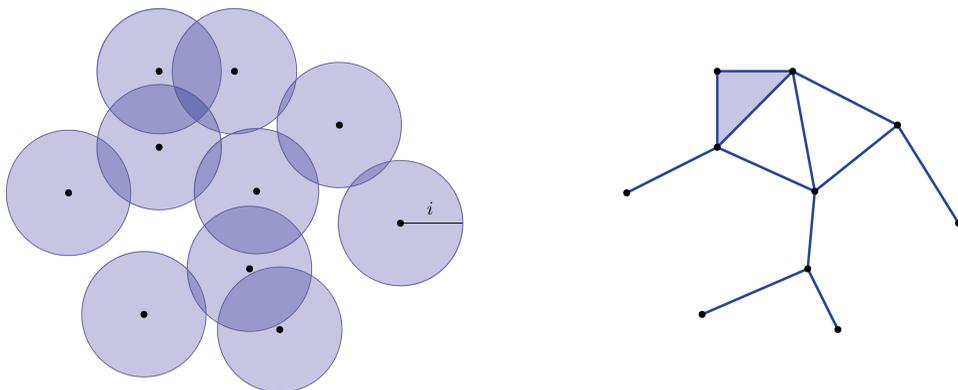


FIGURE 4.11. Left: a point cloud P surrounded by balls of radius i . Right: the corresponding Čech complex $C_i(P)$.

distance ε stays fixed, the sweeter range shrinks until it eventually becomes empty, when μ passes below $\sqrt{\frac{4\varepsilon}{r_\mu(K)-3\varepsilon}}$. This was the phenomenon observed empirically in the example of Figure 4.8.

3. From filtrations to simplicial filtrations

Offsets of compact sets are continuous objects, as such they are not naturally amenable to manipulation on a computer. Nevertheless, only their topology matters in our context, not their geometry. In the special case where the compact set is a finite point cloud P , its offsets are just finite unions of balls, whose topology can be accessed through various combinatorial constructions. The most notorious one among them is the so-called *Čech complex*⁶, defined as follows and illustrated in Figure 4.11:

DEFINITION 4.10. The *Čech complex* of P of parameter $i \in \mathbb{R}$, denoted $C_i(P)$, has one k -simplex per $(k+1)$ -tuple of points of P such that the closed Euclidean balls of radius i about these points have a nonempty common intersection. The *Čech filtration* is the indexed family $\mathcal{C}(P) = \{C_i(P)\}_{i \in \mathbb{R}}$.

Each Čech complex $C_i(P)$ is related to the corresponding offset P_i of P through the *Nerve Lemma*. Given a collection $\mathfrak{U} = \{U_a\}_{a \in A}$ of subspaces of a same topological space, the *nerve* $N(\mathfrak{U})$ has one k -simplex per $(k+1)$ -tuple of elements of \mathfrak{U} that have a nonempty common intersection. In our case the collection is formed by the Euclidean balls of radius i centered at the points of P , and its nerve is precisely the Čech complex $C_i(P)$. There exist many variations of the Nerve Lemma in the literature. They usually claim that if the sets U_a are ‘nice’ in some sense, and if their $(k+1)$ -fold intersections for all $k \in \mathbb{N}$ are either empty or topologically ‘simple’ in some sense, then the nerve $N(\mathfrak{U})$ is homotopy equivalent to the union $\bigcup_{a \in A} U_a$, which in our case is precisely the offset P_i . The most classical variant uses open sets and asks the nonempty intersections to be contractible—see e.g. [152, §4G]:

⁶Named after E. Čech for its relationship to Čech homology [122].

LEMMA 4.11 (Nerve). *Let X be a paracompact space, and let \mathfrak{U} be an open cover of X such that the $(k+1)$ -fold intersections of elements of \mathfrak{U} are either empty or contractible for all $k \in \mathbb{N}$. Then, there is a homotopy equivalence $N(\mathfrak{U}) \rightarrow X$.*

Unfortunately, in our setting the elements of the cover are closed, so the result does not apply directly. In fact, the Nerve Lemma is generally not true for closed covers, and finding sufficient conditions under which it holds is a research topic in its own right, with deep connections to the theory of retracts initiated by Borsuk [35]. The bottomline is that the Nerve Lemma holds for a closed cover \mathfrak{U} as soon as the $(k+1)$ -fold intersections of its elements are neighborhood retracts in X , i.e. each intersection V admits a retraction $r : W \rightarrow V$ from some open neighborhood W [34]. Luckily for us, this is the case for intersections of finitely many Euclidean balls.

Chazal and Oudot [66] have extended Lemma 4.11 so it guarantees not only pointwise homotopy equivalence between the spaces $C_i(P)$ and P_i for all $i \in \mathbb{R}$, but also equivalence between the filtrations $\mathcal{C}(P)$ and \mathcal{P} themselves. The proof is a straightforward adaptation of the one of Lemma 4.11 found in [152, §4G], and the conditions under which it adapts to closed covers are the same⁷:

LEMMA 4.12 (Persistent Nerve). *Let $X \subseteq X'$ be two paracompact spaces, and let $\mathfrak{U} = \{U_a\}_{a \in A}$ and $\mathfrak{U}' = \{U'_{a'}\}_{a' \in A'}$ be open covers of X and X' respectively, based on finite parameter sets $A \subseteq A'$. Assume that $U_a \subseteq U'_a$ for all $a \in A$, and that the $(k+1)$ -fold intersections of elements of \mathfrak{U} (resp. of \mathfrak{U}') are either empty or contractible for all $k \in \mathbb{N}$. Then, the homotopy equivalences $N(\mathfrak{U}) \rightarrow X$ and $N(\mathfrak{U}') \rightarrow X'$ provided by the Nerve Lemma 4.11 commute with the inclusion maps $X \hookrightarrow X'$ and $N(\mathfrak{U}) \hookrightarrow N(\mathfrak{U}')$ at the homology level, that is, the following diagram commutes for each homology dimension p :*

$$\begin{array}{ccc} H_p(X) & \longrightarrow & H_p(X') \\ \uparrow & & \uparrow \\ H_p(N(\mathfrak{U})) & \longrightarrow & H_p(N(\mathfrak{U}')) \end{array}$$

What this lemma entails is that the homotopy equivalences $C_i(P) \rightarrow P_i$ at all indices $i \in \mathbb{R}$, put together, induce an isomorphism between persistence modules $H_p(\mathcal{C}(P)) \rightarrow H_p(\mathcal{P})$ at each homology dimension p . Hence, up to isomorphism, the filtrations $\mathcal{C}(P)$ and \mathcal{P} have the same persistent homology. The guarantees on offsets filtrations obtained in Section 2 extend therefore to Čech filtrations:

COROLLARY 4.13. *Let K be a compact set in \mathbb{R}^d with positive weak feature size. Let P be a point cloud in \mathbb{R}^d such that $d_H(K, P) = \varepsilon < \frac{1}{4} \text{wfs}(K)$. Then, there is a sweet range $T = (\varepsilon, \text{wfs}(K) - \varepsilon)$ whose intersection with the barcode of the Čech filtration $\mathcal{C}(P)$ has the same properties as in Corollary 4.8. If in addition K has positive μ -reach for some $\mu \in (0, 1]$, and $\varepsilon < \frac{\mu^2}{5\mu^2 + 12} r_\mu(K)$, then there is a sweeter range $\left[\frac{4\varepsilon}{\mu^2}, r_\mu(K) - 3\varepsilon \right] \subseteq T$ whose intersection with the barcode of $\mathcal{C}(P)$ has the same properties as in Corollary 4.9.*

⁷Bendich et al. [22] gave a more direct proof of this result in the special case of covers by closed Euclidean balls.

Another notorious combinatorial construction is the so-called α -*complex*, introduced by Edelsbrunner, Kirkpatrick, and Seidel [113], which we will call *i -Delaunay complex* in the following to better emphasize its relationship with the Delaunay triangulation⁸. It is defined as follows.

DEFINITION 4.14. The *i -Delaunay complex* of P , denoted $D_i(P)$, has one k -simplex per $(k + 1)$ -tuple of points of P circumscribed by a ball of radius at most i containing no point of P in its interior. The *Delaunay filtration* is the indexed family $\mathcal{D}(P) = \{D_i(P)\}_{i \in \mathbb{R}}$. It is a filtration of the Delaunay triangulation of P .

As a subcomplex of the Delaunay triangulation of P , the *i -Delaunay complex* $D_i(P)$ embeds linearly into \mathbb{R}^d under the genericity assumption that there are no $d + 2$ cospherical points and no $d + 1$ affinely dependent points in P , which will be assumed implicitly in the following. For simplicity, the image of $D_i(P)$ through the linear embedding into \mathbb{R}^d will also be denoted $D_i(P)$ by a slight abuse of notations. This image is known to be contained in the offset P_i , and the connection between the two is made via a deformation retraction worked out by Edelsbrunner [112]:

LEMMA 4.15. *For any $i \in \mathbb{R}$, the offset P_i deformation retracts onto (the linear image of) $D_i(P)$, so the inclusion map $D_i(P) \hookrightarrow P_i$ is a homotopy equivalence.*

Once again, these homotopy equivalences at all indices $i \in \mathbb{R}$, put together, induce an isomorphism between persistence modules $H_p(\mathcal{D}(P)) \rightarrow H_p(\mathcal{P})$ at each homology dimension p . Hence, up to isomorphism, the filtrations $\mathcal{C}(P)$, \mathcal{P} , and $\mathcal{D}(P)$ have the same persistent homology, so Corollary 4.13 holds the same with the Čech filtration $\mathcal{C}(P)$ replaced by the Delaunay filtration $\mathcal{D}(P)$.

Delaunay filtrations are quite popular in small dimensions ($d = 2$ or 3), where they can be computed efficiently. Indeed, computing the Delaunay triangulation of n points takes $O(n \log n)$ time in the worst case in the plane, and $O(n^2)$ in 3-space. Moreover, in many practical cases, such as when the points lie on a smooth or polygonal surface, the size and computation time of the 3d Delaunay triangulation become near linear [4, 5]. As a matter of fact, the result of Figures 4.9 and 4.10 was computed using the Delaunay filtration.

Unfortunately, the sizes of both the Čech and Delaunay filtrations scale up very badly with the ambient dimension, so these filtrations become quickly intractable in practice. Finding more lightweight filtrations or zigzags that can be used as replacements with theoretical guarantees is the topic of the next chapter.

⁸In this we are following some authors such as Bauer and Edelsbrunner [16].

Topological Inference 2.0

In Chapter 4 we focused mostly on the mathematical foundations of a working pipeline for doing topological inference on a computer. Our aim was to get proof for the existence of simplicial filtrations whose barcodes can reveal the homology of the geometric structures underlying the input data, regardless of the actual computing cost of building these filtrations.

This was also the dominant viewpoint in the early days of the development of the theory. However, the rapid growth in size and complexity of the data sets considered in practice quickly made it clear that algorithmic questions needed to be addressed as well by the theory, or otherwise theory and practice would diverge. More precisely, there was a need for more lightweight filtrations or zigzags and for optimized variants of the persistence algorithm, without too much sacrifice in terms of quality of the output barcodes. The goal was to reduce the overall memory footprint of the approach, by then—and still by now—its main bottleneck. The resulting extension of the theory is referred to as ‘version 2.0’ of the inference pipeline in these pages.

Let us give a concrete example illustrating how important the memory footprint can be. This is only a toy example, but it is quite revealing of the situation in practice. Consider the following variant of the data set of Figure 4.2, called the *Clifford data set* hereafter: 2,000 points evenly spaced along the line $l : y = 20x \bmod 2\pi$ in the 2-d flat torus $(\mathbb{R} \bmod 2\pi)^2$, then mapped onto the Clifford torus in \mathbb{R}^4 via the embedding $f : (u, v) \mapsto (\cos u, \sin u, \cos v, \sin v)$. This data set admits three non-trivial underlying structures: at small scales, the image of l through f , which is a closed helicoidal curve on the torus; at larger scales, the torus itself; at even larger scales, the 3-sphere of radius $\sqrt{2}$ on which the torus is sitting. One can also see the point cloud itself and \mathbb{R}^4 as possible underlying structures at extreme scales. In order to capture the homology of the 3-sphere, the union of Euclidean balls of radius i about the data points must at least cover it, which happens only for $i \geq \sqrt{4 - 2\sqrt{2}} \approx 1.08$. At such parameter values, the Čech filtration has already become huge, that is, the 4-skeleton of the Čech complex of parameter i contains more than 31 billion simplices, a size that is at least 2 orders of magnitude beyond what existing implementations of the persistence algorithm can handle. On a 24-GB machine, one can store the 4-skeleton of the Čech filtration and compute its persistent homology within the main memory up to $i \approx 0.625$ using the C++ library DIONYSUS (<http://www.mrzv.org/software/dionysus/>). The corresponding truncated barcode is given in Figure 5.1. As expected, it shows only the curve and the torus, not the 3-sphere.

Not only is the size of the filtrations a critical issue in its own right, but it also has a direct impact on the running time of the persistence algorithm as we saw in Chapter 2. Thus, the efficiency of the whole inference pipeline is driven by

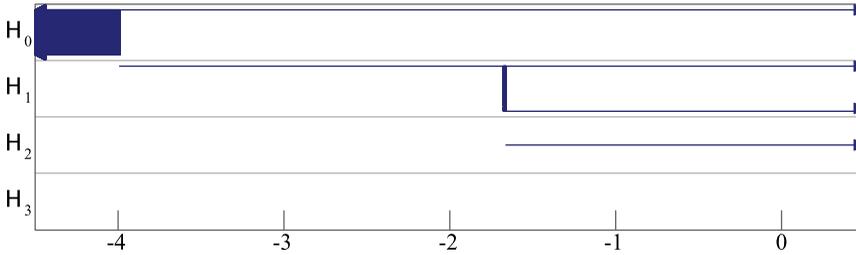


FIGURE 5.1. Truncated barcode of the Čech filtration on the Clifford data set, represented on a logarithmic scale. Due to the truncation at $\log_2(i) = \log_2(0.625) \approx -0.678$, the intervals encoding the homology of the torus become right-infinite, while the interval encoding the 3-homology of the sphere does not appear.

— Reprinted from Oudot and Sheehy [208]: “Zigzag Zoology: Rips Zigzags for Homology Inference”, Foundations of Computational Mathematics, pp. 1–36, ©2014, with kind permission from Springer Science and Business Media.

it. Another important though less critical issue in practice is the complexity of the geometric predicates involved in the construction of the filtrations.

Take for instance the Delaunay filtration. The size of the Delaunay triangulation of an n -points set P in \mathbb{R}^d is known to be $\Theta(n^{\lceil \frac{d}{2} \rceil})$ [189], and the worst-case lower bound is achieved when the points lie on a curve, so even assuming that P samples some low-dimensional structure does not help. Meanwhile, the construction of the Delaunay filtration requires determining whether the circumscribing ball of $d + 1$ affinely independent points includes another given point in its interior. It is answered by evaluating the sign of the determinant of a $(d + 2) \times (d + 2)$ matrix. Assuming fixed-precision entries, exact evaluation with floating-point arithmetic uses an overhead budget of $O(d)$ bits and incurs a cubic (in d) bit complexity [37, 84]. This becomes costly when d increases, especially as the predicate needs to be answered many times.

The Čech complex incurs a similar complexity blowup: as parameter i grows towards infinity, $C_i(P)$ eventually becomes the $(n - 1)$ -simplex of vertex set P , therefore its size grows up to 2^n . When the ambient dimension d is known, it is possible to maintain only the d -skeleton of $C_i(P)$, yet the size still grows like $\Theta(n^d)$ as n tends to infinity while d remains fixed. Meanwhile, the construction of the Čech complex requires determining whether a given collection of balls of same radius r have a nonempty common intersection, which is equivalent to comparing r with the radius of the smallest Euclidean ball enclosing their centers. This geometric predicate is answered exactly in time linear in the number of balls but superpolynomial in d [73, 137, 190, 239]. Even though reasonable timings have been reported for a single instance of the predicate in practice [129, 138], repeated calls to it during the construction of a Čech complex lead to a substantial overhead in the running time.

The impact in practice is that the Čech and Delaunay filtrations become quickly intractable to maintain when d increases beyond 3 or 4, as we saw in the example of Figure 5.1. This is why finding tractable alternatives has been identified as an essential question for applications. In order to make precise complexity claims and fair comparisons, we need to state our data model and objectives formally.

Data model and objectives. Clearly, capturing the topology of arbitrarily large-dimensional structures is out of reach due to the size of the required sampling, which grows exponentially with the dimension of the structures. Therefore, the usual assumption is that the space underlying the input data has small intrinsic dimension, even though the ambient dimension itself may be large. This is a common assumption across the data analysis literature, and to formalize it in our context we will use the concept of *doubling dimension* from metric geometry:

DEFINITION 5.1. The *doubling dimension* of a metric space (X, d) at scale r is the smallest positive integer m such that any d -ball of radius r in X can be covered by 2^m balls of radius $r/2$. The *doubling dimension* of (X, d) is the supremum of the doubling dimensions over all scales $r > 0$.

In the following, the quantity of reference will be the doubling dimension m of the space underlying the input point cloud. In \mathbb{R}^d it is bounded above by $O(d)$, but as we said it will be assumed much smaller in our data model:

Data model: *The input is an n -points set P in Euclidean space \mathbb{R}^d , located ε -close in the Hausdorff distance to some unknown compact set K with positive weak feature size and small (constant) doubling dimension m .*

Under this data model, and in light of the aforementioned challenges, our task is to design filtrations or zigzags that fulfill the following objectives:

- O1** their size (i.e. total number of simplex insertions or deletions) scales up only linearly with n and exponentially with m , typically like $2^{O(m^2)}n$, to reduce the combinatorial complexity of the inference pipeline,
- O2** their construction involves only simple predicates, typically distance comparisons, to reduce the bit complexity of the pipeline,
- O3** their barcode has a sweet range revealing the homology of K among a limited amount of noise.

Note that the size bound in objective O1 is independent of the ambient dimension d , ignoring the space $\Theta(dn)$ needed to store the coordinates of the input points. As we will see, the size bound comes from standard packing arguments involving the doubling dimension of K . Note also the lack of a precise definition of the sweet range in objective O3. Generally speaking, the reader can expect the same properties as in Corollary 4.8, however with some variability in the bounds and amount of noise in the sweet range from one filtration or zigzag to the next.

Strategy at a glance. The main idea is to use approximation. Suppose for instance that when building $\mathcal{C}(P)$ we give up on computing the radii of minimum enclosing balls exactly, but rather we compute them within a multiplicative error $c > 1$, for which fully polynomial-time (both in n and in d) algorithms exist [2, 173]. Then, instead of the exact Čech complex filtration $\mathcal{C}(P)$, we will be building some approximation $\tilde{\mathcal{C}}(P)$ that is interleaved multiplicatively with it:

$$(5.1) \quad \forall i > 0, C_i(P) \subseteq \tilde{C}_{ci}(P) \text{ and } \tilde{C}_i(P) \subseteq C_{ci}(P).$$

This multiplicative interleaving turns into an additive interleaving in the sense of (3.3) on a logarithmic scale. More precisely, reparametrizing the real line by $i \mapsto 2^i$ and calling respectively $\mathcal{C}^{\log}(P)$ and $\tilde{\mathcal{C}}^{\log}(P)$ the resulting filtrations $(C_i^{\log}(P) =$

$C_{2^i}(P)$ and $\tilde{C}_i^{\log}(P) = \tilde{C}_{2^i}(P)$), we have

$$(5.2) \quad \forall i \in \mathbb{R}, C_i^{\log}(P) \subseteq \tilde{C}_{i+\log_2(c)}^{\log}(P) \text{ and } \tilde{C}_i^{\log}(P) \subseteq C_{i+\log_2(c)}^{\log}(P).$$

In other words, $C^{\log}(P)$ and $\tilde{C}^{\log}(P)$ are $\log_2(c)$ -interleaved in the sense of (3.3), therefore their persistence diagrams are $\log_2(c)$ -close in the bottleneck distance according to the Isometry Theorem 3.1. When c is close enough to 1, this approximation result can be combined with Corollary 4.13 to guarantee the existence of sweet ranges in the persistence barcode of the filtration $\tilde{C}^{\log}(P)$. The same approach applies to the Delaunay filtration $\mathcal{D}(P)$ as well. Note the barcode of $\tilde{C}^{\log}(P)$ is nothing but the persistence barcode of $\tilde{C}(P)$ represented on a logarithmic scale—called the *logscale persistence barcode* of $\tilde{C}(P)$ hereafter.

Simplifying the predicates. This question is addressed by replacing the geometric predicates involved in the construction of the Čech and Delaunay filtrations by approximate versions based solely on distance comparisons. This gives rise to two new combinatorial objects, the so-called (*victoris-*)*Rips filtration* and *witness filtration*, which turn out to be interleaved multiplicatively with the Čech filtration as in (5.1), as we will see in Section 1.

Reducing the size. This question is addressed by introducing new filtrations or zigzags that are interleaved with the Čech or Delaunay filtrations, not at the topological level as in (5.1), but at the algebraic level directly. The interleavings can take various forms, sometimes pretty different from the standard one from Definition 3.3, so the Isometry Theorem 3.1 does not always help in their analysis. It is then necessary to develop new theoretical tools for comparing filtrations, or zigzags, or both. Historically, the size was reduced in two steps: first, to linear in the number n of input points, with a factor depending exponentially on the ambient dimension d (Section 2); second, to linear in n with a factor scaling up (exponentially) with the intrinsic dimension m of the data (Section 3).

Signal-to-noise ratio. An unexpected byproduct of these developments has been to improve the signal-to-noise ratio of the barcodes inside their sweet ranges. This is an important contribution as the interpretation of the barcode is still currently left to the user, so the larger the signal-to-noise ratio, the easier the interpretation. This aspect is addressed in Section 3.2.

Conclusions. To conclude Chapters 4 and 5 altogether, we provide in Section 4 a summary of the qualities and drawbacks of the various filtrations and zigzags introduced, with respect to the above three criteria: predicates, size, signal-to-noise ratio. We also report on the behavior of the inference pipeline on real-life data in Section 5. Finally, we address the mostly unexplored question of dealing with the presence of outliers in the input data in Section 6.

Prerequisites. The prerequisites for this chapter are the same as for Chapter 4, except we will also assume familiarity with the design and analysis of algorithms. To the unfamiliar reader we recommend reading [92].

1. Simple geometric predicates

In practice it is desirable to use as simple predicates as possible. Ideally, one aims at reducing them to mere distance comparisons, which can be evaluated in $O(d)$ time in \mathbb{R}^d and performed in arbitrary metric spaces. This is possible using variants of the Čech filtration such as the following one (illustrated in Figure 5.2):

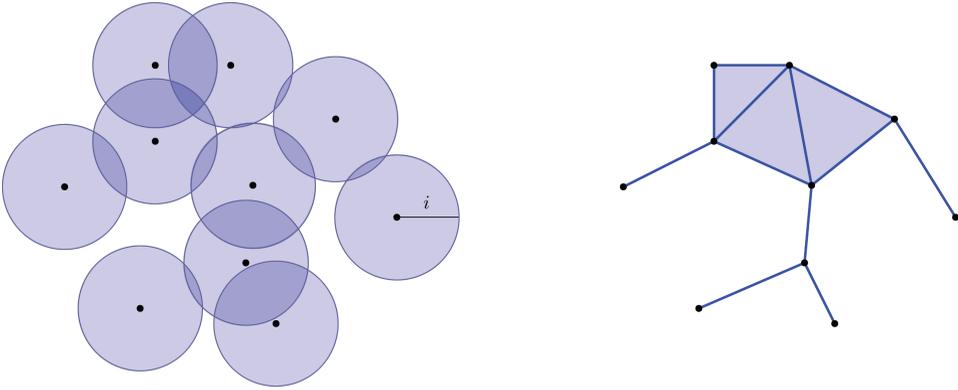


FIGURE 5.2. A point cloud P (left) and its Rips complex $R_{2i}(P)$ (right), which is the same as the clique complex of the 1-skeleton graph of $C_i(P)$ —recall Figure 4.11.

DEFINITION 5.2. The (*Vietoris-*)*Rips complex*¹ of P of parameter $i \in \mathbb{R}$, denoted $R_i(P)$, has one k -simplex per $(k + 1)$ -tuple of points of P whose Euclidean diameter (maximum pairwise Euclidean distance) is at most i . The (*Vietoris-*)*Rips filtration* is the indexed family $\mathcal{R}(P) = \{R_i(P)\}_{i \in \mathbb{R}}$.

An equivalent way to define $R_i(P)$ is as the clique complex (or flag complex) of the graph having P as vertex set and an edge between any pair of points $p, q \in P$ such that $\|p - q\| \leq i$. The Rips complex has then one k -simplex per $(k + 1)$ -clique of the graph, and it is easily seen that the only geometric predicates involved in its construction are distance comparisons.

In general metric spaces, Čech and Rips filtrations are interleaved multiplicatively as follows:

$$(5.3) \quad \forall i > 0, C_i(P) \subseteq R_{2i}(P) \text{ and } R_i(P) \subseteq C_i(P).$$

Observe that the interleaving is not symmetric, so rescaling the Rips parameter by $\sqrt{2}$ gives a new filtration $\tilde{\mathcal{R}}(P)$ that is $\sqrt{2}$ -interleaved multiplicatively with $\mathcal{C}(P)$ in the sense of (5.1).

In Euclidean space \mathbb{R}^d , de Silva and Ghrist [103] worked out tight interleaving bounds, where $\vartheta_d = \sqrt{\frac{d}{2(d+1)}} \in \left[\frac{1}{2}, \frac{1}{\sqrt{2}}\right)$:

$$(5.4) \quad \forall i > 0, C_i(P) \subseteq R_{2i}(P) \text{ and } R_i(P) \subseteq C_{\vartheta_d i}(P).$$

Thus, rescaling the Rips parameter by $\sqrt{\frac{2}{\vartheta_d}}$ gives a new filtration $\tilde{\mathcal{R}}(P)$ that is $\sqrt{2\vartheta_d}$ -interleaved multiplicatively with $\mathcal{C}(P)$ in the sense of (5.1).

Combined with Corollary 4.13, these interleavings imply the existence of a sweet range in the logscale persistence barcode of the Rips filtration $\mathcal{R}(P)$. The details of the proof rely on mere rank arguments and can be found in [66]—see also [136]. Here we reproduce the simplest instance of the result, based on (5.3), which has the least tight constants:

¹Originally introduced by L. Vietoris to define homology groups of the Čech type for compact metric spaces. Later reintroduced by E. Rips for the study of hyperbolic groups, and popularized by Gromov [143] under the name of *Rips complex*.

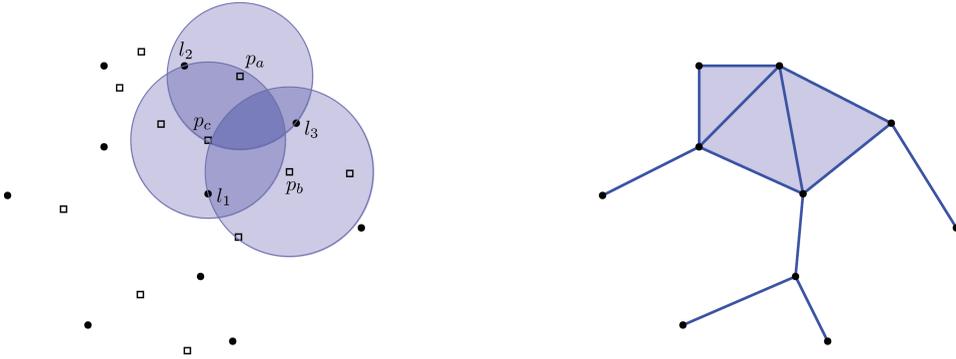


FIGURE 5.3. Left: a set L of landmarks (dots) and a set P of witnesses (squares). Point p_a witnesses vertex l_2 and edge $\{l_2, l_3\}$. Point p_b witnesses $\{l_3\}$ and $\{l_3, l_1\}$. Point p_c witnesses $\{l_1\}$, $\{l_1, l_2\}$ and $\{l_1, l_2, l_3\}$. Therefore, triangle $\{l_1, l_2, l_3\}$ belongs to the witness complex $W_0(L, P)$ shown to the right.

THEOREM 5.3 (Chazal and Oudot [66]). *Let K be a compact set in \mathbb{R}^d with positive weak feature size. Let P be a point cloud in \mathbb{R}^d such that $d_H(K, P) = \varepsilon < \frac{1}{9}\text{wfs}(K)$. Then, there is a sweet range $T = (\log_2(2\varepsilon), \log_2(\text{wfs}(K) - \varepsilon))$ whose intersection with the logscale barcode of the Rips filtration $\mathcal{R}(P)$ has the following properties:*

- *The intervals that span T encode the homology of K , in the sense that their number for each homology dimension p is equal to the dimension of $H_p(K_r)$, for any $r \in (0, \text{wfs}(K))$.*
- *The remaining intervals have length at most 2.*

Notice that the amount of noise in the sweet range is constant and does not go to zero as the sampling density increases. This is a result of the Rips filtration being interleaved with the Čech filtration only by a constant multiplicative factor. In subsequent sections we will see other filtrations that are $(1 + \varepsilon)$ -interleaved multiplicatively with the Čech filtration.

Notice also that the interleaving with the Čech filtration only guarantees the existence of a sweet range in the logscale barcode of the Rips filtration, and not the existence of a sweeter range over which there would be no noise at all.

Another approximation of interest is the so-called *witness filtration*. The idea, illustrated in Figure 5.3, is to use a loose version of the Delaunay predicate, based only on distance comparisons. For this we select landmark points among the point cloud P and build a simplicial complex on top of the landmarks set L using the rest of the points to answer the approximate Delaunay predicate and thus drive the construction. Specifically, a point $p \in P$ is an *i -witness* for a simplex $\sigma \subseteq L$ if

$$(5.5) \quad \|p - l\| \leq \|p - l'\| + i \text{ for all } l \in \sigma \text{ and all } l' \in L \setminus \sigma.$$

Note that the classical Delaunay predicate corresponds to letting $i = 0$ and l' range over the entire landmarks set L instead of $L \setminus \sigma$ in (5.5). In that case, p is called a *strong witness* of σ .

DEFINITION 5.4. Given $i \in \mathbb{R}$, the i -witness complex of the pair (L, P) , denoted $W_i(L, P)$, is the maximal simplicial complex of vertex set L whose simplices are i -witnessed by points of P . The 0-witness complex is called simply the *witness complex* of (L, P) . The *witness filtration* is the indexed family $\mathcal{W}(L, P) = \{W_i(L, P)\}_{i \in \mathbb{R}}$.

The witness complex was introduced by de Silva [101], who proved that it is a subcomplex of the Delaunay triangulation—see [7] for an alternate proof. His idea was to use it as a lightweight alternative to the Delaunay triangulation, and a natural question is whether the witness filtration can be interleaved with the Delaunay filtration in the sense of (3.3). This is not possible unfortunately, as the inclusion of the i -witness complex in the Delaunay triangulation holds only for $i = 0$. Nevertheless, an interleaving with the Čech filtration can be worked out under some sampling conditions:

LEMMA 5.5 (Chazal and Oudot [66]). *Let K be a connected compact set in \mathbb{R}^d , and let $L \subseteq P \subset \mathbb{R}^d$ be finite sets such that*

$$d_H(K, P) \leq d_H(P, L) = \varepsilon < \frac{1}{8} \operatorname{diam}(K),$$

where $\operatorname{diam}(K) = \max\{\|x - y\| \mid x, y \in K\}$ is the Euclidean diameter of K . Then,

$$(5.6) \quad \forall i \geq 2\varepsilon, C_{\frac{i}{4}}(L) \subseteq W_i(L, P) \subseteq C_{8i}(L).$$

The same kind of sweet range as in Theorem 5.3 can then be derived for the logscale barcode of the witness filtration $\mathcal{W}(L, P)$, provided that a relevant choice of landmarks is made:

THEOREM 5.6 (Chazal and Oudot [66]). *Let K be a compact set in \mathbb{R}^d with positive weak feature size. Let P be a point cloud in \mathbb{R}^d , and $L \subseteq P$ a subset of landmarks, such that*

$$(5.7) \quad d_H(K, P) \leq d_H(P, L) = \varepsilon < \min \left\{ \frac{1}{8} \operatorname{diam}(K), \frac{1}{2^{11} + 1} \operatorname{wfs}(K) \right\}.$$

Then, there is a sweet range $T = \left(\log_2(4\varepsilon), \log_2\left(\frac{\operatorname{wfs}(K) - \varepsilon}{8}\right) \right)$ whose intersection with the logscale barcode of the witness filtration $\mathcal{W}(L, P)$ has the following properties:

- The intervals that span T encode the homology of K , in the sense that their number for each homology dimension p is equal to the dimension of $H_p(K_r)$, for any $r \in (0, \operatorname{wfs}(K))$.
- The remaining intervals have length at most 6.

Note that the obtained bounds on the sweet range are certainly not optimal, given the crudeness of the interleaving with the Čech filtration in (5.6). This is fortunate, as the corresponding sampling condition (5.7) requires an immense sampling to be satisfied. In practice, the witness filtration has been observed to produce consistently cleaner barcodes than the Rips filtration. In particular, de Silva and Carlsson [102] conjectured that the topological signal can appear arbitrarily sooner in the barcode, which they presented as a major motivation for using the witness filtration over the Čech filtration. Chazal and Oudot [66] proved this conjecture correct at least in the special case where the compact set K underlying the input point cloud P is a submanifold of \mathbb{R}^d with positive reach:

THEOREM 5.7. *There exist a constant $c > 0$ and a continuous, non-decreasing map $\omega : [0, c) \rightarrow [0, \frac{1}{2})$, with $\omega(0) = 0$, such that, for any compact submanifold K of \mathbb{R}^d with positive reach, for any finite sets $L \subseteq P \subset \mathbb{R}^d$ such that*

$$(5.8) \quad d_H(K, P) \leq d_H(P, L) = \varepsilon < c \operatorname{rch}(K)$$

and L is ε -sparse in the sense that

$$(5.9) \quad \forall l \neq l' \in L, \|l - l'\| \geq \varepsilon,$$

the left bound of the sweet range T from Theorem 5.6 becomes $\log_2(\omega(\frac{\varepsilon}{\operatorname{rch}(K)})^2\varepsilon) + O(1)$.

The properties of ω imply that the quantity $\omega(\frac{\varepsilon}{\operatorname{rch}(K)})^2\varepsilon$ is in $o(\varepsilon)$ as ε goes to zero, so the left bound of the sweet range T is arbitrarily smaller than the one from Theorem 5.6. What this entails is that, given a fixed landmarks set L , the topological signal appears sooner in the barcode of the witness filtration as the superset P of witnesses increases. Conversely, given a fixed set of witnesses P , a careful choice of a subset L of landmarks guarantees the presence of a sweet range with the same lower bound in the barcode at a reduced algorithmic cost—recall that the total size of the witness filtration depends only on $|L|$, not on $|P|$.

The proof of Theorem 5.7 exploits the relationship that exists between the witness complex and the so-called *restricted Delaunay triangulation*, another subcomplex of the Delaunay triangulation that plays a key role in the context of manifold reconstruction [28, 77].

2. Linear size

While the use of Rips and witness filtrations greatly simplifies the geometric predicates, it does not help in terms of size compared to the Čech or Delaunay filtrations. Indeed, as the filtration parameter grows towards infinity, the size of the complexes on an n -points set eventually becomes 2^n . Even knowing the ambient dimension d and computing the d -skeleton of each complex still incurs a $\Theta(n^d)$ space complexity. In practice this means that the filtrations cannot be built and stored in their entirety.

Hudson et al. [161] initiated a new line of work aimed specifically at reducing the filtration size. Given an n -points set P in \mathbb{R}^d , their approach considers the Delaunay filtration $\mathcal{D}(P)$ and modifies its vertex set P so as to avoid the worst-case $\Omega(n^{\lceil \frac{d}{2} \rceil})$ size bound. Specifically, they preprocess P using techniques inspired by Delaunay refinement, iteratively inserting new points of \mathbb{R}^d called *Steiner points* into P until the aspect ratios² of the Voronoi cells of the augmented set $P \cup S$ are good enough. This shape quality criterion on the Voronoi cells guarantees that the size of the Delaunay triangulation of $P \cup S$ drops down to $2^{O(d^2)}(n + |S|)$ when the criterion is met. Furthermore, the number of Steiner points needed to meet the criterion is $2^{O(d)}n$, which makes the size of the final triangulation only $2^{O(d^2)}n$. Once this preprocessing is done, they build a filtration of the Delaunay triangulation $D(P \cup S)$, called the *mesh filtration* of P and denoted $\mathcal{M}(P)$, which departs from the classical Delaunay filtration and captures the topology of the offsets of P rather than that of the offsets of $P \cup S$. It goes without saying that in order to realize the benefits of the refinement in terms of size, the Steiner set S is computed without

²By aspect ratio of the Voronoi cell of a point p is meant the ratio between the radii of the smallest enclosing ball and largest inscribed ball centered at p .

first constructing the Delaunay triangulation of P , as is possible using the Sparse Voronoi Refinement (SVR) algorithm of Hudson, Miller, and Phillips [159].

Details. The SVR algorithm takes the n -points set P as input and computes a Steiner set S such that $P \cup S$ satisfies the following properties:

- (i) $P \cup S$ is a point sampling of some axis-aligned bounding box B of side length $O(\text{diam}(P))$ around the input point cloud P ,
- (ii) the Voronoi cells of the points of $P \cup S$, clipped to B , have aspect ratios bounded from above by an absolute constant $\rho \geq 2$,
- (iii) the subcomplex of $D(P \cup S)$ dual to the clipped Voronoi diagram is equal to the full Delaunay triangulation $D(P \cup S)$,
- (iv) the size of $P \cup S$ is $2^{O(d)} n \log_2 \Delta(P)$, where $\Delta(P)$ denotes the *spread* of P , i.e. the ratio of the largest to smallest interpoint distances among the points of P ,
- (v) the size of $D(P \cup S)$ is $2^{O(d^2)} |P \cup S|$.

As shown by Hudson et al. [160], the extra work needed to fill in the entire bounding box B with point samples is negligible. The SVR algorithm can produce $P \cup S$ together with its Delaunay triangulation in near-optimal $|D(P \cup S)| + 2^{O(d)} n \log_2 \Delta(P)$ time, where the second term arises from point location costs [159]. This bound is dominated by $2^{O(d^2)} n \log_2 \Delta(P)$.

Once both $P \cup S$ and its Delaunay triangulation have been built, we can define a filter $f : D(P \cup S) \rightarrow \mathbb{R}$ as follows, where we slightly abuse notations by identifying each simplex $\sigma \in D(P \cup S)$ with its vertex set:

$$(5.10) \quad \forall \sigma \in D(P \cup S), f(\sigma) = \max\{d_P(q) \mid q \in \sigma\}.$$

Note that if τ is a face of σ then $f(\tau) \leq f(\sigma)$, so the filter is compatible with the incidence relations in $D(P \cup S)$. The mesh filtration $\mathcal{M}(P)$ is then defined formally as the sublevel-sets filtration of f , that is: for every value $i \in \mathbb{R}$, we let F_i be the subcomplex of $D(P \cup S)$ made of the simplices σ such that $d_P(q) \leq i$ for all $q \in \sigma$.

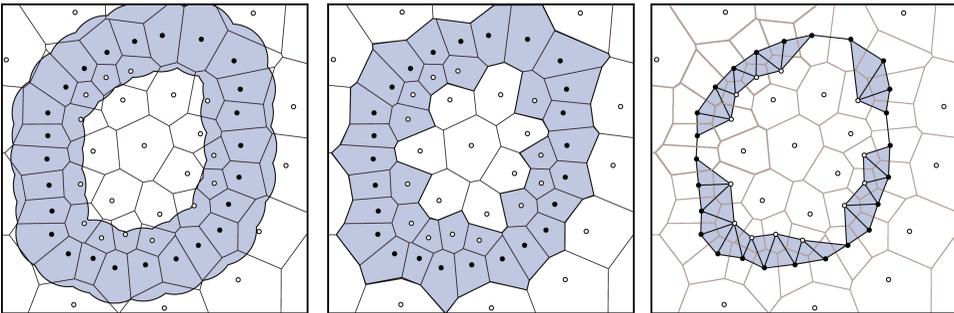


FIGURE 5.4. Relationship between the offsets and mesh filtrations. Black points belong to P while white points belong to S . From left to right: the offset P_i superimposed with the Voronoi diagram of $P \cup S$; the Voronoi cells whose centers lie in P_i ; the corresponding dual subcomplex of $D(P \cup S)$, which is part of the mesh filtration.

— From Hudson et al. [161].

Intuitively, the filter of (5.10) sorts the simplices of $D(P \cup S)$ according to their distances to the input point cloud P , in order to simulate within $D(P \cup S)$ the growth of the offsets of P —see Figure 5.4 for an illustration. The fact that the clipped Voronoi cells have bounded aspect ratios ensures that this simulation process works, i.e. that the mesh filtration (or rather its dual filtration of the Voronoi diagram of $P \cup S$) and the offsets filtration are ρ -interleaved multiplicatively³. As a consequence,

THEOREM 5.8 (Hudson et al. [161]). *The logscale persistence diagram of the mesh filtration $\mathcal{M}(P)$ is $\log_2(\rho)$ -close to the one of the offsets filtration \mathcal{P} in the bottleneck distance. Meanwhile, the size of the mesh filtration is $2^{O(d^2)}|P|\log_2 \Delta(P)$.*

Then, assuming the input point cloud P lies ε -close (in the Hausdorff distance) to some compact set K with positive weak feature size, the same arguments as in Section 1 prove the existence of a sweet range of type

$$T = (\log_2(\varepsilon) + O(1), \log_2(\text{wfs}(K)) - O(1))$$

in the logscale barcode of $\mathcal{M}(P)$, where the constants hidden in the big-O notations depend on the aspect ratio bound ρ .

The problem with Theorem 5.8 is that the guaranteed size bound depends on the spread of P , an unbounded quantity, therefore it does not comply with objective O1 stated in the introduction of the chapter. In fact, Miller, Phillips, and Sheehy [194] showed that it is possible to reduce the size of the superset $P \cup S$ to $2^{O(d)}|P|$ by applying the SVR algorithm recursively to well-chosen subsets of P called *well-paced sets*. The output is different from the one of the simple SVR algorithm, however our filter f can be adapted so as to maintain the multiplicative interleaving between the mesh and offsets filtrations of P . Let us merely state the end result here and refer the interested reader to [161] for the details.

THEOREM 5.9 (Hudson et al. [161]). *The logscale persistence diagram of the modified mesh filtration is $\log_2 \rho$ -close to the one of the offsets filtration \mathcal{P} in the bottleneck distance. Meanwhile, the size of the modified mesh filtration is $2^{O(d^2)}|P|$.*

REMARK. One can even extend the SVR algorithm and adapt the filter so the induced mesh filtration is $(1 + \delta)$ -interleaved multiplicatively with \mathcal{P} , for any given parameter $\delta > 0$. The corresponding Steiner set S has size $(\frac{1}{\delta})^{O(d)}|P|$, while the induced mesh filtration has size $(\frac{1}{\delta})^{O(d^2)}|P|$. The trade-off is then between the size of the data structure and the tightness of the interleaving factor between the offsets and mesh filtrations. Typically, one can take $\delta = \varepsilon$ so the interleaving factor converges to 1 as the sampling density goes to infinity.

Theorem 5.9 is illustrated in Figure 5.5, which shows the output of the method on the Clifford data set described in introduction. As expected from the theory, the barcode contains sweet ranges not only for the curve and the torus, but also for the unit 3-sphere. To the 2,000 input points, the SVR algorithm added approximately 71,000 Steiner points to achieve an aspect ratio $\rho = 3.08$ (a value chosen for technical reasons). The mesh filtration contained a total of 12 million simplices

³There is a catch for small scales because the Delaunay simplices connecting points of P appear already at time $i = 0$ in the mesh filtration. Hudson et al. [161] proposed a slight modification to the filter of (5.10) that allows to get an interleaving over the entire real line. For simplicity we are overlooking this technicality here.

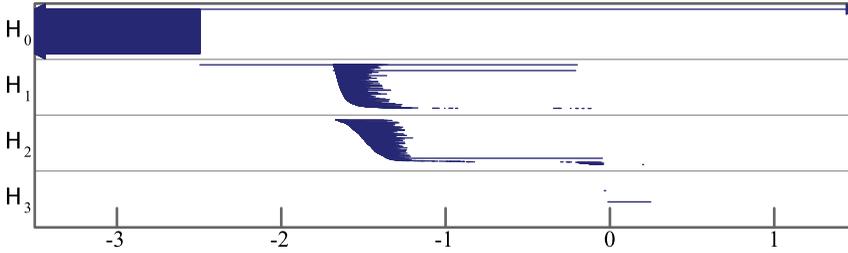


FIGURE 5.5. Logscale barcode of the mesh filtration on the Clifford data set.

— Reprinted from Oudot and Sheehy [208]: “Zigzag Zoology: Rips Zigzags for Homology Inference”, *Foundations of Computational Mathematics*, pp. 1–36, ©2014, with kind permission from Springer Science and Business Media.

and took approximately 1 hour to compute entirely. This is to be compared with the 31 billion simplices required by the Čech filtration to detect the 3-sphere, as mentioned in introduction.

3. Scaling up with the intrinsic dimensionality of the data

The size bound provided by Theorem 5.9 is oblivious to the intrinsic dimensionality m of the input data. Indeed, the size of the Steiner set is exponential in the ambient dimension d regardless of m . As a result, the method from Section 2 becomes quickly intractable as d grows, even though the data points may still live close to some low-dimensional structure.

3.1. Iterative subsampling. Chazal and Oudot [66] proposed a different approach to sparsifying the filtrations, inspired from previous work by Guibas and Oudot [147] on multiscale reconstruction, and based on the following simple idea. Since the sizes of classical filtrations blow up at large scales, just decimate the vertex set as the scale increases. Thus, while the filtration parameter grows, the vertex set shrinks, and ideally some trade-off may be found between keeping a controlled complex size and still capturing the homology of the various structures underlying the input data.

Formally, given an n -points set P in \mathbb{R}^d and a total order (p_1, \dots, p_n) on the points of P , we consider the collection of prefixes $\{p_1, \dots, p_i\}$ together with their associated ‘scales’ $\varepsilon_i = d_H(\{p_1, \dots, p_i\}, P)$. Since the prefixes grow, we have

$$\varepsilon_1 \geq \varepsilon_2 \geq \dots \geq \varepsilon_n = 0.$$

Now, given two parameters $\rho > \eta > 0$, called *multipliers*, for each prefix $\{p_1, \dots, p_i\}$ we build two Rips complexes, one of parameter $\eta\varepsilon_i$, the other of parameter $\rho\varepsilon_i$, then we compute the image of the homomorphism h_i induced by the inclusion map $R_{\eta\varepsilon_i}(\{p_1, \dots, p_i\}) \hookrightarrow R_{\rho\varepsilon_i}(\{p_1, \dots, p_i\})$ at the homology level:

$$h_i : H_*(R_{\eta\varepsilon_i}(\{p_1, \dots, p_i\})) \rightarrow H_*(R_{\rho\varepsilon_i}(\{p_1, \dots, p_i\})).$$

The reason why we use a pair of Rips complexes instead of a single Rips complex at each iteration i stems from the fact that a single Rips complex may fail to capture the topology of the structure underlying the data, as was already observed for unions of balls in the counter-example of Figure 4.8.

The output of the method is the sequence of pairs $(\varepsilon_i, \text{im } h_i)$. For visualization purposes, one can replace the images of the homomorphisms h_i by their ranks, so the output becomes the sequence of pairs $(\varepsilon_i, \text{rank } h_i)$, whose coordinates can be plotted against each other in the so-called *scale-rank plot*. An illustration is given in Figure 5.6, which shows the scale-rank plot obtained on the data set of Figure 4.2.

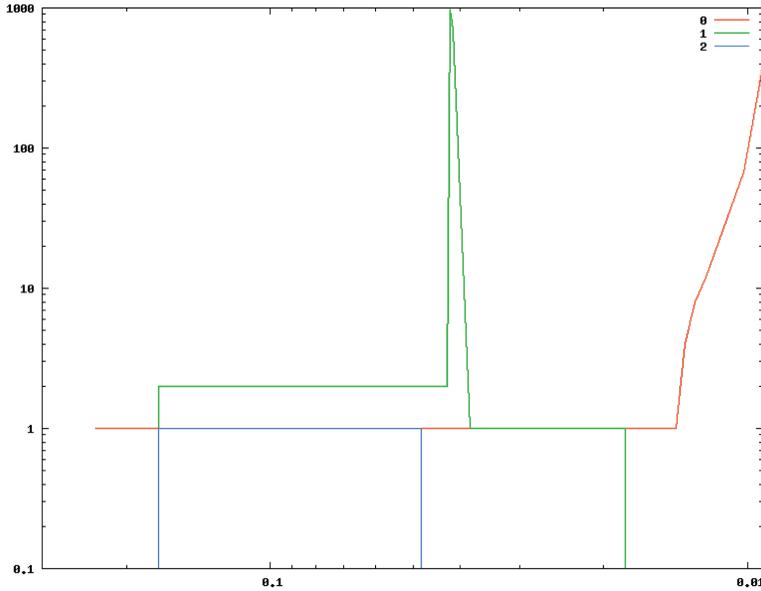


FIGURE 5.6. Scale-rank plot (on a logarithmic scale) obtained from the data set of Figure 4.2. The curves show the rank of h_i for homology dimensions 0, 1 and 2.

— From Chazal and Oudot [66].

The scale-rank plot is readily interpretable: the ranges of scales over which the ranks stabilize play the role of sweet ranges indicating the presence of underlying structures with the corresponding homology. In the example of Figure 5.6, the two main sweet ranges reveal the presence of the curve and of the torus. On the theoretical side, the interpretation of the scale-rank plot is backed up by the following more or less immediate consequence of Theorem 5.3:

COROLLARY 5.10 (Chazal and Oudot [66]). *Let ρ and η be multipliers such that $\rho > 8$ and $2 < \eta \leq \frac{\rho}{4}$. Suppose there is a compact set $K \subset \mathbb{R}^d$ such that $d_H(K, P) = \varepsilon < \frac{\eta-2}{2\rho+\eta} \text{wfs}(K)$. Then, for any iterations $l > k$ such that*

$$\frac{2\varepsilon}{\eta-2} < \varepsilon_l \leq \varepsilon_k < \frac{\text{wfs}(K) - \varepsilon}{\rho+1},$$

the range of iterations $[k, l]$ is a sweet range in the sense that for all $i \in [k, l]$ the rank of h_i is equal to the dimension of $H_(K_r)$ for any $r \in (0, \text{wfs}(K))$.*

In order to analyze the size of the data structure, we need to make an assumption on the order (p_1, \dots, p_n) on the points of P given as input. Indeed, orders that insert nearby points first may lead to local oversampling, which of course is not

desirable. On the contrary, it is better to maintain some kind of uniform subsample of P as long as possible. This is what *furthest-point sampling* provides.

DEFINITION 5.11. Given an n -points set P equipped with a metric, *furthest point sampling* orders the points of P in the following greedy way. It starts with an arbitrary point $p_1 \in P$. Then, at each iteration $i = 2, \dots, n$ it chooses for p_i the point of P lying furthest away from $\{p_1, \dots, p_{i-1}\}$, breaking ties arbitrarily. Upon termination, $P = \{p_1, \dots, p_n\}$ is totally ordered.

It is easily seen that every prefix $\{p_1, \dots, p_i\}$ in this order is ε_i -sparse in the sense of (5.9), i.e.

$$\forall 1 \leq j < k \leq i, \|p_j - p_k\| \geq \varepsilon_i.$$

Then, a standard ball packing argument shows that every vertex is connected to at most $2^{O(m_i)}$ other vertices in the Rips complexes built at iteration i of the algorithm, where m_i denotes the doubling dimension of the finite metric space $(P, \|\cdot\|)$ at scale ε_i . A bound on the number of simplices of the Rips complexes follows trivially:

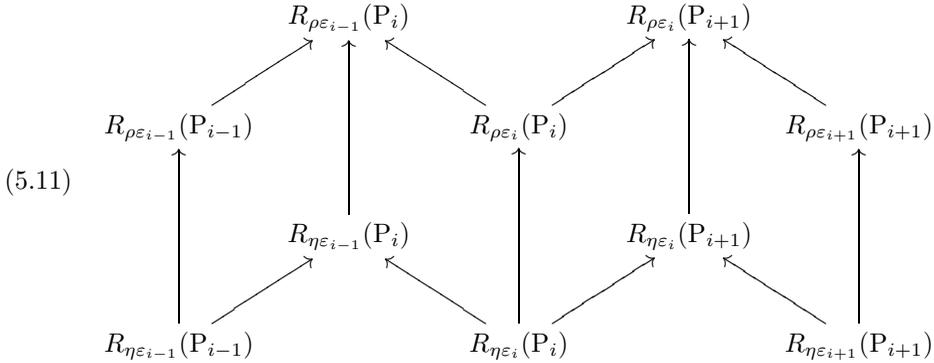
THEOREM 5.12. Assume the order (p_1, \dots, p_n) on the points of P is given by *furthest-point sampling*. Then, at every iteration i of the algorithm the number of k -simplices in the pair of Rips complexes $R_{\eta\varepsilon_i}(\{p_1, \dots, p_i\}) \subseteq R_{\rho\varepsilon_i}(\{p_1, \dots, p_i\})$ is at most $2^{O(km_i)}i$, where m_i denotes the doubling dimension of $(P, \|\cdot\|)$ at scale ε_i .

A few words of explanation are in order here. First of all, the sizes of the Rips complexes remain at most linear in the number n of input points throughout the course of the algorithm. However, this does not mean that only a linear number of simplex insertions or deletions will be needed to maintain the complexes. In the next section we will see how this cost can be amortized. Second, in \mathbb{R}^d the doubling dimension of $(P, \|\cdot\|)$ is bounded by $O(d)$, so if d is known and only the d -skeleton of every Rips complex is built, then the size bound itself is bounded by $2^{O(d^2)}n$, which is as good as the bound from Section 2. Third, if P happens to sample some compact set K of doubling dimension m , then Theorem 5.12 guarantees that the sizes of the $O(m)$ -skeleta of the Rips complexes within the sweet range associated to K will not exceed $2^{O(m^2)}n$. If m is known and small enough, then computing the scale-rank function within the sweet range will be tractable. This does not mean that it will be so outside the range, however it has been observed in practice that the sizes of the Rips complexes tend to increase with the scale, so the part of the scale-rank plot located to the right of the sweet range can be computed entirely. This justifies to run the algorithm backwards, starting with small scales (i.e. large prefixes) and ending with large scales (i.e. small prefixes), stopping when the intrinsic dimensionality of the data at the considered scale becomes too large. Fourth, m is often unknown in applications, so what practitioners commonly do is to fix a maximum simplex dimension a priori and hope that their threshold is large enough. In some cases, reasonable upper bounds on m can be obtained using dimensionality estimation techniques. Fifth, note that the furthest-point sampling takes quadratic time (in n) to compute exactly, but that near linear-time approximations like *net-trees* [149] can be used instead, with roughly the same impact on the space complexity of the approach.

