

Discrete Morse Theory for Filtrations

Vidit Nanda

A Thesis submitted for the degree of Doctor of Philosophy and approved by

Konstantin **Mischaikow**
Steve **Ferry**
Roger **Nussbaum**
Lou **Kondic**

Department of Mathematics,
Rutgers University



October 2012

Acknowledgements

Enumerating all the ways in which I am grateful to Konstantin would essentially double the length of this document, so I'll save all that stuff for my autobiography. But I will note here that he was simultaneously patient, engaged, proactive, and – best of all – ruthlessly determined to refine and sculpt all our vague big-picture ideas into digestible and implementable concrete pieces. He also provided remarkable opportunities to see and do things that are rarely offered to graduate students. What I'm trying to say is, that conference in Hawaii was *amazing*. The odds of me making it through grad school without an advisor of Konstantin's caliber were precisely zero, and it is extremely fortunate that my talent for screwing things up is exceeded only by his ability to fix them. The tales of my mind-bogglingly embarrassing failures over the last few years would also double the length of this document, and I think I'll leave them out of that autobiography as well. But back to the point: Konstantin, thanks so much for everything, and here's to many more years of working together.

Thanks is also due to various Professors at Rutgers, Georgia Tech and elsewhere who have been patient, helpful and (in a few cases) even willing to collaborate on various projects: Doctors Andrew, Carbone, Chafee, Brezis, de Silva, Duke, Ellis, Gangbo, Greenfield, Ferry, Ghrist, Harer, Han, Kalies, Kokubu, Loss, Lutz, Mrozek, Mucha, Mukherjee, Nussbaum, Sahi, Tamaki, Wang, Wanner, Weibel, Weingart, Yi and many others. I've had many great friends while I was in grad school: Ali, Amit, Asya, Avi, Avital, Brent, Brian, David, Edi, Eliane, Emilie, Gene, Glen, Humberto, Jay, Jeff, Jean-Philippe, Justin, Ke, Kellen, Luc, Marcio, Michael, Miro, Priyam, Radmila, Robert, Shaun, Simao, Susovan, Tianling, Yasu, Zin and everyone else I've forgotten (sorry, deadlines are looming). These people have often kept me sane even when I was having the opposite effect on them.

My family has always been (rather inexplicably, if you ask me) wildly supportive of my ambitions. My parents have believed in me even when I have not believed in myself; and while I can't imagine ever being courageous enough to trust *my* firstborn to fly halfway across the world at seventeen, I have benefitted greatly from their faith and counsel. My brother Vipul has been a constant source of love and companionship ever since he decided to infest our family with his indomitable presence, and I wish him all the success in the world. Vast credit is also due to Yusra for being panacea during certain low points in my academic career and for providing in huge quantities the food love that changed my life for the better. My friendship with Tony has improved me in many ways, so straight love and good vibes to him. Uncles, aunts, cousins, nephews and nieces: you know who you are, thank you for everything.

Finally, I convey sincere gratitude to the Rutgers Math staff for making the endless bureaucracy bearable, my buddy Karthik for changing his name every Thursday, the Naidus and Kishores for helping me adjust to new environments, and the 24-hour Starbucks on Route 18 where much of this dissertation was written.

Oh, and Ria: you are the most intelligent creature in the world.

Dedication

To my family, for all the
five-course mornings, helicopter days, Beowulf evenings, and rude² nights.

Contents

Acknowledgements	i
Dedication	vi
Chapter 1. The Topology of Data and its Transformations	1
1.1. Introduction and Motivation	1
1.2. Contributions and Outline	2
Chapter 2. Background and Terminology	5
2.1. Introduction	5
2.2. Elementary Homological Algebra	5
2.3. Cell Complexes and their Homology	7
2.4. Filtrations and Persistent Homology	9
2.5. Combinatorial Topology	13
Chapter 3. Recovering Manifolds and Functions from Data	17
3.1. Introduction	17
3.2. Reconstructing Manifolds from Sampled Data	17
3.3. Simplicial Reconstruction of Functions	19
3.4. Robustness to Conditioned Noise	23
3.5. Multi-Scale Analysis of Data Transformations	26
Chapter 4. Discrete Morse Theory for Filtrations	29
4.1. Introduction	29
4.2. Discrete Morse Theory for Cell Complexes	29
4.3. Filtration-Subordinate Acyclic Matchings	35
Chapter 5. Simplifying Computation of Persistent Homology	39
5.1. Introduction	39
5.2. The MorseReduce Algorithm	39
5.3. Verification of Correctness	42
5.4. Complexity	44
5.5. Experimental Results	46
Chapter 6. Applications to Exact Sequences	49
6.1. Introduction	49
6.2. The Zigzag Lemma	50
6.3. The Long Exact Sequence of a Triple	51
6.4. The Mayer-Vietoris Sequence	53
Bibliography	57

The Topology of Data and its Transformations

What information consumes is rather obvious: it consumes the attention of its recipients. Hence a wealth of information creates a poverty of attention, and a need to allocate that attention efficiently among the overabundance of information sources that might consume it.

–Herbert Simon [34]

1.1. Introduction and Motivation

Large quantities of high-dimensional *data* arise in innumerable contexts. Of particular interest are the experimental sciences whose measuring devices have improved their power, scope and resolution at extraordinary rates. Aside from large size and dimensionality, experimental data is often plagued by *incompleteness* and *noise*. Naturally, as the rapid generation of such data outpaces our ability to extract useful and dependable information from it, we seek efficient and robust data processing tools. These tools must distill qualitative information which yields insight into the processes that have generated this over-abundance of data in the first place.

Here is a dataset which is neither large nor high dimensional, and can therefore be easily visualized:

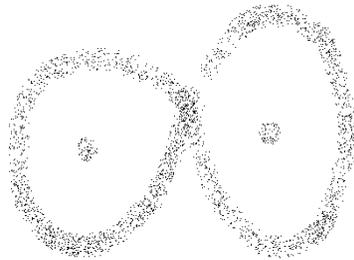


FIGURE 1. Points in the plane

Traditional methods of extracting information from such datasets rely on statistical tools such as regression analysis. These tools provide efficiently computable answers to questions like “*what line best approximates these points?*”, where “best” refers to an optimal solution of a suitably constructed variational problem, such as a minimizer of least square error. Perhaps most importantly for concrete applications, these statistical methods typically provide coarse information which is remarkably robust to measurement errors.

The central focus of *topological* data analysis is the construction of methods that – while also being efficiently computable and robust to noise like their statistical cousins – yield a different type of insight into data. For instance, a coarse topological description of the example dataset from Figure 1 would ideally detect the existence of three point clusters and indicate that one of

these clusters contains two fairly prominent loops. Such information is provided by the theory of *persistent homology*, which we will describe formally later.

Heuristically, a nested family of topological sub-spaces $\mathcal{K} = \{\mathcal{K}_r \mid r \geq 0\}$ of the plane is constructed from the dataset as follows. For each r , \mathcal{K}_r is the union of radius r open balls centered at the data points. For each topological feature (i.e., a connected component or a loop) that exists in any member space of \mathcal{K} , one can associate a *birth* and *death* scale corresponding to the r values where it first appears and disappears respectively. The *persistence* of each feature is then defined to be its lifespan, equalling death minus birth. For instance, the two large holes in the picture will have much greater persistence than the many smaller holes that arise from growing balls around densely clustered points.

The features of persistent homology which make it particularly suitable for analyzing vast datasets have been thoroughly documented in the survey articles of Gunnar Carlsson [7], Robert Ghrist [16], Herbert Edelsbrunner and John Harer [12], so we will not delve too deeply into those advantages here. It should suffice for the purposes of this introduction to state that topological data analysis in general, and persistent homology in particular, have been extremely useful tools in diverse data-driven contexts.

The three basic steps of topological data analysis are as follows:

- (1) **topologize** the data by imposing the structure of a filtered cell complex on it,
- (2) **compute** algebraic invariants of topological features within this complex, and
- (3) **represent** these invariants in a manner that isolates topological features by prominence.

Since we are restricting our attention to persistent homology, the algebraic invariants of interest to us will usually be persistent homology groups and the representations of choice are their associated *persistence diagrams*.

Also of interest – particularly in the topological analysis of dynamical systems – are *transformations* of such data. In this setting, there is no fixed dataset but rather an evolving family of datasets where the object under investigation is the unknown underlying *function* which drives this evolution. Topological analysis of such data transformations involves imposing the structure of a cell complex on each intermediate dataset and approximating the underlying function by a suitable combinatorial map of complexes. In addition to the usual complications imposed by large size and dimensionality, one must also account for *evaluation errors* in determining the image of a given data point under this unknown function. Since algebraic topology is a functorial construction, it associates robust algebraic invariants to continuous functions as well as topological spaces. These invariants play an important role in analyzing the underlying dynamics induced by such unknown continuous functions.

1.2. Contributions and Outline

The central result presented here is an extension of *discrete Morse theory* to filtered cell complexes. This result is from [27] and we cover it here in **Chapter 4**.

Discrete Morse theory was originally developed by Robin Forman [13] for regular CW complexes. The basic idea of this theory is to define a pairing V on some of the cells of a given complex \mathcal{X} . This pairing V , called a *discrete vector field*, induces a new boundary map connecting unpaired, or *critical* cells of adjacent dimension. This new complex consisting only of critical cells is called the *Morse complex* \mathcal{M} associated to V . The central result of discrete Morse theory establishes chain homotopy equivalence of \mathcal{X} and \mathcal{M} , and hence implies an isomorphism on homology. The more cells that V pairs, the smaller \mathcal{M} is, and the easier it becomes to compute its homology groups. This central idea has also been extended to a purely algebraic framework by

Ekkehard Batzies, Volkmar Welker [4] and Emil Sköldbberg [35] for the purposes of constructing minimal cellular resolutions of multi-graded ideals in polynomial rings.

Here, we introduce a version of algebraic discrete Morse theory which is *subordinate* to a given filtration $\mathcal{X}^0 \subset \mathcal{X}^1 \subset \dots \subset \mathcal{X}^K = \mathcal{X}$ of a cell complex \mathcal{X} . This theory produces a filtered Morse complex $\mathcal{M}^0 \subset \dots \subset \mathcal{M}^K = \mathcal{M}$ so that not only is each \mathcal{X}^k chain homotopy equivalent to the corresponding \mathcal{M}^k in the sense of Forman's original result, but also the map on homology induced by including \mathcal{X}^k into \mathcal{X}^{k+1} is naturally equivalent to the corresponding map for the inclusion of \mathcal{M}^k into \mathcal{M}^{k+1} . Thus, *the persistent homology groups of the filtered complexes \mathcal{X} and \mathcal{M} are isomorphic*. As a first application of this theory, we introduce efficient algorithms to pre-process computation of persistent homology groups in **Chapter 5**.

Making the discrete vector field structure subordinate to filtrations naturally raises the question of whether such methods will simplify more general computations in homological algebra. We answer this question affirmatively in **Chapter 6** with a second application of our theory. The context here is one of the most basic constructions of homological algebra: the derived functor. That is, one can apply the Morse pre-processing machinery to significantly reduce the computational cost of building long exact sequences in homology from short exact sequences of free module chain complexes over a principal ideal domain. In particular, we focus our attention here on the constructing the *long exact sequence of a triple* and the *Mayer-Vietoris sequence*.

One important question regarding the application of algebraic topology to datasets is, *assuming that the data has been randomly sampled from an underlying unknown topological space, how much information about that space can one recover from the data?* A result of Partha Niyogi, Steve Smale and Shmuel Weinberger [32] provides explicit bounds on the size of a uniformly sampled dataset lying on or near a compact Riemannian submanifold $M \subset \mathbb{R}^n$ required to recover the homotopy type of M with high confidence. The secondary contribution presented here, which is the main result of [28], is an analogous theorem for Lipschitz-continuous *functions* between such submanifolds. We demonstrate in **Chapter 3** that we can recover the map induced on homology by such a function with high confidence given only a sufficiently large set of uniformly sampled data points and their images.

The algebraic, topological and combinatorial machinery required to get to the central results of this dissertation is minimal; it has been included in **Chapter 2** for completeness.

CHAPTER 2

Background and Terminology

“A little learning is a dangerous thing
 Drink deep, or taste not the Pierian spring:
 There shallow draughts intoxicate the brain,
 And drinking largely sobers us again.”

–Alexander Pope, *An essay on criticism*

2.1. Introduction

We provide a brief synopsis of definitions and results from homological algebra in Section 2.2 and combinatorial topology in Section 2.5 while referring the reader to the excellent textbooks of Charles Weibel [39] and Dmitry Kozlov [23] for complete treatments of these topics. Section 2.3 deals with combinatorial cell complexes as developed by Albert Tucker [37] and Solomon Lefschetz [24]. Section 2.4 surveys the relatively new field of *persistent homology* [15, 7, 40, 10] which is defined for filtered cell complexes.

2.2. Elementary Homological Algebra

Let \mathbb{N} denote the set of natural numbers (including zero), let \mathbb{Z} denote the integers and let \mathbf{R} be a principal ideal domain (PID).

2.2.1. Exact Sequences. By a *sequence* \mathcal{A} of \mathbf{R} -modules we mean a collection of \mathbf{R} -modules A_n and their morphisms $\gamma_n : A_n \rightarrow A_{n-1}$ for $n \in \mathbb{N}$ arranged as follows:

$$\dots \xrightarrow{\gamma_{n+1}} A_n \xrightarrow{\gamma_n} A_{n-1} \xrightarrow{\gamma_{n-1}} \dots \xrightarrow{\gamma_1} A_0$$

We say that \mathcal{A} is *exact* if it is exact at each A_n , i.e., if $\ker \gamma_n = \text{img } \gamma_{n+1}$ for each $n \in \mathbb{N}$. Given $k \in \mathbb{Z}$, a *degree k morphism* from \mathcal{A} to another sequence \mathcal{A}' is a collection of \mathbf{R} -module maps $\Sigma = \{\sigma_n : A_n \rightarrow A'_{n+k}\}$ so that the following diagram commutes for each n and k

$$\begin{array}{ccc}
 A_n & \xrightarrow{\gamma_n} & A_{n-1} \\
 \downarrow \sigma_n & & \downarrow \sigma_{n-1} \\
 A'_{n+k} & \xrightarrow{\gamma'_{n+k}} & A'_{n+k-1}
 \end{array}$$

with the understanding that all negative indices correspond to the trivial module. We call $\Sigma : \mathcal{A} \rightarrow \mathcal{A}'$ an *isomorphism* if it has degree 0 and each σ_n is an isomorphism of \mathbf{R} -modules. We denote by $\text{id}_{\mathcal{A}}$ the degree 0 automorphism of \mathcal{A} where each A_n is identically mapped to itself.

2.2.2. Chain Complexes. A *chain complex* $(\mathcal{C}; \partial)$ over \mathbf{R} is a sequence of \mathbf{R} -modules

$$\dots \xrightarrow{\partial_{n+1}} C_n \xrightarrow{\partial_n} C_{n-1} \xrightarrow{\partial_{n-1}} \dots \xrightarrow{\partial_1} C_0 \rightarrow 0$$

such that $\partial_n \circ \partial_{n+1} \equiv 0$ for each n . An element of C_n is called a n -chain and ∂_n is called the n -th boundary operator. The submodules $Z_n = \ker \partial_n$ and $B_n = \text{img } \partial_{n+1}$ of C_n are called the n -cycles and n -boundaries respectively. The n -th homology module of the chain complex $(\mathcal{C}; \partial)$ is defined to be the quotient

$$H_n(\mathcal{C}; \partial) = \frac{Z_n}{B_n}$$

When the boundary operator ∂ is clear from context, we suppress it and simply denote the homology modules by $H_n(\mathcal{C})$. Recall that exact sequences are chain complexes which have trivial homology modules.

Given a chain complex $(\mathcal{C}; \partial)$, a *chain subcomplex* is a collection of submodules $\{C'_n \subset C_n \mid n \in \mathbb{N}\}$ such that $\partial_n(C'_n) \subset C'_{n-1}$. Note that $(\mathcal{C}'; \partial)$ forms a chain complex in its own right. Given such a subcomplex, we may define a *relative chain complex* $(\mathcal{C}, \mathcal{C}')$ whose n -chains are the quotient modules C_n/C'_n and the boundary maps are naturally induced by ∂ . The homology modules of this relative complex are denoted by $H_*(\mathcal{C}, \mathcal{C}')$.

A *chain map* $\phi : (\mathcal{C}; \partial) \rightarrow (\mathcal{D}; \delta)$ between two chain complexes is a degree 0 morphism of the underlying \mathbf{R} -module sequences, i.e., a sequence of maps $\phi_n : C_n \rightarrow D_n$ such that $\phi_{n-1} \circ \partial_n \equiv \delta_n \circ \phi_n$. It is easy to see that any chain map ϕ maps cycles in \mathcal{C} to cycles in \mathcal{D} and similarly boundaries to boundaries. Therefore, ϕ induces well-defined maps $\phi_n^* : H_n(\mathcal{C}) \rightarrow H_n(\mathcal{D})$ of homology modules.

We say that two chain maps $\alpha, \beta : (\mathcal{C}; \partial) \rightarrow (\mathcal{D}; \delta)$ are *chain homotopic* if there exist \mathbf{R} -module morphisms $\Theta_n : C_n \rightarrow D_{n+1}$ such that

$$\alpha_n - \beta_n \equiv \delta_{n+1} \circ \Theta_n + \Theta_{n-1} \circ \partial_n \text{ on } C_n.$$

for each $n \in \mathbb{N}$. In this case, we call Θ a *chain homotopy* which realizes the chain homotopy equivalence of α and β .

Two chain maps $\phi : (\mathcal{C}; \partial) \rightarrow (\mathcal{D}; \delta)$ and $\psi : (\mathcal{D}; \delta) \rightarrow (\mathcal{C}; \partial)$ are *chain equivalences* if $\phi \circ \psi$ is chain homotopic to the identity map on \mathcal{C} and $\psi \circ \phi$ is chain homotopic to the identity map on \mathcal{D} . It is well-known that chain equivalences induce inverse maps on the homology modules and hence establish an isomorphism $H_*(\mathcal{C}, \partial) \simeq H_*(\mathcal{D}, \delta)$. For a detailed discussion and proofs see [36, Ch. 4] or [31, Ch. 1.13].

2.2.3. Sequences of Chain Complexes. A *sequence of chain complexes* $S = \{\mathcal{C}^m, \gamma^m\}$ consists of

- (1) chain complexes $\mathcal{C}^m = (C^m; \partial^m)$, $m \in \mathbb{N}$, and
- (2) chain maps $\gamma^m : \mathcal{C}^m \rightarrow \mathcal{C}^{m-1}$.

