FP7-ICT-2011.8
Contract no.: 318493

www.toposys.org

TOPOSYS

Deliverable D3.1

Progress and Activity Report on WP3

Deliverable Nature:	Report (R)
Dissemination Level: (Confidentiality)	Public (PU)
Contractual Delivery Date:	M12
Actual Delivery Date:	M12
Version:	1.0

Contents

1 Summary

Category theory is a powerful language for describing mathematics. By phrasing existing mathematical techniques in a categorical setting, focus is placed on the underlying abstract structures, and often a common abstraction suggests itself from placing theories on a categorical footing. In particular, categorical formulations often highlight and even formalize analogies within a field. Among the powerful current applications, we recognize areas such as Quantum Field Theory, and the application of homotopy theory methods to algebra, combinatorics and other areas of mathematics. In particular, the perspective of functoriality was highlighted by Carlsson in Topology and Data as a core point of progress in the applications of persistent homology, and this theme has been further explored in work on functorial clustering as well as in the field of algebraic statistics. The goal of this work package is to explore categories as a fundamental abstraction tool for topological and algebraic approaches to systems. There relate to both categories as direct tools for dynamical systems and to category theory as an abstraction layer for persistent homology:

The goals we concentrate on are

1. Placing persistence into a categorical setting of topoi
2. Defining statistical/machine learning sheaves
3. Categorical dynamical systems

The ultimate goal of this work package is to gain a more fundamental understanding of persistence in its numerous forms, which is a necessity if it is to serve as a foundation for complex systems. The highlights of our progress in this year are

1. Taking important steps developing a topos of sheaves, including defining an underlying Heyting algebra as well as work out initial examples of sheaves on this structure.
2. Developing algorithms based on sheaf-theoretic ideas as well as further as well as developing local structures, which could be considered in a sheaf setting.
3. Two more general notions of persistence (other than the topos foundation), which are related but complementary.

That is why we first present a current state of the foundations of persistence based on [33], followed by a description of the of the Heyting Algebra upon which we will base our further topos. This is currently in the stage of working notes and will the section highlights the construction. We expect a preprint to be finished by the end of the calendar year. We then describe one construction which is sheaf-like and one which could be used with sheaf theory in the context of stratified manifold learning. Finally we describe two generalized views on persistence which have algorithmic and conceptual consequences.

That is why we first present a current state of the foundations of persistence based on "Sketches fo a platypus: persistent homology and its algebraic foundations", (available at `http://arxiv.org/abs/1212.5398`). This is followed by a description of the of the Heyting Algebra upon which we will base our further topos on. This is currently in the stage of working notes and will the section highlights the construction. We expect a preprint to be finished by the end of the calender year.

We then describe one construction which is sheaf-like and one which could be used with sheaf theory in the context of stratified manifold learning. The two papers upon which is based on are

- P. Skraba, M. Vejdemo-Johansson, "Parallel & scalable zig-zag persistent homology," NIPS 2012, Workshop on Algebraic Topology and Machine Learning - (available at `http://www.cs.cmu.edu/~sbalakri/Topology_final_versions/zigzag.pdf`). Note that the paper does not make explicit reference to sheaves, but rather is algorithmic.
- P. Skraba, B. Wang, "Approximating Local Homology from Samples," ACM-SIAM Symposium on Discrete Algorithms (SODA) (accepted), 2014. (please see the attached version - there is an ArXiV version which is out-of-date and which contains a bug)

Finally, we describe two generalized views on persistence which have algorithmic and conceptual consequences. The exposition of towers of persistence is based on the preprint

- H. Edelsbrunner, G. Jablonski and M. Mrozek. "The persistent homology of a self-map." submitted to Found. Comput. Math. (see attached version)

A more complete version of a discussion of k[t]-algebras can be found at

- P. Skraba, M. Vejdemo-Johansson, "Persistence modules: Algebra and algorithms" to be submitted, (available at `http://arxiv.org/abs/1302.2015`)

Note that other than the unpublished notes, all the papers are attached to this report (including those which are not publicly available).

2 Background

Before describing the initial stages of the construction of a topos for persistence we describe the two main genres of foundations of persistence.

Filtered spaces Persistent homology is about the effect of applying the homology functor to a filtration of topological spaces. Invariants describing the resulting homology diagrams help us construct tools for visualization and data analysis eventually allowing for the inference of topological structure for point clouds using specific constructions of filtered complexes that encode properties of point clouds.

Representations of the reals Persistent homology is about studying sublevel sets of real-valued functions on topological spaces. Such sublevel sets have – for nice enough functions and spaces – discretizations that allow us to adapt descriptions of finite diagrams of vector spaces to efficient descriptors. In particular, by using the "distance from a set" family of functions we can support inference of topological structures from point clouds.

Both these choices come with built in benefits as well as drawbacks. They give rise to different generalizations of the fundamental inference problem for point clouds sampled from a topological space, and they support different further constructions and proofs.

In particular, among the results that emerge from the two viewpoints, we will be discussthe following:

Stability The **representations of the reals** viewpoint allows us to prove a Lipschitz-style property for the inference process underlying the theory: there is a metric, the *bottleneck metric*, on the invariants of the diagrams of homology groups such that the distance between the homologies of the sublevel sets of two different functions is bounded by the L_∞-distance between the functions. Evolutions in the exact definitions used for persistence lead to increasingly generous assumptions in this bound.

Sub- and super- and iso-level sets By modifying the constructions used, we can get new constructions that allow us to study sequences of super-level sets, of iso-level sets (or level-sets), and of the result from collapsing sub- or super-level sets to a single point. In particular, this brings us *extended persistence*, where no infinite length intervals occur, and a number of topological features comes into play, including Poincaré duality. Current technologies for iso-level sets tend to rely on zig-zag persistent homology (see below).

Graded modules The kinds of diagrams emerging from the **filtered spaces** viewpoint have the structure of graded modules over the polynomial ring k[t]. This recognition sparked both new algorithms for computing persistent homology with far less assumptions on the chosen coefficient ring, and a number of extensions of the fundamental constructions that we will mention below.

Relax the *filtration* requirement In a seminal paper, [90] proved that the *tameness* of the representation theory of quiver algebras depends only on the corresponding Dynkin diagram, not on the particular orientation of arrows in the quiver. Re-interpreting the diagrams of vector spaces emerging from the **filtered spaces** viewpoint as modules over quiver algebras rather than modules over $\mathrm{k}[t]$ allows for inclusion maps that go both forwards and backwards producing *zig-zag persistent homology*, which has allowed for both a topological approach to statistical bootstrapping and concrete approaches to iso-level set persistent homology.

More directions The work by [39] on the topology of configurations of pixels in natural images relied on being able to vary several independent variables in the construction of the intermediate simplicial complexes studied. This inspired [14] to study how these *multi-dimensional* approaches can be handled. A straight generalization from graded $\mathrm{k}[t]$-modules directs us to study modules over $\mathrm{k}[t_1, t_2, \ldots, t_n]$, which brings a whole range of theoretical and computational problems. Nevertheless, recent research seems promising. [100, 70]

2.1 Persistence barcodes and diagrams

Throughout there is an underlying ideal of what persistent homology should be computing, which the field as a whole agrees on: given a filtered (and parametrized) sequence of topological spaces $\mathbb{X}_*$, the persistent homology $H_j^{a\to b}(\mathbb{X}_*)$ is the image of the induced map $H_j(\mathbb{X}_a) \to H_j(\mathbb{X}_b)$. In nice enough cases the collection of all such homologies has a nice algebraic description as some sort of collection of intervals, and these intervals with their start- and end-parameters can be used to produce diagrams that allow reasoning about the original spaces.

There are two main such diagrams in use – both can be seen in Figure 2. One view is the *persistence barcode* – the sequence of intervals is drawn, stacked on top of another. Such a barcode can be seen in the middle of Figure 2. The rank of any particular $H_j^{a\to b}(\mathbb{X}_*)$ is the number of intervals in the barcode that entirely covers the interval (a, b).

The other diagram in use is the *persistence diagram*: the start- and end-points of an interval in the interval decomposition of the persistent homology are taken to be x- and y-coordinates of points in the upper half of the first quadrant of the plane. An example can be seen to the right of Figure 2. The number of points contained in the quadrant delimited by the horizontal line at height a and the vertical line at width b determines the rank of $H_j^{a,b}(\mathbb{X}_*)$.

Either of these cases is a visualization of the underlying data of a *barcode*, which we can define as [76] as a multiset in $\bar{\mathbb{R}}^2$. The barcode is usually taken to include the uncountably many points along the diagonal of $\bar{\mathbb{R}}^2$ as part of the barcode.

There are numerous metrics available to compare diagrams, many of which have provable stability rsults, however, we do not cover them here.

2.2 Functions on a manifold

The study of persistent homology originates from [18], who first define the term and provide an algorithm for the computation of persistent homology. Taking their inspiration from α-shapes, the authors assume that a filtered simplicial complex is provided as input, and produce a description of its persistent homology. In a slightly later paper, [86] demonstrate that persistent homology can be applied to morse complexes from piecewise linear functions on a manifold – the filtered simplicial complex required is given by combining the morse complex cells with the function values at the critical points witnessing each cell.

From this point and onwards, one strongly present culture in persistent homology remains focused on the role of a function defined on a manifold as the input data for the method. This viewpoint has proven remarkably fruitful in the study of *stability*, and provides the best tools we currently have for justifying topological inferences with persistent homology.

It is worth noting that a point cloud topology point of view fits in this framework: as is illustrated in Figure 1, the distance to a discrete set of points produces a real-valued function on the ambient space of the

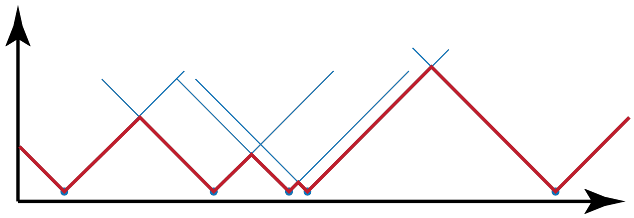

Figure 1: The distance to a set of points, defined for any point as the infimum of individual distances to points in the point cloud, produces a function for use in the functional approach to persistent homology. The points at the bottom of the valleys in the graph are the points of a 1-dimensional point cloud; and the lightly drawn cones emanating from each point correspond to the distance function from that point itself. The lower envelope of these distances forms the distance to the entire set, thus the function for encoding Čech complex persistent homology as a functional persistent homology.

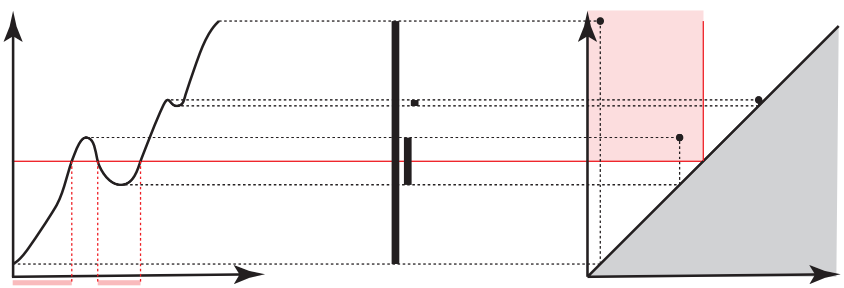

Figure 2: Persistence of H_0 of sublevel sets of a function $\mathbb{R} \to \mathbb{R}$. In black, we see the three components that appear at different times show up – in the middle in a persistence barcode, and to the right as the three points in a persistence diagram. In red, we indicate a particular choice of height ε, at which the sublevel set has two components – drawn below the graph to the left. These two components can be read off in both persistence visualizations – through the two intersected bars in the middle, and through the two points contained in the shaded red region to the right.

points, with a persistent homology corresponding closely to the Čech complex homology of the point cloud itself.

2.2.1 A functional view of persistent homology

With this viewpoint, the fundamental given datum is a geometric object $\mathbb{X}$ and a tame function $f : \mathbb{X} \to \mathbb{R}$. In order to study the behavior of sublevel sets of f, persistent homology is used to measure the filtration of $\mathbb{X}$ given by $\mathbb{X}_\varepsilon = f^{-1}((-\infty, \varepsilon])$.

A function $f : \mathbb{X} \to \mathbb{R}$ is called *tame* if it is continuous, all sublevel sets have homology groups of finite rank, and there are finitely many critical values where the homology groups change.

This viewpoint, and the reasons for some of the choices made in creating algorithms are at their most apparent when considering the 1-dimensional case, where $\mathbb{X} = \mathbb{R}$, and we consider sublevel sets of some function $\mathbb{R} \to \mathbb{R}$.

Consider Figure 2. Critical points of the function correspond to points where the sublevel set topology changes – at minima, a new component is born, and at maxima, two components merge. To reflect these correspondences, we pair up critical points, choosing to pair a maximum with the latest relevant minimum, to reflect that the newer connected component merges in with the older one. The red line gives an example of a particular choice of height; the sublevel sets are split into two components, a fact reflected in the two bars intersected by the red line in the barcode – the number of bars at any given parameter value reflects

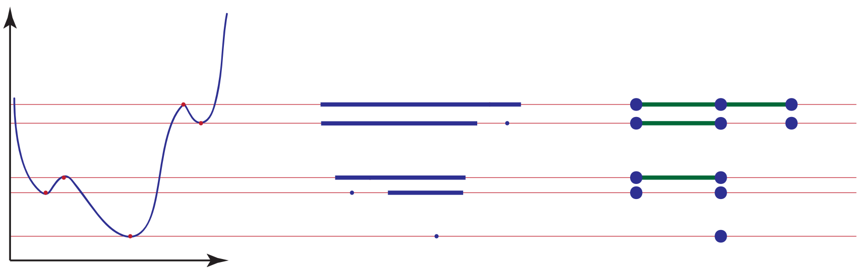

Figure 3: Going from a function on a manifold to a filtered sequence of spaces. Vertices of the Morse complex are given by the local minima, and each local maximum witnesses an edge connecting two neighbouring minima. To the left, we see the function with critical points marked, in the middle the sublevel sets at these points, and to the right, the corresponding filtered Morse complex. The filtered and parametrized structure of both spaces and complexes is clearly visible.

the corresponding Betti number at that stage.

The original persistence algorithm was formulated in terms of filtered complexes, and the functional view is fast to generate a filtered complex from the function under study. The key method to do this is described in [86]: in a Morse theory approach, cells of a cellular complex correspond to critical values of the function, and depending on the index of the critical point, we can read off the dimensionality of the cell.

The Morse theoretic viewpoint gives a translation dictionary between critical points and cells in all dimensions, even where the example given in Figure 3 is working in just one dimension. The fundamental feature to pay attention to is the *index* of a critical point – the number of negative signs in the appropriate quadratic form formulation of the Hessian at the critical point – the higher the index, the higher the dimension of the cell corresponding to that critical point and introduced at the parameter of its function value in a sublevel set filtration.

2.2.2 Tameness conditions

There is a family of tameness conditions enabling different strengths of statements, with inclusions of classes of modules along the arrows:

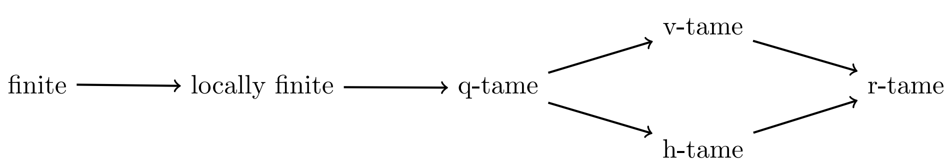

Here, a module M is...

finite if M is a finite direct sum of interval modules.

locally finite if M is a direct sum of interval modules, such that only finitely many span any given $t \in \mathbb{R}$.

q-tame if the measure corresponding to M is finite over every quadrant not touching the diagonal.

h-tame if the measure corresponding to M is finite over every horizontally infinite strip H not touching the diagonal.

v-tame if the measure corresponding to M is finite over every vertically infinite strip V not touching the diagonal.

r-tame if the measure corresponding to M is finite over every finite rectangle not touching the diagonal.

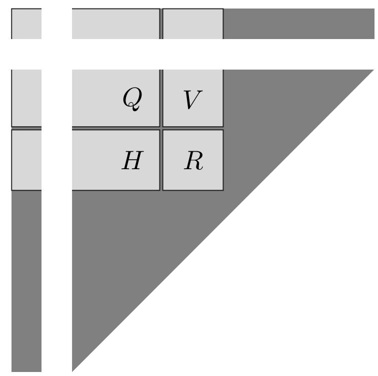

Figure 4: Tameness-conditions schematically illustrated.

These four last cases are sketched out in Figure 4; the horizontal bars above and to the left correspond to interval modules that survive until $+\infty$ and interval modules that were born at $-\infty$ respectively.

2.2.3 Categorification

Work by [63] studies the category of functors $(\mathbb{R}, \leqslant) \to \mathbf{Vect}_{\Bbbk}$, and is able to prove that the category of persistence modules is abelian.

They leverage this to prove a generous stability theorem: for *arbitrary* (not necessarily continuous) functions $\mathbb{X} \to \mathbb{R}$ from a topological space, and any functor H from topological spaces to a category of real-indexed diagrams in an abelian category $\mathcal{D}$, the *interleaving distance between the diagrams* generated by applying H to the sublevel set filtrations of the functions is bounded above by the L_∞*-distance* of the functions.

Furthermore, they prove many of the categories that emerge naturally in persistent homology are abelian.

2.3 Filtered topological spaces

The other culture present in the study of persistent homology focuses on the role of a *filtered topological space* and derived algebraic objects as the fundamental notion. This viewpoint has sparked a wealth of algebraic abstractions and given rise to several different notions of the *shape* of a persistent homology theory.

Connecting this viewpoint with the original study by [18], and indeed with the entire viewpoint present in Section 2.2, one may point out that for any function $f : \mathbb{X} \to \mathbb{R}$, the sublevel sets $f^{-1}((-\infty, x])$ form a filtration of $\mathbb{X}$. For tame enough – finitely many topological critical points, finite rank homology for any sublevel set, and similar conditions – functions, this filtration can be described by a finite filtration, or even a parametrization with finitely many different states.

Since homology is a functor, and inclusions are continuous maps, applying homology to a filtration produces a diagram of homology groups on the shape

$$H_j\mathbb{X}_0 \to H_j\mathbb{X}_1 \to \cdots \to H_j\mathbb{X}_n$$

and by interpreting this diagram as a module in one of a number of different possible module categories, further generalizations are possible.

