Chapter 1
Introduction

We intend this monograph to be a self-contained presentation of the theory of persistence modules over the real line. We give the best proofs we know of the most important results. Each theorem is located at an appropriate level of abstraction (we believe).

- Newcomers will find this to be an accessible introduction to the algebraic foundations of persistence. They will learn what persistence modules are, how to access their internal structure, their different degrees of 'tameness', how to construct their persistence diagrams mathematically, and how to show that those diagrams are stable.
- Experts will find that our methods add precision and power to what they already know. We construct persistence modules in great generality and show that they can be manipulated just as easily as the traditional examples. We give evidence that this greater generality occurs naturally 'in the wild'.

Persistence modules are the mathematical object at the heart of the young, rapidly growing field of topological data analysis. This field—a blend of computer science, algebraic topology and statistics—is founded on the assumption that scientific data sets carry information in their internal structure and that sometimes this internal structure is topological. Persistence modules were designed to carry topological information about a data set at many different scales simultaneously. This information can be extracted in the form of an invariant—the *persistence diagram* or *barcode*—that can be computed effectively and is statistically robust.

New researchers in the field have to come to terms with multiple aspects of persistence. They must learn how data can be turned into geometric objects (usually a nested family of simplicial complexes). They must learn to work with the algorithms that turn these geometric objects into barcodes or persistence diagrams. They need algebraic topology to know how to interpret these barcodes. And, increasingly often nowadays, they must know enough statistical theory to draw valid inferences from the results of these calculations. There is by now a great deal of literature covering

F. Chazal et al., *The Structure and Stability of Persistence Modules*,
SpringerBriefs in Mathematics, DOI 10.1007/978-3-319-42545-0_1

these different facets, written by and intended for researchers across a large range of disciplines.

Our focus is narrower. In this monograph, we are concerned almost exclusively with the mathematical properties of persistence modules. We have several reasons for doing this:

- The existing literature is largely built around particular instances of topological persistence, such as the sublevelset persistent homology 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. We wish to correct this bias.
- Several components of the theory are algebraic in nature. The fact that most of the usual examples come from topology sometimes creates a different impression. Here we give a purely algebraic presentation of the main ingredients: the persistence diagram, tameness, stability. Topological arguments are used when studying the properties of persistence modules obtained from topological data, but are not needed otherwise.
- Variations such as image and kernel persistence can be treated equally.

More specifically, we are interested in persistence modules indexed by a single real parameter. Much of our work goes to show that a continuous parameter can be handled just as effectively as a discrete parameter; this fills a gap in the literature. On the other hand, we make no attempt to discuss multiparameter persistence. It is a complicated subject and at some point will be ready for its own book.

Within its scope, the monograph is almost entirely self-contained. We do appeal to two off-the-shelf technical results from module theory (Theorems 2.7 and 2.8). When discussing persistent *homology*, we assume that the reader is familiar with the standard properties of simplicial complexes and homology functors. Since it is important to remain connected to the larger world of topological data analysis, we briefly present two applications in Sects. 1.3 and 1.4, and a themed reading list in Sect. 1.5.

1.1 A Brief History of Persistence

The early history of persistence is concerned with the quantity

$$\mathrm{r}_t^s = \mathrm{rank}(\mathrm{H}(X_s) \to \mathrm{H}(X_t))$$

for an object X represented at two different scales s, t, and where H is homology. This appeared in the early 1990s in the work of Frosini [34], with different notation and under the name ‘size function’. Independently, a few years later, Robins [46] introduced the term ‘persistent Betti numbers’ for quantities of the form $\mathrm{r}_{\varepsilon+\rho}^{\varepsilon}$, and noted their stability with respect to Hausdorff distance.