In particular, for each $n \in \mathbb{N}$ the corresponding n -chain modules form a \mathbf{R} -module sequence S_n which we call the n -th row of S :

$$\dots \xrightarrow{\gamma_n^{m+1}} C_n^m \xrightarrow{\gamma_n^m} C_n^{m-1} \xrightarrow{\gamma_n^{m-1}} \dots \xrightarrow{\gamma_n^1} C_n^0$$

Let S be a sequence of chain complexes as defined above and let $S' = \{\mathcal{D}^m, \beta^m\}$ be another such sequence. A *morphism* Ω from S to S' is given by chain maps $\omega^m : \mathcal{C}^m \rightarrow \mathcal{D}^m$ so that for each $n \in \mathbb{N}$, the collection $\Omega_n = \{\omega_n^m : C_n^m \rightarrow D_n^m \mid m \in \mathbb{N}\}$ of maps on n -chains determines a morphism of \mathbf{R} -module sequences from S_n to S'_n .

Note that morphisms of chain complex sequences may be composed in the expected way. Given $\Omega : S \rightarrow S'$ and $\Upsilon : S' \rightarrow S''$, the composition $\Upsilon \circ \Omega : S \rightarrow S''$ is defined by the sequence of composed chain maps $\{\nu^m \circ \omega^m\}$. Moreover, the concept of an *identity map* id_S on S is well-defined to be the morphism which acts as the identity map on each row of S .

Note that a morphism $\Omega = \{\omega^m : \mathcal{C}^m \rightarrow \mathcal{D}^m\}$ of chain complex sequences induces maps $\omega_*^m : H_*(\mathcal{C}^m) \rightarrow H_*(\mathcal{D}^m)$ of homology modules. A morphism $\Upsilon : S' \rightarrow S$ determined by maps $\upsilon^m : \mathcal{D}^m \rightarrow \mathcal{C}^m$ is called a *weak inverse* of $\Omega : S \rightarrow S'$ if ω^m and υ^m are chain equivalences for each $m \in \mathbb{N}$. We call Ω a *weak equivalence* of chain complex sequences if it has a weak inverse.

The existence of Υ as a weak inverse to Ω implies the existence of a sequence of chain homotopies $\Theta^m : \mathcal{C}^m \rightarrow \mathcal{C}^m$ and $\Xi^m : \mathcal{D}^m \rightarrow \mathcal{D}^m$ for each $m \in \mathbb{N}$ which realize the chain equivalence of ω^m and υ^m and hence implies isomorphisms $H_*(\mathcal{C}^m) \simeq H_*(\mathcal{D}^m)$ of homology modules for each $m \in \mathbb{N}$. The following stricter notion yields similar results for *relative* homology but requires these chain homotopies to respect the structure of the the rows of S and S' .

DEFINITION 2.2.1. We call $\Omega : S \rightarrow S'$ a *strong equivalence* if there exists a weak inverse $\Upsilon : S' \rightarrow S$ along with chain homotopies $\Theta^m : \mathcal{C}_*^m \rightarrow \mathcal{C}_{*+1}^m$ and $\Xi^m : \mathcal{D}_*^m \rightarrow \mathcal{D}_{*+1}^m$ for each $m \in \mathbb{N}$ such that

- (1) Θ^m and Ξ^m realize the weak equivalence of Ω and Υ , and
- (2) the maps $\{\Theta_n^m : \mathcal{C}_n^m \rightarrow \mathcal{C}_{n+1}^m \mid m \in \mathbb{N}\}$ and $\{\Xi_n^m : \mathcal{D}_n^m \rightarrow \mathcal{D}_{n+1}^m \mid m \in \mathbb{N}\}$ constitute degree 1 self-morphisms of the rows S_n and S'_n respectively for each $n \in \mathbb{N}$.

Υ is called a *strong inverse* of Ω in this case.

The second requirement of the definition above imposes the following additional relations on the chain homotopies which realize a weak equivalence of chain complex sequences:

$$\Theta_n^{m-1} \circ \gamma_n^m \equiv \gamma_{n+1}^m \circ \Theta_n^m \quad \text{and} \quad \Xi_n^{m-1} \circ \beta_n^m \equiv \beta_{n+1}^m \circ \Xi_n^m$$

For each m and $p \in \mathbb{N}$, we denote by $\gamma^{m,p} : \mathcal{C}^m \rightarrow \mathcal{C}^{m+p}$ the composition of chain maps $\gamma^{m+p-1} \circ \dots \circ \gamma^m$.

REMARK 2.2.2. Let $\Omega : S \rightarrow S'$ and $\Upsilon : S' \rightarrow S$ be strong equivalences. Given natural numbers m and p , note that ω^{m+p} restricted to $\gamma^{m,p}(\mathcal{C}^m)$ maps into $\beta^{m,p}(\mathcal{D}^m)$ and hence induces a well-defined relative chain map $\bar{\omega}^{m,p} : (\mathcal{C}^{m+p}, \gamma^{m,p}(\mathcal{C}^m)) \rightarrow (\mathcal{D}^{m+p}, \beta^{m,p}(\mathcal{D}^m))$. Let $\bar{\upsilon}^{m,p}$ be the corresponding map for Υ . Since Ω and Υ are strong equivalences, there exist chain homotopies $\Theta^{m,p}$ and $\Xi^{m,p}$ which induce well-defined relative chain homotopies that in turn establish the chain equivalence of $\bar{\omega}^{m,p}$ and $\bar{\upsilon}^{m,p}$. Thus, we obtain an isomorphism of *relative* homology modules $H_*(\mathcal{C}^{m+p}, \gamma^{m,p}(\mathcal{C}^m)) \simeq H_*(\mathcal{D}^{m+p}, \beta^{m,p}(\mathcal{D}^m))$.

Recall that a sequence S of chain complexes is *short exact* if it has the form

$$\dots \rightarrow 0 \rightarrow \mathcal{C} \xrightarrow{\alpha} \mathcal{D} \xrightarrow{\beta} \mathcal{E} \rightarrow 0 \rightarrow \dots$$

and moreover, if for each $n \in \mathbb{N}$ the n -th row S_n is an exact sequence of \mathbf{R} -modules

$$0 \rightarrow C_n \xrightarrow{\alpha_n} D_n \xrightarrow{\beta_n} E_n \rightarrow 0$$

In particular, exactness of S_n implies that α_n is injective, β_n is surjective and $\ker \beta_n = \text{img } \alpha_n$ for each $n \in \mathbb{N}$.

2.3. Cell Complexes and their Homology

The following notion of a cell complex which dates back to Tucker [37] and Lefschetz [24] provides a general and convenient method of representing the *bases* of free and finitely generated \mathbf{R} -modules which comprise a chain complex. Our presentation here is more closely related Mrozek and Batko's definition of *S-complexes* from [29].

DEFINITION 2.3.1. Consider a finite graded set $\mathcal{X} = \bigsqcup_{n \in \mathbb{N}} \mathcal{X}_n$ along with a function $\kappa : \mathcal{X} \times \mathcal{X} \rightarrow \mathbf{R}$ and denote $\xi \in \mathcal{X}_n$ by $\dim \xi = n$. Then, (\mathcal{X}, κ) is a *complex* if

(i) For each ξ and ξ' in \mathcal{X} ,

$$\kappa(\xi, \xi') \neq 0 \text{ implies } \dim \xi = \dim \xi' + 1. \quad (2.1)$$

(ii) For each ξ and ξ'' in \mathcal{X} ,

$$\sum_{\xi' \in \mathcal{X}} \kappa(\xi, \xi') \cdot \kappa(\xi', \xi'') = 0 \quad (2.2)$$

An element $\xi \in \mathcal{X}$ is called a *cell* and $\dim \xi$ is called the *dimension* of ξ . The function κ is called the *incidence function* for the cell complex (\mathcal{X}, κ) . We denote (\mathcal{X}, κ) simply by \mathcal{X} when there is no possible confusion about the incidence function. The *face partial order* \preceq is induced on the elements of \mathcal{X} by the transitive closure of the generating relation \prec given as follows. For ξ and $\xi' \in \mathcal{X}$

$$\xi' \prec \xi \text{ if } \kappa(\xi, \xi') \neq 0.$$

By (2.1), the function $\dim : \mathcal{X} \rightarrow \mathbb{N}$ is an order-preserving map.

EXAMPLE 2.3.2. The most common types of complexes encountered in the analysis of data are *simplicial* and *cubical* complexes. We observe that both are special cases of the general notion of cell complex from the preceding definition.

(1) Recall that a finite *abstract simplicial complex* \mathcal{X} is a collection of non-empty subsets of a finite universal set \mathcal{U} such that if $\xi \subset \mathcal{U}$ is in \mathcal{X} and $\xi' \subset \xi$ then ξ' is also in \mathcal{X} . Given a cell ξ in \mathcal{X} – called a *simplex* – we define $\dim \xi = \#\xi - 1$ (where $\#$ denotes cardinality as a set). Note that \mathcal{U} is naturally identified with the 0-dimensional simplices \mathcal{X}_0 . Assume that the elements of \mathcal{U} are ordered and let $\xi = \{u_0, \dots, u_k\}$ be a k -simplex where the indices respect the order inherited from \mathcal{U} . Over the integers, the simplicial incidence function $\kappa_\Delta : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{Z}$ is defined by

$$\kappa_\Delta(\xi, \xi') = \begin{cases} (-1)^j & \text{if } \xi' = \xi \setminus \{u_j\} \\ 0 & \text{otherwise} \end{cases}$$

(2) An *elementary interval* I has the form $[p, p+k] \subset \mathbb{R}$ for $p \in \mathbb{Z}$ and $k \in \{0, 1\}$. I is called *degenerate* if $k = 0$. For a non-degenerate I , define the degenerate intervals $I^+ = [p+k, p+k]$ and $I^- = [p, p]$ called the *front* and *back* faces of I respectively. An *elementary cube* ξ of dimension d is the product of elementary intervals $I_1 \times \dots \times I_N$ so that precisely d of these intervals are non-degenerate; the n -th front and back faces ξ_n^+ and ξ_n^- of ξ are the elementary cubes obtained by replacing a non-degenerate factor I_n of ξ by I_n^+ or I_n^- respectively in the product expansion.

A *cubical complex* \mathcal{X} is a finite union of elementary cubes so that for each $\xi \in \mathcal{X}$, all faces of ξ are also in \mathcal{X} . Letting $\xi = I_1 \times \dots \times I_N$ as above, the *cubical incidence function* $\kappa_\square : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{Z}$ is defined by

$$\kappa_\square(\xi, \xi') = \begin{cases} -1 & \text{if } \xi' = \xi_n^+ \text{ for even } n \text{ or } \xi' = \xi_n^- \text{ for odd } n \\ +1 & \text{if } \xi' = \xi_n^+ \text{ for odd } n \text{ or } \xi' = \xi_n^- \text{ for even } n \\ 0 & \text{otherwise.} \end{cases}$$

Given a complex (\mathcal{X}, κ) the *associated chain complex* $\mathcal{C}(\mathcal{X})$ may be written as

$$\dots \xrightarrow{\partial_{n+1}} C_n(\mathcal{X}) \xrightarrow{\partial_n} C_{n-1}(\mathcal{X}) \xrightarrow{\partial_{n-1}} \dots \xrightarrow{\partial_1} C_0(\mathcal{X})$$

where $C_n(\mathcal{X})$ is the free \mathbf{R} -module $\mathbf{R}(\mathcal{X}_n)$ based at the n -dimensional cells and the boundary operator $\partial_n: C_n(\mathcal{X}) \rightarrow C_{n-1}(\mathcal{X})$ is generated by the following basis action

$$\partial_n \xi := \sum_{\xi' \in \mathcal{X}} \kappa(\xi, \xi') \xi'.$$

The n -dimensional cycles, boundaries and homology modules of the cell complex \mathcal{X} are denoted $Z_n(\mathcal{X})$, $B_n(\mathcal{X})$ and $H_n(\mathcal{X})$ respectively and defined to be the n -cycles, boundaries and homology modules of the associated chain complex $\mathcal{C}(\mathcal{X})$.

Consider $\mathcal{X}' \subset \mathcal{X}$ and note that the restriction of κ to $\mathcal{X}' \times \mathcal{X}'$ satisfies (2.1). If for each $\eta \in \mathcal{X}'$ the set $\{\xi \in \mathcal{X} \mid \xi \prec \eta\}$ is contained in \mathcal{X}' , then we say that \mathcal{X}' satisfies the *subcomplex property* and call \mathcal{X}' a *subcomplex* of (\mathcal{X}, κ) . Note that equation (2.2) is automatically satisfied for \mathcal{X}' , and so (\mathcal{X}', κ) is a complex in its own right.

When \mathcal{X}' is a subcomplex of \mathcal{X} , its associated chain complex $\mathcal{C}(\mathcal{X}')$ is a chain subcomplex of $\mathcal{C}(\mathcal{X})$ and the relative homology modules $H_n(\mathcal{X}, \mathcal{X}')$ are defined to be the relative homology modules $H_n(\mathcal{C}(\mathcal{X}), \mathcal{C}(\mathcal{X}'))$ of chain complexes.

2.4. Filtrations and Persistent Homology

Consider a complex (\mathcal{X}, κ) over a PID \mathbf{R} . A *filtration* F of \mathcal{X} is a finite sequence of subcomplexes $\{\mathcal{X}^m \mid 1 \leq m \leq M\}$ of \mathcal{X} such that

$$\mathcal{X}^1 \subset \mathcal{X}^2 \subset \dots \subset \mathcal{X}^M = \mathcal{X}$$

The *associated chain filtration* $F(\mathcal{X})$ is defined to be the following sequence of chain complexes

$$\mathcal{C}(\mathcal{X}^1) \xrightarrow{i^1} \mathcal{C}(\mathcal{X}^2) \xrightarrow{i^2} \dots \xrightarrow{i^{M-1}} \mathcal{C}(\mathcal{X}^M)$$

where the chain maps i^m arise from inclusion of chains. For each m and n in \mathbb{N} , let $C_n(\mathcal{X}^m)$, $Z_n(\mathcal{X}^m)$ and $B_n(\mathcal{X}^m)$ denote the n -dimensional chains, cycles and boundaries of the chain complex $\mathcal{C}(\mathcal{X}^m)$ respectively.

For $p \in \mathbb{N}$, we denote by $i^{m,p}: \mathcal{C}(\mathcal{X}^m) \rightarrow \mathcal{C}(\mathcal{X}^{m+p})$ the composition $i^{m+p-1} \circ \dots \circ i^m$ with the tacit understanding that the composition equals the identity map on $\mathcal{C}(\mathcal{X}^m)$ when $p = 0$ and equals zero when $m + p$ exceeds the length M of the filtration F . The p -persistent q -th homology module of \mathcal{X}^m is defined to be

$$H_n^p(\mathcal{X}^m) := i_{n*}^{m,p}(H_q(\mathcal{X}^k)) = \frac{i^{m,p}(Z_n(\mathcal{X}^m))}{i^{m,p}(Z_n(\mathcal{X}^m)) \cap B_n(\mathcal{X}^{m+p})} \quad (2.3)$$

where the second equality follows by induction on p .

Informally, $H^p(\mathcal{X}^m)$ consists of cycles included from $\mathcal{C}(\mathcal{X}^m)$ into $\mathcal{C}(\mathcal{X}^{m+p})$ modulo boundaries. An algorithm for computing persistent homology modules when the underlying PID \mathbf{R} is a field may be found in [40].

2.4.1. Persistence Diagrams and Stability. Let S be a sequence of chain complexes of finitely generated modules over a PID \mathbf{R} , given by

$$\mathcal{C}^1 \xrightarrow{\gamma^1} \mathcal{C}^2 \xrightarrow{\gamma^2} \dots \xrightarrow{\gamma^{m-1}} \mathcal{C}^m \xrightarrow{\gamma^m} \dots$$

Fix a dimension $n \in \mathbb{N}$ and note that S induces a sequence $H_n(S)$ of n -dimensional homology modules, i.e.,

$$H_n(\mathcal{C}^1) \xrightarrow{\gamma_{n*}^1} H_n(\mathcal{C}^2) \xrightarrow{\gamma_{n*}^2} \dots \xrightarrow{\gamma_{n*}^{m-1}} H_n(\mathcal{C}^m) \xrightarrow{\gamma_{n*}^m} \dots$$

and let $\mathcal{H}_n = \bigoplus_{m \in \mathbb{N}} H_n(\mathcal{C}^m)$ denote the direct sum of \mathbf{R} -modules which comprise $H_n(S)$. Denote the ring of polynomials with \mathbf{R} -coefficients in an indeterminate t by $\mathbf{R}[t]$ and recall that

this ring is naturally graded:

$$\mathbf{R}[t] = \bigoplus_{m \in \mathbb{N}} \mathbf{R}_m[t],$$

where $\mathbf{R}_m[t]$ consists of all polynomials in $\mathbf{R}[t]$ of degree m . We endow $H_n(S)$ with the structure of an $\mathbf{R}[t]$ -module via the following action of t on every generator $x = (x_1, x_2, \dots) \in \mathcal{H}_n$

$$t \cdot (x_1, x_2, \dots) = (0, \gamma_{n*}^1(x_1), \gamma_{n*}^2(x_2), \dots)$$

In general, there is no nice classification theorem for such modules because $\mathbf{R}[t]$ may not be a principal ideal domain. However, if \mathbf{R} happens to be a *field*, then $\mathbf{R}[t]$ is a PID and hence the structure theorem for finitely generated (graded) modules over a (graded) PID applies. Recalling that every graded ideal of $\mathbf{R}[t]$ has the form $t^k \cdot \mathbf{R}[t]$ for some $k \in \mathbb{N}$, this structure theorem decomposes \mathcal{H}_n into *free* and *torsional* parts

$$\mathcal{H}_n \simeq \left[\bigoplus_{i=1}^I t^{a_i} \cdot \mathbf{R}[t] \right] \oplus \left[\bigoplus_{j=1}^J t^{b_j} \mathbf{R}[t]/(t^{d_j}) \right]$$

Thus, the family of natural numbers $\{a_i, b_j, d_j\}$ for $i \in \{1, \dots, I\}$ and $j \in \{1, \dots, J\}$ provides a complete invariant for the isomorphism class of \mathcal{H}_n as a module over $\mathbf{R}[t]$.

Consider now the case where S is the chain filtration associated to a cell complex (\mathcal{X}, κ) filtered by $F = \{\mathcal{X}^m\}$. In this case, the numbers $\{a_i, b_j, d_j\}$ completely classify the n -dimensional persistent homology modules of F . We encode the structural information of \mathcal{H}_n in a *persistence diagram* P_n^F which consists of all the points (a_i, ∞) and (b_j, d_j) – with appropriate multiplicity. In Figure 1, we plot a persistence diagram in the extended plane $\mathbb{R} \times (\mathbb{R} \cup \{\infty\})$. Here, the unique point of the type (a, ∞) occurs when $a = 0$; it has been placed at the top of the diagram above 0 and is represented by a diamond shape (\diamond). By convention, all points of the diagonal $\{(x, x) \mid x \in \mathbb{R}\}$ are also included in the persistence diagram with infinite multiplicity.