Commonly, the geometric filtrations in use in persistent homology really are parametrizations – for any value $\varepsilon \in \mathbb{R}$, there is some resulting space $\mathbb{X}_\varepsilon$ – that happen to generate filtrations:

$$\mathbb{X}_{(-\infty,\varepsilon]} = \bigcup_{\delta\in(-\infty,\varepsilon]} X_\delta \qquad\qquad \varepsilon < \varepsilon' \Rightarrow \mathbb{X}_{(-\infty,\varepsilon]} \subseteq \mathbb{X}_{(-\infty,\varepsilon']}$$

For these cases, it is common to blur the lines between the definitions of filtered spaces and parametrized spaces.

2.3.1 Graded modules over $\mathrm{k}[t]$

The first significant advance in the choice of underlying algebraic structure for persistence modules came from [34]. They observe that a diagram of vector spaces

$$V_0 \to V_1 \to \dots$$

can be modelled as a graded module over the polynomial ring $\mathrm{k}[t]$. The module V_* is taken to have V_d in degree d, and the action of multiplying by t corresponds to the linear map $V_d \to V_{d+1}$. Seeing as homology groups with field coefficients are vector spaces, and the induced maps from the inclusions in the filtration are linear, this construction translates a persistent homology diagram to a graded module.

At this stage, [34] observe that the existence of a barcode decomposition follows directly from the fact that $\mathrm{k}[t]$ is a principal ideal domain, and therefore any module V_* decomposes into a direct sum of cyclic modules. These come in two versions: torsion modules isomorphic to $\mathrm{k}[t]/(t^d)$ for some natural number d, and free modules isomorphic to $\mathrm{k}[t]$. These two classes can be directly translated into free and finite intervals $[a, a+d)$ or $[a, \infty)$.

The work in [34] also demonstrates that the persistence algorithm described by [18] works with the same result for arbitrary field coefficients where the original description required coefficients in the field $\mathbb{Z}/2\mathbb{Z}$.

This work has been extensively cited – to the point where the papers [18, 34] are the standard reference citations for the persistence algorithm, and a number of extensions to the results have been provided, as well as numerous applications to the extension of expressive power the change of fields produces.

Results relying on non-binary fields The most obvious direct usefulness of the graded polynomial ring module approach has been in cases where the dependency of homology on the characteristic of the coefficients matters. This was the case in work by [39].

A study by [95] investigates the statistics of 3×3 pixel patches from naturally occuring images. They find, inter alia, a high density circle in the first few PCA coordinates. This circle, they notice, corresponds closely to linear gradient directions within the dataset.

[39] pick up the same dataset, and study it using persistent homology. They are able to recover two additional, secondary, high-density circle shapes within the dataset. These three circles combine to form a high-density 2-dimensional surface, which after computing persistent homology over both $\mathbb{Z}/2\mathbb{Z}$ and over $\mathbb{Z}/3\mathbb{Z}$ could be identified as the Klein bottle.

The ability to compute persistent homology with coefficients in $\mathbb{Z}/3\mathbb{Z}$ was crucial for this approach, and algorithmically dependent on the graded module over $\mathrm{k}[t]$ approach to persistent homology.

Multi-dimensional persistence With inspiration from the several relevant parameters affecting the analysis in [39], [14] constructed *multidimensional persistence*. The underlying observation is that just as graded modules over $\mathrm{k}[t]$ model singly parametrized topological spaces, adding more parameters corresponds to adding more variables to the polynomial ring. Hence, a d-dimensional parametrization can be modeled in a persistence way by working with graded modules over $\mathrm{k}[t_1, \dots, t_d]$.

The multi-dimensional theory has problems – chief among which is the lack of as useful a decomposition into a small and easy to describe class of indecomposables. The category of graded modules over $\mathrm{k}[t_1, \dots, t_d]$ has no complete discrete invariant, but [14] propose a discrete invariant – the rank invariant – turning out to be incomplete but useful.

The theory has been further studied since:

[12] introduce Grbner basis methods for computing multidimensional persistent homology, demonstrating that for *one-critical multifiltrations*, the rank invariant can be computed in polynomial time. The translation process they use to recast the problem to a Grbner basis computation has potential exponential blowup

behaviours for the general case. [100] demonstrate that by avoiding the mapping telescope and using more refined Grbner basis approaches, the computation can be bounded to polynomial time in general.

The multidimensional approach has received a lot of attention from the Italian size function community. [62, 68, 69, 66] treat multidimensional persistent homology in a size function framework as important tools for image analysis.

Questions of stability for persistence modules have been studied, both in the size function community ([68, 69]) and in the context of persistent homology by [96].

Cohomology and duality Persistent cohomology was mentioned by [43], who immediately use Lefschetz duality to transform it into relative homology. [83], later extended by [98], produce an algorithm for computing persistent cohomology and observe connections to computing intrinsic circle-valued coordinate functions from point cloud datasets.

This work inspired a paper by [46] in which two duality functors – $M_* \mapsto \hom_k(M_*, k)$ and $M_* \mapsto \hom_{k[t]}(M_*, k[t])$ on graded $k[t]$-modules are studied, and how these functors affect both the persistence algorithm itself, the ordering of basis elements in a sorted vector space approach, and how the barcodes are modified. These two functors allow the transport of information between relative and absolute versions of persistent homology and cohomology.

.

2.3.2 Modules over a quiver algebra

Another algebraic model that describes the persistent homology diagrams of vector spaces is given by quiver algebras. A persistent homology diagram of the shape

$$H_k(\mathbb{X}_0) \to H_k(\mathbb{X}_1) \to H_k(\mathbb{X}_2) \to H_k(\mathbb{X}_3) \to H_k(\mathbb{X}_4) \to H_k(\mathbb{X}_5)$$

can be considered as a module over the path algebra kQ for Q the quiver

$$\bullet \to \bullet \to \bullet \to \bullet \to \bullet \to \bullet$$

A theorem by [90] states:

Theorem 2.1 (Gabriel's theorem)**.** *Ein Kcher K hat genau dann nur endlich viele Isomorphieklassen von unzerlegbaren endlichdimensionalen* k*-linearen Darstellungen, wenn K eine "disjunkte Vereinigung" endlich vieler Kcher der Klassen A_e, D_m oder E_n ist, $e \geqslant 1$, $m \geqslant 4$, $6 \leqslant n \leqslant 8$.*

A quiver K has finitely many isomorphism classes of irreducible finite dimensional k*-linear representations if and only if K is a disjoint union of finitely many quivers of the classes A_e, D_m, or E_n for $e \geqslant 1$, $m \geqslant 4$, $6 \leqslant n \leqslant 8$. (translation: Mikael Vejdemo-Johansson)*

In particular, Gabriel goes on to prove that the exact isomorphism classes that show up for the quivers of type A_e – quivers of linear sequences of arrows, possibly alternating in direction – are the interval modules. These have some connected interval along the linear sequence where one-dimensional vector spaces are connected by identity maps – and outside this interval, all maps are and all vector spaces are zeros.

For the case of "classical" persistent homology, this recovers the barcode description for the case of a filtered finite simplicial complex: the persistent homology decomposes into a direct sum of irreducibles, and these irreducibles all are these interval modules. To describe each interval module, it is enough to state its start and end index, which is the exact data that a barcode conveys.

This approach has given rise to further generalization directions.

Zigzag persistence [67] pointed out that Gabriel's theorem has concrete consequences for topological data analysis. In particular, the non-dependency on arrow direction for a quiver to qualify as having type A_e means that we can consider quivers where arrows alternate direction, either occasionally or consistently.

This paper introduces the fundamental idea, provides matrix algorithms for computing zigzag persistence, and provides the *diamond principle*, relating how local changes along the zigzag reflect in changes to the persistence diagram. The paper also suggests several applications where the zigzag naturally arises:

Balancing different parameters In the study by [39], the $p\%$ densest points as computed with a parametrized density estimator were used to determine the topology of the dataset. For studies like this one, it is worth while to try to work with all possible values of the parameter determining the density estimator at once – to replicate the success persistent homology has in sweeping over entire ranges for a parameter.

Writing $\mathbb{X}_r^p$ for the densest $p\%$ of the point cloud $\mathbb{X}$ as measured using the parameter r. Varying r along $r_1 < r_2 < \cdots < r_N$, there is a zigzag

$$\begin{array}{ccccccccccccc} & X_{r_1}^p \cup X_{r_2}^p & & X_{r_2}^p \cup X_{r_3}^p & & X_{r_3}^p \cup X_{r_4}^p & & \dots & & X_{r_{N-1}}^p \cup X_{r_N}^p & \\ \nearrow & & \nwarrow \quad \nearrow & & \nwarrow \quad \nearrow & & \nwarrow \quad \nearrow & & \nwarrow \quad \nearrow & & \nwarrow \\ X_{r_1}^p & & X_{r_2}^p & & X_{r_3}^p & & X_{r_4}^p & & X_{r_{N-1}}^p & & X_{r_N}^p \end{array}$$

For each point cloud in this sequence, compute a geometric complex, and compute its homology – the resulting diagram is a zigzag diagram, and its decomposition into barcodes carries information about the variation of r in a way directly analogous to how persistent homology itself measures homological features over varying values for a parametrization.

Topological bootstrapping Similar to bootstrapping in statistics, one may want to take a sequence of small samples $\mathbb{X}_i$ from a large dataset $\mathbb{X}$ and estimate the topology of each $\mathbb{X}_i$ individually. Doing this, disambiguation between local features of each $\mathbb{X}_i$ and global features of $\mathbb{X}$ is not entirely transparent.

Here, the *union zigzag* provides a method for persisting features across several samples:

$$\begin{array}{ccccccccccc} & X_1 \cup X_2 & & X_2 \cup X_3 & & X_3 \cup X_4 & & \dots & & X_{N-1} \cup X_N & \\ \nearrow & & \nwarrow \quad \nearrow & & \nwarrow \quad \nearrow & & \nwarrow \quad \nearrow & & \nwarrow \quad \nearrow & & \nwarrow \\ X_1 & & X_2 & & X_3 & & X_4 & & X_{N-1} & & X_N \end{array}$$

Features that are local to any one of the point clouds will not persist along the zigzag, while global features will be carried along the zigzag to long barcodes.

This approach was further studied for practical aspects by [103], who give concrete algorithms for the computation of the union zigzag, and demonstrate the computational behaviour on a number of concrete examples, including the images dataset studied in [39].

The union zigzag was also further applied to dynamic network analysis by [91].

Levelset zigzag Given a space $\mathbb{X}$ and a continuous function $f : \mathbb{X} \to \mathbb{R}$, the levelset zigzag would relate the levelsets of f through a zigzag, introduced by [13]:

$$\begin{array}{ccccccccc} & f^{-1}([s_1, s_2]) & & f^{-1}([s_2, s_3]) & & \dots & & f^{-1}([s_{n-1}, s_n]) & \\ \nearrow & & \nwarrow \quad \nearrow & & \nwarrow \quad \nearrow & & \nwarrow \quad \nearrow & & \nwarrow \\ f^{-1}(s_1) & & f^{-1}(s_2) & & f^{-1}(s_3) & & f^{-1}(s_{n-1}) & & f^{-1}(s_n) \end{array}$$

where a_j are the critical values of f, and s_j are picked to satisfy:

$$-\infty < s_0 < a_1 < s_1 < a_2 < \cdots < s_{n-1} < a_n < s_n < \infty$$

This zigzag produces a computational approach to the interval persistence introduced by [84].

[13] also elaborate the diamond principle to connect it with the Mayer-Vietoris long exact sequence, and give a concrete graphical language for modifying barcodes between union and intersection zigzag sequences. This Mayer-Vietoris relation produces a large diagram from the levelset zigzag (see description below) introduced in the paper that connects to extended persistence, and admits a stability theorem.

2.4 Shapes of theories

A dichotomy such as the one we have seen above cries out for a unifying theory – everyone start out with the same underlying problem, and believe they do approximately the same thing, there should be a way to treat all the algebraic foundations in use as aspects of the same underlying theory. While such a unification is not published as this paper is finalized, there are several ongoing efforts in the community that may well lead towards a unifying theory of persistent homology.

3 Topos of Persistence

Our approach to a unifying theory of persistent homology aims to construct a *topos* of sheaves that anchors the parameter variation of persistence at the level of the underlying logic and set theory.

A topos is a category that is similar enough to the category of sets that one can reconstruct constructive logic within the new topos. Topoi have been utilized with success to model fuzzy sets, temporal progressions in quantum mechanics, and to connect geometry, topology and logic to each other.

A corner stone to building a new logic from a topos is a foundational result: the category of sheaves over some fixed underlying space is always a topos, and the resulting set theory takes on the shape of the underlying space of the sheaves. In particular, one can connect the structure of the collection of subsheaves of a single sheaf to the geometry of the underlying space.

A minimal structure that allows for the definition of a category of sheaves in this way is in fact a weaker condition than expecting a topological space: one axiomatization of the minimal conditions required is given by Heyting algebras. By picking a Heyting algebra, one can shape the set theory. We may recover classical logic by picking the Heyting algebra of open sets of the one-point topological space as our underlying space.

Among the payoffs for modeling persistent homology this way would be the immediate translation of a large swathe of results from classical topology to the realm of persistent homology: any results on simplicial complexes that can be proven by intuitionistic logic will automatically hold true for persistent homology. This direction also ties in directly with very recent developments on extending and refining Reeb graph methods by de Silva, Munch and Patel (personal communications, 2013). To model persistent homology, we want to construct a Heyting algebra that produces sets

where elements have lifetimes: an element appears, lives for a while, and then disappears. The elements of such a Heyting algebra would be the possible lifetimes of elements, the set of elements over any given such lifetime would be the collection of elements that exist at least for the entire duration observed.

Hence, the elements we would choose are intervals, and the ambition of this work package is to build a Heyting algebra structure on the collection of intervals, and then develop homology of simplicial complexes within the logic produced by the particular topos of sheaves over that Heyting algebra.

One of the achievements of the project to date is the construction of a promising candidate Heyting algebra.

The inspiration for our approach is to some extent rooted in the exposition by Barr and Wells [3], where sheaves of sets are described as sets with a particular shape. The shape corresponds to the shape of the underlying site, which also describes the shape of the available *truth values* for the corresponding logic. Classical logic and set theory would correspond to having two discrete truth values; fuzzy logic to a continuum of truth values encoding reliability of a statement, and the persistent approach would encode truth as valid over some regions of a persistence parameter, but not other. We are looking for a formulation that would make the topos setting clear and amenable for generalization. Fundamental to such an approach is the formulation of an underlying space, a *site*, such that the sheaves of vector spaces over this site correspond to persistence modules. Defining such a site turns out to be more subtle than we first suspected, and the paper takes us throu We first go through the requisite lattice theory to build a site over *persistence diagrams*.

3.1 Motivation

The similarities in definitions and in algorithms of the different flavors of persistence (i.e. multi-dimensional, zig-zag, extended, etc.) suggest that all three should be instances of a unifying theory; We believe that topos-theory can provide such a unifying theory.

In particular, our inspiration draws from the presentation of *time sheaves* from Barr & Wells [3] as sets with a temporally varying structure. Based on this as an inspirational source, we aim to develop the theory of persistent homology as the internal homology of simplicial complexes over a set theory in which elements have life-times; where elements of sets, simplices of simplicial complexes, and base vectors of vector spaces are born, live, and die, much like the language has evolved for discussing persistent homology.

In such a setting, a filtered topological space corresponds to a topological space where parts of the space come in at later times; the construction of the homology functor immediately provides homology groups where elements come in and go away as time flows.

Sheaves are most often defined as diagrams with the open sets of a topological space as their shape; but more general diagrams can be used to define them. We have chosen to focus on a model for the *site* of sheaves called *Heyting algebra*, much used in logic. Based on this approach, a sheaf is a functor from a particular Heyting algebra to the category of sets, or of topological spaces, or of simplicial complexes, or of vector spaces – such that a categorical version of the sheaf axiom holds.

Our first attempt was to consider as our Heyting algebra the collection of open *intervals* in $\mathbb{R}$, with the analogy of union taken to be the smallest covering interval. This construction unfortunately fails. Instead, as our next and more successful attempt, we study the Heyting algebra of collections of vertices in a persistence diagram. In order to guarantee the axioms necessary to define a sheaf we need to allow *virtual persistence*: points below the diagonal. Classical persistence will emerge as a subcategory of the resulting topos of sheaves.

The fundamental observation is that we have seen numerous cases lately where the *shape* of a persistence theory matters; there has been the classical persistent homology, zigzag persistence, multi-dimensional persistence. In all of these cases, there is a sense of shape to the theory, embodied by a choice of algebra and module category that reflects the kinds of information we can extract from the method.

3.2 Preliminaries

Definition 3.1. *A* lattice *is a poset for which all pairs of elements have an infimum and a supremum. Whenever every subset of a lattice L has a supremum and a infimum, L is named a* complete lattice.

A lattice A can be seen as an algebraic structure $(L; \wedge, \vee)$ with two operations $\wedge$ and $\vee$ satisfying the following properties, for all $x, y, z \in L$:

(i) $x \wedge (y \wedge z) = (x \wedge y) \wedge z$ and $x \vee (y \vee z) = (x \vee y) \vee z$ (associativity);

(ii) $x \wedge x = x = x \vee x$ (idempotence);

(iii) $x \wedge (x \vee y) = x = x \vee (x \wedge y)$ (absorption);

(iv) $x \wedge y = y \wedge x$ and $x \vee y = y \vee x$ (comutativity).

The equivalence between this algebraic perspective of a lattice L and its ordered perspective is given by the following: for all $x, y \in L$, $x \leqslant y$ iff $x \wedge y = x$ iff $x \vee y = y$. The natural numbers form a lattice under the operations of taking the greatest common divisor and least common multiple, with divisibility as the order relation.