3.2. Connecting the short filtrations together. The iterative subsampling algorithm does not quite follow the same philosophy as the previous methods.

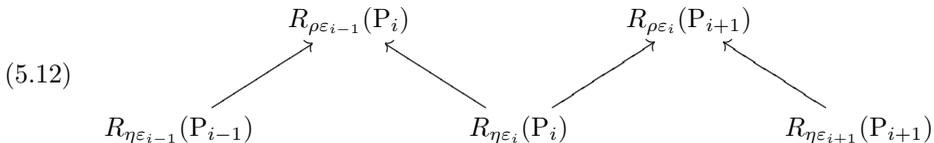
Indeed, instead of building a single filtration or zigzag on top of the input point cloud $P = \{p_1, \dots, p_n\}$, it builds a sequence of short filtrations $R_{\eta\varepsilon_i}(\{p_1, \dots, p_i\}) \subseteq R_{\rho\varepsilon_i}(\{p_1, \dots, p_i\})$ that it processes independently from one another. In particular, homology generators found at consecutive iterations $i, i+1$ of the algorithm are not related, so even if the ranks of the homomorphisms h_i and h_{i+1} are the same, it is unclear whether it is because the topological features at iterations $i, i+1$ are the same or because by chance the same number of features are born and die in-between the two iterations. Such an example can be found in [147].

The challenge is then to ‘connect’ the short filtrations together, in order to obtain a single long zigzag. The most natural way to do so is by observing that when we move from iteration i to $i+1$, one vertex is added to each Rips complex while the scale decreases from ε_i to ε_{i+1} , two operations that can be performed sequentially instead of simultaneously. This gives rise to a big diagram of inclusions relating the short filtrations across all scales, a representative portion of which is depicted below, where P_i is a shorthand for the prefix $\{p_1, \dots, p_i\}$ (not to be confused with the i -offset of P , denoted P_i).



Several zigzags can be extracted from this commutative diagram, including:

- The top row $\dots R_{\rho\varepsilon_{i-1}}(P_{i-1}) \rightarrow R_{\rho\varepsilon_{i-1}}(P_i) \leftarrow R_{\rho\varepsilon_i}(P_i) \rightarrow R_{\rho\varepsilon_i}(P_{i+1}) \leftarrow R_{\rho\varepsilon_{i+1}}(P_{i+1}) \dots$ is called the *Morozov zigzag (M-ZZ)* because it was first proposed by D. Morozov. The bottom row of the diagram is merely the same zigzag with a different multiplier.
- The vertical arrows in the diagram induce a morphism between the bottom and top Morozov zigzags at the homology level. Its image, kernel and cokernel can be considered. Here we will only consider its image, called the *image Rips zigzag (iR-ZZ)* in the following.
- The path $\dots R_{\eta\varepsilon_{i-1}}(P_{i-1}) \rightarrow R_{\rho\varepsilon_{i-1}}(P_{i-1}) \rightarrow R_{\rho\varepsilon_{i-1}}(P_i) \leftarrow R_{\rho\varepsilon_i}(P_i) \leftarrow R_{\eta\varepsilon_i}(P_i) \rightarrow R_{\rho\varepsilon_i}(P_i) \rightarrow R_{\rho\varepsilon_i}(P_{i+1}) \leftarrow R_{\rho\varepsilon_{i+1}}(P_{i+1}) \leftarrow R_{\eta\varepsilon_{i+1}}(P_{i+1}) \dots$ oscillates between the top and bottom rows of the diagram, and by composing adjacent maps with same orientation we obtain the following zigzag called the *oscillating Rips zigzag (oR-ZZ)* hereafter:



Providing theoretical guarantees to these zigzags requires to develop novel tools, since the stability part of the Isometry Theorem 3.1, which was the key ingredient in the analysis of the persistence modules from the previous sections, currently has no counterpart for zigzag modules. The strategy we developed by Oudot and Sheehy [208] is to perform low-level manipulations on the zigzags and their underlying quivers, so as to turn them progressively into other zigzags with a simple algebraic structure. In a way, this is the same approach as the one followed by Bernstein, Gelfand, and Ponomarev [24] to prove Gabriel’s theorem using reflection functors, except here the functors are replaced by more ad-hoc operations. We will now spend some time detailing the manipulations and how they are used to derive guarantees on the barcodes of the Rips zigzags. The unfamiliar reader may safely skip these details and move on directly to the inference results—Theorem 5.15.

Zigzags manipulations. The manipulations take place at the algebraic level directly. They are of two types: reversing a single arrow or removing a single node in the quiver underlying a zigzag module. Here is a formal description:

LEMMA 5.13 (Arrow Reversal).

Let $\mathbb{V} = V_1 \text{---} \cdots \text{---} V_k \xrightarrow{v} V_{k+1} \text{---} \cdots \text{---} V_n$ be a zigzag module. Then, there is a map $V_k \xleftarrow{u} V_{k+1}$ such that $v \circ u|_{\text{im } v} = \mathbb{1}_{\text{im } v}$ and $u \circ v|_{\text{im } u} = \mathbb{1}_{\text{im } u}$, and the zigzag module \mathbb{V}^* obtained from \mathbb{V} by replacing the arrow $V_k \xrightarrow{v} V_{k+1}$ by $V_k \xleftarrow{u} V_{k+1}$ has the same persistence barcode as \mathbb{V} .

When v is injective, u is surjective and $u \circ v$ is the identity over the domain of v . Conversely, when v is surjective, u is injective and $v \circ u$ is the identity over the codomain of v . These properties are useful when \mathbb{V} is part of a commutative diagram because they help preserve the commutativity after the arrow reversal, as will be the case in the following. The next operation preserves the commutativity by construction.

LEMMA 5.14 (Space Removal).

Let \mathbb{V} be a zigzag module containing $V_k \xrightarrow{u} V_{k+1} \xrightarrow{v} V_{k+2}$. Then, replacing $V_k \xrightarrow{u} V_{k+1} \xrightarrow{v} V_{k+2}$ by $V_k \xrightarrow{v \circ u} V_{k+2}$ in \mathbb{V} simply removes the index k from its barcode, that is: intervals $[k, k]$ disappear, $[b, k]$ becomes $[b, k - 1]$, $[k, d]$ becomes $[k + 1, d]$, and all other intervals remain unchanged.

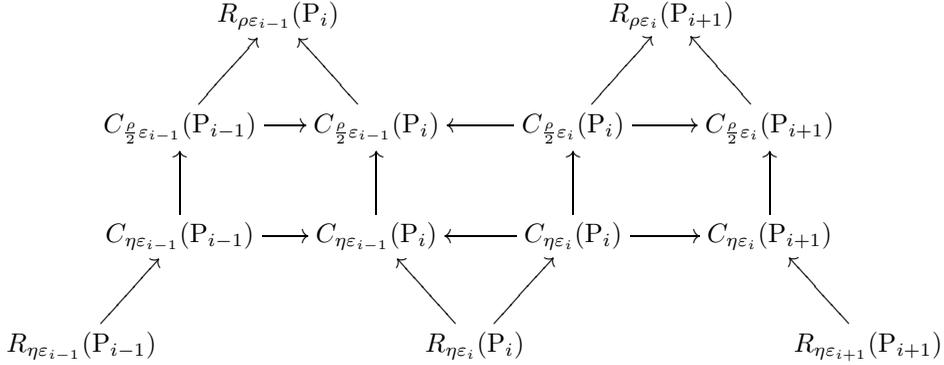
The proofs of these lemmas rely on the Interval Decomposition Theorem 1.9. The outline for Lemma 5.13 goes as follows:

- (1) Decompose \mathbb{V} into indecomposable modules according to Theorem 1.9. The decomposition is unique up to isomorphism and permutation of the terms.
- (2) For each element $\mathbb{W} = W_1 \text{---} \cdots \text{---} W_k \xrightarrow{w} W_{k+1} \text{---} \cdots \text{---} W_n \cong \mathbb{I}[b, d]$ in the decomposition, the map w is either an isomorphism (if $b \leq k < d$) or the zero map. In either case it is reversible without affecting the interval $[b, d]$: simply take w^{-1} if w is an isomorphism, and the zero map otherwise.
- (3) \mathbb{V}^* is obtained as the direct sum of the thus modified indecomposable modules, the reverse map $V_k \xleftarrow{u} V_{k+1}$ being the direct sum of their reverse maps.

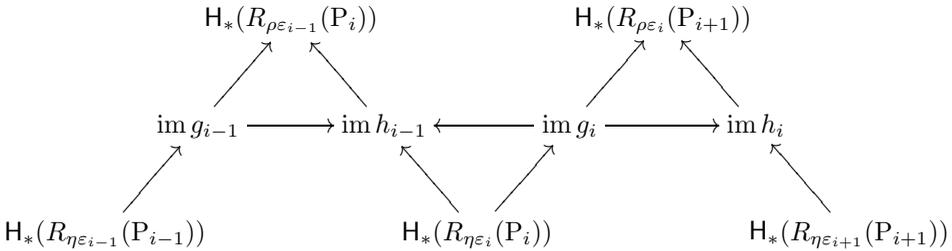
The claimed properties on \mathbb{V}^* and u follow by construction.

The outline for Lemma 5.14 is similar.

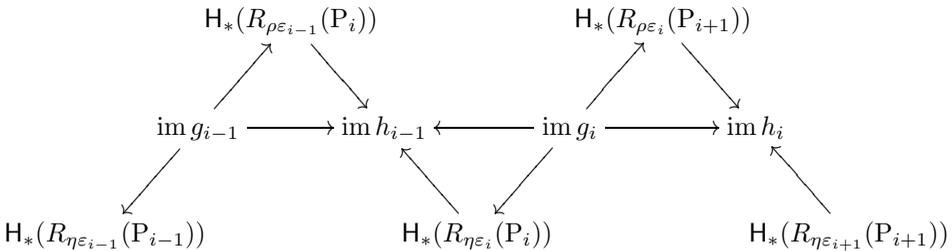
Application to Rips zigzags. The previous low-level zigzag manipulations can be effectively used to relate zigzag modules to one another, in the same spirit as what interleaving does with persistence modules. As an example, let us consider the case of the oscillating Rips zigzag (5.12). Assuming that $\eta < \frac{\rho}{2}$, the inclusion maps between Rips complexes factor through Čech complexes as follows according to (5.3):



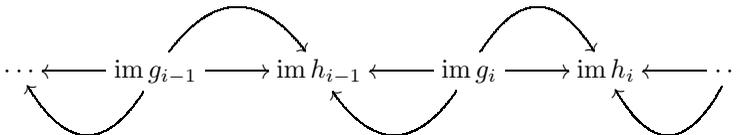
All arrows are inclusions here, so this diagram commutes and therefore induces a commutative diagram at the homology level. Calling $g_j : H_*(C_{\eta\varepsilon_j}(P_j)) \rightarrow H_*(C_{\frac{\rho}{2}\varepsilon_j}(P_j))$ and $h_j : H_*(C_{\eta\varepsilon_j}(P_{j+1})) \rightarrow H_*(C_{\frac{\rho}{2}\varepsilon_j}(P_{j+1}))$ the homomorphisms induced by the vertical arrows at every index j , we have



Let us now apply Lemma 5.13 to reverse every arrow $H_*(R_{\eta\varepsilon_j}(P_j)) \rightarrow \text{im } g_j$ and every arrow $H_*(R_{\rho\varepsilon_j}(P_{j+1})) \leftarrow \text{im } h_j$, so the diagram becomes



Finally, let us apply Lemma 5.14 to remove all the Rips complexes from the oscillating zigzag module. The result is the curved path in the following diagram:



As it turns out, the above sequence of operations preserves the commutativity of the diagrams⁴, therefore the straight-path and curved-path zigzag modules are in fact the same. Thus, by a sequence of low-level manipulations, we have turned the oscillating Rips zigzag (5.12) into a zigzag involving only Čech complexes. By tracking down the changes that occurred in the barcode during the sequence of manipulations, we can relate the barcodes of the two zigzags to each other.

Inference results. Combined with Corollary 4.13, the above manipulations provide the following theoretical guarantee on the existence of a sweet range in the barcode of the oscillating Rips zigzag under sufficient sampling conditions:

THEOREM 5.15. *Let ρ and η be multipliers such that $\rho > 10$ and $3 < \eta < \frac{\rho-4}{2}$. Let $K \subset \mathbb{R}^d$ be a compact set and let $P \subset \mathbb{R}^d$ be such that $d_H(P, K) < \varepsilon$ with*

$$\varepsilon < \min \left\{ \frac{\eta - 3}{6\eta}, \frac{\eta - 3}{3\rho + \eta}, \frac{\rho - 2\eta - 4}{6(\rho - 2\eta)}, \frac{\rho - 2\eta - 4}{5\rho - 2\eta} \right\} \text{wfs}(K).$$

Then, for any indices $l > k$ such that

$$\begin{aligned} \varepsilon_l &\geq \max \left\{ \frac{3\varepsilon}{\eta - 3}, \frac{4\varepsilon}{\rho - 2\eta - 4} \right\}, \text{ and} \\ \varepsilon_k &< \min \left\{ \frac{1}{6} \text{wfs}(K) - \varepsilon, \frac{1}{\rho + 1} (\text{wfs}(K) - \varepsilon) \right\}, \end{aligned}$$

the persistence barcode of the oscillating Rips zigzag (5.12), restricted to the index range $[k, l]$, has two types of intervals:

- *The intervals that span $[k, l]$ encode the homology of K , in the sense that their number for each homology dimension p is equal to the dimension of $H_p(K_r)$, for any $r \in (0, \text{wfs}(K))$.*
- *The remaining intervals are ephemeral, i.e. they have length zero⁵.*

The major difference with the inference results obtained in the previous sections using Čech, Rips, Delaunay, or mesh filtrations, is that here the amount of topological noise is reduced to an ephemeral quantity, so the signal-to-noise ratio within the sweet range is infinite. This is a remarkable property, not even satisfied by the offsets filtration \mathcal{P} itself when the μ -reach of K is zero. A similar theoretical guarantee can be proved for the image Rips zigzag, but not for the Morozov zigzag. For the latter indeed, a sweet range exists throughout which the topological signal persists, however there is currently no known bound on the topological noise, and whether one exists is still an open question. Since this zigzag involves a single multiplier instead of two, intuition suggests that there is not enough slack between different indices to kill the noise, however such a phenomenon has yet to be observed in experiments.

⁴This is true generally for space removals but not for arrow reversals. Here we are anticipating Theorem 5.15, whose sampling assumptions imply that every map $H_*(R_{\rho\varepsilon_j}(P_{j+1})) \leftarrow \text{im } h_j$ in the sweet range is injective, so its reverse counterpart $H_*(R_{\rho\varepsilon_j}(P_{j+1})) \rightarrow \text{im } h_j$ is its left inverse and therefore preserves the commutativity of the corresponding upper triangle. Similarly, every map $H_*(R_{\eta\varepsilon_j}(P_j)) \rightarrow \text{im } g_j$ in the sweet range is surjective, so $H_*(R_{\eta\varepsilon_j}(P_j)) \leftarrow \text{im } g_j$ is its right inverse and therefore preserves the commutativity of the corresponding lower triangle.

⁵Recall our convention from Chapter 1, which is to represent intervals in the finite set $\{1, \dots, n\}$ as closed intervals in \mathbb{R} . By ‘ephemeral interval’ is therefore meant an interval of type $[i, i]$.

Size bounds. Notice that the complexes involved in (5.11) are the same as the ones from Section 3, so the size bound from Theorem 5.12 applies to the Rips zigzags as well. More precisely, the size of the considered Rips complex at any given time remains linear in the input size n and scales up with the intrinsic dimensionality m of the data. For zigzags there is another important quantity to measure: the total number of simplex insertions and deletions. Indeed, this quantity drives the runtime of the zigzag persistence calculation, and it can be arbitrarily large compared to the size of the largest complex in the zigzag, in contrast to filtrations. The same kind of ball packing argument as in Theorem 5.12 applies for this quantity, with the twist that the oscillating Rips zigzag may insert the same simplex multiple times, which is not the case for the other Rips zigzags. All in all, the following bounds are derived:

THEOREM 5.16. *Suppose P is sitting in some metric space of doubling dimension m . Then, for any $k \geq 0$, the total number of k -simplices inserted in the construction of the Morozov zigzag is at most $2^{O(km)}n$, where n is the cardinality of P . The same bound applies to the image Rips zigzag. For the oscillating Rips zigzag, the bound becomes $2^{O(km)}n^2$.*

When P is sitting in \mathbb{R}^d but Hausdorff-close to some compact set K of doubling dimension m , then Theorem 5.16 can be applied to the scales within the sweet range associated to K . Assuming m or a reasonable upper bound is known, the total number of simplices inserted in the Morozov or image Rips zigzag within the sweet range is $2^{O(m^2)}n$, and this quantity becomes $2^{O(m^2)}n^2$ for the oscillating Rips zigzag.

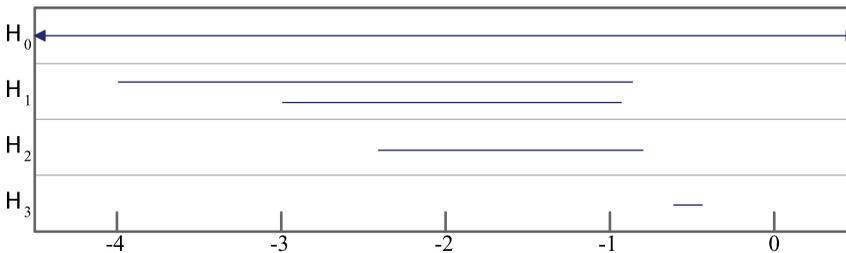


FIGURE 5.7. Logscale barcode of the oscillating Rips zigzag after removal of the ephemeral (length zero) intervals.

Experimental validation. Theorem 5.15 is illustrated in Figure 5.7, which shows the logscale barcode of the oscillating Rips zigzag computed on the Clifford data set described in introduction. The ephemeral (length zero) intervals were removed from the barcode before taking the picture. The barcode obtained with the image Rips zigzag is similar. These results are to be compared with the ones obtained previously using the offsets/Čech and mesh filtrations in Figures 5.1 and 5.5. They can also be compared with the one obtained using the Morozov zigzag in Figure 5.8, which contains a small amount of noise concentrated around the transition between the sweet ranges of the torus and of the 3-sphere. Overall the Morozov zigzag performs well on this particular example, however there currently is no guarantee that it will do so in general.

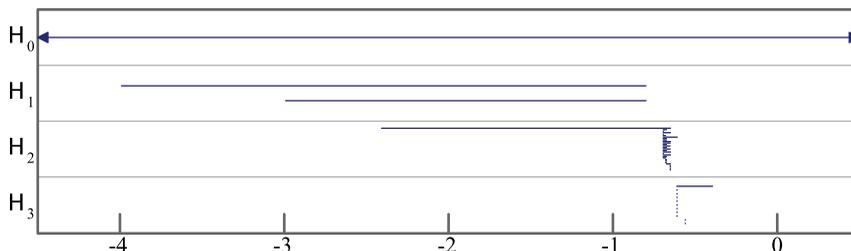


FIGURE 5.8. Logscale barcode of the Morozov zigzag after removal of the ephemeral (length zero) intervals.

— Reprinted from Oudot and Sheehy [208]: “Zigzag Zoology: Rips Zigzags for Homology Inference”, Foundations of Computational Mathematics, pp. 1–36, ©2014, with kind permission from Springer Science and Business Media.

filtration/zigzag	maximum complex size	total number of insertions
Čech	> 31,000,000,000	> 31,000,000,000
mesh	12,000,000	12,000,000
M-ZZ	107,927	398,107
iR-ZZ	174,436	1,003,215
oR-ZZ	174,436	7,252,772

TABLE 5.1. Maximum complex size (number of simplices) and total number of simplex insertions for various filtrations and zigzags on the Clifford data set. For filtrations the two quantities are equal.

Theorem 5.16 is illustrated in Table 5.1. Regarding the memory usage, the relevant quantity for filtrations is the total number of simplex insertions (third column), which is also the total size of the filtration. For zigzags, the relevant quantity is the maximum size of the complex at any given time (second column) because, as we saw in Section 2.2 of Chapter 2, the zigzag persistence algorithm of Carlsson, de Silva, and Morozov [50] proceeds by scanning through the zigzag once from left to right, loading a single complex in main memory at a time. On this front the contribution of the Rips zigzags is eloquent, reducing the number of simplices stored in main memory at a time from more than 31 billion to less than 200,000. The mesh filtration also does reasonably well, however its performance is likely to degrade quickly as the ambient dimension increases, even though the intrinsic dimension stays small. Note that for obvious reasons the memory usage of the iterative subsampling approach from Section 3 should be comparable to the one of the Rips zigzags.

Regarding the computation time, the relevant quantity for both filtrations and zigzags is the total number of simplex insertions (third column). On this front the performances of the Rips zigzags are more contrasted. While the Morozov zigzag is remarkably economical, the image and oscillating Rips zigzags are significantly

more greedy: the first one because it requires some bookkeeping to maintain two Rips complexes at any given time, the second one because it inserts the same simplex multiple times and therefore incurs a quadratic (in n) bound according to Theorem 5.16. In practice, computing any of the Rips zigzags together with its barcode takes a few minutes on a single Intel Xeon CPU core running at 2.40 GHz, meanwhile computing the mesh-based filtration and its barcode takes hours on a similar architecture, and the computation of the standard Rips filtration never completes because the main memory gets saturated.

4. Side-by-side comparison

Let us compare the simplicial filtrations and zigzags introduced in this chapter and the previous one with respect to objectives O1 through O3 stated in the introduction. For a fair comparison, we summarize their respective theoretical guarantees in terms of sweet range, memory usage, and computation cost in Table 5.2.

The first message is that the Rips zigzags perform consistently better than the other constructions on these criteria, and among them, the advantage goes to the iR-ZZ. This is to be taken with a grain of salt of course, given that the bounds are worst-case bounds. In practice it is up to the user to choose the construction that is best adapted to the data. For instance, in small dimensions the Delaunay filtration is an excellent choice, but in high dimensions or in more general metric spaces it is preferable to use constructions based only on distance comparisons.

Strictly speaking, the iR-ZZ is the only one among these constructions that fully fulfills objectives O1 through O3 mentioned in introduction. Others, like the M-ZZ or oR-ZZ, satisfy almost all the constraints and are viable alternatives in practice. In fact, the M-ZZ is particularly attractive due to its light weight, so much that it is recommended as a first approach on new data before trying other more elaborate constructions. The iterative resampling method from Section 3 is also a lightweight alternative, with strong guarantees on the noise level, however as we saw it does not offer a comprehensive picture of the evolution of the topology across scales since it builds multiple independent short filtrations. The other constructions scale up fairly badly with the ambient dimension and are therefore not recommended beyond medium dimensions.

Let us point out the existence of at least two other constructions fulfilling objectives O1 through O3: the *sparse Rips filtration* by Sheehy [221], and its variant by Dey, Fan, and Wang [107]. Their distinctive feature is to induce persistence modules at the homology level that are $(1 + O(\varepsilon))$ -interleaved multiplicatively with the one of the standard Rips filtration, without the need for a multiplicative $(1 + O(\varepsilon))$ -interleaving based on inclusions as in (5.1) at the topological level. Their corresponding figures are shown in the last row of Table 5.2. The fact that the noise level is at least the one of the standard Rips filtration makes the sparse Rips filtration and its variant a less preferred choice in the context of topological inference. Nevertheless, their ability to approximate the Rips filtration arbitrarily well makes them an invaluable tool in contexts where the Rips filtration is the main object of study, such as for instance when it is used as a topological signature for metric spaces. This aspect will be developed in Chapter 7, which gives the formal definition of the sparse Rips filtration.

To conclude this discussion, let us mention an interesting recent contribution of Sheehy [220], who showed how dimensionality reduction via random projections

filtration or zigzag	scale	sweet range		noise	memory usage	computation cost	predicate
		left bound	right bound				
Čech	linear	$O(\epsilon)$	$\Omega(\text{wfs}(K))$	$O(\epsilon)$	$O(2^n)$	$O(2^n)$	miniball
	linear	$O(\epsilon)$	$\Omega(\text{wfs}(K))$	$O(\epsilon)$	$O(n^{\lceil \frac{d}{2} \rceil})$	$O(n^{\lceil \frac{d}{2} \rceil})$	in_sphere
Rips	log	$\log_2(\epsilon) + O(1)$	$\log_2(\text{wfs}(K)) - O(1)$	2	$O(2^n)$	$O(2^n)$	distance
witness	log	$\log_2(\epsilon) + O(1)$	$\log_2(\text{wfs}(K)) - O(1)$	6	$O(2^n)$	$O(2^n)$	distance
mesh	log	$\log_2(\epsilon) + O(1)$	$\log_2(\text{wfs}(K)) - O(1)$	$O(\epsilon)$	$(\frac{1}{\epsilon})^{O(d^2)} n$	$(\frac{1}{\epsilon})^{O(d^2)} n$	in_sphere
M-ZZ	linear	$O(\epsilon)$	$\Omega(\text{wfs}(K))$?	$2^{O(m^2)} n$	$2^{O(m^2)} n$	distance
iR-ZZ	linear	$O(\epsilon)$	$\Omega(\text{wfs}(K))$	0	$2^{O(m^2)} n$	$2^{O(m^2)} n$	distance
oR-ZZ	linear	$O(\epsilon)$	$\Omega(\text{wfs}(K))$	0	$2^{O(m^2)} n$	$2^{O(m^2)} n^2$	distance
sparse Rips	log	$\log_2(\epsilon) + O(1)$	$\log_2(\text{wfs}(K)) - O(1)$	$2 + O(\epsilon)$	$(\frac{1}{\epsilon})^{O(m^2)} n$	$(\frac{1}{\epsilon})^{O(m^2)} n$	distance

TABLE 5.2. Comparison between the simplicial filtrations and zigzags of Chapters 4 and 5. The memory usage is measured by the number of simplices in the current complex at any time, while the computation cost is measured by the total number of simplex insertions.

à la Johnson and Lindenstrauss [166] can be used to map the input point cloud P to some linear space of dimension $O(\log n/\varepsilon^2)$, in such a way that the logscale persistence diagram of the offsets filtration \mathcal{P}' of the image P' remains ε -close to the one of \mathcal{P} in the bottleneck distance. The projection can be applied as a preprocessing step before running the inference pipeline. The complexity bounds of the various filtrations and zigags introduced in Chapter 5 and this one remain the same, except d is now replaced by $O(\log n/\varepsilon^2)$. The Delaunay and mesh-based filtrations are clearly the ones that benefit the most from this operation, however it does not seem sufficient to make them competitive with Rips-based zigags or sparse Rips filtrations in general.

5. Natural images

Let us now check the validity of the inference pipeline against real data. We will focus on one specific example, derived from the statistics of natural images, which is emblematic because it was the one put forward by the community for many years. This example indeed has interesting nontrivial topology, which the inference pipeline has helped uncover. The data set, described by Lee, Pederson, and Mumford [177], consists of 4.2 million high-contrast log-intensity 3×3 image patches, sampled from van Hateren's collection of still greyscaled images [153]. Each patch is represented as a point in Euclidean space \mathbb{R}^9 , with each coordinate giving the intensity at a specific pixel. After a proper contrast normalization, each patch becomes a point on the unit sphere \mathbb{S}^7 sitting in \mathbb{R}^8 .

It turns out that the entire sphere \mathbb{S}^7 is sampled by the data points, yet that the sampling density along the sphere is highly non-uniform. Lee, Pederson, and Mumford [177] set as their goal to understand the structure of the high-density regions, in order to get a better understanding of the cognitive representation of the space of images, and to obtain more realistic priors for applications such as object localization, segmentation, image reconstruction, denoising, and compression.

Direct inspection is made difficult by the dimensionality of the data. Figure 5.9 shows the results of various dimensionality reduction techniques on this data set, after applying a density thresholding (detailed below). These results are consistent across choices of input parameters, and the majority of them agrees on the overall structure of the high-density regions, suggesting that the data concentrate around some 1-d structure with 8 holes: 4 at the front, 4 at the back. The corresponding geometric representation would be 3 circles that intersect pairwise, with 6 points of intersection in total—the result of Isomap is the closest one to that representation.

In order to check this insight, de Silva and Carlsson [102] then Carlsson et al. [54] applied the topological inference pipeline (using witness filtrations) on the data. Due to algorithmic constraints, they had to pre-process the point cloud—called P hereafter—by applying the following statistical filters:

- (1) As before, they thresholded P by density using the k -th nearest neighbor density estimator, keeping only the fraction x of the data points with lowest k -th nearest neighbor distance. Varying k from 1,200 to 24,000 and x from 10% to 30%, they obtained a collection $P_{k,x}$ of high-density subsets of P —in Figure 5.9 we used $k = 1,200$ and $x = 30\%$.
- (2) Considering each set $P_{k,x}$ independently, they selected a subset $Q_{k,x}$ of 5,000 random points sampled uniformly, which they took as witness set.

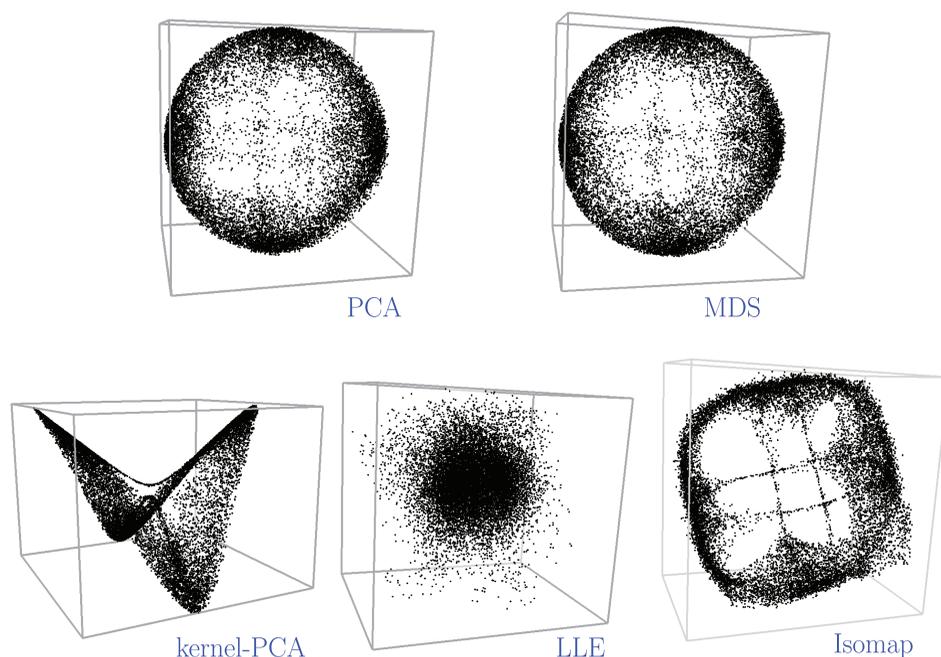


FIGURE 5.9. Results obtained using various dimensionality reduction techniques to project the natural images data down to 3 dimensions. In the top row, linear techniques: Principal Component Analysis (PCA), Multi-Dimensional Scaling (MDS). In the bottom row, non-linear techniques: kernel-PCA with a Gaussian kernel, Locally Linear Embedding (LLE), Isomap.

- (3) Among the points of $Q_{k,x}$ they selected a subset $L_{k,x}$ of 50 landmarks by a furthest-point sampling strategy.

After this preprocessing phase, they computed the barcode of the witness filtration for each pair $(L_{k,x}, Q_{k,x})$ separately. A fraction of their results in 1-homology is reproduced in Figure 5.10, and it show two trends :

- for smaller values of k , the barcodes reveal 5 long bars—Figure 5.10(b),
- for larger values of k , the barcodes reveal only 1 long bar—Figure 5.10(c).

These findings led Carlsson et al. [54] to conjecture that the data are actually concentrated around three circles with only 4 points of intersection instead of 6, as depicted in Figure 5.10(a). Indeed, the 5 long bars obtained with small values of k give the homology of the 3-circles model, while the single long bar obtained with larger values of k suggests that one of the three circles is prevailing over the other two.

These conclusions could be questioned because they rely on a selection of only 50 landmarks among millions of data points. In order to provide further validation, Oudot and Sheehy [208] used Rips zigzags on larger landmarks sets. Starting from the same preprocessed data, they took the whole sets $Q_{k,x}$ as landmarks and computed the barcodes of their associated Morozov zigzags up to homology dimension 3. They then subsampled the data down to 500 points to compute the barcodes

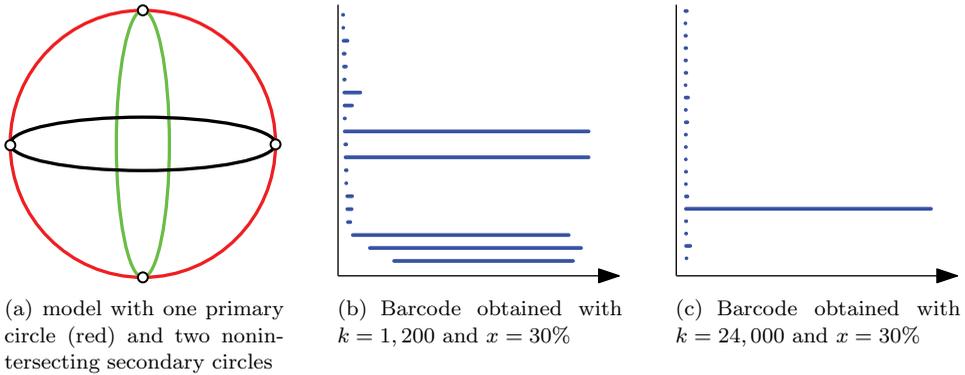


FIGURE 5.10. Experimental results obtained using the witness filtration with 50 landmarks and 5,000 witnesses.

— Based on de Silva and Carlsson [102].

for dimensions 4 to 7. The barcodes obtained for $Q_{k,x}$ with ($k = 1,200$, $x = 30\%$) and ($k = 24,000$, $x = 30\%$) are reported in Figures 5.11 and 5.12 respectively. They corroborate the results of Figure 5.10.

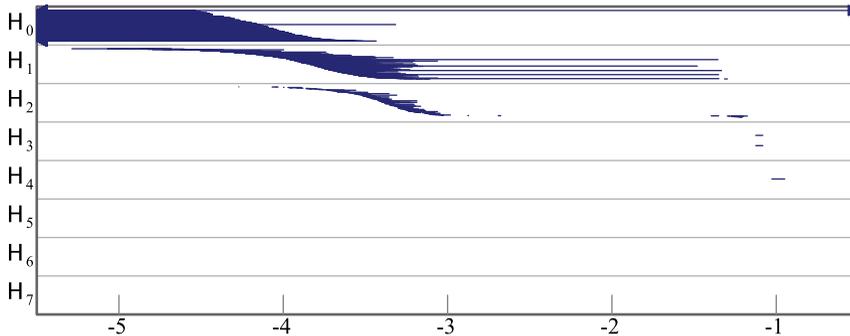


FIGURE 5.11. Logscale barcode of the Morozov zigzag on $Q_{k,x}$ with $k = 1,200$ and $x = 30\%$. Homology was computed with \mathbb{Z}_2 -coefficients.

— Reprinted from Oudot and Sheehy [208]: “Zigzag Zoology: Rips Zigzags for Homology Inference”, Foundations of Computational Mathematics, pp. 1–36, ©2014, with kind permission from Springer Science and Business Media.

The 3-circles model has led researchers to believe that the space of high-contrast 3×3 patches should also contain some higher-dimensional structure. To uncover it, Carlsson et al. [54] proposed a parametrization by a space of polynomials in two variables, which can be represented pictorially as in Figure 5.13. The main circle in the 3-circles model of Figure 5.10(a) is marked in red, while the two secondary circles are marked respectively in green and in black. The parametrization identifies the top and bottom horizontal lines together, and the left and right vertical lines together modulo a rotation by 180 degrees, so it is a parametrization of the Klein bottle.

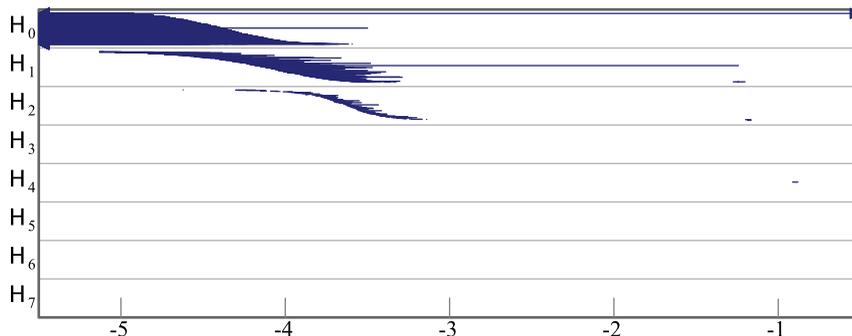


FIGURE 5.12. Logscale barcode of the Morozov zigzag on $Q_{k,x}$ with $k = 24,000$ and $x = 30\%$. Homology was computed with \mathbb{Z}_2 -coefficients.

— Reprinted from Oudot and Sheehy [208]: “Zigzag Zoology: Rips Zigzags for Homology Inference”, *Foundations of Computational Mathematics*, pp. 1–36, ©2014, with kind permission from Springer Science and Business Media.

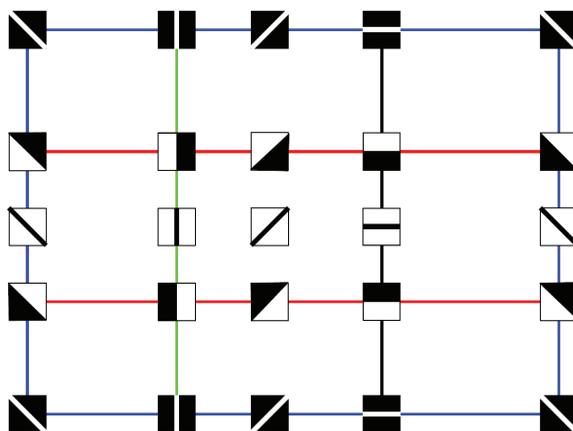


FIGURE 5.13. Parametrization of the high-contrast 3×3 patches.

— Reprinted from Carlsson et al. [54]: “On the Local Behavior of Spaces of Natural Images”, *International Journal of Computer Vision*, 76(1):1–12, ©2007, with kind permission from Springer Science and Business Media.

Validating the Klein bottle model experimentally turned out to be a significantly more difficult task, as the Klein bottle does not seem to appear in the barcode of any of the superlevel-sets of the density⁶. Nevertheless, as reported by Carlsson et al. [54], direct inspection of the data suggests that the Klein bottle is really sitting there but may not be statistically significant. The question remains

⁶Take for instance the results of Figures 5.11 and 5.12. These were obtained by computing persistent homology over the field of coefficients \mathbb{Z}_2 . This field does not make possible to distinguish the torus from the Klein bottle. Therefore, should a sweet range exist for the Klein bottle, the restriction of the barcode to that sweet range would be similar to the one of the barcodes from the previous sections, which is clearly not the case.

open to date, yet the important fact in this story is that the inference pipeline has provided relevant insights into the geometric structures underlying the data, which other techniques such as dimensionality reduction could not do because of the nontrivial topology of these structures.

6. Dealing with outliers

The data model considered so far is limited to bounded Hausdorff noise and therefore does not allow for the presence of outliers in the data. When this happens, such as for instance in the natural images data set of Section 5, the current pipeline leaves the user with no choice but to filter out the outliers in a preprocessing phase, which raises the tricky question of determining the right amount of filtering.

Chazal, Cohen-Steiner, and M erigot [62] proposed a different data model that elegantly takes the presence of outliers into account. Their model replaces the input n -points set P by its empirical measure μ_P , defined for any (Borel) subset of \mathbb{R}^d as

$$\mu_P(B) = \frac{1}{n} |B \cap P|.$$

The distance to P is then replaced by the following distance to its measure μ_P , parametrized by a mass parameter $m \in (0, 1]$:

$$(5.13) \quad d_{\mu_P, m}(x) = \sqrt{\frac{1}{k} \sum_{i=1}^k \|x - p_i(x)\|^2},$$

where $k = mn$ is assumed to be an integer for simplicity, and $p_i(x)$ is the i -th nearest neighbor of x among the points of P . The distance of x to μ_P is thus the ℓ^2 -average of the distances of x to its $k = mn$ nearest neighbors in P . As illustrated in Figure 5.14 (left and center), increasing the mass parameter m ‘smooths’ the sublevel-sets of $d_{\mu_P, m}$ and thus makes isolated outliers or small groups of those disappear from the small sublevel-sets. The choice of m then comes down to a trade-off between the amount of noise filtering and level smoothing versus the approximation accuracy.

More generally, the *distance to a measure* can be defined for any probability measure μ supported in \mathbb{R}^d . Its formula is the continuous analogue of (5.13), where the distance to the i -th nearest neighbor of x for a given $i \leq mn$ is replaced by the infimum $\delta_{\mu, m'}(x)$ of the radii of the Euclidean balls centered at x that contain at least a given fraction $m' = i/n$ of the total mass of μ :

$$d_{\mu, m}(x) = \sqrt{\frac{1}{m} \int_{m'=0}^m \delta_{\mu, m'}(x) dm'}.$$

The distance to a measure satisfies axioms A1 through A3 from Section 1.4 of Chapter 4, so it is a distance-like function. A generalized gradient flow can then be defined, which allows to introduce generalized critical values between which the sublevel sets of $d_{\mu, m}$ are homotopy equivalent. Besides, the distance to a measure is stable with respect to the Wasserstein distance W_2 : for any probability measures μ, ν supported in \mathbb{R}^d , and for a fixed mass parameter $m \in (0, 1]$,

$$\|d_{\mu, m} - d_{\nu, m}\|_\infty \leq \frac{1}{\sqrt{m}} W_2(\mu, \nu).$$

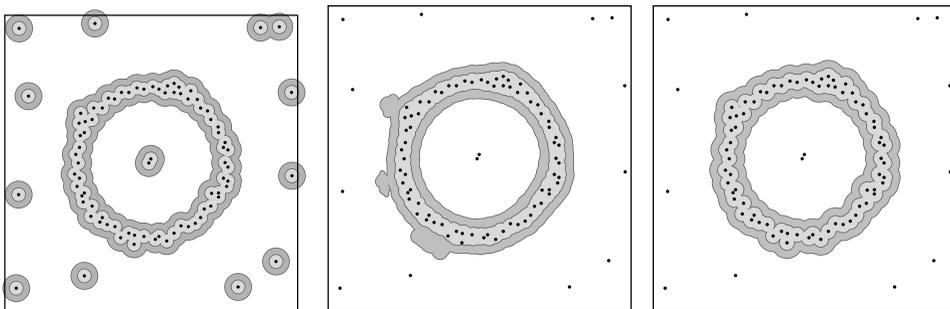


FIGURE 5.14. Sublevel sets of the distance-like functions d_P , $d_{\mu_P, m}$, and $d_{\mu_P, m}^P$ (shown in this order from left to right) for $m = \frac{3}{|P|}$. The first distance is sensitive to noise and outliers. The second one is smoother but its sublevel-sets are difficult to compute. The third one is resilient to noise and its sublevel-sets are both easy to compute and compact to represent.

— From Buchet et al. [42]: “Efficient and Robust Persistent Homology for Measures”, in Proceedings of the ACM/SIAM Symposium on Discrete Algorithms, pp. 168–180, 2015. Reprinted with permission.

This stability result makes it possible to interleave the sublevel-sets of the distances to Wasserstein-close probability measures, and combined with Axioms A1-A3, to derive inference results for probability measures and their support, in the same vein as in Section 2 of Chapter 4. The filtrations involved are sublevel-sets filtrations, therefore once again it is necessary to replace them by equivalent simplicial filtrations in order to make the approach practical.

From sublevel-sets filtrations to simplicial filtrations. The sublevel-sets of $d_{\mu_P, m}$ turn out to be unions of balls, as in the case of the distance d_P . However, the balls are no longer centered at the data points, but at their barycenters. Moreover, their radii are not all the same. Indeed, letting Q denote the set of barycenters of k -points subsets of P , we have equality between $d_{\mu_P, m}$ and the so-called *power distance* to the weighted set $\{(q, w_q) \mid q \in Q\}$, defined by

$$d_{\mu_P, m}(x) = \min_{q \in Q} \sqrt{\|x - q\|^2 + w_q^2},$$

where the weight w_q assigned to the barycenter q of a k -subset $\{p_1, \dots, p_k\} \subseteq P$ is⁷

$$(5.14) \quad w_q = \sqrt{\frac{1}{k} \sum_{i=1}^k \|q - p_i\|^2}.$$

Hence, for any $i \in \mathbb{R}$, the i -sublevel-set of $d_{\mu_P, m}$ is the union of the Euclidean balls centered at the points $q \in Q$ and of respective radii⁸

$$(5.15) \quad r_i(q) = \sqrt{i^2 - w_q^2}.$$

⁷This quantity is different from $d_{\mu_P, m}(q)$ because p_1, \dots, p_k may not be the k nearest neighbors of q in P .

⁸By convention, a ball is empty when its radius is imaginary.

The machinery of Section 3 of Chapter 4 can then be applied to derive simplicial filtrations on Q that are equivalent to the sublevel-sets filtration of $d_{\mu_P, m}$.

Unfortunately, the set Q has size $\Theta(n^k)$, therefore it cannot be computed or manipulated directly. Of course, not all balls participate in the union at each level i , however determining the balls that do participate amounts to computing the order- k Voronoi diagram of P in \mathbb{R}^d , where $k = mn$, which is classically done by computing the order- k' diagrams for all $k' \leq k$ successively in expected time $O(n^{\lceil \frac{d}{2} \rceil} k^{1 + \lfloor \frac{d}{2} \rfloor})$ [201]. This is costly even in small dimensions, since m is typically chosen so k is a constant fraction of n .

The workaround proposed by Guibas, Mérigot, and Morozov [146] is to drop most of the barycenters and keep only a small and easy-to-compute subset Q' . Each point $q \in Q'$ is the barycenter of some data point $p \in P$ and its k -nearest neighbors p_1, \dots, p_k . Its associated weight w_q is the same as in (5.14). The power distance $d_{\mu_P, m}^{Q'}$ to the weighted set $\{(q, w_q) \mid q \in Q'\}$ is called the *witnessed k -distance* by Guibas, Mérigot, and Morozov [146], who use it as a proxy for the true distance to measure $d_{\mu_P, m}$. This approach is justified by the following multiplicative interleaving:

$$(5.16) \quad \forall x \in \mathbb{R}^d, \quad d_{\mu_P, m}(x) \leq d_{\mu_P, m}^{Q'}(x) \leq \sqrt{6} \, d_{\mu_P, m}(x).$$

Then, the machinery of Section 3 of Chapter 4 can be applied to build a simplicial filtration on Q' that is equivalent to the sublevel-sets filtration of $d_{\mu_P, m}^{Q'}$, and whose logscale barcode is therefore $\frac{\log_2(6)}{2}$ -close to the one of $d_{\mu_P, m}$ in the bottleneck distance. Note that this time we are working with the decimated barycenters set Q' of size $|Q'| = n$ instead of the original barycenters set Q , so the approach becomes tractable.

Buchet et al. [42] proposed another decimation scheme that uses the input data points as ball centers instead of their barycenters—see Figure 5.14 (right) for an illustration. The weight associated to each point $p \in P$ is simply $d_{\mu_P, m}(p)$, and the approximating function is the power distance $d_{\mu_P, m}^P$ to the weighted set $\{(p, d_{\mu_P, m}(p)) \mid p \in P\}$. Surprisingly, this modification still allows us to get a constant-factor multiplicative interleaving as in (5.16):

$$(5.17) \quad \forall x \in \mathbb{R}^d, \quad \frac{1}{\sqrt{2}} \, d_{\mu_P, m}(x) \leq d_{\mu_P, m}^P(x) \leq \sqrt{3} \, d_{\mu_P, m}(x).$$

The great advantage of using the input data points as ball centers is that no barycenter calculation is needed any more, so the approach can be applied in arbitrary metric spaces. Note that the upper bound on $d_{\mu_P, m}^P$ in (5.17) holds only in Euclidean space \mathbb{R}^d . In arbitrary metric spaces it is replaced by $\sqrt{5} \, d_{\mu_P, m}(x)$, which is known to be tight [42].

Lightweight simplicial filtrations. Buchet et al. [42] also introduced weighted versions of the Rips and sparse Rips filtrations that allow to reproduce a good fraction of the analysis of this chapter for weighted point sets.

DEFINITION 5.17. Given a weighted point set $\{(p, w_p) \mid p \in P\}$, the *weighted (Vietoris-)Rips complex* of parameter i is the maximal simplicial complex of vertex set P whose 1-skeleton graph has one edge for each pair (p, q) such that $\|p - q\| \leq r_i(p) + r_i(q)$, where $r_i(q)$ (resp. $r_i(p)$) is defined as in (5.15). The *weighted (Vietoris-)Rips filtration* is the nested family of weighted Rips complexes obtained by letting parameter i range over \mathbb{R} from $-\infty$ to $+\infty$.

This filtration is interleaved multiplicatively with the nerve of the union of balls $B(p, r_i(p))$ just like its unweighted counterpart is interleaved multiplicatively with the nerve of the union of balls $B(p, i)$, and in fact the interleaving factors are the same as in (5.3). From this follows an inference result of the same form as Theorem 5.3.

The sparse version of the weighted Rips filtration is defined as the ‘intersection’ of two other filtrations: the sparse unweighted Rips, and the weighted Rips. Specifically, for any $i \in \mathbb{R}$, the sparse weighted Rips complex of parameter i is the intersection of its unweighted counterpart with the weighted Rips complex of same parameter. This definition may sound a bit awkward, but there is a fundamental justification to it, which is that the weighted distance is not a true distance therefore the furthest-point resampling technique from Section 3.1 and its near linear-time approximations like the *net-trees* of Har-Peled and Mendel [149] are not even defined, whereas they are the central building block of the sparse Rips construction. The size and approximation properties of the sparse Rips filtration, which will be described in Chapter 7, are maintained in the weighted setting thanks to this definition.

These contributions to topological inference in the presence of outliers leave many important questions open, such as for instance how many weighted points are required to approximate the true distance to measure within a given error. Mériçot [191] has given a lower bound that is exponential in the ambient dimension in the case of an additive approximation, but the case of a multiplicative approximation remains unanswered. This leaves room for future developments in this mostly unexplored direction of research. In Chapter 8 we will mention other current trends in topological inference, which remains to date the most active area of application of persistence.

Clustering

Unsupervised learning or clustering can be viewed as the zero-dimensional version of topological inference, in which the focus is primarily on understanding the connectivity of the structures hidden in the data. Given a point cloud P , the goal is to partition P into clusters that best reflect this connectivity. As occurs generally in topological inference, obtaining the most relevant clustering is an ill-posed problem, and it is particularly difficult with massive and high-dimensional data sets where visualization techniques fail.

The breadth of the existing work on clustering [150] shows the high interest this topic has aroused among the scientific community. Let us recount a few classical approaches before showing where and how persistence contributes to the problem:

K-means [181] is perhaps the most commonly used approach. Given a fixed number k of clusters, it tries to place cluster centers and define cluster boundaries so as to minimize the sum of the squared distances to the center within each cluster. This minimization problem is known to be NP-hard, so k -means resorts to an iterative expectation-maximization procedure that is guaranteed to converge at least to some local minimum. This minimum is not guaranteed to be global, however. Another issue with k -means and its variants is that they produce bad results on highly non-convex clusters.

Spectral clustering [182] was designed specifically to work on non-convex data. It first computes an embedding of the data set endowed with a diffusion distance between the points, given by a Laplacian of some neighborhood graph. Then, it applies the standard k -means method in the new ambient space. Computing the embedding requires an eigendecomposition of the Laplacian, which may have numerical issues as the size of the data grows. The presence of a gap in the spectrum of the Laplacian gives an indication of the correct number k of clusters. However, problems arise when there are more than a small number of outliers in the data, in which case no such gap may exist.

Density thresholding techniques make the assumption that the data points are drawn from some unknown density function f . Their principle is to threshold the density at some fixed level t , then treat the connected components of the superlevel set $F^t = f^{-1}([t, +\infty))$ as clusters and the rest of the data as noise. In practice, the density f is unknown so its superlevel F^t needs to be approximated from the data, which algorithms like DBSCAN [127, 219] do by various graph-based heuristics. Unfortunately, due to the use of a fixed density threshold t , these techniques do not respond well to hierarchical data sets, in which subtle multi-scale clustering phenomena may occur.

Hierarchical clustering [165, §3.2] has been introduced specifically to cope with multi-scale phenomena, albeit in a purely metric context. It produces a hierarchy of clusterings (see Figure 6.1), in which the bottom level represents the data points

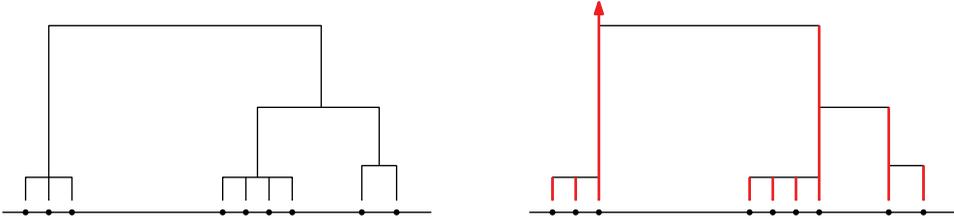


FIGURE 6.1. Left: a point cloud in \mathbb{R} and the corresponding dendrogram produced by single linkage clustering. Right: the hierarchy produced by the 0-dimensional persistence algorithm on the 1-skeleton graph of the Rips filtration, with the 0-dimensional barcode (vertical segments) marked in bold red lines. The two results are the same up to horizontal shifts of the vertical segments, which correspond to choosing arbitrarily the parent and the child each time a negative edge connects two previously independent connected components in the graph.

forming independent clusters, and the top level represents the entire point cloud gathered into a single cluster. The intermediate levels are nested in the sense that clusters cannot be split when travelling up the hierarchy. A convenient representation of the hierarchy is as a rooted tree, also called a *dendrogram*. The most popular algorithms to produce such hierarchies are agglomerative, starting with each point being its own cluster and merging at each step the two clusters that are most alike under some notion of (dis-)similarity. For instance, *single linkage* uses the minimum inter-point distance as dissimilarity, while *complete linkage* and *average linkage* use respectively the maximum and average inter-point distances. The qualities and defects of the produced hierarchies depend largely on the choice of (dis-)similarity between clusters. For instance, hierarchies produced by single linkage tend to suffer from the so-called ‘chaining effect’ caused by a chain of outlier data points connecting clusters too early in the agglomerative process [174]. Generally speaking, the main drawback of hierarchical clustering is to make difficult the recovery from bad clusterings performed in early stages of the construction.