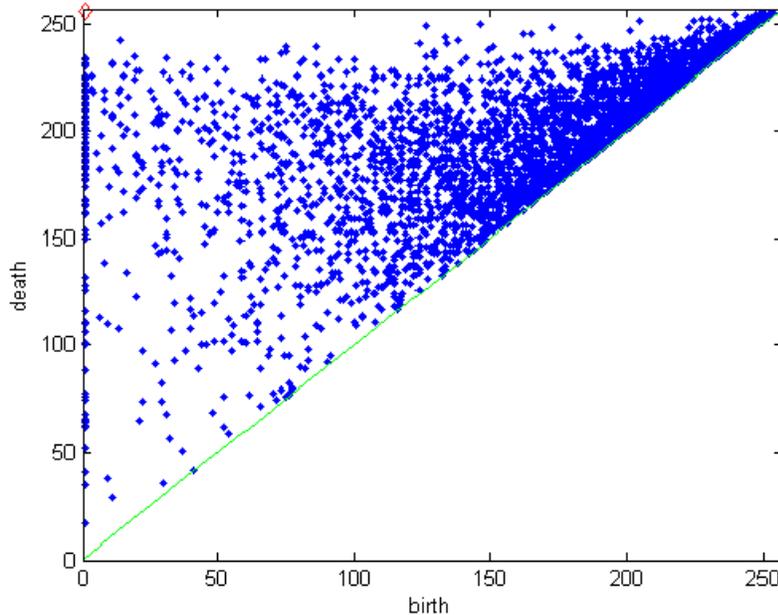


FIGURE 1. A persistence diagram.

The *bottleneck distance* between two persistence diagrams P and P' is defined as follows

$$d_B(P, P') = \inf_{\lambda} \sup_{p \in P} \|p - \lambda(p)\|_{\infty}$$

where λ ranges over all bijections from P to P' , and $\|\cdot\|_{\infty}$ denotes the usual L^{∞} -norm on the plane. Note that since the diagonals are included in both diagrams with infinite multiplicity, it is always possible to pair an excess off-diagonal point with a point on the diagonal. Thus, it is not necessary to restrict the scope of bottleneck distance to only those diagrams which have the same number of off-diagonal points.

The key feature of persistence diagrams which makes them effective invariants of data is their *stability* to perturbations [10] in the following sense. Let $g : X \rightarrow \mathbb{R}$ be a continuous function on a topological space X and for each $a \in \mathbb{R}$ denote by X^a the sub-level set $g^{-1}(-\infty, a)$. We call g *tame* if there are only finitely many real numbers $c \in \mathbb{R}$ for which the inclusion of spaces $X^c \hookrightarrow X^{c+\epsilon}$ does not induce an isomorphism of homology modules for sufficiently small $\epsilon > 0$. These exceptional numbers c are called *homological critical values* of g .

Assume now that the domain X is triangulable and that X^c is a sub-complex of X for each homological critical value c of g . Ordering these critical values by $c_1 \leq c_2 \leq \dots \leq c_M$, we have the following filtration F_g of the triangulated complex X associated to the tame function g :

$$X^{c_1} \subset X^{c_2} \subset \dots \subset X^{c_M} \subset X^{\infty} = X$$

Define P_g to be the persistence diagram associated to F_g with the following simple modification: each point (i, j) of the usual diagram is replaced with the corresponding pair of homological critical values (c_i, c_j) of g . The following theorem from [10] makes precise the stability of such persistence diagrams to perturbation

THEOREM 2.4.1. *Let g and g' be tame functions from a triangulable space X to \mathbb{R} . Then,*

$$d_B(P_g, P_{g'}) \leq \|g - g'\|_{\infty}$$

2.4.2. Morphisms of Filtrations. Using the framework of Section 2.2.3 allows us to impose a categorical structure on filtrations of cell complexes as follows. Let \mathcal{X} and \mathcal{Y} be complexes over a PID \mathbf{R} , and let $F = \{\mathcal{X}^m\}_1^M$ and $G = \{\mathcal{Y}^m\}_1^M$ be respective filtrations. We denote the corresponding chain maps arising from inclusion by $i^m : \mathcal{C}(\mathcal{X}^m) \rightarrow \mathcal{C}(\mathcal{X}^{m+1})$ and $j^m : \mathcal{C}(\mathcal{Y}^m) \rightarrow \mathcal{C}(\mathcal{Y}^{m+1})$. A *filtered chain map* $\Phi = \{\phi^m : \mathcal{C}(\mathcal{X}^m) \rightarrow \mathcal{C}(\mathcal{Y}^m)\}$ from F to G is a morphism $F(\mathcal{X}) \rightarrow G(\mathcal{Y})$ of the associated chain filtrations.

This definition is illustrated by the following commutative diagram:

$$\begin{array}{ccccccc} \mathcal{C}(\mathcal{X}^0) & \xrightarrow{i^0} & \mathcal{C}(\mathcal{X}^1) & \xrightarrow{i^1} & \dots & \xrightarrow{i^{m-1}} & \mathcal{C}(\mathcal{X}^m) & \xrightarrow{i^m} & \dots \\ \downarrow \phi^0 & & \downarrow \phi^1 & & & & \downarrow \phi^m & & \\ \mathcal{C}(\mathcal{Y}^0) & \xrightarrow{j^0} & \mathcal{C}(\mathcal{Y}^1) & \xrightarrow{j^1} & \dots & \xrightarrow{j^{m-1}} & \mathcal{C}(\mathcal{Y}^m) & \xrightarrow{j^m} & \dots \end{array}$$

Note that for each $m, p \in \mathbb{N}$ we may define a chain map $\phi^{m,p} : \mathcal{C}(\mathcal{X}^{m+p}) \rightarrow \mathcal{C}(\mathcal{Y}^{m+p})$ simply by setting $\phi^{m,p} = \phi^{m+p}$, and recall that $i^{m,p} : \mathcal{C}(\mathcal{X}^m) \rightarrow \mathcal{C}(\mathcal{X}^{m+p})$ denotes the usual composition $i^{m+p-1} \circ \dots \circ i^m$.

PROPOSITION 2.4.2. The maps $\phi^{m,p}$ canonically induce morphisms $\phi_*^{m,p} : H_*^p(\mathcal{X}^m) \rightarrow H_*^p(\mathcal{Y}^m)$ of persistent homology modules.

PROOF. Pick a cycle $\mathbf{x}' \in Z(\mathcal{X}^m)$ and include it into $\mathcal{C}(\mathcal{X}^{m+p})$ via $i^{m,p}$. To avoid confusion, call the included chain $\mathbf{x} = i^{m,p}(\mathbf{x}')$. Now let $[\mathbf{x}]$ be the equivalence class of \mathbf{x} in $H^p(\mathcal{X}^m)$. We will

provisionally define $\phi_*^{m,p}([\mathbf{x}])$ to be the equivalence class $[\phi^{m,p}(\mathbf{x})]$ in $H^p(\mathcal{Y}^m)$. First, we show that the image lies in $H^p(\mathcal{Y}^m)$ as desired, and then we will show this map is well-defined on persistent homology in the sense of being independent of the choice of \mathbf{x}' (and hence \mathbf{x}) up to elements of $B(\mathcal{X}^{m+p})$.

First, $\phi^{m,p}(\mathbf{x}) \in C(\mathcal{Y}^{m+p})$ can be re-expressed as a cycle in $C(\mathcal{Y}^{m+p})$ by commuting the chain maps ϕ^* past the inclusions i^* by using the fact that Φ is a morphism of chain filtrations. More precisely,

$$\begin{aligned} \phi^{m,p}(\mathbf{x}) &= \phi^{m,p} \circ i^{m,p}(\mathbf{x}') \text{ by definition of } \mathbf{x} \\ &= \phi^{m,p} \circ i^{m+p-1} \circ \dots \circ i^m(\mathbf{x}') \text{ by definition of } i^{m,p} \\ &= \phi^{m+p} \circ i^{m+p-1} \circ i^{m+p-2} \circ \dots \circ i^m(\mathbf{x}') \text{ by definition of } \phi^{m,p} \\ &= j^{m+p-1} \circ \phi^{m+p-1} \circ i^{m+p-2} \circ \dots \circ i^m(\mathbf{x}') \text{ commuting } \phi \text{ past } i \\ &= \dots \text{ repeatedly commute} \\ &= j^{m+p-1} \circ \dots \circ j^m \circ \phi^m(\mathbf{x}') \\ &= j^{m,p} \circ \phi^m(\mathbf{x}') \text{ by definition of } j^{m,p} \end{aligned}$$

Since $\phi^m : C(\mathcal{X}^m) \rightarrow C(\mathcal{Y}^m)$ is a chain map, we recognize $\phi^m(\mathbf{x}')$ as a cycle in $C(\mathcal{Y}^m)$ which is then included via $j^{m,p}$ to a cycle in $C(\mathcal{Y}^{m+p})$, as desired. For well-definedness, consider a boundary $\mathbf{y} \in B(\mathcal{X}^{m+p})$. Then, $\exists \mathbf{z} \in C^{m+p}(\mathcal{X})$ so that $\mathbf{y} = \partial^{m+p}(\mathbf{z})$. Immediately, we see that

$$\begin{aligned} \phi^{m,p}(\mathbf{y}) &= \phi^{m+p}(\mathbf{y}) \text{ by definition} \\ &= \phi^{m+p} \circ \partial^{m+p} \mathbf{z} \text{ by definition of } \mathbf{y} \\ &= \delta^{m+p} \circ \phi^{m+p}(\mathbf{z}) \text{ since } \phi \text{ is a chain map} \\ &\in B(\mathcal{Y}^{m+p}) \end{aligned}$$

Thus, $\phi^{m,p}(\mathbf{y})$ lies in $B(\mathcal{Y}^{m+p})$ and is consequently trivial on the level of persistent homology as desired. \square

In light of the preceding proposition, we will unambiguously refer to *the morphisms on persistent homology modules induced by a filtered chain map*. We will say that filtered chain maps Φ from F to G and Ψ from G to F are *filtered chain equivalent* if they constitute weak inverse morphisms of the associated chain filtrations. The following proposition is a direct consequence of these basic definitions.

PROPOSITION 2.4.3. If Φ and Ψ are filtered chain equivalences between filtrations F and G , then they induce isomorphisms on persistent homology modules.

PROOF. Let $\{\Theta_n^m : C_n(\mathcal{X}^m) \rightarrow C_{n+1}(\mathcal{X}^m)\}$ be any sequence of chain homotopies which realize a weak equivalence between the composition $\Psi \circ \Phi$ and the identity id_F . We will pick arbitrary m and p and demonstrate an isomorphism between the persistent homology modules $H^p(\mathcal{X}^m)$ and $H^p(\mathcal{Y}^m)$. The central idea is to make use of the chain homotopy property of Θ^{m+p} , namely

$$\psi^{m+p} \circ \phi^{m+p} - \text{id}_F \equiv \Theta^{m+p} \circ \partial^{m+p} + \partial^{m+p} \circ \Theta^{m+p} \quad (2.4)$$

Pick any cycle $\mathbf{x}' \in Z(\mathcal{X}^m)$, include it via $i^{m,p}$ to $Z(\mathcal{X}^{m+p})$ and call this included cycle \mathbf{x} . Letting $[\mathbf{x}]$ denote the equivalence class of \mathbf{x} in the persistent homology module $H^p(\mathcal{X}^m)$, we also know that $\phi_*^{m,p}([\mathbf{x}])$ is a well-defined element of $H^p(\mathcal{Y}^m)$. But now,

$$\begin{aligned}
\psi_*^{m,p} \circ \phi_*^{m,p}([\mathbf{x}]) &= [\psi^{m,p} \circ \phi^{m,p}(\mathbf{x})] \\
&= [\psi^{m+p} \circ \phi^{m+p}(\mathbf{x})], \text{ by definitions of } \phi^{m,p} \text{ and } \psi^{m,p} \\
&= [\mathbf{x} + \Theta^{m+p} \circ \partial^{m+p} \mathbf{x} + \partial^{m+p} \circ \Theta^{m+p} \mathbf{x}], \text{ by (2.4)} \\
&= [\mathbf{x} + \partial^{m+p} \circ \Theta^{m+p} \mathbf{x}] \text{ since } \partial^{m+p} \mathbf{x} = 0 \text{ for } \mathbf{x} \in Z(\mathcal{X}^{m+p}) \\
&= [\mathbf{x}], \text{ since the latter term lies in } B(\mathcal{X}^{m+p})
\end{aligned}$$

Reversing the order of composition and using a chain homotopy relating $\Phi \circ \Psi$ to $\text{id}_{\mathcal{G}}$ in a similar argument establishes that $\phi_*^{m,p}$ and $\psi_*^{m,p}$ are indeed inverses, and hence that the persistent homology modules $H^p(\mathcal{X}^m)$ and $H^p(\mathcal{Y}^m)$ are isomorphic. \square

2.5. Combinatorial Topology

We review some classical results related to nerves (used both to provide representations of spaces and to generate associated chain complexes), carriers (used to obtain chain maps between the spaces), and outer approximations (used to approximate continuous functions).

2.5.1. The Nerve Lemma. Let M be a topological space and let \mathcal{U} be a finite open cover of M . Recall [31, 36] that the *nerve* $\mathcal{N}(\mathcal{U})$ is the abstract simplicial complex where each simplex of dimension k corresponds to a non-empty intersection of $k+1$ distinct elements of \mathcal{U} .

Given the simplex σ defined by the intersection of a subcollection $\mathcal{S} \subset \mathcal{U}$, the *support* of σ is defined to be the non-empty intersection

$$[\sigma] = \bigcap_{U \in \mathcal{S}} U$$

The support of a simplicial subcomplex $\mathcal{K} \subset \mathcal{N}(\mathcal{U})$ is the union of the supports of the constituent simplices: $[\mathcal{K}] = \bigcup_{\sigma \in \mathcal{K}} [\sigma]$. On the other hand, given any subset $S \subset M$, we define its *hull* $\mathbf{h}_{\mathcal{U}}[S]$ to be the simplicial subcomplex of $\mathcal{N}(\mathcal{U})$ generated by the collection of simplices whose supports intersect S , i.e., $\{\sigma \in \mathcal{N}(\mathcal{U}) \mid S \cap [\sigma] \neq \emptyset\}$.

We call $\mathcal{N}(\mathcal{U})$ a *contractible nerve* if the support of each simplex $\sigma \in \mathcal{N}(\mathcal{U})$ is contractible to a single point. If M lies in a topological vector space and if each $U \in \mathcal{U}$ is convex, then the non-empty intersections are also convex and hence contractible. Thus, the nerve of a convex open cover is always contractible.

Ordering the elements of \mathcal{U} orders the 0-simplices of $\mathcal{N}(\mathcal{U})$ and hence induces an orientation on $\mathcal{N}(\mathcal{U})$. The k -chains $C_k(\mathcal{U})$ of $\mathcal{N}(\mathcal{U})$ are the free abelian group generated by the oriented k -simplices. The *support* of a chain $\mathbf{x} = \sum_j m_j \sigma_j$ is the union of the supports of the underlying simplices:

$$[\mathbf{x}] = \bigcup_j [\sigma_j]$$

The *simplicial boundary operator* $\partial_k : C_k(\mathcal{U}) \rightarrow C_{k-1}(\mathcal{U})$ is defined on the k -simplex $\sigma_{\mathcal{S}}$ corresponding to the ordered set $\mathcal{S} = \{U_1, \dots, U_{k+1}\}$ by the usual formula (see Example 2.3.2)

$$\partial_k(\sigma_{\mathcal{S}}) = \sum_{j=1}^{k+1} (-1)^j \sigma_{\mathcal{S} \setminus \{U_j\}}$$

It is straightforward to check that $\partial \circ \partial \equiv 0$, and so one obtains the chain complex

$$\dots \xrightarrow{\partial_{k+1}} C_k(\mathcal{U}) \xrightarrow{\partial_k} C_{k-1}(\mathcal{U}) \xrightarrow{\partial_{k-1}} \dots \xrightarrow{\partial_1} C_0(\mathcal{U})$$

The simplicial homology groups of $\mathcal{N}(\mathcal{U})$ are defined to be the usual quotient

$$H_k^\Delta(\mathcal{N}(\mathcal{U})) = \frac{\ker \partial_k}{\text{img } \partial_{k+1}}$$

Finally, we have the *nerve lemma* (see [23, Thm. 15.21] for the proof).

LEMMA 2.5.1. *Let M be a paracompact topological space and \mathcal{U} a finite open cover of M . If the associated nerve $\mathcal{N}(\mathcal{U})$ is acyclic, then M is homotopy equivalent to $|\mathcal{N}(\mathcal{U})|$.*

Since homology is a homotopy invariant, we obtain as an immediate corollary a canonical isomorphism $H_*^\Delta(\mathcal{N}(\mathcal{U})) \simeq H_*(M)$, where $H_*(M)$ denotes the singular homology groups of M .

2.5.2. The Acyclic Carrier Theorem. Consider two finite abstract simplicial complexes Σ and Γ with corresponding simplicial boundary operators ∂^Σ and ∂^Γ . Let $\langle \cdot, \cdot \rangle_\Sigma$ and $\langle \cdot, \cdot \rangle_\Gamma$ be the usual inner products on chains obtained by treating the simplices as an orthonormal basis. Denote by \prec the partial order on simplices in Σ generated by the transitive closure of the following face relation

$$\sigma' \prec \sigma \text{ if } \langle \partial^\Sigma \sigma, \sigma' \rangle_\Sigma \neq 0$$

DEFINITION 2.5.2. An *acyclic carrier* $F : \Sigma \rightrightarrows \Gamma$ is a map from Σ to subcomplexes of Γ which satisfies the following requirements

- (1) **acyclicity:** for every $\sigma \in \Sigma$, $F(\sigma)$ is a (non-empty) closed acyclic simplicial subcomplex of Γ , and
- (2) **semicontinuity:** for simplices $\sigma, \sigma' \in \Sigma$ with $\sigma \prec \sigma'$, we have the inclusion $F(\sigma) \subset F(\sigma')$ of simplicial complexes.