Example 3.2. *[5] The sets $O(X)$ and $C(X)$ of open sets and closed sets of a topological space are both examples of complete lattices where, in $O(X)$ the join and meet of a subfamily $\{ U_i \mid i \in I \}$ are given by*

$$\bigvee_{i\in I} U_i = \bigcup_{i\in I} U_i \quad \bigwedge_{i\in I} U_i = interior(\bigcap_{i\in I} U_i)$$

while in $C(x)$ the join and the meet of a subfamily $\{ A_i \mid i \in I \}$ are given by

$$\bigvee_{i\in I} A_i = closure(\bigcup_{i\in I} U_i) \quad \bigwedge_{i\in I} U_i = \bigcap_{i\in I} U_i.$$

Lemma 3.3. *[22] The following are useful characterizations of the distributivity of a lattice L:*

(i) *L is modular iff all x, y, z with $y \leqslant z$ are such that $x \wedge y = x \wedge z$ and $x \vee y = x \vee z$ imply $y = z$;*

(ii) *L is distributive iff A lattice L is distributive iff all x, y, z are such that $x \wedge y = x \wedge z$ and $x \vee y = x \vee z$ imply $y = z$;*

(iii) *L is modular iff if it does not have embedded any copy of the pentagon $\mathbf{N}_5$;*

(iv) *L is distributive iff if it does not have embedded any copy of the diamond $\mathbf{M}_3$ or the pentagon $\mathbf{N}_5$.*

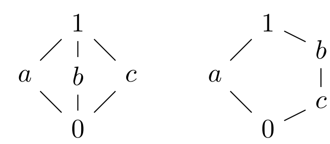

$\mathbf{M}_3$ $\mathbf{N}_5$

Definition 3.4. *A bounded lattice L is a* Heyting algebra *if, for all $a, b \in L$ there is a greatest element $x \in L$ such that $a \wedge x \leqslant b$. This element is the* relative pseudo-complement *of a with respect to b denoted by $a \Rightarrow b$. A subalgebra of an Heyting algebra is thus closed to the usual lattice operations $\wedge$ and $\vee$, and to $\Rightarrow$. An homomorphism between Heyting algebras must preserve both lattice operations as well as the implication operation and both top and bottom elements. Heyting algebras and respective morphisms determine an algebraic category.*

Example 3.5. *The lattice of open sets of a topological space X forms a Heyting algebras under the operations of union $\cup$, empty set $\varnothing$, intersection $\cap$, whole space X, and the operation*

$$U \Rightarrow V = \text{ interior of } (X - U) \cup V.$$

Remark 3.6. *All the finite nonempty chains (that are bounded and complete) constitute Heyting algebras, where $a \Rightarrow b$ equals b whenever $a > b$, and 1 otherwise. In general, every complete distributive lattice constitutes an Heyting algebra. Every Boolean algebra is a Heyting algebra, with $a \Rightarrow b$ given by $\neg a \vee b$. Heyting algebras are examples of distributive lattices that are not Boolean algebras. Any finite distributive lattice is the reduct of a unique Heyting algebra. More generally, the same result holds for any complete and completely distributive lattice.*

The following propositions show that Heyting algebras can be equationally defined and present several algebraic properties of these algebras.

Proposition 3.7. *[5] Let L be a bounded lattice and $\Rightarrow$ a binary operation on L. Then L is a Heyting algebra iff the following equations hold:*

(i) $a \Rightarrow a = 1$,

(ii) $(a \wedge (a \Rightarrow b) = a \wedge b$,

(iiii) $b \wedge (a \Rightarrow b) = b$,

(iv) $a \Rightarrow (b \wedge c) = (a \Rightarrow b) \wedge (a \Rightarrow c)$.

Definition 3.8. *The* implication *operation on a Heyting algebra, sends each pair of elements a and b to the element $a \Rightarrow b$ and sends each element a to the element $a^* = a \Rightarrow 0$, called the* pseudocomplement *of a. The operation of equivalence $\Leftrightarrow$ can be defined by $a \Leftrightarrow b = (a \Rightarrow b) \wedge (b \Rightarrow a)$.*

Proposition 3.9. *[5] Let L be a Heyting algebra and $a, b, c \in L$. Then,* *(i)* $a \Rightarrow (b \Rightarrow c) = (a \wedge b) \Rightarrow c$,

(ii) $a \Rightarrow b = 1$ *iff* $a \leqslant b$,

(iii) $a \Leftrightarrow b = 1$ *iff* $a = b$,

(iv) *if* $b \leqslant c$ *then* $(a \Rightarrow b) \leqslant (a \Rightarrow c)$,

(v) $a \wedge (a \Rightarrow b) \leqslant b$,

(vi) $b \leqslant a^*$ *iff* $b \wedge a = 0$ *iff* $a \leqslant b^*$,

(vii) $a \subseteq a^{**}$ *and* $a^{***} = a^*$,

(viii) $(a \vee b)^* = a^* \wedge b^*$,

Definition 3.10. *A* complete Heyting algebra *is a Heyting algebra which constitutes a complete lattice.*

Proposition 3.11. *[5] Let L be a complete Heyting algebra. Then the following identity holds:*

$$x \wedge \bigvee_{i \in I} y_i = \bigvee_{i \in I} (x \wedge y_i).$$

Conversely, any complete lattice L satisfying the identity above is a complete Heyting algebra, defining the implication operation by

$$x \Rightarrow y = \bigvee \{ x \in L \mid x \wedge a \leqslant b \}.$$

Proposition 3.12. *[5] Let P be a poset. Then, P is a complete Heyting algebra iff the following conditions hold:*

HA-1 There is a top element in P;

HA-2 Each pair of elements $x, y \in P$ *has an infimum,* $x \wedge y$*;*

HA-3 Every subset $\{ x_i \}_{i \in I}$ *of elements of P has a supremum,* $\bigvee_{i \in I} x_i$*;*

HA-4 For every element $x \in P$ *and every subset* $x_{i i \in I} \subseteq P$, $x \wedge (\bigvee_{i \in I} x_i) = \bigvee_{i \in I} (x \wedge x_i)$.

Proposition 3.13. *[21] A complete Heyting algebra is any complete lattice satisfying the infinite distributive law defined in Proposition 3.11.*

Definition 3.14. *A* presheaf *is any contravariant functor $\mathcal{F}$ from the opposite category of a given category $\mathcal{C}$ to the category of sets Set. A presheaf E over an Heyting algebra H is a* sheaf *if, whenever* $\{ x_i \}_{i \in I}$ *is a subset of H with supremum x,* $x = \bigvee_{i \in I} x_i$ *implies that*

$$E(x) \xrightarrow{e} \Pi_{i \in I} E(x_i) \underset{d}{\overset{c}{\rightrightarrows}} \Pi_{i,j \in I} E(x_i \wedge x_j)$$

is an equalizer (cf. [3]).

Theorem 3.15 (Stone Spaces 1977)**.** *The category of sheaves on a Heyting algebra is a topos.*

Remark 3.16. *For any topos E and any object A in E, the sub objects of A form a Heyting lattice. The subobject classifier in a topos is always a Heyting algebra. The underlying Heyting lattice structure in a topos is mainly interesting because it gives a way to observe the internal logic of a topos.*

3.3 Persistence Diagrams Lattice Representation

In the following we are going to discuss a third case where our elements are not bars but representations of bars in the persistence diagram. In fact, as the persistence diagram has the same topological information as the barcode, this structure has just what we need from it. Moreover, it does not include disjoint bars but includes instead what we could think of ”negative” bars in the sense that the birth time is after its death time (cf. [6]). We shall see that such a diagram constitutes a complete Heyting algebra. Let us assume a persistence diagram bounded by $(0, 0)$ and $(\varepsilon_1, \varepsilon_2)$.

3.4 Persistence diagrams as distributive lattices

Consider now the representation of barcodes in a persistence diagram. Let $A = [a_1, a_2]$ and $B = [b_1, b_2]$, where a_1 and b_1 are birth times, and a_2 and b_2 are death times. These intervals are represented in the persistence diagram by the points $A(a_1, a_2)$ and $B(b_1, b_2)$. There is a natural way to define operations between these intervals in the barcode, presented in the following:

Definition 3.17. *Let A and B be intervals in the barcode represented in the persistence diagram $\mathbb{P}$. Consider the operations*

$$A \wedge B = (max\{\, a_1, b_1 \,\}, min\{\, a_2, b_2 \,\})$$

$$A \vee B = (min\{\, a_1, b_1 \,\}, max\{\, a_2, b_2 \,\})$$

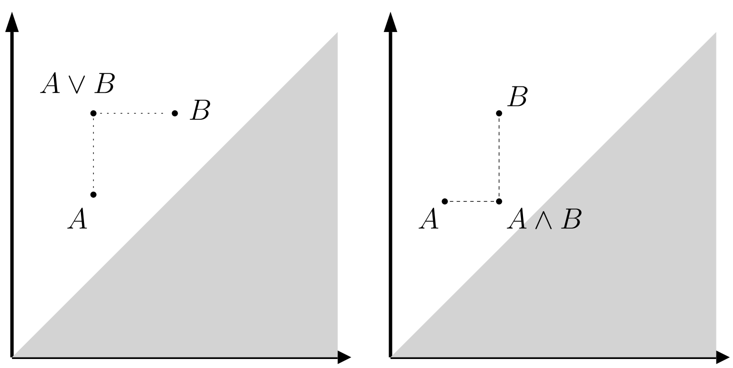

With the operations above we can define a lattice of intervals in the barcode, named *barcode lattice.* In the following paragraphs we discuss several of its nice structural properties.

Proposition 3.18. *The persistence diagram $\mathbb{P}$ together with the operations above determines a lattice.*

This is the lattice structure of all possible bars in the persistence diagram. Thus, every persistence diagram is a part of this lattice and, as any poset, can be completed into a complete Heyting algebra In fact, the barcode lattice is itself a complete Heyting algebra so that all persistence diagrams are indeed its subalgebras as we shall see later in this section.

Remark 3.19. *In this case the counter-example is now an example for which distributivity holds: take $A(0, 1)$, $B(0, 2)$ and $C(2, 3)$ and observe that $A \vee C = (0, 3) = B \vee C$ but $A \wedge C = (2, 1) \neq (2, 2) = B \wedge C$. In fact, $(A \vee B) \wedge C = (0, 2) \wedge (2, 3) = (2, 2) = (2, 1) \vee (2, 2) = (A \wedge C) \vee (A \wedge C)$.*

Theorem 3.20. *The barcode lattice is a distributive lattice.*

The partial order assumed here is determined by the above lattice operations as follows:

Lemma 3.21. *$A \leqslant B$ iff $b_1 \leqslant a_1$ and $a_2 \leqslant b_2$.*

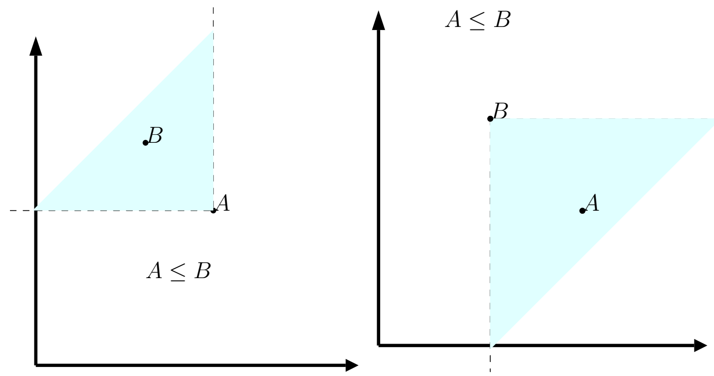

Proposition 3.22. *$A \leqslant B$ iff $A \subseteq B$.*

Remark 3.23. *In the sense of Lemma 3.21, a bar A is below a bar B in the barcode iff B is above and on the right of A in the persistence diagram. Dually, B is above A iff B is on the right and below A.*

Hence, whenever the diagram $A < B < C < D$ is represented as below, it constitutes a chain.

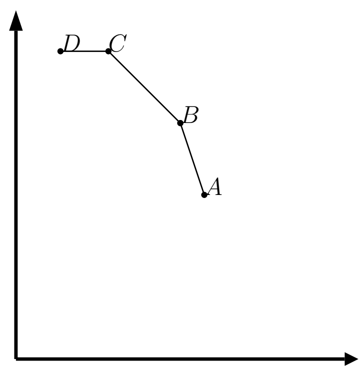

On the other hand, this is not a total order. *In fact, unrelated bars $A(a_1, a_2)$ and $B(b_1, b_2)$ are of such sort that $a_1 \leqslant b_1$ and $a_2 \leqslant b_2$. Moreover,* the smallest element 0 *is in the right lower corner of the diagram, correspondent to the point $(\varepsilon_1, 0)$, while* the biggest element 1 *is on the left upper corner, correspondent to the point $(0, \varepsilon_2)$. Furthermore,* if one of the coordinates is equal, then the correspondent bars are always related. *This can be interpreted by the next proposition.*

Proposition 3.24. *Bars with equal death time or birth time must be related.*

3.5 Irreducibility

Irreducible elements are of great importance for representation theories and decomposition theorems.

Proposition 3.25. *The join-irreducible elements of $\mathbb{P}$ are all the elements with coordinates $(x_1, 0)$ or (ε_1, x_2). Dually, the meet-irreducible elements are all the elements with coordinates $(0, x_2)$ and (x_1, ε_2).*

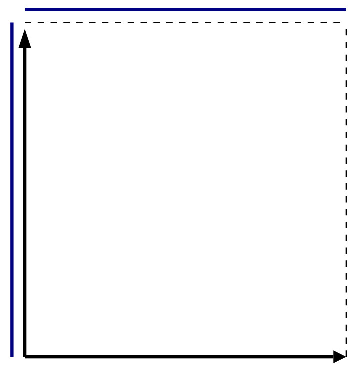

3.6 The barcode algebra

By construction, the lattice operations defined in Theorem 3.20 can naturally be generalized from pairs of elements to any set of elements as follows:

Proposition 3.26. *The persistence diagram $\mathbb{P}$, together with the operations below determines a complete lattice.*

$\bigwedge_i A_i = (max\{\, a_{1i}\,\}, min\{\, a_{2i}\,\})$

$\bigvee_i A_i = (min\{\, a_{1i}\,\}, max\{\, a_{2i}\,\})$

Due to the above the next result is of great interest to the construction of a topos foundation based on time intervals. It is a direct consequence of completeness and distributivity, due to Theorems 3.26 and 3.20, respectively.

Corollary 3.27. *The persistence diagram $\mathbb{P}$, together with the operations defined in Theorem 3.26 is a complete Heyting algebra. In particular, the following identity holds:*

$$X \wedge \bigvee_{i \in I} Y_i = \bigvee_{i \in I} (X \wedge Y_i).$$

To the order structure determined in Corollary 3.27 we will call *barcode algebra* (or *barcode lattice* when we want to emphasize its order structure properties).

Usually, the infinite distributive law of Corollary 3.27 is compatible with the idea of the finite meets and a arbitrary joins in topological spaces. This will permit the dual of the infinite distributive law to hold, as well as it will make possible the existence of complete Heyting algebra homomorphisms.

Proposition 3.28. *The barcode algebra constitutes a completely distributive lattice.*

Corollary 3.29. *The barcode algebra L can be embedded into a direct product of chains $[0, 1]$ by an order embedding that preserves arbitrary meets and joins. Moreover, both L and its dual order L^{op} are continuous posets.*

Remark 3.30. *As for homomorphisms, they shall preserve arbitrary meets and joins as well as the arrow operation. Translations are examples of isomorphisms between barcode algebras: let $f_\epsilon : X \to X$ defined by $f(x_1, x_2) = (x_1 + \epsilon, x_2)$. Clearly $f_\epsilon(x \wedge y) = f_\epsilon(x) \wedge f_\epsilon(y)$ and, similarly, $f_\epsilon(x \vee y) = f_\epsilon(x) \vee f_\epsilon(y)$. This shows that isomorphisms can be established between distinct barcode algebras.*

Proposition 3.31. *Let $f : L \to M$ be a map between barcode lattices preserving sups. Then $f(L)$ is closed under sups in M and is a complete lattice in itself.*

Remark 3.32. *Observe that the barcode lattice is not a Boolean algebra: indeed, take $A(a_1, a_2)$ such that $a_1, a_2 \neq 0$, $a_1 \neq \varepsilon_1$ and $a_2 \neq \varepsilon_2$ where ε_1 and ε_2 are the biggest values on the X coordinates and Y coordinates. Then, for all $B(b_1, b_2)$, $A \wedge B = 0$ implies $\varepsilon_1 = max\{a_1, b_1\}$ and $0 = min\{a_2, b_2\}$, that is, $\varepsilon_1 = b_1$ and $0 = b_2$. Similarly, $A \vee B = 1$ implies $0 = min\{a_1, b_1\}$ and $\varepsilon_2 = max\{a_2, b_2\}$, that is, $0 = b_1$ and $\varepsilon_2 = b_2$, contradicting the latter statement. Note that the only pairs of elements with complement are 0 and 1, and $(0, 0)$ and $(\varepsilon_1, \varepsilon_2)$, the pairs of opposite corners of the persistence diagram.*

3.7 Internal logic aspects

Let us now have a look at the arrow operation. The following result describes this operation for any of the possible four cases.

Theorem 3.33. *Let $A, B \in Pe$. Then, if A and B are (order) related,*

$$A \Rightarrow B = \begin{cases} 1 = (0, \varepsilon_2) & \text{, if } b_1 \leqslant a_1 \text{ and } a_2 \leqslant b_2 \\ B = (b_1, b_2) & \text{, if } a_1 \leqslant b_1 \text{ and } b_2 \leqslant a_2 \end{cases}.$$

Otherwise,

$$A \Rightarrow B = \begin{cases} (b_1, \varepsilon_2) & \text{, if } a_1 \leqslant b_1 \text{ and } a_2 \leqslant b_2 \\ (0, b_2) & \text{, if } b_1 \leqslant a_1 \text{ and } b_2 \leqslant a_2 \end{cases}.$$

Remark 3.34. *In most cases the pseudocomplement is 0. Though, the pseudocomplement of $(0, 0)$ is $(\varepsilon_1, \varepsilon_2)$. In fact, this seems to be, together with $1 = (0, \varepsilon_2)$ and $0 = (\varepsilon_1, 0)$, the only pair of elements with complement. Hence, the derived Boolean algebra is*

$$\{x \mid x = \neg\neg x\} = \{0, (0, 0), (\varepsilon_1, \varepsilon_2), 1\}.$$

Remark 3.35. *The results given in Propositions 3.7 and 3.9 give us a list of properties that we can always use: for all $A, B, C \in \mathbb{S}(I)$, the following statements hold,*

(i) $A \Rightarrow A = 1$,

(ii) $(A \wedge (A \Rightarrow B) = A \wedge B$,

(iii) $B \wedge (A \Rightarrow B) = B$,

(iv) $A \Rightarrow (B \wedge C) = (A \Rightarrow B) \wedge (A \Rightarrow C)$,

(v) $A \Rightarrow (B \Rightarrow C) = (A \wedge B) \Rightarrow C$,

(vi) $A \Rightarrow B = 1$ *iff* $A \leqslant B$,

(vii) $A \Leftrightarrow B = 1$ *iff* $A = B$,

(viii) *if* $B \leqslant C$ *then* $(A \Rightarrow B) \leqslant (A \Rightarrow C)$,

(ix) $A \wedge (A \Rightarrow B) \leqslant B$,

(x) $B \leqslant A^*$ *iff* $B \wedge A = 0$ *iff* $A \leqslant B^*$,

(xi) $A \subseteq A^{**}$ *and* $A^{***} = A^*$,

(xiii) $(A \vee B)^* = A^* \wedge B^*$,

These axioms will show to be useful for the construction of algorithmic applications. Recall that Heyting algebras are the model structures for intuitionistic logic and, therefore, this particular topos is coherent with all constructive persistence.