Mode-seeking is another popular approach. Like density thresholding, it assumes the data points to be drawn from some unknown density function f . The approach consists in detecting the local peaks of f in order to use them as cluster centers, and in partitioning the data according to their *basins of attraction*. The precise notion of the basin of attraction B_p of a peak p varies between references, yet the bottomline is that B_p corresponds to the subset of the data points that eventually reach p by some greedy hill-climbing procedure. This line of work started in the 70’s, with for instance the graph based hill-climbing algorithm of Koontz, Narendra, and Fukunaga [170], and was followed by numerous variants and extensions, including the ambient gradient ascent algorithm Mean-Shift [91] and its successors [222, 233]. A common issue faced by these techniques is that the gradient and extremal points of a density function are notoriously unstable, so their approximation from a density estimator can lead to unpredictable results. This is why methods such as Mean-Shift adopt a proactive strategy that consists in smoothing the estimator before launching the hill-climbing procedure, which in turn raises the

difficult question of how much smoothing is needed to remove the noise without affecting the signal, and to obtain the correct number of clusters.

Prerequisites. For more background on clustering we recommend reading Chapter 14 of [151]. The rest of the material used here comes from Part 1.

1. Contributions of persistence

Persistence contributes in two ways to the literature on clustering. First, as a theoretical tool to improve or generalize the analysis of existing clustering methods. This is the case e.g. for hierarchical clustering. Second, persistence can be used as a mean to improve the methods themselves and to make them more efficient, more general, or more stable. This is the case e.g. for mode seeking.

1.1. Persistence versus hierarchical clustering. As observed already in the general introduction, persistence is naturally tied to hierarchical clustering. Indeed, the graph maintained by single linkage is the same as the 1-skeleton graph of the Rips filtration from Definition 5.2. Consequently, the dendrogram produced by single linkage contains the same information as the hierarchy induced by the persistence pairs computed by the 0-dimensional persistence algorithm—recall Figure 2.6 and related text in Chapter 2 for the definition of the hierarchy, and see Figure 6.1 for an illustration of the correspondence with the dendrogram. Hence, the inference results derived in the previous chapters, once specialized to 0-dimensional homology, provide guarantees on the existence of a *sweet range* in the dendrograms produced by single linkage, within which the ‘correct’ number of clusters is inferred.

Another important property of the dendrograms produced by single linkage is to be stable with respect to perturbations of the input data, as proven by Carlsson and Mémoli [51]. To formalize their result, we need to elaborate on the connection between dendrograms and ultrametrics on the input point cloud P —see e.g. [165, §3.2.3]. There is indeed a way to turn any dendrogram computed by single linkage into an ultrametric on P , by assigning to each pair of points of P the height of their lowest common ancestor in the dendrogram. This map turns out to be a bijection. The resulting ultrametric u_P may differ from the original metric d_P , however it is stable in the following sense, where d_{GH} denotes the so-called *Gromov-Hausdorff distance* between metric spaces:

THEOREM 6.1 (Carlsson and Mémoli [51]). *For any finite metric spaces (P, d_P) and (Q, d_Q) ,*

$$d_{\text{GH}}((P, u_P), (Q, u_Q)) \leq d_{\text{GH}}((P, d_P), (Q, d_Q)).$$

As we will see in Chapter 7 (until which we are deferring the formal definition of d_{GH}), part of this result¹ is a consequence of a more general stability property of Rips filtrations over compact metric spaces.

Unfortunately, these nice properties—*inference of the number of clusters and stability*—are only proven under the bounded Hausdorff noise model, to which the Gromov-Hausdorff distance is related. This model is very restrictive in the context of clustering. In particular, it forbids the aforementioned ‘chaining effect’ caused by outliers, which nonetheless happens in practical situations. The line of work on

¹Namely, the stability of the vertical part of the dendrogram (the barcode). The result of Carlsson and Mémoli [51] shows that the horizontal part of the dendrogram is stable as well in some sense.

topological inference in the presence of outliers, presented in Section 6 of Chapter 5, can be used to extend the guarantees to more general noise models, at the price of substantial modifications in the filtration and corresponding clustering hierarchy.

1.2. Persistence versus mode seeking. As we saw, mode seeking assumes the data points to be drawn from some unknown density function f , and it defines the clusters according to the basins of attraction of the peaks of f . In practice, f is approximated by some estimator \tilde{f} , whose gradient vector field may be noisy compared to the one of f , thus leading to unpredictable results. Persistence makes it possible to detect and merge the unstable clusters after their computation, so as to regain some stability. To be more specific, the persistence diagram of the superlevel-sets filtration $\tilde{\mathcal{F}}_{\geq}$ of \tilde{f} provides a measure of *prominence* for every peak of \tilde{f} , defined as the persistence of that peak as a 0-dimensional feature in the superlevel-sets filtration $\tilde{\mathcal{F}}_{\geq}$, and measured by the vertical distance of the corresponding diagram point to the diagonal $y = x$. Furthermore, the persistence algorithm applied to $\tilde{\mathcal{F}}_{\geq}$ organizes the peaks of \tilde{f} into a hierarchy that tells which connected component of the superlevel set was merged into which other component during the course of the computation. These two pieces of information—measure of prominence and hierarchy of peaks—can be used to drive the merging process on the clusters produced by mode seeking.

Several instances of this strategy have been developed in the literature. The most general one is due to Chazal et al. [70]. The method, called ToMATo for *Topological Mode Analysis Tool*, comes with theoretical guarantees on the number and location of the produced clusters. We propose to focus on it in the rest of the chapter. Section 2 provides the details of the algorithm, Section 3 reviews its theoretical guarantees, Section 4 gives experimental evidence of its practicality, finally Section 5 presents an extension of the method that can also capture higher-dimensional topological features in the data, at the price of an increased complexity.

REMARK. In order to simplify the exposition, from now on and until the end of the chapter we will use the notation $\text{dgm}_0(f)$ as a shorthand for the 0-dimensional part of the undecorated persistence diagram of the superlevel-sets filtration \mathcal{F}_{\geq} , that is,

$$\text{dgm}_0(f) = \text{dgm}(\text{H}_0(\mathcal{F}_{\geq})).$$

Unless otherwise mentioned, we will abuse terms and refer to this object simply as the *persistence diagram of f* .

2. ToMATo

The method combines the graph-based hill-climbing algorithm of Koontz, Narendra, and Fukunaga [170] with a cluster merging step guided by persistence. As illustrated in Figure 6.2(b), hill-climbing is very sensitive to perturbations of the density function f that arise from a density estimator \tilde{f} . Computing the persistence diagram of \tilde{f} makes it possible to quantify the prominences of its peaks and, in favorable cases, to distinguish those that correspond to peaks of the true density f from those that are inconsequential. In Figure 6.2(c) for instance, we can see 2 points (pointed to by arrows) that are further from the diagonal than the other points: these correspond to the 2 prominent peaks of \tilde{f} (one of them is at $y = -\infty$ since the highest peak never dies). To obtain the final clustering, it is enough to merge every cluster of prominence less than a given threshold τ into its

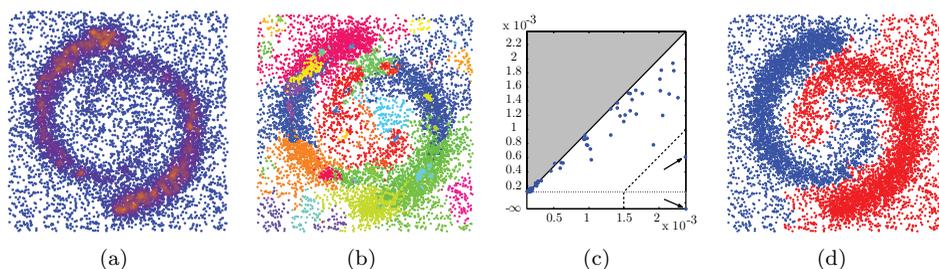


FIGURE 6.2. ToMATo in a nutshell: (a) estimation of the underlying density function f at the data points; (b) result of the basic graph-based hill-climbing step; (c) approximate persistence diagram showing 2 points far off the diagonal corresponding to the 2 prominent peaks of f ; (d) final result obtained after merging the clusters of non-prominent peaks.

— From Chazal et al. [70].

parent cluster in the hierarchy of peaks. As shown in Figures 6.2(c) and 6.2(d), the persistence diagram gives a precise understanding of the relationship between the choice of τ and the number of obtained clusters.

2.1. The algorithm. ToMATo takes as input an unweighted graph G , whose vertex set represents the data points and whose edges connect the points according to some user-defined proximity rule. Each vertex i of G must be assigned a non-negative value $\tilde{f}(i)$ corresponding to the estimated density at that point. In addition, ToMATo takes in a *merging* parameter $\tau \geq 0$. The algorithm proceeds in two steps. The first step implements the graph-based hill-climbing method of Koontz, Narendra, and Fukunaga [170], which simulates a gradient ascent in the graph G and partitions the vertex set according to the influence regions of the local peaks of \tilde{f} in G . The second step implements a variant of the union-find algorithm for computing 0-dimensional persistence (described in Section 2.2.3 of Chapter 2): in addition to building a hierarchy of the peaks of \tilde{f} , this variant uses it to merge the clusters of prominence less than τ into the other clusters. Here are the details:

- (1) (Mode-seeking) To compute the initial clusters, ToMATo iterates over the vertices of G sorted by decreasing \tilde{f} -values: at each vertex i , it simulates the effect of the gradient of the underlying density function by connecting i to its neighbor in G with highest \tilde{f} -value, if that value is higher than $\tilde{f}(i)$. Otherwise, all neighbors of i have lower values, so i is declared a peak of \tilde{f} . The resulting collection of pseudo-gradient edges forms a spanning forest of the graph, and each tree in this forest can be viewed as the analogue within G of the basin of attraction of a peak of the true density function in the underlying continuous domain.
- (2) (Merging) To handle merges between trees, ToMATo iterates over the vertices of G again, in the same order, while maintaining a union-find data structure \mathcal{U} , where each entry corresponds to a union of trees of the spanning forest. Let us call *root* of an entry e , or $r(e)$ for short, the vertex contained in e whose \tilde{f} -value is highest. By definition, this vertex is the

root of one of the trees contained in e , that is, a local peak of \tilde{f} in G . During the iteration process, two different situations may occur when a vertex i is considered:

- (a) Vertex i is a peak of \tilde{f} within G , *i.e.* the root of some tree T . Then, i creates a new entry e in \mathcal{U} , in which T is stored. Let $r(e) = i$.
- (b) Vertex i is not a peak and therefore belongs to some tree stored in an existing entry e_i of \mathcal{U} (of which i is not the root). Then, compute the set \mathcal{E} of the entries of \mathcal{U} that contain neighbors of i in G . Iterate over this set in any order, and for each entry $e \in \mathcal{E}$ considered, check whether $e \neq e_i$ and $\min\{\tilde{f}(r(e)), \tilde{f}(r(e_i))\} < \tilde{f}(i) + \tau$, that is, whether the two entries differ and at least one of them has a less-than- τ -prominent root. If so, then merge e and e_i into a single entry $e \cup e_i$ in \mathcal{U} , and let $r(e \cup e_i) = \arg \max_{\{r(e), r(e_i)\}} \tilde{f}$, so in effect the entry with the lower root is merged into the one with the higher root, as per the *elder rule*.

Upon termination, the (merged) clusters stored in the entries of the union-find data structure \mathcal{U} form a partition of the vertex set of G , and their roots are the peaks of \tilde{f} of prominence at least τ within the graph. The output of ToMATo is then the subset of this collection of clusters that is stored in those entries e such that $\tilde{f}(r(e)) \geq \tau$. The rest of the data points is stored in entries with roots lower than τ , so it is treated as background noise and discarded from the data set².

In addition to the clustering, ToMATo outputs the lifespans of all the entries that have been created in the union-find data structure during the merging phase. More precisely, an entry is considered ‘born’ when it is created in \mathcal{U} with a single tree attached to it as described in scenario (a) above, and ‘dead’ when it gets merged into another entry with higher root as described in scenario (b). For ease of visualization, the lifespan is represented as a point (x, y) in the plane, where x is the birth time and y the death time of the entry ($y = -\infty$ if the entry never gets merged into another one). It is easy to see that the thus obtained planar diagram of points coincides with the persistence diagram of the scalar field \tilde{f} over the graph G when parameter τ is set to $+\infty$, as in this case the condition $\min\{\tilde{f}(r(e)), \tilde{f}(r(e_i))\} < \tilde{f}(i) + \tau$ in scenario (b) is trivially satisfied so the merging rule is simply the one prescribed by persistence. When $\tau < +\infty$, the entries whose roots are at least τ -prominent never get merged into other entries, so their corresponding points in the output diagram are projected down vertically onto the horizontal line $y = -\infty$.

REMARK. The output of ToMATo can also be viewed as a hierarchical clustering, in which every value of $\tau \geq 0$ is assigned a particular set of clusters. The family of clusterings obtained for τ ranging from 0 to $+\infty$ is nested as desired, however the bottom level contains the clusters produced by the mode-seeking step instead of a collection of singletons. In this respect, ToMATo is really a combination of mode seeking and hierarchical clustering. In practice ToMATo must be run twice. In the first run, set $\tau = +\infty$ to merge all clusters and compute the persistence

²This extra filtering step stems from the observation that the data points may not be densely sampled over the entire space. Depending on the proximity rule used in the definition of the neighborhood graph G , the sparseness of the data in low-density regions may create independent connected components that give birth to spurious clusters with infinite prominence — see Figure 6.7 for an illustrative example.

diagram $\text{dgm}_0(\tilde{f})$. Then, using the persistence diagram, choose a value for τ —which amounts to selecting the number of clusters—and then re-run the algorithm to obtain the final clustering.

2.2. Implementation details and complexity. In practice the mode-seeking and merging procedures can be run simultaneously during a single pass over the vertices of the graph G : for each considered vertex i , the pseudo-gradient edge at i is computed, then the possible merges in the union-find data structure \mathcal{U} are performed—these involve only previously visited vertices. The corresponding pseudo-code is given in Algorithm 6.1.

Algorithm 6.1: ToMATo

Input: graph G with n vertices, n -dimensional vector \tilde{f} , parameter $\tau \geq 0$.

- 1 Sort the vertex indices $\{1, 2, \dots, n\}$ so that $\tilde{f}(1) \geq \tilde{f}(2) \geq \dots \geq \tilde{f}(n)$;
- 2 Initialize a union-find data structure \mathcal{U} and two vectors g, r of size n ;
- 3 for $i = 1$ to n do
- 4 Let \mathcal{N} be the set of neighbors of i in G that have indices lower than i ;
- 5 if $\mathcal{N} = \emptyset$ then // vertex i is a peak of \tilde{f} within G
- 6 Create a new entry e in \mathcal{U} and attach vertex i to it;
- 7 $r(e) \leftarrow i$; // $r(e)$ stores the root vertex of entry e
- 8 else // vertex i is not a peak of \tilde{f} within G
- 9 $g(i) \leftarrow \arg \max_{j \in \mathcal{N}} \tilde{f}(j)$; // g stores pseudo-gradient edges
- 10 $e_i \leftarrow \mathcal{U}.\text{find}(g(i))$;
- 11 Attach vertex i to the entry e_i ;
- 12 for $j \in \mathcal{N}$ do
- 13 $e \leftarrow \mathcal{U}.\text{find}(j)$;
- 14 if $e \neq e_i$ and $\min\{\tilde{f}(r(e)), \tilde{f}(r(e_i))\} < \tilde{f}(i) + \tau$ then
- 15 $\mathcal{U}.\text{union}(e, e_i)$;
- 16 $r(e \cup e_i) \leftarrow \arg \max_{\{r(e), r(e_i)\}} \tilde{f}$;
- 17 $e_i \leftarrow e \cup e_i$;
- 18 end
- 19 end
- 20 end
- 21 end

Output: the collection of entries e of \mathcal{U} such that $\tilde{f}(r(e)) \geq \tau$.

The mode-seeking phase takes a linear time in the size of G once the vertices have been sorted. As for the merging phase, it makes $O(n)$ **union** and $O(m)$ **find** queries to the union-find data structure \mathcal{U} , where n and m are respectively the number of vertices and the number of edges of G . If an appropriate representation is used for \mathcal{U} (e.g. a disjoint-set forest [92]), and if the heads of the pseudo-gradient edges and the entry roots are stored in separate containers with constant-time access (e.g. arrays), then the worst-case time complexity of Algorithm 6.1 becomes $O(n \log n + m\alpha(n))$, where α stands for the inverse Ackermann function.

As for the space complexity, note that the graph G does not have to be stored entirely in main memory, since only the neighborhood of the current vertex i is

involved at the i -th iteration of the `clustering` procedure. The main memory usage is thus reduced to $O(n)$. The total space complexity still remains $O(n + m)$ though, as the graph needs to be stored somewhere, e.g. on the disk.

3. Theoretical guarantees

Let X be an m -dimensional Riemannian manifold with positive convexity radius $\varrho(X)$. Recall that the convexity radius at a point $x \in X$ is the supremum over the values $r \geq 0$ such that any geodesic ball of center x and radius $r' < r$ is geodesically convex, that is, any two points in that ball are joined by a unique geodesic of length less than $2r'$, and this geodesic is contained in the ball. The convexity radius $\varrho(X)$ is the infimum of this quantity for x ranging over X .

Let $f : X \rightarrow \mathbb{R}$ be a probability density with respect to the m -dimensional Hausdorff measure. Assume that f is c -Lipschitz in the geodesic distance, and that the input data set P has been sampled over X according to f in i.i.d. fashion. Assume further that the density estimator \tilde{f} approximates the true density f within an error $\eta \geq 0$ over P , that is,

$$\forall p \in P, |\tilde{f}(p) - f(p)| \leq \eta.$$

For simplicity, assume also that the geodesic distances in X between the data points are known exactly. Similar guarantees exist (with slightly degraded constants) for when they are known within a small error—see e.g. [70] for the details. Take for G the so-called δ -neighborhood graph of P , which has P as vertex set and one edge for each pair of points lying within geodesic distance δ of each other. When $X = \mathbb{R}^d$, this graph is the same as the 1-skeleton graph of the Rips complex of parameter δ on P , as defined in Chapter 5 (Definition 5.2).

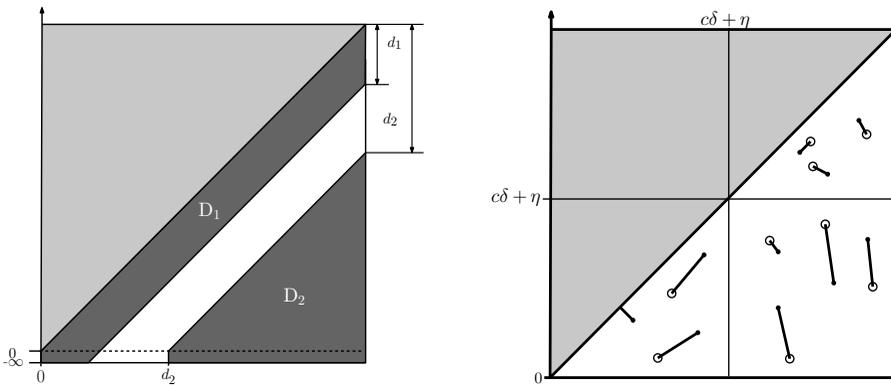


FIGURE 6.3. Left: The separation of the persistence diagram $\text{dgm}_0(f)$ between prominent peaks (region D_2) and topological noise (region D_1). Right: The partial matching between $\text{dgm}_0(f)$ and $\text{dgm}_0(\tilde{G})$ induced by the partial interleaving between \mathcal{F}_{\geq} and \tilde{G} : the bottleneck cost is fully controlled in the upper-right quadrant, while it is controlled only along the x -axis in the bottom-right quadrant, and it is uncontrolled in the lower-left quadrant.

— From Chazal et al. [70].

3.1. Number of clusters. The main guarantee is about the number of clusters output by the algorithm, which is related to the number of prominent peaks of f . The guarantee relies on the following condition on the persistence diagram of f :

DEFINITION 6.2. Given two values $d_2 > d_1 \geq 0$, the persistence diagram $\text{dgm}_0(f)$ is called (d_1, d_2) -separated if every point of $\text{dgm}_0(f)$ lies either in the region D_1 above the diagonal line $y = x - d_1$, or in the region D_2 below the diagonal line $y = x - d_2$ and to the right of the vertical line $x = d_2$.

This condition formalizes the intuitive notion that the points of $\text{dgm}_0(f)$ can be separated between prominent peaks (region D_2) and topological noise (region D_1), as illustrated in Figure 6.3 (left). In this respect, it acts very similarly to a signal-to-noise ratio condition: the larger the prominence gap $d_2 - d_1$, the more clearly the prominent peaks are separated from the noise. In the limit case where $d_1 = 0$, all peaks of f are at least d_2 -prominent and none of them is viewed as noise. The additional condition that the points of D_2 must lie to the right of the vertical line $x = d_2$ follows the description of the extra filtering step performed by the algorithm after the merging phase, and it stems from the fact that only some superlevel set of the density f can be densely sampled by the data points.

THEOREM 6.3 (Chazal et al. [70]). *Assume $\text{dgm}_0(f)$ is (d_1, d_2) -separated, with $d_2 - d_1 > 5\eta$. Then, for any positive parameter $\delta < \min\{\varrho(X), \frac{d_2 - d_1 - 5\eta}{5c}\}$ and any threshold $\tau \in (d_1 + 2(c\delta + \eta), d_2 - 3(c\delta + \eta))$, the number of clusters computed by the algorithm on an input of n sample points drawn at random according to f in an i.i.d. fashion is equal to the number of peaks of f of prominence at least d_2 with probability at least $1 - e^{-\Omega((c\delta + \eta)^n)}$, where the constant hidden in the big- Ω notation depends only on certain geometric quantities (e.g. volumes of geodesic balls) on the manifold X .*

Thus, for right choices of parameters τ and δ , the algorithm outputs the correct number of clusters with high probability. The larger the prominence gap $d_2 - d_1$, the larger the range of admissible values for τ . Meanwhile, the smaller δ , the larger the range but also the smaller the probability of success³. The proof of the theorem, detailed in [70], is rather technical but boils down to the following intuitive steps:

- (1) By a simple application of Boole's inequality, a bound is derived on the probability that the superlevel set $F^{c\delta + \eta}$ of f is densely sampled by a subset of the points of P .
- (2) Assuming that $F^{c\delta + \eta}$ is indeed densely sampled, an interleaving is worked out between the superlevel-sets filtration \mathcal{F}_{\geq} of the density f and the filtration $\tilde{\mathcal{G}}$ of the neighborhood graph G induced by the estimator⁴ \tilde{f} . The interleaving happens at the 0-dimensional homology level directly. It holds only for filtration parameters $t \geq c\delta + \eta$, therefore we call it a *partial interleaving*.
- (3) The stability part of the Isometry Theorem 3.1 is invoked to derive a partial matching between the persistence diagrams $\text{dgm}_0(f)$ and $\text{dgm}_0(\tilde{\mathcal{G}})$

³This follows the intuition that a minimum point density is required for the connectivity of the δ -neighborhood graph to reflect the one of some superlevel set of f .

⁴The time of appearance of a vertex i in $\tilde{\mathcal{G}}$ is $\tilde{f}(i)$, while the time of appearance of an edge $[i, j]$ is $\min\{\tilde{f}(i), \tilde{f}(j)\}$. This filtration is also known as the *upper-star filtration* of \tilde{f} in G —see [115, §V.3].

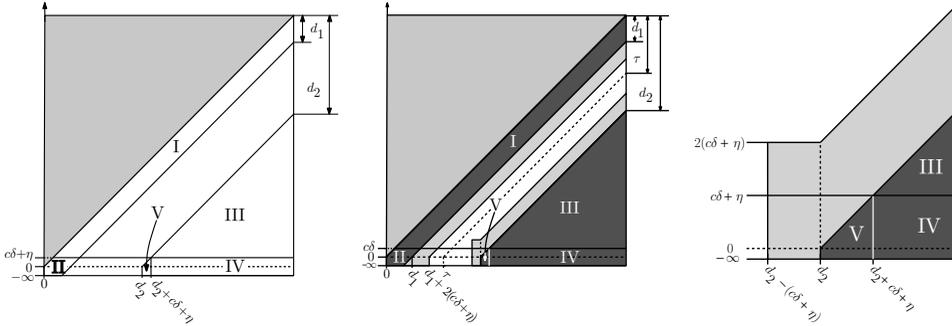


FIGURE 6.4. Influence of the partial matching on the regions D_1 and D_2 . Left: the initial regions, partitioned as $D_1 = I \cup II$ and $D_2 = III \cup IV \cup V$. Center: the perturbed regions. Right: zoom on part V of region D_2 .

— From Chazal et al. [70].

from the partial interleaving between \mathcal{F}_{\geq} and $\tilde{\mathcal{G}}$. As illustrated in Figure 6.3 (right), the bottleneck cost of the matching is controlled in the upper-right quadrant $[c\delta + \eta, +\infty] \times [c\delta + \eta, +\infty]$, which is sufficient for our purpose.

- (4) The partial matching perturbs the regions D_1 and D_2 of Figure 6.3 (left) as depicted in Figure 6.4. For a small enough choice of $\delta > 0$ compared to the initial separation $(d_2 - d_1)$, the perturbed regions remain disjoint, so they can still be separated using an appropriate value for the merging parameter τ . For such values, the algorithm outputs the correct number of clusters.

3.2. Geometric approximation. Another guarantee is about the geometric approximation of the basins of attraction of the peaks of f by the output clusters. Obtaining such a guarantee in full generality is hopeless, since the basins of attraction are notoriously unstable. There are indeed many examples of very close functions having very different basins of attraction, and clearly the algorithm cannot provably-well approximate the unstable parts of the basins, as illustrated in Figures 6.5 and 6.6. Yet, it can approximate the stable parts, as shown by the following result:

THEOREM 6.4 (Chazal et al. [70]). *Under the hypotheses of Theorem 6.3, it holds with the same probability that for every point $p \in D_2$ the algorithm outputs a cluster C such that $C \cap F^t = B_p \cap P \cap F^t$ for all values $t \in [t_p + d_1 + \frac{5}{2}(c\delta + \eta), f(m_p))$, where m_p is the peak of the density f (in the underlying manifold X) corresponding to the diagram point p , where B_p denotes the basin of attraction of m_p in the underlying manifold X , and where t_p is the first value of t at which B_p gets connected to the basin of attraction of another peak of f of prominence at least τ in the superlevel set F^t .*

Here, by *basin of attraction of m_p* is meant the union of the ascending regions (or stable manifolds) of all the peaks of f of prominence less than τ that are subordinate to m_p in the persistence hierarchy. It depends on the choice of parameter τ , which is omitted in the notation for simplicity. In plain words, the theorem says

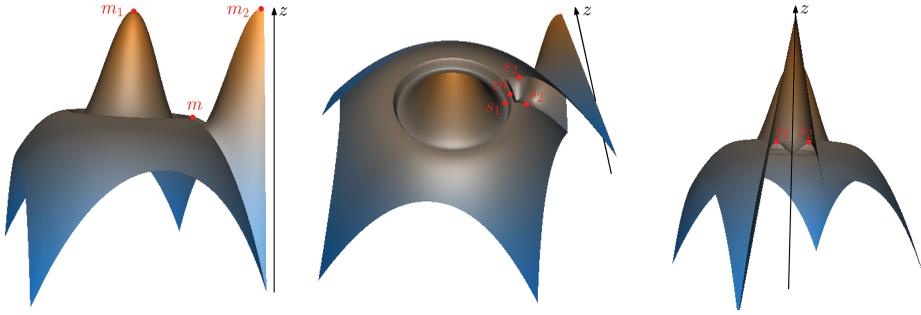


FIGURE 6.5. A function $f : [0, 1]^2 \rightarrow \mathbb{R}$ with unstable basins of attraction. The three peaks m, m_1, m_2 have respective prominences $f(m) - f(s_2)$, $f(m_1) - f(s_1)$, and $+\infty$. When $\tau > f(m) - f(s_2)$, the basin of attraction of m is merged into that of m_2 at the value $t = f(s_2)$. However, since $f(s_2) - f(s_1)$ can be made arbitrarily small compared to $f(m_1) - f(m)$, arbitrarily small perturbations of f compared to the prominence gap $f(m_1) - f(m) + f(s_2) - f(s_1)$ merge the basin of attraction of m into that of m_1 instead.

— From Chazal et al. [70].

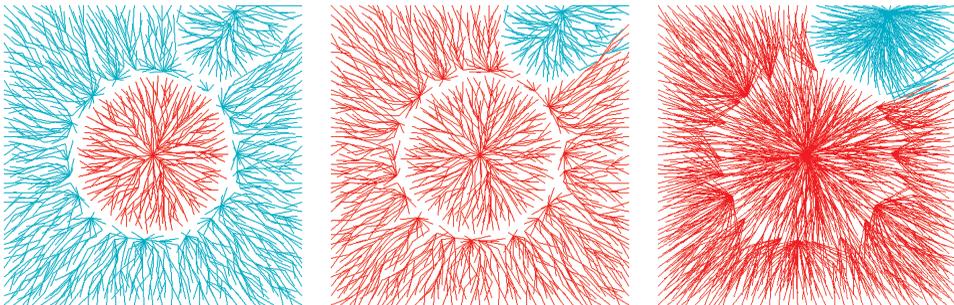


FIGURE 6.6. Outputs of the algorithm obtained from a uniform sampling of the unit square endowed with the function f of Figure 6.5. The chosen value for τ gives two clusters. The result is shown for three increasing values of the neighborhood radius. Notice how some values of δ induce a correct merge of the basin of attraction of m into that of m_2 , while others induce an incorrect merge into that of m_1 .

— From Chazal et al. [70].

that cluster C is the *trace* of B_p over the point cloud P , until (approximately) the value t_p at which B_p meets the basin of attraction of another τ -prominent peak of f . Below that value, the cluster may start diverging from the basin, which itself may start being unstable, as illustrated in Figures 6.5 and 6.6.

The proof of the theorem is technical but the underlying intuition is simple: above t_p , a point of $P \cap B_p$ cannot escape B_p when following the gradient of f or any ascending path of f . Indeed, even if the point eventually flows to another peak of f than m_p , that peak must have prominence less than τ therefore its ascending

region gets merged into B_p . Slightly above t_p , even the approximate ascending path given by the pseudo-gradient edges in the neighborhood graph cannot escape B_p . We refer the interested reader to [70] for the details of the proof.

4. Experimental results

As an illustration, let us show some experimental results on two types of inputs: (1.) a structured synthetic data set in \mathbb{R}^2 , where direct data inspection allows us to check the results visually; (2.) simulated alanine-dipeptide protein conformations in \mathbb{R}^{21} , where the knowledge of the intrinsic parameters of the simulation allows us to check the results *a posteriori*. Two density estimators are used in the experiments: a (truncated) Gaussian kernel estimator, and the so-called *weighted k-nearest neighbor (k-nn) estimator*⁵ proposed by Biau et al. [26]. Each of these estimators takes in one parameter controlling the amount of smoothing.

4.1. Synthetic data. The data set consists of 10k points sampled from two twin spirals in the unit square, as shown in Figure 6.2 (a). Using a δ -neighborhood graph with $\delta = 0.04$, and the weighted k -nn density estimator, we obtain the persistence diagram in Figure 6.2(c). Choosing τ within the prominence gap we obtain the clustering shown in Figure 6.2(d). A smaller neighborhood radius, $\delta = 0.02$, gives many infinitely persistent components (Figure 6.7(a)), with all but one appearing late in the persistence diagram (near the lower-left corner). Components in this part of the persistence diagram are discarded by the extra filtering step performed by the algorithm after completion of the merging phase, which removes much of the background noise—Figure 6.7(b).

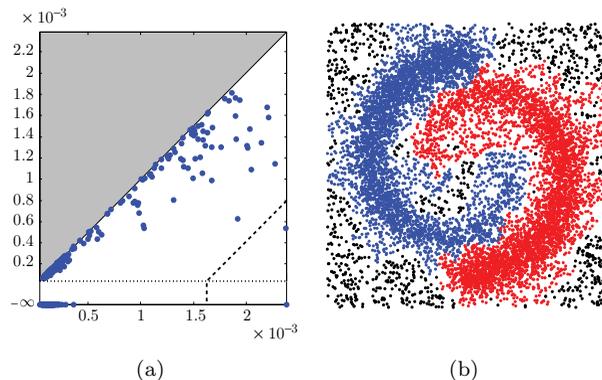


FIGURE 6.7. The twin spirals data set from Figure 6.2, processed using a smaller neighborhood radius: (a) the persistence diagram; (b) the final clustering with late appearing connected components filtered out (in black).

— From Chazal et al. [70].

⁵Which corresponds to taking the inverse of (a power of) the distance to the empirical measure (5.13).

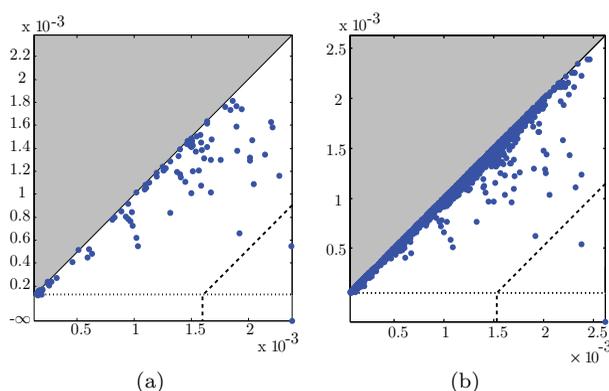


FIGURE 6.8. Persistence diagrams obtained on the twin spirals data set of Figure 6.2 using (a) the k -nn graph with $k = 35$ and (b) the Delaunay graph. The resulting clusterings are virtually the same as in Figure 6.2(d).

— From Chazal et al. [70].

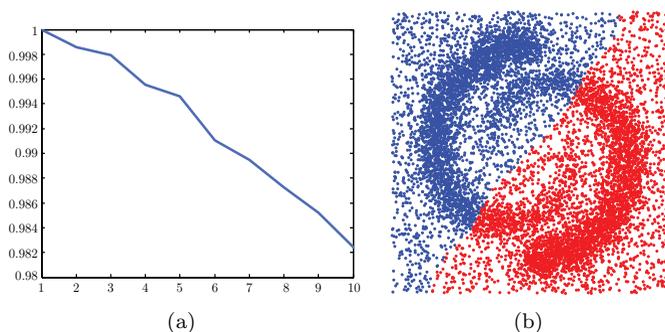


FIGURE 6.9. Result of spectral clustering on the twin spirals data set with 10k samples: (a) plot of the first 10 eigenvalues, and (b) the resulting clustering.

— From Chazal et al. [70].

Results obtained with the k -nn graph (taking $k = 35$) and the Delaunay graph are shown in Figure 6.8. Although not identical, they share the same overall structure with 2 prominent peaks, and the resulting clusterings are virtually identical to the one in Figure 6.2(d).

For comparison, Figure 6.9 shows the result obtained by spectral clustering [76] using the k -nn graph on the twin spirals data set with 10k samples. The result is consistent across choices of input parameters. It is explained by the effect of the background noise on the k -means procedure in eigenspace.

These results illustrate the good behavior of ToMATo in practice. This is a general trend of mode-seeking enhanced by persistence, whose specialized instances, such as the one developed by Paris and Durand [209] for image segmentation, are reported to behave as well on their respective inputs.

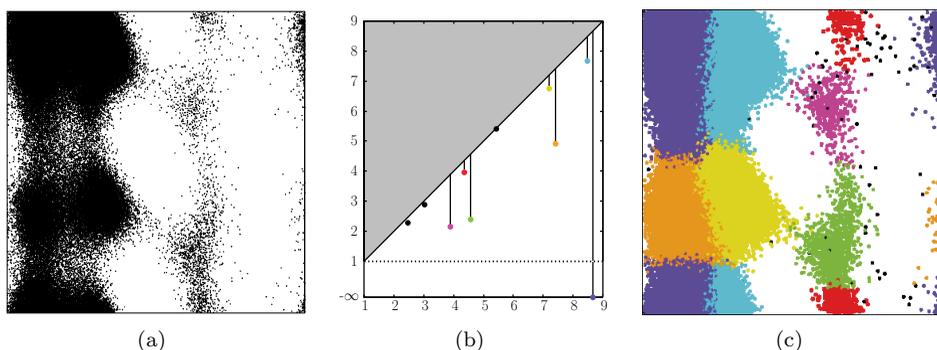


FIGURE 6.10. Biological data set: (a) input point cloud, projected down to the (ϕ, ψ) domain for visualization purposes; (b) output persistence diagram represented on a log-log scale with distinct colors for the 7 prominent peaks; (c) output clustering with 7 clusters and matching colors.

— From Chazal et al. [70].

4.2. Alanine-dipeptide conformations. The data consist of short trajectories of conformations generated by atomistic simulations of the alanine dipeptide [79]. Accurate simulation by molecular dynamics must be done at the atomic scale, generally limiting the length of simulations to picoseconds because of the small time steps needed to integrate stiff bond length and angle potentials. Biologically interesting dynamics, however, often occur on the scale of milliseconds. One solution to this issue is to generate a coarser model using *metastable* states [162]. These are conformational clusters between which transitions are infrequent and independent. Such coarser representations are tractable using Markovian models [78–80] while still allowing for useful simulations. A key problem is the discovery of these metastable states.

The alanine-dipeptide is a good example because its dynamics are relatively well-understood: it is a notorious fact that there are only two relevant degrees of freedom, and these are known *a priori*. This makes it possible to visualize the data and clustering results by projecting the points onto these coordinates which are referred to as ϕ and ψ (Ramachandran plots), as shown in Figure 6.10(a). In early work on the subject [79], clustering into 6 clusters was done manually. Subsequent work [78] tried to automatically recover these 6 clusters, as we do here using ToMATo.

Our input consists of 960 trajectories, each one made of 200 protein conformations, each conformation being represented as a 21-dimensional vector with 3 coordinates per atom of the protein backbone. For the experiments we take the trajectories and treat the conformations as 192,000 independent samples in \mathbb{R}^{21} . The metric used on this point cloud is the root-mean-squared deviation (RMSD) after the best possible rigid matching computed using the method of Theobald [231]. The RMSD distance matrix is our only input. The output is shown in Figure 6.10.

It appears from the persistence diagram that there could be anywhere between 4 and 7 clusters. Indeed, while the purple, orange, green and pink clusters are by far the most prominent, the following 3 clusters (blue, yellow and red) are

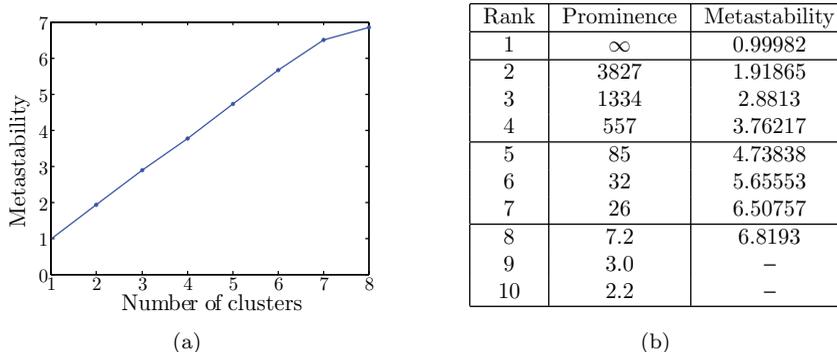


FIGURE 6.11. Quantitative evaluation of the quality of the output of ToMATo on the biological data set: (a) metastability of the obtained clustering versus the number of clusters; (b) corresponding intervals sorted by decreasing prominence.

— From Chazal et al. [70].

still much more prominent than the rest. To confirm this insight, let us come back to the original problem of finding clusters that maximize the metastability as defined in [162]: we compute the metastabilities of all our candidate clusterings, and we report them in the table and plot of Figure 6.11. These results show that the metastability increases linearly with the number of clusters, up to 7 clusters, after which it starts leveling off. So, choosing 4, 5, 6 or 7 clusters does not affect the metastability significantly, thus confirming the observations made from the persistence diagram. This is an example of a practical scenario in which the insights provided by the persistence diagram can be validated *a posteriori* by exploiting further application-specific information on the data.

5. Higher-dimensional structure

So far we have been interested in the connectivity of the data set, not on its higher-dimensional topology. Suppose one of the clusters has the shape of an annulus, as in the example of Figure 6.12. How can we extract this higher-level type of information? A natural approach is first to detect the clusters, and second to run the topological inference pipeline from the previous chapters on every cluster separately. Beside the aforementioned limitations inherent to the Hausdorff noise model used in topological inference, this approach depends critically on the quality of the initial clustering.

A more appealing approach is to generalize ToMATo so that it approximates the full persistence diagram $\text{dgm}(f)$, which captures the entire homological structure of the superlevel sets of the density f . To this end, assume the data points are sitting in \mathbb{R}^d , and recall that the δ -neighborhood graph G is the same as the 1-skeleton graph of the Rips complex $R_\delta(P)$, so the filtration $\tilde{\mathcal{G}}$ of G built by ToMATo is made of the 1-skeleton graph of $R_\delta(\tilde{F}^t \cap P)$ at every filtration level t , where \tilde{F}^t denotes the superlevel set of the estimator \tilde{f} . The idea is then to build the full Rips complex $R_\delta(\tilde{F}^t \cap P)$ at every level t , and to consider the resulting filtration

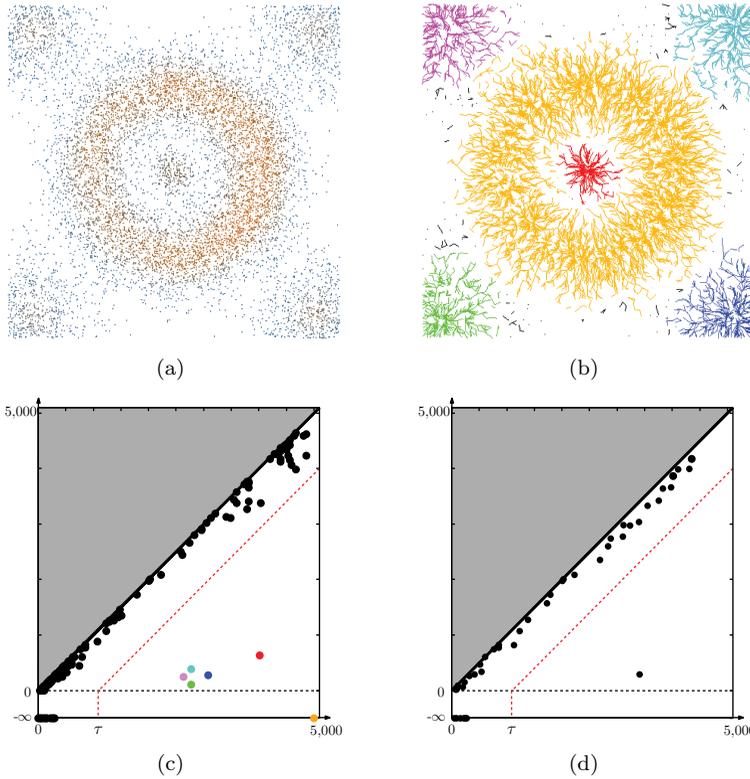


FIGURE 6.12. From the input point cloud equipped with estimated density values (a), the generalized ToMATo computes a clustering (b) together with approximations to the persistence diagrams of the density for all homology dimensions (0 for (c), 1 for (d)), thus revealing the presence of an annulus-shaped cluster. For better readability we used matching colors in (b) and (c).

$\mathcal{R}_\delta(P, \tilde{f})$ —which differs from the Rips filtration of P in that the Rips parameter δ remains fixed while the vertex set evolves with the function level.

As we saw in Section 3, under the hypotheses of Theorem 6.3 there is a partial interleaving between $\mathcal{R}_\delta(P, \tilde{f})$ and the superlevel-sets filtration \mathcal{F}_\geq of f at the 0-dimensional homology level. Do such interleavings exist at higher-dimensional homology levels? This question has been investigated by Chazal et al. [67]. Their answer is no in general because a single Rips complex $R_\delta(\tilde{F}^t \cap P)$ may fail to capture the full homological type of the corresponding superlevel set F^t of f (even when $\tilde{f} = f$), in the same way as a single Rips complex could fail to capture the full homological type of the space underlying the input data in the context of topological inference. However, as we saw in Section 3.1 of Chapter 5, a pair of Rips complexes may succeed. This observation led Chazal et al. [67] to using a pair of Rips complexes $R_\delta(\tilde{F}^t \cap P) \leftrightarrow R_{2\delta}(\tilde{F}^t \cap P)$ at every filtration level t , and studying the image of the induced morphism between persistence modules $H_*(\mathcal{R}_\delta(P, \tilde{f})) \rightarrow H_*(\mathcal{R}_{2\delta}(P, \tilde{f}))$ at the homology level. The image is a submodule

of $H_*(\mathcal{R}_{2\delta}(P, \tilde{f}))$, and its full persistence diagram can be computed in time cubic in the size of $R_{2\delta}(P)$ using the variant of the persistence algorithm for kernels, images and cokernels developed by Cohen-Steiner et al. [89]—recall Chapter 2, Sections 1.4 and 2.2. This is the generalized version of ToMATo.

Steps 1 through 3 of the proof of Theorem 6.3 unfold as before, except the interleaving worked out at step 2 is somewhat more technical to define and to analyze for the module $\text{im } H_*(\mathcal{R}_\delta(P, \tilde{f})) \rightarrow H_*(\mathcal{R}_{2\delta}(P, \tilde{f}))$ than it is for $H_0(\mathcal{R}_\delta(P, \tilde{f}))$. The technical details can be found in [67].

THEOREM 6.5 (Chazal et al. [67]). *Under the hypotheses of Theorem 6.3 (except the one on the separability of $\text{dgm}_0(f)$, which is irrelevant here), it holds with the same probability that the bottleneck distance between the full persistence diagrams $\text{dgm}(f)$ and $\text{dgm}(\text{im } H_*(\mathcal{R}_\delta(P, \tilde{f})) \rightarrow H_*(\mathcal{R}_{2\delta}(P, \tilde{f})))$ is at most $2c\delta + \eta$ in the upper-right quadrant $[2c\delta + \eta, +\infty] \times [2c\delta + \eta, +\infty]$.*

Thus, the generalized ToMATo is able to detect higher-dimensional structure in the data. The full persistence diagram of the density, as approximated by the algorithm, tells us something about (1) the topology of the individual clusters, and (2) their interconnectivity in the ambient space, although the two types of information are mixed up in the diagram and may require some extra work to be distinguished. Such knowledge is valuable in a range of practical contexts, for instance in the study of the conformation space of a protein, for which not only the metastable states but also the transitions between them matter.

On the downside, connections between the basins of attractions of the density f may happen in highly unstable areas or areas of very low density, and may therefore require humongous amounts of data points to be detected reliably. This is a general problem with higher-dimensional structures being more delicate than the clusters themselves. ToMATo does nothing particular to cope with it.

To conclude, let us mention that Theorem 6.5 has been generalized by Chazal et al. [67] to arbitrary Lipschitz continuous scalar fields (not just probability densities) under a sufficiently dense sampling of their domain. Clustering is but one of the possible applications of this result, and in fact, historically, ToMATo was introduced as a special (0-dimensional) case of a general method for analyzing scalar fields over point cloud data.

Signatures for Metric Spaces

In the previous chapters we were analyzing data sets independently of one another, putting the emphasis on discovering and characterizing their underlying geometric structure. Another way of getting insights into a data set is to compare it against other available data within a collection. This approach provides information not only on the structure of each data set taken separately, but also on the structure of the whole collection. This is all the more interesting as the recent years have seen great advances in the fields of data acquisition and simulation, and huge collections of digital data have emerged. With the goal of organizing these collections, there is a need for meaningful notions of *similarity* between data sets that exhibit invariance to different transformations of the objects they represent. Problems of this nature arise in areas such as molecular biology, metagenomics, face recognition, matching of articulated objects, graph matching, and pattern recognition in general.

Let us give a concrete example: in the context of 3d shape analysis, one typically wishes to be able to discriminate digital shapes under various notions of invariance. Many approaches have been proposed for the problem of (pose invariant) shape classification and recognition, including the work of Hilaga et al. [156], the *shape contexts* of Belongie, Malik, and Puzicha [21], the *integral invariants* of Manay et al. [185], the *eccentricity functions* of Hamza and Krim [148], the *shape distributions* of Osada et al. [207], the *canonical forms* of Elad and Kimmel [124], and the *shape DNA* and *global point signatures* based spectral methods of Reuter, Wolter, and Peinecke [215] and Rustamov [218], respectively. Their common underlying idea revolves around the computation and comparison of certain metric invariants, or *signatures*, so as to ascertain whether two given digital models represent in fact the same shape under a certain class of transformations.

More generally, signatures can be used to compare features across data, and to measure the amount by which two data sets (or their underlying objects) differ from each other. It is then desirable to devise families of signatures that are able to signal proximity or (dis-)similarity of data in a reasonable way. Although central, this question is rarely addressed from a formal viewpoint. In particular, the degree by which two data sets with similar signatures are forced to be similar is in general not well understood. Conversely, one can ask the more basic question of whether the similarity between two data sets forces their signatures to be similar. These questions cannot be completely well formulated until one agrees on: (1) a notion of *equality* between data sets, and (2) a notion of *(dis)similarity* between data sets. The same questions arise for the continuous objects underlying the data, when such objects exist.

Metric spaces and Gromov-Hausdorff distance. By regarding data sets as finite metric spaces, one can use the Gromov-Hausdorff distance [144] as a measure of dissimilarity between them. This distance generalizes the Hausdorff distance (4.1)

to pairs of sets sitting in different metric spaces, by reporting the smallest achievable Hausdorff distance through simultaneous isometric embeddings of the sets into a common space:

DEFINITION 7.1. The *Gromov-Hausdorff distance* between two metric spaces (X, d_X) and (Y, d_Y) is:

$$(7.1) \quad d_{\text{GH}}((X, d_X), (Y, d_Y)) = \inf_{Z, \gamma_X, \gamma_Y} d_{\text{H}}^Z(\gamma_X(X), \gamma_Y(Y)),$$

where Z ranges over all metric spaces and γ_X, γ_Y over all isometric embeddings of X, Y into (Z, d_Z) , and where d_{H}^Z denotes the Hausdorff distance in Z .

By endowing the considered data sets with different kinds of metrics, one obtains a great deal of flexibility in the various degrees of invariance that can be encoded in the measure of dissimilarity. For instance, for a point cloud sitting in \mathbb{R}^d , using the Euclidean metric makes the Gromov-Hausdorff distance invariant under ambient rigid isometries. In contrast, using intrinsic metrics within the object underlying the data makes the Gromov-Hausdorff distance blind to intrinsic isometries of that object, such as when a same animal is represented in different poses—see Figure 7.1. One can also combine intrinsic and extrinsic metrics to increase the discriminating power of the signatures, as suggested by Bronstein, Bronstein, and Kimmel [38] in the context of 3d shape comparison.

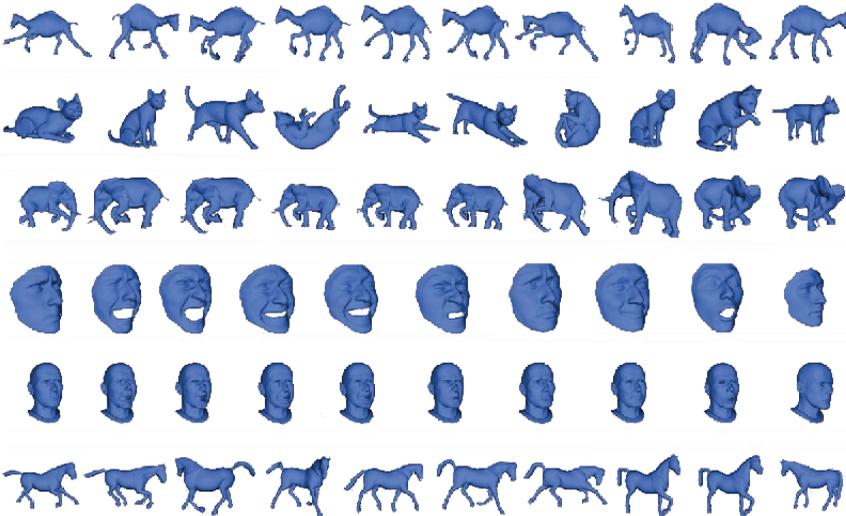


FIGURE 7.1. A database containing sixty 3d shapes divided into six classes.

— Based on Sumner and Popović [229].

Once the finite metric space view on data sets has been adopted, the challenge is to be able to estimate the Gromov-Hausdorff distance between finite metric spaces, either exactly or approximately. Since direct computation is notoriously difficult, leading to hard combinatorial problems akin to the *quadratic assignment problem* (which is NP hard), easily computable alternatives such as distances between signatures must be considered. The question becomes then to understand

how a given family of signatures behaves under perturbations of the data in the Gromov-Hausdorff distance. Formally, given such a family S together with a metric d between the signatures, the goal is to evaluate how d relates to d_{GH} , which means working out tight bounds $b, B : \mathbb{R}^+ \rightarrow \mathbb{R}^+$ such that the following inequalities

$$b(d_{\text{GH}}((X, d_X), (Y, d_Y))) \leq d(s(X, d_X), s(Y, d_Y)) \leq B(d_{\text{GH}}((X, d_X), (Y, d_Y)))$$

hold for any finite metric spaces (X, d_X) and (Y, d_Y) and any signature $s \in S$. Intuitively, the lower bound b measures the discriminative power of the signatures, while the upper bound B measures their stability. Ideally, a positive b and a finite B would be desired, to guarantee that replacing the metric spaces by their signatures is harmless. However such guarantees remain generally out of reach, so it is common to reduce one's expectations and merely ask that B be finite (in fact linear), so the signatures are at least guaranteed to be stable.