The modern theory of persistence is built on three pillars:

- The persistence diagram, and an algorithm for computing it, were introduced by Edelsbrunner, Letscher and Zomorodian [30]. This gives a compact representation of the size function and an effective way to compute it.
- Zomorodian and Carlsson [50] defined persistence modules in the abstract, indexed by the natural numbers and viewed as graded modules over the polynomial ring $\mathbf{k}[t]$. This introduced tools from commutative algebra.
- Cohen-Steiner, Edelsbrunner and Harer [19] formulated and proved the stability theorem, which guarantees that the persistence diagram is robust to changes in the input data. Robustness is measured in terms of a 'bottleneck distance' between persistence diagrams.

All three papers make the assumption that the data is essentially finite. This is understandable from the perspective of computer science: a finite machine can only handle a finite amount of data. And mathematically it is the natural place to begin. In the realm of continuous topology it is common to make finiteness assumptions: a continuous function on a smooth manifold may be approximated by a Morse function, and on a polyhedron may be approximated by a piecewise-linear map.

The finiteness restrictions were lifted in [15] (and its published conference version [14]), which generalises the main results to persistence modules indexed over the real line, under the relatively mild assumption that $\mathrm{r}_t^s < \infty$ for $s < t$. In the present work, we call these modules 'q-tame'. It turns out that this is a natural condition; large classes of examples are q-tame. Moreover, the formulation of the stability theorem in [15] is purely algebraic, and not tied to any particular geometric situation. The only drawback is that some of the arguments are rather complicated.

In this monograph, we carry out the program of [15] with new arguments. The proofs are now very clean and the methods are versatile. Let us say a few words about why we think it worth the effort to work with continuous-parameter persistence modules. Here are our two main reasons:

- Real-world data sets are always finite, but they may be statistical samples from an underlying continuous object or process. Ideally the persistent homology of a sample will be an approximation to the persistent homology of the continuous model. Formulating this requires a theory of continuous-parameter persistence.
- Continuous-parameter persistence extends the applicability of the theory from finite data sets to more general objects such as compact metric spaces. This widens the applicability of persistence within pure mathematics.

In support of this last point, we observe that in recent years the discrete form of persistence has seen application in various branches of pure mathematics. For example, Ellis and King [33] use persistence to study p-groups; and Pakianathan and Winfree [45] have reformulated a number of famous problems in number theory, including the Riemann Hypothesis, in terms of the persistent homology of certain filtered simplicial complexes.

We draw attention to three recent papers which share our goal of understanding continuous-parameter persistence modules:

- Lesnick [42] gives an extensive algebraic treatment of modules over one or more real parameters. The converse stability inequality, and hence the isometry theorem, appears for the first time in his work. Our present work was carried out largely independently, with one salient exception: it was from Lesnick [42] we learned of results of Webb [48] that resolved a sticking-point for us.
- Bubenik and Scott [5] develop the category-theoretical view of persistence modules. This allows them to formulate and prove stability theorems in great generality. Categories and functors surface occasionally in the present monograph, and can be used to streamline some of the work very effectively and non-trivially.
- Bauer and Lesnick [3] give a completely original proof of the stability theorem for q-tame continuous-parameter persistence modules, making no use of the interpolation lemma that is so crucial to our approach (which is based on the original proof of Cohen-Steiner, Edelsbrunner and Harer [19]). Their work depends on strong results on the decomposition of persistence modules.

All of these works appeared during the writing of this monograph.

1.2 Main Contributions

Many authors have studied persistence modules in recent years, and many of the theorems presented here are not original in themselves. The originality lies in the methods that we use. Our main innovations are these:

- We construct persistence diagrams using measure theory. The existence of a diagram is equivalent to the existence of a certain kind of measure on rectangles in the plane.
- We introduce 'decorated' real numbers for two related purposes: to remove the ambiguity about the endpoints of persistence intervals, and to get the measure theory to work.
- We introduce a special notation for calculations on quiver representations. This considerably simplifies the linear algebra (for instance, in proving the 'box lemma').
- We define several kinds of 'tameness' for a persistence module. These occur naturally in practice. The most restrictive of these, finite type, is what is normally seen in the literature. We show how to work effectively with the less restrictive hypotheses.
- We give a clean proof of the algebraic interpolation theorem of [15].
- We rewrite the algebraic stability theorem of [15] as a theorem about measures. Among other consequences, this leads to diagram stability results for even quite badly behaved persistence modules.