Let $F : \Sigma \rightrightarrows \Gamma$ be fixed throughout the rest of this section. Given another acyclic carrier $G : \Sigma \rightrightarrows \Gamma$, we write $F \subseteq G$ to indicate $F(\sigma) \subset G(\sigma)$ for each $\sigma \in \Sigma$. It is easy to check that \subseteq induces a partial order on the set of all acyclic carriers from Σ to Γ . We say that F *carries* the abelian group morphisms $\omega : C_k(\Sigma) \rightarrow C_\ell(\Gamma)$ if for each $\sigma \in \Sigma$, we have

$$\{\gamma \in \Gamma \mid \langle \omega(\sigma), \gamma \rangle_\Gamma \neq 0\} \subset F(\sigma)$$

The *acyclic carrier theorem* guarantees existence and uniqueness (up to chain homotopy) of chain maps carried by acyclic carriers. See [31, Ch 1, Thm 13.3] for details.

THEOREM 2.5.3. *Let $F : \Sigma \rightrightarrows \Gamma$ be an acyclic carrier. Then, we have*

- (1) **existence:** *there exists a chain map carried by F , and*
- (2) **uniqueness:** *if F carries two chain maps ϕ and ψ , then F also carries a chain homotopy between ϕ and ψ .*

As an immediate corollary, we find that given any two chain maps ϕ and ψ which share an acyclic carrier, the induced maps on homology satisfy $\phi_* \equiv \psi_*$.

2.5.3. Outer Approximations. Throughout this section, let M and N be topological spaces with finite open covers \mathcal{U} and \mathcal{V} respectively so that the nerves $\mathcal{N}(\mathcal{U})$ and $\mathcal{N}(\mathcal{V})$ are acyclic. Also fix a continuous function $f : M \rightarrow N$. We will follow the strategy of [21, Ch 6] for approximating f by a suitable acyclic carrier.

We say that an acyclic carrier $F : \mathcal{N}(\mathcal{U}) \rightrightarrows \mathcal{N}(\mathcal{V})$ is an *outer approximation* of f if for each $\sigma \in \mathcal{N}(\mathcal{U})$, we have the basic containment $f(|\sigma|) \subset |F(\sigma)|$. In general, it is not true that f must have an outer approximation because there is no guarantee that given a $\sigma \in \mathcal{N}(\mathcal{U})$ the image $f(|\sigma|)$ is contained in the support of some acyclic subcomplex of $\mathcal{N}(\mathcal{V})$.

Define the correspondence \mathbf{M}_f from simplices of $\mathcal{N}(\mathcal{U})$ to simplicial subcomplexes of $\mathcal{N}(\mathcal{V})$ in the following way:

$$\mathbf{M}_f(\sigma) = \mathbf{h}_{\mathcal{V}}[f(\lfloor \sigma \rfloor)]$$

where $\mathbf{h}_{\mathcal{V}}$ denotes the hull in $\mathcal{N}(\mathcal{V})$.

We say that f is *approximable* if \mathbf{M}_f is an outer approximation of f , and in this case we call \mathbf{M}_f the *minimal* outer approximation of f . The key requirement for approximability is that the hull $\mathbf{h}_{\mathcal{V}}[f(\lfloor \sigma \rfloor)]$ must be an acyclic subcomplex of $\mathcal{N}(\mathcal{V})$ for each simplex $\sigma \in \mathcal{N}(\mathcal{U})$, since the containment $f(\lfloor \sigma \rfloor) \subset \lfloor \mathbf{M}_f(\sigma) \rfloor$ and the semicontinuity are both automatic from the definition of \mathbf{M}_f . We have the following proposition [21, Prop. 6.29, 6.36].

PROPOSITION 2.5.4. *Assume that $F : \mathcal{N}(\mathcal{U}) \rightrightarrows \mathcal{N}(\mathcal{V})$ is an outer approximation of an approximable function f . Then, we have*

- (1) **minimality:** $\mathbf{M}_f \subseteq F$, and consequently
- (2) **universality:** If \mathbf{M}_f carries a chain map ϕ , then so does F .

Thus, if $f : M \rightarrow N$ is approximable then by the preceding proposition and the acyclic carrier theorem, any chain map ϕ carried by any outer approximation F of f descends to the same module morphism $\phi_* : H_*^{\Delta}(\mathcal{N}(\mathcal{U})) \rightarrow H_*^{\Delta}(\mathcal{N}(\mathcal{V}))$. We call ϕ_* a *simplicial reconstruction* of f_* .

Note that while ϕ_* is independent of the choices of F and ϕ , there remains an explicit dependence on the choice of covers \mathcal{U} and \mathcal{V} which generate the acyclic nerves. However, we have the following commutative diagram for any such \mathcal{U} and \mathcal{V} :

$$\begin{array}{ccc} H_*(M) & \xrightarrow{f_*} & H_*(N) \\ \cong \downarrow & & \downarrow \cong \\ H_*^{\Delta}(\mathcal{N}(\mathcal{U})) & \xrightarrow{\phi_*} & H_*^{\Delta}(\mathcal{N}(\mathcal{V})) \end{array}$$

where the vertical isomorphisms arise from the nerve lemma.

CHAPTER 3

Recovering Manifolds and Functions from Data

*Fácilis descensus Averni:
noctes atque dies patet atri ianua Ditis;
sed revocare gradium superasque evadere ad auras.
Hoc opus, hic labor est.*

–Virgil, *The Aenid*

3.1. Introduction

Perhaps the most ubiquitous form of experimental data is a *point cloud*, i.e., a finite collection of points in Euclidean space \mathbb{R}^n . Such data may arise from *sampling* an unknown topological subspace S of \mathbb{R}^n according to some distribution concentrated on (or near) S . As mentioned in Chapter 1, the first step of topological analysis is the construction of a cell complex \mathcal{K} around this data. There are a variety of well-known methods for constructing a simplicial complex whose vertices coincide with a given point cloud, and we refer the interested reader to [8] for an overview of the most common constructions. Since the homology modules of a cell complex are algorithmically computable, a natural question to ask is *when are the simplicial homology modules $H_*^\Delta(\mathcal{K})$ of \mathcal{K} isomorphic to the singular homology modules $H_*(S)$ of S ?* Under these conditions, it becomes possible to know S up to homology using only a finite amount of sampled data.

Sufficient conditions for homology-preserving simplicial reconstruction of compact Riemannian submanifolds of Euclidean space were provided by Partha Niyogi, Steve Smale and Shmuel Weinberger in [32]. Section 3.2 contains a brief overview of their results. Section 3.3 presents the theory and algorithms from [28] for achieving simplicial reconstruction of Lipschitz-continuous *functions* between such manifolds. In Section 3.4, we prove that under suitable hypotheses this reconstruction of functions is robust to sampling noise. Finally, we note in Section 3.5 that the methodology of simplicial reconstruction is useful even when various assumptions requiring manifold structure and Lipschitz continuity are relaxed. Here, we relate *transformations of data* to *morphisms of persistent homology modules* in the absence of any strong assumptions.

3.2. Reconstructing Manifolds from Sampled Data

Let $\mathcal{X} \subset \mathbb{R}^n$ be a compact k -dimensional Riemannian submanifold. The *condition number* $1/\tau_{\mathcal{X}}$ of \mathcal{X} is defined as follows: $\tau_{\mathcal{X}}$ is the largest positive real number such that for any $r \in (0, \tau_{\mathcal{X}})$, the normal bundle of radius r about \mathcal{X} can be embedded in \mathbb{R}^n .

Given $x \in \mathbb{R}^n$ and $r > 0$, denote by $\mathcal{B}_r(x)$ the n -dimensional Euclidean open ball of radius r centered at x . Define the one-parameter family $\mathcal{N}_\epsilon(X)$ to be the nerve complexes generated by the the collections $\mathcal{U}_\epsilon = \{\mathcal{B}_\epsilon(\xi) \mid \xi \in X\}$ of open subsets of \mathbb{R}^n .

DEFINITION 3.2.1. The *bounding function* $\beta_{\mathcal{X}} : \mathbb{R}^+ \times (0, 1] \rightarrow \mathbb{R}$ is given by

$$\beta_{\mathcal{X}}(\epsilon, \delta) = \beta_1 [\log(\beta_2) - \log(1/\delta)], \quad (3.1)$$

where

$$\beta_1 = \frac{\text{vol}(\mathcal{X})}{\cos^k(\sin^{-1}(\frac{\epsilon}{8\tau_{\mathcal{X}}})) \cdot \text{vol}(\mathcal{B}_{\epsilon/4})}$$

$$\beta_2 = \frac{\text{vol}(\mathcal{X})}{\cos^k(\sin^{-1}(\frac{\epsilon}{16\tau_{\mathcal{X}}})) \cdot \text{vol}(\mathcal{B}_{\epsilon/8})}$$

and $\text{vol}(\mathcal{B}_{\epsilon})$ denotes the usual volume of the standard k -dimensional ϵ -ball.

The following notion of density is standard in metric subspace topology: given $\epsilon > 0$, we say that a set $S \subset \mathbb{R}^n$ is ϵ -dense in \mathcal{X} if for every $x \in \mathcal{X}$ there exists some $s \in S$ so that $x \in \mathcal{B}_{\epsilon}(s)$. The next proposition enables the reconstruction of the manifold \mathcal{X} up to homology from a finite point set $X \subset \mathbb{R}^n$ provided that X is sufficiently dense in \mathcal{X} relative to $\tau_{\mathcal{X}}$.

PROPOSITION 3.2.2. ([32, Prop 3.1]) *Assume $\epsilon \in (0, \sqrt{3/5} \tau_{\mathcal{X}})$ and that X is $\epsilon/2$ -dense in \mathcal{X} . Then, the canonical projection map $\pi_{\mathcal{X}} : [\mathcal{N}_{\epsilon}(X)] \rightarrow \mathcal{X}$ defined by*

$$\pi_{\mathcal{X}}(w) = \arg \min_{x \in \mathcal{X}} \|w - x\|_{\mathbb{R}^n} \quad (3.2)$$

is a deformation-retract.

As a corollary to this proposition one obtains the isomorphisms

$$H_*(M) \simeq H_*([\mathcal{N}_{\epsilon}(X)]) \simeq H_*^{\Delta}(\mathcal{N}_{\epsilon}(X))$$

The first isomorphism comes from the fact that deformation retracts preserve homotopy type. The second isomorphism results from applying the nerve lemma: since \mathcal{U}_{ϵ} is a convex cover of $[\mathcal{N}_{\epsilon}(X)]$ for each ϵ , the associated nerve is contractible. A key step in proving Proposition 3.2.2 involves understanding the fibers $\pi_{\mathcal{X}}^{-1}(x)$ for each $x \in \mathcal{X}$. The following lemma is useful in this regard.

LEMMA 3.2.3. ([32, Lem 4.1]) *Assume the hypotheses of Proposition 3.2.2 and let $\xi \in X$. Given some $w \in \mathcal{B}_{\epsilon}(\xi)$ with $\pi_{\mathcal{X}}(w) \notin \mathcal{B}_{\epsilon}(\xi)$, we must have $\pi_{\mathcal{X}}(w) \in \mathcal{B}_{\epsilon^2/\tau_{\mathcal{X}}}(\xi)$.*

Note that since $\epsilon < \tau_{\mathcal{X}}/2$ by hypothesis, we have $\epsilon^2/\tau_{\mathcal{X}} < \epsilon/2$. Thus, if we assume the hypotheses of Lemma 3.2.3 then we may conclude that $\|\pi_{\mathcal{X}}(w) - \xi\| < 3\epsilon/2$ by a simple application of the triangle inequality. This observation yields the following corollary.

COROLLARY 3.2.4. *Given $\sigma \in \mathcal{N}_{\epsilon}(X)$ determined by a subcollection \mathcal{S} of \mathcal{U}_{ϵ} , we have*

$$\pi_{\mathcal{X}}([\sigma]) \subset \bigcap_{\xi \in \mathcal{S}} \mathcal{B}_{\frac{3\epsilon}{2}}(\xi),$$

where $[\sigma]$ is the support of σ in the nerve $\mathcal{N}_{\epsilon}(X)$.

The following proposition assumes that X is obtained by uniform i.i.d. sampling on \mathcal{X} and provides a lower bound on the sample size $\#X$ which guarantees – with high confidence – the $\epsilon/2$ -density needed by the previous proposition.

PROPOSITION 3.2.5. ([32, Prop 3.2]) *Choose $\epsilon \in (0, \tau_{\mathcal{X}}/2)$ and the probability parameter $\delta \in (0, 1]$. Assume that X is obtained by i.i.d. uniform samplings from \mathcal{X} . If $\#X > \beta_{\mathcal{X}}(\epsilon, \delta)$, then X is $\epsilon/2$ -dense in \mathcal{X} with probability exceeding $(1 - \delta)$.*

Propositions 3.2.2 and 3.2.5 lead directly to the following main theorem of [32].

THEOREM 3.2.6. *Let \mathcal{X} be a compact k -dimensional Riemannian submanifold of \mathbb{R}^n with condition number $1/\tau_{\mathcal{X}}$. Given*

- (1) some probability parameter $\delta \in (0, 1]$,
- (2) a radius $\epsilon < \tau_x/2$, and
- (3) a finite set $X \subset \mathcal{X}$ of independent and identically distributed (i.i.d.) uniformly sampled points,

let $\mathcal{N}(X)$ denote the nerve generated by n -dimensional open ϵ -balls centered at the points in X . If the sample size $\#X$ is larger than the bounding value $\beta_x(\epsilon, \delta)$, then $\mathcal{N}(X)$ deformation-retracts to X with probability exceeding $(1 - \delta)$.

3.3. Simplicial Reconstruction of Functions

In this section we provide a proof of the following Theorem.

THEOREM 3.3.1. *Let $X \subset \mathbb{R}^n$ and $Y \subset \mathbb{R}^m$ be compact Riemannian submanifolds with condition numbers $1/\tau_x$ and $1/\tau_y$ respectively and let $f : X \rightarrow Y$ be a Lipschitz continuous function with Lipschitz constant κ . Given*

- (1) probability parameters $\delta_x, \delta_y \in (0, 1]$,
- (2) radii $\epsilon_x < \tau_x/2$ and $\epsilon_y < \tau_y/2$ satisfying $3\kappa \cdot \epsilon_x < \epsilon_y$, and
- (3) finite sets $X \subset \mathcal{X}$ and $Y \subset \mathcal{Y}$ of independent and identically distributed (i.i.d.) uniformly sampled points,

let $\mathcal{N}(X)$ and $\mathcal{N}(Y)$ denote the nerve complexes generated by open balls of radius ϵ_x and ϵ_y around X and Y respectively and assume that $\#X > \beta_x(\epsilon_x, \delta_x)$ and $\#Y > \beta_y(\epsilon_y, \delta_y)$. Then, there exist algorithms which take X, Y and $f(X)$ as input and produce the following output with probability exceeding $(1 - \delta_x) \cdot (1 - \delta_y)$:

- (1) An acyclic outer approximation $F : \mathcal{N}(X) \rightrightarrows \mathcal{N}(Y)$ of f , and
- (2) A chain map $\phi : \mathcal{N}(X) \rightarrow \mathcal{N}(Y)$ carried by F .

3.3.1. Hypotheses and Assumptions. Observe that the hypothesis consists of a variety of assumptions and a-priori choices of parameters. To clarify their respective roles we present them via the following exhaustive list.

Cnd: $X \subset \mathbb{R}^n$ and $Y \subset \mathbb{R}^m$ are compact Riemannian submanifolds with condition numbers $1/\tau_x$ and $1/\tau_y$, respectively.

Lip: $f : X \rightarrow Y$ is a Lipschitz continuous function with Lipschitz constant κ .

Rad: The radii $\epsilon_x \in (0, \tau_x/2)$ and $\epsilon_y \in (0, \tau_y/2)$ satisfy $3\kappa \cdot \epsilon_x < \epsilon_y$.

Prb: The probability parameters $\delta_x, \delta_y \in (0, 1]$.

Smp: We know the finite point sets $X \subset \mathcal{X}$ and $Y \subset \mathcal{Y}$ obtained by i.i.d. uniform sampling from \mathcal{X} and \mathcal{Y} respectively. Furthermore, $\#X > \beta_x(\epsilon_x, \delta_x)$ and $\#Y > \beta_y(\epsilon_y, \delta_y)$.

Img: We know the pairs $\{(\xi, f(\xi)) \in X \times Y \mid \xi \in X\}$ which catalogs the images under f of the data sample X .

It is important to note that neither **Smp** nor **Img** imply that sampled points map to sampled points, so in general $f(X) \not\subset Y$. Since ϵ_x and ϵ_y are now fixed by the choices in **Rad**, we simplify the notation by declaring that $\mathcal{N}(X) := \mathcal{N}_{\epsilon_x}(X)$ and $\mathcal{N}(Y) := \mathcal{N}_{\epsilon_y}(Y)$. Also set

$$\rho = \epsilon_y - \frac{3}{2}\kappa\epsilon_x \quad (3.3)$$

3.3.2. Constructing an Acyclic Outer Approximation. We begin with two elementary algorithms. The first – called Approximator – constructs a correspondence $\mathcal{L} : X \rightarrow 2^Y$ which induces an acyclic outer approximation $F : \mathcal{N}(X) \rightrightarrows \mathcal{N}(Y)$ of f . The second algorithm – called Selector – leads to the selection of a chain map $\phi_{\#}$ carried by F .

Approximator relies on the input data from **Smp** and **Img** and constructs the correspondence $\mathcal{L} : X \rightarrow 2^Y$ defined by

$$\mathcal{L}(\xi) = \{\eta \in Y \mid \|\eta - f(\xi)\|_{\mathbb{R}^m} < \rho\} \quad (3.4)$$

for each $\xi \in X$. The triangle inequality implies that

$$\eta \in \mathcal{L}(\xi) \text{ if and only if } \mathcal{B}_{\frac{3}{2}\kappa \cdot \epsilon_X}(f(\xi)) \subset \mathcal{B}_{\epsilon_Y}(\eta) \quad (3.5)$$

\mathcal{L} is constructed in the obvious way by iterating over Y once for each element of X , and so the algorithm terminates after at most $\#X \cdot \#Y$ total iterations.

Observe that Approximator returns a failed status precisely when there exists some $\xi \in X$ with $\mathcal{L}(\xi) = \emptyset$. We begin our analysis of this algorithm by bounding the probability of success from below.

TABLE 1. **Algorithm:** Approximator

Input: $(X, f(X), Y)$; **Output:** $(status, \mathcal{L})$

1	for each $\xi \in X$
2	define $\mathcal{L}(\xi) = \emptyset$
3	for each $\eta \in Y$
4	if $\ f(\xi) - \eta\ _{\mathbb{R}^m} < \rho$
5	add η to $\mathcal{L}(\xi)$
6	end if
7	end for (η)
8	if $\mathcal{L}(\xi) = \emptyset$
9	return $(failure, \mathcal{L})$
10	end if
11	end for (ξ)
12	return $(success, \mathcal{L})$

PROPOSITION 3.3.2. Approximator returns a successful status on termination with probability exceeding $(1 - \delta_Y)$.