Persistence based signatures. In this context, persistence provides a theoretically sound way of devising signatures for finite metric spaces that are stable with respect to Gromov-Hausdorff perturbations of the spaces. This approach was first investigated in the context of *size theory* [96, 97], exclusively for 0-dimensional homology and for a restricted class of metric space transformations. The modern view of persistence makes it possible to generalize the approach to higher-dimensional homology and arbitrary transformations of the metric spaces. This new line of work was initiated by Carlsson et al. [47].

The ideas guiding the construction of persistence based signatures come from the work on topological inference presented in Chapters 4 and 5. When the considered data set X is a finite point cloud sampled from a compact set K in Euclidean space, the Euclidean distance to X is known to approximate the distance to K . Then, the persistence diagram of the distance to X serves as a proxy for the one of the distance to K . This is a very specific instance of our problem, in which two metric spaces— X and K —that are nearby in the (Gromov-)Hausdorff distance have nearby signatures—their persistence diagrams—in the bottleneck distance.

For a general finite metric space (X, d_X) , the lack of a proper ambient space makes the persistence diagram of the distance function meaningless. However, we can still use some of the combinatorial constructions from Chapter 5 instead, in particular the Rips filtration, whose definition depends on the metric put on X but not on the ambient space from which this metric is induced. It turns out that the Rips construction can be adapted to arbitrary metric spaces, whether finite or infinite. Furthermore, so do the constructions of the Čech and witness filtrations. Thus, their persistence diagrams, whenever defined, can be used as signatures for general metric spaces. This is the subject of Section 1.

Stability. Persistence based signatures can be proven stable with respect to perturbations of the metric spaces in the Gromov-Hausdorff distance. This result is a direct consequence of the stability part of the Isometry Theorem 3.1 once an appropriate interleaving between filtrations has been worked out. This is the subject of Sections 2 and 3. Note that deriving stability results for a general class of metric spaces and not just finite spaces is useful in applications, e.g. to guarantee that the signatures computed from finite samplings of a continuous space (e.g. a 3d shape) are meaningful in the sense that they approximate the signatures of that space.

Interestingly enough, the proof of stability in the case where the space (X, d_X) is finite is simple and has a nice geometric interpretation. In the general case however, the proof has a more combinatorial flavor and requires a condition of *total boundedness* on (X, d_X) . This means that, for every $\varepsilon > 0$, (X, d_X) admits a finite ε -*sample*, i.e. a finite subset P_ε such that $d_H(P_\varepsilon, X) \leq \varepsilon$. In particular, every compact metric space is totally bounded. Such spaces are approximable at every resolution by finite metric spaces, which explains their good behavior with respect to persistence. Beyond this class of spaces, the persistence based signatures can behave wildly, as we will see in Section 3.3.

These stability results have close connections to the line of work initiated by Hausmann [154], whose aim is to study the properties of single Rips complexes on various types of metric spaces. When (X, d_X) is a closed Riemannian manifold, Hausmann [154] proved that if $i > 0$ is sufficiently small then the Rips complex of X of parameter i is homotopy equivalent to X . This result was later generalized by Latschev [175], who proved that if (Y, d_Y) is sufficiently close to (X, d_X) in the Gromov-Hausdorff distance, then there exists $i > 0$ such that the Rips complex of Y of parameter i is homotopy equivalent to X . Recently, Attali, Lieutier, and Salinas [9] adapted these results to a class of sufficiently regular compact subsets of Euclidean spaces. For larger classes of compact subsets of Riemannian manifolds, the homology and homotopy of such sets are encoded in nested pairs of Rips complexes as we saw in Chapter 5, yet it remains still open whether or not a single Rips complex can carry this topological information.

Metric spaces equipped with functions. The construction of persistence based signatures extends naturally to metric spaces equipped with functions, without any extra conditions on the spaces. The previous stability guarantees adapt as well, under an extended notion of Gromov-Hausdorff distance. This is the subject of Section 4.

Computing the signatures. As we saw in Chapter 5, the full computation of Rips or Čech or witness filtrations becomes quickly intractable in practice. Fortunately, there are ways to reduce the computational cost significantly, at the price of a (controlled) degradation in the quality of the signatures. Section 5 presents two such ways: one is based on truncating the filtrations (which means computing simplified signatures), the other is based on building sparsified yet interleaved filtrations (which means computing approximate signatures). In both cases the stability properties of the signatures are maintained.

Prerequisites. Most of the material used in the chapter comes from Part 1. We will be assuming familiarity with simplicial complexes, simplicial maps, and the concept of contiguity between such maps—see Chapter 1 of [202] for an introduction to these topics.

A concrete example. Before diving into the technical details, and as an appetizer for the reader, let us give an example of application of the persistence based signatures to shape classification. This example comes from [58].

Figure 7.2 shows a few toy shapes and some approximations to their persistence based signatures, computed from uniform samplings. Several interesting observations can be made:

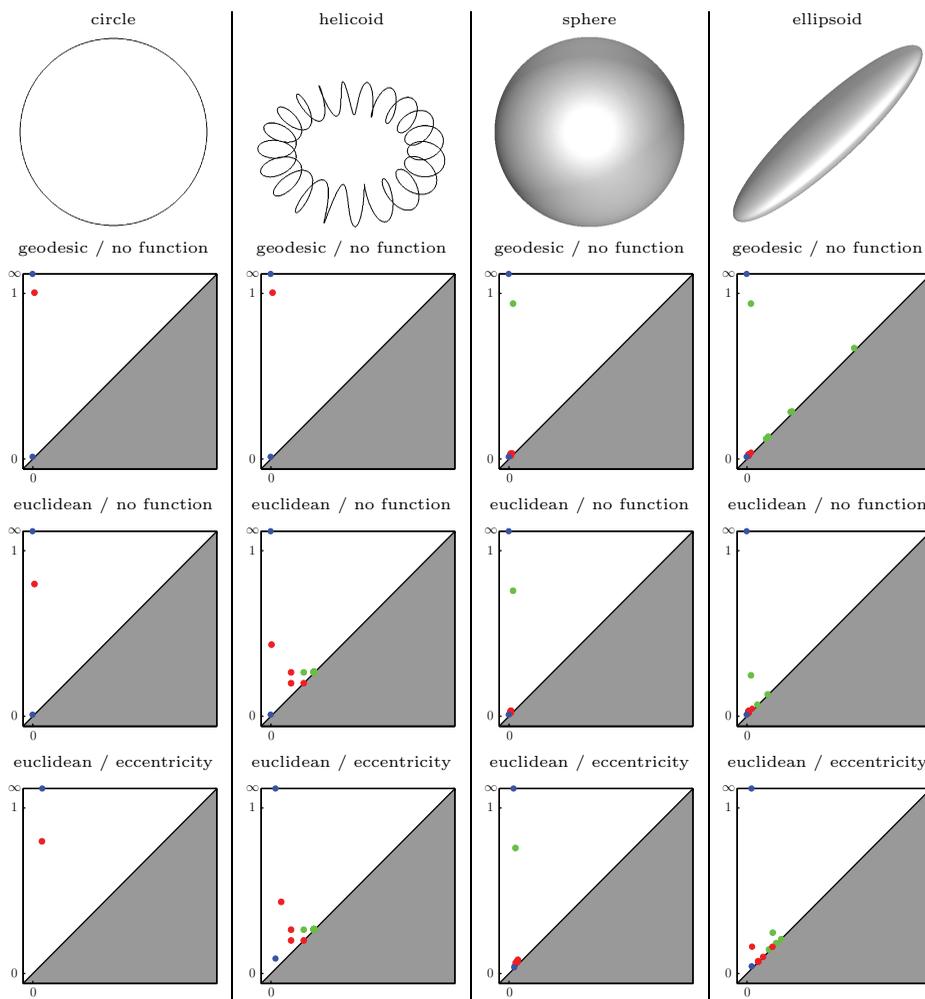


FIGURE 7.2. Some toy examples, from left to right: the unit circle, a helical curve of same length 2π drawn on a torus, the unit sphere, an ellipsoid whose smallest equatorial ellipse has same length 2π as the equator of the sphere. On each shape, a uniform 0.0125-sample has been generated, on top of which various Rips filtrations have been constructed using different metrics and functions: geodesic distance and no function (second row), Euclidean distance and no function (third row), Euclidean distance and 0.2 times the eccentricity (fourth row). The corresponding persistence diagrams are presented in rows 2 through 4, with 0-dimensional features marked as blue points, 1-dimensional features as red points, and 2-dimensional features as green points.

— From Chazal et al. [58].

- The first two shapes (circle and helicoidal curve on a torus) have the same length and are therefore intrinsically isometric. As a result, signatures

obtained from geodesic distances are identical whereas signatures based on Euclidean distances differ. This illustrates the importance of the choice of metric in practice.

- All the shapes have been ε -sampled for a same value ε , therefore all the zero-dimensional diagrams look the same in rows 2 and 3. Discrimination between the shapes is made possible by the higher-dimensional diagrams, which illustrates the advantage of using persistence over size theory.
- The sphere and ellipsoid cannot be discriminated using their Rips filtrations alone, because their diagrams are the same up to rescaling and small noise (rows 2 and 3). By equipping the shapes with functions such as the eccentricity, and by computing the corresponding signatures, we can discriminate between the two shapes. The reason is that the 1- and 2-dimensional diagrams of the ellipsoid are significantly affected by the eccentricity, whereas the ones for the sphere hardly change due to the eccentricity being constant over it (row 4).

Let us now turn to a real shape classification problem. The input is the database of triangulated shapes shown in Figure 7.1. This is an excerpt from the full database gathered by Sumner and Popović [229]. It comprises 60 shapes from six different classes: `camel`, `cat`, `elephant`, `face`, `head` and `horse`. Each class contains 10 different poses of the same shape. These poses are richer than just rigid isometries. The number of vertices in the models ranges from $7K$ to $30K$.

Let us normalize each model X_i and equip it with an approximation d_i of its geodesic distance, computed using Dijkstra's algorithm on the 1-skeleton graph of the mesh¹. Let us then compute a persistence based signature for the metric space (X_i, d_i) using a truncated Rips filtration. The signatures are compared against one another in the bottleneck distance. The resulting distance matrix M is shown in Figure 7.3 (left). In order to evaluate the discriminative power contained in M , let us consider a classification task as follows: we randomly select one shape from each class, form a training set T and use it for performing 1-nearest neighbor classification (where *nearest* is with respect to the metric defined by M) on the remaining shapes. By simple comparison between the class predicted by the classifier and the actual class to which the shape belongs, we obtain an estimate of the probability $P_e(M)$ of mis-classification. We repeat this procedure for $2K$ random choices of the training set. Using the same randomized procedure we obtain an estimate of the *confusion matrix* for this problem, whose entry (i, j) is the probability that the classifier will assign class j to a shape when the actual class is i — see Figure 7.3 (center). We obtain $P_e(M) = 2\%$. As Figure 7.3 (right) shows, the classes are well separated in signature space, which explains the low mis-classification probability.

1. Simplicial filtrations for arbitrary metric spaces

We need to extend the simplicial filtrations introduced in Chapters 4 and 5 to arbitrary metric spaces, so that we can use their persistence diagrams (whenever defined) as signatures. Throughout the section the metric space (X, d_X) is arbitrary.

¹In this particular application we take advantage of the fact that a mesh is available for each of the shapes. In full generality, it can be replaced by some neighborhood graph in the ambient Euclidean metric, with similar guarantees.

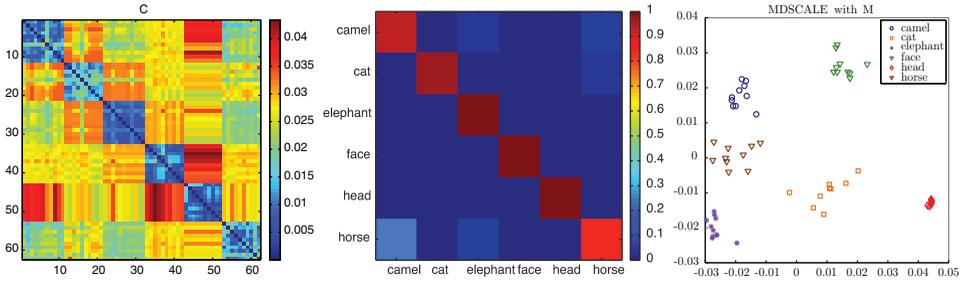


FIGURE 7.3. Left: estimate of the Gromov-Gausdorff distance computed on the database of Figure 7.1. Center: estimated confusion matrix for the 1-nearest neighbor classification problem. Right: MDS plot of the matrix M with labels corresponding to each class. The overall error is estimated to be 2%.

— From Chazal et al. [58].

1.1. (Vietoris-)Rips filtration. For $i \in \mathbb{R}$, we define a simplicial complex $R_i(X, d_X)$ on the vertex set X by the following condition:

$$\sigma \in R_i(X, d_X) \iff d_X(x, y) \leq i \text{ for all } x, y \in \sigma.$$

Note that $R_i(X, d_X)$ is empty for $i < 0$ and consists of the vertex set X alone for $i = 0$. There is a natural inclusion $R_i(X, d_X) \subseteq R_j(X, d_X)$ whenever $i \leq j$. Thus, the simplicial complexes $R_i(X, d_X)$ together with these inclusion maps define a simplicial filtration $\mathcal{R}(X, d_X)$ with vertex set X , called the (Vietoris-)Rips filtration of (X, d_X) .

1.2. Ambient and intrinsic Čech filtrations. Let L, W be subsets (‘landmarks’ and ‘witnesses’) of X . For $i \in \mathbb{R}$, consider the complex with vertices L and simplices determined by:

$$\sigma \in C_i(L, W, d_X) \iff \exists w \in W \text{ such that } d_X(w, l) \leq i \text{ for all } l \in \sigma.$$

The resulting simplicial filtration is denoted by $\mathcal{C}(L, W, d_X)$ and called the ambient Čech filtration. The intrinsic Čech filtration of (X, d_X) is $\mathcal{C}(X, X, d_X)$, also denoted by $\mathcal{C}(X, d_X)$ for simplicity.

1.3. Witness filtration. Let L, W be once again subsets of X . For any finite subset $\sigma \subseteq L$, and any $w \in W$ and $i \in \mathbb{R}$, we say that w is an i -witness for the simplex σ if

$$d_X(w, l) \leq d_X(w, l') + i \text{ for all } l \in \sigma \text{ and } l' \in L \setminus \sigma.$$

We can then define for any $i \in \mathbb{R}$ a simplicial complex $W_i(L, W, d_X)$ by

$$\sigma \in W_i(L, W, d_X) \iff \forall \tau \subseteq \sigma, \exists w \in W \text{ such that } w \text{ is an } i\text{-witness for } \tau.$$

An i -witness is obviously a j -witness for any $j \geq i$, so there is a natural inclusion $W_i(L, W, d_X) \subseteq W_j(L, W, d_X)$. The simplicial complexes $W_i(L, W, d_X)$ together with these inclusion maps define a simplicial filtration $\mathcal{W}(L, W, d_X)$ with vertex set L , called the witness filtration. Note that this filtration has nontrivial behaviour for $i < 0$, unlike the Rips and Čech filtrations.

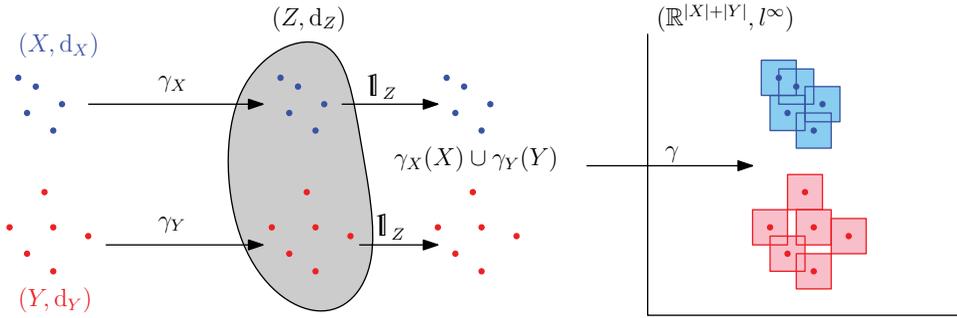


FIGURE 7.4. Overview of the proof of Theorem 7.2. The input metric spaces (X, d_X) and (Y, d_Y) are mapped isometrically to a common space (Z, d_Z) , then to $(\mathbb{R}^{|X|+|Y|}, \ell^\infty)$ by Lemma 7.3. There, the point sets are Hausdorff close, therefore their distance functions are close in the supremum norm, and their persistence diagrams are close in the bottleneck distance. These diagrams are also the ones of the nerves of the unions of ℓ^∞ -balls that form the sublevel sets of the distance functions. These nerves happen to be also Rips complexes, by Lemma 7.4.

2. Stability for finite metric spaces

In the special case of finite metric spaces, the proof of stability for Rips based signatures is simple and has a nice geometric interpretation—illustrated in Figure 7.4. We therefore reproduce it in full extent below. Signatures based on Čech or witness filtrations enjoy similar guarantees, however their analysis turns out to be more subtle. This is somehow related to the way they interact with the ambient space, as appears from Section 1. The only stability proofs for these signatures that we are aware of use the machinery developed in Section 3 for more general metric spaces, so let us defer their treatment until then.

THEOREM 7.2. *For any finite metric spaces (X, d_X) and (Y, d_Y) ,*

$$(7.2) \quad d_b(\text{dgm}(\mathcal{R}(X, d_X)), \text{dgm}(\mathcal{R}(Y, d_Y))) \leq 2 d_{\text{GH}}((X, d_X), (Y, d_Y)).$$

The proof relies on the following well-known embedding for finite metric spaces—see e.g. example 3.5.3 and exercise 3.5.4 in [44]:

LEMMA 7.3. *Any finite metric space of cardinality n can be isometrically embedded into $(\mathbb{R}^n, \ell^\infty)$.*

The proof of Theorem 7.2 relies also on a remarkable equality between the Čech and Rips filtrations of subsets of $(\mathbb{R}^n, \ell^\infty)$. This equality holds up to a rescaling of the Rips parameter by a factor of 2, and it is to be compared with the tight interleaving (5.4) obtained in the ℓ^2 -norm:

LEMMA 7.4 (Ghrist and Muhammad [140]). *For any $X \subseteq \mathbb{R}^n$ and $i \in \mathbb{R}$,*

$$C_i(X, \mathbb{R}^n, \ell^\infty) = R_{2i}(X, \ell^\infty).$$

PROOF OF THEOREM 7.2—SEE FIGURE 7.4 FOR A PICTORIAL OVERVIEW.

Let $\varepsilon > d_{\text{GH}}((X, d_X), (Y, d_Y))$. By Definition 7.1, there is a metric space (Z, d_Z)

and two isometric embeddings $\gamma_X : X \rightarrow Z$ and $\gamma_Y : Y \rightarrow Z$ such that

$$d_{\mathbb{H}}^Z(\gamma_X(X), \gamma_Y(Y)) \leq \varepsilon.$$

Equip the subset $\gamma_X(X) \sqcup \gamma_Y(Y)$ with the induced metric d_Z . This new metric space is finite, therefore it can be embedded isometrically into $(\mathbb{R}^n, \ell^\infty)$, where $n = |X| + |Y|$, by Lemma 7.3. Let γ be the isometric embedding. We then have

$$d_{\mathbb{H}}^{\mathbb{R}^n}(\gamma \circ \gamma_X(X), \gamma \circ \gamma_Y(Y)) = d_{\mathbb{H}}^Z(\gamma_X(X), \gamma_Y(Y)) \leq \varepsilon.$$

Hence, inside $(\mathbb{R}^n, \ell^\infty)$, the distance function δ_X to $\gamma \circ \gamma_X(X)$ and the distance function δ_Y to $\gamma \circ \gamma_Y(Y)$ are ε -close in the supremum norm, so their sublevel-sets filtrations are ε -interleaved in the sense of (3.3). Moreover, they are \mathfrak{q} -tame by Proposition 2.3 (ii), therefore their persistence diagrams are well-defined. The stability part of the Isometry Theorem 3.1 implies then that

$$(7.3) \quad d_b(\mathbf{dgm}(\delta_X), \mathbf{dgm}(\delta_Y)) \leq \varepsilon.$$

Now, for all $i \in \mathbb{R}$, the i -sublevel set of δ_X is the union of the closed ℓ^∞ -balls of same radius i about the points of $\gamma \circ \gamma_X(X)$. Since ℓ^∞ -balls are hypercubes, they are convex and therefore their intersections are either empty or contractible. Hence, the Persistent Nerve Lemma 4.12 applies² and ensures that

$$(7.4) \quad \begin{aligned} \mathbf{dgm}(\delta_X) &= \mathbf{dgm}(\mathcal{C}(\gamma \circ \gamma_X(X), \mathbb{R}^n, \ell^\infty)), \\ \mathbf{dgm}(\delta_Y) &= \mathbf{dgm}(\mathcal{C}(\gamma \circ \gamma_Y(Y), \mathbb{R}^n, \ell^\infty)). \end{aligned}$$

In addition, Lemma 7.4 guarantees that

$$(7.5) \quad \begin{aligned} \mathcal{C}(\gamma \circ \gamma_X(X), \mathbb{R}^n, \ell^\infty) &= \bar{\mathcal{R}}(\gamma \circ \gamma_X(X), \ell^\infty), \\ \mathcal{C}(\gamma \circ \gamma_Y(Y), \mathbb{R}^n, \ell^\infty) &= \bar{\mathcal{R}}(\gamma \circ \gamma_Y(Y), \ell^\infty), \end{aligned}$$

where $\bar{\mathcal{R}}$ denotes the Rips filtration with indices rescaled by a factor of 2. The rescaling implies

$$(7.6) \quad \begin{aligned} d_b(\mathbf{dgm}(\bar{\mathcal{R}}(\gamma \circ \gamma_X(X), \ell^\infty)), \mathbf{dgm}(\bar{\mathcal{R}}(\gamma \circ \gamma_Y(Y), \ell^\infty))) &= \\ 2 d_b(\mathbf{dgm}(\bar{\mathcal{R}}(\gamma \circ \gamma_X(X), \ell^\infty)), \mathbf{dgm}(\bar{\mathcal{R}}(\gamma \circ \gamma_Y(Y), \ell^\infty))). \end{aligned}$$

Finally, since $\gamma \circ \gamma_X$ and $\gamma \circ \gamma_Y$ are isometric embeddings, we have

$$(7.7) \quad \mathcal{R}(\gamma \circ \gamma_X(X), \ell^\infty) = \mathcal{R}(X, d_X) \quad \text{and} \quad \mathcal{R}(\gamma \circ \gamma_Y(Y), \ell^\infty) = \mathcal{R}(Y, d_Y).$$

Combining (7.3) through (7.7) together, we obtain

$$d_b(\mathbf{dgm}(\mathcal{R}(X, d_X)), \mathbf{dgm}(\mathcal{R}(Y, d_Y))) \leq 2\varepsilon.$$

Since ε was chosen greater than $d_{\text{GH}}((X, d_X), (Y, d_Y))$ but arbitrarily close to it in the first place, we conclude that

$$d_b(\mathbf{dgm}(\mathcal{R}(X, d_X)), \mathbf{dgm}(\mathcal{R}(Y, d_Y))) \leq 2 d_{\text{GH}}((X, d_X), (Y, d_Y))$$

as desired. □

Let us emphasize that the lower bound on the Gromov-Hausdorff distance given in (7.2) is worst-case tight. For instance, take for X a set of two points at distance 2 and for Y a set of two points at distance $2 + 2\varepsilon$. Then, (X, d_X) and (Y, d_Y) can be isometrically embedded into the real line, with X mapped to $\{0, 2\}$ and Y mapped to $\{-\varepsilon, 2 + \varepsilon\}$, which shows that their Gromov-Hausdorff distance is at

²The conditions under which this lemma applies to closed covers are satisfied here, since nonempty intersections of hypercubes are convex and therefore neighborhood retracts in \mathbb{R}^n .

most ε . Now, the persistence diagram of the Rips filtration of (X, d_X) is made of two points, namely $(0, +\infty)$ and $(0, 2)$, while the persistence diagram of the Rips filtration of (Y, d_Y) is made of $(0, +\infty)$ and $(0, 2 + 2\varepsilon)$, hence their bottleneck distance is 2ε .

3. Stability for totally bounded metric spaces

Theorem 7.2 extends to totally bounded metric spaces as follows. The worst-case tightness of the bound in (7.8) follows from the one in (7.2).

THEOREM 7.5. *For any totally bounded metric spaces (X, d_X) and (Y, d_Y) , the Rips filtrations $\mathcal{R}(X, d_X)$ and $\mathcal{R}(Y, d_Y)$ are \mathfrak{q} -tame, and*

$$(7.8) \quad d_b(\text{dgm}(\mathcal{R}(X, d_X)), \text{dgm}(\mathcal{R}(Y, d_Y))) \leq 2 d_{\text{GH}}((X, d_X), (Y, d_Y)).$$

The proof is more elaborate than in the finite case, however it better emphasizes the mechanisms at work, therefore we also include it here³. It relies on a different but equivalent definition of the Gromov-Hausdorff distance, introduced in Section 3.1 and based on correspondences between metric spaces. The concept of a correspondence leads quite naturally to an interleaving between Rips filtrations that holds in a remarkably general setting, as shown in Section 3.2. One can obtain (7.8) as an immediate consequence of that interleaving provided that the persistence diagrams of the Rips filtrations are well-defined. It is indeed not always the case, but at least for totally bounded metric spaces it is, as shown in Section 3.3. The analysis can be extended to work with signatures based on Čech and witness filtrations, as discussed in Section 3.4.

3.1. Gromov-Hausdorff distance via correspondences. Correspondences are best explained through the concept of multivalued maps.

Multivalued maps. In short, a multivalued map is a ‘one-to-many’ function.

DEFINITION 7.6. A *multivalued map* $C : X \rightrightarrows Y$ from a set X to a set Y is a subset of $X \times Y$, also denoted C , that projects surjectively onto X through the canonical projection $\pi_X : X \times Y \rightarrow X$. The *image* $C(\sigma)$ of a subset σ of X is the canonical projection of $\pi_X^{-1}(\sigma)$ onto Y .

A (single-valued) map f from X to Y is *subordinate* to C if we have $(x, f(x)) \in C$ for every $x \in X$. In that case we write $f : X \xrightarrow{C} Y$. The *composition* of two multivalued maps $C : X \rightrightarrows Y$ and $D : Y \rightrightarrows Z$ is the multivalued map $D \circ C : X \rightrightarrows Z$ defined by:

$$(x, z) \in D \circ C \Leftrightarrow \exists y \in Y \text{ such that } (x, y) \in C \text{ and } (y, z) \in D.$$

The *transpose* of C , denoted C^T , is the image of C through the symmetry map $(x, y) \mapsto (y, x)$. Although C^T is well-defined as a subset of $Y \times X$, it is not always a multivalued map because it may not project surjectively onto Y .

³The uncomfortable reader may safely skip Sections 3.2 and 3.3.

Correspondences. These are multivalued maps whose graph projects surjectively onto the codomain.

DEFINITION 7.7. A multivalued map $C : X \rightrightarrows Y$ is a *correspondence* if the canonical projection $C \rightarrow Y$ is surjective, or equivalently, if C^T is also a multivalued map.

We immediately deduce, if C is a correspondence, that the identity maps $\mathbb{1}_X$ and $\mathbb{1}_Y$ are subordinate respectively to the compositions $C^T \circ C$ and $C \circ C^T$:

$$(7.9) \quad \mathbb{1}_X : X \xrightarrow{C^T \circ C} X, \quad \mathbb{1}_Y : Y \xrightarrow{C \circ C^T} Y.$$

Gromov-Hausdorff distance. Given two metric spaces (X, d_X) , (Y, d_Y) , the *metric distortion* induced by a correspondence $C : X \rightrightarrows Y$ is defined as the quantity

$$(7.10) \quad \text{dist}_m(C) = \sup_{(x,y),(x',y') \in C} |d_X(x, x') - d_Y(y, y')|.$$

The Gromov-Hausdorff distance $d_{\text{GH}}((X, d_X), (Y, d_Y))$ is then defined as half the smallest possible metric distortion induced by correspondences $X \rightrightarrows Y$:

$$(7.11) \quad d_{\text{GH}}((X, d_X), (Y, d_Y)) = \frac{1}{2} \inf_{C: X \rightrightarrows Y} \text{dist}_m(C).$$

This quantity is known to be equivalent to (7.1)—see e.g. Theorem 7.3.25 in [44].

3.2. Interleavings. The key to proving Theorem 7.5 is the following bound on the interleaving distance, which holds for general metric spaces:

PROPOSITION 7.8. *For any metric spaces (X, d_X) and (Y, d_Y) ,*

$$d_i(\mathbf{H}_*(\mathcal{R}(X, d_X)), \mathbf{H}_*(\mathcal{R}(Y, d_Y))) \leq 2 d_{\text{GH}}((X, d_X), (Y, d_Y)),$$

where d_i denotes the interleaving distance from Chapter 3.

PROOF. Given any $\varepsilon > 2 d_{\text{GH}}((X, d_X), (Y, d_Y))$, let $C : X \rightrightarrows Y$ be a correspondence such that $\text{dist}_m(C) \leq \varepsilon$. Then, it follows easily from (7.10) that any choice of subordinate map $f : X \xrightarrow{C} Y$ induces a simplicial map $R_i(X, d_X) \rightarrow R_{i+\varepsilon}(Y, d_Y)$ at each $i \in \mathbb{R}$, and that these maps commute with the inclusions $R_i(X, d_X) \hookrightarrow R_j(X, d_X)$ and $R_{i+\varepsilon}(Y, d_Y) \hookrightarrow R_{j+\varepsilon}(Y, d_Y)$ for all $i \leq j$. Thus, f induces a degree- ε morphism $\phi \in \text{Hom}^\varepsilon(\mathbf{H}_*(\mathcal{R}(X, d_X)), \mathbf{H}_*(\mathcal{R}(Y, d_Y)))$ as per (3.4).

Now, any two subordinate maps $f_1, f_2 : X \xrightarrow{C} Y$ induce simplicial maps $R_i(X, d_X) \rightarrow R_{i+\varepsilon}(Y, d_Y)$ that are contiguous for all $i \in \mathbb{R}$, therefore their induced degree- ε morphisms ϕ_1, ϕ_2 are equal [202, theorems. 12.4 and 12.5]. Thus, the correspondence C induces a canonical degree- ε morphism

$$\phi \in \text{Hom}^\varepsilon(\mathbf{H}_*(\mathcal{R}(X, d_X)), \mathbf{H}_*(\mathcal{R}(Y, d_Y))).$$

In the same way, its transpose C^T induces another canonical degree- ε morphism

$$\psi \in \text{Hom}^\varepsilon(\mathbf{H}_*(\mathcal{R}(Y, d_Y)), \mathbf{H}_*(\mathcal{R}(X, d_X))).$$

The compositions $C^T \circ C$ and $C \circ C^T$ induce then respectively

$$\begin{aligned} \psi \circ \phi &\in \text{Hom}^{2\varepsilon}(\mathbf{H}_*(\mathcal{R}(X, d_X)), \mathbf{H}_*(\mathcal{R}(X, d_X))) \text{ and} \\ \phi \circ \psi &\in \text{Hom}^{2\varepsilon}(\mathbf{H}_*(\mathcal{R}(Y, d_Y)), \mathbf{H}_*(\mathcal{R}(Y, d_Y))), \end{aligned}$$

which, by (7.9), are equal respectively to the 2ε -shift morphisms $\mathbb{1}_{H_*(R(X, d_X))}^{2\varepsilon}$ and $\mathbb{1}_{H_*(R(Y, d_Y))}^{2\varepsilon}$. Hence, $H_*(R(X, d_X))$ and $H_*(R(Y, d_Y))$ are ε -interleaved as per Definition 3.3 (rephrased). Since ε was chosen greater than $2 d_{\text{GH}}((X, d_X), (Y, d_Y))$ but arbitrarily close to it in the first place, the conclusion of the proposition follows. \square

3.3. Tameness of the Rips filtration. With Proposition 7.8 at hand, proving Theorem 7.5 becomes merely a matter of showing that the Rips filtrations are \mathfrak{q} -tame (see below), then applying the stability part of the Isometry Theorem 3.1.

PROPOSITION 7.9. *The Rips filtration of a totally bounded metric space is always \mathfrak{q} -tame.*

PROOF. Let (X, d_X) be a totally bounded metric space. We must show that the linear map $v_i^j : H_*(R_i(X, d_X)) \rightarrow H_*(R_j(X, d_X))$ induced by the inclusion $R_i(X, d_X) \hookrightarrow R_j(X, d_X)$ has finite rank whenever $i < j$. Let $\varepsilon = \frac{j-i}{2} > 0$. Since X is totally bounded, there exists a finite $\frac{\varepsilon}{2}$ -sample P of X . The set

$$C = \left\{ (x, p) \in X \times P \mid d_X(x, p) \leq \frac{\varepsilon}{2} \right\}$$

is then the graph of a correspondence of metric distortion at most ε , so Proposition 7.8 guarantees that there exists an ε -interleaving between $H_*(\mathcal{R}(X, d_X))$ and $H_*(\mathcal{R}(P, d_X))$. Using the interleaving maps, v_i^j factorizes as

$$H_*(R_i(X, d_X)) \rightarrow H_*(R_{i+\varepsilon}(P, d_X)) \rightarrow H_*(R_{i+2\varepsilon}(X, d_X)) = H_*(R_j(X, d_X)).$$

The second space in the sequence is finite-dimensional since $R_{i+\varepsilon}(P, d_X)$ is a finite simplicial complex. Therefore, v_i^j has finite rank. \square

Note that \mathfrak{q} -tameness is the best one can hope for on totally bounded metric spaces. Indeed, it is easy to construct an example of a compact metric space (X, d_X) such that the homology group $H_1(R_1(X, d_X))$ has an infinite (even uncountable) dimension. For instance, consider the union X of two parallel segments in \mathbb{R}^2 defined by

$$X = \{(x, 0) \in \mathbb{R}^2 \mid x \in [0, 1]\} \cup \{(x, 1) \in \mathbb{R}^2 \mid x \in [0, 1]\}$$

and equipped with the Euclidean metric. Then, for any $x \in [0, 1]$, the edge $e_x = [(x, 0), (x, 1)]$ belongs to $R_1(X, d_X)$ but there is no triangle in $R_1(X, d_X)$ that contains e_x in its boundary. As a consequence, for $x \in (0, 1]$ the cycles $\gamma_x = [(0, 0), (x, 0)] + e_x + [(x, 1), (0, 1)] - e_0$ are not homologous to 0 and are linearly independent in $H_1(R_1(X, d_X))$. Thus, $H_1(R_1(X, d_X))$ has uncountable dimension. Note that 1 is the only value of the Rips parameter i for which the homology group $H_1(R_i(X, d_X))$ fails to be finite-dimensional in this example. In fact, it is possible to construct examples of compact metric spaces on which the set of ‘bad’ values is arbitrarily large, even though the filtration itself remains \mathfrak{q} -tame by virtue of Proposition 7.9—see [64] for such an example.

When the space (X, d_X) is not totally bounded, the Rips filtration can behave wildly, in particular it may cease to be \mathfrak{q} -tame. For instance, take the space

$$X = \{y^2 \mid y \in \mathbb{N}\} \subset \mathbb{N}$$

equipped with the absolute-value metric. Then, for any $i \geq 0$, the points $x \in X$ such that $x \geq \left(\frac{i+1}{2}\right)^2$ form independent connected components in $R_i(X, d_X)$. Hence, all

the linear maps in the persistence module $H_0(\mathcal{R}(X, d_X))$ have infinite rank beyond index $i = 0$.

3.4. Extension to Čech and witness based signatures. Theorem 7.5 extends to Čech and witness filtrations in a fairly straightforward manner. The proof of \mathfrak{q} -tameness (Proposition 7.9) is in fact the same, while the proof of ε -interleaving (Proposition 7.8) is roughly the same except that the initial step is slightly more subtle. This step claims that any linear map subordinate to a correspondence $C : X \rightrightarrows Y$ of metric distortion at most ε induces a simplicial map between the complexes considered. While this is obvious for Rips filtrations, it is somewhat less so for Čech and witness filtrations, even though the proof still relies on simple applications of the triangle inequality. Let us refer the reader to [64] for the details, and state the results without proofs:

THEOREM 7.10. *Let (X, d_X) and (Y, d_Y) be totally bounded metric spaces, and let $L, L', W, W' \subseteq X$. Then,*

(i) *The intrinsic Čech filtrations $\mathcal{C}(X, d_X)$ and $\mathcal{C}(Y, d_Y)$ are \mathfrak{q} -tame, and*

$$d_b(\text{dgm}(\mathcal{C}(X, d_X)), \text{dgm}(\mathcal{C}(Y, d_Y))) \leq 2 d_{\text{GH}}((X, d_X), (Y, d_Y)).$$

(ii) *The ambient Čech filtrations $\mathcal{C}(L, W, d_X)$ and $\mathcal{C}(L', W', d_X)$ are \mathfrak{q} -tame, and*

$$d_b(\text{dgm}(\mathcal{C}(L, W, d_X)), \text{dgm}(\mathcal{C}(L', W', d_X))) \leq d_H(L, L') + d_H(W, W').$$

(iii) *The witness filtrations $\mathcal{W}(L, W, d_X)$ and $\mathcal{W}(L, W', d_X)$ are \mathfrak{q} -tame, and*

$$d_b(\text{dgm}(\mathcal{W}(L, W, d_X)), \text{dgm}(\mathcal{W}(L, W', d_X))) \leq 2 d_H(W, W').$$

It is worth pointing out that (ii) is well known in the special case where $L, L' \subseteq W = W' = X = \mathbb{R}^n$. The usual argument is based on the Nerve Lemma, as in the proof of Theorem 7.2, so it relies on the local topological properties of finite-dimensional normed vector spaces and does not work in general. The above result shows that the dependence on the Nerve Lemma is in fact not necessary.

Let us also emphasize that $L = L'$ in (iii). There is indeed no equivalent of (ii) for witness complexes, as the persistent homology of witness filtrations turns out to be unstable with respect to perturbations of the set of landmarks, even if the set of witnesses is constrained to stay fixed ($W = W'$). Here is a counterexample. On the real line, consider the sets $W = L = \{0, 1\}$ and $L' = \{-\delta, 0, 1, 1 + \delta\}$, where $\delta \in (0, \frac{1}{2})$ is arbitrary. Then,

$$W_i(L, W, \ell^2) = \{[0], [1], [0, 1]\}$$

for all $i \geq 0$, whereas

$$W_i(L', W, \ell^2) = \{[-\delta], [0], [1], [1 + \delta], [-\delta, 0], [1, 1 + \delta]\}$$

for all $i \in [\delta, 1 - \delta)$. Thus, $H_*(\mathcal{W}(L, W, \ell^2))$ and $H_*(\mathcal{W}(L', W, \ell^2))$ are not ε -interleaved for any $\varepsilon < 1 - 2\delta$, whereas $d_H(L, L') = \delta$ can be made arbitrarily small compared to $1 - 2\delta$. Note that the set of witnesses is fairly sparse compared to the set of landmarks in this example. This raises several interesting questions, such as whether densifying W (e.g. taking the full real line) would allow us to regain some stability.

4. Signatures for metric spaces equipped with functions

The previous stability results extend to metric spaces equipped with real-valued functions. The corresponding metric is the following. Given two metric spaces (X, d_X) and (Y, d_Y) , equipped respectively with functions $f_X : X \rightarrow \mathbb{R}$ and $f_Y : Y \rightarrow \mathbb{R}$, the *functional distortion* induced by a correspondence $C : X \rightrightarrows Y$ is defined as the quantity

$$\text{dist}_f(C) = \sup_{(x,y) \in C} |f_X(x) - f_Y(y)|.$$

The Gromov-Hausdorff distance adapts naturally to this new setting:

$$(7.12) \quad d_{\text{GH}}((X, d_X, f_X), (Y, d_Y, f_Y)) = \frac{1}{2} \inf_{C: X \rightrightarrows Y} \max\{\text{dist}_m(C), \text{dist}_f(C)\}.$$

Note that we recover the usual Gromov-Hausdorff distance when f_X and f_Y are equal constant functions (say zero). The Rips based signature for (X, d_X) is modified as follows for (X, d_X, f_X) : consider the family of simplicial complexes $R_i(f_X^{-1}((-\infty, i]), d_X)$ for i ranging over \mathbb{R} , call this family $\mathcal{R}(X, d_X, f_X)$, and use its persistence diagram as signature. In view of (7.11)-(7.12), the proof of Proposition 7.8 adapts in a straightforward manner, so we get an interleaving between the modules $H_*(\mathcal{R}(X, d_X, f_X))$ and $H_*(\mathcal{R}(Y, d_Y, f_Y))$. Moreover, under the extra condition that the real-valued functions f_X, f_Y are Lipschitz continuous, the proof of Proposition 7.9 also goes through, proving that the two modules are \mathfrak{q} -tame. Hence⁴,

THEOREM 7.11. *Let (X, d_X) and (Y, d_Y) be totally bounded metric spaces, equipped with functions $f_X : X \rightarrow \mathbb{R}$ and $f_Y : Y \rightarrow \mathbb{R}$. Then,*

$$d_i(\mathcal{R}(X, d_X, f_X), \mathcal{R}(Y, d_Y, f_Y)) \leq 2 d_{\text{GH}}((X, d_X, f_X), (Y, d_Y, f_Y)).$$

Moreover, if f_X and f_Y are Lipschitz continuous, then the modules are \mathfrak{q} -tame, so

$$d_b(\text{dgm}(\mathcal{R}(X, d_X, f_X)), \text{dgm}(\mathcal{R}(Y, d_Y, f_Y))) \leq 2 d_{\text{GH}}((X, d_X, f_X), (Y, d_Y, f_Y)).$$

The Čech and witness based signatures adapt in the same way, with similar stability guarantees.

5. Computations

Given a finite metric space (X, d_X) with n points, the goal is to compute the signatures from the previous sections efficiently. We will restrict the focus to the Rips based signature $\text{dgm}(\mathcal{R}(X, d_X))$, as the challenges and solutions for the other signatures are similar.

We know from Section 2 of Chapter 5 that the Rips filtration contains 2^n simplices in total, so building it explicitly requires at least 2^n operations and therefore becomes quickly intractable in practice. There are two ways around this issue: first, one may compute only a fraction of the signature by truncating the filtration; second, one may compute an approximation of the signature by building a sparser yet interleaved filtration. In both cases we seek theoretical guarantees: in the first case, we want to maintain the stability properties of the signatures; in the second case, we want to control the approximation error.

⁴Theorem 7.5 can be viewed as a special case of Theorem 7.11 in which $f_X = f_Y = 0$.

5.1. Partial persistence diagrams via truncated filtrations. Let us start with the truncation strategy, which was introduced formally by Chazal et al. [58]. We will present it in full generality first, then we will specialize it to the case of the Rips filtration. Given a space F and a filtration $\mathcal{F} = \{F_t\}_{t \in \mathbb{R}}$ of that space, whose persistence diagram is well-defined, we wish to truncate \mathcal{F} at a prescribed parameter value $t = t_0$. The naive way is as follows:

$$(7.13) \quad \forall t \in \mathbb{R}, \overline{\overline{F}}_t^{t_0} = \begin{cases} F_t & \text{if } t < t_0, \\ F_{t_0} & \text{if } t \geq t_0. \end{cases}$$

Unfortunately, the resulting filtration $\overline{\overline{F}}^{t_0} = \{\overline{\overline{F}}_t^{t_0}\}_{t \in \mathbb{R}}$ does not retain the stability properties of the original filtration \mathcal{F} . Indeed, the points of $\text{dgm}(\mathcal{F})$ that lie just above the line $y = t_0$ are projected vertically at infinity in $\text{dgm}(\overline{\overline{F}}^{t_0})$, whereas small ℓ^∞ -perturbations may fall right below the line $y = t_0$ and therefore stay in place after the truncation. We therefore modify the truncation as follows:

$$(7.14) \quad \forall t \in \mathbb{R}, \overline{F}_t^{t_0} = \begin{cases} F_t & \text{if } t < t_0, \\ F & \text{if } t \geq t_0. \end{cases}$$

Let $\overline{\mathcal{F}}^{t_0} = \{\overline{F}_t^{t_0}\}_{t \in \mathbb{R}}$ be the resulting filtration. The effect of the new truncation is to project the persistence diagram of \mathcal{F} as follows—see Figure 7.5 for an illustration:

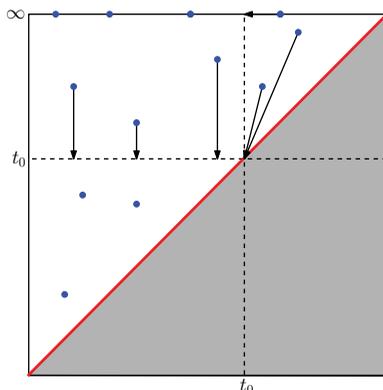


FIGURE 7.5. Effect of truncating a filtration as per (7.14) on its persistence diagram.

— From Chazal et al. [58].

LEMMA 7.12. *There is a matching $\gamma : \text{dgm}(\mathcal{F}) \rightarrow \text{dgm}(\overline{\mathcal{F}}^{t_0})$ such that:*

- the restriction of γ to the lower-left quadrant $(-\infty, t_0] \times (-\infty, t_0]$ is the identity;
- the restriction of γ to the upper-left quadrant $(-\infty, t_0] \times (t_0, +\infty)$ is the vertical projection onto the line $y = t_0$;
- the restriction of γ to the line $y = +\infty$ is the leftward horizontal projection onto the half-line $(-\infty, t_0] \times \{+\infty\}$;
- finally, the restriction of γ to the half-plane $(t_0, +\infty) \times \mathbb{R}$ is the projection onto the diagonal point (t_0, t_0) .

Since the projection is non-expansive in the bottleneck distance, for any filtered spaces (F, \mathcal{F}) and (G, \mathcal{G}) as above we have:

$$(7.15) \quad d_b(\text{dgm}(\overline{\mathcal{F}}^{t_0}), \text{dgm}(\overline{\mathcal{G}}^{t_0})) \leq d_b(\text{dgm}(\mathcal{F}), \text{dgm}(\mathcal{G})).$$

This means that the stability results from the previous sections still hold when the filtrations are replaced by their truncated versions as per (7.14).

In the special case where \mathcal{F} is the Rips filtration of a finite metric space (X, d_X) , Lemma 7.12 provides a way of computing $\text{dgm}(\overline{\mathcal{R}}^{t_0}(X, d_X))$ without building the Rips filtration $\mathcal{R}(X, d_X)$ entirely. This is not obvious, as the definition of the truncation in (7.14) involves the full simplex over the point set X . The fact is, the full simplex has trivial reduced homology, therefore the Rips filtration $\mathcal{R}(X, d_X)$ has only one essential homology class $[c]$, easily identified as the connected component created by one of the points of X (chosen arbitrarily). Then, $\text{dgm}(\overline{\mathcal{R}}^{t_0}(X, d_X))$ is obtained by (1) computing the naive truncated filtration $\overline{\mathcal{R}}^{t_0}(X, d_X)$ as per (7.13), and (2) projecting its persistence diagram (except the point corresponding to $[c]$) onto the lower-left quadrant $(-\infty, t_0] \times (-\infty, t_0]$ as described in Lemma 7.12 and depicted in Figure 7.5. The overall running time of the procedure is of the same order as the time needed to compute $\overline{\mathcal{R}}^{t_0}(X, d_X)$, which of course depends on the choice of threshold t_0 but diminishes (eventually becoming linear in the number of points in X) as t_0 gets closer to 0. The price to pay is that the signature becomes less informative as t_0 decreases.

5.2. Approximate persistence diagrams via sparse filtrations. Let us now focus on the approximation strategy. We will present a sparsification technique that turns the entire Rips filtration $\mathcal{R}(X, d_X)$ into a linear-sized filtration, with a control over the bottleneck distance between the persistence diagrams of the two filtrations. This strategy was introduced by Sheehy [221] and has since inspired other contributions [107], including the Rips zigzags from Chapter 5.

The intuition underlying the sparse Rips construction is best described in terms of unions of balls. Suppose X is a finite point set in \mathbb{R}^d , equipped with the Euclidean distance. Assume our target object is the persistence diagram of the ambient Čech filtration. By the Persistent Nerve Lemma 4.12, this diagram is the same as the one of the offsets filtration \mathcal{X} .

The principle governing the sparsification is to remove the balls one by one from the offset $X_t = \bigcup_{x \in X} B(x, t)$ as parameter t increases towards infinity. A simple criterion to remove a ball $B(x, t)$ safely is to make sure that it is covered by the other balls. Indeed, in this case the union stays the same after the removal, so the inclusion map $\bigcup_{y \in X \setminus \{x\}} B(y, t) \hookrightarrow \bigcup_{y \in X} B(y, t) = X_t$ is the identity and therefore induces an isomorphism at the homology level.

The problem with this criterion is that some balls may never get covered by the others. In fact, the top row in Figure 7.6 shows a simple example where none of the balls ever gets covered by the other ones, so the sparsification process does nothing. The fix consists in perturbing the metric in such a way that the balls that contribute little to the union eventually get covered by the other ones, as illustrated in the bottom row of Figure 7.6.

To describe the metric perturbation and the rest of the construction formally, we now go back to the general setting of a finite metric space (X, d_X) . Let

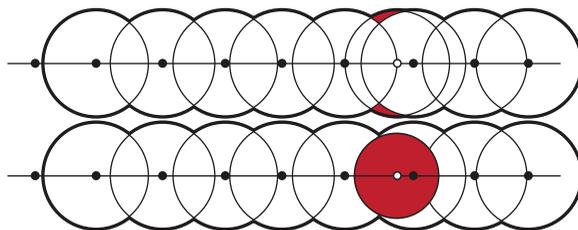


FIGURE 7.6. Top row: some points on a line. The ball centered at the white point contributes little to the union of balls. Bottom row: using the relaxed distance, the new ball is completely contained in the union of the other balls. This property is key to allowing for a safe removal of this ball without changing the topology of the union.

— Reprinted from Sheehy [221]: “Linear-Size Approximations to the Vietoris-Rips Filtration”, *Discrete & Computational Geometry*, 49(4):778–796, ©2013, with kind permission from Springer Science and Business Media.

(x_1, \dots, x_n) be a total order on the points of X , and let $X_i = \{1, \dots, i\}$ and $\varepsilon_i = d_H(X_i, X)$ for $i = 1, \dots, n$. Given a target approximation error $\varepsilon \in (0, \frac{1}{2})$, we perturb the metric d_X by incorporating additive weights that grow with the filtration parameter t :

$$(7.16) \quad \forall 1 \leq i, j \leq n, \quad d_X^t(x_i, x_j) = d_X(x_i, x_j) + s_t(x_i) + s_t(x_j),$$

where each weight $s_t(x_k)$ is defined by

$$s_t(x_k) = \begin{cases} 0 & \text{if } t \leq \frac{\varepsilon_{k-1}}{\varepsilon} \\ \frac{1}{2} \left(t - \frac{\varepsilon_{k-1}}{\varepsilon} \right) & \text{if } \frac{\varepsilon_{k-1}}{\varepsilon} \leq t \leq \frac{\varepsilon_{k-1}}{\varepsilon(1-2\varepsilon)} \\ \varepsilon t & \text{if } \frac{\varepsilon_{k-1}}{\varepsilon(1-2\varepsilon)} \leq t \end{cases}$$

The weight function $t \mapsto s_t(x_k)$ is depicted in Figure 7.7. It is clearly $\frac{1}{2}$ -Lipschitz, nonnegative, and its effect on the metric is to take x_k away from the other points artificially, so its distance to the rest of the point set looks greater than it is in reality. The consequence is that its metric balls in the perturbed metric are smaller than the ones in the original metric, and the choice of parameters makes it so that they get covered by the other balls eventually, following the intuition from Figure 7.6.

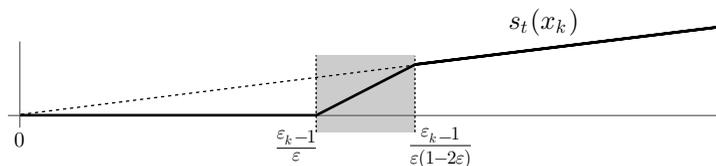


FIGURE 7.7. The additive weight function $t \mapsto s_t(x_k)$ associated to point $x_k \in X$.

— Reprinted from Sheehy [221]: “Linear-Size Approximations to the Vietoris-Rips Filtration”, *Discrete & Computational Geometry*, 49(4):778–796, ©2013, with kind permission from Springer Science and Business Media.

To simulate the ball growth and removal process in the discrete metric space (X, d_X) , the sparse Rips construction relies on Rips complexes. Set $\rho = \frac{1}{\varepsilon(1-2\varepsilon)}$ and consider the Morozov zigzag (M-ZZ) in the perturbed metric d_X^t —recall the definition of the M-ZZ from the top row in (5.11):

$$(7.17) \quad \begin{array}{ccccccc} \cdots & & R_{\rho\varepsilon_{i-1}}(X_i, d_X^{\rho\varepsilon_{i-1}}) & & R_{\rho\varepsilon_i}(X_{i+1}, d_X^{\rho\varepsilon_i}) & & \cdots \\ & \swarrow & \nearrow & \swarrow & \nearrow & \swarrow & \nearrow \\ R_{\rho\varepsilon_{i-1}}(X_{i-1}, d_X^{\rho\varepsilon_{i-1}}) & & R_{\rho\varepsilon_i}(X_i, d_X^{\rho\varepsilon_i}) & & R_{\rho\varepsilon_{i+1}}(X_{i+1}, d_X^{\rho\varepsilon_{i+1}}) & & \end{array}$$

This construction is the combinatorial counterpart of the continuous ball growth and removal process described previously. Specifically, growing the balls corresponds to increasing the Rips parameter, while removing a ball from the union corresponds to removing a vertex and its cofaces from the Rips complex. In order to avoid actual simplex removals and keep a standard filtration, the sparse Rips construction proceeds by merging the top row of (7.17), that is to say, for every filtration parameter $t > 0$ it takes the union

$$S_t(X, d_X, \varepsilon) = R_t(X_i, d_X^t) \cup \bigcup_{j=i+1}^n R_{\rho\varepsilon_{j-1}}(X_j, d_X^{\rho\varepsilon_{j-1}}),$$

where $i \in \{1, \dots, n\}$ is the value such that⁵ $t \in (\rho\varepsilon_i, \rho\varepsilon_{i-1}]$. This is what we call the *sparse Rips complex* of parameter $t > 0$. In the continuous setting, the corresponding construction stops the growth of the balls at the time when they should be removed from the union, instead of actually removing them. The effect in terms of topological change is the same, while the size overhead remains under control.