Our goal in introducing these ideas is to enable our readers to define persistence diagrams cleanly and effectively in a wide variety of situations. In the earlier work [15], continuous-parameter persistence diagrams are constructed using a careful limiting process through ever-finer discretisations of the parameter. Unfortunately the limiting arguments turn out to be quite complicated, and the resulting diagrams are difficult to work with. Our approach in the present monograph gives the best of both worlds; we are able to work with broader classes of persistence modules, and we can reason about their diagrams in a clean way using arguments of a finite nature.

1.3 Application: Stable Descriptors for Metric Spaces

In this section and the next, we illustrate the usefulness of the persistence diagram for problems in data analysis and machine learning. The key point is that the persistence diagram is a robust invariant of the underlying geometric situation.

It often happens in classification tasks that the objects to be classified cannot be compared against one another directly but only through descriptors or 'signatures'. Consider the problem of organizing a database of 3-dimensional objects into meaningful classes, as illustrated in Fig. 1.1. Each object—or '3d shape'—is represented as some part of the bounding surface of a 3-dimensional region, and it can be abstracted mathematically as a compact metric space $(P, \mathrm{d_P})$ where the set P is the surface itself and the metric is the geodesic distance along the surface. Comparing two 3d shapes amounts to comparing the corresponding metric spaces.

The natural distance on the space of compact metric spaces is the so-called *Gromov–Hausdorff distance* $\mathrm{d_{GH}}$, a generalisation of the classical Hausdorff distance defined as follows:

$$\mathrm{d_{GH}}(P, Q) = \inf_{S,f,g} \mathrm{d_H}(f(P), g(Q))$$

where $(P, \mathrm{d_P})$ and $(Q, \mathrm{d_Q})$ are the two given compact metric spaces, where $(S, \mathrm{d_S})$ ranges over all metric spaces, where $f : P \to S$ and $g : Q \to S$ range over all isometric embeddings of P and Q into S, and where $\mathrm{d_H}$ denotes the Hausdorff distance in $(S, \mathrm{d_S})$.

The issue with this distance is that its direct computation leads to a quadratic assignment problem that is hard to solve in practice. Upper bounds can be obtained easily, but lower bounds are more difficult. A workaround is to map the metric spaces P, Q to some space of signatures in which distances are easier to compute. Ideally, one would like the distance between signatures to be the same as the Gromov–Hausdorff distance between the initial metric spaces, but this is usually too much to ask. We can at least require that the distance between signatures provide a lower bound on the distance between the metric spaces, so in particular the signatures are provably stable under small perturbations of the spaces.

The general stability theorem stated in the present monograph (Theorem 5.25) makes it possible to derive such a stability guarantee, when the signature of a compact