3.8 On a Topos for Persistence

Complete Heyting algebras form a variety of algebras (thus closed for subalgebras, products and homomorphic images). The class of these order structures also constitutes three relevant categories that share the same objects and differ only on their morphisms. They are the category of *complete Heyting algebras*, denoted by $cHey$, for which their morphisms preserve infinite joins and meets as well as the implication operation; the category of *frames*, denoted by Frm, where the morphisms are necessarily monotone functions that preserve finite meets and arbitrary joins; and most important for this paper, the category of *Locales*, denoted by *Loc*, for which the morphisms are the opposite morphisms of Frm.

The relation of locales and their homomorphisms to topological spaces and continuous functions is well known (cf. [26]). In fact, the set of all open sets of a topological space X constitutes a locale $\Omega(X)$ that this is a part of the power set Boolean algebra $\mathcal{P}(X)$. In such a model, the operations $\wedge$ and $\vee$ correspond to intersection and union, respectively, while $U \Rightarrow V$ is the largest open set W such that $W \cap U \subseteq V$, and $\neg U$ is the interior of the complement of U being the largest open set disjoint from U. Furthermore, to each locale $\mathcal{L}$ corresponds a topological space for which the points $pt(\mathcal{L})$ are equivalently described by:

- a frame morphism from $\mathcal{L}$ to $\mathbf{2}$;
- a principal prime ideal of $\mathcal{L}$;
- a $\wedge$-prime element of $\mathcal{L}$.

and for which the appropriate topology is given by the open sets

$$\phi(a) = \{\, p \in pt(\mathcal{L}) \mid p(a) = 1 \,\}$$

for all $a \in \mathcal{L}$. This permits us to construct the space of bars correspondent to each barcode, in the following.

In the case of Boolean algebras, the corresponding points are their atoms. Though, for Heyting algebras, the points correspond to the prime ideals as in Esakia duality.

3.9 From local to global

All sheaves on locales are constituted of elements defined at various levels (cf. [10], Chapter 2).

Let $\mathcal{L}$ be a locale. The classical construction of sheaves over Heyting algebras considers a presheaf $\mathcal{F} : \mathcal{L} \to Set$ to be a contravariant functor. When an element is defined at some level $u \in \mathcal{L}$, it can be restricted to any smaller level $v \leqslant u$ and, moreover, if elements are defined at levels $u_i \leqslant u$ (with $i \in I$) in a sufficiently compatible way, they can be glued together to become an element defined at the level u.

Formally, given $v \leqslant u$ in $\mathcal{L}$, $\rho^u_v : \mathcal{F}(u) \to \mathcal{F}(v)$ defined by $\rho^u_v(x) = x_{|v}$. Functionality implies $x_{|u} = x$ and $x_{|w} = (x_{|v})_{|w}$; compatibility implies $x_{i|u_i \wedge u_j} = x_{j|u_i \wedge u_j}$; and separability implies $x = y$ whenever $x_{|u_i} = y_{|u_i}$ and $u = \bigvee_{i \in I} u_i$.

The sheaf condition for locales establishes that, if $u = \bigvee_{i \in I} u_i$ is a compatible family in $\mathcal{L}$ and $x_i \in \mathcal{F}(u_i)$, then there exists a unique $x \in \mathcal{F}(u)$ such that $x_{|u_i} = x_i$, i.e., x is the gluing of the family $(x_i)_{i \in I}$.

Alternatively, the sheaf condition can be characterized by: given $\mathcal{F}$ a separated presheaf and $u = \bigcup_{i \in I} u_i$, then every compatible family $(x_i \in \mathcal{F}(u_i))_{i \in I}$ can be glued into an element $x \in F(u)$ such that for every $i \in I$, $x_{|u_i} = x_i$.

Remark 3.36. *Here the open sets considered are the elements of the complete Heyting algebra of bars, i.e., the space of all persistence diagrams. In that sense, the restriction of an element x in u to an open v such that $v \leqslant u$, $\rho^u_v(x) = x_{|v}$, regards the restriction to a smaller interval $V \subseteq U$, where u and v are the correspondent representations of the bars U and V in the persistence diagram.*

The adopted definition of sheaf in Definition 3.14 seen as an equalizer permits us a construction motivated by the relations between elements in the lattice. Consider a presheaf $E : H \to Set$ that assigns to each element $z \in H$ a set $E(z)$, and assume $x = \bigvee_i x_i$. Take to bars A and B and their representations a and

b in the persistence diagram. From the relation $a \wedge b \leqslant a, b \leqslant a \vee b$ one can derive the restriction maps $E(a \vee b) \to E(a) \to E(a \wedge b)$ and $E(a \vee b) \to E(b) \to E(a \wedge b)$ Let $c \in L$ be such that $a, b \leqslant d$. Then $a \vee b \leqslant c$ with the restriction map $E(c) \to E(a \vee b)$ so that $E(a \vee b)$ is the smallest with maps into both $E(a)$ and $E(b)$. Thus, $E(a \vee b)$ is the pullback of the maps $E(a) \to E(a \wedge b)$ and $E(b) \to E(a \wedge b)$. It seems that, in an analogous way, $E(a \wedge b)$ is the pullback of the maps $E(a \vee b) \to E(a)$ and $E(a \vee b) \to E(b)$. In order to deal with more then a pair of elements of the lattice we shall talk about equalizers instead of pullbacks and coequalizers instead of pushouts.

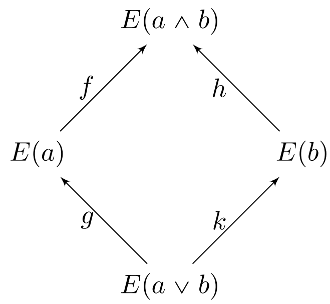

A topos is a category with certain properties characteristic of the category of sets. The theory of sheaves on a locale appears as a generalized set theory, provided with an equality which takes for possible truth values all the elements of the locale. It is well known that sheaves over locales constitute a topos (cf. [2], Theor. 2.4). In general, for any small category C, the category of diagrams $Set^{C^{op}}$ is a topos since it has limits and exponentials, with subobject classifier as follows: given any object E and subobject $U \rightarrowtail E$, define $u : E \to \Omega$ at any object $C \in \mathbf{C}$ by the sieve of arrows[1] into C that take $e \in E(C)$ back into the subobject U, that is,

$$u_C(e) = \{\, f : D \to C \mid f^*(e) \in U(D) \rightarrowtail E(D) \,\}$$

for any $e \in E(C)$.

4 Sheaves

An alternative, intuitive view of a sheaf is that it assigns to open sets some algebraic structure, such that this local information can be glued together. The limit of the these sheaves represent global sections. First we show how in the zig-zag case, this can be exploited to compute certain zig-zags in parallel. Although we do not give an explicit sheaf construction, the idea is sheaf-like. In the second part, we give a contruction for approximating local homology. This is an important step in stratified manifold learning. In the future we will use this construction along with sheaf theory to compute strata.

4.1 Pullback

We present a new viewpoint on zig-zag persistent homology and with this viewpoint achieve a naturally parallelizable and distributed algorithm whose performance scales with the number of processors available up to a worst case running time of $O(m^\omega \log^2 n)$, which in some cases is better than other methods currently known. Further to the benefit of the proposed algorithm, it is simple to implement, and we give some preliminary experimental results. Its approach is sheaf-like. Recall that the global section can be phrased in terms of categorical limits. In our context, the limits are simply pullbacks. This viewpoint in some cases leads to gains in efficiency.

The efficiency of our approach relies on the problem being correspondingly well-conditioned. Our motivating example is the case of union zig-zags of multiple independent samplings of some space, and in this type of setting performance gains are likely. This particular approach has gained interest in the applied topology

[1] i.e. a set of arrows that all have C as their target

community as a topological form of boosting. With the introduction of large data sets, parallelization and distributed computation of homology are increasingly important.

We make use of the pullback and pushout constructions:

The **pullback** of a pair of linear maps $A \xrightarrow{f} X \xleftarrow{g} B$ with the same target vector space is a subspace $P \subset A \oplus B$ consisting of $\{(a, b) \in P : f(a) = g(b)\}$. It comes naturally equipped with a pair of maps $\pi_A : P \to A$ and $\pi_B : P \to B$ with the property that $f\pi_A = g\pi_B$.

The pullback can be computed as the kernel of a related linear map – the map $f \oplus -g : A \oplus B \to X$ maps $(a, b) \mapsto f(a) - g(b)$, and the kernel of $f \oplus -g$ is the pullback P.

The **pushout** of a pair of linear maps $A \xrightarrow{f} X \xleftarrow{g} B$ is the dual construction to the pullback; in practice it is given by the quotient of $A \oplus B$ by the images of $f \oplus -g$: in the pullback, we impose the smallest equivalence relation such that $f(x) = g(x)$ for any $x \in X$.

We say a zig-zag module is *normalized* if all arrows alternate in direction. Any zig-zag module can be transformed into a normalized zig-zag module by introducing copies of modules equipped with identity maps in the appropriate directions. This adds at most a factor of 2 to the number of spaces in the zig-zag. We assume w.l.o.g. we are given a normalized zig-zag.

Suppose $Z : Z_0 \to Z_{01} \leftarrow Z_1 \to Z_{12} \leftarrow \cdots \leftarrow Z_n$ is a finite length normalized zig-zag module[2]. Any barcodes between points Z_i and Z_j can be computed using only pullbacks and localized computations.

As a first step, the homology is computed separately at each of the points $Z_0, Z_{01}, Z_1, \ldots$. The induced maps between the homology modules also are computed. The computation at each node takes $O(m^\omega)$, where m is the size of the complex at each point.

Next, the pullback $P_{i,i+1}$ of the induced maps on homology are computed – independently – for each of the spans $H_*Z_i \to H_*Z_{i,i+1} \leftarrow H_*Z_{i+1}$. Each such pullback, after the computation, represents all classes in Z_i and in Z_{i+1} that mapped to homologous classes in $Z_{i,i+1}$. In other words, each basis element in P_{01} corresponds to some homology class that has a bar including at least the interval (Z_i, Z_{i+1}), but potentially more.

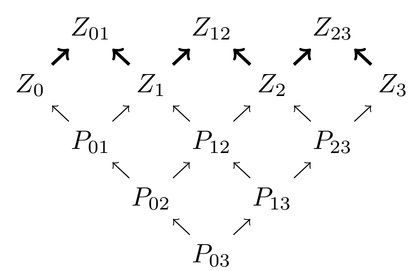

Figure 5: The pullback zig-zag computational scheme

These pullbacks are repeated – and in each new layer, pullbacks represent bars in the zig-zag barcode that extend one step further: a basis element in P_{02} contains elements that persist at least from Z_0 to Z_2, possibly all the way to Z_3.

Finally, representatives for the exact zig-zag barcode may be computed by a sequence of cokernel computations: P_{0n} contains representatives for all bars that extend all the way from Z_0 to Z_n. The cokernels of $P_{0n} \to P_{0,n-1}$ and of $P_{0n} \to P_{1n}$ correspond to bars that only extended from Z_0 to Z_{n-1} and from Z_1 to Z_n respectively – the elements that were present due to a bar existing in P_{0n} have been removed by the cokernel computation. Similarly, at any stage, the cokernel of the sum of the incoming maps removes any bars that are also present at longer stages, leaving the essential bars of each extent in place.

4.1.1 Half-step intervals

Careful examination shows that although we can recover the rank of any bar in the diagram, not all the bars are represented. In particular, "long" bars which end at a $Z_{i,i+1}$, can not be distinguished from ones which end at a Z_i.

When the $Z_{i,i+1}$ are merely a way of connecting different spaces – such as the levelset zigzag or the union zigzag – this is unimportant. However, there is a simple way to recover the missing information, should we need it.

We augment the initial zig-zag quiver by introducing the images of the maps between the spaces. Let f_i denote $\mathrm{im}(Z_i \to Z_{i,i+1})$, g_j denote $\mathrm{im}(Z_j \to Z_{j-1,j})$ and $\mathrm{po}(f, g)$ as the pushout of maps f and g. We

[2]Note that the number of spaces is actually $2n$

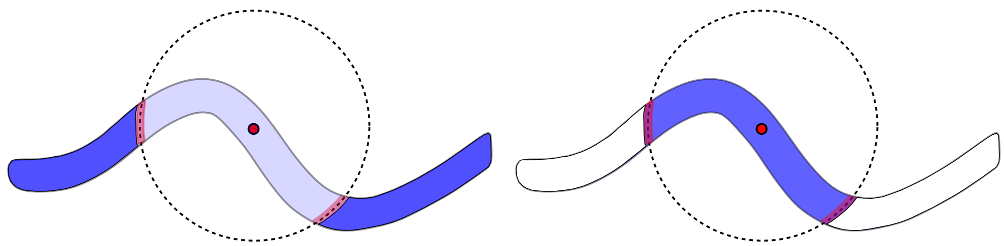

Figure 6: Local homology as the direct limit, $\lim_{r\to 0} \mathrm{H}(\mathbb{X}, \mathbb{X} \cap B^r)$ (left) or $\lim_{r\to 0} \mathrm{H}(\mathbb{X} \cap B_r, \mathbb{X} \cap \partial B_r)$ (right).

can construct a new zig-zag quiver of the form[3] $f_i \to \mathrm{po}(f_i, g_{i+1}) \leftarrow g_{i+1}$. From the decomposition of **this** quiver, in addition to the ranks of the individual spaces, we can recover the missing information. The computation has four phases:

1. compute homology locally. This step is embarrassingly parallel.
2. compute induced maps locally. The induced map $H_* Z_i \to H_* Z_{i,i+1}$ depends only on the two homology groups, all such maps can be computed in parallel.
3. compute pullbacks of all induced maps. This step has dependencies upwards in the graph – P_{ab} depends on $P_{a,b-1}$ and $P_{a+1,b}$; where P_{aa} represents a space with maps from step 2 above.
4. Finally we compute cokernels of the computed maps: this step for P_{ab} only depends on the two covering pullbacks $P_{a-1,b}$ and $P_{a,b+1}$.

4.2 Local Homology

One view of sheaves is to attach an algebraic object to an open set such that it agrees on intersections of open sets. This allows us to "glue" together different pieces in a unique way to get a global picture. Here we present one such structure which we will explore for use in stratification learning.

In stratification learning (or mixed manifold learning), a point cloud is assumed to be sampled from a mixture of (possibly intersecting) manifolds. The objective is to recover the different pieces, often treated as clusters, of the data associated with different manifolds of varying dimensions. Stratified spaces has been studied extensively in mathematics, see the seminal work in [121, 126].

The *local homology groups* at a point $x \in \mathbb{X}$ are defined as the relative homology groups $\mathrm{H}(\mathbb{X}, \mathbb{X} - x)$ ([124], page 126). In this paper, we assume that the topological space $\mathbb{X}$ is embedded in some Euclidean space $\mathbb{R}^d$ [4]. Let $d_x : \mathbb{R}^d \to \mathbb{R}$ be the Euclidean distance function from a fixed $x \in \mathbb{X}$, $d_x(y) := d(x, y) = ||y - x||$. Let $B_r = B_r(x) = d_x^{-1}[0, r]$ and $B^r = B^r(x) = d_x^{-1}[r, \infty)$ be the sublevel sets and superlevel sets of d_x. Taking a small enough r, the local homology groups in questions are in fact the *direct limit* of relative homology groups, $\lim_{r\to 0} \mathrm{H}(\mathbb{X}, \mathbb{X} \cap B^r)$, or alternatively $\lim_{r\to 0} \mathrm{H}(\mathbb{X} \cap B_r, \mathbb{X} \cap \partial B_r)$[109], see Fig. 6.

For a fixed $\alpha \geqslant 0$, let $\mathbb{X}_\alpha$ be the "thickened" or "offset" version of $\mathbb{X}$, that is, the space of points in $\mathbb{R}^d$ at Euclidean distance at most α from $\mathbb{X}$. Suppose L is a finite set of points sampled from $\mathbb{X}$ [5], where $L \subset \mathbb{X}$ and $L_\alpha = \cup_{x \in L} B_\alpha(x)$. In subsequent sections, we put further restrictions on L where we suppose L is an ϵ*-sample* of $\mathbb{X}$, that is,

$$\forall x \in \mathbb{X}, \quad d(x, L) := \inf_{y \in L} d(x, y) \leqslant \epsilon.$$

The persistence module representing this local homology can then be attached to an open set forming a sheaf.

[3] Some details as to the procedure at the ends is omitted due to space constraints.

[4] This assumption can be relaxed in several ways, but this setting is most common in our applications.

[5] Our results would hold with minor modifications in the setting of sampling with noise, where elements of L lie on or near $\mathbb{X}$.

The α-*filtration* [6] (Fig. 7) is a sequence of relative homology groups connected by inclusion, constructed by fixing r and varying α, for $\alpha < \alpha'$,

$$\cdots \to \mathrm{H}(\mathbb{X}_\alpha \cap B_r, \mathbb{X}_\alpha \cap \partial B_r) \to \dots$$
$$\to \mathrm{H}(\mathbb{X}_{\alpha'} \cap B_r, \mathbb{X}_{\alpha'} \cap \partial B_r) \to \cdots .$$

Its discrete counterpart built on a set of points L sampled from $\mathbb{X}$ is,

$$\cdots \to \mathrm{H}(\mathbb{L}_\alpha \cap B_r, \mathbb{L}_\alpha \cap \partial B_r) \to \dots$$
$$\to \mathrm{H}(\mathbb{L}_{\alpha'} \cap B_r, \mathbb{L}_{\alpha'} \cap \partial B_r) \to \cdots .$$

Here, we fix the size of the ball which defines the locality, i.e. the size r of the local neighborhood, and we vary the scale α at which we analyze the

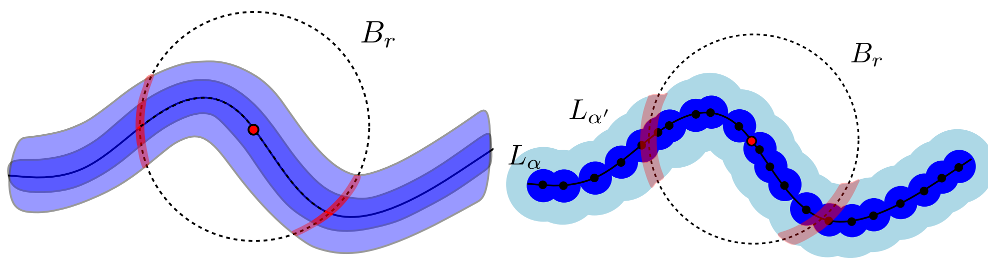

Figure 7: The α-filtration for space $\mathbb{X}$ and its offset (left), and on the right, the same filtration built on a set of points L, sampled from $\mathbb{X}$.