PROOF. Noting that the assumptions **Cnd** and **Smp** and the choices **Rad** and **Prb** satisfy the hypotheses of Proposition 3.2.5 for \mathcal{Y} , we see that Y is $\epsilon_Y/2$ -dense in \mathcal{Y} with probability exceeding $(1 - \delta_Y)$.

The following argument shows that that this density suffices to guarantee a non-empty $\mathcal{L}(\xi)$ for each $\xi \in X$. For any $\xi \in X$ there exists $\eta \in Y$ such that $f(\xi) \in \mathcal{B}_{\epsilon_Y/2}(\eta)$. Thus

$$\mathcal{B}_{\frac{3}{2}\kappa \epsilon_X}(f(\xi)) \subset \mathcal{B}_{\frac{3}{2}\kappa \epsilon_X + \frac{1}{2}\epsilon_Y}(\eta) \subset \mathcal{B}_{\epsilon_Y}(\eta)$$

where the last inclusion follows from **Rad**. Thus, $\eta \in \mathcal{L}(\xi) \neq \emptyset$ by (3.5). Hence Approximator terminates successfully since the failing condition in line 5 is never satisfied. \square

Assuming that Approximator has terminated with successful status producing \mathcal{L} we now turn to the question of constructing an acyclic outer approximation. For convenience of presentation let $\mathbf{g}_X : \mathcal{N}(X) \rightarrow 2^X$ denote the injection taking each simplex $\sigma \in \mathcal{N}(X)$ to its *defining subset* $S \subset X$ such that σ has the support $[\sigma] = \bigcap_{\xi \in S} \mathcal{B}_{\epsilon_X}(\xi)$. Note that $\sigma \prec \sigma'$ in $\mathcal{N}(X)$ if and only if we have the inclusion $\mathbf{g}_X(\sigma) \subset \mathbf{g}_X(\sigma')$. Define $\mathbf{g}_Y : \mathcal{N}(Y) \rightarrow 2^Y$ similarly.

\mathcal{L} induces a correspondence $\mathbf{G} : \mathcal{N}(X) \rightarrow 2^Y$ via the following definition

$$\mathbf{G}(\sigma) = \{\eta \in Y \mid \eta \in \mathcal{L}(\xi) \text{ for some } \xi \in \mathbf{g}_X(\sigma)\} \quad (3.6)$$

PROPOSITION 3.3.3. *With probability exceeding $(1 - \delta_X)$, the following holds. For each $\sigma \in \mathcal{N}(X)$ and $\eta \in \mathbf{G}(\sigma)$, we have*

$$\emptyset \neq f \circ \pi_X([\sigma]) \subset \mathcal{B}_{\epsilon_Y}(\eta) \quad (3.7)$$

where $\pi_X : [\mathcal{N}(X)] \rightarrow \mathcal{X}$ is the canonical projection from (3.2).

PROOF. By assumptions **Cnd** and **Smp** and choices **Rad** and **Prb**, Proposition 3.2.5 holds for \mathcal{X} . Therefore, with probability exceeding $(1 - \delta_X)$ we are guaranteed the requisite density of X in \mathcal{X} to apply Proposition 3.2.2. Assuming this density, pick $w \in [\sigma]$. Set $x = \pi_X(w) \in \mathcal{X}$ and note from Corollary 3.2.4 that $x \in \bigcap_{\xi \in \mathbf{g}_X(\sigma)} \mathcal{B}_{\frac{3}{2}\epsilon_X}(\xi)$. By **Lip**,

$$f(x) \in \bigcap_{\xi \in \mathbf{g}_X(\sigma)} \mathcal{B}_{\frac{3}{2}\kappa\epsilon_X}(f(\xi)). \quad (3.8)$$

By definition, $\eta \in \mathbf{G}(\sigma)$ implies the existence of some $\xi_* \in \mathbf{g}_X(\sigma)$ so that $\eta \in \mathcal{L}(\xi_*)$. By (3.5), we now have

$$\mathcal{B}_{\frac{3}{2}\kappa\epsilon_X}(f(\xi_*)) \subset \mathcal{B}_{\epsilon_Y}(\eta).$$

Using (3.8) we see that $f(x) \in \mathcal{B}_{\frac{3}{2}\kappa\epsilon_X}(f(\xi_*))$, and so by **Rad** $f(x) \in \mathcal{B}_{\epsilon_Y}(\eta)$ as desired. \square

COROLLARY 3.3.4. *With probability exceeding $(1 - \delta_X)$, the following holds. For each $\sigma \in \mathcal{N}(X)$, the set $\mathbf{G}(\sigma) \subset Y$ defines a simplex in $\mathcal{N}(Y)$*

PROOF. We must show that the intersection $\bigcap_{\eta \in \mathbf{G}(\sigma)} \mathcal{B}_{\epsilon_Y}(\eta)$ is non-empty. By Proposition 3.3.3, the non-empty set $f \circ \pi_X([\sigma]) \subset \mathcal{B}_{\epsilon_Y}(\eta)$ for each $\eta \in \mathbf{G}(\sigma)$, and so $\mathbf{G}(\sigma)$ defines a simplex in $\mathcal{N}(Y)$. \square

We now assume that **Approximator** terminates with a successful status and that (3.7) holds. Observe that the probability of these conditions be satisfied simultaneously exceeds $(1 - \delta_X) \cdot (1 - \delta_Y)$.

For each $\sigma \in \mathcal{N}(X)$ define γ_σ in $\mathcal{N}(Y)$ by

$$\gamma_\sigma = \mathbf{g}_Y^{-1}(\mathbf{G}(\sigma)). \quad (3.9)$$

Define the correspondence F from $\mathcal{N}(X)$ to subcomplexes of $\mathcal{N}(Y)$ by

$$F(\sigma) = \{\gamma \in \mathcal{N}(Y) \mid \gamma \preceq \gamma_\sigma\}. \quad (3.10)$$

PROPOSITION 3.3.5. $F : \mathcal{N}(X) \rightrightarrows \mathcal{N}(Y)$ is an acyclic carrier.

PROOF. As indicated in Definition 2.5.2 there are two conditions to check: **acyclicity** and **semicontinuity**. To verify acyclicity observe that by (3.10) $F(\sigma)$ consists of the sub complex generated by the single simplex γ_σ . For semicontinuity, consider another simplex $\sigma' \in \mathcal{N}(X)$ so that $\sigma \prec \sigma'$, in which case $\mathbf{g}_X(\sigma) \subset \mathbf{g}_X(\sigma')$. By (3.6), $\mathbf{G}(\sigma) \subset \mathbf{G}(\sigma')$. Thus, $\gamma_\sigma \prec \gamma_{\sigma'}$ in $\mathcal{N}(Y)$ and so we have the desired inclusion $F(\sigma) \subset F(\sigma')$. \square

We would like to claim that F is an acyclic outer approximation of f . Unfortunately, for the following technical reasons this is not true: the domain of f is \mathcal{X} which is a strict subset of $[\mathcal{N}(X)]$. However, we can use the retraction π_X to expand the domain of f without altering the map f induces on homology. Define the function $\tilde{f} : [\mathcal{N}(X)] \rightarrow \mathcal{Y} \subset [\mathcal{N}(Y)]$ by $\tilde{f} \equiv f \circ \pi_X$ and note that

- (1) $\tilde{f}|_{\mathcal{X}} \equiv f$, and
- (2) \tilde{f} is homotopic to f via $\Theta : [0, 1] \times [\mathcal{N}(X)] \rightarrow [\mathcal{N}(Y)]$ defined by $\Theta(t, z) = f \circ \pi_X((1 - t) \cdot z + t \cdot \pi_X(z))$.

THEOREM 3.3.6. F is an acyclic outer approximation of $\tilde{f} : [\mathcal{N}(X)] \rightarrow [\mathcal{N}(Y)]$.

PROOF. Consider $\sigma \in \mathcal{N}(X)$ and note that $\tilde{f}([\sigma]) = f \circ \pi_X([\sigma])$ is contained in the intersection $\bigcap_{\eta \in \mathbf{G}(\sigma)} \mathcal{B}_{\epsilon_Y}(\eta)$ by Proposition 3.3.3. But by definition, this intersection is precisely $[\gamma_\sigma]$. Since γ_σ is a simplex in $F(\sigma)$ by (3.10), we have $[\gamma_\sigma] \subset [F(\sigma)]$ as desired. \square

Observe that this completes the proof of the first conclusion of Theorem 3.3.1.

3.3.3. Constructing a Representative Simplicial Map. We now turn to the second part which is the construction of a chain map carried by F . We begin by considering the second algorithm `Selector` which requires the output \mathcal{L} of `Approximator` as defined in (3.4) and constructs a function $h : X \rightarrow Y$ with following property: for each $\xi \in X$,

$$h(\xi) \in \mathcal{L}(\xi). \quad (3.11)$$

`Selector` constructs h by choosing an arbitrary η in $\mathcal{L}(\xi)$ for each $\xi \in X$ and hence is guaranteed to terminate after $\#X$ iterations of the `for` loop.

TABLE 2. **Algorithm:** `Selector`

Input: $(X, Y, \mathcal{L} : X \rightarrow 2^Y)$; **Output:** h

1	for each $\xi \in X$
2	choose any $\eta \in \mathcal{L}(\xi)$
3	define $h(\xi) = \eta$
4	end for
5	return h

Our goal is to use h to define a simplicial map $\phi : \mathcal{N}(X) \rightarrow \mathcal{N}(Y)$. Recall that given a simplex $\sigma \in \mathcal{N}(X)$, $\mathbf{g}_X(\sigma)$ identifies the vertices of σ with elements of X . Let $\xi \in \mathbf{g}_X(\sigma)$ and let $\eta = h(\xi)$. By (3.11) we have $\eta \in \mathcal{L}(\xi)$, and so by definition $\eta \in \mathbf{G}(\sigma)$. Therefore, for all $\sigma \in \mathcal{N}(X)$

$$h(\mathbf{g}_X(\sigma)) \subset \mathbf{G}(\sigma). \quad (3.12)$$

PROPOSITION 3.3.7. *Define $\phi : \mathcal{N}(X) \rightarrow \mathcal{N}(Y)$ by*

$$\phi(\sigma) = \mathbf{g}_Y^{-1}[h(\mathbf{g}_X(\sigma))].$$

Then ϕ is a simplicial map.

PROOF. Observe that as defined ϕ maps vertices to vertices. By Corollary 3.3.4, for any simplex $\sigma \in \mathcal{N}(X)$, $\mathbf{G}(\sigma)$ is a simplex in $\mathcal{N}(Y)$. Thus by (3.12), ϕ takes the vertices of the simplex σ into a simplex and therefore ϕ is a simplicial map. \square

Let $\phi_\#$ denote the chain map generated by the simplicial map ϕ . More precisely, here is the action of $\phi_\#$ on each simplex $\sigma \in \mathcal{N}(X)$

$$\phi_\#(\sigma) = \begin{cases} \phi(\sigma) & \text{if } \dim \sigma = \dim \phi(\sigma) \\ 0 & \text{otherwise.} \end{cases}$$

PROPOSITION 3.3.8. *The chain map $\phi_\#$ is carried by the acyclic carrier F .*

PROOF. Pick $\sigma \in \mathcal{N}(X)$ and assume without loss of generality that the image $\phi_\#(\sigma)$ is non-trivial. From Proposition 3.3.7 and (3.12), we have $\phi_\#(\sigma) = \phi(\sigma) \prec \gamma_\sigma$ and so $\phi(\sigma)$ is a simplex in $F(\sigma)$ by (3.10). \square

3.4. Robustness to Conditioned Noise

As is indicated in the Introduction, we extend the results of the previous section to the case where data samples are assumed to lie near, rather than on, the underlying manifolds. This error in sampling further cascades into imprecise knowledge of the images of the samples, particularly if the evaluation of the function is also characterized by some inherent errors.

3.4.1. Conditioned Noise. We adopt the model of noise from [32, Sec. 7].

DEFINITION 3.4.1. Given $r > 0$, a probability measure μ on \mathbb{R}^n is called r -conditioned about \mathcal{X} if it is

- (1) **Locally Supported:** the support of μ is contained in the tubular neighborhood $\text{Tub}_r(\mathcal{X})$ of radius r about \mathcal{X} , and
- (2) **Regular:** for each $s \in (0, r)$, there exists some *regularity constant* $\Omega_s > 0$ so that

$$\inf_{x \in \mathcal{X}} \mu(\mathcal{B}_s^n(x)) > \Omega_s$$

where $\mathcal{B}_s^n(x)$ denotes the n -dimensional open ball of radius s about x .

As in the previous case, the following fundamental results concerning sampling of the manifolds are taken from [32]. Let \mathcal{X} be a compact Riemannian submanifold of \mathbb{R}^n with condition number $1/\tau_{\mathcal{X}}$. For $r > 0$ define the functions

$$\Gamma_{\mathcal{X}}^{\pm}(r) = \frac{(r + \tau_{\mathcal{X}}) \pm \sqrt{\tau_{\mathcal{X}}^2 + r^2 - 6\tau_{\mathcal{X}}r}}{2}$$

and note that $0 < \Gamma_{\mathcal{X}}^- < \Gamma_{\mathcal{X}}^+$ when the quantity under the square root is strictly positive. It is straightforward to check that this positivity holds for $r < (3 - \sqrt{8})\tau_{\mathcal{X}}$. Pick such an r and assume that $\hat{X} \subset \mathbb{R}^n$ is a finite point set contained in $\text{Tub}_r(\mathcal{X})$. For each $\epsilon > 0$, define $\mathcal{N}_{\epsilon}(\hat{X})$ to be the acyclic nerve complex generated by open balls of radius ϵ about the points in \hat{X} .

3.4.2. Reconstructing Manifolds from Noisy Data. The following result is the noisy analogue of Proposition 3.2.2.

PROPOSITION 3.4.2. ([32, Prop 7.1]) *Assume that \hat{X} is r -dense in \mathcal{X} for some $0 < r < (3 - \sqrt{8})\tau_{\mathcal{X}}$ and choose a radius ϵ satisfying $\Gamma_{\mathcal{X}}^-(r) < \epsilon < \Gamma_{\mathcal{X}}^+(r)$. Let $\mathcal{N}_{\epsilon}(\hat{X})$ denote the nerve complex generated by open balls of radius ϵ about the points in \hat{X} . Then, the canonical projection map $\pi_{\mathcal{X}} : \lfloor \mathcal{N}_{\epsilon}(\hat{X}) \rfloor \rightarrow \mathcal{X}$ as defined in (3.2) is a retraction.*

As before, assuming the hypotheses of this proposition yields the isomorphisms of homology groups

$$H_*(\mathcal{X}) \simeq H_*\left(\lfloor \mathcal{N}_{\epsilon}(\hat{X}) \rfloor\right) \simeq H_*^{\Delta}(\mathcal{N}_{\epsilon}(\hat{X})).$$

Recall that given $r > 0$, the r -covering number of \mathcal{X} – denoted $\Lambda_r(\mathcal{X})$ – is defined to be the minimum possible $q \in \mathbb{N}$ satisfying the following property: there exists some point set $S \subset \mathcal{X}$ of cardinality q such that the collection $\{\mathcal{B}_r(s) \mid s \in S\}$ of n -dimensional open balls covers \mathcal{X} . Given an r -conditioned probability measure μ about \mathcal{X} with regularity constants $\{\Omega_s \mid 0 < s < r\}$ and a parameter $\delta > 0$, define the new bounding function $\hat{\beta}_{\mathcal{X}}$ as follows:

$$\hat{\beta}_{\mathcal{X}}(\mu, \delta) = \frac{1}{\Omega_{r/2}} (\log(\Lambda_{r/2}(\mathcal{X})) + \log(1/\delta)) \quad (3.13)$$

The next result replaces Proposition 3.2.5 in the setting of conditioned noise.

PROPOSITION 3.4.3. ([32, Prop 7.2]) *Assume that the positive real numbers r and ϵ satisfy $0 < r < (3 - \sqrt{8})\tau_x$ and $\Gamma_x^-(r) < \epsilon < \Gamma_x^+(r)$. Let μ be any r -conditioned probability measure about \mathcal{X} and assume that a point set \widehat{X} is drawn from \mathbb{R}^n in i.i.d. fashion with respect to μ . Given a parameter $\delta \in (0, 1]$, if $\#\widehat{X} > \widehat{\beta}_x(\mu, \delta)$ then \widehat{X} is r -dense in \mathcal{X} with probability exceeding $(1 - \delta)$.*

Combining the preceding propositions yields the main result of [32] as adapted for conditioned noise.

THEOREM 3.4.4. *Let $\mathcal{X} \subset \mathbb{R}^n$ be a compact Riemannian submanifold with condition number $1/\tau_x$. Fix $r \in (0, (3 - \sqrt{8})\tau_x)$ and choose a radius ϵ satisfying $\Gamma_x^-(r) < \epsilon < \Gamma_x^+(r)$. Assume that μ is an r -conditioned probability measure about \mathcal{X} and that $\widehat{X} \subset \mathbb{R}^n$ is a point set obtained by μ -i.i.d. sampling. Denote by $\mathcal{N}_\epsilon(\widehat{X})$ the nerve complex generated by open ϵ -balls in \mathbb{R}^n centered at points in \widehat{X} . If $\#\widehat{X} > \widehat{\beta}_x(\mu, \delta)$ for some $\delta \in (0, 1]$, then $[\mathcal{N}_\epsilon(\widehat{X})]$ retracts onto \mathcal{X} with probability exceeding $(1 - \delta)$.*

3.4.3. Accounting for Conditioned Noise. Introducing the model of conditioned noise requires the following modifications to the assumptions underlying our algorithms.

Lip': $f : \mathcal{X} \rightarrow \mathcal{Y}$ is a Lipschitz continuous function with constant κ satisfying $2\kappa \cdot \tau_x < \tau_y$.

Nse': Choose positive noise bounds $r_x < \alpha \cdot \tau_x$ and $r_y < \alpha \cdot \tau_y$ where $\alpha = (3 - \sqrt{8})$. μ_x and μ_y are r_x and r_y -conditioned probability measures about \mathcal{X} and \mathcal{Y} , respectively.

Rad': Choose radii ϵ_x and ϵ_y satisfying $\Gamma_x^-(r_x) < \epsilon_x < \Gamma_x^+(r_x)$ and $\Gamma_y^-(r_y) < \epsilon_y < \Gamma_y^+(r_y)$ so that

$$\kappa \cdot (2\epsilon_x + r_x) < (\epsilon_y - r_y). \quad (3.14)$$

Smp': We know the finite point sets $\widehat{X} \subset \mathbb{R}^n$ and $\widehat{Y} \subset \mathbb{R}^m$ obtained by i.i.d. μ_x and μ_y sampling respectively. We require that $\#\widehat{X} > \widehat{\beta}_x(\mu_x, \delta_x)$ and $\#\widehat{Y} > \widehat{\beta}_y(\mu_y, \delta_y)$.