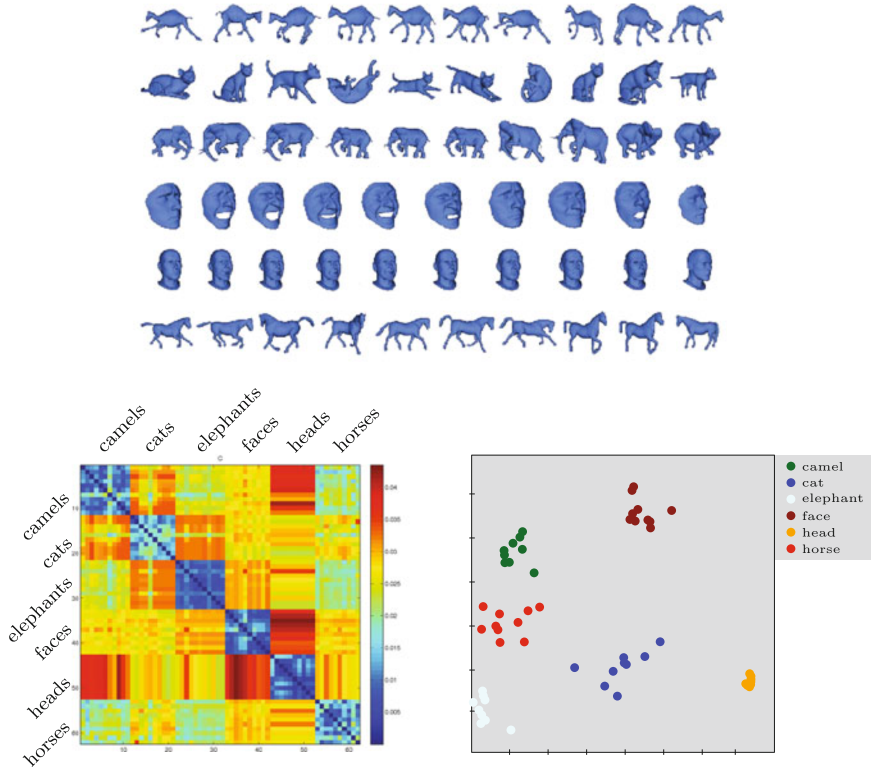

Fig. 1.1 (From [13]) An unsupervised classification task using persistence-based signatures. *Top:* The collection of 60 shapes to be classified into 6 classes (labels unknown). *Bottom-left:* The distance matrix in signature space, with color-coded values (each row and column corresponds to a single shape in the collection). *Bottom-right:* The signatures are embedded into the Euclidean plane using multidimensional scaling. The objects are then classified by a simple k-means clustering procedure applied to this embedding. Label names can be extrapolated to classes if some of the individual objects have known labels

metric space is taken to be the persistence diagram of the homology of its Vietoris–Rips complexes[1]:

Theorem 1.1 ([13, 16]) *For any compact metric spaces* (P, d_P) *and* (Q, d_Q),

$$\mathrm{d_b}(\mathsf{dgm}(\mathrm{H}_*(\mathbb{R}\mathrm{ips}(P))),\ \mathsf{dgm}(\mathrm{H}_*(\mathbb{R}\mathrm{ips}(Q)))) \leq \mathrm{d_{GH}}(P, Q).$$

[1]We do not give the details of this construction here; see [13, 16] for instance. What matters is that the signature $\mathsf{dgm}(\mathrm{H}_*(\mathbb{R}\mathrm{ips}(P)))$ is easily computed, the distance between two signatures is easily computed, the distance is robust in the sense of Theorem 1.1, and that the lower bound in the theorem is sufficiently tight to solve the learning problem under consideration.

As desired, this result provides a lower bound on the Gromov–Hausdorff distance between the metric spaces in terms of the bottleneck distance d_b between their signatures.

It turns out that signatures of this type are rich enough to be used effectively in machine learning applications such as the one depicted in Fig. 1.1. In such applications, the continuous shapes themselves are replaced by finite samples for practical purposes. One can exploit Theorem 1.1 to prove minimax-optimal upper bounds on the convergence rate of the sample signatures to the signatures of the underlying continuous objects [17]. It is also possible to define local versions of the signatures, with similar stability guarantees, for use in partial comparison and matching applications [11].

1.4 Application: Stable Clustering Using Persistence

Unsupervised learning or clustering is an important tool for understanding and interpreting data. Among the wealth of existing approaches, mode seeking is the one most closely related to Morse theory and persistence. The approach assumes that the data points have been drawn from some unknown probability distribution with density f. The idea is to detect the local peaks of f and use them as cluster centers, grouping each data point with the local peak that it eventually reaches by following the gradient vector field of f uphill—assuming f has a well-behaved gradient flow.