Čech and Vietoris-Rips Complexes. Suppose L is a finite point set in $\mathbb{R}^d$ and $L_\alpha = \cup_{x \in L} B_\alpha(x)$. The nerve of L_α is the simplicial complex induced by all the non-empty intersections of subcollections of balls in L_α and is called the *Čech complex* of L, denoted as $\mathcal{C}_\alpha = \mathcal{C}_\alpha(L)$ (omitting L from the notation unless necessary). The *Vietoris-Rips complex* of L is denoted as $\mathcal{R}_\alpha$, whose simplices correspond to non-empty subsets of L of diameter less than α. For Euclidean metric space, we have, $\forall \alpha > 0$, $\mathcal{C}_{\alpha/2} \subseteq \mathcal{R}_\alpha \subseteq \mathcal{C}_\alpha \subseteq \mathcal{R}_{2\alpha}$ [7]. This implies that the persistence modules $\{\mathrm{H}(\mathcal{C}_\alpha)\}_\alpha$ and $\{\mathrm{H}(\mathcal{R}_\alpha)\}_\alpha$ are α-interleaved [8].

4.2.1 Approximating the α-Filtration

In the α-filtration, since we will be computing relative persistent homology, there are certain requirements on the pairs, such that the maps of the relative filtration are well-defined. Two persistence modules, $\mathcal{A} = \{A_\alpha\}_{\alpha \in \mathbb{R}}$ and $\mathcal{F} = \{F_\alpha\}_{\alpha \in \mathbb{R}}$ are called *compatible* if for all $\alpha \leqslant \beta$, the following diagram commutes:

$$\begin{array}{ccc} A_\alpha & \longrightarrow & F_\alpha \\ \downarrow & & \downarrow \\ A_\beta & \longrightarrow & F_\beta. \end{array}$$

This ensures that the relative persistence module is well-defined. In our context, all the maps are induced by inclusions hence the above diagram commutes. We highlight steps involved to obtain our approximation results:

[6] Technically, the α-filtration is a persistence module that arise from their corresponding filtrations, we refer to them as such for simplicity.

[7] Jung's Theorem gives a tighter relation between Vietoris-Rips and Čech complexes. We use the slightly looser relation in our paper for simplicity.

[8] We emphasize that for our results the interleaving parameter does depend on the parameter of the filtration. In other words, for a fixed filtration scale parameter, we have a certain "constant" interleaving.

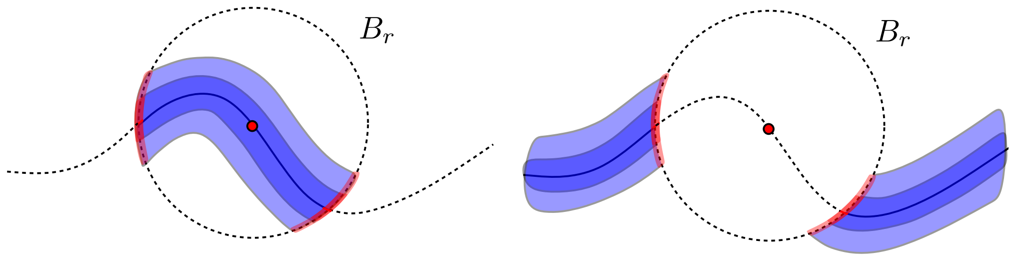

Figure 8: Left: the α-filtration with respect to the pair $(\mathbb{X}_\alpha \cap B_r, \mathbb{X}_\alpha \cap \partial B_r)$. Right: the filtration with respect to the pair $(\mathbb{X}_\alpha, \mathbb{X}_\alpha - \text{int}B_r)$.

- First, we show that under certain conditions, the relative homology of a ball modulo its boundary is isomorphic to that of the entire space modulo the subspace outside the ball.
- Second, we prove that if we have two compatible persistence modules $\mathcal{F}$ and $\mathcal{A}$ which are respectively interleaved with $\mathcal{G}$ and $\mathcal{B}$, the relative persistent homology $\mathrm{H}(\mathcal{F}, \mathcal{A})$ is approximated by $\mathrm{H}(\mathcal{G}, \mathcal{B})$. This result may be of independent interest.
- Last, we prove a series of interleavings to show that both filtrations in our case can be interleaved with a Vietoris-Rips construction on the samples.

We first show that the following two filtrations are equivalent (where $\alpha < \alpha'$):

$$0 \to \mathrm{H}(\mathbb{X}_\alpha \cap B_r, \mathbb{X}_\alpha \cap \partial B_r) \to \mathrm{H}(\mathbb{X}_{\alpha'} \cap B_r, \mathbb{X}_{\alpha'} \cap \partial B_r) \to \ldots \to \mathrm{H}(B_r, \partial B_r), \tag{1}$$

$$0 \to \mathrm{H}(\mathbb{X}_\alpha, \mathbb{X}_\alpha - \text{int}B_r) \to \mathrm{H}(\mathbb{X}_{\alpha'}, \mathbb{X}_{\alpha'} - \text{int}B_r) \to \ldots\ldots \to \mathrm{H}(\mathbb{R}^n, \mathbb{R}^n - \text{int}B_r). \tag{2}$$

Note that $\mathbb{X}_\alpha - \text{int}B_r = \mathbb{X}_\alpha - (\mathbb{X}_\alpha \cap \text{int}B_r)$. Unless otherwise specified, $\alpha, \alpha' \in [0, \infty)$. Graphically, these filtrations are shown in Fig. 8. As it turns out, it is easier to argue about the filtration in Fig. 8(right) than Fig. 8(left), as shown in the following lemma.

Lemma 4.1. *Assuming that spaces $\mathbb{X}_\alpha$ and $\mathbb{X}_\alpha - \text{int}B_r$ form a* good pair*, then $\mathrm{H}(\mathbb{X}_\alpha \cap B_r, \mathbb{X}_\alpha \cap \partial B_r) \cong \mathrm{H}(\mathbb{X}_\alpha, \mathbb{X}_\alpha - \text{int}B_r)$.*

This follows from the Excision Theorem ([122], Theorem 15.1, page 82) and the Excision Extension Theorem ([122], Theorem 15.2, page 82).

We now show that we can approximate local homology at multi-scale via the α-filtration using sample points. We begin with sequence (2). Specifically, we first consider the persistence module corresponding to the whole space $\{\mathbb{X}_\alpha\}$, and then the persistence module corresponding to the subspace we quotient by, $\{\mathbb{X}_\alpha - \text{int}B_r\}$. The key is a technical result described in Theorem 4.2 which says that if we can interleave persistence modules independently, we can interleave their corresponding quotient persistence modules. We consider this result to be of independent interest.

Theorem 4.2. *Consider two pairs of compatible persistence modules. Let $\mathcal{A} = \{A_\alpha\}_{\alpha\in\mathbb{R}}$ be compatible with $\mathcal{F} = \{F_\alpha\}_{\alpha\in\mathbb{R}}$ and $\mathcal{B} = \{B_\alpha\}_{\alpha\in\mathbb{R}}$ be compatible with $\mathcal{G} = \{G_\alpha\}_{\alpha\in\mathbb{R}}$. If the modules $\mathcal{A}$ and $\mathcal{B}$ are ϵ_1-interleaved and $\mathcal{F}$ and $\mathcal{G}$ are ϵ_2-interleaved, then the relative modules $\{(F_\alpha, A_\alpha)\}_{\alpha\in\mathbb{R}}$ and $\{(G_\alpha, B_\alpha)\}_{\alpha\in\mathbb{R}}$ are ϵ-interleaved, where $\epsilon = \max\{\epsilon_1, \epsilon_2\}$.*

Without loss of generality, assume $\epsilon_1 = \epsilon_2 = \epsilon$. Each pair, $\{(F, A)\}$ and $\{(G, B)\}$, gives rise to a long exact sequence. The two sequences are related by interleaving maps yielding the commutative diagram in Figure 9.

To prove that the interleavings between individual modules imply an interleaving between $\{(F, A)\}$ and $\{(G, B)\}$, we would need some careful diagram chasing at the chain level. That is, we need to prove each

$$
\begin{array}{ccccccccc}
\mathrm{H}_n(A_\alpha) & \xrightarrow{i_n^\alpha} & \mathrm{H}_n(F_\alpha) & \xrightarrow{j_n^\alpha} & \mathrm{H}_n(F_\alpha, A_\alpha) & \xrightarrow{k_n^\alpha} & \mathrm{H}_{n-1}(A_\alpha) & \xrightarrow{i_{n-1}^\alpha} & \mathrm{H}_{n-1}(F_\alpha) \\
\downarrow{\scriptstyle \phi_n^\alpha} & & \downarrow{\scriptstyle f_n^\alpha} & & \downarrow{\scriptstyle \mu_n^\alpha} & & \downarrow{\scriptstyle \phi_{n-1}^\alpha} & & \downarrow{\scriptstyle f_{n-1}^\alpha} \\
\mathrm{H}_n(B_{\alpha+\epsilon}) & \xrightarrow{p_n^{\alpha+\epsilon}} & \mathrm{H}_n(G_{\alpha+\epsilon}) & \xrightarrow{q_n^{\alpha+\epsilon}} & \mathrm{H}_n(G_{\alpha+\epsilon}, B_{\alpha+\epsilon}) & \xrightarrow{r_n^{\alpha+\epsilon}} & \mathrm{H}_{n-1}(B_{\alpha+\epsilon}) & \xrightarrow{p_{n-1}^{\alpha+\epsilon}} & \mathrm{H}_{n-1}(G_{\alpha+\epsilon}) \\
\downarrow{\scriptstyle \psi_n^{\alpha+\epsilon}} & & \downarrow{\scriptstyle g_n^{\alpha+\epsilon}} & & \downarrow{\scriptstyle \nu_n^{\alpha+\epsilon}} & & \downarrow{\scriptstyle \psi_{n-1}^{\alpha+\epsilon}} & & \downarrow{\scriptstyle g_{n-1}^{\alpha+\epsilon}} \\
\mathrm{H}_n(A_{\alpha+2\epsilon}) & \xrightarrow{i_n^{\alpha+2\epsilon}} & \mathrm{H}_n(F_{\alpha+2\epsilon}) & \xrightarrow{j_n^{\alpha+2\epsilon}} & \mathrm{H}_n(F_{\alpha+2\epsilon}, A_{\alpha+2\epsilon}) & \xrightarrow{k_n^{\alpha+2\epsilon}} & \mathrm{H}_{n-1}(A_{\alpha+2\epsilon}) & \xrightarrow{i_{n-1}^{\alpha+2\epsilon}} & \mathrm{H}_{n-1}(F_{\alpha+2\epsilon})
\end{array}
$$

Figure 9: Commuting diagrams for the long exact sequence involving two pairs of filtrations.

of the four diagrams needed for interleaving commutes, i.e. diagrams in Fig. 10 commute. The key issue is that although each row is exact, maps between persistence modules do not split — therefore we may have one persistent relative class without a persistent class in either component filtrations.

$$
\begin{array}{ccc}
\mathrm{H}_n(F_\alpha, A_\alpha) & \longrightarrow & \mathrm{H}_n(F_{\alpha+2\epsilon}, A_{\alpha+2\epsilon}) \\
\searrow & & \nearrow \\
& \mathrm{H}_n(G_{\alpha+\epsilon}, B_{\alpha+\epsilon}) &
\end{array}
$$

$$
\begin{array}{ccc}
& \mathrm{H}_n(F_{\alpha+\epsilon}, A_{\alpha+\epsilon}) & \\
\nearrow & & \searrow \\
\mathrm{H}_n(G_\alpha, B_\alpha) & \longrightarrow & \mathrm{H}_n(G_{\alpha+2\epsilon}, B_{\alpha+2\epsilon})
\end{array}
$$

Figure 10: Commuting diagrams for ϵ-interleaved persistence modules.

With Theorem 4.2 in hand, we can begin to prove the main result. We would like to construct a persistence module (based upon Vietoris-Rips filtration) that interleaves with the α-filtration $\{(\mathbb{X}_\alpha, \mathbb{X}_\alpha - \mathrm{int}B_r)\}$. The straightforward approach is to consider $\{(L_\alpha, L_\alpha - \mathrm{int}B_r)\}$, as illustrated in Figure 11. Such a construction is possible with careful geometric considerations through interleaving with an intermediate complex described below. We obtain our main result by proving the following key steps:

1. $\{L_\alpha\}$ and $\{\mathbb{X}_\alpha\}$ are ϵ-interleaved.

2. $\{L_\alpha - \mathrm{int}B_r\}$ and $\{\mathbb{X}_\alpha - \mathrm{int}B_r\}$ are ϵ-interleaved.

3. (Nerve Lemma for $\{L_\alpha - \mathrm{int}B_r\}$) For $\alpha < r$, $\mathcal{N}L_\alpha - \mathrm{int}B_r$ is homotopic to $L_\alpha - \mathrm{int}B_r$.

4. For a fixed α, we give an algorithm that compute the 0- and 1-skeleton of $\mathcal{N}L_\alpha - \mathrm{int}B_r$. Then we build a Vietoris-Rips complex based on these 0- and 1-skeletons by filling in the high-dimensional co-faces. This is in fact a flag complex, or equivalently, the clique complex of the 1-skeleton of $\mathcal{N}L_\alpha - \mathrm{int}B_r$. As α-varies, we refer to such a filtration as $\tilde{\mathcal{R}}_\alpha(L)$. We show that $\{L_\alpha - \mathrm{int}B_r\}$ is α-interleaved with $\tilde{\mathcal{R}}_\alpha(L)$.

5. We now arrive at the main result (Theorem 4.3): $\{(\mathcal{R}_\alpha(\mathbb{L}), \tilde{\mathcal{R}}_\alpha(\mathbb{L}))\}$ is $(\alpha + 2\epsilon)$-interleaved with the α-filtration $\{(\mathbb{X}_\alpha, \mathbb{X}_\alpha - \mathrm{int}B_r)\}$.

Theorem 4.3. *Suppose $\mathbb{L}$ is an ϵ-sample of $\mathbb{X}$. For $\alpha < r$, the relative module $\{(\mathcal{R}_\alpha(\mathbb{L}), \tilde{\mathcal{R}}_\alpha(\mathbb{L}))\}$ is $(\alpha+2\epsilon)$-interleaved with the α-filtration $\{(\mathbb{X}_\alpha, \mathbb{X}_\alpha - \mathrm{int}B_r)\}$.*

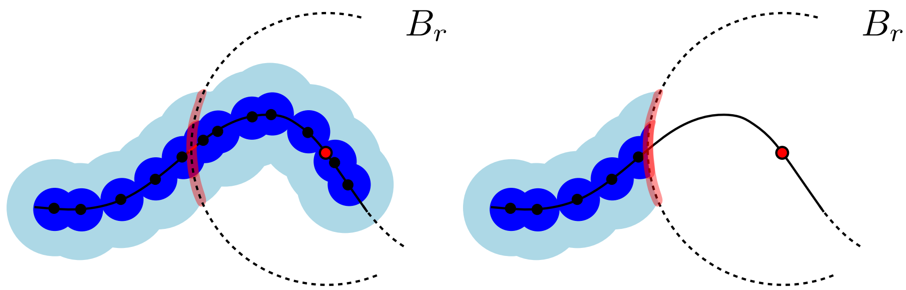

Figure 11: Illustration of filtrations, $\{L_\alpha\}$ (left) and $\{L_\alpha - \mathrm{int} B_r\}$ (right).

5 Generalizations of Persistence

5.1 Towers and Persistence

In this section, we study subcategories defined by paths in a category. Parallel paths in **Vect** and **Mch** will naturally lead to the concept of persistent homology, now within a more general framework than in the traditional literature.

Paths and categories of paths in a category. A *tower* is a path with finitely many non-zero objects in a category. More formally, it consists of objects X_i and arrows $\xi_i : X_i \to X_{i+1}$, for all $i \in \mathbb{Z}$, in which all but a finite number of the X_i are the zero object. We denote this tower as $\mathcal{X} = (X_i, \xi_i)$. In later discussions, we will refer to compositions of the arrows, so we write ξ_i^i for the identity arrow of X_i and define

$$\xi_i^j \quad := \quad \xi_{j-1} \ldots \xi_{i+1} \xi_i : \; X_i \to X_j, \tag{3}$$

for $i < j$. Suppose $\mathcal{Y} = (Y_i, \eta_i)$ is a second tower in the same category, and there is a vector of arrows $\varphi = (\varphi_i)$, with $\varphi_i : X_i \to Y_i$, such that $\eta_i \varphi_i = \varphi_{i+1} \xi_i$ for all i. Referring to this vector of arrows as a *morphism*, we denote this by $\varphi : \mathcal{X} \to \mathcal{Y}$. To verify that the towers and morphisms form a new category, we note that the identity morphism is the vector of identity arrows, and that morphisms compose naturally. The zero object is the tower consisting solely of zero objects. Finally, an *isomorphism* is an invertible morphism; it consists of invertible arrows between objects that commute with the arrows of the towers. We remark that arrows and morphisms are alternative terms for the same notion in category theory. We find it convenient to use both so we can emphasize different levels of the construction.

Persistence in a tower of matchings. As a first concrete case, we consider a tower $\mathcal{A} = (A_i, \alpha_i)$ of matchings. Recall that $\mathsf{rank}\, \alpha_i^j$ is the number of pairs in α_i^j. We formalize this notion by introducing the *rank function* $\mathsf{a} : \mathbb{Z} \times \mathbb{N}_0 \to \mathbb{Z}$ defined by $\mathsf{a}_i^j := \mathsf{a}(i, j-i) := \mathsf{rank}\, \alpha_i^j$, where $\mathbb{N}_0$ is the set of non-negative integers. It can alternatively be understood by counting intervals, as we now explain. Letting A be the disjoint union of the A_i, we call a non-empty partial function $a : \mathbb{Z} \twoheadrightarrow A$ an *interval* in $\mathcal{A}$ if the domain of a is an interval $[k, \ell]$ in $\mathbb{Z}$, $a_i = a(i)$ belongs to A_i for every $k \leqslant i \leqslant \ell$, and $a_{i+1} = \alpha_i(a_i)$ whenever $k \leqslant i < \ell$; see Figure 12. The interval is *maximal* if it cannot be extended at either end. Specifically, a is maximal iff $a_k \notin \mathsf{im}\, \alpha_{k-1}$ and $a_\ell \in \ker \alpha_\ell$. Finally, by a *persistence interval*[9] we mean the domain of a maximal interval in $\mathcal{A}$. It should be clear that a_k^ℓ is the number of intervals in $\mathcal{A}$ with domain $[k, \ell]$. Equivalently, it is the number of maximal intervals with domains that contain $[k, \ell]$. We can also reverse the relationship and compute the number

[9] We note a difference in convention to most of the related literature, in which persistence intervals are defined half-open. In particular, $[k, \ell]$ in our notation corresponds to $[k, \ell + 1)$ in [48]. Reading them as intervals in $\mathbb{Z}$, they are the same.