Img': We know the pairs $\{(\xi, \widehat{f}(\xi)) \in \widehat{X} \times \mathbb{R}^m \mid \xi \in \widehat{X}\}$. Furthermore, for each $\xi \in \widehat{X}$

$$\widehat{f}(\xi) \in \mathcal{B}_d(f \circ \pi_x(\xi))$$

where π_x is the canonical projection map from (3.2) and

$$d < \frac{(\epsilon_y - r_y) - \kappa \cdot (2\epsilon_x + r_x)}{2}. \quad (3.15)$$

In particular, the assumptions **Cnd** and **Prb** remain unchanged. Note that in the assumption **Img'** we do not assume knowledge of the true image $f \circ \pi_x(\xi) \in \mathcal{Y}$. The inequality (3.14) is a constraint that involves the Lipschitz constant, the models for the noise, and the radii for the nerves. It guarantees that the restriction (3.15) is always positive. The following result provides conditions on the manifolds and the map under which (3.14) can be satisfied.

PROPOSITION 3.4.5. *If $2\kappa \cdot \tau_x < \tau_y$, then there exist valid choices of ϵ_x and ϵ_y which satisfy (3.14).*

PROOF. First, we check that $(2\epsilon_x + r_x) < 2\tau_x$ on the domain $0 < r_x < (3 - \sqrt{8})\tau_x$ imposed by **Nse'**. Recall that $\epsilon_x < \Gamma_x^+(r_x)$ by **Rad'** and consider the function

$$2\Gamma_x^+(r_x) + r_x = (\tau_x + 2r_x) + \sqrt{\tau_x^2 + r_x^2 - 6\tau_x r_x}$$

This function has no local maximum in its domain and attains a maximum value of $2\tau_x$ at the left endpoint. Since (3.14) imposes a lower bound of $\kappa \cdot (2\epsilon_x + r_x) + r_y$ on ϵ_y , it suffices to show that the over-estimate $2\kappa \cdot \tau_x + r_y$ of this lower bound is smaller than the upper bound $\Gamma_y^+(r_y)$ imposed on ϵ_y by **Rad'**. Equivalently, we must show that $\Gamma_y^+(r_y) - r_y > 2\kappa \cdot \tau_x$. Observe that

the function

$$\Gamma_y^+(r_y) - r_y = \frac{(\tau_y - r_y) + \sqrt{\tau_y^2 + r_y^2 - 6\tau_y r_y}}{2}$$

has no local minima on the domain $(0, (3 - \sqrt{8})\tau_y)$ imposed by **Nse'** and attains a minimum value of τ_y at the left endpoint. Thus, it is possible to satisfy (3.14) if $2\kappa \cdot \tau_x < \tau_y$. \square

3.4.4. Reconstructing Functions from Noisy Data. The main result of this section is the following theorem.

THEOREM 3.4.6. *Assume **Cnd**, **Lip'**, **Nse'**, **Rad'**, **Prb**, **Smp'**, and **Img'**. If $\#\widehat{X} > \widehat{\beta}_x(\mu_x, \delta_x)$ and $\#\widehat{Y} > \widehat{\beta}_y(\mu_y, \delta_y)$ then there exists an algorithm which takes as its input $(\widehat{X}, \widehat{f}(\widehat{X}), \widehat{Y})$ and produces the following output with probability exceeding $(1 - \delta_x) \cdot (1 - \delta_y)$:*

- (1) an acyclic outer approximation F of f , and
- (2) a chain map $\phi_{\#}$ carried by F .

For the most part the proof of this theorem is analogous to that of Theorem 3.3.1. We note the significant differences in the arguments below.

PROPOSITION 3.4.7. *Set $\rho = r_y + d$. With probability exceeding $(1 - \delta_y)$, Approximator terminates with a successful status when called with the input $(\widehat{X}, \widehat{f}(\widehat{X}), \widehat{Y})$.*

PROOF. By **Cnd**, **Nse'**, **Rad'** and **Prb**, we may apply Proposition 3.4.3 to \mathcal{Y} and assume that \widehat{Y} is r_y -dense in \mathcal{Y} with high probability. So for each $\xi \in \widehat{X}$, there exists some $\eta \in \widehat{Y}$ which is within distance r_y of the true image $f \circ \pi_x(\xi) \in \mathcal{Y}$. By **Img'**, this η is within $r_y + d$ of the sampled image $\widehat{f}(\xi)$, whence $\eta \in \mathcal{L}(\xi)$. Thus, the output $\mathcal{L} : \widehat{X} \rightarrow 2^{\widehat{Y}}$ satisfies $\mathcal{L}(\xi) \neq \emptyset$ for each $\xi \in \widehat{X}$ as desired. \square

Define $\mathbf{G} : \mathcal{N}(\widehat{X}) \rightarrow 2^{\widehat{Y}}$ using \mathcal{L} as in (3.6). The following result is the noisy analogue of Proposition 3.3.3.

PROPOSITION 3.4.8. *With probability exceeding $(1 - \delta_x)$, the following is true. For each $\sigma \in \mathcal{N}(\widehat{X})$ and $\eta \in \mathbf{G}(\sigma)$, we have $f \circ \pi_x(\lfloor \sigma \rfloor) \subset \mathcal{B}_{\varepsilon_y}(\eta)$.*

PROOF. Choose $w \in \lfloor \sigma \rfloor$ and note that $w \in \bigcap_{\xi \in \mathbf{g}_{\widehat{X}}(\sigma)} \mathcal{B}_{\varepsilon_x}(\xi)$ from the definition of support. By Lemma 3.2.3, the distance from w to $\pi_x(w)$ is smaller than ε_x^2/τ_x which is bounded above by ε_x from **Rad'** and (3.13). Thus, the distance to $\pi_x(w)$ from each $\xi \in \mathbf{g}_{\widehat{X}}(\sigma)$ is at most $2\varepsilon_x$ by the triangle inequality. Therefore, we have

$$\pi_x(w) \in \bigcap_{\xi \in \mathbf{g}_{\widehat{X}}(\sigma)} \mathcal{B}_{2\varepsilon_x}(\xi)$$

By **Nse'** and **Smp'**, each element of \widehat{X} is no more than r_x away from \mathcal{X} , and in particular each ξ in $\mathbf{g}_{\widehat{X}}(\sigma)$ satisfies $\|\xi - \pi_x(\xi)\|_{\mathbb{R}^n} < r_x$. We use this in the expression above to get

$$\pi_x(w) \in \bigcap_{\xi \in \mathbf{g}_{\widehat{X}}(\sigma)} \mathcal{B}_{2\varepsilon_x + r_x}(\pi_x(\xi))$$

Using **Lip'**,

$$f \circ \pi_x(w) \in \bigcap_{\xi \in \mathbf{g}_{\widehat{X}}(\sigma)} \mathcal{B}_{\kappa \cdot (2\varepsilon_x + r_x)}(f \circ \pi_x(\xi))$$

Now if $\eta \in \mathbf{G}(\sigma)$, then by definition $\eta \in \mathcal{L}(\xi_*)$ for some $\xi_* \in \mathbf{g}_{\widehat{X}}(\sigma)$. We have the following bounds for distances in \mathbb{R}^m :

- (1) $\|f \circ \pi_X(w) - f \circ \pi_X(\xi_*)\| < \kappa \cdot (2\epsilon_X + r_X)$ as above,
- (2) $\|f \circ \pi_X(\xi_*) - \widehat{f}(\xi_*)\| \leq d$ from **Img'**, and
- (3) $\|\widehat{f}(\xi_*) - \eta\| < \rho = r_Y + d$ from Proposition 3.4.7.

We conclude with the observation that

$$\|f \circ \pi_X(w) - \eta\| < \kappa \cdot (2\epsilon_X + r_X) + r_Y + 2d < \epsilon_Y$$

where the first inequality is obtained using the triangle inequality and second follows from (3.15). \square

3.5. Multi-Scale Analysis of Data Transformations

In this section, we utilize the *method* of building simplicial reconstructions purely on data points without requiring most of the assumptions and hypotheses from Section 3.3.1. Of course, in this setting it is impossible to recover the homology of manifolds or functions: we have not assumed that the data comes from a smooth process in the first place.

In practice, one frequently does not have knowledge of regularity, condition numbers or Lipschitz constants when confronted with experimental data. In the absence of such theoretical guarantees, it is nevertheless of practical importance to have means of computing topological invariants at various scales from transformations of data.

Let $X \subset \mathbb{R}^n$ and $Y \subset \mathbb{R}^m$ be finite point sets, and assume the existence of a map $T : X \rightarrow \mathbb{R}^m$, called a *transformation* of X , which assigns to each $\xi \in X$ an *image* $T(\xi)$. We make no further assumptions about X , Y or T .

DEFINITION 3.5.1. The *density* of Y in $T(X)$ is defined to be the smallest $r \geq 0$ such that for each $\xi \in X$ there exists some $\eta \in Y$ with $\|T(\xi) - \eta\|_{\mathbb{R}^m} \leq r$.

Let ρ be the density of Y in $T(X)$, and note that $\rho = 0$ if and only if $T(X) \subset Y$, i.e., if T maps the points of X to points of Y .

DEFINITION 3.5.2. Define the *local Lipschitz function* $\kappa : \mathbb{R}^+ \rightarrow \mathbb{R}^+$ of the transformation T as follows:

$$\kappa(\epsilon) = \sup \left\{ \frac{\|T(\xi) - T(\xi')\|_{\mathbb{R}^m}}{\|\xi - \xi'\|_{\mathbb{R}^n}} \mid \xi \neq \xi' \in X \text{ with } \|\xi - \xi'\|_{\mathbb{R}^n} \leq \epsilon \right\}$$

with the understanding that $\kappa(\epsilon) = 0$ whenever $\epsilon < \inf_{\xi, \xi' \in X} \|\xi - \xi'\|_{\mathbb{R}^n}$

For each $\epsilon \geq 0$ we have the following Lipschitz-like behavior by definition of κ :

$$\|T(\xi) - T(\xi')\|_{\mathbb{R}^m} \leq \kappa(\epsilon) \cdot \|\xi - \xi'\|_{\mathbb{R}^n} \quad (3.16)$$

where ξ, ξ' in X are distinct points such that $\|\xi - \xi'\|_{\mathbb{R}^n} < \epsilon$.

Following the general procedure of Section 3.3, we may use Approximator to construct the correspondence $\mathcal{L} : X \rightarrow 2^Y$ as follows:

$$\mathcal{L}(\xi) = \{\eta \in Y \mid \|T(\xi) - \eta\|_{\mathbb{R}^m} \leq \rho\} \quad (3.17)$$

and note that $\mathcal{L}(\xi) \neq \emptyset$ for each $\xi \in X$ by the ρ -density of Y in $T(X)$. Next, one can define the map $h : X \rightarrow Y$ using Selector. More precisely, $h(\xi)$ is defined to be any point from $\mathcal{L}(\xi)$. Although this choice is arbitrary, we will assume that it has been fixed throughout the sequel. The essential property of h by construction is:

$$\|h(\xi) - T(\xi)\|_{\mathbb{R}^m} \leq \rho \quad (3.18)$$

For $Z \in \{X, Y\}$ and $\epsilon > 0$, let $\mathcal{U}_\epsilon(Z)$ be the union of balls $\bigcup_{\zeta \in Z} \mathcal{B}_\epsilon(\zeta)$ and let $\mathcal{N}_\epsilon(Z)$ be the nerve associated to the cover $\{\mathcal{B}_\epsilon(\zeta) \mid \zeta \in Z\}$ of $\mathcal{U}_\epsilon(Z)$.

PROPOSITION 3.5.3. *Given $\epsilon_X > 0$, the function $h : X \rightarrow Y$ determines a simplicial map from $\mathcal{N}_{\epsilon_X}(X)$ to $\mathcal{N}_{\epsilon_Y}(Y)$ provided $\epsilon_Y > \rho + 2\epsilon_X \cdot \kappa(\epsilon_X)$.*

PROOF. Let $\sigma \in \mathcal{N}_{\epsilon_X}(X)$ be a simplex with vertex set $V \subset X$. It suffices to show that there exists some $y \in \mathbb{R}^m$ whose distance from each point in $h(V) \subset Y$ is smaller than $\rho + 2\epsilon_X \cdot \kappa(\epsilon_X)$. By definition of the nerve, there exists some $x \in \mathbb{R}^n$ so that $\|x - \xi\|_{\mathbb{R}^n} < \epsilon_X$ for each $\xi \in V$, and so by the triangle inequality we obtain $\|\xi - \xi'\| < 2\epsilon_X$ for any pair of vertices ξ, ξ' of σ . By (3.16), we have $\|T(\xi) - T(\xi')\|_{\mathbb{R}^m} \leq 2\epsilon_X \cdot \kappa(\epsilon_X)$, and so the set $T(V)$ consists of finitely many points in \mathbb{R}^m that have pairwise distance at most $2\epsilon_X \cdot \kappa(\epsilon_X)$. Assume that there are K such distinct points, and define $y \in \mathbb{R}^m$ by

$$y = \frac{1}{K} \sum_{\zeta \in T(V)} \zeta$$

Clearly, we have $\|y - T(\xi)\|_{\mathbb{R}^m} \leq 2\epsilon_X \cdot \kappa(\epsilon_X)$ for each vertex ξ of σ . Finally, by (3.18) and the triangle inequality we obtain $\|y - h(\xi)\|_{\mathbb{R}^m} \leq \rho + 2\epsilon_X \cdot \kappa(\epsilon_X)$ for each vertex ξ of σ . Thus, if ϵ_Y exceeds $\rho + 2\epsilon_X \cdot \kappa(\epsilon_X)$ then y lies within ϵ_Y of each point in the set $h(V)$. Therefore, $h(V)$ determines a simplex of the nerve $\mathcal{N}_{\epsilon_Y}(Y)$ as desired. \square

Since there are only finitely many points in X , there are only finitely many values of $\epsilon_X \geq 0$ such that the homology of the nerve $\mathcal{N}_{\epsilon_X}(X)$ changes. We arrange these *critical* values of ϵ_X in ascending order as follows:

$$0 \leq \epsilon_X^{(1)} < \epsilon_X^{(2)} < \dots < \epsilon_X^{(M)} \leq \text{diam}(X)$$

where $\text{diam}(X)$ is the maximum distance between any two points in X . For notational convenience, define the nerves $\mathcal{X}^m = \mathcal{N}_{\epsilon_X^{(m)}}(X)$ for $m \in \{1, \dots, M\}$ and note that we have a filtration F of \mathcal{X}^M by simplicial subcomplexes:

$$\mathcal{X}^1 \subset \mathcal{X}^2 \subset \dots \subset \mathcal{X}^M$$

From Definition 3.5.2, it is clear that κ is a monotone increasing function and so it is possible to choose an increasing sequence

$$0 \leq \epsilon_Y^{(1)} < \epsilon_Y^{(2)} < \dots < \epsilon_Y^{(M)}$$

so that for each $m \in \{1, \dots, M\}$, we have

$$\epsilon_Y^{(m)} > \rho + 2\epsilon_X^{(m)} \cdot \kappa(\epsilon_X^{(m)}). \quad (3.19)$$

Setting $\mathcal{Y}^m = \mathcal{N}_{\epsilon_Y^{(m)}}(Y)$ for each such m , we have the following filtration G of \mathcal{Y}^M :

$$\mathcal{Y}^1 \subset \mathcal{Y}^2 \subset \dots \subset \mathcal{Y}^M$$

By Proposition 3.5.3 and the assumption (3.19), $h : X \rightarrow Y$ induces a simplicial map $\phi^m : \mathcal{X}^m \rightarrow \mathcal{Y}^m$ for each $1 \leq m \leq M$. Moreover, since each ϕ^m is completely determined by $h : X \rightarrow Y$ independent of m , it is clear that $\phi^{m+1}|_{\mathcal{X}^m} \equiv \phi^m$. Thus, the chain maps $\phi_{\#}^m : \mathcal{C}(\mathcal{X}^m) \rightarrow \mathcal{C}(\mathcal{Y}^m)$ associated to ϕ^m fit into the following commutative diagram

$$\begin{array}{ccccccc} \mathcal{C}(\mathcal{X}^1) & \longrightarrow & \mathcal{C}(\mathcal{X}^2) & \longrightarrow & \dots & \longrightarrow & \mathcal{C}(\mathcal{X}^M) \\ \downarrow \phi_{\#}^1 & & \downarrow \phi_{\#}^2 & & & & \downarrow \phi_{\#}^M \\ \mathcal{C}(\mathcal{Y}^1) & \longrightarrow & \mathcal{C}(\mathcal{Y}^2) & \longrightarrow & \dots & \longrightarrow & \mathcal{C}(\mathcal{Y}^M) \end{array}$$

where the horizontal maps arise from inclusion. Thus, we have proved the following main result of this section.

PROPOSITION 3.5.4. *The transformation T canonically induces a morphism of filtrations from F to G .*

As a consequence of results from Section 2.4.2, T induces homomorphisms of persistent homology modules of F and G ; thus, it is possible to systematically track how T maps topological features at various scales arising from nerves centered at X to the corresponding features in nerves centered at Y .

Discrete Morse Theory for Filtrations

Every mathematician has a secret weapon. Mine is Morse theory.

–Raoul Bott, [5]

4.1. Introduction

Traditionally, Morse theory studies the connection between a smooth compact manifold M and a suitably generic smooth map $f : M \rightarrow \mathbb{R}$ – called a *Morse function* – by considering the sub-level sets $M^a = \{m \in M \mid f(m) \leq a\}$ for different values of $a \in \mathbb{R}$. It turns out that there is no change in homotopy type between M^a and M^b for $a < b$ if there exist no critical values of f between a and b . On the other hand, if such a critical value c exists then the homotopy type of M^b may be captured by attaching a single disc of suitable dimension along its boundary to M^a . Thus, it is possible to recover M up to homotopy type by understanding how the topology of the sub-level sets of a generic \mathbb{R} -valued function changes at the critical points.

John Milnor’s book [26] provides a clear and wonderful exposition of smooth Morse theory. Various flavors of this theory have been developed since Marston Morse’s original work. The survey article of Raoul Bott [6]¹ provides an entertaining overview of Morse-Smale, Morse-Bott and Morse-Witten theories as well as their triumphs.