By convention we let $S_0(X, d_X, \varepsilon) = X$ and $S_t(X, d_X, \varepsilon) = \emptyset$ for $t < 0$. The resulting indexed family $\mathcal{S}(X, d_X, \varepsilon) = \{S_t(X, d_X, \varepsilon)\}_{t \in \mathbb{R}}$ is called the *sparse Rips filtration*. One can check that it is indeed a finite simplicial filtration⁶ on the point set X . Assuming the order (x_1, \dots, x_n) on the points of X is obtained by furthest-point sampling as in Definition 5.11, the same ball packing argument as the one used in Section 3.2 of Chapter 5 to bound the total number of simplex insertions in the M-ZZ allows us to get a linear bound on the size of $\mathcal{S}(X, d_X, \varepsilon)$:

THEOREM 7.13. *Suppose (X, d_X) has doubling dimension m and the order (x_1, \dots, x_n) on the points of X is obtained by furthest-point sampling. Then, for any $k \geq 0$, the total number of k -simplices in the sparse Rips filtration $\mathcal{S}(X, d_X, \varepsilon)$ is at most $(\frac{1}{\varepsilon})^{O(km)} n$.*

REMARK. As pointed out at the end of Section 3.1 in Chapter 5, it takes quadratic time to compute a furthest-point sampling order naively. However, near linear-time approximations like *net-trees* [149] can be used instead, with roughly the same size complexity and approximation guarantees.

Concerning now the approximation properties of the sparse Rips filtration, there is an $O(\varepsilon)$ -proximity between the logscale persistence diagrams of $\mathcal{R}(X, d_X)$ and $\mathcal{S}(X, d_X, \varepsilon)$. The proof proceeds by working out an explicit multiplicative

⁵By convention we let $\varepsilon_0 = d_H(\emptyset, X) = +\infty$.

⁶The key point is to show that $R_t(X_i, d_X^t) \subseteq R_{t'}(X_i, d_X^{t'})$ for all $t \leq t'$, using the fact that the map $t \mapsto s_t(x_k)$ is $\frac{1}{2}$ -Lipschitz for all $k = 1, \dots, n$.

interleaving between the filtrations, not in the sense of inclusions as in (3.3) but in the sense of contiguous maps as in the proof of Proposition 7.8. Specifically, it shows that for all $t > 0$ the following diagram of simplicial complexes and simplicial maps commutes at the simplicial level, i.e. oriented paths sharing the same source and target complexes induce contiguous maps:

$$(7.18) \quad \begin{array}{ccc} R_t(X, d_X) & \longrightarrow & R_{t(1+2\varepsilon)}(X, d_X) \\ \uparrow & \searrow & \uparrow \\ S_t(X, d_X, \varepsilon) & \longrightarrow & S_{t(1+2\varepsilon)}(X, d_X, \varepsilon) \end{array}$$

The horizontal maps are the inclusions involved in the filtrations $\mathcal{R}(X, d_X)$ and $\mathcal{S}(X, d_X, \varepsilon)$. The vertical maps are also inclusions, since by definition d_X^t is at least d_X so

$$\begin{aligned} S_t(X, d_X, \varepsilon) &= R_t(X_i, d_X^t) \cup \bigcup_{j=i+1}^n R_{\rho\varepsilon_{j-1}}(X_j, d_X^{\rho\varepsilon_{j-1}}) \\ &\subseteq R_t(X_i, d_X) \cup \bigcup_{j=i+1}^n R_{\rho\varepsilon_{j-1}}(X_j, d_X) \subseteq R_t(X, d_X), \end{aligned}$$

where $i \in \{1, \dots, n\}$ is such that $t \in (\rho\varepsilon_i, \rho\varepsilon_{i-1}]$. The diagonal arrow $S_t(X, d_X, \varepsilon) \rightarrow R_{t(1+2\varepsilon)}(X, d_X)$ is also an inclusion, therefore it makes its two incident triangles commute. The other diagonal arrow $R_t(X, d_X) \rightarrow S_{t(1+2\varepsilon)}(X, d_X, \varepsilon)$ is defined as the simplicial map induced by the projection π_t onto X_i in the perturbed metric⁷:

$$\pi_t(x_k) = \begin{cases} x_k & \text{if } k \leq i, \\ \arg \min_{1 \leq l \leq i} d_X^t(x_k, x_l) & \text{otherwise.} \end{cases}$$

The fact that this projection indeed induces a simplicial map from $R_t(X, d_X)$ to $S_{t(1+2\varepsilon)}(X, d_X, \varepsilon)$ is a consequence of the fact that it is non-expansive in the perturbed metric d_X^t . The fact that its incident triangles in (7.18) define contiguous maps (and therefore commute at the homology level) follows from the fact that all the simplicial maps induced by projections π_t with $t \in (\rho\varepsilon_i, \rho\varepsilon_{i-1}]$ are contiguous in $S_{\rho\varepsilon_{i-1}}(X, d_X, \varepsilon)$ and $R_{\rho\varepsilon_{i-1}}(X, d_X)$, which is also a consequence of the projection π_t being non-expansive. Let us refer the reader to [221] for the details of the proofs of these claims.

Thus, (7.18) induces a commutative diagram at the homology level, which defines a multiplicative $(1+2\varepsilon)$ -interleaving between the persistence modules $H_*(\mathcal{R}(X, d_X))$ and $H_*(\mathcal{S}(X, d_X, \varepsilon))$. On a logarithmic scale, this turns into an additive $\log_2(1 + 2\varepsilon)$ -interleaving, which by the Isometry Theorem 3.1 gives the claimed approximation guarantee:

THEOREM 7.14. *The bottleneck distance between the logscale persistence diagrams of the filtrations $\mathcal{R}(X, d_X)$ and $\mathcal{S}(X, d_X, \varepsilon)$ is bounded by $\log_2(1+2\varepsilon) = O(\varepsilon)$.*

REMARK. As pointed out at the end of Section 4 in Chapter 5, approximations to the Rips filtration such as the sparse Rips filtration can also be used in the context of topological inference, however they are best suited for scenarios in which

⁷The case $k \leq i$ is treated separately in the definition of the projection because in principle one might have $d_X^t(x_k, x_l) < d_X^t(x_k, x_k)$, the perturbed metric d_X^t being not a true metric.

the Rips filtration is the target object, such as this one. For topological inference, the Rips zigzags from Chapter 5 give the best signal-to-noise ratio in the barcodes, at a smaller algorithmic cost (the complexity bounds do not involve the quantity $\frac{1}{\varepsilon}$).

Part 3

Perspectives

New Trends in Topological Data Analysis

Topological inference has been the main practical motivation for the development of persistence theory for the last decade or so. As such, it has been assigned a special place in Part 2 of this book. Yet, another lesson to be learnt from Part 2 is that the recent years have seen the emergence of novel applications of persistence, which have defined new needs and raised new challenges for the theory. Generally speaking, persistence barcodes or diagrams can be used as compact topological descriptors for geometric data, to be combined with other descriptors for effective interpretation or comparison. Experience has shown that topological descriptors do provide complementary information to the other existing descriptors¹, which makes them a valuable addition to the data analysis literature. Yet, there still remains a long way to go before they become part of the standard toolbox.

The main bottleneck so far has been the space and time complexities of the pipeline: as we saw in Section 5 of Chapter 5, we are currently able to process a few thousands of points in a single data set. While this is enough to obtain relevant inference results (topology is indeed fairly robust to undersampling compared to other quantities like curvature), it is not enough to make these results statistically significant and to validate them *a posteriori*. In this respect, being able to handle inputs of much larger sizes (say by 2 or 3 orders of magnitude) would be desirable—see Section 1 below. This is already the case for the 0-dimensional version of the inference pipeline, which, for instance, has been successfully applied to cluster millions of points, as we saw in Chapter 6.

Another bottleneck that is becoming obvious now is the lack of a sound statistical framework for the interpretation of the barcodes or diagrams. Currently, these descriptors are used as an exploratory tool for data mining, and their interpretation relies entirely on the user. It is desirable to develop visualization and analysis techniques to help the user ‘read’ the barcodes, identify their relevant parts (the ones corresponding to relevant scales at which to look at the data), and assess their statistical significance—see Section 2.

Another bottleneck is the lack of a proper way to apply supervised learning techniques on collections of barcodes or diagrams. Currently, the signatures are mostly used in unsupervised learning tasks (e.g. unsupervised shape classification, as illustrated in Chapter 7), simply because the space of diagrams is not naturally amenable to the use of linear classifiers. Finding ways of defining kernels for diagrams that would be both meaningful and algorithmically tractable, would be desirable. This would be a first step towards combining topological data analysis with machine learning—see Section 3.

¹See e.g. [56] for experimental evidence of this phenomenon.

1. Optimized inference pipeline

There are essentially two ways of optimizing the pipeline: either by reducing the size of the filtrations or zigzags involved, or by improving the complexity of computing their persistence diagrams. Both directions have been investigated, as reported in Chapters 2 and 5. Handling inputs of much larger sizes (by 2 or 3 orders of magnitude) will require to bring in some new ideas.

Beating the matrix multiplication time? As we saw in Chapter 2, the persistence algorithm proceeds by a mere Gaussian elimination. As such it can be optimized using fast matrix multiplication, which reduces the complexity from $O(m^3)$ down to $O(m^\omega)$, where m is the number of simplices in the input filtration and $\omega \in [2, 2.373)$ is the best known exponent for multiplying two $m \times m$ matrices. The same holds for zigzag persistence modulo some adaptations [197].

Whether matrix multiplication time is the true asymptotic complexity of the persistence computation problem remains open, and to date there is still a large gap with the currently best known lower bound, which is only linear in m . It was recently shown that computing the Betti numbers of a (2-dimensional) simplicial complex with m simplices is equivalent to computing the rank of a sparse matrix with m nonzero entries [121]. However, it is not known whether the sparsity of matrices can help in computing the ranks, and more generally whether computing Betti numbers is easier than computing persistent homology.

OPEN QUESTIONS. *Can the question of the asymptotic complexity of the persistence computation problem be settled? Perhaps not in full generality, but in special cases such as when the filtrations or zigzags are the ones used for homology inference?*

Indeed, worst-case examples such as the ones in [121, 199] are far from the typical complexes built for homology inference. The latter have a specific structure inherited from the geometry, which could possibly be exploited using for instance space partition techniques [149].

Memory-efficient and distributed persistence computation. A simple idea to reduce the memory footprint of the persistence algorithm is to cut the boundary matrix into blocks to be stored on the disk and loaded separately in main memory. This is the *divide-and-conquer* idea underlying the fast matrix multiplication approach, and it could be pushed further towards an algorithm whose memory usage is fully controlled by the user via the choice of a maximum block size. The main issue is the book-keeping: although each block can easily be reduced locally, merging the results on nearby blocks may require non-local operations. This is because homology itself is global, and the homology generators may be highly non-localized.

The standard tool to merge local homological information into global information is the Mayer-Vietoris sequence. It was used e.g. by Zomorodian and Carlsson [244] to localize topological attributes of a topological space relative to a given cover of that space. Similar ideas were used by Bauer, Kerber, and Reininghaus [17] to distribute the persistence calculation. These approaches share a common limitation, which is that a global step is needed to finalize the calculation, so the actual memory usage is not controlled.

OPEN QUESTIONS. *Is it possible to force all the calculations to remain ‘local’ in the sense that only a constant number of (possibly noncontiguous) blocks of the*

boundary matrix are concerned at each step? What fraction of these operations could be performed in parallel? Can the implementation be adapted for a large-scale deployment, either on a local computer cluster or on the cloud?

Bridging the gap between standard persistence and zigzag persistence computations. While zigzag constructions such as the Rips zigzags of Chapter 5 can be proven both lightweight and efficient for inference purposes, their main drawback is that the current state of the art in zigzag persistence computation is nowhere as optimized as the one in standard persistence. To give an idea, the best implementations of the standard persistence algorithm [18, 30] are able to process millions of simplex insertions per second on a recent machine, whereas in the experiments with Rips zigzags conducted by Oudot and Sheehy [208], the software was only able to process a few thousands of simplex insertions per second².

OPEN QUESTIONS. *To what extent can persistence for zigzags be optimized? Can the gap in performance between standard persistence and zigzag persistence computations be closed? These questions are becoming essential as zigzags are gaining more and more importance in applications.*

A possible strategy is to turn the zigzags into standard filtrations, or rather, to turn their induced zigzag modules at the homology level into standard persistence modules. This can be done via sequences of arrow reflections, as in the proof of existence of interval decompositions for zigzag modules (see Appendix A). Maria and Oudot [186] started investigating this direction from an algorithmic point of view, combining arrow reflections with arrow transpositions to compute a compatible homology basis for the input zigzag. Preliminary experiments show that this approach is competitive versus the state of the art and promising in terms of future optimizations. In particular:

OPEN QUESTIONS. *Can the approach be adapted so as to work with cohomology instead of homology? Can it benefit then from the recent optimizations for standard persistence computation, thereby reducing the gap between standard persistence and zigzag persistence in terms of practical performance?*

Lightweight data structures. Somewhat orthogonal to the previous questions is the choice of data structures in practical implementations. The main challenge is to find the right trade-off between lightweight and efficiency. For the storage of simplicial complexes and filtrations, the current trends are: on the one hand, constant-factor optimizations compared to the full Hasse diagram representation, with efficient faces and cofaces queries, as offered e.g. by the *simplex tree* [33]; on the other hand, exponential improvements on specific types of complexes (typically, Rips complexes) but with slow query times, as offered e.g. by the *blockers* data structure [6]. For the persistence computation itself, alternate representations of the boundary matrix are being explored, such as for instance the annotation matrix [107] and its compressed version [30].

OPEN QUESTIONS. *What is the right trade-off between size reduction and query optimization in the design of new simplicial complex representations? What would be the influence on the design and performance of the persistence algorithm?*

²And far fewer deletions due to a known limitation in the internal data representation of the zigzag persistence package in the `Dionysus` library at the time of the experiment.

Let us point out that there currently is no reference library for persistence computation. Several implementations are competing, including PLEX (<http://comptop.stanford.edu/u/programs/plex.html>) and its successor JPLEX (<http://comptop.stanford.edu/u/programs/jplex/>), DIONYSUS (<http://www.mrzv.org/software/dionysus/>), PHAT (<https://code.google.com/p/phat/>), PERSEUS (<http://www.sas.upenn.edu/~vnanda/perseus/index.html>), and GUDHI (<http://pages.saclay.inria.fr/vincent.rouvreau/gudhi/>). Each one of them implements a fraction of the approaches listed in Chapter 2 and possesses its own strengths and weaknesses. In the future, it would be desirable to unify all these implementations into a single reference library, for better visibility outside the field.

2. Statistical topological data analysis

Model selection for scale estimation. The general question of choosing relevant geometric scales at which to process the input data for inference can be cast into a *model selection* problem. Broadly speaking, the goal of model selection is to choose between several predictive models within a given family. A classical answer is, first, to estimate the prediction error for each model, and second, to select the model that minimizes this criterion among the family. To perform the second step, cross-validation or penalization techniques have been well studied and intensively used with linear models [188].

The approach was adapted by Caillerie and Michel [46] to work with families of simplicial complexes, which are piecewise linear models. This was done by introducing a least-squares penalized criterion to choose a complex. The method is limited to minimizing the ℓ^2 -distance between the immersed simplicial complex and the compact set underlying the input point cloud.

OPEN QUESTIONS. *Can the method be adapted to minimizing the Hausdorff distance between the immersed complex and the compact set? Can topological information such as Betti numbers or persistence barcodes be incorporated into the penalty criterion?*

Statistics on the space of persistence diagrams. The question of assessing the quality of the produced barcodes is usually addressed by running the inference pipeline repeatedly on resampled data, and by computing various statistical quantities on its outputs. The problem is that the space of persistence diagrams (equipped with the bottleneck distance) is not a flat space, so basic quantities like the mean of a probability distribution over the persistence diagrams are not naturally defined. Mileyko, Mukherjee, and Harer [193] gave sufficient conditions on the distribution so its Fréchet mean exists. Moreover, Turner et al. [232] gave an algorithm that converges to a local minimum of the second moment of the distribution under some conditions. Nevertheless, they acknowledged that the mean may not exist in general, and that when it does it may not necessarily be unique. The ‘right’ conditions still remain to be found.

OPEN QUESTION. *Can the existence and uniqueness of a mean persistence diagram be guaranteed under realistic conditions on the distribution over the diagrams?*

Alternate representations of persistence diagrams. A workaround is to map the persistence barcodes or diagrams to a different space where classical statistical quantities are well-defined—and preferably also easy to compute. This is the approach taken by Bubenik [39], who introduced the concept of *persistence landscapes*. Roughly speaking, these are obtained by taking the size functions associated to the persistence diagrams (recall Figure 0.5 from the general introduction) and by parametrizing their discontinuity lines along the diagonal. This defines an injective mapping to the space of square-integrable functions over the real line. The gain is that means can be easily computed in that space, however the loss is that they may fail to be images of persistence diagrams through the mapping. In other words, statistical quantities become well-defined and computable, but they no longer have a meaning in terms of persistence.

OPEN QUESTION. *Can we find other mappings (possibly to other function spaces) that allow to map means back to the space of persistence diagrams in a meaningful way?*

Convergence rates for persistence diagram estimators. Another option is to avoid using means at all. A setting in which this is possible is when the distribution of persistence diagrams is composed of persistence based signatures (e.g. Rips based signatures, as defined in Chapter 7) of more and more refined point samples of a fixed probability distribution μ over some metric space. In this case indeed, there is a ground truth, namely the persistence based signature of the support of μ , and the goal is to study the convergence of the estimator to the ground truth as the number of sample points goes to infinity. Taking advantage of the stability properties of persistence based signatures established in Chapter 7, Chazal et al. [68] derived bounds on the convergence rate of the estimator that are optimal in the minimax sense, up to some logarithmic factors.

OPEN QUESTIONS. *Can similar bounds be derived for a more general class of persistence diagram estimators? How can the stability properties of the weighted Rips filtration (Definition 5.17) be exploited in this context?*

Confidence intervals. In a similar context, Balakrishnan et al. [15] then Chazal et al. [69] derived confidence intervals for persistence diagrams. Such intervals are useful for discriminating the topological signal from the noise in the persistence diagrams. Several approaches were investigated for this purpose, including subsampling, concentration of measure, and bootstrapping. The key point was to use them for a slightly different problem, namely the one of deriving confidence bounds for the Hausdorff approximation of the support of the probability measure from which the input points were sampled. The bounds in Hausdorff distance between samples and support were then turned into bounds in bottleneck distance between the persistence diagrams, by means of the stability part of the Isometry Theorem—or rather its functional version stated in Corollary 3.6. The main problem with this strategy is that it focuses exclusively on the variance terms in the estimation. Moreover, it misses somewhat the point by estimating the support of the measure rather than the measure itself. Finally, the analysis generally relies on hypotheses (e.g. *(a, b)-standardness*) involving quantities that remain hidden in practice, thus making the calibration of the methods difficult.

OPEN QUESTIONS. *Can the strategy be adapted so as to give control also over the bias terms? Can the distance to the measure itself be considered instead of the distance to its support in the approximation? How can the hidden quantities be estimated in practice?*

3. Topological data analysis and machine learning

Kernels for persistence barcodes and diagrams. As we saw, the space of persistence diagrams is not a flat space, which prevents the use of linear classifiers directly. However, the so-called ‘kernel trick’ can be employed. Indeed, to any positive-definite kernel defined over a space \mathcal{X} corresponds a map from \mathcal{X} to some Hilbert space \mathcal{H} , such that the value of the kernel for a pair of elements of \mathcal{X} matches with the scalar product between their images in \mathcal{H} . Thus, standard algorithms relying solely on scalar product evaluations, such as linear classification, k-means, or Principal Component Analysis, can easily be applied. However, this approach requires to be able to define kernels that are positive-definite in the first place. Moreover, they should be fast enough to compute regardless of the size and dimensionality of the data, so the approach can handle large and high-dimensional data sets.

A standard technique to design positive-definite kernels is to plug the distance function into a Gaussian kernel. However, this requires the distance function itself to be negative-definite, which is not the case of the bottleneck distance, whose behavior is similar to the one of an ℓ^∞ -norm. More generally, it is not the case for any of the Wasserstein distances between persistence diagrams.

A possible workaround is to change the representation of the diagrams. For instance, we already mentioned the *persistence landscapes* [39], which map the diagrams to square-integrable functions over the real line. Another possible embedding to function space, proposed by Reininghaus et al. [214], consists in viewing each diagram as a measure (sum of Dirac masses) and to apply a heat diffusion process to obtain a square-integrable function over the upper half-plane above the diagonal. A third possibility, proposed by Carrière, Oudot, and Ovsjanikov [56], is to see each diagram itself as a finite metric space, and to map it to its so-called *shape context* (i.e. the distribution of its sorted pairwise distances), which gives a mapping to a finite-dimensional vector space. In all three cases, classical kernels (linear, polynomial, Gaussian, etc.) can be applied in the new ambient space. Moreover, each mapping is provably stable, so the stability properties of persistence diagrams are preserved. The downside is that the induced kernels on the space of persistence diagrams are only guaranteed to be positive semidefinite, not positive definite, thus resulting in a potential loss in discriminative power.

OPEN QUESTION. *Can we design other mappings of persistence diagrams to normed vector spaces, which would be both stable and with a controlled loss of discriminative power compared to their originating diagrams?*

Meaningful metrics or (dis-)similarities via signatures. Descriptors that are both stable and informative can in turn be used to define relevant distances between data sets. For instance, recall from Chapter 7 that the natural metric between 3d shapes—the Gromov-Hausdorff distance—is hard to compute or approximate directly, which motivates the use of signatures that are easier to compare. When the chosen signatures (or group thereof) are both stable and informative, their pairwise

distances can serve as a reliable proxy for the true Gromov-Hausdorff distance. This is what happens e.g. in the example of Figure 7.3, where the bottleneck distance between persistence diagrams separates the various classes of shapes nicely. The question is whether this property holds more generally.

OPEN QUESTIONS. *Can we derive general lower bounds on the Gromov-Hausdorff distance based on the bottleneck distance between persistence diagrams? Do such guarantees on the discriminative power of persistence diagrams hold for other learning applications?*

Other topological descriptors. Persistence barcodes and diagrams are but one class of topological descriptors to be used for data interpretation and comparison. Among the other existing descriptors, the graphs produced by *Mapper* [225] have gained a lot of interest among practitioners. The underlying intuition is the following one. Given a topological space X and an open cover \mathfrak{U} of that space, we can use the nerve of that cover as a proxy for the topological structure of X . This proxy will be correct for instance when the conditions of the Nerve Lemma 4.11 are satisfied. But how do we choose a relevant cover? This is where the term ‘Mapper’ reveals itself: consider a continuous map f to another topological space Y , choose a cover \mathfrak{U}' of that space, and pull it back via f to a cover of X . More precisely, define \mathfrak{U} to be composed of the sets $U = f^{-1}(U')$, for U' ranging over the elements of \mathfrak{U}' . The interest of this pull-back operation is that we can choose a space Y with a much simpler structure (e.g. a Euclidean space) to serve as ‘parameter space’ for X . When $Y = \mathbb{R}$, the construction is closely related to the one of the Reeb graph of f .

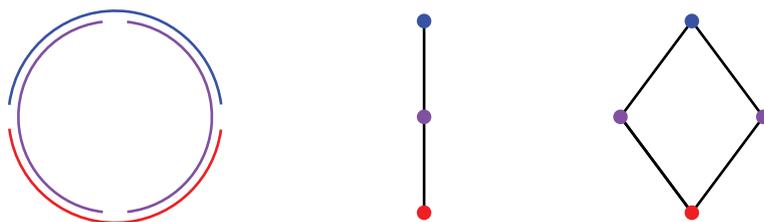


FIGURE 8.1. Left: a cover \mathfrak{U} of the circle $x^2 + y^2 = 1$ pulled back from the cover $\mathfrak{U}' = \{[-1, -\varepsilon); (-1 + \varepsilon, 1 - \varepsilon); (\varepsilon, 1]\}$ of the interval $Y = [-1, 1] \subset \mathbb{R}$ via the height function $(x, y) \mapsto y$. Center: the nerve of \mathfrak{U} . Right: the nerve of \mathfrak{U} after refinement.

— Based on Carlsson [48].

Now, it is not always the case that the pulled-back cover \mathfrak{U} satisfies the conditions of the Nerve Lemma 4.11. In particular, an element $U \in \mathfrak{U}$ may have several connected components. In that case, it is reasonable to replace U by its various connected components in \mathfrak{U} , as illustrated in Figure 8.1. The nerve of the thus refined cover is closely related to the concept of *multinerve* introduced by Colin de Verdière, Ginot, and Goaoc [90]. In the discrete setting, i.e. when X is replaced by a finite point sample P , the connected components of each set $U \in \mathfrak{U}$ are replaced by the connected components of the Rips complex of $U \cap P$, for some user-defined Rips parameter.

OPEN QUESTIONS. *Can the framework of multinerves, in particular the Homological Multinerve Theorem of Colin de Verdière, Ginot, and Goaoc [90], be used to derive theoretical guarantees on the topology of the graphs produced by Mapper? In the context where the data points are sampled at random from some probability distribution on X , can we derive bounds on the probability of correctness of the construction? Can we use statistics to choose a relevant cover of the parameter space Y to start with, for instance in the case where $Y = \mathbb{R}$? Finally, can we employ the graphs produced by Mapper in unsupervised learning tasks, e.g. by defining a distance or (dis-)similarity between the graphs? Or in supervised learning tasks, e.g. by defining positive-definite kernels over the space of such graphs?*

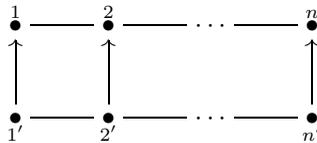
Further prospects on the theory

The theory itself has had some exciting recent contributions towards greater generality and flexibility. Pursuing these efforts is important for the future, to find and explore new applications of the theory.

1. Persistence for other types of quivers

As mentioned in the discussion sections of Chapters 1 and 3, Bubenik and Scott [41] then Bubenik, de Silva, and Scott [40] proposed to generalize the concept of persistence module to representations of arbitrary posets. By doing so, they lost the ability to define complete discrete invariants like persistence barcodes in a systematic way. Nevertheless, they were able to generalize the concept of interleaving to this setting, and to derive ‘soft’ stability results that bound the interleaving distance in terms of the distance between the topological objects the persistence modules are derived from originally. This is a promising direction toward a generalization of the theory, and a strong incentive to look beyond linear quivers and their infinite extensions.

Indeed, there are other quivers of finite type beyond the Dynkin quivers, thanks to the relations (path identifications) put on them. Escolar and Hiraoka [126] distinguished several such quivers among a family called the *commutative ladders*, depicted as follows and equipped with commutativity relations in every quadrant and the constraint that the arrow orientations in the top sequence must be the same as those in the bottom sequence:



Interestingly, up to horizontal arrow orientations, this quiver is the same as the one considered in persistence for kernels, images and cokernels, introduced in Section 1.4 of Chapter 2. However, here we are interested in the full representation structure, not just in the kernel, image or cokernel of the morphism between zigzag modules induced by the vertical arrows. What Escolar and Hiraoka [126] found is that Gabriel’s theorem extends to the commutative ladders as long as $n \leq 4$, so there is a complete discrete invariant for their finite-dimensional representations. They also provided an algorithm to compute this invariant, based on the so-called *Auslander-Reiten quivers* [13] associated to the commutative ladders. Their approach has found applications in materials science.

OPEN QUESTIONS. *Are there other classes of quivers (possibly with relations) that are of finite type? Can the associated complete discrete invariants be computed efficiently?*

There is also the question of the tame quivers, possibly with relations. Beyond the classical case of the Euclidean quivers, presented in Appendix A (Figure A.3), there are examples of quivers with relations that are of tame type. Yakovleva [240] gave such an example, composed of a chain of commutative squares glued along common vertices instead of common edges as above. Tame-type quivers have no complete discrete invariant, however their set of isomorphism classes of indecomposable representations is parametrized by a countable family of 1-dimensional varieties (Theorem A.30). This means that the decomposition of a finite-dimensional representation into indecomposables can be represented as a finite set of parameter values.

OPEN QUESTION. *Can we design algorithms to decompose the representations of tame-type quivers into indecomposables, and to represent the decompositions in the aforementioned parametrization?*

This is in fact what Burghelea and Dey [45] did for a certain class of Euclidean quivers of type \tilde{A}_{2n} ($n > 0$), whose arrow orientations along the cycle are alternating. They gave an algorithm to decompose any finite-dimensional representation into indecomposables, which in this particular setting are intervals and Jordan block matrices¹. They also showed how their quiver representations and associated decompositions apply to the context of persistence for circle-valued maps $f : X \rightarrow \mathbb{S}^1$.

But perhaps the previous question is too much asking in general. When it is, we are reduced to designing incomplete discrete invariants. This is e.g. what Edelsbrunner, Jabłonski, and Mrozek [116] did recently for the Euclidean quiver of type \tilde{A}_0 , which is made of one node and one loop, and whose isomorphism classes of indecomposable representations correspond to the conjugacy classes of Jordan block matrices²—see Example A.26 in Appendix A. Given a representation $V \xrightarrow{v} V$ of the quiver, the invariant used by Edelsbrunner, Jabłonski, and Mrozek [116] is the persistence diagram of a certain persistence module defined on the eigenspaces of v . Although incomplete, this invariant is shown to be easy to compute and informative in practical scenarios.

The situation for wild-type quivers seems more desperate. Due to the difficulty of the classification problem, we are once again reduced to deriving some incomplete discrete invariant. This is what Carlsson and Zomorodian [53] did for the poset \mathbb{N}^n equipped with the order relation $x \preceq y$ iff $x_i \leq y_i$ for all $i \in \{1, \dots, n\}$, which appears naturally in applications where scale is not the only relevant parameter for the considered problem, other parameters like density being also important. Just like the 1-dimensional poset \mathbb{N} is related to the category of graded modules over the polynomial ring $\mathbf{k}[t]$ (as pointed out after Theorem 1.4), the n -dimensional poset \mathbb{N}^n is related to the category of n -graded modules over the polynomial ring $\mathbf{k}[t_1, \dots, t_n]$, whose classification was shown to be wild by Carlsson and Zomorodian [53]. This motivated the introduction of the (discrete and incomplete) rank invariant, which records the ranks of the maps v_x^y for all indices $x \preceq y \in \mathbb{N}^n$. Carlsson, Singh, and

¹Assuming that the field of coefficients is algebraically closed.

²Assuming once again that the field of coefficients is algebraically closed.

Zomorodian [52] showed how to compute this invariant using Groebner bases, and they demonstrated its practicality.

OPEN QUESTIONS. *How can we assess the degree to which a given discrete invariant is incomplete? Generally speaking, what properties should a ‘good’ discrete invariant satisfy when it is incomplete?*

2. Stability for zigzags

While the literature on the stability of persistence modules is becoming abundant, there still lacks a version of the Isometry Theorem 3.1 for zigzags. Even though their persistence diagrams can be proven to carry relevant information in the context of topological inference, as we saw in Chapter 5, this information is still not guaranteed to be fully stable with respect to small perturbations of the input. Deriving such stability guarantees would be interesting e.g. for using the persistence diagrams of Rips zigzags as signatures for metric spaces, as we did with the ones of Rips filtrations and their sparse variants in Chapter 7.

A major limitation with zigzags is that the index set has to be finite, or at the very least countable. Indeed, there is no clear notion of a zigzag indexed over \mathbb{R} . This is an important limitation as the concept of module interleaving and its associated stability properties presented in Chapter 3 apply only to modules indexed over \mathbb{R} .

In fact, as representations of posets, zigzag modules can be interleaved in a categorical sense, as defined in [40, 41], and a bound on their interleaving distance can be obtained in terms of the distances between the topological objects the zigzags are derived from. From there, the question is whether the bound on the interleaving distance implies a bound on the bottleneck distance between their persistence diagrams.

OPEN QUESTION. *Can we relate the categorical interleaving distance of [40, 41] between zigzag modules and the bottleneck distance between their persistence diagrams?*

A different approach is to connect zigzags to standard filtrations, in order to benefit from the stability properties of the latter. This is what level-set persistence does through the pyramid construction and its associated Pyramid Theorem 2.11. The challenge is to be able to draw the same kind of connection at the algebraic level directly, so that the algebraic framework developed for the stability of standard persistence modules can be applied to zigzag modules by proxy. One possible strategy would be to turn a zigzag module into a standard persistence module by iteratively reversing its backward arrows, either through the Arrow Reversal Lemma 5.13, or through the Reflection Functors Theorem A.15 presented in Appendix A. The question is whether any of these operations can be made canonical, so that the derived stability result satisfies some kind of universality property.

OPEN QUESTION. *Can we define a canonical way to turn a zigzag module into a standard persistence module, while preserving (most of) its algebraic structure?*

3. Simplification and reconstruction

Finally, let us go back to one of the early motivations of persistence, which is the topological simplification of real-valued functions [118]. Given a function f

from a simplicial complex K to \mathbb{R} , and a tolerance parameter $\delta \geq 0$, the goal is to construct a δ -simplification of f , i.e. a function f_δ such that $\|f_\delta - f\|_\infty \leq \delta$ and the persistence diagram of (the sublevel-sets filtration of) f_δ is the same as the one of f but with all points within ℓ^∞ -distance δ of the diagonal removed. Edelsbrunner, Morozov, and Pascucci [120] gave a constructive proof that such δ -simplifications always exist when the underlying space of the complex K is a 2-manifold, possibly with boundary. They also showed that $\|f_\delta - f\|_\infty$ may have to be arbitrarily close to δ . In this context, Attali et al. [11] then Bauer, Lange, and Wardetzky [19] gave near-linear time (in the complex size) algorithms to compute such δ -approximations, by exploiting the duality between the 0-dimensional and 1-dimensional parts of the persistence diagrams—recall the EP Symmetry Theorem 2.9 from Chapter 2. Attali et al. [11] also showed that no such δ -simplification (or any constant-factor approximation) may exist when the underlying space of the complex K is a 3-manifold, even as simple as the 3-sphere. Thus, topological simplification turns out to be among those topological problems that are easy for surfaces but hard—if at all feasible—for manifolds of higher dimensions.

Attali et al. [10] pointed out an interesting connection to the reconstruction problem. Given a point cloud P in \mathbb{R}^d , supposedly sampled from some unknown object X , we want not only to infer the topology of X from P , but also to build a ‘faithful’ reconstruction. Typically, the reconstruction would have to be in the same homeomorphism class as X , which is useful e.g. for deriving global parametrizations of X . However, producing such reconstructions usually requires stringent conditions on X (such as being a submanifold with positive reach) and an immense sampling size [28, 31, 77]. In weaker versions of the problem, the reconstruction is merely asked to have the same homotopy type as X , or even just the same homology. Chazal and Oudot [66] showed how to produce a nested pair of simplicial complexes $K(P) \subseteq K'(P)$ (in fact, Rips or witness complexes) such that the inclusion map induces a morphism of the same rank as the corresponding homology group of X —recall Theorem 5.3 from Chapter 5. The question then was whether one can find an intermediate complex $K(P) \subseteq K'' \subseteq K'(P)$ with the same homology as X . Attali et al. [10] showed that such an intermediate complex K'' does not always exist, and that determining whether it does is NP-hard, even in \mathbb{R}^3 , by a reduction from 3-SAT. They also reduced the problem to the one of simplifying real-valued functions on \mathbb{S}^3 , thus deriving a NP-hardness result for that other problem as well.

OPEN QUESTIONS. *What would be a good trade-off between full homeomorphic reconstruction and mere homology inference? How can persistence help in achieving this trade-off?*

Introduction to Quiver Theory with a View Toward Persistence

Known as the *graphic method* within the larger picture of representation theory for associative algebras, quiver theory has been studied for more than four decades now, so the reader will find many introductions in the literature¹. These introductions are meant to be very general, with an emphasis on the richness of the subject and on its multiple connections to other areas of mathematics—excluding persistence, which is a comparatively recent topic. By contrast, this appendix is meant to be more focused, and to give a view on the subject that is clearly oriented toward persistence. This means in particular introducing only the concepts and results that are most relevant to persistence, illustrating them through carefully-chosen examples, and proving the results in some special cases once again related to persistence.

On the whole, the format of the appendix is that of a short course rather than a survey, which means that the exposition goes into a certain level of detail. This is required in order to exhibit some of the connections between persistence and quiver theory. For instance, the *reflection functors* [24] from quiver theory are closely related to the *Diamond Principle* [49] from persistence theory, but as we will see in Example A.16, the connection happens fairly deep down in the algebra.

The appendix is organized as follows. We define quivers in Section 1 and their representations in Section 2. We then introduce the classification problem and state Gabriel's theorem in Section 3. In Section 4 we give a simple proof of the theorem in the special case of A_n -type quivers. We then present a proof in the general case in Section 5, mostly for completeness. Finally, in Section 6 we present various extensions of the theorem to more general classes of quivers and representations. The progression is the same as in Section 1 of Chapter 1, from more general to more specifically connected to persistence.

Prerequisites. The content of this appendix takes on a very algebraic flavor, so the learning curve may look somewhat steep to some. Nevertheless, no prior exposure to quiver representation theory is required. Only a reasonable background in abstract and commutative algebra is needed, corresponding roughly to Parts I through III of [111]. Also, some basic notions of category theory, corresponding roughly to Chapters I and VIII of [184], can be helpful although they are not strictly required.

¹The content of this appendix was compiled from a number of sources, including [36, 93, 104, 172]. See also [106] for a high-level introduction.

1. Quivers

DEFINITION A.1. A *quiver* \mathbb{Q} consists of two sets Q_0, Q_1 and two maps $h, t : Q_1 \rightarrow Q_0$. The elements of Q_0 are called the *vertices* of \mathbb{Q} , while those of Q_1 are called the *arrows*. The *head map* h and *tail map* t assign a head h_a and a tail t_a to every arrow $a \in Q_1$.

Graphically, \mathbb{Q} can be represented as a directed graph with one vertex per element in Q_0 and one edge (t_a, h_a) per element $a \in Q_1$. Note that there are no restrictions on the sets Q_0, Q_1 , so technically this graph is a multigraph (see Figure A.1 for an illustration), possibly infinite.



FIGURE A.1. Left: the quiver with $Q_0 = \{1, 2, 3\}$, $Q_1 = \{a, b, c, d, e\}$, $h : (a, b, c, d, e) \mapsto (2, 2, 3, 1, 1)$, $t : (a, b, c, d, e) \mapsto (1, 2, 2, 3, 3)$. Right: the underlying undirected graph.

\mathbb{Q} is called *finite* if both sets Q_0, Q_1 are finite. It is a common practice to identify the quiver \mathbb{Q} with its graph representation, which we will do in the following to simplify the exposition. We denote by $\bar{\mathbb{Q}}$ the underlying undirected graph of \mathbb{Q} . If $\bar{\mathbb{Q}}$ is (dis-)connected, then \mathbb{Q} is called a (dis-)connected quiver. If $\bar{\mathbb{Q}}$ is one of the diagrams of Figure A.2, then \mathbb{Q} is called a *Dynkin* quiver. As we will see throughout the appendix, Dynkin quivers play a distinguished role in the theory.

REMARK. The diagrams of Figure A.2 are but a subset of the Dynkin diagrams, which include also B_n ($n \geq 1$), C_n ($n \geq 1$), F_4 and G_2 —these are the Dynkin diagrams that contain multiple edges. Interestingly, Dynkin diagrams also play a prominent role in the classification of simple Lie groups, as do a number of concepts introduced in the following pages.

2. The category of quiver representations

DEFINITION A.2. A *representation* of \mathbb{Q} over a field \mathbf{k} is a pair $\mathbb{V} = (V_i, v_a)$ consisting of a set of \mathbf{k} -vector spaces $\{V_i \mid i \in Q_0\}$ together with a set of \mathbf{k} -linear maps $\{v_a : V_{t_a} \rightarrow V_{h_a} \mid a \in Q_1\}$.

Note that the vector spaces and linear maps in \mathbb{V} can be arbitrary. In particular, no composition law is enforced, so if $a, b, c \in Q_1$ are such that $t_c = t_a$, $h_c = h_b$ and $h_a = t_b$, then v_c does not have to be equal to the composition $v_b \circ v_a$. In addition, the spaces V_i can be infinite-dimensional. We say that \mathbb{V} is *finite-dimensional* if the sum of the dimensions of its constituent spaces V_i is finite.

DEFINITION A.3. $\mathbb{W} = (W_i, w_a)$ is a *subrepresentation* of $\mathbb{V} = (V_i, v_a)$ if W_i is a subspace of V_i for all $i \in Q_0$ and if w_a is the restriction of the map v_a to the subspace W_{t_a} for all $a \in Q_1$.

Morphisms between representations of \mathbb{Q} are defined as follows.

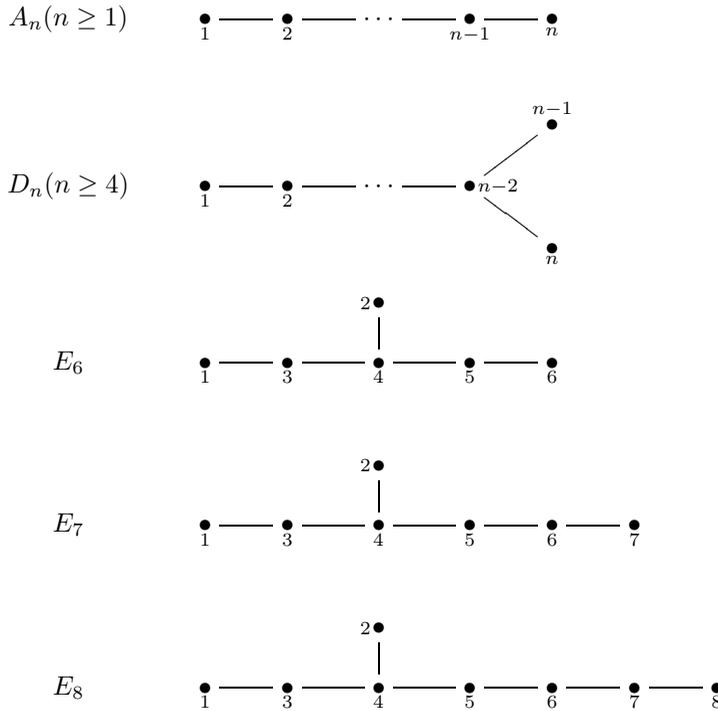


FIGURE A.2. The Dynkin diagrams.

DEFINITION A.4. A *morphism* ϕ between two \mathbf{k} -representations \mathbb{V}, \mathbb{W} of \mathbb{Q} is a set of \mathbf{k} -linear maps $\phi_i : V_i \rightarrow W_i$ such that the following diagram commutes for every arrow $a \in Q_1$:

$$\begin{array}{ccc}
 V_{t_a} & \xrightarrow{v_a} & V_{h_a} \\
 \phi_{t_a} \downarrow & & \downarrow \phi_{h_a} \\
 W_{t_a} & \xrightarrow{w_a} & W_{h_a}
 \end{array}$$

The morphism is called a *monomorphism* if every linear map ϕ_i is injective, an *epimorphism* if every ϕ_i is surjective, and an *isomorphism* (denoted \cong) if every ϕ_i is bijective.

Morphisms between representations are composed *pointwise*, by composing the linear maps at each vertex i independently. That is, the composition of $\phi : \mathbb{U} \rightarrow \mathbb{V}$ with $\psi : \mathbb{V} \rightarrow \mathbb{W}$ is the morphism $\psi \circ \phi : \mathbb{U} \rightarrow \mathbb{W}$ defined by $(\psi \circ \phi)_i = \psi_i \circ \phi_i$ for all $i \in Q_0$. Moreover, for each representation \mathbb{V} we have the *identity morphism* $\mathbb{1}_{\mathbb{V}} : \mathbb{V} \rightarrow \mathbb{V}$ defined by $(\mathbb{1}_{\mathbb{V}})_i = \mathbb{1}_{V_i}$ for all $i \in Q_0$. This turns the representations of \mathbb{Q} into a category, called $\text{Rep}_{\mathbf{k}}(\mathbb{Q})$. The subcategory of finite-dimensional representations is called $\text{rep}_{\mathbf{k}}(\mathbb{Q})$. Both categories are abelian², in particular:

- They contain a zero object, called the *trivial representation*, with all spaces and all maps equal to 0.

²This follows from the facts that the category of vector spaces over \mathbf{k} is itself abelian, and that compositions of morphisms within $\text{Rep}_{\mathbf{k}}(\mathbb{Q})$ are done *pointwise*.

- They have a biproduct, called a *direct sum*, and defined for any \mathbb{V}, \mathbb{W} as the representation $\mathbb{V} \oplus \mathbb{W}$ with spaces $V_i \oplus W_i$ for $i \in Q_0$ and maps $v_a \oplus w_a = \begin{pmatrix} v_a & 0 \\ 0 & w_a \end{pmatrix}$ for $a \in Q_1$. A nontrivial representation \mathbb{V} is called *decomposable* if it is isomorphic to the direct sum of two nontrivial representations (called *summands*), and *indecomposable* otherwise.
- Every morphism $\phi : \mathbb{V} \rightarrow \mathbb{W}$ has a *kernel*, defined by $(\ker \phi)_i = \ker \phi_i$ for all $i \in Q_0$. Similarly, ϕ has an *image* and a *cokernel*, also defined *pointwise*.
- A morphism ϕ is a monomorphism if and only if $\ker \phi = 0$, an epimorphism if and only if $\operatorname{coker} \phi = 0$, and an isomorphism if and only if ϕ is both a monomorphism and an epimorphism.

Thus, many of the known properties of vector spaces carry over to quiver representations. Among the ones that do not carry over, let us mention semisimplicity: indeed, not all subrepresentations of a given representation \mathbb{V} may be summands of \mathbb{V} . For instance, if $\mathbb{V} = \mathbf{k} \xrightarrow{1} \mathbf{k}$ and $\mathbb{W} = 0 \xrightarrow{0} \mathbf{k}$, then \mathbb{W} is a subrepresentation of \mathbb{V} yet it can be easily checked that there is no subrepresentation \mathbb{U} such that $\mathbb{V} = \mathbb{U} \oplus \mathbb{W}$. Broadly speaking, such obstructions are consequences of quiver representations being modules over certain types of associative algebras (see Section 6.2), and as such not enjoying all the properties of vector spaces, including semisimplicity. This is what makes the classification of quiver representations a challenging problem, as we will see next.

3. Classification of quiver representations

One of the central questions in quiver theory is to classify the representations of a given quiver up to isomorphism. Much of the literature on the subject focuses on finite-dimensional representations of finite quivers, as even in this simple setting things go wild pretty quickly, as we shall see in Section 5.2. Extensions of the theory to infinite quivers and/or infinite-dimensional representations exist, some of which will be mentioned briefly in Section 6, and in fact this is still an active area of research. For now and until Section 6, every quiver \mathbf{Q} will be finite and $\operatorname{rep}_{\mathbf{k}}(\mathbf{Q})$ will be the category under consideration. We will also assume without loss of generality that \mathbf{Q} is connected, since otherwise (i.e. when \mathbf{Q} is the disjoint union of two quivers \mathbf{Q}' and \mathbf{Q}'') $\operatorname{rep}_{\mathbf{k}}(\mathbf{Q})$ is isomorphic to the product category $\operatorname{rep}_{\mathbf{k}}(\mathbf{Q}') \times \operatorname{rep}_{\mathbf{k}}(\mathbf{Q}'')$, that is, any representation of \mathbf{Q} is the same as a pair of representations, one of \mathbf{Q}' , the other of \mathbf{Q}'' , and any morphism between representations of \mathbf{Q} is the same as a pair of morphisms acting on each component separately.

Decompositions. Let the vertex set of \mathbf{Q} be $Q_0 = \{1, \dots, n\}$. Given a representation $\mathbb{V} \in \operatorname{rep}_{\mathbf{k}}(\mathbf{Q})$, we define its *dimension vector* $\underline{\dim} \mathbb{V}$ and its *dimension* $\dim \mathbb{V}$ as follows:

$$\underline{\dim} \mathbb{V} = (\dim V_1, \dots, \dim V_n)^\top,$$

$$\dim \mathbb{V} = \|\underline{\dim} \mathbb{V}\|_1 = \sum_{i=1}^n \dim V_i.$$

An easy induction on the dimension shows that \mathbb{V} has a Remak decomposition, that is, \mathbb{V} can be decomposed into a direct sum of finitely many indecomposable representations. Moreover, it can be shown that this decomposition is unique up

to isomorphism and permutation of the terms in the direct sum, so the category $\text{rep}_{\mathbf{k}}(\mathbf{Q})$ has the Krull-Schmidt property.

THEOREM A.5 (Krull, Remak, Schmidt). *Assuming \mathbf{Q} is finite, for any $\mathbb{V} \in \text{rep}_{\mathbf{k}}(\mathbf{Q})$ there are indecomposable representations $\mathbb{V}_1, \dots, \mathbb{V}_r$ such that $\mathbb{V} \cong \mathbb{V}_1 \oplus \dots \oplus \mathbb{V}_r$. Moreover, for any indecomposable representations $\mathbb{W}_1, \dots, \mathbb{W}_s$ such that $\mathbb{V} \cong \mathbb{W}_1 \oplus \dots \oplus \mathbb{W}_s$, one has $r = s$ and there is a permutation σ such that $\mathbb{V}_i \cong \mathbb{W}_{\sigma(i)}$ for $1 \leq i \leq r$.*

REMARK. The decomposability of a quiver representation \mathbb{V} is closely tied to the locality³ of its endomorphism ring $\text{End}(\mathbb{V})$. Indeed, when \mathbb{V} is decomposable into a direct-sum of nontrivial subrepresentations $\mathbb{U} \oplus \mathbb{W}$, the canonical projections onto \mathbb{U} and \mathbb{W} are non-invertible endomorphisms of \mathbb{V} whose sum is $\mathbf{1}_{\mathbb{V}}$ and therefore invertible, thus $\text{End}(\mathbb{V})$ is not a local ring. Conversely, the locality of $\text{End}(\mathbb{V})$ when \mathbb{V} is indecomposable and finite-dimensional follows from Fitting's decomposition theorem—see e.g. [172, §2] for the details, and note that the finite dimensionality of \mathbb{V} is essential for the proof. From there, the uniqueness of the decomposition in Theorem A.5 is a direct consequence of Azumaya's theorem⁴ [14].

Theorem A.5 turns the original classification problem into that of classifying the indecomposable representations of \mathbf{Q} . The difficulty of the task resides in characterizing the indecomposable representations of \mathbf{Q} , which is different from (and significantly more difficult than) characterizing the representations with no proper subrepresentation. Indeed, as we saw earlier, not all subrepresentations are summands, so there are indecomposable representations with proper subrepresentations. Recall for instance our previous example $\mathbb{V} = \mathbf{k} \xrightarrow{\mathbf{1}} \mathbf{k}$, which is indecomposable yet admits $\mathbb{W} = 0 \xrightarrow{0} \mathbf{k}$ as a subrepresentation (but not as a summand).

The first major advance on this classification problem was made by Gabriel [133], and it was the starting point for the development of the quiver representation theory. Gabriel's result characterizes the so-called *finite type* quivers, i.e. the quivers that have a finite number of isomorphism classes of indecomposable finite-dimensional representations.

THEOREM A.6 (Gabriel I). *Let \mathbf{Q} be a finite connected quiver and let \mathbf{k} be a field. Then, \mathbf{Q} is of finite type if and only if \mathbf{Q} is Dynkin, i.e. if and only if $\bar{\mathbf{Q}}$ is one of the diagrams of Figure A.2.*

The fascinating thing about this result is that it does not rely on a particular choice of base field \mathbf{k} or arrow orientations. It introduces a dichotomy on the finite connected quivers, between those (very few) that have finite type and those that do not⁵.

This is but the first part of Gabriel's result. The second part (Theorem A.11) is a refinement of the 'if' statement, providing a complete characterization of the

³A ring is called local when it has a unique maximal proper left or right ideal, or equivalently, when the sum of any two non-invertible elements in the ring is itself non-invertible.

⁴Invoking Azumaya's theorem in the finite-dimensional setting is somewhat excessive, because the theorem holds in greater generality and can be replaced by a more direct dimension argument in this special case [172, §2]. Nevertheless, it explains why locality is desirable for endomorphism rings of indecomposable representations in general.

⁵There is in fact a trichotomy, as we will see in Section 6.1.

isomorphism classes of indecomposable representations of a Dynkin quiver Q by identifying them with the elements of a certain root system that we will now describe. The identification happens through the map $V \mapsto \underline{\dim} V$, which takes values in \mathbb{N}^n , so the analysis takes place in \mathbb{Z}^n , the free abelian group having Q_0 as basis.

Euler and Tits forms. A vector in \mathbb{Z}^n is called *positive* if it belongs to $\mathbb{N}^n \setminus \{0\}$, i.e. if it is nonzero and its coordinates are nonnegative. The dimension vectors of nontrivial representations of Q are such vectors.

DEFINITION A.7. The *Euler form* of Q is the bilinear form $\langle -, - \rangle_Q : \mathbb{Z}^n \times \mathbb{Z}^n \rightarrow \mathbb{Z}$ defined by

$$(A.1) \quad \langle x, y \rangle_Q = \sum_{i \in Q_0} x_i y_i - \sum_{a \in Q_1} x_{t_a} y_{h_a}.$$

Its symmetrization

$$(x, y)_Q = \langle x, y \rangle_Q + \langle y, x \rangle_Q$$

is called the *symmetric Euler form*.

Treating elements in \mathbb{Z}^n as column vectors, we can rewrite the symmetric Euler form as follows:

$$(x, y)_Q = x^\top C_Q y,$$

where $C_Q = (c_{ij})_{i,j \in Q_0}$ is the symmetric matrix whose entries are

$$c_{ij} = \begin{cases} 2 - 2|\{\text{loops at } i\}| & \text{if } i = j; \\ -|\{\text{arrows between } i \text{ and } j\}| & \text{if } i \neq j. \end{cases}$$

DEFINITION A.8. The *Tits form* is the quadratic form q_Q associated with the Euler form, that is:

$$q_Q(x) = \langle x, x \rangle_Q = \frac{1}{2} (x, x)_Q = \frac{1}{2} x^\top C_Q x.$$

Note that neither the symmetric Euler form nor the Tits form depend on the orientations of the arrows in Q . Their interest for us lies in their ability to distinguish Dynkin quivers from the rest of the connected quivers:

THEOREM A.9. *A finite connected quiver Q is Dynkin if and only if its Tits form q_Q is positive definite (i.e. $q_Q(x) > 0$ for any nonzero vector $x \in \mathbb{Z}^n$).*

PROOF OUTLINE. The “only if” part of the statement is proved easily by inspection. For instance, considering a quiver Q of type A_n (see Figure A.2), we have:

$$(A.2) \quad \begin{aligned} q_Q(x) &= \sum_{i \in Q_0} x_i^2 - \sum_{a \in Q_1} x_{t_a} x_{h_a} = \sum_{i=1}^n x_i^2 - \sum_{i=1}^{n-1} x_i x_{i+1} \\ &= \sum_{i=1}^{n-1} \frac{1}{2} (x_i - x_{i+1})^2 + \frac{1}{2} x_1^2 + \frac{1}{2} x_n^2, \end{aligned}$$

which is non-negative, and which is equal to zero only when all the terms in the sum are zero, that is, when $x_1 = x_2 = \dots = x_{n-1} = x_n = 0$.