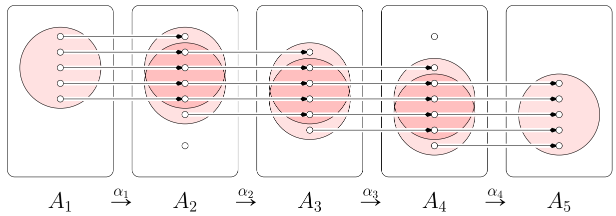

Figure 12: A tower of matchings with maximal intervals $[1,2]$, $[1,3]$, $[1,4]$, twice $[1,5]$, $[2,2]$, $[2,5]$, $[3,5]$, $[4,4]$, $[4,5]$.

of persistence intervals from the rank function. Letting $\#_{[i,j]} = \#_{[i,j]}(\mathcal{A})$ denote the number of maximal intervals with domain $[i,j]$, we have

$$\#_{[i,j]} \quad = \quad \mathsf{a}_i^j - \mathsf{a}_{i-1}^j - \mathsf{a}_i^{j+1} + \mathsf{a}_{i-1}^{j+1}. \tag{4}$$

To see this, we consider again A and identify elements if they belong to a common pair in any of the α_i. After identification, every point a in A corresponds to a maximal interval. The domain of this maximal interval is $[i,j]$ iff a belongs to the domain of α_i^j but not to the image of α_{i-1} and not to the domain of α_i^{j+1}. The relation (4) now follows from $\alpha = \alpha_{i-1}$, $\beta = \alpha_i^j$, $\gamma = \alpha_j$.

This relationship motivates us to introduce the *persistence diagram* of $\mathcal{A}$ as the multiset of persistence intervals, which we denote as $\mathrm{Dgm}(\mathcal{A})$. Note that $\#_{[i,j]}$ is the multiplicity of $[i,j]$. The number of intervals in the persistence diagram, counted with multiplicities, is therefore $\#\mathrm{Dgm}(\mathcal{A}) = \sum_{i \leqslant j} \#_{[i,j]}$. It is important to observe that the persistence diagram characterizes the tower up to isomorphisms.

1 (Equivalence Theorem A)**.** *Letting $\mathcal{A}$ and $\mathcal{B}$ be towers in **Mch**, the following conditions are equivalent:*

- *(i) $\mathcal{A}$ and $\mathcal{B}$ are isomorphic;*
- *(ii) the rank functions of $\mathcal{A}$ and $\mathcal{B}$ coincide;*
- *(iii) the persistence diagrams of $\mathcal{A}$ and $\mathcal{B}$ are the same.*

Proof. (i) $\Rightarrow$ (ii). Since $\mathcal{A}$ and $\mathcal{B}$ are isomorphic, we have invertible arrows $\theta_i : A_i \to B_i$ that commute with the arrows in $\mathcal{A}$ and $\mathcal{B}$. It follows that $\mathsf{rank}\,\alpha_i^j = \mathsf{rank}\,\beta_i^j$, for all $i \leqslant j$.

(ii) $\Rightarrow$ (iii). The rank function determines the multiplicities of the intervals in the persistence diagram by (4).

(iii) $\Rightarrow$ (i). We use induction over the cardinality of the common diagram. If the cardinality is zero, then both towers consist of empty matchings only, so they are isomorphic. Assuming the implication for cardinality at most n, we consider $\#\mathrm{Dgm}(\mathcal{A}) = \#\mathrm{Dgm}(\mathcal{B}) = n+1$. Picking up a persistence interval in the common diagram and removing its images in $\mathcal{A}$ and $\mathcal{B}$, we obtain towers $\mathcal{A}'$, $\mathcal{B}'$ with coinciding persistence diagrams of cardinality at most n. Thus, the towers are isomorphic. The extension of the isomorphisms of $\mathcal{A}'$ and $\mathcal{B}'$ to isomorphisms of $\mathcal{A}$ and $\mathcal{B}$ is now straightforward.

Persistence in a tower of linear maps. We return to assuming a fixed field, and consider a tower $\mathcal{U} = (U_i, \upsilon_i)$ in the category of vector spaces over this field.[10] For each i, let A_i be a basis of U_i. Restricting υ_i to A_i and A_{i+1}, we get a partial function $\alpha_i : A_i \twoheadrightarrow A_{i+1}$, again for every i.

2 (Definition)**.** *We call the tower of partial functions $\mathcal{A} = (A_i, \alpha_i)$ a* basis *of the tower $\mathcal{U}$ if α_i is a matching and $\mathsf{rank}\,\alpha_i = \mathsf{rank}\,\upsilon_i$, for every i.*

[10] Part of the theory in this section can be developed for the more general case of finitely generated modules over a principal ideal domain. For reasons of simplicity, and because the crucial connection to matchings relies on stronger algebraic properties, we limit this discussion to vector spaces over a field right from the start.

For each $\upsilon_i : U_i \to U_{i+1}$, there are bases A_i and A_{i+1} of the two vector spaces such that the implied partial function $\alpha_i : A_i \nrightarrow A_{i+1}$ is a matching that satisfies $\mathsf{rank}\,\alpha_i = \mathsf{rank}\,\upsilon_i$; We will show shortly that such bases exist for all vector spaces in the tower simultaneously; see the Basis Lemma below. For now, we assume that $\mathcal{A} = (A_i, \alpha_i)$ is such a basis, deferring the proof to later. Considering compositions α_i^j, we note that $\mathsf{rank}\,\alpha_i^j = \mathsf{rank}\,\upsilon_i^j$. Consequently, the rank functions of $\mathcal{A}$ and $\mathcal{U}$ are the same. The basis of $\mathcal{U}$ is not necessarily unique, but the rank function does not depend on the choice. Thus, we can define the *persistence diagram* of $\mathcal{U}$ as the persistence diagram of a basis, $\mathrm{Dgm}(\mathcal{U}) := \mathrm{Dgm}(\mathcal{A})$. We also write $\#_{[i,j]}(\mathcal{U}) := \#_{[i,j]}(\mathcal{A})$ for the multiplicity of the interval $[i,j]$ in the persistence diagram of $\mathcal{U}$. Writing u_i^j for the rank of υ_i^j, we thus get

$$\#_{[i,j]}(\mathcal{U}) \quad = \quad \mathsf{u}_i^j - \mathsf{u}_{i-1}^j - \mathsf{u}_i^{j+1} + \mathsf{u}_{i-1}^{j+1} \tag{5}$$

from (4). Similarly, we can generalize the Equivalence Theorem A to the case of linear maps between finite-dimensional vector spaces.

3 (Equivalence Theorem B)**.** *Letting $\mathcal{U}$ and $\mathcal{V}$ be towers in* ***Vect****, the following conditions are equivalent:*

(i) $\mathcal{U}$ and $\mathcal{V}$ are isomorphic;
(ii) the rank functions of $\mathcal{U}$ and $\mathcal{V}$ coincide;
(iii) the persistence diagrams of $\mathcal{U}$ and $\mathcal{V}$ are the same.

Tower bases. We now prove the technical result assumed above to get Equivalence Theorem B.

4 (Basis Lemma)**.** *Every tower in* ***Vect*** *has a basis.*

Proof. We construct the basis in two phases of an algorithm, as sketched in Figure 13. Let $\mathcal{U} = (U_i, \upsilon_i)$ be a tower in **Vect**, and for each i, let M_i be the matrix that represents the map υ_i in terms of the given bases of U_i and U_{i+1}.

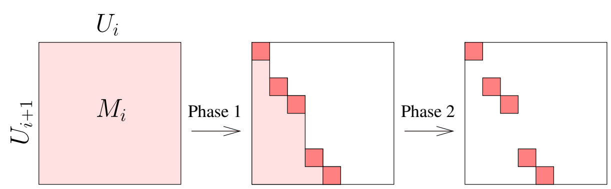

Figure 13: Two-phase reduction of the matrix. The shaded areas contain zeros and non-zeros, the white areas contain only zeros, and all dark squares are 1.

In Phase 1, we use column operations to turn M_i into column echelon form, as sketched in Figure 13 in the middle. We get a strictly descending staircase of non-zero entries, with zeros and non-zeros below and zeros above the staircase. Here, we call the collection of topmost non-zero entries in the columns the *staircase*, and we multiply with inverses so that all entries in the staircase are equal to 1. By definition, each column contains at most one element of the staircase, and by construction, each row contains at most one element of the staircase. The reduction to echelon form is done from right to left in the sequence of matrices; that is, in the order of decreasing index i. Indeed, every column operation in M_i changes the basis of U_i, so we need to follow up with the corresponding row operation in M_{i-1}. Since M_{i-1} has not yet been transformed to echelon form, there is nothing else to do.

In Phase 2, we use row operations to turn the column echelon into the normal form, as sketched in Figure 13 on the right. Equivalently, we preserve the staircase and turn all non-zero entries below it into zeros. To do this for a single column, we add multiples of the row of its staircase element to lower rows. Processing the columns from left to right, this requires no backtracking. The reduction to normal form is done from left to right in the sequence of matrices; that is, in the order of increasing index i. Each row operation in M_i changes the basis of U_{i+1}, so we need to follow up with the corresponding column operation in M_{i+1}.

This operation is a right-to-left column addition, which preserves the echelon form. Since M_{i+1} has not yet been transformed to normal form, there is nothing else to do.

In summary, we have an algorithm that turns each matrix M_i into a matrix in which every row and every column contains at most one non-zero element, which is 1. This is the matrix of a matching. Since we use only row and column operations, the ranks of the matrices are the same as at the beginning. Each column operation in M_i has a corresponding operation on the basis of U_i. Similarly, each row operation in M_i has a corresponding operation on the basis of U_{i+1}. By performing these operations on the bases of the vector spaces, we arrive at a basis of the tower.

Persistent homology and derivations. Persistent homology is a special case of the persistence of towers of vector spaces. To see this, let $\mathcal{C} = (C_i, \gamma_i)$ be a tower of chain complexes, with γ_i the inclusion of C_i in C_{i+1}, and obtain $\mathcal{H} = (H_i, \eta_i)$ by applying the homology functor. Assuming coefficients in a field, the latter is a tower of vector spaces. The *persistent homology groups* are the images of the η_i^j. Zig-zag persistent homology studies a sequence of vector spaces U_i and linear maps, some of which go forward, from U_i to U_{i+1}, while others go backward, from U_{i+1} to U_i. This is a *zigzag module* if we have exactly one map between any two contiguous vector spaces. It turns out that the theory of persistence generalizes to this setting. In view of our approach based on matchings, this is not surprising. Indeed, the inverse of a matching is again a matching, so that there is no difference at all in the category of matchings. To achieve the same in the category of vector spaces, we only need to adapt the above algorithm to obtain the zigzag generalization of the Basis Lemma. The adaptation is also straightforward, running the algorithm on a sequence of matrices that are the original matrices for the forward maps and the transposed matrices for the backward maps.

There are several ways one can derive towers from towers, and we discuss some of them. Letting $\mathcal{U} = (U_i, \upsilon_i)$ and $\mathcal{V} = (V_i, \nu_i)$ be towers in **Vect**, we call $\mathcal{V}$ a *subtower* of $\mathcal{U}$ if $V_i \subseteq U_i$ and ν_i is the restriction of υ_i to V_i and V_{i+1}, for each i. Given $\mathcal{U}$ and a subtower $\mathcal{V}$, we can take quotients and define the *quotient tower*, $\mathcal{U}/\mathcal{V} = (U_i/V_i, \varrho_i)$, where ϱ_i is the induced map from U_i/V_i to U_{i+1}/V_{i+1}. Similarly, we can construct towers from a morphism $\varphi : \mathcal{U} \to \mathcal{V}$, where we no longer assume that $\mathcal{V}$ is a subtower of $\mathcal{U}$. Taking kernels and images, we get the *tower of kernels*, which is a subtower of $\mathcal{U}$, and the *tower of images*, which is a subtower of $\mathcal{V}$. Taking the quotients, $U_i/\mathsf{ker}\,\varphi_i$ and $V_i/\mathsf{im}\,\varphi_i$, we furthermore define the *towers of coimages* and *of cokernels*. In [44], towers of kernels are used in the analysis of sampled stratified spaces and introduced along with the towers of images and cokernels. The benefit of the general framework presented in this section is that persistence is now defined for all these towers, without the need to prove or define anything else.

Suppose now that $\mathcal{V} = \mathcal{U}$, and let $\phi : \mathcal{U} \to \mathcal{U}$ be an endomorphism. We can iterate ϕ and thus obtain a sequence of endomorphisms. The *generalized kernel* of the component $\phi_i : U_i \to U_i$ is the union of the kernels of its iterated compositions. Similarly, the *generalized image* is the intersection of the images of the iterated compositions:

$$\mathsf{gker}\,\phi_i \quad := \quad \bigcup_{k=1}^{\infty} \mathsf{ker}\,\phi_i^{\circ k}, \tag{6}$$

$$\mathsf{gim}\,\phi_i \quad := \quad \bigcap_{k=1}^{\infty} \mathsf{im}\,\phi_i^{\circ k}, \tag{7}$$

where $\phi_i^{\circ k}$ is the k-fold composition of ϕ_i with itself. Similar to before, we define two subtowers of $\mathcal{U}$: the *tower of generalized kernels*, denoted as $\mathsf{gker}\,\phi$, and the *tower of generalized images*, denoted as $\mathsf{gim}\,\phi$. By assumption, U_i has finite dimension, which implies that both the generalized kernel and the generalized image are already defined by finite compositions of ϕ_i. Furthermore, the ranks of $\mathsf{gker}\,\phi_i$ and $\mathsf{gim}\,\phi_i$ add up to the rank of U_i. A trivial example of an element in the generalized image is an eigenvector of ϕ_i, but $\mathsf{gim}\,\phi_i$ also contains the eigenvectors of the iterated compositions of ϕ_i and the spaces they span.

Of particular interest are the quotient towers, $\mathcal{U}/\mathsf{gker}\,\phi$ and $\mathcal{U}/\mathsf{gim}\,\phi$, because of their relation to the Leray functor [130] and Conley index theory [131, 129]. They may be of interest in the future study of the persistence of the Conley index applied to sampled dynamical systems.

Tower of eigenspaces. Of particular interest to this paper is the tower of eigenspaces. When studying the eigenvectors of the endomorphisms, we do this for each eigenvalue in turn. To begin, we note that $\phi : \mathcal{U} \to \mathcal{U}$ is a tower in the category **Endo**(**Vect**). Indeed, each $\phi_i : U_i \to U_i$ is an object, and $\upsilon_i : U_i \to U_{i+1}$ commutes with ϕ_i and ϕ_{i+1}. Applying the eigenspace functor, E_t, we get the tower $\mathcal{E}_t(\phi) = (E_t(\phi_i), \delta_{t,i})$ in **Vect**. Its objects are the eigenspaces, $E_t(\phi_i)$, and its arrows are the restrictions, $\delta_{t,i}$, of the υ_i to $E_t(\phi_i)$ and $E_t(\phi_{i+1})$. We refer to it as the *eigenspace tower* of ϕ for eigenvalue t.

Much of the technical challenge we face in this paper derives from the difficulty in constructing linear self-maps from sampled self-maps. This motivates us to extend the above construction to a pair of morphisms. Let $\mathcal{V} = (V_i, \nu_i)$ be a second tower in **Vect**, let $\varphi, \psi : \mathcal{U} \to \mathcal{V}$ be morphisms between the two towers, and recall that this gives a tower in **Pairs**(**Vect**). Its objects are the pairs $\varphi_i, \psi_i : U_i \to V_i$, and its arrows are the commutative diagrams with vertical maps υ_i and ν_i. Similar to the single-map case, we apply the eigenspace functor, E_t, now from **Pairs**(**Vect**) to **Vect**. This gives the tower $\mathcal{E}_t(\varphi, \psi) = (E_t(\varphi_i, \psi_i), \epsilon_{t,i})$ in **Vect**. Its objects are the eigenspaces, and its arrows are the linear maps that map $[u] \in E_t(\varphi_i, \psi_i)$ to $[\upsilon_i(u)] \in E_t(\varphi_{i+1}, \psi_{i+1})$. We refer to it as the *eigenspace tower of the pair* (φ, ψ) for the eigenvalue t. This is the main new tool in our study of self-maps. Of particular interest will be the persistence module of this tower.

5.2 k[t]-modules

Recall that a persistence module can be seen as a graded ring over k$[t]$. We now give algorithms for a host of constructions to be built on top of this algebra. This gives us much greater flexibilty with the type of computations which are feasible.

5.2.1 Presentation of Modules

Suppose M is a finitely presented module over a ring R. Then, M is given by some quotient R^d/K for K a finitely generated submodule of R^d. If R has global dimension at most 1, K is a free module, and thus has a finite basis expressed as elements of R^d.

Hence, to track M as a module, and to compute with elements of M, it is enough to keep track of d and of a basis of K. If the basis of K is maintained as a Grobner basis, then we can compute normal forms for any element of M that have the property that if normal forms of m, n are equal, then the elements are equal in the quotient module.

Quite often, we shall meet constructions where the module M is more naturally expressed as a subquotient of some semantically relevant free module, possibly of higher than necessary rank. In this case, we shall track two sets of data to enable computation in M: a basis of the relations submodule K, and an extension of that basis to a basis of the generators submodule G, resulting in two bases that track the behaviour of $M = G/K$. Again, as long as K has a basis maintained as a Grobner basis, it is enough to compute this normal form to compare elements.

In particular, this is the case where the module M is the homology of some chain complex. There, the free module is the module of chains, and the submodules G and K correspond to the modules of *cycles* and *boundaries* respectively. We observe that this reflects practice among algorithms for persistent homology: to compute the persistent homology with a basis of a filtered simplicial complex, we maintain a cycle basis and a boundary basis, reducing modulo the boundary at each step to find out if the boundary $\partial\sigma$ of a newly introduced simplex σ is already a boundary or not.

5.2.2 Useful Forms

The Smith normal form transforms a matrix representing a map between free modules by basis changes in the source and target modules until the matrix is diagonal; all while respecting the grading of the modules.

Another fundamentally useful form that we will be using a lot is the *Row or Column echelon form* of a matrix. For a matrix F representing a map $f : M \to N$ between free graded k$[t]$-modules, a row echelon form changes basis for N to eliminate all redundancies in the information contained in the matrix, while a column echelon form changes basis for M to achieve the same goal. Certainly, a Smith normal form is an echelon form, but the converse does not necessarily hold true.

The properties most interesting to us for an echelon form – and we shall state these for the row echelon form; column echelon form holds mutatis mutandis — are:

- All completely zero rows are at the bottom of the matrix
- The leading coefficient of a non-zero row (the *pivot* of that row) is strictly to the right of the leading coefficient of the row above it
- All entries in a column below a leading coefficient of some row are all zero.