Our focus here is on a different type of Morse theory, much closer in spirit to the piece-wise linear Morse theory described in [5] by Mladen Bestvina and most importantly, the discrete Morse theory of Robin Forman [13] developed for regular CW complexes. Forman’s original description of the theory involves actual *discrete Morse functions* $\mu : \mathcal{X} \rightarrow \mathbb{R}$ associating a real value to each cell in the regular complex \mathcal{X} such that μ increases with dimension, allowing at most one exception per cell. In almost all contexts, the actual function values are inconsequential and one only has to catalog pairs of cells that are exceptions to this rule. This leads to Manoj Chari’s reformulation [9] of discrete Morse theory in terms of *acyclic matchings*.

Our definition of an acyclic matching $(\mathcal{A}, w : \mathcal{Q} \rightarrow \mathcal{K})$ coincides with that of Chari. For our purposes, the w -paired cells are important, and so we provide explicit labels for them along the lines of Dmitry Kozlov’s presentation [22, 23]. In order to optimize the notation and terminology for extending this theory to *filtered* cell complexes in Section 4.3, we provide a self-contained introduction to discrete Morse theory as it applies to complexes in Section 4.2.

4.2. Discrete Morse Theory for Cell Complexes

Let (\mathcal{X}, κ) be a complex over the PID \mathbf{R} and denote by \prec the generating relation of the face partial order \preceq on \mathcal{X} which is defined by $\xi \prec \eta$ if and only if $\kappa(\eta, \xi) \neq 0$.

DEFINITION 4.2.1. A *partial matching* μ of (\mathcal{X}, κ) consists of a partition of \mathcal{X} into three sets \mathcal{A} , \mathcal{K} and \mathcal{Q} along with a bijection $w : \mathcal{Q} \rightarrow \mathcal{K}$, such that for each $q \in \mathcal{Q}$ the incidence $\kappa(w(q), q)$ is a unit in \mathbf{R} .

¹aptly titled “Morse theory indomitable”!

Let $\mu = (\mathcal{A}, w : \mathcal{Q} \rightarrow \mathcal{X})$ be a matching on (\mathcal{X}, κ) .

Observe that by Definition 2.3.1(i) and the unit incidence requirement, we have $\dim w(q) = \dim q + 1$ and $q \prec w(q)$ for each $q \in \mathcal{Q}$. The cells in \mathcal{A} are called the *critical* cells of μ and the remaining cells are said to be *paired* by μ .

A *path* of μ is a sequence of cells

$$\rho = (q_1, w(q_1), q_2, w(q_2), \dots, q_d, w(q_d))$$

such that $q_j \in \mathcal{Q}$ and $q_j \neq q_{j+1} \prec w(q_j)$ for each j . We denote the starting and ending cells of this path by $\mathbf{s}_\rho = q_1$ and $\mathbf{e}_\rho = w(q_d)$. We call ρ a *cycle* if it has length $d > 1$ and $\mathbf{s}_\rho \prec \mathbf{e}_\rho$. The matching μ is called *acyclic* if none of its paths is a cycle; we will assume throughout the remainder of this chapter that μ is acyclic.

The *index* of ρ is defined as

$$\text{ind}(\rho) = \frac{\prod_{j=1}^{d-1} \kappa(w(q_j), q_{j+1})}{\prod_{j=1}^d -\kappa(w(q_j), q_j)}$$

And for any pair \mathbf{a}, \mathbf{a}' of critical cells, the *multiplicity* of ρ from \mathbf{a} to \mathbf{a}' is defined to be

$$\text{mul}(\mathbf{a} \xrightarrow{\rho} \mathbf{a}') = \kappa(\mathbf{a}, \mathbf{s}_\rho) \cdot \text{ind}(\rho) \cdot \kappa(\mathbf{e}_\rho, \mathbf{a}')$$

Define the following \mathbf{R} -valued function on $\mathcal{A} \times \mathcal{A}$:

$$\kappa_\mu(\mathbf{a}, \mathbf{a}') = \kappa(\mathbf{a}, \mathbf{a}') + \sum_{\rho} \text{mul}(\mathbf{a} \xrightarrow{\rho} \mathbf{a}') \quad (4.1)$$

where the sum is taken over all paths ρ of the matching μ . The remainder of this section is devoted to proving the following central theorem.

THEOREM 4.2.2. *Let (\mathcal{X}, κ) be a complex over a PID \mathbf{R} and let $\mu = (\mathcal{A}, w : \mathcal{Q} \rightarrow \mathcal{X})$ be an acyclic matching of \mathcal{X} . Then, $(\mathcal{A}, \kappa_\mu)$ is also a complex over \mathbf{R} , and*

$$H_*(\mathcal{X}) \simeq H_*(\mathcal{A}).$$

The complex $(\mathcal{A}, \kappa_\mu)$ is called the *Morse complex* associated to μ and κ_μ is called the associated *Morse incidence function*.

4.2.1. The Reduction Step. Let (\mathcal{X}, κ) be a complex with associated chain complex $(\mathcal{C}(\mathcal{X}); \partial)$. Let μ be an acyclic matching $(\mathcal{A}, w : \mathcal{Q} \rightarrow \mathcal{X})$ of \mathcal{X} . Fix $q \in \mathcal{Q}$ and define $\mathcal{X}' \subset \mathcal{X}$ by

$$\mathcal{X}' := \mathcal{X} \setminus \{q, w(q)\}$$

and the function $\kappa' : \mathcal{X}' \times \mathcal{X}' \rightarrow \mathbf{R}$ by

$$\kappa'(\eta, \xi) = \kappa(\eta, \xi) - \frac{\kappa(\eta, q) \cdot \kappa(w(q), \xi)}{\kappa(w(q), q)} \quad (4.2)$$

Note that κ' is constructed by a sequence of row and column operations on the matrix representation of the boundary operator ∂ which make the unit incidence of q and $w(q)$ a pivot.

PROPOSITION 4.2.3. *For any η and η' in \mathcal{X}' , we have*

$$\sum_{\xi \in \mathcal{X}'} \kappa(\eta, \xi) \cdot \kappa(\xi, \eta') = -\kappa(\eta, w(q)) \cdot \kappa(w(q), \eta') - \kappa(\eta, q) \cdot \kappa(q, \eta')$$

Moreover, at most one of the two terms on the right side can be non-zero.

PROOF. We only need to use (2.2) along with the fact that \mathcal{X}' contains all the cells in \mathcal{X} except q and $w(q)$. The non-zerosness assertion follows from property (i) of Definition 2.3.1 and the fact that $\dim w(q) = \dim q + 1 \neq \dim q$. \square

A direct computation using this calculation establishes that κ' is an incidence function on \mathcal{X}' , and so (\mathcal{X}', κ') is a complex. We also show that the new κ' is not very different from κ in the following sense

PROPOSITION 4.2.4. *Let $q' \neq q$ be an element of \mathcal{Q} . Then, $\kappa'(w(q'), q') = \kappa(w(q'), q')$ and in particular, the unit incidence of paired cells is preserved in \mathcal{X}' .*

PROOF. Recall that

$$\kappa'(w(q'), q') - \kappa(w(q'), q') = \frac{\kappa(w(q'), q) \cdot \kappa(w(q), q')}{\kappa(w(q), q)}$$

by the defining formula (4.2). Assume for contradiction that this quantity is non-zero. Since \mathbf{R} is a PID, each factor comprising the numerator must be non-zero. Thus, we have $q \prec w(q')$ which implies that q' precedes q in the following path of μ :

$$\rho = (q', w(q'), q, w(q)).$$

But by non-zerosness of that numerator, we also have $q' \prec w(q)$ which in turn forces ρ to be a cycle. Since the matching μ is acyclic by our standing assumption, we must have $q = q'$, a contradiction. \square

Observe that the acyclic matching μ on \mathcal{X} induces an acyclic matching μ' on \mathcal{X}' of the form $(\mathcal{A}, w: \mathcal{Q}' \rightarrow \mathcal{K}')$ where we have $\mathcal{Q}' = \mathcal{Q} \setminus \{q\}$ and $\mathcal{K}' = \mathcal{K} \setminus \{w(q)\}$. The following result which guarantees that the Morse incidence function κ_μ remains unaffected by the reduction step.

PROPOSITION 4.2.5. *Let κ'_μ denote the Morse incidence function of the induced acyclic matching μ' on the reduced complex \mathcal{X}' . Then, $\kappa'_\mu \equiv \kappa_\mu$ on $\mathcal{A} \times \mathcal{A}$.*

PROOF. Note that $\kappa'_\mu(w(q'), q') = \kappa(w(q'), q')$ for any $q' \in \mathcal{Q}'$ by Proposition 4.2.4.

Fix critical cells a and a' in \mathcal{A} and let $\rho = (q_1, \dots, w(q_d))$ be a connection of μ' from a to a' . We make the simplifying assumptions that $q \not\prec a$ and $a' \not\prec w(q)$, because the argument is very similar to the sequel when one or both of these assumptions is revoked. Now, we have $\kappa'(a, a') = \kappa(a, a')$ by (4.2), so we only need to show that the sum-over-connections term of (4.1) is the same for \mathcal{X} with the matching μ and \mathcal{X}' with the matching μ' .

Note that there is at most one $j \in 1, \dots, (d-1)$ with $\kappa'(w(q_j), q_{j+1}) \neq \kappa(w(q_j), q_{j+1})$ by the following contradiction. If $j < j'$ both satisfy this inequality, then an argument similar to the proof of Proposition 4.2.4 allows us to conclude that the following path is a cycle of μ :

$$(q, w(q), q_{j+1}, \dots, q_{j'}, w(q_{j'}))$$

If there is no such j , then the index of ρ – and hence its multiplicity as a connection between any two critical cells – is the same in both \mathcal{X} and \mathcal{X}' . On the other hand, if there is such a j , then there exists a unique augmented connection ρ^+ from a to a' in \mathcal{X} given by

$$\rho^+ = (q_1, w(q_1), \dots, w(q_j), q, w(q), q_{j+1}, \dots, q_d, w(q_d))$$

and it is readily verified by definition that the index of ρ in \mathcal{X}' equals the sum of the indices of ρ and ρ^+ in \mathcal{X} . Thus, the sum over all connections of the multiplicities is preserved in the reduced complex (\mathcal{X}', κ') . \square

4.2.2. Constructing Chain Equivalences. Define an \mathbf{R} -linear map ψ from $C_*(\mathcal{X})$ to the reduced chain complex $C_*(\mathcal{X}')$ by the following basis action

$$\psi(\eta) = \begin{cases} 0, & \text{if } \eta = w(q) \\ -\sum_{\xi \in \mathcal{X}'} \frac{\kappa(w(q), \xi)}{\kappa(w(q), q)} \xi, & \text{if } \eta = q \\ \eta, & \text{otherwise} \end{cases} \quad (4.3)$$

Let ∂' be the boundary operator on $\mathcal{C}(\mathcal{X}')$ associated to the incidence function κ' from (4.2).

PROPOSITION 4.2.6. ψ is a chain map, i.e., $\psi \circ \partial \equiv \partial' \circ \psi$.

PROOF. For any $\eta \in \mathcal{X}$, we have

$$\psi \circ \partial \eta = \psi \left(\sum_{\xi \in \mathcal{X}} \kappa(\eta, \xi) \xi \right) \text{ by definition of } \partial$$

Since $\mathcal{X} = \mathcal{X}' \cup \{w(q), q\}$, we have

$$\begin{aligned} \psi \circ \partial \eta &= \psi \left(\left(\sum_{\xi \in \mathcal{X}'} \kappa(\eta, \xi) \xi \right) + \kappa(\eta, w(q)) w(q) + \kappa(\eta, q) q \right) \\ &= \sum_{\xi \in \mathcal{X}'} \kappa(\eta, \xi) \psi(\xi) + \kappa(\eta, w(q)) \psi(w(q)) + \kappa(\eta, q) \psi(q) \\ &= \sum_{\xi \in \mathcal{X}'} \kappa(\eta, \xi) \xi + \kappa(\eta, q) \psi(q), \text{ by (4.3)} \\ &= \sum_{\xi \in \mathcal{X}'} \left[\kappa(\eta, \xi) - \frac{\kappa(\eta, q)}{\kappa(w(q), q)} \kappa(w(q), \xi) \right] \xi, \text{ also by (4.3)} \end{aligned}$$

Recognizing the formula (4.2) for κ' in the summand and recalling that it may only differ from κ when $\dim \eta = \dim w(q) \neq \dim q$, we have

$$\psi \circ \partial(\eta) = \begin{cases} 0, & \eta = w(q) \\ \partial q, & \eta = q \\ \partial' \eta & \text{otherwise} \end{cases}$$

On the other hand, if $\eta \notin \{w(q), q\}$ then the right side $\partial' \circ \psi(\eta)$ of the desired equality also equals $\partial' \eta$ since ψ is just the identity map in this case. So we only need to check separately when $\eta = w(q)$ or $\eta = q$. Clearly, $\partial' \circ \psi(w(q)) = 0$ by the trivial action of ψ on $w(q)$. Finally,

$$\begin{aligned} \partial' \circ \psi(q) &= \partial' \left[-\sum_{\xi \in \mathcal{X}'} \frac{\kappa(w(q), \xi)}{\kappa(w(q), q)} \xi \right], \text{ by (4.3)} \\ &= -\sum_{\xi \in \mathcal{X}'} \frac{\kappa(w(q), \xi)}{\kappa(w(q), q)} \partial' \xi \end{aligned}$$

Applying the definition of ∂' to this expression, we have

$$\begin{aligned}\partial' \circ \psi(q) &= - \sum_{\xi \in \mathcal{X}'} \frac{\kappa(w(q), \xi)}{\kappa(w(q), q)} \left(\sum_{\zeta \in \mathcal{X}'} \kappa'(\xi, \zeta) \zeta \right) \\ &= \frac{-1}{\kappa(w(q), q)} \sum_{\zeta \in \mathcal{X}'} \sum_{\xi \in \mathcal{X}'} \kappa(w(q), \xi) \cdot \kappa'(\xi, \zeta) \zeta\end{aligned}$$

We know that if $\kappa(w(q), \xi) \neq 0$ then $\dim \xi = \dim q$ and so $\kappa'(\xi, \zeta) \neq 0$ forces $\dim \zeta = \dim q - 1$. But now $\kappa'(\xi, \zeta) = \kappa(\xi, \eta)$ for these dimensions and so

$$\partial' \circ \psi(q) = \frac{-1}{\kappa(w(q), q)} \sum_{\zeta \in \mathcal{X}'} \sum_{\xi \in \mathcal{X}'} \kappa(w(q), \xi) \cdot \kappa(\xi, \zeta) \zeta$$

An application of Proposition 4.2.3 on the inner sum yields

$$\partial' \circ \psi(q) = \sum_{\zeta \in \mathcal{X}'} \kappa(q, \zeta) \zeta$$

Finally, we may as well let ζ range over all of \mathcal{X} instead of just \mathcal{X}' since $\kappa(q, *) = 0$ for $* = w(q)$ and $* = q$ by dimension requirements. At last, we have $\partial' \circ \psi(q) = \partial q$ as desired. \square

We now require a chain map in the other direction. To this end, we define another \mathbf{R} -linear map ϕ from $C_*(\mathcal{X}')$ to $C_*(\mathcal{X})$ by

$$\phi'(\eta) = \eta - \frac{\kappa(\eta, q)}{\kappa(w(q), q)} w(q) \quad (4.4)$$

PROPOSITION 4.2.7. ϕ is a chain map, i.e., $\partial \circ \phi \equiv \phi \circ \partial'$.

PROOF. For any $\eta' \in \mathcal{X}'$,

$$\begin{aligned}\partial \circ \phi(\eta') &= \partial \left(\eta' - \frac{\kappa(\eta', q)}{\kappa(w(q), q)} w(q) \right), \text{ by (4.4)} \\ &= \partial \eta' - \frac{\kappa(\eta', q)}{\kappa(w(q), q)} \partial w(q) \\ &= \sum_{\xi \in \mathcal{X}} \left[\kappa(\eta', \xi) - \frac{\kappa(\eta', q)}{\kappa(w(q), q)} \kappa(w(q), \xi) \right] \xi, \text{ by definition of } \partial\end{aligned}$$

Compare the summand to (4.2) and note that it equals zero for $\xi = q$. Using $\mathcal{X} = \mathcal{X}' \cup \{w(q), q\}$ and $\kappa(w(q), w(q)) = 0$ by dimension requirements, we have

$$\begin{aligned}\partial \circ \phi(\eta') &= \sum_{\xi \in \mathcal{X}'} \kappa'(\eta', \xi) \xi + \kappa(\eta', w(q)) w(q) \\ &= \partial' \eta' + \kappa(\eta', w(q)) w(q)\end{aligned}$$

Thus, $\partial \circ \phi(\eta') = \partial' \eta'$ except when $\kappa(\eta', w(q)) \neq 0$; but this inequality may only hold when $\dim \eta' = \dim w(q) + 1$. On the other hand, we note that $\phi \circ \partial(\eta')$ is also just $\partial \eta'$ whenever $\dim \eta' \neq \dim w(q) + 1$ since in those cases ϕ reduces to the identity map. So assume without

loss of generality that $\dim \eta' = \dim w(q) + 1$ and observe

$$\begin{aligned}
\phi \circ \partial' \eta' &= \phi \left(\sum_{\xi \in \mathcal{X}'} \kappa(\eta', \xi) \xi \right), \text{ by definition of } \partial' \\
&= \sum_{\xi \in \mathcal{X}'} \kappa(\eta', \xi) \phi(\xi) \\
&= \sum_{\xi \in \mathcal{X}'} \kappa(\eta', \xi) \left(\xi - \frac{\kappa(\xi, q)}{\kappa(w(q), q)} w(q) \right), \text{ by (4.4)} \\
&= \partial' \eta' - \frac{1}{\kappa(w(q), q)} \left[\sum_{\xi \in \mathcal{X}'} \kappa(\eta', \xi) \cdot \kappa(\xi, q) \right] w(q)
\end{aligned}$$

Using Proposition 4.2.3 on the sum and noting that $\kappa(\eta', q) = 0$ by dimension requirements gives

$$\phi \circ \partial' \eta' = \partial' \eta' + \kappa(\eta', w(q)) w(q)$$

Comparing this to the expression for $\partial \circ \phi(\eta')$ obtained above concludes the argument. \square

LEMMA 4.2.8. *The maps ψ and ϕ are chain equivalences.*

PROOF. A direct computation shows that the composition $\psi \circ \phi$ is the identity map on $\mathcal{C}(\mathcal{X}')$. It remains to be shown that $\phi \circ \psi$ is chain homotopic to the identity on \mathcal{X} . Define a collection Θ of \mathbf{R} -module maps $\{\Theta: C_*(\mathcal{X}) \rightarrow C_{*+1}(\mathcal{X})\}$ by the following action on cells

$$\Theta(\eta) = \begin{cases} \frac{1}{\kappa(w(q), q)} w(q) & \text{if } \eta = q \\ 0 & \text{otherwise.} \end{cases} \quad (4.5)$$