The “if” part of the statement is proved by contraposition: assuming Q is not Dynkin, we find a nonzero vector $x \in \mathbb{Z}^n$ such that $q_Q(x) \leq 0$. The key observation is that the underlying undirected graph \bar{Q} must contain one of the diagrams of Figure A.3 as a subgraph. Now, for every such diagram, letting Q' be

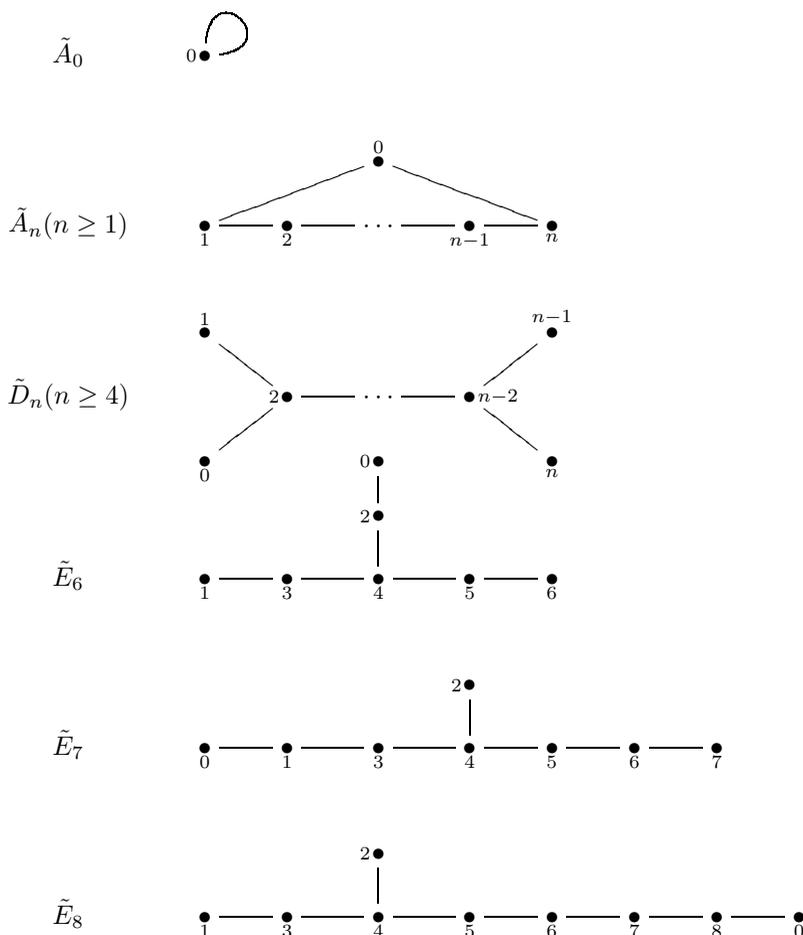


FIGURE A.3. The Euclidean diagrams.

the corresponding subquiver of \mathbb{Q} , it is easy to check by inspection the existence of a nonzero vector $x' \in \mathbb{Z}^{|Q'_0|}$ such that $q_{Q'}(x') = 0$ (take for instance $x = (1, \dots, 1)^\top$ when $\bar{Q}' = \tilde{A}_n$). Then, one of the following three scenarios occurs:

- $Q' = \mathbb{Q}$, in which case letting $x = x'$ gives $q_{\mathbb{Q}}(x) = q_{Q'}(x') = 0$, and so $q_{\mathbb{Q}}$ is not positive definite.
- $Q'_0 = Q_0$ and $Q'_1 \subsetneq Q_1$, in which case letting $x = x'$ gives $q_{\mathbb{Q}}(x) < q_{Q'}(x') = 0$ by definition of the Tits form, so $q_{\mathbb{Q}}$ is indefinite.
- $Q'_0 \subsetneq Q_0$, in which case let i be a vertex of $\mathbb{Q} \setminus Q'$ that is connected to Q' by an edge a (this vertex exists since \mathbb{Q} is connected), and let $x = 2x' + b_i$ where b_i is the i -th basis vector in \mathbb{Z}^n . This gives $q_{\mathbb{Q}}(x) \leq 4q_{Q'}(x') + x_i^2 - x_{t_a}x_{h_a} = 0 + 1 - 2 < 0$, so once again $q_{\mathbb{Q}}$ is indefinite.

□

A simple adaptation of the proof shows the following result as well—see [172, §4]:

Assuming \mathbf{Q} is finite and connected, its Tits form $q_{\mathbf{Q}}$ is positive semidefinite (i.e. $q_{\mathbf{Q}}(x) \geq 0$ for all $x \in \mathbb{Z}^n$) if and only if $\bar{\mathbf{Q}}$ belongs to the diagrams of Figures A.2 and A.3.

Combined with Theorem A.9, this result classifies the finite connected quivers into 3 categories: the *Dynkin* quivers, whose underlying undirected graph is Dynkin and whose Tits form is positive definite; the *tame* quivers, whose underlying undirected graph is Euclidean and whose Tits form is positive semidefinite but not definite; the rest, called *wild* quivers, whose Tits form is indefinite. The second and third categories are considered as one in Gabriel's theorem, however they play an important part in the theory beyond, as we will see in Section 6.1 where the terms *tame* and *wild* find their justification.

Roots. A *root* of $q_{\mathbf{Q}}$ is any nonzero vector $x \in \mathbb{Z}^n$ such that $q_{\mathbf{Q}}(x) \leq 1$, and we denote by $\Phi_{\mathbf{Q}}$ the set of roots of \mathbf{Q} . When \mathbf{Q} is a Dynkin quiver, $q_{\mathbf{Q}}$ is positive definite, so $\Phi_{\mathbf{Q}} = \{x \in \mathbb{Z}^n \mid q_{\mathbf{Q}}(x) = 1\}$ and $\Phi_{\mathbf{Q}}$ enjoys the properties of a root system, among which the following one plays a key part in Gabriel's theorem:

PROPOSITION A.10. *If \mathbf{Q} is Dynkin, then $\Phi_{\mathbf{Q}}$ is finite.*

PROOF OUTLINE. We can view $q_{\mathbf{Q}}$ indifferently as a quadratic form on \mathbb{Z}^n , \mathbb{Q}^n or \mathbb{R}^n . Since $q_{\mathbf{Q}}$ is positive definite on \mathbb{Z}^n , it is also positive definite on \mathbb{Q}^n , and by taking limits it is positive semidefinite on \mathbb{R}^n . But since $q_{\mathbf{Q}}$ is positive definite on \mathbb{Q}^n , its matrix is invertible in \mathbb{Q} and therefore also in \mathbb{R} , so $q_{\mathbf{Q}}$ is also positive definite on \mathbb{R}^n . Then, its sublevel set $\{x \in \mathbb{R}^n \mid q_{\mathbf{Q}}(x) \leq 1\}$ is an ellipsoid, which implies that it is bounded and so its intersection with the integer lattice \mathbb{Z}^n is finite. \square

THEOREM A.11 (Gabriel II). *Suppose \mathbf{Q} is a Dynkin quiver with n vertices. Then, the map $\mathbb{V} \mapsto \underline{\dim} \mathbb{V}$ induces a bijection between the set of isomorphism classes of indecomposable representations of \mathbf{Q} and the set $\Phi_{\mathbf{Q}} \cap \mathbb{N}^n$ of positive roots of $q_{\mathbf{Q}}$. In particular, the set of isomorphism classes of indecomposable representations of \mathbf{Q} is finite, by Proposition A.10.*

This is a refinement of the “if” part of Theorem A.6. Once again, the result is independent of the choices of base field \mathbf{k} and arrow orientations. It implies in particular that there is at most one isomorphism class of indecomposable representations per dimension vector. Intuitively, this means that once the dimensions of the vector spaces in the representation are fixed, there is only one choice (at most) of linear maps between the spaces to make the representation indecomposable, up to isomorphism.

EXAMPLE A.12 (A_n). Let us use Theorem A.11 to identify the isomorphism classes of indecomposable representations of a quiver \mathbf{Q} of type A_n . It follows from (A.2) that every positive root of $q_{\mathbf{Q}}$ is of the form $(0, \dots, 0, 1, \dots, 1, 0, \dots, 0)^{\top}$, with the first and last 1's occurring at some positions $b \leq d$ within the range $[1, n]$. The corresponding indecomposable representations \mathbb{V} of \mathbf{Q} have a space isomorphic to \mathbf{k} at every vertex $i \in [b, d]$, and a zero space at every other vertex. The maps from or to zero spaces are trivially zero, while the maps between consecutive copies of \mathbf{k} are isomorphisms because otherwise \mathbb{V} could be further decomposed. Thus, every such \mathbb{V} is isomorphic to the following representation, called an *interval*

representation and denoted $\mathbb{I}_Q[b, d]$:

$$(A.3) \quad \underbrace{0 \xrightarrow{0} \dots \xrightarrow{0} 0}_{[1, b-1]} \xrightarrow{0} \underbrace{k \xrightarrow{1} \dots \xrightarrow{1} k}_{[b, d]} \xrightarrow{0} \underbrace{0 \xrightarrow{0} \dots \xrightarrow{0} 0}_{[d+1, n]}$$

Sections 4 and 5 below are devoted to the proof of Gabriel’s theorem. Several proofs coexist in the literature, including Gabriel’s original proof. The one presented below is due to Bernstein, Gelfand, and Ponomarev [24] and is of a special interest to us. First of all, it emphasizes clearly the fact that the indecomposable representations are determined by their sole dimension vectors. Second, it uses reflection functors, which we know have close connections to persistence.

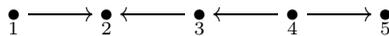
In Section 4 we introduce the reflection functors and give a simple proof of Gabriel’s theorem in the special case of A_n -type quivers, which we connect to the work of Carlsson and de Silva [49]. For the interested reader, Section 5 provides a proof of Gabriel’s theorem in the general case.

4. Reflections

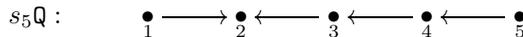
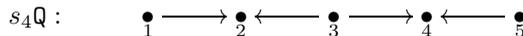
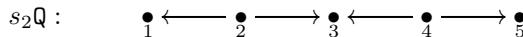
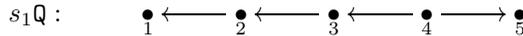
Let Q be a finite connected quiver. We call a vertex $i \in Q_0$ a *sink* if all arrows incident to i are incoming, that is, if there is no arrow $a \in Q_1$ such that $t_a = i$. Symmetrically, we call i a *source* if all arrows incident to i are outgoing, that is, if there is no arrow $a \in Q_1$ such that $h_a = i$. In particular, an isolated vertex is both a sink and a source.

For each sink or source vertex $i \in Q_0$, we denote by s_i the *reflection* at vertex i . More precisely, $s_i Q$ is the quiver obtained from Q by reversing the direction of all arrows incident to i .

EXAMPLE A.13. Let Q be the following quiver:



Then, vertices 1 and 4 are sources, vertices 2 and 5 are sinks, and vertex 3 is neither a source nor a sink. The corresponding reflections give the following quivers:



Notice how s_2 and s_4 turned vertex 3 into a sink and a source respectively. This property will be useful for defining Coxeter functors in Section 4.2.

4.1. Reflection functors. Let $i \in Q_0$ be a sink. For every representation $\mathbb{V} = (V_i, v_a) \in \text{rep}_{\mathbf{k}}(\mathbb{Q})$, we define a representation $\mathcal{R}_i^+ \mathbb{V} = (V'_i, v'_a) \in \text{rep}_{\mathbf{k}}(s_i \mathbb{Q})$ as follows. For all $j \neq i$, set $V'_j = V_j$, and define V'_i to be the kernel of the map

$$\xi_i : \bigoplus_{a \in Q_1^i} V_{t_a} \longrightarrow V_i, \quad (x_{t_a})_{a \in Q_1^i} \longmapsto \sum_{a \in Q_1^i} v_a(x_{t_a}),$$

where $Q_1^i = \{a \in Q_1 \mid h_a = i\}$ is the subset of arrows of \mathbb{Q} that are incident to vertex i (these are all incoming arrows). Define now linear maps between the spaces V'_j as follows. For each arrow $a \in Q_1$, if $a \notin Q_1^i$ then set $v'_a = v_a$; if $a \in Q_1^i$, then let b denote the reverse arrow, and set v'_b to be the composition

$$V'_{t_b} = V'_i = \ker \xi_i \hookrightarrow \bigoplus_{c \in Q_1^i} V_{t_c} \longrightarrow V_{t_a} = V'_{t_a} = V'_{h_b},$$

where the map before the direct sum is the canonical inclusion, and the map after the direct sum is the canonical projection onto its component V_{t_a} . Given now a morphism $\phi : \mathbb{V} \rightarrow \mathbb{W}$ between two representations $\mathbb{V}, \mathbb{W} \in \text{rep}_{\mathbf{k}}(\mathbb{Q})$, the morphism $\phi' = \mathcal{R}_i^+ \phi : \mathcal{R}_i^+ \mathbb{V} \rightarrow \mathcal{R}_i^+ \mathbb{W}$ is defined by $\phi'_j = \phi_j$ for all $j \neq i$ and by ϕ'_i being the restriction of the map

$$\bigoplus_{a \in Q_1^i} \phi_{t_a} : \bigoplus_{a \in Q_1^i} V_{t_a} \rightarrow \bigoplus_{a \in Q_1^i} W_{t_a}$$

to $V'_i = \ker \xi_i$. It can be readily checked that these mappings define a functor from the category $\text{rep}_{\mathbf{k}}(\mathbb{Q})$ to $\text{rep}_{\mathbf{k}}(s_i \mathbb{Q})$.

Dually, given a source vertex $i \in Q_0$, to every representation $\mathbb{V} = (V_i, v_a) \in \text{rep}_{\mathbf{k}}(\mathbb{Q})$ we associate a representation $\mathcal{R}_i^- \mathbb{V} = (V'_i, v'_a)$ as follows. For all $j \neq i$, let $V'_j = V_j$, and define V'_i as the cokernel of the map

$$\zeta_i : V_i \longrightarrow \bigoplus_{a \in Q_1^i} V_{h_a}, \quad x \longmapsto (v_a(x))_{a \in Q_1^i},$$

where $Q_1^i = \{a \in Q_1 \mid t_a = i\}$ denotes once again the subset of arrows of \mathbb{Q} that are incident to i (these are all outgoing arrows). Define now linear maps between the spaces V'_j as follows. For each arrow $a \in Q_1$, if $a \notin Q_1^i$ then set $v'_a = v_a$; if $a \in Q_1^i$, then let b denote the reverse arrow, and set v'_b to be the composition

$$V'_{t_b} = V'_{h_a} = V_{h_a} \hookrightarrow \bigoplus_{c \in Q_1^i} V_{h_c} \longrightarrow \text{coker } \zeta_i = V'_i = V'_{h_b},$$

where the map before the direct sum is the canonical inclusion, and the map after the direct sum is the canonical quotient map (i.e. the quotient map modulo the image of ζ_i). Finally, given a morphism $\phi : \mathbb{V} \rightarrow \mathbb{W}$ between two representations $\mathbb{V}, \mathbb{W} \in \text{rep}_{\mathbf{k}}(\mathbb{Q})$, the morphism $\phi' = \mathcal{R}_i^- \phi : \mathcal{R}_i^- \mathbb{V} \rightarrow \mathcal{R}_i^- \mathbb{W}$ is defined by $\phi'_j = \phi_j$ for all $j \neq i$ and by ϕ'_i being the map induced by

$$\bigoplus_{a \in Q_1^i} \phi_{h_a} : \bigoplus_{a \in Q_1^i} V_{h_a} \rightarrow \bigoplus_{a \in Q_1^i} W_{h_a}$$

on the quotient space $V'_i = \text{coker } \zeta_i$. Once again, these mappings define a functor from the category $\text{rep}_{\mathbf{k}}(\mathbb{Q})$ to $\text{rep}_{\mathbf{k}}(s_i \mathbb{Q})$.

The question that comes to mind is whether the functors \mathcal{R}_i^+ and \mathcal{R}_i^- are mutual inverses. It turns out that this is not the case, as each one of them nullifies the so-called *simple representation* \mathbb{S}_i , which has zero vector spaces everywhere except at vertex i where it has the space \mathbf{k} . Nevertheless, \mathcal{R}_i^+ and \mathcal{R}_i^- do become mutual inverses if we restrict ourselves to the full subcategory of $\text{rep}_{\mathbf{k}}(\mathbb{Q})$ consisting of the representations of \mathbb{Q} having no summand isomorphic to \mathbb{S}_i . Let us illustrate these claims through a simple example.

EXAMPLE A.14. Let \mathbb{Q} be the quiver of Example A.13, and let \mathbb{V} be some representation of \mathbb{Q} :

$$V_1 \xrightarrow{v_a} V_2 \xleftarrow{v_b} V_3 \xleftarrow{v_c} V_4 \xrightarrow{v_d} V_5$$

Applying \mathcal{R}_5^+ to \mathbb{V} , we get the following representation:

$$V_1 \xrightarrow{v_a} V_2 \xleftarrow{v_b} V_3 \xleftarrow{v_c} V_4 \xrightarrow{\hookrightarrow} \ker v_d$$

Applying now \mathcal{R}_5^- to $\mathcal{R}_5^+\mathbb{V}$, we obtain:

$$V_1 \xrightarrow{v_a} V_2 \xleftarrow{v_b} V_3 \xleftarrow{v_c} V_4 \xrightarrow{\text{mod } \ker v_d} V_4 / \ker v_d$$

Thus, by the first isomorphism theorem, we have $\mathcal{R}_5^- \mathcal{R}_5^+ \mathbb{V} \cong \mathbb{V}$ whenever $V_5 = \text{im } v_d$. When this is not the case, we have $\mathbb{V} \cong \mathcal{R}_5^- \mathcal{R}_5^+ \mathbb{V} \oplus \mathbb{S}_5^r$, where \mathbb{S}_5^r is made of $r = \dim \text{coker } v_d$ copies of the simple representation \mathbb{S}_5 . In particular, if \mathbb{V} is composed only of copies of \mathbb{S}_5 , then $\mathcal{R}_5^- \mathcal{R}_5^+ \mathbb{V}$ is the trivial representation.

Let us now apply \mathcal{R}_2^+ to \mathbb{V} . The result is the top row in the following commutative diagram, where π_1, π_3 denote the canonical projections from the direct sum $V_1 \oplus V_3$ to its components:

$$\begin{array}{ccccccc} V_1 & \xleftarrow{\quad} & \ker v_a + v_b & \xrightarrow{\quad} & V_3 & \xleftarrow{v_c} & V_4 \xrightarrow{v_d} V_5 \\ & \searrow \pi_1 & \downarrow & \nearrow \pi_3 & & & \\ & & V_1 \oplus V_3 & & & & \end{array}$$

Applying now \mathcal{R}_2^- to $\mathcal{R}_2^+\mathbb{V}$, we obtain the top row in the following commutative diagram, where the vertical arrow is the canonical quotient map:

$$\begin{array}{ccccccc} V_1 & \xrightarrow{\quad} & \frac{V_1 \oplus V_3}{\ker v_a + v_b} & \xleftarrow{\quad} & V_3 & \xleftarrow{v_c} & V_4 \xrightarrow{v_d} V_5 \\ & \searrow (-,0) & \downarrow & \nearrow (0,-) & & & \\ & & \ker v_a + v_b & \xrightarrow{\quad} & V_1 \oplus V_3 & \xrightarrow{v_a + v_b} & V_2 \end{array}$$

Again, by the first isomorphism theorem, there is a unique injective map $\phi_2 : \frac{V_1 \oplus V_3}{\ker v_a + v_b} \rightarrow V_2$ such that the diagram augmented with ϕ_2 still commutes. Letting $\phi_i = \mathbf{1}_{V_i}$ for all $i \neq 2$, we obtain a monomorphism $\phi = (\phi_i)_{i \in Q_0} : \mathcal{R}_2^- \mathcal{R}_2^+ \mathbb{V} \rightarrow \mathbb{V}$. This monomorphism becomes an isomorphism whenever $V_2 = \text{im } v_a + v_b$, and when this is not the case, we have $\mathbb{V} \cong \mathcal{R}_2^- \mathcal{R}_2^+ \mathbb{V} \oplus \mathbb{S}_2^r$ where $r = \dim \text{coker } v_a + v_b$.

Bernstein, Gelfand, and Ponomarev [24] proved these empirical findings to be general properties of the reflection functors \mathcal{R}_i^\pm . Their proof is based on the same argument as the one used in Example A.14, and the interested reader may refer directly to their paper for the details, or to [104, theorem 1.18] for a more recent treatment. The statement of the theorem is reproduced below in a form that will be practical to us in the following.

THEOREM A.15 (Reflection Functors). *Let \mathbf{Q} be a finite connected quiver and let \mathbb{V} be a representation of \mathbf{Q} . If $\mathbb{V} \cong \mathbb{U} \oplus \mathbb{W}$, then for any source or sink $i \in Q_0$, $\mathcal{R}_i^\pm \mathbb{V} \cong \mathcal{R}_i^\pm \mathbb{U} \oplus \mathcal{R}_i^\pm \mathbb{W}$. If now \mathbb{V} is indecomposable:*

1. *If $i \in Q_0$ is a sink, then two cases are possible:*

- $\mathbb{V} \cong \mathbb{S}_i$: *in this case, $\mathcal{R}_i^+ \mathbb{V} = 0$.*
- $\mathbb{V} \not\cong \mathbb{S}_i$: *in this case, $\mathcal{R}_i^+ \mathbb{V}$ is nonzero and indecomposable, $\mathcal{R}_i^- \mathcal{R}_i^+ \mathbb{V} \cong \mathbb{V}$, and the dimension vectors x of \mathbb{V} and y of $\mathcal{R}_i^+ \mathbb{V}$ are related to each other by the following formula:*

$$(A.4) \quad y_j = \begin{cases} x_j & \text{if } j \neq i; \\ -x_i + \sum_{\substack{a \in Q_1 \\ h_a = i}} x_{t_a} & \text{if } j = i. \end{cases}$$

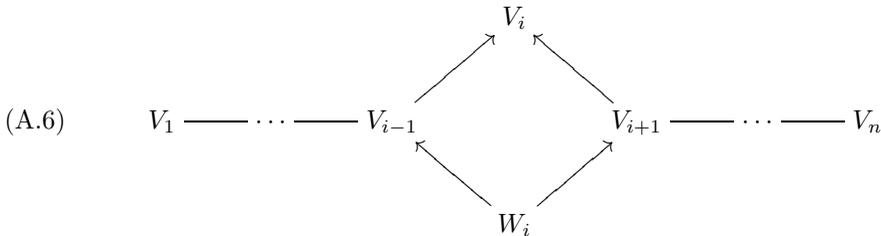
2. *If $i \in Q_0$ is a source, then once again two cases are possible:*

- $\mathbb{V} \cong \mathbb{S}_i$: *in this case, $\mathcal{R}_i^- \mathbb{V} = 0$.*
- $\mathbb{V} \not\cong \mathbb{S}_i$: *in this case, $\mathcal{R}_i^- \mathbb{V}$ is nonzero and indecomposable, $\mathcal{R}_i^+ \mathcal{R}_i^- \mathbb{V} \cong \mathbb{V}$, and the dimension vectors x of \mathbb{V} and y of $\mathcal{R}_i^- \mathbb{V}$ are related to each other by the following formula:*

$$(A.5) \quad y_j = \begin{cases} x_j & \text{if } j \neq i; \\ -x_i + \sum_{\substack{a \in Q_1 \\ t_a = i}} x_{h_a} & \text{if } j = i. \end{cases}$$

Before we proceed any further, let us see what Theorem A.15 entails for quivers of type A_n . In the example below we are assuming Gabriel’s theorem holds for the sake of the exposition.

EXAMPLE A.16. Let \mathbf{Q} be a quiver of type A_n , and let $i \in Q_0$ be a sink. Then, any representation $\mathbb{V} \in \text{rep}_k(\mathbf{Q})$ and its reflection $\mathbb{W} = \mathcal{R}_i^+ \mathbb{V}$ can be represented as follows since they share the same spaces $V_j = W_j$ for $j \neq i$:



By Gabriel’s theorem, we know that $\mathbb{V} \cong \bigoplus_{j=1}^r \mathbb{I}_{\mathbf{Q}}[b_j, d_j]$, where each $\mathbb{I}_{\mathbf{Q}}[b_j, d_j]$ is the interval representation associated with some interval $[b_j, d_j]$ according to (A.3). Then, Theorem A.15 says that $\mathcal{R}_i^+ \mathbb{V} \cong \bigoplus_{j=1}^r \mathcal{R}_i^+ \mathbb{I}_{\mathbf{Q}}[b_j, d_j]$, where each summand is

either zero or an interval representation, determined according to the following set of rules:

$$(A.7) \quad \mathcal{R}_i^+ \mathbb{I}_{\mathbb{Q}}[b_j, d_j] = \begin{cases} 0 & \text{if } i = b_j = d_j; \\ \mathbb{I}_{s_i \mathbb{Q}}[i + 1, d_j] & \text{if } i = b_j < d_j; \\ \mathbb{I}_{s_i \mathbb{Q}}[i, d_j] & \text{if } i + 1 = b_j < d_j; \\ \mathbb{I}_{s_i \mathbb{Q}}[b_j, i - 1] & \text{if } b_j < d_j = i; \\ \mathbb{I}_{s_i \mathbb{Q}}[b_j, i] & \text{if } b_j \leq d_j = i - 1; \\ \mathbb{I}_{s_i \mathbb{Q}}[b_j, d_j] & \text{otherwise.} \end{cases}$$

As we will see in Section 4.4, in the language of persistence theory, the quadrangle in (A.6) is called a *diamond*, and the above set of conversion rules between the decompositions of \mathbb{V} and of its reflection $\mathcal{R}_i^+ \mathbb{V}$ is known as the *Diamond Principle*.

A consequence of Theorem A.15 is that the value of the Tits form at the dimension vectors of indecomposable representations is either preserved or sent to zero by \mathcal{R}_i^\pm .

COROLLARY A.17. *Let \mathbb{Q} be a finite connected quiver and let \mathbb{V} be an indecomposable representation of \mathbb{Q} . Then, for any source or sink $i \in Q_0$,*

- *either $\mathbb{V} \cong \mathbb{S}_i$, in which case $q_{s_i \mathbb{Q}}(\underline{\dim} \mathcal{R}_i^\pm \mathbb{V}) = 0$,*
- *or $q_{s_i \mathbb{Q}}(\underline{\dim} \mathcal{R}_i^\pm \mathbb{V}) = q_{\mathbb{Q}}(\underline{\dim} \mathbb{V})$.*

The proof is a short calculation, which we reproduce below for completeness.

PROOF. If $\mathbb{V} \cong \mathbb{S}_i$, then by Theorem A.15 we have $\mathcal{R}_i^\pm \mathbb{V} = 0$, therefore $q_{s_i \mathbb{Q}}(\underline{\dim} \mathcal{R}_i^\pm \mathbb{V}) = 0$. Assume now that $\mathbb{V} \not\cong \mathbb{S}_i$, and let i be a sink without loss of generality, the case of a source being similar. Let also x denote $\underline{\dim} \mathbb{V}$ and y denote $\underline{\dim} \mathcal{R}_i^\pm \mathbb{V}$ for simplicity. Then, by (A.4) and Definition A.8 we have:

$$\begin{aligned} q_{s_i \mathbb{Q}}(y) &= \sum_{j \in Q_0} y_j^2 - \sum_{a \in Q_1} y_{t_a} y_{h_a} = y_i^2 + \sum_{j \neq i} y_j^2 - \sum_{\substack{a \in Q_1 \\ h_a = i}} y_{t_a} y_i - \sum_{\substack{a \in Q_1 \\ h_a \neq i}} y_{t_a} y_{h_a} \\ &= \left(-x_i + \sum_{\substack{a \in Q_1 \\ h_a = i}} x_{t_a} \right)^2 + \sum_{j \neq i} x_j^2 - \sum_{\substack{a \in Q_1 \\ h_a = i}} x_{t_a} \left(-x_i + \sum_{\substack{a \in Q_1 \\ h_a = i}} x_{t_a} \right) - \sum_{\substack{a \in Q_1 \\ h_a \neq i}} x_{t_a} x_{h_a} \\ &= x_i^2 - 2x_i \sum_{\substack{a \in Q_1 \\ h_a = i}} x_{t_a} + \left(\sum_{\substack{a \in Q_1 \\ h_a = i}} x_{t_a} \right)^2 + \sum_{j \neq i} x_j^2 + x_i \sum_{\substack{a \in Q_1 \\ h_a = i}} x_{t_a} - \left(\sum_{\substack{a \in Q_1 \\ h_a = i}} x_{t_a} \right)^2 - \sum_{\substack{a \in Q_1 \\ h_a \neq i}} x_{t_a} x_{h_a} \\ &= \sum_{j \in Q_0} x_j^2 - \sum_{a \in Q_1} x_{t_a} x_{h_a} = q_{\mathbb{Q}}(x). \end{aligned}$$

□

4.2. Coxeter functors. From now on we assume that the quiver \mathbb{Q} is acyclic, i.e. that it contains no oriented cycle. Then, \mathbb{Q} represents a partial order on its vertex set Q_0 . Take an arbitrary total order on Q_0 that is compatible with the partial order given by \mathbb{Q} , and relabel the elements in $Q_0 = \{1, \dots, n\}$ according to

that total order, so that we have $t_a < h_a$ for every arrow $a \in Q_1$. Consider now the following combinations of reflections on \mathbb{Q} , where the order of operations is from right to left:

$$c^+ = s_1 s_2 \cdots s_{n-1} s_n;$$

$$c^- = s_n s_{n-1} \cdots s_2 s_1.$$

Observe that vertex i is a sink in $(s_{i+1} \cdots s_n)\mathbb{Q}$ and a source in $(s_1 \cdots s_{i-1})\mathbb{Q}$, therefore c^+ and c^- are well-defined. Observe also that $c^\pm \mathbb{Q} = \mathbb{Q}$ since every arrow in \mathbb{Q} gets reversed twice. Thus, c^+ and c^- do not modify the quiver \mathbb{Q} , and their corresponding functors,

$$C^+ = \mathcal{R}_1^+ \mathcal{R}_2^+ \cdots \mathcal{R}_{n-1}^+ \mathcal{R}_n^+, \text{ and}$$

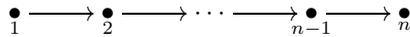
$$C^- = \mathcal{R}_n^- \mathcal{R}_{n-1}^- \cdots \mathcal{R}_2^- \mathcal{R}_1^-,$$

are endofunctors $\text{rep}_{\mathbf{k}}(\mathbb{Q}) \rightarrow \text{rep}_{\mathbf{k}}(\mathbb{Q})$. They are called *Coxeter functors* because of their connection to the Coxeter transformations in Lie group theory. The following property is an immediate consequence of Theorem A.15 and Corollary A.17:

COROLLARY A.18. *Let \mathbb{Q} be a finite, connected and acyclic quiver, and let C^\pm be defined as above. Then, for any indecomposable representation \mathbb{V} of \mathbb{Q} , either $C^\pm \mathbb{V}$ is indecomposable or $C^\pm \mathbb{V} = 0$. In the first case we have $q_{\mathbb{Q}}(\underline{\dim} C^\pm \mathbb{V}) = q_{\mathbb{Q}}(\underline{\dim} \mathbb{V})$, while in the second case we have $q_{\mathbb{Q}}(\underline{\dim} C^\pm \mathbb{V}) = 0$.*

Let us now see the effects of the Coxeter functors on the indecomposable representations of quivers of type A_n . For simplicity, let us begin with the quiver with all arrows oriented forwards, which is called the *linear quiver* and denoted L_n .

EXAMPLE A.19 (L_n). Consider the linear quiver L_n :



The natural order on the integers is the only one compatible with L_n . Given an indecomposable representation \mathbb{V} of L_n , let $(x_1, x_2, \dots, x_{n-1}, x_n)^\top = \underline{\dim} \mathbb{V}$, and apply the Coxeter functor C^+ corresponding to the sequence of reflections $c^+ = s_1 \cdots s_n$. By Theorem A.15, we have:

$$\underline{\dim} \mathcal{R}_n^+ \mathbb{V} = 0 \text{ or } (x_1, x_2, \dots, x_{n-1}, x_{n-1} - x_n)^\top,$$

$$\underline{\dim} \mathcal{R}_{n-1}^+ \mathcal{R}_n^+ \mathbb{V} = 0 \text{ or } (x_1, x_2, \dots, x_{n-2} - x_n, x_{n-1} - x_n)^\top,$$

$$\dots$$

$$\underline{\dim} \mathcal{R}_2^+ \cdots \mathcal{R}_{n-1}^+ \mathcal{R}_n^+ \mathbb{V} = 0 \text{ or } (x_1, x_1 - x_n, \dots, x_{n-2} - x_n, x_{n-1} - x_n)^\top,$$

$$\underline{\dim} C^+ \mathbb{V} = \underline{\dim} \mathcal{R}_1^+ \mathcal{R}_2^+ \cdots \mathcal{R}_{n-1}^+ \mathcal{R}_n^+ \mathbb{V} = 0 \text{ or } (-x_n, x_1 - x_n, \dots, x_{n-2} - x_n, x_{n-1} - x_n)^\top.$$

Thus, either $\underline{\dim} C^+ \mathbb{V} = 0$ or $x_n = 0$, since a dimension vector must have nonnegative coordinates. In the latter case, $\underline{\dim} C^+ \mathbb{V} = (0, x_1, x_2, \dots, x_{n-2}, x_{n-1})^\top$. By

iterating this process further, we get:

$$\begin{aligned}
 \underline{\dim} \mathcal{C}^+ \mathbb{V} &= 0 \text{ or } (0, x_1, x_2, \dots, x_{n-2}, x_{n-1})^\top, \\
 \underline{\dim} \mathcal{C}^+ \mathcal{C}^+ \mathbb{V} &= 0 \text{ or } (0, 0, x_1, \dots, x_{n-3}, x_{n-2})^\top, \\
 &\dots \\
 \underline{\dim} \underbrace{\mathcal{C}^+ \dots \mathcal{C}^+}_{n-1 \text{ times}} \mathbb{V} &= 0 \text{ or } (0, 0, 0, \dots, 0, x_1)^\top, \\
 \text{(A.8)} \quad \underline{\dim} \underbrace{\mathcal{C}^+ \dots \mathcal{C}^+}_n \mathbb{V} &= 0.
 \end{aligned}$$

Thus, $\mathcal{C}^+ \dots \mathcal{C}^+ \mathbb{V}$ eventually goes to 0 when \mathbb{V} is an indecomposable representation of the linear quiver. Bernstein, Gelfand, and Ponomarev [24] proved that this property carries over to A_n -type quivers, and beyond that, to any Dynkin quiver \mathbb{Q} —see Proposition A.24 in Section 5. However, one does not need such a strong result to proceed with the proof of Gabriel’s theorem. It is enough to find a combination of reflection functors that sends the indecomposable representations of \mathbb{Q} to 0. When \mathbb{Q} is of type A_n , such a sequence can be obtained easily through the following reduction to the linear quiver.

EXAMPLE A.20 (A_n). Let \mathbb{Q} be of type A_n . If \mathbb{Q} is not the linear quiver, then let $i_1 < i_2 < \dots < i_r$ be the heads of backward arrows. By performing the sequence of reflections $s_1 \dots s_{i_1}$ on \mathbb{Q} , we obtain the same quiver except that the arrow between i_1 and $i_1 + 1$ is now a forward arrow. By repeating this process for i_2, \dots, i_r , we finally get the linear quiver. The corresponding sequence of reflection functors sends every indecomposable representation \mathbb{V} of \mathbb{Q} either to 0 or to an indecomposable representation of the linear quiver. We can now apply the functor \mathcal{C}^+ a sufficient number of times to send the representation to 0, as in (A.8). To summarize:

$$\text{(A.9)} \quad \underbrace{\mathcal{C}^+ \dots \mathcal{C}^+}_n \mathcal{R}_1^+ \dots \mathcal{R}_{i_r}^+ \dots \mathcal{R}_1^+ \dots \mathcal{R}_{i_2}^+ \mathcal{R}_1^+ \dots \mathcal{R}_{i_1}^+ \mathbb{V} = 0.$$

4.3. Proof of Gabriel’s theorem for A_n -type quivers. We now have the required material to prove the “if” part of Theorem A.6 in the special case where \mathbb{Q} is a quiver of type A_n .

Let \mathbb{V} be an indecomposable representation of \mathbb{Q} . According to (A.9), there is a sequence of indices i_1, \dots, i_s , possibly with repetitions, such that $\mathcal{R}_{i_s}^+ \dots \mathcal{R}_{i_1}^+ \mathbb{V} = 0$. Assume without loss of generality that this is a minimal such sequence, so $\mathcal{R}_{i_{s-1}}^+ \dots \mathcal{R}_{i_1}^+ \mathbb{V} \neq 0$. Then, by Theorem A.15, $\mathcal{R}_{i_{s-1}}^+ \dots \mathcal{R}_{i_1}^+ \mathbb{V}$ is isomorphic to the simple representation \mathbb{S}_{i_s} , and every representation $\mathcal{R}_{i_j}^+ \dots \mathcal{R}_{i_1}^+ \mathbb{V}$ for $j < s$ is indecomposable. It follows by Corollary A.17 that

$$\begin{aligned}
 q_{\mathbb{Q}}(\underline{\dim} \mathbb{V}) &= q_{s_{i_1} \mathbb{Q}}(\underline{\dim} \mathcal{R}_{i_1}^+ \mathbb{V}) = \dots \\
 &\dots = q_{s_{i_{s-1}} \dots s_{i_1} \mathbb{Q}}(\underline{\dim} \mathcal{R}_{i_{s-1}}^+ \dots \mathcal{R}_{i_1}^+) = q_{s_{i_{s-1}} \dots s_{i_1} \mathbb{Q}}(\underline{\dim} \mathbb{S}_{i_s}) = 1.
 \end{aligned}$$

Thus, $\underline{\dim} \mathbb{V}$ is a positive root of $q_{\mathbb{Q}}$, which, according to (A.2), implies that

$$\underline{\dim} \mathbb{V} = (0, \dots, 0, 1, \dots, 1, 0, \dots, 0)^\top,$$

with the first and last 1’s occurring at some positions $b \leq d$ within the range $[1, n]$. So, \mathbb{V} has a space isomorphic to \mathbf{k} at every vertex $i \in [b, d]$ and a zero space at every other vertex. The maps from or to zero spaces are trivially zero, while the

maps between consecutive copies of \mathbf{k} are isomorphisms because otherwise \mathbb{V} could be further decomposed. Thus, \mathbb{V} is isomorphic to the interval representation $\mathbb{I}_{\mathbb{Q}}[b, d]$ introduced in (A.3).

To summarize, every isomorphism class of indecomposable representations of \mathbb{Q} contains at least one of the (finitely many) interval representations, while of course every interval representation belongs to at most one of the isomorphism classes, so the “if” part of Theorem A.6 is proved.

From there we can also prove Theorem A.11, by identifying the set of isomorphism classes of indecomposable representations of \mathbb{Q} with the set of interval representations. This boils down to showing (i) that different interval representations cannot be isomorphic, and (ii) that every interval representation is indecomposable. Item (i) is a direct consequence of the fact that different interval representations have different dimension vectors. Item (ii) follows from the fact that the endomorphism ring of an interval representation is isomorphic to the field \mathbf{k} and therefore local⁶.

4.4. Connection to the proof of Carlsson and de Silva [49]. We will first introduce the Diamond Principle and establish a direct link to reflection functors. We will then focus on the proof of Gabriel’s theorem and relate it to the one given in Section 4.3.

REMARK. Interestingly enough, reflection functors were introduced in [24] as a tool to prove Gabriel’s theorem, whereas the Diamond Principle came as a byproduct of the decomposition theorem in [49]. The version in [24] is stronger because it does not assume the existence of a decomposition.

4.4.1. *Diamond Principle.* Let \mathbb{Q} be an A_n -type quiver that has a sink i , and let $s_i\mathbb{Q}$ denote the quiver obtained from \mathbb{Q} by applying the reflection s_i at node i . Given two finite-dimensional representations, $\mathbb{V} \in \text{rep}_{\mathbf{k}}(\mathbb{Q})$ and $\mathbb{W} \in \text{rep}_{\mathbf{k}}(s_i\mathbb{Q})$, that differ only by the spaces V_i, W_i and their incident maps, we can form the following diagram where the central quadrangle is called a *diamond*:

$$(A.10) \quad \begin{array}{ccccccc} & & & & V_i & & \\ & & & & \swarrow & & \searrow \\ & & & & v_c & & v_d \\ & & & & & & \\ V_1 & \cdots & \cdots & V_{i-1} & & V_{i+1} & \cdots & \cdots & V_n \\ & & & \swarrow & & \searrow & & & \\ & & & w_a & & w_b & & & \\ & & & & W_i & & & & \end{array}$$

The diamond is said to be *exact* if $\text{im } f = \ker g$ in the following sequence

$$W_i \xrightarrow{f} V_{i-1} \oplus V_{i+1} \xrightarrow{g} V_i$$

where⁷ $f : x \mapsto (w_a(x), w_b(x))$ and $g : (x, y) \mapsto v_c(x) + v_d(y)$. Carlsson and de Silva [49] proved the following result:

⁶Item (ii) can also be proven directly using elementary linear algebra. See e.g. Proposition 2.4 in [49].

⁷For convenience we are departing slightly from [49], by letting g be an addition instead of a subtraction. This modification has no incidence on the result, since one can replace the map v_d by $-v_d$ in the representation \mathbb{V} and get an isomorphic representation.

THEOREM A.21 (Diamond Principle). *Given \mathbb{V} and \mathbb{W} as above, suppose that the diamond in (A.10) is exact. Then, the interval decompositions of \mathbb{V} and \mathbb{W} are related to each other through the following matching rules:*

- *summands $\mathbb{I}_{\mathbb{Q}}[i, i]$ and $\mathbb{I}_{s_i\mathbb{Q}}[i, i]$ are unmatched,*
- *summands $\mathbb{I}_{\mathbb{Q}}[i, d]$ are matched with summands $\mathbb{I}_{s_i\mathbb{Q}}[i + 1, d]$, and $\mathbb{I}_{\mathbb{Q}}[i + 1, d]$ with $\mathbb{I}_{s_i\mathbb{Q}}[i, d]$,*
- *summands $\mathbb{I}_{\mathbb{Q}}[b, i]$ are matched with summands $\mathbb{I}_{s_i\mathbb{Q}}[b, i - 1]$, and $\mathbb{I}_{\mathbb{Q}}[b, i - 1]$ with $\mathbb{I}_{s_i\mathbb{Q}}[b, i]$,*
- *every other summand $\mathbb{I}_{\mathbb{Q}}[b, d]$ is matched with $\mathbb{I}_{s_i\mathbb{Q}}[b, d]$.*

The similarity of this result with the one from Example A.16 is striking. In fact, it is easy to prove that both results imply each other in a natural way (see below), so the Diamond Principle can really be viewed as the expression of the Reflection Functors Theorem A.15 in the special case of A_n -type quivers.

PROOF. Proving that Theorem A.21 implies the result from Example A.16 is just a matter of checking that the diamond in (A.6) is exact, which is immediate using the definition of $\mathcal{R}_i^+\mathbb{V}$.

Let us now prove Theorem A.21 using the result from Example A.16. Given \mathbb{V}, \mathbb{W} as in (A.10) and assuming the diamond is exact, the claim is that \mathbb{W} is isomorphic to $\mathbb{U} \oplus \mathbb{K}$, where $\mathbb{U} = \mathcal{R}_i^+\mathbb{V}$ and where $\mathbb{K} = \bigoplus_{j=1}^r \mathbb{I}_{s_i\mathbb{Q}}[i, i]$ is the representation of $s_i\mathbb{Q}$ made of $r = \dim \ker f$ copies of the interval representation $\mathbb{I}_{s_i\mathbb{Q}}[i, i]$. The theorem follows from this claim and (A.7).

To prove the claim, recall from the definition of $\mathcal{R}_i^+\mathbb{V}$ and from our exactness hypothesis that we have $U_j = V_j$ for all $j \neq i$ and $U_i = \ker g = \text{im } f$. Therefore, given an arbitrary complement C of $\ker f$ in W_i , the first isomorphism theorem states that $f|_C$ is an isomorphism onto its image U_i . Now, let us pick an arbitrary isomorphism of vector spaces $h : \ker f \rightarrow K_i$, and let us define an isomorphism of representations $\phi : \mathbb{W} \rightarrow \mathbb{U} \oplus \mathbb{K}$ as follows:

$$\phi_j = \begin{cases} \mathbb{1}_{V_j} & \text{if } j \neq i; \\ f|_C \oplus h = \begin{pmatrix} f|_C & 0 \\ 0 & h \end{pmatrix} & \text{if } j = i. \end{cases}$$

Checking that ϕ is indeed a morphism composed of invertible linear maps is straightforward. □

4.4.2. Proof of Gabriel’s theorem for A_n -type quivers. Let us fix an A_n -type quiver \mathbb{Q} and focus on the finite-dimensional representations of \mathbb{Q} . Such a representation \mathbb{V} is called *(right-)streamlined* if every forward map is injective and every backward map is surjective. For instance, the representation shown in (1.11) is streamlined if every map v_{2k+1} is injective and every map v_{2k+2} is surjective.

Intuitively, for an interval-decomposable representation \mathbb{V} , to be streamlined means that there can be no interval summand of \mathbb{V} that ends before the final index n . This intuition is formalized in the following result by Carlsson and de Silva [49].

PROPOSITION A.22.

- *An interval representation $\mathbb{I}_{\mathbb{Q}}[b, d]$ is streamlined if and only if $d = n$.*
- *A direct sum $\bigoplus_{i=1}^k \mathbb{V}^i$ of representations of \mathbb{Q} is streamlined if and only if every summand \mathbb{V}^i is streamlined.*

Thus, an interval-decomposable representation \mathbb{V} of \mathbb{Q} is streamlined if and only if its interval decomposition is of type $\bigoplus_i \mathbb{I}_{\mathbb{Q}}[b_i, n]$.

By extension, a representation \mathbb{V} is streamlined *up to index i* if every forward map $V_j \rightarrow V_{j+1}$ ($j < i$) is injective and every backward map $V_k \leftarrow V_{k+1}$ ($k < i$) is surjective. In particular, every representation is streamlined up to index 1.

The proof of Carlsson and de Silva [49] proceeds in two steps:

- (1) it shows that every representation \mathbb{V} of \mathbb{Q} decomposes as a direct sum $\mathbb{V}^1 \oplus \dots \oplus \mathbb{V}^k$ where each summand \mathbb{V}^i is streamlined up to index i and trivial (with zero spaces and maps) beyond;
- (2) it shows that each summand \mathbb{V}^i decomposes in turn as a direct sum of interval representations $\bigoplus_j \mathbb{I}_{\mathbb{Q}}[b_j, i]$.

Step 1 is proved by a simple induction on the length n of \mathbb{Q} and does not bring any particular insight into the problem. Let us therefore refer to [49] for the details. Step 2 is more interesting. It boils down to showing that any streamlined representation $\mathbb{V} = V_1 \xrightarrow{v_1} V_2 \xrightarrow{v_2} \dots \xrightarrow{v_{n-1}} V_n$ of \mathbb{Q} decomposes as a direct sum $\bigoplus_i \mathbb{I}_{\mathbb{Q}}[b_i, n]$. For this, Carlsson and de Silva [49] introduce a new tool that proves very handy: the *right filtration* of \mathbb{V} .

DEFINITION A.23. the *right filtration* of \mathbb{V} , noted $\mathcal{R}_{\mathbb{V}}$, is a filtration (i.e. a nested sequence of subspaces) of the vector space V_n , defined recursively as follows:

- if $n = 1$, then $\mathcal{R}_{\mathbb{V}} = (0, V_1)$;
- else ($n > 1$),

$$(A.11) \quad \mathcal{R}_{\mathbb{V}} = \begin{cases} (v_{n-1}(R_0), \dots, v_{n-1}(R_{n-1}), V_n) & \text{if } v_{n-1} : V_{n-1} \rightarrow V_n, \\ (0, v_{n-1}^{-1}(R_0), \dots, v_{n-1}^{-1}(R_{n-1})) & \text{if } v_{n-1} : V_{n-1} \leftarrow V_n, \end{cases}$$

where $0 = R_0 \subseteq \dots \subseteq R_{n-1} = V_{n-1}$ is the right filtration of

$$V_1 \xrightarrow{v_1} V_2 \xrightarrow{v_2} \dots \xrightarrow{v_{n-2}} V_{n-1} .$$

The k -th element in the right filtration of \mathbb{V} is written $\mathcal{R}_{\mathbb{V}}[k]$. Note that the recursive definition maintains the filtration property, that is,

$$0 = \mathcal{R}_{\mathbb{V}}[0] \subseteq \mathcal{R}_{\mathbb{V}}[1] \subseteq \dots \subseteq \mathcal{R}_{\mathbb{V}}[n-1] \subseteq \mathcal{R}_{\mathbb{V}}[n] = V_n .$$

The transformation $\mathbb{V} \rightarrow \mathcal{R}_{\mathbb{V}}$ defines a functor from $\text{rep}_{\mathbf{k}}(\mathbb{Q})$ to the category of finite-dimensional n -filtered \mathbf{k} -vector spaces. This functor induces a ring homomorphism from the endomorphism ring of \mathbb{V} to the one of $\mathcal{R}_{\mathbb{V}}$. The crux of the matter is that this ring homomorphism turns out to be an isomorphism when \mathbb{V} is streamlined. Hence, in this case, the decomposition structures of \mathbb{V} and $\mathcal{R}_{\mathbb{V}}$ are the same, since, as we saw after Theorem A.5, direct summands correspond to idempotents in the endomorphism ring—see also [164].

Thus, step 2 of the proof reduces to showing that the right filtration $\mathcal{R}_{\mathbb{V}}$ admits a decomposition (a standard result), and so the following correspondence is easily established by Carlsson and de Silva [49]:

$$(A.12) \quad \mathbb{V} \cong \bigoplus_{i=1}^n \mathbb{I}_{\mathbb{Q}}[i, n]^{c_i} \quad \text{where } c_i = \dim \mathcal{R}_{\mathbb{V}}[i] / \mathcal{R}_{\mathbb{V}}[i-1] .$$

In other words, the multiplicities of the interval summands $\mathbb{I}_{\mathbb{Q}}[i, n]$ in the decomposition of \mathbb{V} are given by the dimensions of the quotients of consecutive elements in the right filtration $\mathcal{R}_{\mathbb{V}}$.

This proof is constructive in that it induces the following high-level algorithm for decomposing an arbitrary representation $\mathbb{V} \in \text{rep}_{\mathbf{k}}(\mathbf{Q})$ into interval summands. The algorithm proceeds from left to right along the quiver \mathbf{Q} , considering indices 1 through n sequentially. At every index i , it restricts the focus to the subquiver $\mathbf{Q}[1, i]$ spanned by nodes 1 through i , and to the representation $\mathbb{V}[1, i]$ of $\mathbf{Q}[1, i]$ induced by \mathbb{V} . The right filtration of $\mathbb{V}[1, i]$ is obtained from the one of $\mathbb{V}[1, i-1]$ using (A.11), and the multiplicities of the interval summands $\mathbb{I}_{\mathbf{Q}}[j, i]$ for $j = 1, \dots, i$ are computed from the right filtration using (A.12). These summands are then pruned out of the direct-sum decomposition of \mathbb{V} (this step is important to maintain the invariant that $\mathbb{V}[1, i]$ is streamlined at the time when index i is reached). Upon termination, all the interval summands $\mathbb{I}_{\mathbf{Q}}[b, d]$ of \mathbb{V} have been extracted, in the lexicographical order on the pairs (d, b) .

This procedure is closely related to the one based on reflection functors, presented in Section 4.3. To see this, take the linear quiver L_n as our quiver \mathbf{Q} , and recall Example A.19. Each application of the Coxeter functor $\mathcal{C}^+ = \mathcal{R}_1^+ \mathcal{R}_2^+ \cdots \mathcal{R}_{n-1}^+ \mathcal{R}_n^+$ on \mathbb{V} prunes out the interval representations $\mathbb{I}_{\mathbf{Q}}[i, n]$ ($1 \leq i \leq n$) from the direct-sum decomposition of \mathbb{V} , in the decreasing order of their left endpoints i , while shifting the other interval representations $\mathbb{I}_{\mathbf{Q}}[i, j]$ ($1 \leq i \leq j < n$) upwards to $\mathbb{I}_{\mathbf{Q}}[i+1, j+1]$. After applying the Coxeter functor n times, all the interval summands $\mathbb{I}_{\mathbf{Q}}[b, d]$ of \mathbb{V} have been extracted, in an order that is the inverse of the one from [49].

When \mathbf{Q} is an arbitrary A_n -type quiver, the procedure from Section 4.3 is a combination of the previous one with a novel ingredient from [50]. Recall from Example A.20 that we apply the following sequence of reflections, where $i_1 < \cdots < i_r$ are the heads of the backward arrows in \mathbf{Q} , in order to turn \mathbf{Q} into the linear quiver L_n :

$$(A.13) \quad s_1 \cdots s_{i_r} s_1 \cdots s_{i_{r-1}} \cdots s_1 \cdots s_{i_1}.$$

By doing so, we actually travel down a big pyramid-shaped lattice, shown in Figure A.4 for the case $n = 5$. Every maximal x -monotone path in this pyramid corresponds to a unique A_n -type quiver. Each reflection in the sequence (A.13) makes us travel down one diamond in the pyramid, and the corresponding reflection functor modifies our quiver accordingly. Starting from \mathbf{Q} , there may be different sequences of reflections that lead to the linear quiver L_n at the bottom of the pyramid, but it turns out that the corresponding sequences of reflection functors give isomorphic results [24, lemma 1.2]. Thus, the transformation of our initial representation into a representation of L_n is canonical, even though the sequence of reflections itself is not.