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 (Fig. 1.2, top-right). One possible workaround is to smoothe the estimator before launching the hill-climbing procedure; this raises the difficult question of how much smoothing is needed to remove the noise without affecting the signal and still obtain the correct number of clusters. A different workaround, proposed in [18], is to do the hill-climbing anyway and afterwards detect and merge the unstable clusters to regain some stability. This is where persistence comes into play.

The persistence diagram of the density estimator $\hat{f}$ provides a measure of prominence for the peaks of $\hat{f}$, through the distances of their corresponding diagram points to the diagonal. Whenever the diagram satisfies a 'sufficient separation' condition, it is easy to pick a prominence threshold that will separate the relevant peaks of $\hat{f}$ from the irrelevant ones. The clusters associated to the irrelevant peaks can then be merged into the clusters of the relevant peaks, using the hierarchy built by the persistence algorithm. In this way one obtains the correct number of clusters: see Fig. 1.2.

Theorem 1.2 ([18]) *Let c be the Lipschitz constant of the true density f, let η be the approximation error (in the supremum norm) of the estimator $\hat{f}$, and let δ be the neighborhood size parameter used in the hill-climbing procedure. Assume that the peaks of the true density f have prominences at least d. Then, for any prominence threshold within the range $(2(c\delta + \eta),\ \ d - 3(c\delta + \eta))$, the number of clusters*

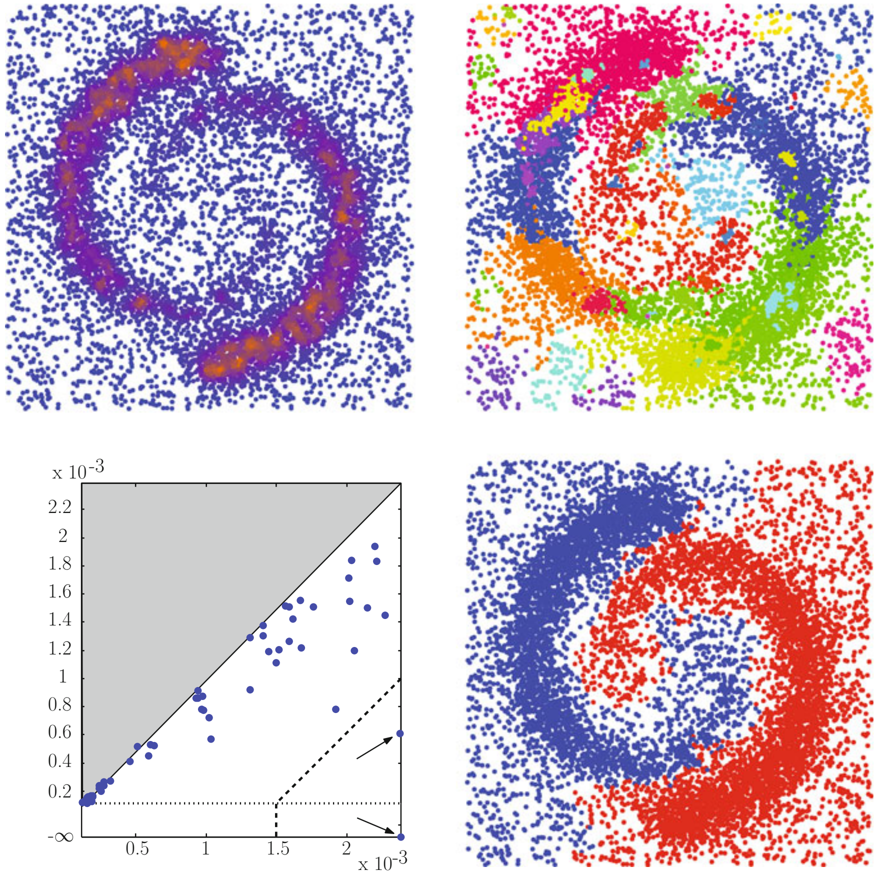

Fig. 1.2 (From [18]) Persistence-based clustering in a nutshell. *Top-left:* The underlying density function f is estimated at the data points. *Top-right:* The hill-climbing procedure applied to the estimated density $\hat{f}$ leads to a large number of unreliable clusters. *Bottom-left:* The persistence diagram of $\hat{f}$ shows 2 points far off the diagonal corresponding to the 2 peaks of the true density f. *Bottom-right:* The final result is obtained by merging the clusters of the other peaks of $\hat{f}$