The power of the echelon forms come in what they imply when you use a matrix to pick out a basis of a submodule: if the basis elements of a submodule of a free module are in the rows of a matrix, and this matrix is placed on row echelon form, then several problems concerning module membership and coefficient choice become easy to work on.

To determine if an element is in the submodule, add the element to the bottom row of the matrix, and put the matrix on row echelon form again. This, by force, will eliminate any entries in pivot columns of the new row; and modify elements to the right of the pivots that can reduce. Eliminating all column entries that occur as pivots of the previous basis puts the new element on a *normal form* with respect to the basis – and this normal form is sufficient to determine equality modulo the submodule represented by the basis.

5.2.3 Two Modules

When given pairs of persistence modules, we must be able to represent the maps between them. This material follows the treatment by [49, Chapter 15], specialized to the case of persistence modules. We begin with presentations for modules P and Q:

$$\begin{array}{ccccccccc}
0 & \longrightarrow & G_P & \xrightarrow{i_P} & F_P & \xrightarrow{p_P} & P & \longrightarrow & 0 \\
 & & \downarrow{\scriptstyle \varphi|_{G_P}} & & \downarrow{\scriptstyle \varphi} & & \downarrow{\scriptstyle f} & & \\
0 & \longrightarrow & G_Q & \xrightarrow{i_Q} & F_Q & \xrightarrow{p_Q} & Q & \longrightarrow & 0
\end{array}$$

Such a presentation exhibits P and Q as quotients of free modules, i.e. $P = F_P/i_P G_P$ and $Q = F_Q/i_Q G_Q$. We call F_P and F_Q modules of generators of P and Q respectively, and G_P and G_Q modules of relations.

Over nice enough rings[11], both G_P and G_Q are free for all modules P. These rings include all Euclidean domains.

Now, an arbitrary map $P \xrightarrow{f} Q$ can be represented by a map from the generators of P to the generators of Q, or in other words a map $F_P \xrightarrow{\varphi} F_Q$. For such a map between the generator modules to represent a map between the quotient modules it needs to obey one condition: $\varphi \cdot i_P G_P \subseteq i_Q \cdot G_Q$. In other words, any vector in F_P that represents a relation in P, and therefore represents the 0-element in P, has to be mapped to a relation in Q so that its image still represent the 0-element. Subject to this condition, any map between the free modules F_P and F_Q will represent a map between P and Q.

By maintaining F_P and G_P we can compute just about anything to do with persistence modules.

5.2.4 Barcodes as module presentations

For the particular case of a persistence barcode, the setting at hand is very specific. The modules are $\mathbb{N}$-graded modules over $\mathrm{k}[t]$, and in the presentation of a persistence module P as

$$0 \to G_P \xrightarrow{i_P} F_P \to P \to 0$$

[11] Global dimension less than or equal to 1, which includes all PIDs and therefore all Euclidean domains

we maintain G_P and F_P as free submodules of a global *chain module*, with a global specific basis given by the simplices in the underlying simplicial complex. In this setting, we further maintain a few invariants, in particular we ensure that G_P and F_P are the results of the appropriate basis changes S and T from the Smith normal form in order to guarantee that i_P is the corresponding diagonal matrix representation of the Smith normal form.

Hence, from an arbitrary presentation of a finitely presented persistence module

$$0 \to G_P \xrightarrow{i_P} F_P \to P \to 0$$

we get the barcode presentation by writing $\iota_P = S i_P T$ and then replacing this presentation above with

$$0 \to T^{-1} G_P \xrightarrow{\iota_P} F_P S^{-1} \to P \to 0$$

One concrete benefit of this approach is that in the new bases for $T^{-1}G_P$ and $F_P S^{-1}$, we can match up basis elements between the relations and the generators into birth/death pairs of barcode intervals, and the corresponding diagonal entries in ι_P are exactly on the form kt^α yielding a bar of length α in the barcode.

5.2.5 Direct Sum

The first construction is the direct sum of P and Q. This is the simplest and most basic construction, but will be used in all the other constructions. We construct it by taking the direct sum across the relations and free part.

$$0 \to G_P \oplus G_Q \xrightarrow{\begin{pmatrix} i_P & 0 \\ 0 & i_Q \end{pmatrix}} F_P \oplus F_Q \xrightarrow{\begin{pmatrix} p_P & 0 \\ 0 & p_Q \end{pmatrix}} P \oplus Q \to 0$$

This is still a short exact sequence since with a direct sum the maps can be composed as a direct sum of the component maps.

In the case where the persistence module is represented as a matrix encoding a cycle basis and another matrix encoding a boundary operator, the new representation is simply a block matrix with the component matrices along the diagonal. Consider the matrix equations

$$R_P = D_P C_P$$
$$R_Q = D_Q C_Q$$

The direct sum is given by

$$\begin{bmatrix} R_P & 0 \\ 0 & R_Q \end{bmatrix} = \begin{bmatrix} D_P & 0 \\ 0 & D_Q \end{bmatrix} \begin{bmatrix} C_P & 0 \\ 0 & C_Q \end{bmatrix} \tag{8}$$

Assuming that the two inclusion maps i_P and i_Q are already on Smith normal form, the Smith normal form of the presentation map $i_P \oplus i_Q$ is given by the block matrix

$$\begin{pmatrix} i_P & 0 \\ 0 & i_Q \end{pmatrix}$$

5.2.6 Image

Given a map, we will want to construct the image of the map between the two modules. This is a relatively straightforward construction. We first compute the image of the map between the free parts: $\operatorname{im} \varphi : \operatorname{im}(F_P \to F_Q)$. The finite presentation is then

$$0 \to G_Q \xrightarrow{i_Q} G_Q + \operatorname{im} \varphi \to \operatorname{im}(P \xrightarrow{f} Q) \to 0$$

Note that the resulting image generators should be reduced with respect to the relations to ensure that they are representative generating elements.

5.2.7 Cokernel

To compute a cokernel, we merely include the images of the generators in F_P among the relations for the cokernel as a quotient module of F_Q.

$$0 \to G_Q \oplus F_P \xrightarrow{i_P+\varphi} F_Q \to \operatorname{coker}(P \xrightarrow{f} Q) \to 0$$

To see why this is correct, first consider the cokernel of the free modules. In this presentation if it is in the image, since we have listed it among the relations, it maps to 0 in the quotient module. One imaginable complication would be if F_P has a basis element representing 0 in P, but mapping to a basis element in F_Q which is not killed by any element in G_Q. For a constellation like this, we would in the presentation above erroneously kill the generator in F_Q by the presence in F_P; but the constellation is impossible, since $\varphi \cdot i_P \cdot G_P \subseteq i_Q \cdot G_Q$.

5.2.8 Kernel

The kernel of a map between finitely presented modules has to be constructed in a two-step process. A generator of the kernel of $P \xrightarrow{\varphi} Q$ is an element of F_P such that its image in F_Q is also the image of something in G_Q. This is to say that the free module of generators of $\ker(P \xrightarrow{\varphi} Q)$ is given by the kernel of the map between free modules $F_P \oplus G_Q \to F_Q$ given by $(f, g) \mapsto (\varphi(f) - i_Q(g))$.

This takes care of the module F_K of generators of $\ker(P \xrightarrow{\varphi} Q)$. For a complete presentation we also need the relations of this module. Since the kernel module is a submodule of P, the relations are the restriction of the relations in P to the kernel K. These can be computed by recognizing that they are given precisely by the elements in F_K such that their image of F_K in F_P coincides with images of elements in G_P. Thus, we setup another map $F_K \oplus G_P \to F_P$, defined by $(f, g) \mapsto f_P - i_P(g)$, and compute the kernel of this map. This kernel is the module of relations we need.

5.2.9 Free Pullback

Before continuing into new constructions, we recount a useful construction on free modules: the free pullback. Given two maps, f and g in the following diagram:

$$\begin{array}{ccc} & & A \\ & & \downarrow f \\ B & \xrightarrow{g} & C \end{array}$$

we construct the pullback P such that the following diagram commutes

$$\begin{array}{ccc} P & \dashrightarrow & A \\ \vdots & & \downarrow f \\ B & \xrightarrow{g} & C \end{array}$$

To compute P, we must set up a kernel computation: namely find a basis of $\ker(f \oplus -g)$.

With this construction we can restate the kernel construction. The generators of the kernel F_K are given by the pullback

$$\begin{array}{ccc} F_K & \overset{\pi_P}{\dashrightarrow} & F_P \\ \vdots\,\pi_Q & & \downarrow \varphi \\ G_Q & \xrightarrow{i_Q} & F_Q \end{array}$$

The relations G_K are given by a second pullback

$$\begin{array}{ccc} G_K & \xrightarrow{i_K} & G_P \\ {\scriptstyle \pi_Q}\downarrow & & \downarrow{\scriptstyle i_P} \\ F_K & \xrightarrow{\pi_P} & F_P \end{array}$$

giving the presentation $G_K \xrightarrow{i_K} F_K \to \ker(\varphi)$.

5.2.10 Pullback

Given module presentations

$$0 \to G_P \xrightarrow{i_P} F_P \xrightarrow{p_P} P \to 0$$
$$0 \to G_Q \xrightarrow{i_Q} F_Q \xrightarrow{p_Q} Q \to 0$$
$$0 \to G_R \xrightarrow{i_R} F_R \xrightarrow{p_R} R \to 0$$

and maps $P \xrightarrow{f} R$ and $Q \xrightarrow{g} R$ given as maps $F_P \xrightarrow{\varphi} F_R$ and $F_Q \xrightarrow{\psi} F_R$ such that $\varphi(G_P) \subseteq G_R$ and $\psi(G_Q) \subseteq G_R$, we can define the *pullback*. This is the submodule of $P \oplus Q$ such that for any element (p, q), $\varphi p = \psi q$.

Pullbacks are useful when dealing with modules – they can express very many algebraic concepts neatly and concisely, such as inverse images of maps, fibre products, kernels, and so on.

To compute the pullback, we set this up analogously as with the free pullback, as a kernel computation. The pullback is a submodule of $P \oplus Q$, and is carved out exactly as the kernel of the map $P \oplus Q \to R$ given by $(p, q) \mapsto \varphi p - \psi q$. Since this is the kernel of a module map, we can compute this using the construction of Section 5.2.9.

$$\begin{array}{ccccccccc} 0 & \longrightarrow & G_P \oplus G_Q & \longrightarrow & F_P \oplus F_Q & \longrightarrow & P \oplus Q & \longrightarrow & 0 \\ & & \downarrow{\scriptstyle \varphi|_{G_P} - \psi|_{G_Q}} & & \downarrow{\scriptstyle \varphi - \psi} & & \downarrow{\scriptstyle f - g} & & \\ 0 & \longrightarrow & G_R & \longrightarrow & F_R & \longrightarrow & R & \longrightarrow & 0 \end{array}$$

Expanding the construction, we compute the following two free pullbacks

$$\begin{array}{ccc} F_{PB} & \xrightarrow{\pi_P} & F_P \\ {\scriptstyle \pi_Q}\downarrow & & \downarrow{\scriptstyle \varphi} \\ F_Q & \xrightarrow{i_Q} & F_R \end{array} \qquad \begin{array}{ccc} G_{PB} & \xrightarrow{\pi_\oplus} & G_P \oplus G_Q \\ {\scriptstyle i_{PB}}\downarrow & & \downarrow{\scriptstyle i_\oplus} \\ F_{PB} & \xrightarrow{\pi_P} & F_P \oplus F_Q \end{array}$$

Here we begin to see that with presentation algorithms can be broken down into a few key constructions.

5.2.11 Pushout

Pushouts are the dual construction to pullbacks. Some of their most important uses is to glue together things that overlap slightly – producing almost but not quite a direct sum.

Given module presentations

$$0 \to G_P \xrightarrow{i_P} F_P \xrightarrow{p_P} P \to 0$$
$$0 \to G_Q \xrightarrow{i_Q} F_Q \xrightarrow{p_Q} Q \to 0$$
$$0 \to G_R \xrightarrow{i_R} F_R \xrightarrow{p_R} R \to 0$$

and maps $R \xrightarrow{f} P$ and $R \xrightarrow{g} Q$ given as maps $F_R \xrightarrow{\varphi} F_P$ and $F_R \xrightarrow{\psi} F_Q$ such that $\varphi(G_R) \subseteq G_P$ and $\psi(G_R) \subseteq G_Q$, we can define the *pushout*.

This is the cokernel of the map $R \to P \times Q$ given by $r \mapsto (\varphi r, -\psi r)$. This way, any image of an element R in either P or Q can "move across the $\times$" in the pushout module $P \times_R Q$.

5.2.12 Tensor products

The tensor product, and its various associated constructions, are easiest to describe if we focus on using the presentation map $G \to F$ instead of explicit chain representations for the relations. We assume that readers have seen tensor products at least in the context of vector spaces and we recommend [49] as a reference.

We fix a basis B_P for F_P and B_Q for F_Q. A basis for $M \otimes N$ is given by $B_P \otimes B_Q$.

Relations are all generated by elements on the form $r \otimes g$ or $g \otimes r$ where r is a relation and g is a generator.

Theorem 5.1. *Assume that the underlying coefficient ring is a graded PID.*

Suppose P is presented by $i : G_P \to F_P$ given on Smith normal form, with B_P the corresponding basis of F_P, and Q is presented by $j : G_Q \to F_Q$ also given on Smith normal form, with B_Q its corresponding basis of F_Q.

Then a Smith normal form presentation of $P \otimes Q$ is given by a basis $B_P \times B_Q$ for the generating module, and a basis element in the relations module for each $p \otimes q$ for $p \in B_P, q \in B_Q$ where at least one of αp and βq is a relation. Let γ be the generator of the ideal $\langle \alpha, \beta \rangle$. Then $\gamma p \otimes q$ is the sole relation influencing $p \otimes q$ in the Smith normal form presentation of $P \otimes Q$.

Proof. First, if all relations for both P and Q are given as Smith normal forms, any relation on the form $\alpha p \otimes q$ for p, q images of the relations basis are going to be products of basis elements; and thus not expand bi-linearly. Hence, each relation induced by the relations construction above is already a coefficient times a basis element.

To ascertain a graded Smith normal form, we also need to verify that each basis element occurs in only one relation. This is not immediately guaranteed from the construction above – there is one small correction step needed. Certainly, if $\alpha p \in iG_P$ and $\beta q \in jG_Q$, then $\gamma p \otimes q$ has to be a relation. However, if both $\alpha \neq 0$ and $\beta q \neq 0$, then both $\alpha p \otimes q$ and $\beta p \otimes q$ will show up from the construction above; for these cases, replacing the two candidates by their common generator $\gamma p \otimes q$ is needed to construct a Smith normal form presentation. ☐

We may note that the tensor product of the presentation is longer than the presentation suggested above: it would be

$$0 \to G_P \otimes G_Q \to G_P \otimes F_Q \oplus F_P \otimes G_Q \to F_P \otimes F_Q \to P \otimes Q \to 0$$

This is a *free resolution* of $P \otimes Q$; but not a minimal one — the method in the proof of Theorem 5.1 tells us exactly how to reduce away the redundancy represented by the syzygies in $G_P \otimes G_Q$ to get a minimal presentation — of length corresponding to the homological dimension of the module category.

For an arbitrary size tensor product $P_1 \otimes \cdots \otimes P_n$ with P_j having relations G_j, generators F_j, Smith normal form presentation map i_j, and a basis of F_j denoted by B_j, a basis for the generators is given by n-tuples in $B_1 \times \cdots \times B_n$, and relations are given from the generator of the ideal generated by all the coefficients of relations of factor basis elements of the relation.

We can summarize the algorithm above as follows:

Tensor Product Algorithm

Input: Two module presentations P and Q with reduced bases in Smith Normal Form

1. The generators are the tensor product of the generators: $F_{P \otimes Q} = F_P \otimes F_Q$ — generating elements are of the form $(f_P(i), f_Q(j))$ for all generator elements in P and Q
2. Create a non-minimal list of relations by pairing relations and generators — $G_{P \otimes Q} = F_P \otimes G_Q \oplus F_Q \otimes G_P$

3. Create a minimal representation of relations — reduce the relations modulo $G_P \otimes G_Q$ — consider the pairs (g_p, f_Q) and (f_p, g_Q) such that $g_p \to f_p$ and $g_Q \to f_Q$ and only keep the relation which occurs first.

Output: A presentation of the tensor product

Algorithm 1 Tensor Product Algorithm

1: Input: Two module presentations P and Q with reduced bases in SNF:
$G_P \to F_P \to P$
$G_Q \to F_Q \to Q$
2: Compute generators: $F_{P\otimes Q} = F_P \otimes F_Q = (f_P(i), f_Q(j)) \forall f_P(i) \in F_P, f_Q(j) \in F_Q$
3: Compute non-minimal relations: $G_{P\otimes Q} = F_P \otimes G_Q \oplus F_Q \otimes G_P = (f_P(i), g_Q(j)) \oplus (g_P(k), f_Q(\ell)) \quad \forall f_P(i) \in F_P, f_Q(\ell) \in F_Q, g_P(j) \in G_P, f_Q(k) \in F_Q,$
4: For pairs (g_p, f_Q) and (f_p, g_Q) such that $g_p \to f_p$ and $g_Q \to f_Q$, keep the relation which occurs first.
5: Output: Relations, generators and a map: $G_{P\otimes Q} \to F_{P\otimes Q}$

References

[1] Aaron Adcock, Erik Carlsson, and Gunnar Carlsson. The ring of algebraic functions on persistence bar codes. *Arxiv preprint arXiv:1304.0530v1*, 2013.

[2] Michael Barr and Charles Wells. *Toposes, triples and theories*, volume 278. Springer-Verlag New York, 1985.

[3] Michael Barr and Charles Wells. *Category theory for computing science*. Michael Barr and Charles Wells, October 1995.

[4] Andrej Bauer, Karin Cvetko-Vah, Mai Gehrke, Samuel J van Gool, and Ganna Kudryavtseva. A non-commutative priestley duality. *Topology and its Applications*, 2013.

[5] John L. Bell. *Boolean-valued Models and Independence Proofs*. Clarendon Press Oxford, 2005.

[6] Paul Bendich, Sergio Cabello, and Herbert Edelsbrunner. A point calculus for interlevel set homology. *Pattern Recognition Letters*, 33(11):1463–1444, 2012.

[7] G. Bezhanishvili. Bitopological duality for distributive lattices and heyting algebras. *Mathematical Structures in Computer Science*, 20(3):359–393, 2010.

[8] Guram Bezhanishvili and Nick Bezhanishvili. Profinite heyting algebras. *Order*, 25(3):211–277, 2008.