This approach is very similar to the one used by Carlsson, de Silva, and Morozov [50] to relate the so-called *level-sets zigzags* and *extended persistence* (see Section 1.3 in Chapter 2), except here we work in a general setting where we can ‘lose’ some summands in the interval decomposition of our representation along the way down the pyramid. Once we reach the bottom, the Coxeter functor can be applied n times as before to prune out the remaining interval summands.

All in all, the summands $\mathbb{I}_{\mathbf{Q}}[b, i_1], \dots, \mathbb{I}_{\mathbf{Q}}[b, i_r]$ are extracted first, in an order that depends on the order in which the diamonds of the pyramid are visited⁸, and

⁸For instance, with the sequence (A.13), the summands $\mathbb{I}_{\mathbf{Q}}[b, i_1]$ are pruned first, then the summands $\mathbb{I}_{\mathbf{Q}}[b, i_2]$, etc.

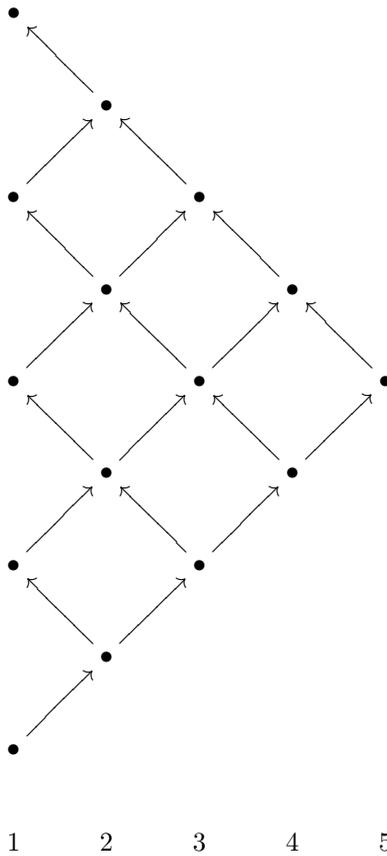


FIGURE A.4. The pyramid travelled down when turning an A_5 -type quiver into L_5 .

then the rest of the summands are extracted in an order that is the inverse of the one from [49].

5. Proof of Gabriel’s theorem: the general case

This section provides a proof of Gabriel’s theorem (Theorems A.6 and A.11) in the general case. The proof is divided into two parts: the “if” part, which shows that for every Dynkin quiver there are finitely many isomorphism classes of indecomposable representations; the “only if” part, which shows that this finiteness property holds only for Dynkin quivers.

5.1. The “if” part. We will present the proof of Bernstein, Gelfand, and Ponomarev [24] based on reflection functors. The outline is the same as in Section 4.3, but the details are more subtle because the indecomposable representations of a general Dynkin quiver are not naturally identified with interval representations.

A short detour. Before proceeding with the proof itself, let us recall the result on the nilpotency of Coxeter functors, which was mentioned in passing between Examples A.19 and A.20 in Section 4.2.

PROPOSITION A.24 ([24]). *Let \mathbf{Q} be a Dynkin quiver, and let \mathbb{V} be an indecomposable representation of \mathbf{Q} . Then, for any total order on Q_0 that is compatible with \mathbf{Q} , there is a finite r such that $\underbrace{\mathcal{C}^+ \cdots \mathcal{C}^+}_{r \text{ times}} \mathbb{V} = 0$.*

The proof of this result considers the subgroup of automorphisms $\mathbb{Z}^n \rightarrow \mathbb{Z}^n$ generated by the maps (A.4) and (A.5) for i ranging over $Q_0 = \{1, 2, \dots, n\}$. This group is called the *Weyl group*, denoted by $W_{\mathbf{Q}}$. It preserves the Tits form $q_{\mathbf{Q}}$, so the set $\Phi_{\mathbf{Q}}$ of roots of $q_{\mathbf{Q}}$ is stable under its action. When \mathbf{Q} is Dynkin, $\Phi_{\mathbf{Q}}$ is finite, so each element of $W_{\mathbf{Q}}$ induces a permutation on the set $\Phi_{\mathbf{Q}}$. Moreover, the simple representations \mathbb{S}_i verify $q_{\mathbf{Q}}(\dim \mathbb{S}_i) = 1$, therefore $\Phi_{\mathbf{Q}}$ contains the entire basis of \mathbb{Z}^n . As a result, $W_{\mathbf{Q}}$ can be embedded as a subgroup of the permutation group of $\Phi_{\mathbf{Q}}$ and is therefore finite.

The following result is an easy consequence of the finiteness of $W_{\mathbf{Q}}$. The proof is based on a fixed-point argument—see [24, lemma 2.3] for the details. In the statement we abuse notations and write \mathcal{C}^+x for the image of a vector $x \in \mathbb{Z}^n$ under the action of the element of $W_{\mathbf{Q}}$ corresponding to the Coxeter functor \mathcal{C}^+ (we will do the same with reflection functors in the following):

LEMMA A.25 ([24]). *Let \mathbf{Q} be a Dynkin quiver, and let x be a positive root of $q_{\mathbf{Q}}$. Then, for any total order on Q_0 that is compatible with \mathbf{Q} , there is a finite r such that the vector $\underbrace{\mathcal{C}^+ \cdots \mathcal{C}^+}_{r \text{ times}} x$ is non-positive.*

Letting $x = \dim \mathbb{V}$ in the lemma implies that $\dim \mathbb{V}$ is eventually sent to some non-positive vector under the repeated action of \mathcal{C}^+ . At that stage, the image of \mathbb{V} itself must be 0, which concludes the proof of Proposition A.24.

REMARK. Bernstein, Gelfand, and Ponomarev [24] also showed that $\mathcal{C}^{\pm} \mathbb{V}$ and $\mathcal{C}^{\pm}x$ are in fact independent of the choice of total order on Q_0 , as long as that order is compatible with \mathbf{Q} . Thus, \mathcal{C}^+ and \mathcal{C}^- are uniquely defined for each quiver \mathbf{Q} , and the clause “for any total order on Q_0 that is compatible with \mathbf{Q} ” can be safely removed from the statements of Proposition A.24 and Lemma A.25. This fact is not used in the proof of Gabriel’s theorem as it is presented here though.

The proof. Given a Dynkin quiver \mathbf{Q} and an indecomposable representation \mathbb{V} , we know from Proposition A.24 that there is a sequence of indices i_1, \dots, i_s , possibly with repetitions, such that $\mathcal{R}_{i_s}^+ \cdots \mathcal{R}_{i_1}^+ \mathbb{V} = 0$. Assuming without loss of generality that this is a minimal such sequence, we deduce as in Section 4.3 that $\mathcal{R}_{i_{s-1}}^+ \cdots \mathcal{R}_{i_1}^+ \mathbb{V}$ is the simple representation \mathbb{S}_{i_s} , from which, using Corollary A.17, we conclude that the dimension vector of \mathbb{V} is a positive root of $q_{\mathbf{Q}}$. Thus, $\mathbb{V} \mapsto \dim \mathbb{V}$ maps indecomposable representations of \mathbf{Q} to positive roots of $q_{\mathbf{Q}}$.

To prove the “if” part of Theorem A.6, we need to show that the map $\mathbb{V} \mapsto \dim \mathbb{V}$ is an injection from the isomorphism classes of indecomposable representations of \mathbf{Q} to the set of positive roots of $q_{\mathbf{Q}}$. Suppose \mathbb{V} and \mathbb{W} are nonisomorphic indecomposable representations sharing the same dimension vector. Then, the sequence of reflections $\mathcal{R}_{i_{s-1}}^+ \cdots \mathcal{R}_{i_1}^+$ that sends \mathbb{V} to \mathbb{S}_{i_s} also sends $\dim \mathbb{V}$ to $\dim \mathbb{S}_{i_s}$, and therefore it sends \mathbb{W} to an indecomposable representation with the same dimension vector as \mathbb{S}_{i_s} . But since \mathbb{S}_{i_s} is obviously the only representation with that dimension vector (up to isomorphism), we have $\mathcal{R}_{i_{s-1}}^+ \cdots \mathcal{R}_{i_1}^+ \mathbb{W} \cong \mathbb{S}_{i_s}$. Applying

now the mirrored sequence of reflections, we obtain:

$$\begin{aligned} \mathbb{W} &\cong \mathcal{R}_{i_1}^- \cdots \mathcal{R}_{i_{s-1}}^- \mathcal{R}_{i_{s-1}}^+ \cdots \mathcal{R}_{i_1}^+ \mathbb{W} \cong \mathcal{R}_{i_1}^- \cdots \mathcal{R}_{i_{s-1}}^- \mathbb{S}_{i_s} \\ &\cong \mathcal{R}_{i_1}^- \cdots \mathcal{R}_{i_{s-1}}^- \mathcal{R}_{i_{s-1}}^+ \cdots \mathcal{R}_{i_1}^+ \mathbb{V} \cong \mathbb{V}. \end{aligned}$$

Thus, the map $\mathbb{V} \mapsto \underline{\dim} \mathbb{V}$ is injective. Combined with Proposition A.10, this fact implies that there are only finitely many isomorphism classes of indecomposable representations of \mathbb{Q} , hereby proving the “if” part of Theorem A.6.

To prove Theorem A.11, we also need to show that the map $\mathbb{V} \mapsto \underline{\dim} \mathbb{V}$ is surjective. Let $x \in \mathbb{N}^n$ be a root of $q_{\mathbb{Q}}$. According to Lemma A.25, there is a sequence of indices i_1, \dots, i_s , possibly with repetitions, such that the sequence of reflections $\mathcal{R}_{i_s}^+ \cdots \mathcal{R}_{i_1}^+$ sends x to a non-positive root of $q_{\mathbb{Q}}$. Assuming without loss of generality that this is a minimal such sequence, we deduce that $\mathcal{R}_{i_{s-1}}^+ \cdots \mathcal{R}_{i_1}^+$ sends x to the vector $(0, \dots, 0, 1, 0, \dots, 0)^\top$ with a single 1 at position i_s . This is the dimension vector of the simple representation \mathbb{S}_{i_s} , which we know is indecomposable. Therefore, $\mathcal{R}_{i_1}^- \cdots \mathcal{R}_{i_{s-1}}^- \mathbb{S}_{i_s}$ is also indecomposable (because nonzero) and has x as dimension vector. This proves that the map $\mathbb{V} \mapsto \underline{\dim} \mathbb{V}$ is surjective from the isomorphism classes of indecomposable representations of \mathbb{Q} to the set of positive roots of $q_{\mathbb{Q}}$.

5.2. The “only if” part. We will give two different proofs of the “only if” part of Theorem A.6. The first one is both direct and self-contained, requiring no extra mathematical background. However, it is based on an exhaustive enumeration, so it is quite lengthy and tedious, and not really enlightening. This is why we also give a second proof, which on the contrary is both compact and elegant, with a more geometric flavor, but whose details require some notions in algebraic geometry. Gabriel himself adopted the first proof and mentioned the second one briefly in his original paper [133].

5.2.1. *Direct inspection.* The proof goes by contraposition. Let \mathbb{Q} be a finite connected quiver, and assume \mathbb{Q} is not Dynkin. Then, as we observed already in the proof of Theorem A.9, $\bar{\mathbb{Q}}$ must contain one of the diagrams of Figure A.3 as a subgraph. Denoting $\tilde{\mathbb{Q}}$ the restriction of \mathbb{Q} to that subgraph, we have that any representation of $\tilde{\mathbb{Q}}$ can be extended to a representation of \mathbb{Q} by assigning trivial spaces to the remaining vertices and trivial maps to the remaining arrows. This process turns the indecomposable representations of $\tilde{\mathbb{Q}}$ into indecomposable representations of \mathbb{Q} in a way that preserves the isomorphism classes. Hence, if we can exhibit an infinite family of pairwise nonisomorphic indecomposable representations for $\tilde{\mathbb{Q}}$, then we can conclude that \mathbb{Q} itself also has a similar family.

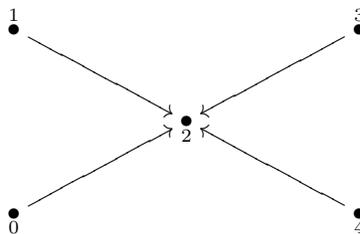
Thus, the proof of the “only if” part of Theorem A.6 reduces to enumerating all the quivers whose underlying graphs are Euclidean, and for each one of them, exhibiting an infinite family of pairwise nonisomorphic indecomposable representations. This is what Gabriel did in his original proof, and we refer the interested reader to [133] for the complete enumeration—see also [104, theorem 1.23] for a more recent treatment. Here we restrict ourselves to a few noteworthy excerpts.

EXAMPLE A.26 (\tilde{A}_0). Let \mathbb{Q} be the quiver with one vertex and one loop. Fix the space V_0 to be \mathbf{k}^n , so a representation is given by a linear map $\mathbf{k}^n \rightarrow \mathbf{k}^n$ or, equivalently, by an $n \times n$ matrix M . Two representations with matrices M, N are isomorphic if and only if there is an invertible matrix B such that $NB = BM$,

meaning that M, N lie in the same conjugacy class. Therefore, the Jordan block matrices $J_{\lambda, n}$ for $\lambda \in \mathbf{k}$ give pairwise nonisomorphic indecomposable representations. If \mathbf{k} is an infinite field, then we get an infinite number of isomorphism classes of indecomposable representations. Otherwise, it suffices to let n range over $\mathbb{N} \setminus \{0\}$ to get an infinite number. Note that when \mathbf{k} is algebraically closed, the isomorphism classes of indecomposable representations correspond exactly to the conjugacy classes of the Jordan block matrices.

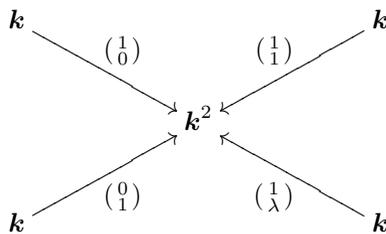
EXAMPLE A.27 ($\tilde{A}_n, n \geq 1$). Let \mathbf{Q} be a quiver of type $\tilde{A}_n, n \geq 1$. We put space \mathbf{k} at every vertex of \mathbf{Q} , and identity maps at all the arrows except one where we put λ times the identity for an arbitrary scalar λ . Then, the various choices of $\lambda \in \mathbf{k}$ give pairwise nonisomorphic indecomposable representations. More generally, given $r \in \mathbb{N} \setminus \{0\}$, we put \mathbf{k}^r at every vertex, and identity maps at all the arrows except one where we put the Jordan block matrix $J_{\lambda, r}$ for an arbitrary scalar λ . The various choices of $\lambda \in \mathbf{k}$ once again give pairwise nonisomorphic indecomposable representations.

EXAMPLE A.28 (\tilde{D}_n). Consider the following quiver \mathbf{Q} of type \tilde{D}_4 :

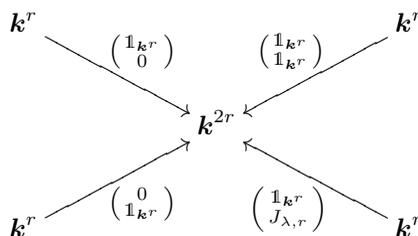


This quiver has a historical interest, since it is for its study that Gelfand and Ponomarev [139] introduced the reflection functors.

Consider the following representation of \mathbf{Q} , where λ is an arbitrary scalar:



The various choices of $\lambda \in \mathbf{k}$ give pairwise nonisomorphic indecomposable representations of \mathbf{Q} . More generally, for any $r \in \mathbb{N} \setminus \{0\}$ and $\lambda \in \mathbf{k}$, the following representations are indecomposable and pairwise nonisomorphic:



Indecomposable representations for \tilde{D}_n can be obtained from this one by inserting a string of spaces \mathbf{k}^{2r} linked by identity maps in place of vertex 2.

REMARK. This is but one possible choice of arrow orientations in the quiver \tilde{D}_n . All the other choices have to be considered as well, thus making the enumeration tedious. In fact, reflection functors give a way of switching between arrow orientations, with a predictable effect on the representations of the corresponding quivers. This greatly helps reduce the length of the enumeration, since basically a single choice of arrow orientations needs to be considered for every diagram⁹ of Figure A.3. Let us emphasize however that the enumeration proof does not rely on the reflection functors in any essential way.

5.2.2. *Tits' geometric argument.* Tits proposed a simple and elegant proof with a geometric flavor for the “only if” part of Theorem A.6. This proof stresses the role played by the Tits form in Gabriel’s theorem. Although the details use certain facts from algebraic geometry, the outline does not require to introduce any additional concepts. We refer the reader to [36] for a formal treatment.

Assume as before that \mathbf{Q} is a finite connected quiver with vertex set $Q_0 = \{1, \dots, n\}$. Let us fix a nonzero vector $x = (x_1, \dots, x_n)^\top \in \mathbb{N}^n$ and restrict the focus to those representations of \mathbf{Q} that have x as dimension vector. For any such representation \mathbb{V} , we choose arbitrary bases so that each space V_i is identified with \mathbf{k}^{x_i} and each linear map v_a is identified with an $x_{h_a} \times x_{t_a}$ matrix. Thus, \mathbb{V} is viewed as an element of the following space (an algebraic variety over the field \mathbf{k}):

$$\text{rep}_{\mathbf{k}}(\mathbf{Q}, x) = \bigoplus_{a \in Q_1} \text{Mat}_{x_{h_a} \times x_{t_a}}(\mathbf{k}).$$

Given any isomorphism $\phi : \mathbb{V} \rightarrow \mathbb{W}$, we have $\phi_i \in \text{GL}_{x_i}(\mathbf{k})$ for all $i \in Q_0$, and by Definition A.4 we have $w_a \circ \phi_{t_a} = \phi_{h_a} \circ v_a$, or rather $w_a = \phi_{h_a} \circ v_a \circ \phi_{t_a}^{-1}$, for all $a \in Q_1$. Thus, ϕ can be viewed as an element of the following (algebraic) group:

$$\text{GL}_x(\mathbf{k}) = \prod_{i=1}^n \text{GL}_{x_i}(\mathbf{k})$$

which acts (algebraically) on $\mathbb{V} \in \text{rep}_{\mathbf{k}}(\mathbf{Q}, x)$ by conjugation, that is:

$$(A.14) \quad \forall a \in Q_1, (\phi \cdot \mathbb{V})_a = \phi_{h_a} \circ v_a \circ \phi_{t_a}^{-1}.$$

This group action is key to understanding the structure of the space of representations of \mathbf{Q} , as the isomorphism classes in $\text{rep}_{\mathbf{k}}(\mathbf{Q}, x)$ coincide with the orbits of the action:

$$\forall \mathbb{V}, \mathbb{W} \in \text{rep}_{\mathbf{k}}(\mathbf{Q}, x), \mathbb{V} \cong \mathbb{W} \iff \mathbb{W} \in \text{GL}_x(\mathbf{k}) \cdot \mathbb{V}.$$

These orbits satisfy the following orbit-stabilizer type theorem:

$\forall \mathbb{V} \in \text{rep}_{\mathbf{k}}(\mathbf{Q}, x), \dim \text{GL}_x(\mathbf{k}) - \dim \text{GL}_x(\mathbf{k}) \cdot \mathbb{V} = \dim \{\phi \in \text{GL}_x(\mathbf{k}) \mid \phi \cdot \mathbb{V} = \mathbb{V}\}$, where the stabilizer $\{\phi \in \text{GL}_x(\mathbf{k}) \mid \phi \cdot \mathbb{V} = \mathbb{V}\}$ has dimension at least 1 in our context since it contains the set $\{\lambda \mathbf{1}_{\text{GL}_x(\mathbf{k})} \mid \lambda \in \mathbf{k}\}$, whose members are acting trivially on \mathbb{V} according to (A.14). Thus,

$$\forall \mathbb{V} \in \text{rep}_{\mathbf{k}}(\mathbf{Q}, x), \dim \text{GL}_x(\mathbf{k}) - \dim \text{GL}_x(\mathbf{k}) \cdot \mathbb{V} \geq 1.$$

⁹Except \tilde{A}_n ($n \geq 0$), which has undirected loops and is therefore treated differently, as described in Examples A.26 and A.27.

Assume now that the quiver \mathbf{Q} is of finite type. This implies by Theorem A.5 that there can only be finitely many isomorphism classes in $\text{rep}_{\mathbf{k}}(\mathbf{Q}, x)$. Equivalently, $\text{rep}_{\mathbf{k}}(\mathbf{Q}, x)$ decomposes into finitely many orbits of $\text{GL}_x(\mathbf{k})$. It follows¹⁰ that at least one of these orbits must have full dimension. Thus,

$$\dim \text{GL}_x(\mathbf{k}) - \dim \text{rep}_{\mathbf{k}}(\mathbf{Q}, x) \geq 1,$$

which can be rewritten as

$$q_{\mathbf{Q}}(x) \geq 1$$

since the dimensions of $\text{GL}_x(\mathbf{k})$ and $\text{rep}_{\mathbf{k}}(\mathbf{Q}, x)$ are $\sum_{i \in Q_0} x_i^2$ and $\sum_{a \in Q_1} x_{t_a} x_{h_a}$ respectively. This proves that, when \mathbf{Q} is of finite type, $q_{\mathbf{Q}}(x)$ is positive for any dimension vector $x \in \mathbb{N}^n \setminus \{0\}$. We extend this property to any vector $x \in \mathbb{Z}^n \setminus \{0\}$ by observing that $q_{\mathbf{Q}}(x) \geq q_{\mathbf{Q}}(|x_1|, \dots, |x_n|^\top)$. It follows then from Theorem A.9 that \mathbf{Q} is Dynkin.

6. Beyond Gabriel's theorem

Four decades of active research in quiver theory have resulted in the development of a vast and rich literature on the subject, with many connections to other areas of mathematics. Here we restrict the exposition to few selected topics that are relevant to persistence theory. For more comprehensive overviews, see for instance [94, 104].

6.1. Tame and wild quivers. The correspondence between the isomorphism classes of indecomposable representations of a quiver and the positive roots of its Tits form given in Theorem A.11 was generalized by Kac [168] to arbitrary quivers, provided the base field \mathbf{k} is algebraically closed.

THEOREM A.29 (Kac I). *Assuming the base field is algebraically closed, the set of dimension vectors of indecomposable representations of an arbitrary finite connected quiver \mathbf{Q} is precisely the set of positive roots of its Tits form. In particular, this set is independent of the arrow orientations in \mathbf{Q} .*

The difference with the Dynkin case is that the correspondence is no longer one-to-one. To be more precise, Kac's result distinguishes between two types of positive roots x of $q_{\mathbf{Q}}$:

- those that satisfy $q_{\mathbf{Q}}(x) = 1$ are called *real roots*,
- those that satisfy $q_{\mathbf{Q}}(x) \leq 0$ are called *imaginary roots*.

In the Dynkin case, all positive roots are real, and Gabriel's theorem asserts that the correspondence between these and the dimension vectors of indecomposable representations is one-to-one. In the general case, Kac's theorem asserts the following.

THEOREM A.30 (Kac II). *Assume the base field is algebraically closed. Given an arbitrary finite connected quiver \mathbf{Q} , to each real root of $q_{\mathbf{Q}}$ corresponds exactly one isomorphism class of indecomposable representations of \mathbf{Q} , and to each imaginary root x of $q_{\mathbf{Q}}$ corresponds an $r(x)$ -dimensional variety of isomorphism classes of indecomposable representations of \mathbf{Q} , where $r(x) = 1 - q_{\mathbf{Q}}(x)$.*

¹⁰This implication is not trivial. Nevertheless, it follows the intuition from affine geometry that a finite union of subspaces cannot cover the ambient space unless at least one element in the union has full dimension.

In light of Theorem A.9 and its following remark, this result classifies the finite connected quivers into 3 categories, thus refining Gabriel's dichotomy into a trichotomy:

- The Dynkin quivers, whose underlying graphs are the diagrams of Figure A.2 and whose Tits forms are positive definite. Those have finitely many isomorphism classes of indecomposable representations, one per positive real root. This coincides with Gabriel's theorem.
- The tame quivers, whose underlying graphs are the diagrams of Figure A.3 and whose Tits forms are positive semidefinite. Those have 1-dimensional varieties of isomorphism classes of indecomposable representations, one per positive imaginary root. Moreover, it can be shown¹¹ that their Tits forms have a 1-dimensional radical and therefore only countably many imaginary roots. As a consequence, such quivers have countably many 1-dimensional varieties of isomorphism classes of indecomposable representations. For instance, the quiver of type \tilde{A}_0 has a uniformly zero Tits form and \mathbb{Z} as radical. As depicted in Example A.26, each dimension vector $(n) \in \mathbb{Z}$ gives rise to a family of isomorphism classes of indecomposable representations parametrized by $\lambda \in \mathbf{k}$, one isomorphism class per Jordan block matrix $J_{\lambda, n}$.
- The wild quivers, whose Tits forms are indefinite. Those admit higher-dimensional varieties of isomorphism classes of indecomposable representations. For instance, the quiver with 1 node and 2 loops admits an $(n^2 + 1)$ -dimensional family of isomorphism classes of indecomposable representations per vector space dimension n . The fact that the number of parameters needed to describe the family can be arbitrarily large means that describing explicitly the set of indecomposable representations is essentially an impossible task. This is in fact true for any wild quiver (hence the name), as it turns out that the problem will be at least as hard as it is for this particular quiver.

REMARK. Tame quivers have also been studied in their own right, and a complete classification of their indecomposable representations was provided independently by Donovan and Freislich [110] and Nazarova [204]. This happened around the same time as Gabriel's theorem, and almost a decade before Kac's theorem.

6.2. Quivers as algebras. An important aspect of quiver theory is its connection to the representation theory of associative algebras. In some sense, a quiver can be viewed as an algebra, and its representations can be viewed as modules over that algebra. Let us describe the connection briefly.

Paths and cycles in a quiver \mathbf{Q} are understood in the same way as in (directed) graph theory. A *path* in \mathbf{Q} is a finite sequence of arrows $a_r \cdots a_1$ such that $h_{a_i} = t_{a_{i+1}}$ for $1 \leq i < r$. If $h_{a_r} = t_{a_1}$, then the path is called a *cycle*. Here, r denotes the *length* of the path. To each vertex $i \in Q_0$ we associate a *trivial path* e_i of length zero.

DEFINITION A.31. The *path algebra* of a quiver \mathbf{Q} , noted \mathbf{kQ} , is the \mathbf{k} -algebra having as basis the set of all paths in \mathbf{Q} . The product in the algebra is defined by linearity and by the following product rule for two paths $a_r \cdots a_1$ and $b_s \cdots b_1$:

$$(A.15) \quad (a_r \cdots a_1) \cdot (b_s \cdots b_1) = \begin{cases} a_r \cdots a_1 b_s \cdots b_1 & \text{if } t_{a_1} = h_{b_s}, \\ 0 & \text{otherwise.} \end{cases}$$

¹¹See [172, lemma 4.1.3] for a proof.

The following properties are straightforward. First, $\mathbf{k}\mathbf{Q}$ is an associative algebra. Second, when Q_0 is finite, $\mathbf{k}\mathbf{Q}$ has a unit element, namely $1 = \sum_{i \in Q_0} e_i$. Third, $\mathbf{k}\mathbf{Q}$ is finite-dimensional (as a \mathbf{k} -vector space) if and only if there are only finitely many different paths in \mathbf{Q} , i.e. if and only if \mathbf{Q} is finite and acyclic.

Representations of a finite quiver \mathbf{Q} can be viewed as left modules over its path algebra. Indeed, given $\mathbb{V} = (V_i, v_a) \in \text{Rep}_{\mathbf{k}}(\mathbf{Q})$, equip $X = \bigoplus_{i \in Q_0} V_i$ with the scalar multiplication induced by \mathbf{k} -linearity from the rule:

$$(a_r \cdots a_1) \cdot x = v_{a_r} \circ \cdots \circ v_{a_1}(x_{t_{a_1}}) \in V_{h_{a_r}} \subseteq X$$

for any path $a_r \cdots a_1$ in \mathbf{Q} and any vector $x = \sum_{i \in Q_0} x_i \in X$. This turns \mathbb{V} into a $\mathbf{k}\mathbf{Q}$ -module. Conversely, any left $\mathbf{k}\mathbf{Q}$ -module M can be turned into a representation by letting $V_i = e_i M$ for each $i \in Q_0$, where e_i is the trivial path at vertex i , and by letting each \mathbf{k} -linear map $v_a : V_{t_a} \rightarrow V_{h_a}$ be induced by the $\mathbf{k}\mathbf{Q}$ -module structure on M . It can be shown¹² that these conversions define an equivalence of categories between $\text{Rep}_{\mathbf{k}}(\mathbf{Q})$ and the category of $\mathbf{k}\mathbf{Q}$ -modules.

PROPOSITION A.32. *For a finite quiver \mathbf{Q} , $\text{Rep}_{\mathbf{k}}(\mathbf{Q})$ is equivalent to the category of left $\mathbf{k}\mathbf{Q}$ -modules.*

Thus, classifying the representations of a finite quiver \mathbf{Q} is the same as classifying the left modules over its path algebra $\mathbf{k}\mathbf{Q}$. As we saw earlier, when \mathbf{Q} is acyclic, $\mathbf{k}\mathbf{Q}$ is a finite-dimensional algebra over a field, therefore it is an Artin algebra. The representation theory of Artin algebras [13] can then be used to classify its left modules, including the infinite-dimensional ones. Of particular interest to us is the following result, proven independently by Auslander [12] and Ringel and Tachikawa [217]:

THEOREM A.33. *Let A be an Artin algebra. If A has only a finite number of isomorphism classes of finitely generated indecomposable left modules, then every indecomposable left A -module is finitely generated, and every left A -module is a direct sum of indecomposable modules.*

Gabriel's theorem ensures that the hypothesis of Theorem A.33 is satisfied by $\mathbf{k}\mathbf{Q}$ when \mathbf{Q} is Dynkin. In this case, both theorems combined together give the following:

COROLLARY A.34. *For a Dynkin quiver \mathbf{Q} , every indecomposable representation in $\text{Rep}_{\mathbf{k}}(\mathbf{Q})$ has finite dimension, and every representation in $\text{Rep}_{\mathbf{k}}(\mathbf{Q})$ is a direct sum of indecomposable representations. In particular, \mathbf{Q} has finitely many isomorphism classes of indecomposable representations, and all of them are finite-dimensional.*

Moreover, as we saw in Section 3, the finite-dimensional indecomposable representations of \mathbf{Q} have local endomorphism rings, so Azumaya's theorem applies and the decomposition of any representation in $\text{Rep}_{\mathbf{k}}(\mathbf{Q})$ is unique up to isomorphism and permutation of the terms.

REMARK. These properties hold for Dynkin quivers, but not for finite (even acyclic) quivers in general. There are indeed examples of infinite-dimensional representations of tame (acyclic) quivers that are indecomposable, and whose endomorphism rings are not local. See e.g. [172, §2] for such an example.

¹²See [13, theorem III.1.5] for a proof.

In the special case where \mathbf{Q} is of type A_n , we know from Example A.12 that the indecomposable finite-dimensional representations of \mathbf{Q} are the interval representations $\mathbb{I}_0[b, d]$, so Corollary A.34 and its following paragraph assert that every representation of \mathbf{Q} , whether of finite or infinite dimension, decomposes essentially uniquely as a direct sum of interval representations.

6.3. Quivers as categories. In categorical terms, quivers are defined as functors from the Kronecker category $\bullet \rightrightarrows \bullet$ having 2 objects and 2 morphisms (plus identities) to the category of sets. This is really the categorical version of Definition A.1.

A quiver \mathbf{Q} itself can be turned into a category \mathbf{Q} , where the objects are the vertices of \mathbf{Q} and where there is one morphism per (possibly trivial) path in \mathbf{Q} , the composition rule for morphisms being induced by the product rule (A.15) for paths¹³. Given two vertices $i, j \in Q_0$, the number of morphisms in $\text{Hom}(i, j)$ is equal to the number of different paths from i to j in \mathbf{Q} . It is finite for all $i, j \in Q_0$ if and only if \mathbf{Q} is acyclic. A representation of \mathbf{Q} over some field \mathbf{k} is then a functor from \mathbf{Q} to the category of \mathbf{k} -vector spaces, and a morphism between representations is just a natural transformation between functors.

Among the (small) categories that can be constructed this way, the category \mathbf{Z} corresponding to the poset (partially ordered set) (\mathbb{Z}, \leq) is of particular interest to persistence. It has one object per integer and one morphism per couple (i, j) such that $i \leq j$. It is derived from the (infinite) quiver \mathbf{Z} having \mathbb{Z} as vertex set and an edge $i \rightarrow i + 1$ for every integer i , using the previous path construction. Webb [238] has extended Gabriel's theorem to this quiver. The uniqueness of the decomposition follows once again from Azumaya's theorem, since the endomorphism ring of any interval representation is isomorphic to the field \mathbf{k} and therefore local.

THEOREM A.35. *Any representation \mathbb{V} of the quiver \mathbf{Z} such that the spaces $(V_i)_{i \in \mathbb{Z}}$ are finite-dimensional is a direct sum of interval representations.*

REMARK. Webb's proof does not work with representations of the quiver \mathbf{Z} directly, but rather with \mathbb{Z} -graded modules over the graded ring of polynomials $\mathbf{k}[t]$. Indeed, any representation $(V_i, v_i)_{i \in \mathbb{Z}}$ of \mathbf{Z} can be viewed as a \mathbb{Z} -graded module $X = \bigoplus_{i \in \mathbb{Z}} V_i$ over $\mathbf{k}[t]$, the grading being induced by the action of t as follows:

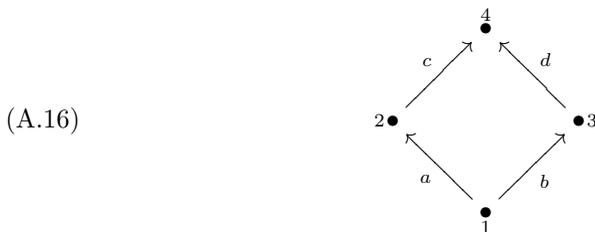
$$\forall i \in \mathbb{Z}, \forall x_i \in V_i, t \cdot x_i = v_i(x_i) \in V_{i+1}.$$

Conversely, any module $X = \bigoplus_{i \in \mathbb{Z}} V_i$ over $\mathbf{k}[t]$ that is equipped with a grading $(v_i : V_i \rightarrow V_{i+1})_{i \in \mathbb{Z}}$ can be viewed as a representation of the quiver \mathbf{Z} . This correspondence is known to induce an equivalence of categories between $\text{Rep}_{\mathbf{k}}(\mathbf{Z})$ and the category of \mathbb{Z} -graded modules over $\mathbf{k}[t]$. In light of this equivalence, Theorem 1.4 is a generalization of the classical structure theorem for finitely generated (graded) modules over a (graded) principal ideal domain [163].

6.4. Quivers as posets. Every poset (P, \preceq) gives rise to a category \mathbf{P} , in which the objects are the elements of P and there is a unique morphism per couple $x \preceq y$. As it turns out, not all poset categories can be derived from a quiver using the construction of Section 6.3. For instance, the poset (P, \preceq) whose Hasse diagram

¹³Notice how this construction is similar in spirit to that of the path algebra $\mathbf{k}\mathbf{Q}$. It is a common practice to identify a quiver with its path algebra and with its corresponding category.

is the following gives a category \mathbf{P} in which the two morphisms $c \circ a$ and $d \circ b$ are identified.



By contrast, the construction of Section 6.3 on the quiver (A.16) gives a category with the same objects as \mathbf{P} but with two morphisms $1 \rightarrow 4$, the paths ca and db being considered different. The concept that unifies these two constructions is the one of a *quiver with relations*, also called a *bound quiver*. In short, this is a quiver \mathbf{Q} in which some of the paths with common sources and targets are considered equal, or more generally, some linear combinations of these paths in the path algebra \mathbf{kQ} are zeroed out.

DEFINITION A.36. A *relation* on a quiver \mathbf{Q} is a \mathbf{k} -linear combination of paths sharing the same source and target vertices. A *quiver with relations* is a pair (\mathbf{Q}, I) where \mathbf{Q} is a quiver and I is an ideal of \mathbf{kQ} spanned by relations. The quotient algebra \mathbf{kQ}/I is called the *path algebra* of (\mathbf{Q}, I) .

Every poset (P, \preceq) is equivalent (as a category) to some quiver with relations. To see this, take the quiver having one vertex per element in P and one directed edge $i \rightarrow j$ per couple $i \prec j$ (do not connect i to itself), and equip that quiver with the relations induced by the transitivity of \preceq . These relations identify all the paths sharing the same source and target vertices, so the resulting quiver with relations is equivalent (as a category) to (P, \preceq) .

A representation \mathbb{V} of a quiver with relations (\mathbf{Q}, I) is defined formally as a left module over the path algebra \mathbf{kQ}/I . Intuitively, the linear relations between paths encoded in the ideal I induce linear relations between compositions of linear maps in \mathbb{V} . When (\mathbf{Q}, I) is equivalent (as a category) to a poset (P, \preceq) , \mathbb{V} is also called a *representation* of (P, \preceq) as it defines a functor from the corresponding category \mathbf{P} to the category of \mathbf{k} -vector spaces.

The posets that are most relevant to us are the sets $T \subseteq \mathbb{R}$ equipped with the natural order \leq on real numbers. As a functor from the poset category \mathbf{T} to the category of \mathbf{k} -vector spaces, a representation of (T, \leq) defines \mathbf{k} -vector spaces $(V_i)_{i \in T}$ and \mathbf{k} -linear maps $(v_i^j : V_i \rightarrow V_j)_{i \leq j \in T}$ satisfying the following constraints:

$$(A.17) \quad \begin{aligned} v_i^i &= \mathbf{1}_{V_i} && \text{for every } i \in T, \text{ and} \\ v_i^k &= v_j^k \circ v_i^j && \text{for every } i \leq j \leq k \in T. \end{aligned}$$

Crawley-Boevey [93] has extended Theorem A.35 to this setting. His proof technique corresponds to a specialized version of the *functorial filtration* method developed in quiver theory [216]. The uniqueness of the decomposition once again follows from Azumaya's theorem.

THEOREM A.37. *Given $T \subseteq \mathbb{R}$, any representation \mathbb{V} of (T, \leq) such that the spaces $(V_i)_{i \in T}$ are finite-dimensional is a direct sum of interval representations.*

REMARK. Definition A.36 follows Gabriel [132]. More recent references such as [13] add the extra condition that the relations in the quiver should only involve paths of length at least 2. This condition is justified in the context of the representation theory for associative algebras. There is indeed a deep connection between finite-dimensional algebras over algebraically closed fields and quivers with relations, the bottomline being that every such algebra A is *equivalent*¹⁴ to the path algebra of an essentially unique quiver with relations (\mathbf{Q}, I) . Uniqueness relies in a critical way on the assumption that the relations in I involve no path of length less than 2, as otherwise extra arrows can be added to \mathbf{Q} arbitrarily and zeroed out in the relations.

¹⁴In the sense that the left modules over A and over $k\mathbf{Q}/I$ form equivalent categories. This is known as *Morita equivalence* in the literature.

Bibliography

- [1] A. Adcock, D. Rubin, and G. Carlsson. “Classification of hepatic lesions using the matching metric”. In: *Computer Vision and Image Understanding* 121 (2014), pp. 36–42 (cit. on p. 8).
- [2] P. K. Agarwal, S. Har-Peled, and K. R. Varadarajan. “Geometric approximation via coresets”. In: *Combinatorial and computational geometry* 52 (2005), pp. 1–30 (cit. on p. 87).
- [3] N. Amenta and M. Bern. “Surface Reconstruction by Voronoi Filtering”. In: *Discrete Comput. Geom.* 22.4 (1999), pp. 481–504 (cit. on p. 72).
- [4] D. Attali and J.-D. Boissonnat. “A Linear Bound on the Complexity of the Delaunay Triangulation of Points on Polyhedral Surfaces”. In: *Discrete and Comp. Geometry* 31 (2004), pp. 369–384 (cit. on p. 83).
- [5] D. Attali, J.-D. Boissonnat, and A. Lieutier. “Complexity of the Delaunay Triangulation of Points on Surfaces: the Smooth Case”. In: *Proc. 19th Annu. ACM Sympos. Comput. Geom.* 2003, pp. 201–210 (cit. on p. 83).
- [6] D. Attali, A. Lieutier, and D. Salinas. “Efficient data structure for representing and simplifying simplicial complexes in high dimensions”. In: *Proc. Symposium on Computational Geometry*. 2011, pp. 501–509 (cit. on p. 157).
- [7] D. Attali, H. Edelsbrunner, and Y. Mileyko. “Weak Witnesses for Delaunay Triangulations of Submanifolds”. In: *Proceedings of the 2007 ACM Symposium on Solid and Physical Modeling*. SPM ’07. Beijing, China: ACM, 2007, pp. 143–150. DOI: 10.1145/1236246.1236267 (cit. on p. 91).
- [8] D. Attali, A. Lieutier, and D. Salinas. “Efficient data structure for representing and simplifying simplicial complexes in high dimensions”. In: *International Journal of Computational Geometry & Applications* 22.04 (2012), pp. 279–303 (cit. on p. 48).
- [9] D. Attali, A. Lieutier, and D. Salinas. “Vietoris-Rips complexes also provide topologically correct reconstructions of sampled shapes”. In: *Computational Geometry: Theory and Applications* 46.4 (2012), pp. 448–465 (cit. on p. 136).
- [10] D. Attali, U. Bauer, O. Devillers, M. Glisse, and A. Lieutier. “Homological reconstruction and simplification in \mathbb{R}^3 ”. In: *Computational Geometry: Theory and Applications* (2014). To appear (cit. on p. 166).
- [11] D. Attali, M. Glisse, S. Hornus, F. Lazarus, and D. Morozov. “Persistence-sensitive simplification of functions on surfaces in linear time”. In: *Presented at TOPOINVIS* 9 (2009), pp. 23–24 (cit. on p. 166).
- [12] M. Auslander. “Representation theory of Artin algebras II”. In: *Communications in Algebra* 1 (1974), pp. 269–310 (cit. on pp. 17, 27, 193).
- [13] M. Auslander, I. Reiten, and S. O. Smalø. *Representation theory of Artin algebras*. Vol. 36. Cambridge Studies in Advanced Mathematics. Cambridge University Press, 1995 (cit. on pp. 163, 193, 196).

- [14] G. Azumaya. “Corrections and supplementaries to my paper concerning Krull-Remak-Schmidt’s theorem”. In: *Nagoya Mathematical Journal* 1 (1950), pp. 117–124 (cit. on pp. 16, 171).
- [15] S. Balakrishnan, B. Fasy, F. Lecci, A. Rinaldo, A. Singh, and L. Wasserman. *Statistical inference for persistent homology*. Tech. rep. arXiv:1303.7117. 2013 (cit. on p. 159).
- [16] U. Bauer and H. Edelsbrunner. “The Morse theory of Čech and Delaunay filtrations”. In: *Proceedings of the Annual Symposium on Computational Geometry*. To appear. 2014 (cit. on p. 83).
- [17] U. Bauer, M. Kerber, and J. Reininghaus. *Clear and Compress: Computing Persistent Homology in Chunks*. Research Report arXiv:1303.0477 [math.AT]. 2013 (cit. on pp. 45, 156).
- [18] U. Bauer, M. Kerber, and J. Reininghaus. “Clear and Compress: Computing Persistent Homology in Chunks”. In: *Proc. TopoInVis*. 2013 (cit. on p. 157).
- [19] U. Bauer, C. Lange, and M. Wardetzky. “Optimal topological simplification of discrete functions on surfaces”. In: *Discrete & Computational Geometry* 47.2 (2012), pp. 347–377 (cit. on p. 166).
- [20] U. Bauer and M. Lesnick. “Induced Matchings of Barcodes and the Algebraic Stability of Persistence”. In: *Proc. Annual Symposium on Computational Geometry*. 2014 (cit. on pp. 28, 49, 61, 62).
- [21] S. Belongie, J. Malik, and J. Puzicha. “Shape Matching and Object Recognition Using Shape Contexts”. In: *IEEE Trans. Pattern Anal. Mach. Intell.* 24.4 (2002), pp. 509–522. DOI: <http://dx.doi.org/10.1109/34.993558> (cit. on p. 133).
- [22] P. Bendich, D. Cohen-Steiner, H. Edelsbrunner, J. Harer, and D. Morozov. “Inferring local homology from sampled stratified spaces”. In: *Foundations of Computer Science, 2007. FOCS’07. 48th Annual IEEE Symposium on*. IEEE. 2007, pp. 536–546 (cit. on pp. 36, 39, 82).
- [23] P. Bendich, S. Chin, J. Clarke, J. deSena, J. Harer, E. Munch, A. Newman, D. Porter, D. Rouse, N. Strawn, et al. “Topological and Statistical Behavior Classifiers for Tracking Applications”. In: *arXiv preprint arXiv:1406.0214* (2014) (cit. on p. 8).
- [24] I. N. Bernstein, I. M. Gelfand, and V. A. Ponomarev. “Coxeter functors and Gabriel’s theorem”. In: *Russian Mathematical Surveys* 28.2 (1973), pp. 17–32 (cit. on pp. 14, 27, 99, 167, 175, 178, 181, 182, 185–187).
- [25] J. Berwald, M. Gidea, and M. Vejdemo-Johansson. “Automatic recognition and tagging of topologically different regimes in dynamical systems”. In: *arXiv preprint arXiv:1312.2482* (2013) (cit. on p. 8).
- [26] G. Biau, F. Chazal, D. Cohen-Steiner, L. Devroye, and C. Rodríguez. “A Weighted k-Nearest Neighbor Density Estimate for Geometric Inference”. In: *Electronic Journal of Statistics* 5 (2011), pp. 204–237 (cit. on p. 126).
- [27] C. M. Bishop. *Neural networks for pattern recognition*. Oxford university press, 1995 (cit. on p. 68).
- [28] J.-D. Boissonnat, L. J. Guibas, and S. Y. Oudot. “Manifold Reconstruction in Arbitrary Dimensions using Witness Complexes”. In: *Discrete and Computational Geometry* 42.1 (2009), pp. 37–70 (cit. on pp. 92, 166).
- [29] J.-D. Boissonnat, F. Chazal, and M. Yvinec. “Computational Topology Inference”. <http://geometrica.saclay.inria.fr/team/Fred.Chazal/>

- papers/CGLcourseNotes/main.pdf. Notes from the course *Computational Geometry Learning* given at the *Parisian Master of Research in Computer Science* (MPRI). 2014 (cit. on p. 42).
- [30] J.-D. Boissonnat, T. K. Dey, and C. Maria. “The Compressed Annotation Matrix: an Efficient Data Structure for Computing Persistent Cohomology”. In: *Proc. European Symposium on Algorithms*. 2013, pp. 695–706 (cit. on pp. 45, 157).
- [31] J.-D. Boissonnat and A. Ghosh. “Manifold reconstruction using tangential Delaunay complexes”. In: *Discrete & Computational Geometry* 51.1 (2014), pp. 221–267 (cit. on p. 166).
- [32] J.-D. Boissonnat and C. Maria. “Computing persistent homology with various coefficient fields in a single pass”. In: *Algorithms-ESA 2014*. Springer, 2014, pp. 185–196 (cit. on p. 44).
- [33] J.-D. Boissonnat and C. Maria. “The simplex tree: An efficient data structure for general simplicial complexes”. In: *Proc. European Symposium on Algorithms*. 2012, pp. 731–742 (cit. on pp. 45, 157).
- [34] K. Borsuk. “On the imbedding of systems of compacta in simplicial complexes”. eng. In: *Fundamenta Mathematicae* 35.1 (1948), pp. 217–234 (cit. on p. 82).
- [35] K. Borsuk. “Sur les rétractes”. In: *Fund. Math* 17 (1931), pp. 152–170 (cit. on p. 82).
- [36] M. Brion. *Lecture Notes — Representations of Quivers*. Summer School “Geometric Methods in Representation Theory”. 2008 (cit. on pp. 167, 190).
- [37] H. Brönnimann and M. Yvinec. “Efficient exact evaluation of signs of determinants”. In: *Algorithmica* 27.1 (2000), pp. 21–56 (cit. on p. 86).
- [38] A. M. Bronstein, M. M. Bronstein, and R. Kimmel. “Topology-invariant similarity of nonrigid shapes”. In: *Intl. Journal of Computer Vision (IJCV)* 81.3 (2009), pp. 281–301 (cit. on p. 134).
- [39] P. Bubenik. “Statistical Topological Data Analysis using Persistence Landscapes”. In: *Journal of Machine Learning Research* 16 (2015), pp. 77–102 (cit. on pp. 159, 160).
- [40] P. Bubenik, V. de Silva, and J. Scott. *Metrics for generalized persistence modules*. Research Report arXiv:1312.3829 [math.AT]. 2013 (cit. on pp. 27, 62, 163, 165).
- [41] P. Bubenik and J. Scott. “Categorification of Persistent Homology”. In: *Discrete and Computational Geometry* (2014). To appear (cit. on pp. 27, 49, 62, 163, 165).
- [42] M. Buchet, F. Chazal, S. Y. Oudot, and D. Sheehy. “Efficient and Robust Persistent Homology for Measures”. In: *Proceedings of the ACM-SIAM symposium on Discrete algorithms*. 2015, pp. 168–180. DOI: 10.1137/1.9781611973730.13 (cit. on pp. 111, 112).
- [43] J. R. Bunch and J. Hopcroft. “Factorization and Inversion by Fast Matrix Multiplication”. In: *Mathematics of Computation* 28.125 (1974), pp. 231–236 (cit. on p. 46).
- [44] D. Burago, Y. Burago, and S. Ivanov. *A Course in Metric Geometry*. Vol. 33. Graduate Studies in Mathematics. Providence, RI: American Mathematical Society, 2001 (cit. on pp. 140, 143).

- [45] D. Burghilea and T. K. Dey. “Topological persistence for circle-valued maps”. In: *Discrete & Computational Geometry* 50.1 (2013), pp. 69–98 (cit. on p. 164).
- [46] C. Caillerie and B. Michel. “Model Selection for Simplicial Approximation”. In: *Foundations of Computational Mathematics* 11.6 (2011), pp. 707–731 (cit. on p. 158).
- [47] G. Carlsson, A. Zomorodian, A. Collins, and L. J. Guibas. “Persistence Barcodes for Shapes”. In: *International Journal of Shape Modeling* 11.2 (2005), pp. 149–187 (cit. on p. 135).
- [48] G. Carlsson. “Topology and data”. In: *Bulletin of the American Mathematical Society* 46.2 (2009), pp. 255–308 (cit. on pp. vii, 69, 161).
- [49] G. Carlsson and V. de Silva. “Zigzag Persistence”. In: *Foundations of Computational Mathematics* 10.4 (2010), pp. 367–405 (cit. on pp. 14, 27, 39, 44, 167, 175, 182–186).
- [50] G. Carlsson, V. de Silva, and D. Morozov. “Zigzag Persistent Homology and Real-valued Functions”. In: *Proceedings of the Twenty-fifth Annual Symposium on Computational Geometry*. SCG ’09. 2009, pp. 247–256 (cit. on pp. 27, 37–39, 44, 45, 103, 185).
- [51] G. Carlsson and F. Mémoli. “Characterization, Stability and Convergence of Hierarchical Clustering Methods”. In: *J. Machine Learning Research* 11 (Aug. 2010), pp. 1425–1470 (cit. on p. 117).
- [52] G. Carlsson, G. Singh, and A. Zomorodian. “Computing multidimensional persistence”. In: *Algorithms and computation*. Springer, 2009, pp. 730–739 (cit. on p. 164).
- [53] G. Carlsson and A. Zomorodian. “The Theory of Multidimensional Persistence”. In: *Discrete and Computational Geometry* 42.1 (May 2009), pp. 71–93 (cit. on p. 164).
- [54] G. Carlsson, T. Ishkhanov, V. de Silva, and A. Zomorodian. “On the Local Behavior of Spaces of Natural Images”. In: *Int. J. Comput. Vision* 76.1 (Jan. 2008), pp. 1–12. DOI: 10.1007/s11263-007-0056-x (cit. on pp. 8, 106–109).
- [55] H. Carr, J. Snoeyink, and U. Axen. “Computing contour trees in all dimensions”. In: *Proceedings of the eleventh annual ACM-SIAM symposium on Discrete algorithms*. Society for Industrial and Applied Mathematics. 2000, pp. 918–926 (cit. on p. 7).
- [56] M. Carrière, S. Y. Oudot, and M. Ovsjanikov. “Stable Topological Signatures for Points on 3D Shapes”. In: *Computer Graphics Forum (proc. Symposium on Geometry Processing)* (2015) (cit. on pp. 155, 160).
- [57] F. Chazal and A. Lieutier. “The λ -medial axis”. In: *Graphical Models* 67.4 (2005), pp. 304–331 (cit. on pp. 75, 77).
- [58] F. Chazal, D. Cohen-Steiner, L. J. Guibas, F. Mémoli, and S. Y. Oudot. “Gromov-Hausdorff Stable Signatures for Shapes using Persistence”. In: *Computer Graphics Forum (proc. Symposium on Geometry Processing)* (2009), pp. 1393–1403. DOI: 10.1111/j.1467-8659.2009.01516.x (cit. on pp. 8, 136, 137, 139, 147).
- [59] F. Chazal and D. Cohen-Steiner. “Geometric Inference”. In: *Tessellations in the Sciences*. To appear. Springer-Verlag, 2013 (cit. on pp. 71, 76).