computed by the above procedure, on an input of n random sample points drawn i.i.d. according to f, is equal to the number of peaks of f with probability at least $1 - e^{-\Omega((c\delta+\eta)n)}$, where the constant in the big-Ω notation depends only on geometric quantities (e.g. volumes of balls) associated with the ambient space.

The proof of this result relies on a partial notion of interleaving and a version of the stability theorem for such partial interleavings (Theorem 6.1). Both follow easily from the framework developed in this monograph.

1.5 Recommended Reading

There is by now a substantial literature on topological persistence, launched by the work of Edelsbrunner, Letscher and Zomorodian [30], with antecedents in papers of Frosini [34] and Robins [46]. Beyond these historical documents, we have some themed suggested readings for the reader seeking orientation in the larger field of topological data analysis.

Applied Algebraic Topology

A substantial grounding in the broad field of applied algebraic topology may be found in the following books, each reflecting the particular tastes of its author(s).

- Tomasz Kaczynski, Konstantin Mischaikow, and Marian Mrozek. *Computational Homology*, volume 157 of *Applied Mathematical Sciences*. Springer, 2004.
- Afra Zomorodian. *Topology for Computing*, volume 16 of *Cambridge Monographs on Applied and Computational Mathematics*. Cambridge University Press, 2005.
- Herbert Edelsbrunner and John L. Harer. *Computational topology: an introduction*. American Mathematical Society, 2010.
- Robert Ghrist. *Elementary Applied Topology*. CreateSpace Independent Publishing Platform, September 2014.
- Steve Y. Oudot. *Persistence Theory: from quiver representations to data analysis*, volume 209 of *Mathematical Surveys and Monographs*. American Mathematical Society, 2015.

Topological Data Analysis

The following survey articles provide different perspectives on the application of topological persistence to data analysis. These articles provide a good introduction for readers new to the field. In particular, they explain how persistence diagrams (or, equivalently, barcodes) are used in practice.

- Robert Ghrist. Barcodes: the persistent topology of data. *Bulletin of the American Mathematical Society*, 45(1):61–75, 2008.
- Gunnar Carlsson. Topology and data. *Bulletin of the American Mathematical Society*, 46(2):255–308, 2009.
- Herbert Edelsbrunner and Dmitriy Morozov. Persistent homology: theory and practice. In *European Congress of Mathematics*, pages 31–50. European Mathematical Society, 2012.

The Persistence Algorithm

The following articles deal specifically with the details of the persistence algorithm, an aspect that is not covered in the present book.

- Herbert Edelsbrunner, David Letscher, and Afra Zomorodian. Topological persistence and simplification. *Discrete & Computational Geometry*, 28:511–533, 2002.

- Afra Zomorodian and Gunnar Carlsson. Computing persistent homology. *Discrete & Computational Geometry*, 33(2):249–274, 2005.
- Vin de Silva, Dmitriy Morozov, and Mikael Vejdemo-Johansson. Dualities in persistent (co)homology. *Inverse Problems*, 27:124003, 2011.

Stability Theorems

The following articles present various versions of the proof of stability for persistence diagrams. We are omitting [15], which served as a basis for the present work.

- David Cohen-Steiner, Herbert Edelsbrunner, and John Harer. Stability of persistence diagrams. *Discrete & Computational Geometry*, 37(1):103–120, 2007.
- David Cohen-Steiner, Herbert Edelsbrunner, John Harer, and Yuriy Mileyko. Lipschitz functions have Lp-stable persistence. *Foundations of Computational Mathematics*, 10(2):127–139, 2010.
- Ulrich Bauer and Michael Lesnick. Induced matchings of barcodes and the algebraic stability of persistence. In *Proceedings of the 30th Annual Symposium on Computational Geometry (SoCG)*, pages 355–364, Kyoto, Japan, June 2014.