[9] G. Birkhoff. *Lattice theory*, volume 5. AMS Colloquium Publications, Providence RI, third edition, 1940.

[10] Francis Borceux. *Handbook of Categorical Algebra: Volume 3, Sheaf Theory*. Cambridge University Press, December 1994.

[11] G. Carlsson. Topology and data. *American Mathematical Society*, 46(2):255–308, 2009.

[12] G. Carlsson, G. Singh, and A. Zomorodian. Computing multidimensional persistence. *Algorithms and computation*, pages 730–739, 2009.

[13] Gunnar Carlsson, Vin de Silva, and Dmitriy Morozov. Zigzag persistent homology and real-valued functions. In *Proceedings of the 25th annual symposium on Computational geometry*, SCG '09, pages 247–256, New York, NY, USA, 2009. ACM.

[14] Gunnar Carlsson and Afra Zomorodian. The theory of multidimensional persistence. *Discrete & Computational Geometry*, 42(1):71–93, April 2009.

[15] William H Cornish. A sheaf representation of distributive pseudocomplemented lattices. 57:11–15, 1976.

[16] J M Curry. *Sheaves, Cosheaves and Applications*. PhD thesis, 2013.

[17] B. A. Davey and H. A. Priestley. *Introduction to Lattices and Order*. Cambridge University Press, 2 edition, 2002.

[18] H. Edelsbrunner, D. Letscher, and A. Zomorodian. Topological persistence and simplification. *Discrete and Computational Geometry*, 28(4):511–533, 2002.

[19] M P Fourman, C J Mulvey, and Dana S Scott. *Applications of sheaves*. proceedings of the Research Symposium on Applications of Sheaf Theory to Logic, Algebra, and Analysis, Durham, July 9-21, 1977. Springer Verlag, 1979.

[20] G Gierz, K H Hofmann, K Keimel, J D Lawson, M Mislove, and D S Scott. *Continuous lattices and domains*. Cambridge University Press, 1998.

[21] Gerhard Gierz, Karl Heinrich Hofmann, Klaus Keimel, Jimmie D Lawson, Michael W Mislove, and Dana S Scott. *A compendium of continuous lattices*. Springer-Verlag, Berlin, 1980.

[22] G. Grätzer. *Lattice theory*. WH Freeman and Co, San Francisco, 1971.

[23] C J Isham and J Butterfield. Some possible roles for topos theory in quantum theory and quantum gravity. *Foundations of Physics*, 30(10):1707–1735, 2000.

[24] Chris J Isham. Topos Methods in the Foundations of Physics. *arXiv*, quant-ph, April 2010.

[25] A W Jankowski and M Zawadowski. Sheaves over Heyting lattices. *Stud Logica*, 44(3):237–256, 1985.

[26] Peter T Johnstone. *Stone Spaces*. Cambridge University Press, August 1986.

[27] Goro Kato. Elemental principles of t-topos. *EPL (Europhysics Letters)*, 68(4):467, October 2004.

[28] Goro C Kato. Temporal Topos and U-Singularities. pages 1–12, April 2012.

[29] Michael Kinyon, Jonathan Leech, and Joao Pita Costa. Distributivity in skew lattices. *arXiv:1306.5598*, 2013.

[30] Saunders MacLane and I Moerdijk. *Sheaves in Geometry and Logic A First Introduction to Topos Theory*. Springer-Verlag, January 1992.

[31] Amit Patel. A Continuous Theory of Persistence for Mappings Between Manifolds. *arXiv*, math.AT, February 2011.

[32] Allen Dudley Shepard. A cellular description of the derived category of a stratified space. pages 1–161, June 1980.

[33] Mikael Vejdemo-Johansson. Sketches of a platypus: persistent homology and its algebraic foundations. *arXiv*, December 2012.

[34] A. Zomorodian and G. Carlsson. Computing persistent homology. *Discrete and Computational Geometry*, 33(2):249–274, 2005.

[35] H. Adams, A. Atanasov, and G. Carlsson. Morse theory in topological data analysis, 2011.

[36] A. Adcock, D. Rubin, and G. Carlsson. Classification of hepatic lesions using the matching metric, 2012.

[37] 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*, pages 536–546. IEEE, 2007.

[38] P. Bendich, B. Wang, and S. Mukherjee. Local homology transfer and stratification learning. In *Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1355–1370. SIAM, 2012.

[39] G. Carlsson, T. Ishkhanov, V. Silva, and A. Zomorodian. On the local behavior of spaces of natural images. *International Journal of Computer Vision*, 76:1–12, 2008.

[40] 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 annual symposium on Computational geometry*, pages 237–246. ACM, 2009.

[41] F. Chazal, D. Cohen-Steiner, L. J. Guibas, F. Mémoli, and S. Y. Oudot. Gromov-hausdorff stable signatures for shapes using persistence. *Computer Graphics Forum (proc. SGP 2009)*, 2009.

[42] C. Chen, D. Freedman, and C.H. Lampert. Enforcing topological constraints in random field image segmentation. In *Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on*, pages 2089 –2096, June 2011.

[43] D. Cohen-Steiner, H. Edelsbrunner, and J. Harer. Extending persistence using poincaré and lefschetz duality. *Foundations of Computational Mathematics*, 9(1):79–103, 2009.

[44] 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*, pages 1011–1020. Society for Industrial and Applied Mathematics, 2009.

[45] Y. Dabaghian, F. Mmoli, L. Frank, and G. Carlsson. A topological paradigm for hippocampal spatial map formation using persistent homology. *PLoS Computational Biology*, 8(8), 2012.

[46] Vin de Silva, Dmitriy Morozov, and Mikael Vejdemo-Johansson. Dualities in persistent (co)homology. *Inverse Problems*, 27(12):124003, December 2011.

[47] H. Edelsbrunner and J. Harer. Persistent homology-a survey. *Contemporary mathematics*, 453:257–282, 2008.

[48] H. Edelsbrunner and J. Harer. *Computational Topology: an Introduction.* AMS Press, 2009.

[49] D. Eisenbud. *Commutative algebra with a view toward algebraic geometry*, volume 150. Springer, 1995.

[50] Paul G. Goerss and John F. Jardine. *Simplicial homotopy theory*, volume 174 of *Progress in Mathematics.* Birkhuser Verlag, Basel, 1999.

[51] 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. *Bioinformatics*, 23(14):1753–1759, 2007.

[52] David Lipsky, Primoz Skraba, Mikael Vejdemo-Johansson, and Dmitriy Morozov. A spectral sequence for parallelized persistence. 2011.

[53] S. Martin, A. Thompson, E. A. Coutsias, and J.P. Watson. Topology of cyclo-octane energy landscape. *J. Chem. Phys*, 132, 2010.

[54] Dmitriy Morozov. Persistence algorithm takes cubic time in worst case. *BioGeometry News, Dept. Comput. Sci., Duke Univ*, 2005.

[55] 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. *PNAS*, 108(17):7265–7270, 2011.

[56] P. Niyogi, S. Smale, and S. Weinberger. Finding the homology of submanifolds with high confidence from random samples. *Discrete Comput. Geom.*, 39(1):419–441, March 2008.

[57] P. Niyogi, S. Smale, and S. Weinberger. A topological view of unsupervised learning from noisy data. *SIAM Journal on Computing*, 40(3):646–663, 2011.

[58] P. Skraba, M. Ovsjanikov, F. Chazal, and L. Guibas. Persistence-based segmentation of deformable shapes. In *CVPR Workshop on Non-Rigid Shape Analysis and Deformable Image Alignment*, page to appear, June 2010.

[59] B. Walker. Using persistent homology to recover spatial information from encounter traces. In *Proceedings of the 9th ACM international symposium on Mobile ad hoc networking and computing*, MobiHoc '08, pages 371–380. ACM, 2008.

[60] B. Wang, B. Summa, V. Pascucci, and M. Vejdemo-Johansson. Branching and circular features in high dimensional data. *Visualization and Computer Graphics, IEEE Transactions on*, 17(12):1902–1911, 2011.

[61] A.J. Zomorodian. *Topology for computing*, volume 16. Cambridge Univ Pr, 2005.

[62] S. Biasotti, A. Cerri, P. Frosini, D. Giorgi, and C. Landi. Multidimensional size functions for shape comparison. *Journal of Mathematical Imaging and Vision*, 32(2):161179, 2008.

[63] P. Bubenik and J. A. Scott. Categorification of persistent homology. *arXiv:1205.3669*, May 2012.

[64] D. Burghelea and T. K. Dey. Persistence for circle valued maps. *arXiv:1104.5646*, April 2011.

[65] D. Burghelea, T. K. Dey, and Du Dong. Defining and computing topological persistence for 1-cocycles. *arXiv:1012.3763*, December 2010.

[66] F. Cagliari and C. Landi. Finiteness of rank invariants of multidimensional persistent homology groups. *Applied Mathematics Letters*, 24(4):516518, 2011.

[67] G. Carlsson and V. de Silva. Zigzag persistence.

[68] A. Cerri, B. Di Fabio, M. Ferri, P. Frosini, and C. Landi. Multidimensional persistent homology is stable. *arXiv preprint arXiv:0908.0064*, 2009.

[69] A. Cerri, B. Di Fabio, M. Ferri, P. Frosini, and C. Landi. Betti numbers in multidimensional persistent homology are stable functions. 2010.

[70] W. Chacholski, M. Scolamiero, and F. Vaccarino. Combinatorial resolutions of multigraded modules and multipersistent homology.

[71] F. Chazal, V. de Silva, M. Glisse, and S. Oudot. The structure and stability of persistence modules.

[72] F. Chazal, V. de Silva, and S. Oudot. Persistence stability for geometric complexes. *arXiv preprint arXiv:1207.3885*, 2012.

[73] F. Chazal and S. Oudot. Towards persistence-based reconstruction in euclidean spaces. In *Proceedings of the twenty-fourth annual symposium on Computational geometry*, SCG '08, page 232241. ACM.

[74] C. Chen and M. Kerber. An output-sensitive algorithm for persistent homology. In *Proceedings of the 27th annual symposium on Computational geometry.*

[75] C. Chen and M. Kerber. Persistent homology computation with a twist. In *27th European Workshop on Computational Geometry (EuroCG 2011)*, 2011.

[76] D. Cohen-Steiner, H. Edelsbrunner, and J. Harer. Stability of persistence diagrams. 37:103120.

[77] D. Cohen-Steiner, H. Edelsbrunner, J. Harer, and Y. Mileyko. Lipschitz functions have l_p-stable persistence. 10:127139.

[78] D. Cohen-Steiner, H. Edelsbrunner, and D. Morozov. Vines and vineyards by updating persistence in linear time. In *Proceedings of the twenty-second annual symposium on Computational geometry*, SCG '06, page 119126. ACM.

[79] W. Crawley-Boevey. Decomposition of pointwise finite-dimensional persistence modules. *arXiv:1210.0819*, October 2012.

[80] M. d'Amico, P. Frosini, and C. Landi. Optimal matching between reduced size functions.

[81] V. de Silva. Persistence measure and stability, October 2012.

[82] V. de Silva and G. Carlsson. Topological estimation using witness complexes.

[83] V. de Silva and M. Vejdemo-Johansson. Persistent cohomology and circular coordinates. In J. Hershberger and E. Fogel, editors, *Proceedings of the 25th annual symposium on Computational geometry*, page 227236, Aarhus, June 2009. Journal version [98].

[84] T. K. Dey and R. Wenger. Stability of critical points with interval persistence. *Discrete & Computational Geometry*, 38(3):479512, 2007.

[85] P. Donovan and M. R. Freislich. The representation theory of finite graphs and associated algebras. *Carleton Mathematical Lecture Notes*, 5, 1974.

[86] H. Edelsbrunner, J. Harer, and A. Zomorodian. Hierarchical morse complexes for piecewise linear 2-manifolds. In *Proceedings of the seventeenth annual symposium on Computational geometry*, SCG '01, page 7079. ACM.

[87] H. Edelsbrunner, D. Kirkpatrick, and R. Seidel. On the shape of a set of points in the plane. 29:551 – 559.

[88] H. Edelsbrunner and E. Mcke. Three-dimensional alpha shapes. In *Proceedings of the 1992 workshop on Volume visualization*, VVS '92, page 7582. ACM.

[89] G. Ellis and S. King. Persistent homology of groups. *Journal of Group Theory*, 14(4):575–587, 2011.

[90] P. Gabriel. Unzerlegbare darstellungen i. 6:71–103.

[91] J. Gamble, H. Chintakunta, and H. Krim. Applied topology in static and dynamic sensor networks. In *2012 International Conference on Signal Processing and Communications (SPCOM)*, pages 1 –5, July 2012.

[92] R. Ghrist. Barcodes: the persistent topology of data. 45:61 – 75.

[93] R. Ghrist and Y. Hiraoka. Applications of sheaf cohomology and exact sequences to network coding. *preprint*, 2011.

[94] P. T. Johnstone. *Sketches of an Elephant: a topos theory compendium.* Number 43 in Oxford Logic Guides. Oxford University Press, 2002.

[95] A. B. Lee, K. S. Pedersen, and D. Mumford. The nonlinear statistics of high-contrast patches in natural images. *International Journal of Computer Vision*, 54(1):83103, 2003.

[96] M. Lesnick. The optimality of the interleaving distance on multidimensional persistence modules. *arXiv:1106.5305*, June 2011.

[97] D. Lipsky, D. Morozov, P. Skraba, and M. Vejdemo-Johansson. A spectral sequence for parallelized persistence. 12 2011.

[98] D. Morozov, V. de Silva, and M. Vejdemo-Johansson. Persistent cohomology and circular coordinates. *Discrete and Computational Geometry*, 45(4):737—759, 2011.

[99] L. A. Nazarova. Representations of quivers of infinite type. *Mathematics of the USSR-Izvestiya*, 7(4):749, 2007.

[100] A. Patriarca, M. Scolamiero, and F. Vaccarino. A presentation of general multipersistence modules computable in polynomial time? *arXiv:1210.1932*, October 2012.

[101] V. Robins. *Computational Topology for Point Data: Betti Numbers of α-Shapes*, volume 600 of *Lecture Notes in Physics*, pages 261–274. Springer Berlin / Heidelberg.

[102] P. Skraba and M. Vejdemo-Johansson. Persistence modules: algebra and algorithms. Manuscript, 2013.

[103] A. Tausz and G. Carlsson. Applications of zigzag persistence to topological data analysis.

[104] M. Vejdemo-Johansson. Interleaved equivalence of categories of persistence modules. *arXiv:1210.7913*, October 2012.

[105] C. Webb. Decomposition of graded modules. 94:565–571. ArticleType: research-article / Full publication date: Aug., 1985 / Copyright 1985 American Mathematical Society.

[106] S. Weinberger. What is...persistent homology? *Notices of the American Mathematical Society*, pages 36–39, January 2011.

[107] Dominique Attali, André Lieutier, and David Salinas. Efficient data structure for representing and simplifying simplicial complexes in high dimensions. *SoCG*, pages 501–509, 2011.

[108] Dominique Attali, André Lieutier, and David Salinas. Vietoris-Rips complexes also provide topologically correct reconstructions of sampled shapes. *SoCG*, pages 491–500, 2011.

[109] Paul Bendich. *Analyzing Stratified Spaces Using Persistent Versions of Intersection and Local Homology*. PhD thesis, Duke University, 2008.

[110] Paul Bendich, David Cohen-Steiner, Herbert Edelsbrunner, John Harer, and Dmitriy Morozov. Inferring local homology from sampled stratified spaces. *IEEE FoCS*, pages 536–546, 2007.

[111] Paul Bendich and John Harer. Persistent intersection homology. *FoCM*, 11:305–336, 2011.

[112] Paul Bendich, Bei Wang, and Sayan Mukherjee. Local homology transfer and stratification learning. *ACM-SIAM SoDA*, pages 1355–1370, 2012.

[113] J. Boissonnat and A. Ghosh. Manifold reconstruction using tangential delaunay complexes. *SoCG*, pages 324–333, 2010.

[114] Frédéric Cazals and Joachim Giesen. Delaunay triangulation based surface reconstruction: Ideas and algorithms. In *Eff. Comp. Geom. for Curves and Surfaces*, pages 231–273, 2006.

[115] Frédéric Chazal, Leonidas J. Guibas, Steve Y. Oudot, and Primoz Skraba. Analysis of scalar fields over point cloud data. *ACM-SIAM SoDA*, pages 1021–1030, 2009.

[116] Frédéric Chazal, Primoz Skraba, and Amit Patel. Computing well diagrams for vector fields on $\mathbb{R}^n$. Manuscript, http://www.akpatel.org/Homepage/Amit_Patel_files/fp.pdf, 2012.

[117] V. de Silva and G. Carlsson. Topological estimation using witness complexes. *Symposium on Point-Based Graphics*, pages 157–166, 2004.

[118] Tamal K. Dey. *Curve and Surface Reconstruction.* Cambridge University Press, 2007.

[119] H. Edelsbrunner and E. Mucke. Three-dimensional alpha shapes. *Proceedings workshop on Volume Visualization*, (75-82), 1992.

[120] M Goresky and R MacPherson. Intersection homology i. *Topology*, 19:135–162, 1982.

[121] Mark Goresky and Robert MacPherson. *Stratified Morse Theory.* Springer-Verlag, 1988.

[122] Marvin J. Greenberg and John R. Harper. *Algebraic Toplogy A First Course.* Addison-Wesley, 1981.

[123] Allen Hatcher. *Algebraic Topology.* Cambridge University Press, 2002.

[124] James R. Munkres. *Elements of algebraic topology.* Addison-Wesley, 1984.

[125] Donald R. Sheehy. Linear-size approximations to the Vietoris-Rips filtration. *SoCG*, pages 239–248, 2012.

[126] Shmuel Weinberger. The topological classification of stratified spaces. *Chicago Lectures in Mathematics*, 1994.

[127] Afra J. Zomorodian. Fast construction of the Vietoris-Rips complex. *Computers & Graphics*, 34(3):263–271, 2010.

[128] Afra J. Zomorodian. The tidy set: a minimal simplicial set for computing homology of clique complexes. *Proceedings ACM Symposium of Computational Geometry*, pages 257–266, 2010.

[129] K. Mischaikow and M. Mrozek. The Conley index theory. In *Handbook of Dynamical Systems III: Towards Applications*, eds.: B. Fiedler, G. Iooss and N. Kopell, Elsevier, Singapore 2002, 393–460.

[130] M. Mrozek. Leray functor and the cohomological Conley index for discrete time dynamical systems. *Trans. AMS* **318** (1990), 149–178.

[131] C. Conley. *Isolated Invariant Sets and the Morse Index.* Amer. Math. Soc., CBMS Regional Conf. Series Math. **38**, 1978.