- [60] F. Chazal, D. Cohen-Steiner, and A. Lieutier. “A Sampling Theory for Compact Sets in Euclidean Space.” In: *Discrete & Computational Geometry* 41.3 (2009), pp. 461–479 (cit. on pp. 74, 77).
- [61] F. Chazal, D. Cohen-Steiner, and A. Lieutier. “Normal Cone Approximation and Offset Shape Isotopy”. In: *Comput. Geom. Theory Appl.* 42.6-7 (2009), pp. 566–581 (cit. on p. 78).
- [62] F. Chazal, D. Cohen-Steiner, and Q. Mérigot. “Geometric Inference for Probability Measures”. In: *J. Foundations of Computational Mathematics* 11.6 (2011), pp. 733–751 (cit. on pp. 78, 110).
- [63] F. Chazal, W. Crawley-Boevey, and V. de Silva. *The observable structure of persistence modules*. Research Report arXiv:1405.5644 [math.RT]. 2014 (cit. on pp. 28, 63, 203).
- [64] F. Chazal, V. de Silva, and S. Oudot. “Persistence stability for geometric complexes”. English. In: *Geometriae Dedicata* (2013), pp. 1–22. DOI: 10.1007/s10711-013-9937-z (cit. on pp. 144, 145).
- [65] F. Chazal and A. Lieutier. “Stability and Computation of Topological Invariants of Solids in \mathbb{R}^n ”. In: *Discrete & Computational Geometry* 37.4 (2007), pp. 601–617 (cit. on p. 75).
- [66] F. Chazal and S. Y. Oudot. “Towards Persistence-Based Reconstruction in Euclidean Spaces”. In: *Proceedings of the 24th ACM Symposium on Computational Geometry*. 2008, pp. 232–241. DOI: 10.1145/1377676.1377719 (cit. on pp. 62, 82, 89–91, 95, 96, 166).
- [67] F. Chazal, L. Guibas, S. Oudot, and P. Skraba. “Analysis of Scalar Fields over Point Cloud Data”. In: *Discrete & Computational Geometry* 46.4 (2011), pp. 743–775 (cit. on pp. 130, 131).
- [68] F. Chazal, M. Glisse, C. Labruère, and B. Michel. “Convergence rates for persistence diagram estimation in Topological Data Analysis”. In: *Proceedings of The 31st International Conference on Machine Learning*. 2014, pp. 163–171 (cit. on p. 159).
- [69] F. Chazal, B. T. Fasy, F. Lecci, A. Rinaldo, A. Singh, and L. Wasserman. “On the Bootstrap for Persistence Diagrams and Landscapes”. In: *arXiv preprint arXiv:1311.0376* (2013) (cit. on p. 159).
- [70] F. Chazal, L. J. Guibas, S. Oudot, and P. Skraba. “Persistence-Based Clustering in Riemannian Manifolds”. In: *J. ACM* 60.6 (Nov. 2013), 41:1–41:38. DOI: 10.1145/2535927 (cit. on pp. 8, 118, 119, 122–129).
- [71] F. Chazal, D. Cohen-Steiner, M. Glisse, L. J. Guibas, and S. Y. Oudot. “Proximity of Persistence Modules and their Diagrams”. In: *Proceedings of the 25th Symposium on Computational Geometry*. 2009, pp. 237–246 (cit. on pp. 25, 26, 55–57, 62).
- [72] F. Chazal, V. de Silva, M. Glisse, and S. Y. Oudot. *The structure and stability of persistence modules*. Research Report arXiv:1207.3674 [math.AT]. To appear as a volume of the SpringerBriefs in Mathematics. 2012 (cit. on pp. vii, 22–27, 32, 33, 49, 51, 55, 58–60, 62).
- [73] B. Chazelle and J. Matoušek. “On linear-time deterministic algorithms for optimization problems in fixed dimension”. In: *Journal of Algorithms* 21.3 (1996), pp. 579–597 (cit. on p. 86).

- [74] J. Cheeger. “Critical points of distance functions and applications to geometry”. In: *Geometric topology: recent developments*. Lecture Notes in Mathematics 1504. Springer, 1991, pp. 1–38 (cit. on p. 75).
- [75] C. Chen and M. Kerber. “An Output-Sensitive Algorithm for Persistent Homology”. In: *Computational Geometry: Theory and Applications* 46.4 (2013), pp. 435–447 (cit. on p. 46).
- [76] W.-Y. Chen, Y. Song, H. Bai, C.-J. Lin, and E. Y. Chang. *PSC: Parallel Spectral Clustering*. <http://www.cs.ucsb.edu/~wychen/sc>. 2008 (cit. on p. 127).
- [77] S.-W. Cheng, T. K. Dey, and E. A. Ramos. “Manifold reconstruction from point samples”. In: *Proc. 16th Sympos. Discrete Algorithms*. 2005, pp. 1018–1027 (cit. on pp. 92, 166).
- [78] J. D. Chodera, N. Singhal, V. S. Pande, K. A. Dill, and W. C. Swope. “Automatic discovery of metastable states for the construction of Markov models of macromolecular conformational dynamics”. In: *The Journal of Chemical Physics* 126.15, 155101 (2007), p. 155101 (cit. on p. 128).
- [79] J. D. Chodera, W. C. Swope, J. W. Pitera, and K. A. Dill. “Long-Time Protein Folding Dynamics from Short-Time Molecular Dynamics Simulations”. In: *Multiscale Modeling & Simulation* 5.4 (2006), pp. 1214–1226. DOI: 10.1137/06065146X (cit. on p. 128).
- [80] J. D. Chodera, W. C. Swope, J. W. Pitera, C. Seok, and K. A. Dill. “Use of the Weighted Histogram Analysis Method for the Analysis of Simulated and Parallel Tempering Simulations”. In: *Journal of Chemical Theory and Computation* 3.1 (2007), pp. 26–41. DOI: 10.1021/ct0502864. eprint: <http://pubs.acs.org/doi/pdf/10.1021/ct0502864> (cit. on p. 128).
- [81] T. Y. Chow. “You Could Have Invented Spectral Sequences”. In: *Notices of the AMS* (2006), pp. 15–19 (cit. on p. 7).
- [82] M. K. Chung, P. Bubenik, and P. T. Kim. “Persistence diagrams of cortical surface data”. In: *Information Processing in Medical Imaging*. Springer, 2009, pp. 386–397 (cit. on p. 8).
- [83] F. H. Clarke. *Optimization and nonsmooth analysis*. Vol. 5. Siam, 1990 (cit. on p. 75).
- [84] K. L. Clarkson. “Safe and effective determinant evaluation”. In: *Foundations of Computer Science, 1992. Proceedings., 33rd Annual Symposium on*. IEEE, 1992, pp. 387–395 (cit. on p. 86).
- [85] D. Cohen-Steiner, H. Edelsbrunner, and D. Morozov. “Vines and Vineyards by Updating Persistence in Linear Time”. In: *Proceedings of the Annual Symposium on Computational Geometry*. 2006, pp. 119–126 (cit. on p. 44).
- [86] D. Cohen-Steiner, H. Edelsbrunner, and J. Harer. “Extending persistence using Poincaré and Lefschetz duality”. In: *Foundations of Computational Mathematics* 9.1 (2009), pp. 79–103 (cit. on pp. 7, 34, 36, 44).
- [87] D. Cohen-Steiner, H. Edelsbrunner, and J. Harer. “Stability of Persistence Diagrams”. In: *Discrete Comput. Geom.* 37.1 (Jan. 2007), pp. 103–120. DOI: 10.1007/s00454-006-1276-5 (cit. on pp. 26, 32, 49, 61, 62, 75).
- [88] D. Cohen-Steiner, H. Edelsbrunner, J. Harer, and Y. Mileyko. “Lipschitz functions have L_p -stable persistence”. In: *Foundations of Computational Mathematics* 10.2 (2010), pp. 127–139 (cit. on p. 63).

- [89] D. Cohen-Steiner, H. Edelsbrunner, J. Harer, and D. Morozov. “Persistent Homology for Kernels, Images, and Cokernels”. In: *Proceedings of the Twentieth Annual ACM-SIAM Symposium on Discrete Algorithms*. SODA '09. New York, New York: Society for Industrial and Applied Mathematics, 2009, pp. 1011–1020 (cit. on pp. 39, 44, 131).
- [90] É. Colin de Verdière, G. Ginot, and X. Goaoc. “Helly numbers of acyclic families”. In: *Advances in Mathematics* 253 (2014), pp. 163–193 (cit. on pp. 161, 162).
- [91] D. Comaniciu and P. Meer. “Mean Shift: A Robust Approach Toward Feature Space Analysis”. In: *IEEE Trans. on Pattern Analysis and Machine Intelligence* 24.5 (2002), pp. 603–619 (cit. on p. 116).
- [92] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein. *Introduction to Algorithms*. 2nd. Cambridge, MA: MIT Press, 2001 (cit. on pp. 47, 88, 121).
- [93] W. Crawley-Boevey. *Decomposition of pointwise finite-dimensional persistence modules*. Research Report arXiv:1210.0819 [math.RT]. 2012 (cit. on pp. 19, 27, 167, 195).
- [94] W. Crawley-Boevey. *Lectures on Representations of Quivers*. Notes for a course of eight lectures given in Oxford in Spring 1992. 1992 (cit. on p. 191).
- [95] W. Crawley-Boevey. *Personal notes, published as part of [63]*. 2013 (cit. on pp. 27, 28).
- [96] M. d’Amico, P. Frosini, and C. Landi. *Natural pseudo-distance and optimal matching between reduced size functions*. Tech. rep. 66. Univ. degli Studi di Modena e Reggio Emilia, Italy.: DISMI, 2005 (cit. on p. 135).
- [97] M. d’Amico, P. Frosini, and C. Landi. “Using matching distance in size theory: A survey”. In: *IJIST* 16.5 (2006), pp. 154–161 (cit. on pp. 7, 135).
- [98] V. de Silva and R. Ghrist. “Homological sensor networks”. In: *Notices of the American Mathematical Society* 54.1 (2007), pp. 10–17 (cit. on pp. 8, 68).
- [99] V. de Silva, D. Morozov, and M. Vejdemo-Johansson. “Dualities in persistent (co)homology”. In: *Inverse Problems* 27 (2011), p. 124003 (cit. on pp. 42, 45).
- [100] V. de Silva, D. Morozov, and M. Vejdemo-Johansson. “Persistent cohomology and circular coordinates”. In: *Discrete Comput. Geom.* 45 (2011), pp. 737–759 (cit. on p. 45).
- [101] V. de Silva. “A weak characterisation of the Delaunay triangulation”. In: *Geometriae Dedicata* 135.1 (2008), pp. 39–64 (cit. on p. 91).
- [102] V. de Silva and G. Carlsson. “Topological estimation using witness complexes”. In: *IEEE Symposium on Point-based Graphic* (2004), pp. 157–166 (cit. on pp. 91, 106, 108).
- [103] V. de Silva and R. Ghrist. “Coverage in Sensor Networks via Persistent Homology”. In: *Algebraic and Geometric Topology* 7 (2007), pp. 339–358 (cit. on p. 89).
- [104] B. Deng, J. Du, B. Parshall, and J. Wang. *Finite Dimensional Algebras and Quantum Groups*. Vol. 150. Mathematical Surveys and Monographs. American Mathematical Society, 2008 (cit. on pp. 167, 178, 188, 191).
- [105] M.-L. Dequeant, S. Ahnert, H. Edelsbrunner, T. M. Fink, E. F. Glynn, G. Hattem, A. Kudlicki, Y. Mileyko, J. Morton, A. R. Mushegian, et al. “Comparison of pattern detection methods in microarray time series of the segmentation clock”. In: *PLoS One* 3.8 (2008), e2856 (cit. on p. 8).

- [106] H. Derksen and J. Weyman. “Quiver representations”. In: *Notices of the American Mathematical Society* 52.2 (2005), pp. 200–206 (cit. on p. 167).
- [107] T. K. Dey, F. Fan, and Y. Wang. “Computing Topological Persistence for Simplicial Maps”. In: *Proc. 30th Annual Symposium on Computational Geometry*. 2014 (cit. on pp. 39, 45, 104, 148, 157).
- [108] T. K. Dey and Y. Wang. “Reeb graphs: approximation and persistence”. In: *Discrete & Computational Geometry* 49.1 (2013), pp. 46–73 (cit. on p. 7).
- [109] T. K. Dey, J. Giesen, E. A. Ramos, and B. Sadri. “Critical Points of the Distance to an ε -sampling of a Surface and Flow-complex-based Surface Reconstruction”. In: *Proceedings of the Twenty-first Annual Symposium on Computational Geometry*. 2005, pp. 218–227 (cit. on p. 77).
- [110] P. Donovan and M. R. Freislich. *The representation theory of finite graphs and associated algebras*. Carleton Math. Lecture Notes 5. 1973 (cit. on p. 192).
- [111] D. S. Dummit and R. M. Foote. *Abstract algebra*. Vol. 1999. Prentice Hall Englewood Cliffs, 1991 (cit. on pp. 13, 167).
- [112] H. Edelsbrunner. “The union of balls and its dual shape”. In: *Discrete & Computational Geometry* 13.1 (1995), pp. 415–440 (cit. on p. 83).
- [113] H. Edelsbrunner, D. G. Kirkpatrick, and R. Seidel. “On the shape of a set of points in the plane”. In: *IEEE Trans. Inform. Theory* IT-29 (1983), pp. 551–559 (cit. on p. 83).
- [114] H. Edelsbrunner and J. Harer. “Persistent homology — a survey”. In: *Contemporary mathematics* 453 (2008), pp. 257–282 (cit. on p. vii).
- [115] H. Edelsbrunner and J. L. Harer. *Computational topology: an introduction*. AMS Bookstore, 2010 (cit. on pp. vii, viii, 8, 42, 43, 63, 123).
- [116] H. Edelsbrunner, G. Jabłonski, and M. Mrozek. “The persistent homology of a self-map”. Preprint. 2013 (cit. on p. 164).
- [117] H. Edelsbrunner and M. Kerber. “Alexander Duality for Functions: The Persistent Behavior of Land and Water and Shore”. In: *Proceedings of the Twenty-eighth Annual Symposium on Computational Geometry*. SoCG '12. Chapel Hill, North Carolina, USA: ACM, 2012, pp. 249–258. DOI: 10.1145/2261250.2261287 (cit. on p. 36).
- [118] H. Edelsbrunner, D. Letscher, and A. Zomorodian. “Topological Persistence and Simplification”. In: *Discrete and Computational Geometry* 28 (2002), pp. 511–533 (cit. on pp. 26, 30, 43, 46, 165).
- [119] H. Edelsbrunner and D. Morozov. “Persistent Homology: Theory and Practice”. In: *Proceedings of the European Congress of Mathematics*. 2012 (cit. on p. vii).
- [120] H. Edelsbrunner, D. Morozov, and V. Pascucci. “Persistence-sensitive simplification of functions on 2-manifolds”. In: *Proceedings of the twenty-second annual symposium on Computational geometry*. 2006, pp. 127–134 (cit. on p. 166).
- [121] H. Edelsbrunner and S. Parsa. “On the Computational Complexity of Betti Numbers: Reductions from Matrix Rank”. In: *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms*. SIAM. 2014 (cit. on p. 156).
- [122] S. Eilenberg and N. Steenrod. *Foundations of algebraic topology*. Vol. 952. Princeton NJ, 1952 (cit. on p. 81).

- [123] D. J. Eisenstein, D. H. Weinberg, E. Agol, H. Aihara, C. Allende Prieto, S. F. Anderson, J. A. Arns, É. Aubourg, S. Bailey, E. Balbinot, and et al. “SDSS-III: Massive Spectroscopic Surveys of the Distant Universe, the Milky Way, and Extra-Solar Planetary Systems”. In: *The Astronomical Journal* 142.3, 72 (2011), p. 72. DOI: 10.1088/0004-6256/142/3/72 (cit. on p. 68).
- [124] A. Elad and R. Kimmel. “On Bending Invariant Signatures for Surfaces.” In: *IEEE Trans. Pattern Anal. Mach. Intell.* 25.10 (2003), pp. 1285–1295 (cit. on p. 133).
- [125] K. J. Emmett and R. Rabadan. “Characterizing Scales of Genetic Recombination and Antibiotic Resistance in Pathogenic Bacteria Using Topological Data Analysis”. In: *arXiv preprint arXiv:1406.1219* (2014) (cit. on p. 8).
- [126] E. Escobar and Y. Hiraoka. *Persistence Modules on Commutative Ladders of Finite Type*. Research Report arXiv:1404.7588 [math.AT]. 2014 (cit. on p. 163).
- [127] M. Ester, H.-P. Kriegel, J. Sander, and X. Xu. “A density-based algorithm for discovering clusters in large spatial databases with noise”. In: *Proc. 2nd Internat. Conf. on Knowledge Discovery and Data Mining (KDD-96)*. AAAI Press, 1996, pp. 226–231 (cit. on p. 115).
- [128] H. Federer. “Curvature Measures”. In: *Trans. Amer. Math. Soc.* 93 (1959), pp. 418–491 (cit. on pp. 71, 72).
- [129] K. Fischer, B. Gärtner, and M. Kutz. “Fast smallest-enclosing-ball computation in high dimensions”. In: *Algorithms-ESA 2003*. Springer, 2003, pp. 630–641 (cit. on p. 86).
- [130] R. Forman. “A user’s guide to discrete Morse theory”. In: *Sém. Lothar. Combin* 48 (2002), B48c (cit. on pp. 7, 47).
- [131] P. Frosini. “Discrete Computation of Size Functions”. In: *Journal of Combinatorics, Information, and System Science* 17 (1992), pp. 232–250 (cit. on p. 7).
- [132] P. Gabriel. “Auslander-Reiten sequences and representation-finite algebras”. In: *Representation theory I*. Springer, 1980, pp. 1–71 (cit. on p. 196).
- [133] P. Gabriel. “Unzerlegbare Darstellungen I”. In: *Manuscripta Mathematica* 6 (1972), pp. 71–103 (cit. on pp. 16, 171, 188).
- [134] J. Gamble and G. Heo. “Exploring uses of persistent homology for statistical analysis of landmark-based shape data”. In: *Journal of Multivariate Analysis* 101.9 (2010), pp. 2184–2199 (cit. on p. 8).
- [135] M. Gameiro, Y. Hiraoka, S. Izumi, M. Kramar, K. Mischaikow, and V. Nanda. “A Topological Measurement of Protein Compressibility”. Preprint. 2013 (cit. on p. 8).
- [136] J. Gao, L. J. Guibas, S. Oudot, and Y. Wang. “Geodesic Delaunay Triangulation and Witness Complex in the Plane”. In: *Transactions on Algorithms* 6.4 (2010), pp. 1–67 (cit. on pp. 8, 89).
- [137] B. Gärtner. “A subexponential algorithm for abstract optimization problems”. In: *SIAM Journal on Computing* 24.5 (1995), pp. 1018–1035 (cit. on p. 86).
- [138] B. Gärtner. “Fast and robust smallest enclosing balls”. In: *Algorithms-ESA ’99*. Springer, 1999, pp. 325–338 (cit. on p. 86).
- [139] I. M. Gelfand and V. A. Ponomarev. “Problems of linear algebra and classification of quadruples of subspaces in a finite-dimensional vector space”. In:

- Colloquia Mathematica Societatis Ianos Bolyai*. Vol. 5 (Hilbert space operators). 1970, pp. 163–237 (cit. on p. 189).
- [140] R. Ghrist and A. Muhammad. “Coverage and hole-detection in sensor networks via Homology”. In: *Proc. the 4th International Symposium on Information Processing in Sensor Networks (IPSN’05)*. 2005, pp. 254–260 (cit. on p. 140).
- [141] R. Ghrist. “Barcodes: the persistent topology of data”. In: *Bulletin of the American Mathematical Society* 45.1 (2008), pp. 61–75 (cit. on p. vii).
- [142] R. Ghrist. *Elementary Applied Topology*. CreateSpace Independent Publishing Platform, 2014 (cit. on p. viii).
- [143] M. Gromov. “Hyperbolic groups”. In: *Essays in group theory*. Ed. by S. M. Gersten. Vol. 8. MSRI Publications. Springer-Verlag, 1987, pp. 75–263 (cit. on p. 89).
- [144] M. Gromov. *Metric structures for Riemannian and non-Riemannian spaces*. Vol. 152. Progress in Mathematics. Boston, MA: Birkhäuser Boston Inc., 1999, pp. xx+585 (cit. on p. 133).
- [145] K. Grove. “Critical point theory for distance functions”. In: *Proc. of Symp. in Pure Mathematics*. Vol. 54.3. 1993 (cit. on p. 75).
- [146] L. Guibas, Q. Mérigot, and D. Morozov. “Witnessed k -distance”. In: *Discrete & Computational Geometry* 49.1 (2013), pp. 22–45 (cit. on p. 112).
- [147] L. J. Guibas and S. Y. Oudot. “Reconstruction Using Witness Complexes”. In: *Discrete and Computational Geometry* 40.3 (2008), pp. 325–356 (cit. on pp. 95, 98).
- [148] A. B. Hamza and H. Krim. “Geodesic Object Representation and Recognition”. In: *Lecture Notes in Computer Science*. Vol. 2886. Springer, 2003, pp. 378–387 (cit. on p. 133).
- [149] S. Har-Peled and M. Mendel. “Fast construction of nets in low-dimensional metrics and their applications”. In: *SIAM Journal on Computing* 35.5 (2006), pp. 1148–1184 (cit. on pp. 97, 113, 150, 156).
- [150] J. Hartigan. *Clustering Algorithms*. Wiley, 1975 (cit. on p. 115).
- [151] T. Hastie, R. Tibshirani, and J. H. Friedman. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer Series in Statistics. Second edition. Springer, 2009 (cit. on pp. 70, 117).
- [152] A. Hatcher. *Algebraic Topology*. Cambridge Univ. Press, 2001 (cit. on pp. 45, 71, 81, 82).
- [153] J. H. van Hateren and A. van der Schaaff. “Independent component filters of natural images compared with simple cells in primary visual cortex.” In: *Proceedings of the Royal Society, London, B* 265 (1997), pp. 359–366 (cit. on p. 106).
- [154] J.-C. Hausmann. “On the Vietoris-Rips complexes and a cohomology theory for metric spaces”. In: *Prospects in topology, Ann. of Math. Stud.* 138 (1995), pp. 175–188 (cit. on p. 136).
- [155] G. Heo, J. Gamble, and P. T. Kim. “Topological analysis of variance and the maxillary complex”. In: *Journal of the American Statistical Association* 107.498 (2012), pp. 477–492 (cit. on p. 8).
- [156] M. Hilaga, Y. Shinagawa, T. Kohmura, and T. L. Kunii. “Topology matching for fully automatic similarity estimation of 3D shapes”. In: *SIGGRAPH ’01: Proceedings of the 28th annual conference on Computer graphics and*

- interactive techniques*. New York, NY, USA: ACM, 2001, pp. 203–212. DOI: <http://doi.acm.org/10.1145/383259.383282> (cit. on p. 133).
- [157] M. Hilbert and P. López. “The world’s technological capacity to store, communicate, and compute information”. In: *Science* 332.6025 (2011), pp. 60–65 (cit. on p. 67).
- [158] D. Horak, S. Maletić, and M. Rajković. “Persistent homology of complex networks”. In: *Journal of Statistical Mechanics: Theory and Experiment* 2009.03 (2009), P03034 (cit. on p. 8).
- [159] B. Hudson, G. Miller, and T. Phillips. “Sparse Voronoi Refinement”. In: *Proceedings of the 15th International Meshing Roundtable*. Long version available as Carnegie Mellon University Technical Report CMU-CS-06-132. Birmingham, Alabama, 2006, pp. 339–356 (cit. on p. 93).
- [160] B. Hudson, G. L. Miller, T. Phillips, and D. R. Sheehy. “Size Complexity of Volume Meshes vs. Surface Meshes”. In: *SODA: ACM-SIAM Symposium on Discrete Algorithms*. 2009 (cit. on p. 93).
- [161] B. Hudson, G. L. Miller, S. Y. Oudot, and D. R. Sheehy. “Topological inference via meshing”. In: *Proceedings of the 2010 ACM Symposium on Computational geometry*. Snowbird, Utah, USA: ACM, 2010, pp. 277–286. DOI: 10.1145/1810959.1811006 (cit. on pp. 92–94).
- [162] W. Huisinga and B. Schmidt. “Advances in Algorithms for Macromolecular Simulation, chapter Metastability and dominant eigenvalues of transfer operators”. In: *Lecture Notes in Computational Science and Engineering*. Springer (2005) (cit. on pp. 128, 129).
- [163] N. Jacobson. *Basic algebra I*. DoverPublications, 2009 (cit. on pp. 18, 194).
- [164] N. Jacobson. *Basic algebra II*. DoverPublications, 2009 (cit. on p. 184).
- [165] A. K. Jain and R. C. Dubes. *Algorithms for clustering data*. Prentice-Hall, Inc., 1988 (cit. on pp. 115, 117).
- [166] W. B. Johnson and J. Lindenstrauss. “Extensions of Lipschitz mappings into a Hilbert space”. In: *Contemporary mathematics* 26.189-206 (1984), p. 1 (cit. on p. 106).
- [167] M. Joswig and M. E. Pfetsch. “Computing optimal Morse matchings”. In: *SIAM Journal on Discrete Mathematics* 20.1 (2006), pp. 11–25 (cit. on p. 48).
- [168] V. Kac. “Infinite root systems, representations of graphs and invariant theory”. In: *Inventiones math.* 56 (1980), pp. 57–92 (cit. on p. 191).
- [169] P. M. Kasson, A. Zomorodian, S. Park, N. Singhal, L. J. Guibas, and V. S. Pande. “Persistent voids: a new structural metric for membrane fusion”. In: *Bioinformatics* 23.14 (2007), pp. 1753–1759 (cit. on p. 8).
- [170] W. L. Koontz, P. M. Narendra, and K. Fukunaga. “A Graph-Theoretic Approach to Nonparametric Cluster Analysis”. In: *IEEE Trans. on Computers* 24 (1976), pp. 936–944 (cit. on pp. 116, 118, 119).
- [171] M. Kramar, A. Goulet, L. Kondic, and K. Mischaikow. “Persistence of force networks in compressed granular media”. In: *Physical Review E* 87.4 (2013), p. 042207 (cit. on p. 8).
- [172] H. Krause. “Representations of Quivers via Reflection Functors”. In: *arXiv preprint arXiv:0804.1428* (2010) (cit. on pp. 167, 171, 173, 192, 193).

- [173] P. Kumar, J. S. Mitchell, and E. A. Yildirim. “Approximate minimum enclosing balls in high dimensions using core-sets”. In: *Journal of Experimental Algorithmics (JEA)* 8 (2003), pp. 1–1 (cit. on p. 87).
- [174] G. N. Lance and W. T. Williams. “A General Theory of Classificatory Sorting Strategies 1. Hierarchical Systems”. In: *The Computer Journal* 9.4 (1967), pp. 373–380 (cit. on p. 116).
- [175] J. Latschev. “Vietoris-Rips complexes of metric spaces near a closed Riemannian manifold”. In: *Archiv der Mathematik* 77.6 (2001), pp. 522–528 (cit. on p. 136).
- [176] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner. “Gradient-based learning applied to document recognition”. In: *Proceedings of the IEEE* 86.11 (1998), pp. 2278–2324 (cit. on p. 67).
- [177] A. B. Lee, K. S. Pederson, and D. Mumford. “The Nonlinear Statistics of High-contrast Patches in Natural Images”. In: *Internat. J. of Computer Vision* 54.1-3 (2003), pp. 83–103 (cit. on pp. 68, 106).
- [178] J. Leskovec and A. Krevl. *SNAP Datasets: Stanford Large Network Dataset Collection*. <http://snap.stanford.edu/data>. June 2014 (cit. on p. 67).
- [179] M. Lesnick. *The Theory of the Interleaving Distance on Multidimensional Persistence Modules*. Research Report arXiv:1106.5305 [cs.CG]. 2011 (cit. on pp. 27, 49, 53, 62, 63).
- [180] A. Lieutier. “Any open bounded subset of \mathbb{R}^n has the same homotopy type as its medial axis”. In: *Computer-Aided Design* 36.11 (2004), pp. 1029–1046 (cit. on pp. 71, 74).
- [181] S. P. Lloyd. “Least squares quantization in pcm”. In: *IEEE Trans. on Information Theory* 28.2 (1982), pp. 129–136 (cit. on p. 115).
- [182] U. von Luxburg. “A Tutorial on Spectral Clustering”. In: *Statistics and Computing* 17.4 (2007), pp. 395–416 (cit. on p. 115).
- [183] L. J. van der Maaten, E. O. Postma, and H. J. van den Herik. “Dimensionality reduction: A comparative review”. In: *Journal of Machine Learning Research* 10.1-41 (2009), pp. 66–71 (cit. on p. 3).
- [184] S. Mac Lane. *Categories for the Working Mathematician*. Graduate Texts in Mathematics 5. Springer-Verlag, 1971 (cit. on pp. 13, 167).
- [185] S. Manay, D. Cremers, B. Hong, A. Yezzi, and S. Soatto. “Integral Invariants for Shape Matching”. In: *PAMI* 28.10 (2006), pp. 1602–1618 (cit. on p. 133).
- [186] C. Maria and S. Y. Oudot. “Zigzag Persistence via Reflections and Transpositions”. In: *Proc. ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 2015, pp. 181–199 (cit. on p. 157).
- [187] S. Martin, A. Thompson, E. A. Coutsias, and J.-P. Watson. “Topology of cyclo-octane energy landscape”. In: *The journal of chemical physics* 132.23 (2010), p. 234115 (cit. on p. 68).
- [188] P. Massart and J. Picard. *Concentration inequalities and model selection*. *Ecole d’Été de Probabilités de Saint-Flour XXXIII*. Vol. 1896. Lecture Notes in Mathematics. Springer, 2007 (cit. on p. 158).
- [189] P. McMullen. “The maximal number of faces of a convex polytope”. In: *Mathematika* 17 (1970), pp. 179–184 (cit. on p. 86).
- [190] N. Megiddo. “Linear programming in linear time when the dimension is fixed”. In: *Journal of the ACM (JACM)* 31.1 (1984), pp. 114–127 (cit. on p. 86).

- [191] Q. Mérigot. “Lower Bounds for k -distance Approximation”. In: *Proceedings of the Twenty-ninth Annual Symposium on Computational Geometry*. SoCG '13. Rio de Janeiro, Brazil: ACM, 2013, pp. 435–440. DOI: 10.1145/2462356.2462367 (cit. on p. 113).
- [192] *Michael Lesnick*. Personal communication. 2013 (cit. on p. 21).
- [193] Y. Mileyko, S. Mukherjee, and J. Harer. “Probability measures on the space of persistence diagrams”. In: *Inverse Problems* 27.12 (2011), p. 124007 (cit. on p. 158).
- [194] G. L. Miller, T. Phillips, and D. R. Sheehy. “Linear-size meshes”. In: *CCCG: Canadian Conference in Computational Geometry*. 2008 (cit. on p. 94).
- [195] J. W. Milnor. *Morse Theory*. Princeton, NJ: Princeton University Press, 1963 (cit. on pp. 6, 29, 75).
- [196] N. Milosavljevic, D. Morozov, and P. Skraba. *Zigzag Persistent Homology in Matrix Multiplication Time*. Research Report RR-7393. INRIA, 2010 (cit. on p. 46).
- [197] N. Milosavljevic, D. Morozov, and P. Skraba. “Zigzag Persistent Homology in Matrix Multiplication Time”. In: *Proc. 27th Annual Symposium on Computational Geometry*. 2011 (cit. on pp. 46, 156).
- [198] K. Mischaikow and V. Nanda. “Morse theory for filtrations and efficient computation of persistent homology”. In: *Discrete & Computational Geometry* 50.2 (2013), pp. 330–353 (cit. on pp. 7, 48).
- [199] D. Morozov. “Homological Illusions of Persistence and Stability”. Ph.D. dissertation. Department of Computer Science, Duke University, 2008 (cit. on pp. 43–46, 156).
- [200] M. Mrozek and B. Batko. “Coreduction homology algorithm”. In: *Discrete & Computational Geometry* 41.1 (2009), pp. 96–118 (cit. on p. 48).
- [201] K. Mulmuley. “On levels in arrangements and Voronoi diagrams”. In: *Discrete & Computational Geometry* 6.1 (1991), pp. 307–338 (cit. on p. 112).
- [202] J. R. Munkres. *Elements of algebraic topology*. Vol. 2. Addison-Wesley Reading, 1984 (cit. on pp. 29, 42, 136, 143).
- [203] S. F. Nayar, S. A. Nene, and H. Murase. *Columbia object image library (COIL-100)*. Tech. Rep. CUCS-006-96. Columbia University, 1996 (cit. on pp. 3, 67).
- [204] L. A. Nazarova. “Representations of quivers of infinite type”. In: *Izv. Akad. Nauk SSSR Ser. Mat.* 37 (4 1973), pp. 752–791 (cit. on p. 192).
- [205] M. Nicolau, A. J. Levine, and G. Carlsson. “Topology based data analysis identifies a subgroup of breast cancers with a unique mutational profile and excellent survival”. In: *Proceedings of the National Academy of Sciences* 108.17 (2011), pp. 7265–7270 (cit. on p. 8).
- [206] P. Niyogi, S. Smale, and S. Weinberger. “Finding the Homology of Submanifolds with High Confidence from Random Samples”. In: *Discrete Comput. Geom.* 39.1 (Mar. 2008), pp. 419–441 (cit. on p. 72).
- [207] R. Osada, T. Funkhouser, B. Chazelle, and D. Dobkin. “Shape distributions”. In: *ACM Trans. Graph.* 21.4 (2002), 807–832, New York, NY, USA. DOI: <http://doi.acm.org/10.1145/571647.571648> (cit. on p. 133).
- [208] S. Y. Oudot and D. R. Sheehy. “Zigzag Zoology: Rips Zigzags for Homology Inference”. In: *Foundations of Computational Mathematics* (2014), pp. 1–36.

- DOI: 10.1007/s10208-014-9219-7 (cit. on pp. 39, 86, 95, 99, 103, 107–109, 157).
- [209] S. Paris and F. Durand. “A Topological Approach to Hierarchical Segmentation using Mean Shift”. In: *Proc. IEEE conf. on Computer Vision and Pattern Recognition*. 2007 (cit. on pp. 8, 127).
- [210] G. Petri, M. Scalamiero, I. Donato, and F. Vaccarino. “Topological strata of weighted complex networks”. In: *PloS one* 8.6 (2013), e66506 (cit. on p. 8).
- [211] A. Petrunin. “Semiconcave functions in Alexandrov’s geometry”. In: *Surveys in differential geometry. Vol. XI*. Vol. 11. Surv. Differ. Geom. Int. Press, Somerville, MA, 2007, pp. 137–201 (cit. on p. 71).
- [212] G. Plonka and Y. Zheng. “Relation between total variation and persistence distance and its application in signal processing”. Preprint. 2014 (cit. on p. 8).
- [213] G. Reeb. “Sur les points singuliers d’une forme de pfaff complement integrable ou d’une fonction numerique”. In: *Comptes Rendus Acad. Sciences Paris* 222 (1946), pp. 847–849 (cit. on p. 7).
- [214] J. Reininghaus, S. Huber, U. Bauer, and R. Kwitt. “A Stable Multi-Scale Kernel for Topological Machine Learning”. In: 2015 (cit. on p. 160).
- [215] M. Reuter, F.-E. Wolter, and N. Peinecke. “Laplace spectra as fingerprints for shape matching”. In: *SPM ’05: Proceedings of the 2005 ACM symposium on Solid and physical modeling*. Cambridge, Massachusetts: ACM Press, 2005, pp. 101–106. DOI: <http://doi.acm.org/10.1145/1060244.1060256> (cit. on p. 133).
- [216] C. M. Ringel. “The indecomposable representations of the dihedral 2-groups”. In: *Math. Ann.* 214 (1975), pp. 19–34 (cit. on pp. 19, 195).
- [217] C. M. Ringel and H. Tachikawa. “ QF -3 rings”. In: *Journal für die Reine und Angewandte Mathematik* 272 (1975), pp. 49–72 (cit. on pp. 17, 27, 193).
- [218] R. M. Rustamov. “Laplace-Beltrami eigenfunctions for deformation invariant shape representation”. In: *Symposium on Geometry Processing*. 2007, pp. 225–233 (cit. on p. 133).
- [219] J. Sander, M. Ester, H.-P. Kriegel, and X. Xu. “Density-Based Clustering in Spatial Databases: The Algorithm GDBSCAN and Its Applications”. In: *Data Mining and Knowledge Discovery* 2.2 (1998), pp. 169–194 (cit. on p. 115).
- [220] D. Sheehy. “The Persistent Homology of Distance Functions under Random Projection.” In: *Symposium on Computational Geometry*. 2014, p. 328 (cit. on p. 104).
- [221] D. R. Sheehy. “Linear-Size Approximations to the Vietoris-Rips Filtration”. In: *Discrete & Computational Geometry* 49.4 (2013), pp. 778–796 (cit. on pp. 39, 104, 148, 149, 151).
- [222] Y. A. Sheikh, E. Khan, and T. Kanade. “Mode-seeking by Medoidshifts”. In: *Proc. 11th IEEE Internat. Conf. on Computer Vision (ICCV 2007)*. 2007 (cit. on p. 116).
- [223] P. Shilane, P. Min, M. Kazhdan, and T. Funkhouser. “The Princeton shape benchmark”. In: *Shape modeling applications, 2004. Proceedings. IEEE*. 2004, pp. 167–178 (cit. on p. 67).

- [224] G. Singh, F. Memoli, T. Ishkhanov, G. Sapiro, G. Carlsson, and D. L. Ringach. “Topological analysis of population activity in visual cortex”. In: *Journal of vision* 8.8 (2008) (cit. on p. 8).
- [225] G. Singh, F. Mémoli, and G. E. Carlsson. “Topological Methods for the Analysis of High Dimensional Data Sets and 3D Object Recognition.” In: *Proc. Symposium on Point Based Graphics*. 2007, pp. 91–100 (cit. on p. 161).
- [226] T. Sousbie. “The persistent cosmic web and its filamentary structure—I. Theory and implementation”. In: *Monthly Notices of the Royal Astronomical Society* 414.1 (2011), pp. 350–383 (cit. on p. 8).
- [227] T. Sousbie, C. Pichon, and H. Kawahara. “The persistent cosmic web and its filamentary structure—II. Illustrations”. In: *Monthly Notices of the Royal Astronomical Society* 414.1 (2011), pp. 384–403 (cit. on p. 8).
- [228] E. H. Spanier. *Algebraic topology*. Vol. 55. Springer, 1994 (cit. on p. 76).
- [229] R. W. Sumner and J. Popović. “Deformation transfer for triangle meshes”. In: *SIGGRAPH '04: ACM SIGGRAPH 2004 Papers*. Los Angeles, California: Association for Computing Machinery, 2004, pp. 399–405. DOI: <http://doi.acm.org/10.1145/1186562.1015736> (cit. on pp. 134, 138).
- [230] J. B. Tenenbaum, V. De Silva, and J. C. Langford. “A global geometric framework for nonlinear dimensionality reduction”. In: *Science* 290.5500 (2000). Supplemental material., pp. 2319–2323 (cit. on p. 68).
- [231] D. L. Theobald. “Rapid calculation of RMSDs using a quaternion-based characteristic polynomial.” In: *Acta Crystallographica Section A: Foundations of Crystallography* 61.4 (2005), pp. 478–480 (cit. on p. 128).
- [232] K. Turner, Y. Mileyko, S. Mukherjee, and J. Harer. *Fréchet Means for Distributions of Persistence diagrams*. Research Report arXiv:1206.2790. 2012 (cit. on p. 158).
- [233] A. Vedaldi and S. Soatto. “Quick Shift and Kernel Methods for Mode Seeking”. In: *Proc. European Conf. on Computer Vision (ECCV)*. 2008 (cit. on p. 116).
- [234] M. Vejdemo-Johansson. “Sketches of a platypus: a survey of persistent homology and its algebraic foundations”. In: *Algebraic Topology: Applications and New Directions*. Ed. by U. Tillmann, S. Galatius, and D. Sinha. Vol. 620. Contemporary Mathematics. American Mathematical Soc., 2014 (cit. on p. vii).
- [235] C. Villani. *Topics in optimal transportation*. 58. American Mathematical Soc., 2003 (cit. on p. 63).
- [236] J. Wang. *Geometric structure of high-dimensional data and dimensionality reduction*. Springer, 2012 (cit. on p. 70).
- [237] Y. Wang, H. Ombao, and M. K. Chung. “Persistence Landscape of Functional Signal and Its Application to Epileptic Electroencephalogram Data”. Preprint. 2014 (cit. on p. 8).
- [238] C. Webb. “Decomposition of graded modules”. In: *Proceedings of the American Mathematical Society* 94.4 (1985), pp. 565–571 (cit. on pp. 18, 21, 27, 194).
- [239] E. Welzl. “Smallest enclosing disks (balls and ellipsoids)”. In: *New Results and New Trends in Computer Science*. Ed. by H. Maurer. Vol. 555. Lecture Notes in Computer Science. Springer Berlin Heidelberg, 1991, pp. 359–370. DOI: [10.1007/BFb0038202](https://doi.org/10.1007/BFb0038202) (cit. on p. 86).

- [240] A. Yakovleva. “Representation of a quiver with relations”. In: *Journal of Mathematical Sciences* 89.2 (1998), pp. 1172–1179 (cit. on p. 164).
- [241] R. Zafarani and H. Liu. *Social Computing Data Repository at ASU*. <http://socialcomputing.asu.edu>. 2009 (cit. on p. 67).
- [242] A. Zomorodian. “The tidy set: a minimal simplicial set for computing homology of clique complexes”. In: *Proceedings of the twenty-sixth annual symposium on Computational geometry*. ACM. 2010, pp. 257–266 (cit. on p. 48).
- [243] A. Zomorodian and G. Carlsson. “Computing Persistent Homology”. In: *Discrete Comput. Geom.* 33.2 (2005), pp. 249–274 (cit. on pp. 26, 27, 43).
- [244] A. Zomorodian and G. Carlsson. “Localized homology”. In: *Computational Geometry* 41.3 (2008), pp. 126–148 (cit. on p. 156).
- [245] A. J. Zomorodian. *Topology for computing*. Vol. 16. Cambridge University Press, 2005 (cit. on p. vii).

List of Figures

0.1 A planar point set with several underlying geometric structures	1
0.2 The dendrogram produced by <i>single-linkage</i> on the data set of Figure 0.1	2
0.3 The barcode produced by persistence on the data set of Figure 0.1	2
0.4 A collection of images from the Columbia Object Image Library	3
0.5 A smooth function and a piecewise linear approximation	6
1.1 The four decorated points corresponding to intervals $[b_j^\pm, d_j^\pm]$	22
1.2 A classical example in persistence theory	22
1.3 A decorated point (b^\pm, d^\pm) belonging to a rectangle R	23
1.4 Additivity of μ_V under vertical and horizontal splitting	24
1.5 A quadrant, horizontal strip, vertical strip, and finite rectangle	24
2.1 A simplicial filtration of the octahedron	31
2.2 The full persistence diagram of the height function from Example 1.11	33
2.3 The extended persistence diagram of the height function from Example 1.11	35
2.4 The pyramid for $n = 3$	38
2.5 Run of the matrix reduction algorithm on a filtration of a solid triangle	42
2.6 A filtered planar graph and the spanning tree of negative edges	47
3.1 A smooth function, a piecewise linear approximation, and their diagrams	51
3.2 The snapping rule	54
3.3 Tracking down point $\mathbf{p} \in \text{dgm}(V)$ through its successive matchings	56
3.4 Identifying ϕ, ψ with the families of vertical and horizontal morphisms	58
3.5 Effects of various interpolation parameters on the persistence diagrams	60
4.1 Distribution of Matter in the Universe	68
4.2 Sampling of a curve winding around a torus in \mathbb{R}^3	69
4.3 Examples of medial axes	71
4.4 Some compact sets with positive reach	73
4.5 The generalized gradient of the distance to K	74
4.6 Generalized gradient of d_K and its associated flow	74
4.7 A compact set that is not homotopy equivalent to its small offsets	76
4.8 Example with no homologically valid offset of the point cloud	76

4.9	Barcode of the offsets filtration of the point cloud from Figure 4.2	79
4.10	The corresponding persistence diagrams	80
4.11	Union of balls and corresponding Čech complex	81
5.1	Truncated barcode of the Čech filtration on the Clifford data set	86
5.2	A point cloud P and its Rips complex $R_{2i}(P)$	89
5.3	Landmarks set, witnesses set, and corresponding witness complex	90
5.4	Relationship between the offsets and mesh filtrations	93
5.5	Logscale barcode of the mesh filtration on the Clifford data set	95
5.6	Scale-rank plot obtained from the data set of Figure 4.2	96
5.7	Logscale barcode of the oscillating Rips zigzag on the Clifford data set	102
5.8	Logscale barcode of the Morozov zigzag on the Clifford data set	103
5.9	Projecting the natural images data down using dimensionality reduction	107
5.10	Experimental results obtained using the witness filtration	108
5.11	Logscale barcode of the Morozov zigzag on $Q_{k,x}$, $k = 1, 200$, $x = 30\%$	108
5.12	Logscale barcode of the Morozov zigzag on $Q_{k,x}$, $k = 24, 000$, $x = 30\%$	109
5.13	Parametrization of the space of high-contrast 3×3 patches	109
5.14	Sublevel sets of the distance-like functions d_P , $d_{\mu_P, m}$, and $d_{\mu_P, m}^P$	111
6.1	Point cloud, dendrogram, and hierarchy produced by persistence	116
6.2	ToMATo in a nutshell	119
6.3	Separation of $\text{dgm}_0(f)$ and partial matching with $\text{dgm}_0(\tilde{\mathcal{G}})$	122
6.4	Influence of the partial matching on the regions D_1 and D_2	124
6.5	A function f with unstable basins of attraction	125
6.6	Outputs of the algorithm on a uniform sampling of the domain of f	125
6.7	The twin spirals data set processed using a smaller neighborhood radius	126
6.8	Persistence diagrams obtained on the twin spirals data set	127
6.9	Result of spectral clustering on the twin spirals data set	127
6.10	Biological data set	128
6.11	Quantitative evaluation of the quality of the output of ToMATo	129
6.12	The generalized ToMATo in a nutshell	130
7.1	A database of 3d shapes with six classes	134
7.2	Some toy examples	137
7.3	Bottleneck distance matrix, confusion matrix, and MDS embedding	139
7.4	Overview of the proof of Theorem 7.2	140
7.5	Effect of truncating a filtration as per (7.14) on its persistence diagram	147
7.6	An example where no ball of the offset is ever covered	149
7.7	The additive weight function $t \mapsto s_t(x_k)$ associated to point $x_k \in X$	149
8.1	A cover of the unit circle, and its nerve before and after refinement	161

A.1 Example of a quiver and its underlying undirected graph	168
A.2 The Dynkin diagrams	169
A.3 The Euclidean diagrams	173
A.4 The pyramid travelled down when turning an A_5 -type quiver into L_5	186

Index

- μ -reach, 77
- q (uadrant)-tame, 25, 30, 32
- Artin algebra, 17, 193
- complex
 - α -, 83
 - i -Delaunay, 83
 - (Vietoris-)Rips, 89, 139
 - Čech, 81, 139
 - nerve, 81
 - sparse Rips, 150
 - weighted Rips, 112
 - witness, 90, 139
- correspondence, 143
- Coxeter functors, 179
- decorated
 - multiset, 25
 - number, 21
 - persistence diagram, 22
 - point, 22
- diamond, 179, 182
 - Mayer-Vietoris, 38
 - principle, 178, 183
- dimension
 - doubling, 87
- distance
 - bottleneck, 50
 - distance-like function, 78
 - Gromov-Hausdorff, 134, 143
 - Hausdorff, 69
 - interleaving, 52
 - power, 111
 - to compact set, 69
 - to measure, 110
 - witnessed k -, 112
- duality
 - EP, 36
- elder rule, 120
- elder rule, 6, 31
- endomorphism ring, 171
- filter, 31
- filtered space, 30
- filtration, 29
 - relative, 34
 - right-, 184
 - sublevel-sets, 31
 - superlevel-sets, 31
 - truncated, 147
- interleaving, 52
 - multiplicative, 87
 - sublevel-sets, 52
 - weak, 55
- interpolation lemma, 57
- interval, 20
 - (in)decomposable, 21
 - decomposition, 20
 - module, 20
 - representation, 17, 175
 - rule, 21
- local endomorphism ring, 171
- medial axis, 71
- morphism
 - of degree ε , 53
 - of representations, 15
- Morse-type function, 37
- multivalued map, 142
- natural images, 106
- nerve, 81
 - lemma, 81
 - persistent, 82
- observable category, 28, 63
- offset, 69
- outliers, 110
- partial matching, 50
- path, 192
 - algebra, 17, 192
- persistence
 - barcode, 22, 26
 - diagram, 22, 26
 - (d_1, d_2) -separated, 123

- extended, 34
- hierarchy, 2, 47, 116
- kernels, images, cokernels, 39
- measure, 23
- module, 20
 - streamlined, 183
- zigzag, 36
- persistent
 - (co-)homology, 30, 32, 34
 - (co-)homology group, 30
- poset representation, 19, 194
- pyramid, 37, 185
 - theorem, 39
- quiver, 14, 168
 - A_n -type, 14, 169
 - \tilde{A}_n -type, 173
 - N, Z , 18, 194
 - algebra, 17, 192
 - category, 19, 194
 - Dynkin, 16, 169
 - Euclidean, 173
 - linear L_n , 14, 180
 - poset, 19, 194
 - tame, 17, 174, 191
 - wild, 17, 174, 191
 - with relations, 19, 195
- quiver representation, 14, 168
 - (in-)decomposable, 16, 170
 - category, 15, 168
 - classification, 16, 170
 - dimension, 170
 - dimension vector, 170
 - direct sum, 15
 - interval, 17, 175
 - morphism, 15, 169
 - kernel, image, cokernel, 16
 - shift, 53
 - pointwise finite-dimensional, 18
 - trivial, 15
- reach, 72
- reflection, 175
 - functor, 176
 - functors theorem, 178
- root, 174
- sampling
 - ε -sample, 136
 - furthest-point, 97
 - subsampling, 95
- signatures
 - persistence-based, 138
- snapping
 - lemma, 55
 - principle, 54
- Sparse Voronoi Refinement, 93
- subrepresentation, 16, 168
- symmetry
 - EP, 36
- theorem
 - converse stability, 49
 - Crawley-Boevey, 19, 195
 - Gabriel, 17
 - Gabriel I, 171
 - Gabriel II, 174
 - interval decomposition, 20
 - isometry, 49
 - Kac, 191
 - Krull-Remak-Schmidt, 16, 171
 - pyramid, 39
 - reflection functors, 178
 - stability, 49, 61, 63
 - Webb, 18, 194
- time of appearance, 34, 39
- Tits form, 172
- ToMATo, 118
- topological feature, 30
- totally bounded space, 136
- undecorated
 - persistence diagram, 22
- weak feature size, 75
- zigzag, 36
 - image Rips (iR-ZZ), 98
 - level-sets, 37
 - manipulations, 99
 - module, 20
 - Morozov (M-ZZ), 98
 - oscillating Rips (oR-ZZ), 98

Selected Published Titles in This Series

- 209 **Steve Y. Oudot**, Persistence Theory: From Quiver Representations to Data Analysis, 2015
- 205 **Pavel Etingof, Shlomo Gelaki, Dmitri Nikshych, and Victor Ostrik**, Tensor Categories, 2015
- 204 **Victor M. Buchstaber and Taras E. Panov**, Toric Topology, 2015
- 203 **Donald Yau and Mark W. Johnson**, A Foundation for PROPs, Algebras, and Modules, 2015
- 202 **Shiri Artstein-Avidan, Apostolos Giannopoulos, and Vitali D. Milman**, Asymptotic Geometric Analysis, Part I, 2015
- 201 **Christopher L. Douglas, John Francis, André G. Henriques, and Michael A. Hill, Editors**, Topological Modular Forms, 2014
- 200 **Nikolai Nadirashvili, Vladimir Tkachev, and Serge Vlăduț**, Nonlinear Elliptic Equations and Nonassociative Algebras, 2014
- 199 **Dmitry S. Kaliuzhnyi-Verbovetskyi and Victor Vinnikov**, Foundations of Free Noncommutative Function Theory, 2014
- 198 **Jörg Jahnel**, Brauer Groups, Tamagawa Measures, and Rational Points on Algebraic Varieties, 2014
- 197 **Richard Evan Schwartz**, The Octagonal PETs, 2014
- 196 **Silouanos Brazitikos, Apostolos Giannopoulos, Petros Valettas, and Beatrice-Helen Vritsiou**, Geometry of Isotropic Convex Bodies, 2014
- 195 **Ching-Li Chai, Brian Conrad, and Frans Oort**, Complex Multiplication and Lifting Problems, 2014
- 194 **Samuel Herrmann, Peter Imkeller, Ilya Pavlyukevich, and Dierk Peithmann**, Stochastic Resonance, 2014
- 193 **Robert Rumely**, Capacity Theory with Local Rationality, 2013
- 192 **Messoud Efendiev**, Attractors for Degenerate Parabolic Type Equations, 2013
- 191 **Grégory Berhuy and Frédérique Oggier**, An Introduction to Central Simple Algebras and Their Applications to Wireless Communication, 2013
- 190 **Aleksandr Pukhlikov**, Birationally Rigid Varieties, 2013
- 189 **Alberto Elduque and Mikhail Kochetov**, Gradings on Simple Lie Algebras, 2013
- 188 **David Lannes**, The Water Waves Problem, 2013
- 187 **Nassif Ghoussoub and Amir Moradifam**, Functional Inequalities: New Perspectives and New Applications, 2013
- 186 **Gregory Berkolaiko and Peter Kuchment**, Introduction to Quantum Graphs, 2013
- 185 **Patrick Iglesias-Zemmour**, Diffeology, 2013
- 184 **Frederick W. Gehring and Kari Hag**, The Ubiquitous Quasidisk, 2012
- 183 **Gershon Kresin and Vladimir Maz'ya**, Maximum Principles and Sharp Constants for Solutions of Elliptic and Parabolic Systems, 2012
- 182 **Neil A. Watson**, Introduction to Heat Potential Theory, 2012
- 181 **Graham J. Leuschke and Roger Wiegand**, Cohen-Macaulay Representations, 2012
- 180 **Martin W. Liebeck and Gary M. Seitz**, Unipotent and Nilpotent Classes in Simple Algebraic Groups and Lie Algebras, 2012
- 179 **Stephen D. Smith**, Subgroup Complexes, 2011
- 178 **Helmut Brass and Knut Petras**, Quadrature Theory, 2011
- 177 **Alexei Myasnikov, Vladimir Shpilrain, and Alexander Ushakov**, Non-commutative Cryptography and Complexity of Group-theoretic Problems, 2011
- 176 **Peter E. Kloeden and Martin Rasmussen**, Nonautonomous Dynamical Systems, 2011

For a complete list of titles in this series, visit the
AMS Bookstore at www.ams.org/bookstore/survseries/.

Persistence theory emerged in the early 2000s as a new theory in the area of applied and computational topology. This book provides a broad and modern view of the subject, including its algebraic, topological, and algorithmic aspects. It also elaborates on applications in data analysis. The level of detail of the exposition has been set so as to keep a survey style, while providing sufficient insights into the proofs so the reader can understand the mechanisms at work.



The book is organized into three parts. The first part is dedicated to the foundations of persistence and emphasizes its connection to quiver representation theory. The second part focuses on its connection to applications through a few selected topics. The third part provides perspectives for both the theory and its applications. The book can be used as a text for a course on applied topology or data analysis.



For additional information
and updates on this book, visit

www.ams.org/bookpages/surv-209

AMS *on the Web*
www.ams.org

ISBN 978-1-4704-2545-6



SURV/209