Geometric Complexes on Compact Metric Spaces

The following paper uses the tools introduced in the present book to derive stable topological signatures for compact metric spaces.

- Frédéric Chazal, Vin de Silva, and Steve Oudot. Persistence stability of geometric complexes. *Geometriae Dedicata*, 173:193–214, 2014.

It follows previous work in the non-persistent context:

- Jean-Claude Hausmann. On the Vietoris-Rips complexes and a cohomology theory for metric spaces. In *Prospects in Topology*, volume 138 of *Annals of Mathematical Studies*, pages 175–188. Princeton University Press, Princeton, New Jersey, 1995.
- Janko Latschev. Vietoris-Rips complexes of metric spaces near a closed Riemannian manifold. *Archiv der Mathematik*, 77:522–528, 2001.

Variations on Persistence

The following articles introduce several variations of 1-dimensional persistence: vineyards, extended persistence, image persistence, zigzag persistence. The tools introduced in the present book can be used to simplify some parts of their analysis.

- David Cohen-Steiner, Herbert Edelsbrunner, and Dmitriy Morozov. Vines and vineyards by updating persistence in linear time. In *Proceedings of the 22nd Annual Symposium on Computational Geometry (SoCG)*, pages 119–126, 2006.
- David Cohen-Steiner, Herbert Edelsbrunner, and John Harer. Extending persistence using Poincaré and Lefschetz duality. *Foundations of Computational Mathematics*, 9:79–103, 2008.

- David Cohen-Steiner, Herbert Edelsbrunner, John Harer, and Dmitriy Morozov. Persistent homology for kernels, images, and cokernels. In *Proceedings of the Twentieth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1011–1020, 2009.
- Gunnar Carlsson and Vin de Silva. Zigzag persistence. *Foundations of Computational Mathematics*, 10(4):367–405, 2010.

Multidimensional Persistence

The following articles extend the theory to persistence modules indexed over multidimensional index sets. There the situation is much less clear than for 1-dimensional index sets.

- Gunnar Carlsson and Afra Zomorodian. The theory of multidimensional persistence. In *Proceedings of the 23rd Annual Symposium on Computational Geometry (SoCG)*, Gyeongju, South Korea, June 2007.
- Gunnar Carlsson, Gurjeet Singh, and Afra Zomorodian. Computing multidimensional persistence. *Journal of Computational Geometry*, 1(1):72–100, 2010.
- Michael Lesnick. The theory of the interleaving distance on multidimensional persistence modules. *Foundations of Computational Mathematics*, 15(3):613–650, 2015.

Categorified Persistence

Finally, the following articles build more abstract versions of the theory using the language of category theory. They complement the ideas developed in the present book, providing an abstract framework that some readers may find congenial.

- Peter Bubenik and Jonathan A. Scott. Categorification of persistent homology. *Discrete & Computational Geometry*, 51:600–627, 2013.
- Peter Bubenik, Vin de Silva, and Jonathan Scott. Metrics for generalized persistence modules. *Foundations of Computational Mathematics*, 15(6):1501–1531, 2015.

1.6 Organisation

The remaining chapters of the monograph are organised as follows.

Chapter 2 introduces the basic properties of persistence modules. These can be defined over any partially ordered set; we are primarily interested in persistence modules over the real line. In the best case a persistence module can be expressed as a direct sum of ‘interval modules’, which can be thought of as the atomic building blocks of the theory. Decorated real numbers are introduced here. Not all persistence modules decompose into interval modules, so we spend much of the monograph developing techniques that work without this assumption. These techniques depend on a thorough understanding of finitely-indexed persistence modules known as 'A_n-quiver representations' [35, 27]. We introduce a special notation for performing calculations on these quiver representations. This 'quiver calculus' is used throughout.[2]

Chapter 3 addresses the question of how to define the diagram of a persistence module. This is easy for modules which decompose into intervals. To handle the general case, we establish an equivalence between diagrams and a certain kind of measure defined on rectangles in the plane. Whenever a persistence diagram is sought, therefore, it suffices to construct the corresponding persistence measure. Theorems about a diagram can be replaced by simpler-to-prove theorems about its measure. The diagram exists wherever the measure takes finite values. This leads to several different notions of 'tameness'. There are large classes of examples of naturally occurring persistence modules which are tame enough for their diagrams to be defined everywhere or almost everywhere. Some elementary 'vanishing lemmas' facilitate the explicit calculation of persistence modules. We finish by showing how our abstractly defined diagrams agree with the diagrams that are produced by the standard algorithms [30, 50] when working in finite situations derived from real data.

Chapter 4 is concerned with interleavings. An interleaving is an approximate isomorphism between two persistence modules. They occur naturally in applications when the input data are known only up to some bounded error. After presenting the basic properties, we give a clean proof of the technical lemma (from [15]) that two interleaved modules can be interpolated by a 1-parameter family.

Chapter 5 is devoted to the isometry theorem, which asserts that the interleaving distance between two persistence modules is equal to the bottleneck distance between their persistence diagrams. The two inequalities that comprise this result are treated separately. One direction is the celebrated stability theorem of [19]. The more recent converse inequality appears in [42]. We formulate the stability theorem as a statement about measures and their diagrams. The proof of this more abstract result closely follows the original proof in [19]. Our version of the isometry theorem supposes that the persistence modules are q-tame. We also provide a more general version of the stability theorem which allows us to compare diagrams of persistence modules with no assumptions on their tameness: wherever the two diagrams are defined, they must be close to each other.

Chapter 6, finally, contains two worked examples. We show how the theory developed in this monograph can be used in practice to define various forms of persistence and prove the needed theorems and lemmas. We hope these examples illustrate the strength and flexibility of our approach.

[2] Readers who wish to adopt our notation are invited to contact us for the LaTeX macros.

1.7 Multisets

Persistence diagrams are multisets rather than sets. For our purposes, a multiset is a pair $\mathsf{A} = (S, \mathrm{m})$ where S is a set and

$$\mathrm{m} : S \to \{1, 2, 3, \dots\} \cup \{\infty\}$$

is the multiplicity function, which tells us how many times each element of S occurs in A. Here are our conventions regarding multisets:

- The cardinality of $\mathsf{A} = (S, \mathrm{m})$ is defined to be

  $$\operatorname{card} \mathsf{A} = \sum_{s \in S} \mathrm{m}(s)$$

  which takes values in $\{0, 1, 2, \dots\} \cup \{\infty\}$. We do not distinguish between different infinite cardinals.
- We never form the intersection of two multisets, but we will sometimes restrict a multiset A to a set B:

  $$\mathsf{A}|_B = (S \cap B, \mathrm{m}|_{S \cap B})$$

  We may write this as $\mathsf{A} \cap B$ when $\mathsf{A}|_B$ is typographically inconvenient.
- A pair (B, m) where

  $$\mathrm{m} : B \to \{0, 1, 2, \dots\} \cup \{\infty\}$$

  is implicitly regarded as defining a multiset $\mathsf{A} = (S, \mathrm{m}|_S)$ where $S = B - \mathrm{m}^{-1}(0)$ is the support of m.
- If $\mathsf{A} = (S, \mathrm{m})$ is a multiset and $f : S \to B$ where B is a set, then the notation

  $$\{f(a) \mid a \in \mathsf{A}\}$$

  is interpreted as the multiset in B with multiplicity function

  $$\mathrm{m}'(b) = \sum_{s \in f^{-1}(b)} \mathrm{m}(s)$$

Except in definitions like these, we seldom refer explicitly to S.