When $\eta \in \mathcal{X}' \subset \mathcal{X}$, we see that $\psi(\eta) = \eta$ and so $\phi \circ \psi(\eta) = \eta - \frac{\kappa(\eta, q)}{\kappa(w(q), q)} w(q)$. Therefore, we have

$$[\text{id}_{\mathcal{C}(\mathcal{X})} - \phi \circ \psi](\eta) = \eta - \phi \circ \psi(\eta) = \frac{\kappa(\eta, q)}{\kappa(w(q), q)} w(q)$$

On the other hand, $\Theta(\eta) = 0$ since $\eta \neq q$, and so we have

$$[\partial \circ \Theta + \Theta \circ \partial](\eta) = \Theta \circ \partial(\eta) = \kappa(\eta, q) \circ \Theta(q) = \frac{\kappa(\eta, q)}{\kappa(w(q), q)} w(q)$$

so it is clear that $\text{id}_{\mathcal{C}(\mathcal{X})} - \phi \circ \psi \equiv \partial \circ \Theta + \Theta \circ \partial$ on \mathcal{X}' and we only need to check the cases $\eta \in \{w(q), q\}$. First we handle $w(q)$: immediately, we see that $w(q) - \phi \circ \psi(w(q)) = w(q)$ since ψ is trivial on $w(q)$, and

$$\begin{aligned}
\partial \circ \Theta(w(q)) + \Theta \circ \partial(w(q)) &= \Theta \circ \partial(w(q)) \\
&= \sum_{\xi \in \mathcal{X}} \kappa(w(q), \xi) \Theta(\xi) \\
&= \kappa(w(q), q) \Theta(q) \\
&= w(q)
\end{aligned}$$

as desired. On the other hand, for q , we have $\partial \circ \Theta(q) + \Theta \circ \partial(q)$ equals $\partial \circ \Theta(q)$ and hence $\frac{1}{\kappa(w(q), q)} \partial w(q)$. Moreover,

$$\begin{aligned} q - \phi \circ \psi(q) &= q + \phi \left(\sum_{\xi \in \mathcal{X}'} \frac{\kappa(w(q), \xi)}{\kappa(w(q), q)} \xi \right) \\ &= q + \left(\sum_{\xi \in \mathcal{X}'} \frac{\kappa(w(q), \xi)}{\kappa(w(q), q)} \phi(\xi) \right) \end{aligned}$$

We know that $\kappa(w(q), \xi) \neq 0$ forces $\dim \xi = \dim q \neq \dim w(q)$ by the second property of Definition 2.3.1 and so $\phi(\xi) = \xi$. Therefore, we have $q - \phi \circ \psi(q) = q + \sum_{\xi \in \mathcal{X}'} \frac{\kappa(w(q), \xi)}{\kappa(w(q), q)} \xi$. Using $\mathcal{X} = \mathcal{X}' \cup \{w(q), q\}$ and $\kappa(w(q), w(q)) = 0$ gives us

$$q + \sum_{\xi \in \mathcal{X}'} \frac{\kappa(w(q), \xi)}{\kappa(w(q), q)} \xi = \frac{1}{\kappa(w(q), q)} \sum_{\xi \in \mathcal{X}} \kappa(w(q), \xi) \xi = \frac{1}{\kappa(w(q), q)} \partial(w(q))$$

as desired. This concludes the proof. \square

An immediate consequence of Lemma 4.2.8 is the existence of an isomorphism $H_*(\mathcal{X}) \simeq H_*(\mathcal{X}')$ induced by ψ and ϕ . Finally, we provide a brief proof of the central theorem of discrete Morse theory.

PROOF OF THEOREM 4.2.2. Let $\{q_j \mid j = 1, \dots, J\}$ denote the set of all cells in \mathcal{Q} indexed arbitrarily. Consider the sequence of cell complexes $\mathcal{X}(j)$ defined inductively as follows: $\mathcal{X}(0) = \mathcal{X}$, and for each $j > 0$ we construct $\mathcal{X}(j)$ from $\mathcal{X}(j-1)$ by removing the paired cells q_j and $w(q_j)$ via the reduction step of Section 4.2.1. By Propositions 4.2.7 and 4.2.6, we know that there are chain maps ψ_j and ϕ_j between the chain complexes of $\mathcal{X}(j)$ and $\mathcal{X}(j-1)$. By Lemma 4.2.8 we know that these chain maps are in fact chain equivalences. Since all the paired cells are eventually removed, we know that $\mathcal{X}(J) = \mathcal{A}$. By Proposition 4.2.5, the incidence functions converge to κ_μ . Define the chain maps $\psi: \mathcal{C}(\mathcal{X}) \rightarrow \mathcal{C}(\mathcal{A})$ and $\phi: \mathcal{C}(\mathcal{A}) \rightarrow \mathcal{C}(\mathcal{X})$ by the compositions

$$\psi := \prod_{j=1}^J \psi_j \quad \text{and} \quad \phi := \prod_{j=J}^1 \phi_j$$

Note that $(\mathcal{A}, \kappa_\mu)$ is a complex by induction: we have already assumed that (\mathcal{X}, κ) is a complex as the base case. The reduction step does not alter the cells in \mathcal{A} and preserves the complex property. Finally, Lemma 4.2.8 guarantees the chain equivalence of ψ and ϕ . \square

4.3. Filtration-Subordinate Acyclic Matchings

Consider a filtration F of a cell complex (\mathcal{X}, κ) over the PID \mathbf{R} :

$$\mathcal{X}^1 \subset \mathcal{X}^2 \subset \dots \subset \mathcal{X}^m \subset \dots \subset \mathcal{X}^M = \mathcal{X}$$

and let $\mu = (\mathcal{A}; w: \mathcal{Q} \rightarrow \mathcal{K})$ be an acyclic matching on \mathcal{X} . For each $m \in \{1, \dots, M\}$ and $\mathcal{D} \in \{\mathcal{A}, \mathcal{K}, \mathcal{Q}\}$, define $\mathcal{D}^m = \mathcal{D} \cap \mathcal{X}^m$.

DEFINITION 4.3.1. We say that μ is *subordinate* to F if for each $m \in \{1, \dots, M\}$, the restriction $w^m := w|_{\mathcal{Q}^m}$ is a bijection from \mathcal{Q}^m to \mathcal{K}^m .

We check that if μ is F -subordinate, then each subcomplex \mathcal{X}^m inherits an acyclic matching from μ .

PROPOSITION 4.3.2. *If μ is F-subordinate, then it induces an acyclic matching $\mu^m = (\mathcal{A}^m, w^m : \mathcal{Q}^m \rightarrow \mathcal{X}^m)$ on each subcomplex \mathcal{X}^m , where w^m is defined by the restriction of w to \mathcal{Q}^m .*

PROOF. Let $\kappa^m : \mathcal{X}^m \times \mathcal{X}^m \rightarrow \mathbf{R}$ be the incidence function on \mathcal{X}^m obtained by restricting κ . For any $q \in \mathcal{Q}^m$, we have $\kappa^m(w^m(q), q) = \kappa(w(q), q)$ which is a unit in \mathbf{R} since μ is a partial matching on \mathcal{X} . If a path ρ of μ^m is a cycle, then it is also a cycle of μ , a contradiction. \square

We write $\mu^1 \subset \mu^2 \subset \dots \subset \mu^M = \mu$ to indicate that μ is F-subordinate.

4.3.1. The Filtered Morse Complex. Assume throughout that μ is F-subordinate.

PROPOSITION 4.3.3. *Let ρ be a path of μ and let $m \in \{1, \dots, M\}$. If the first cell \mathbf{s}_ρ lies in \mathcal{X}^m , then so do all the other cells of ρ .*

PROOF. Let $\rho = (q_1, w(q_1), \dots, q_d, w(q_d))$ and assume that $\mathbf{s}_\rho = q_1$ satisfies $q_1 \in \mathcal{X}^m$. Since μ is F-subordinate, the pairing w respects the filtration, and so $w(q_1)$ also lies in $\mathcal{X}^m \subset \mathcal{X}^m$. By definition of a path, $q_2 \prec w(q_1)$ and hence $q_2 \in \mathcal{X}^m$ by the subcomplex property. Proceeding in this way, we see that every cell of ρ lies in \mathcal{X}^m . \square

Denoting the Morse complex associated to each μ^m on \mathcal{X}^m by $(\mathcal{A}^m, \kappa_\mu^m)$, we see that the \mathcal{A}^m s constitute a filtration of \mathcal{A} .

PROPOSITION 4.3.4. $F_\mu := \{\mathcal{A}^m \mid m = 1, \dots, M\}$ is a filtration of the Morse complex $(\mathcal{A}, \kappa_\mu)$ associated to μ .

PROOF. First we show that $\kappa_\mu \upharpoonright_{\mathcal{A}^m \times \mathcal{A}^m} \equiv \kappa_\mu^m$ for each m . Given a in \mathcal{A}^m and an arbitrary $a' \in \mathcal{A}$, it suffices to check by (4.1) that there are no connections in $\mathcal{X} \setminus \mathcal{X}^m$ from a to a' . To see this, observe that any such connection ρ must have its first cell \mathbf{s}_ρ satisfy $\mathbf{s}_\rho \prec a$ and so $\mathbf{s}_\rho \in \mathcal{X}^m$ by the subcomplex property of \mathcal{X}^m . By Proposition 4.3.3, each element of ρ lies in \mathcal{X}^m as desired.

Now assume that $\kappa_\mu(a, a') \neq 0$ for some $a \in \mathcal{A}^m$. We will show that $a' \in \mathcal{A}^m$, thus proving the desired subcomplex property for $\mathcal{A}^m \subset \mathcal{A}$. From (4.1) we see that either $\kappa(a, a') \neq 0$ or there exists at least one connection ρ from a to a' with $\text{mul}(a \xrightarrow{\rho} a') \neq 0$. In the first case we have $a' \in \mathcal{X}^m$ by the fact that \mathcal{X}^m is a subcomplex of \mathcal{X} , so without loss of generality we assume that there exists some connection ρ from a to a' . Again, since $\mathbf{s}_\rho \prec a$, we know that each cell of ρ lies in \mathcal{X}^m . In particular, the last cell \mathbf{e}_ρ lies in \mathcal{X}^m . Since $a' \prec \mathbf{e}_\rho$ by definition of a connection, we see that $a' \in \mathcal{X}^m$ by the subcomplex property. \square

We call F_μ the *Morse filtration* associated to the F-subordinate acyclic matching μ . By Theorem 4.2.2, the homology modules of the Morse complex $(\mathcal{A}, \kappa_\mu)$ are isomorphic to those of (\mathcal{X}, κ) . We now extend this result to the level of persistent homology.

4.3.2. Obtaining a Filtered Chain Equivalence. The goal of this section is to prove the following main theorem.

THEOREM 4.3.5. *Let $F = \{\mathcal{X}^m \mid m = 1, \dots, M\}$ be a filtration of a complex (\mathcal{X}, κ) . Let μ be an F-subordinate acyclic matching on \mathcal{X} . Let $(\mathcal{A}, \kappa_\mu)$ be the associated Morse complex with Morse filtration $F_\mu = \{\mathcal{A}^m \mid m = 1, \dots, M\}$. Then for each m in $\{1, \dots, M\}$, n and p in \mathbb{N} , there exists an isomorphism*

$$H_n^p(\mathcal{X}^m) \simeq H_n^p(\mathcal{A}^m). \quad (4.6)$$

In order to make the proof of this theorem transparent, consider the function $b : \mathcal{X} \rightarrow \mathbb{N}$ given by

$$b(\xi) := \min\{m \in \mathbb{N} \mid \xi \in \mathcal{X}^m\}. \quad (4.7)$$

The two important properties of b are:

- (1) By Definition 4.3.1, $b(q) = b(w(q))$ for each $q \in \mathcal{Q}$, and
- (2) By the subcomplex property, $b(\xi) \leq b(\eta)$ whenever $\xi \prec \eta$.

Let $\{q_j \mid j = 1, \dots, J\}$ denote the set of all cells in \mathcal{Q} with the following additional constraint:

$$\text{if } b(q_j) > b(q_i) \text{ then } j > i.$$

This gives us positive integers $J_1 \leq J_2 \leq \dots \leq J_M = J$ such that $q_j \in \mathcal{X}^m$ if and only if $j \leq J_m$. Define chain maps $\psi^m: \mathcal{C}(\mathcal{X}^m) \rightarrow \mathcal{C}(\mathcal{A}^m)$ and $\phi^m: \mathcal{C}(\mathcal{A}^m) \rightarrow \mathcal{C}(\mathcal{X}^m)$ by the compositions

$$\psi^m := \prod_{j=1}^{J_m} \psi_j \quad \text{and} \quad \phi^m := \prod_{j=J_m}^1 \phi_j \quad (4.8)$$

where ψ_j and ϕ_j are the chain equivalences from the proof of Theorem 4.2.2.

Let $\Psi = \{\psi^m: \mathcal{C}(\mathcal{X}^m) \rightarrow \mathcal{C}(\mathcal{A}^m)\}$ and $\Phi = \{\phi^m: \mathcal{C}(\mathcal{A}^m) \rightarrow \mathcal{C}(\mathcal{X}^m)\}$. Note that these collections constitute maps between $F(\mathcal{X})$ and $F_\mu(\mathcal{A})$. The proof of Theorem 4.3.5 concludes with the following result.

PROPOSITION 4.3.6. *The maps Ψ and Φ are filtered chain equivalences.*

PROOF. Lemma 4.2.8 implies that ψ^m and ϕ^m are chain equivalences for all $m \in \{1, \dots, M\}$. Thus, it suffices to show that Ψ and Φ maps of chain filtrations $F(\mathcal{X})$ and $F_\mu(\mathcal{A})$. That is, given any $m \in \{1, \dots, M-1\}$ we will show that the following diagrams commute.

$$\begin{array}{ccc} \mathcal{C}(\mathcal{X}^m) & \longrightarrow & \mathcal{C}(\mathcal{X}^{m+1}) & & \mathcal{C}(\mathcal{X}^m) & \longrightarrow & \mathcal{C}(\mathcal{X}^{m+1}) \\ \psi^m \downarrow & & \downarrow \psi^{m+1} & & \phi^m \uparrow & & \uparrow \phi^{m+1} \\ \mathcal{C}(\mathcal{A}^m) & \longrightarrow & \mathcal{C}(\mathcal{A}^{m+1}) & & \mathcal{C}(\mathcal{A}^m) & \longrightarrow & \mathcal{C}(\mathcal{A}^{m+1}) \end{array}$$

Here the horizontal arrows represent the usual chain maps arising from inclusion.

Fix some $q \in \mathcal{Q}$ with $b(q) \geq m+1$. Note from the defining equation (4.3) that the map ψ associated to the removal of q differs from the identity only on q and $w(q)$, both of which have b -values exceeding $k+1$ by our explicit assumption on q and the first observed property of the function b . Therefore, $\psi^{m+1}|_{\mathcal{C}(\mathcal{X}^m)}$ is the identity map. Thus, Ψ is a filtered chain map by (4.8).

Similarly, we show that the map ϕ from (4.4) associated to the removal of q is the identity map on $\mathcal{C}(\mathcal{X}^m)$ whenever $b(q) \geq m+1$. Note by definition that $\phi(\eta)$ may differ from η only when $\kappa(\eta, Q) \neq 0$, i.e., when $q \prec \eta$. By the second observed property of the function b , we must have $b(\eta) \geq b(q) = m+1$ and so $\eta \in \mathcal{X} \setminus \mathcal{X}^m$ as desired. Thus, Φ is a filtered chain map as well by (4.8). \square

REMARK 4.3.7. In the homological algebraic language of Section 2.2, Ψ and Φ are weak inverse morphisms of the chain filtrations $F(\mathcal{X})$ and $F_\mu(\mathcal{A})$. In fact, it is easy to see that they are *strong* inverses in the sense of Definition 2.2.1. Since $\Psi \circ \Phi$ is the identity automorphism of $F_\mu(\mathcal{A})$, it suffices to observe that the chain homotopy Θ_q from (4.5) associated to the removal of some $q \in \mathcal{Q}$ with $b(q) = j$ also satisfies $\Theta_q(\eta) = 0$ whenever $b(\eta) \leq j-1$. But this is immediate from the definition: $\Theta_q(\eta)$ is only non-zero when $\eta = q$.

Simplifying Computation of Persistent Homology

◇

◇

Exitus acta probat.

– Ovid, *Heroides*

5.1. Introduction

The objective of this chapter is to describe an algorithm called `MorseReduce`, which takes a filtered complex \mathcal{X} as input, efficiently imposes a filtration-subordinate acyclic matching μ on \mathcal{X} , and outputs the filtered Morse complex \mathcal{A} associated to μ . By the results of Chapter 4, this filtered Morse complex has persistent homology isomorphic to that of \mathcal{X} , and hence the persistent homology of \mathcal{X} can be computed up to isomorphism by using the standard persistence algorithm [40] on the filtered Morse complex.

The popularity of persistent homology as a tool for understanding large datasets has led to a variety of highly efficient implementations and pre-processing algorithms (see [17], [38] for instance). To the best of our understanding, these approaches rely heavily on the efficient storability of cubical datasets of low dimension, and there appears to be little hope of similar techniques succeeding on other types of complexes. Since our approach with `MorseReduce` only requires face relations on the input filtered complex as encoded by the underlying incidence function, it applies in significantly broader contexts. The coreduction-based strategy of [30] has similarly wide applicability but it only pairs those cells which have gradient paths descending to unpaired cells of dimension zero, and therefore results in the reduction of much fewer cells when compared to `MorseReduce`.

5.2. The `MorseReduce` Algorithm

Consider a cell complex (\mathcal{X}, κ) over a PID \mathbf{R} filtered by $F = \{\mathcal{X}^m \mid m = 1, \dots, M\}$.

Theorem 4.3.5 implies that it is possible to compute the persistent homology groups of F by applying the algorithm of [40] to a potentially smaller Morse filtration $F_\mu = \{\mathcal{A}^m \mid m = 1, \dots, M\}$ associated with a Morse complex $(\mathcal{A}, \kappa_\mu)$ corresponding to some choice of F -subordinate acyclic matching μ on \mathcal{X} . The usefulness of this approach depends upon having an efficient algorithm for constructing the filtration F_μ and the Morse incidence function κ_μ , or equivalently the boundary operator ∂_μ on \mathcal{A} .

The filtration F_μ and incidence function κ_μ naturally depend on the choice of acyclic matching $\mu = (\mathcal{A}, w: \mathcal{Q} \rightarrow \mathcal{K})$. The trivial matching given by $\mathcal{A} = \mathcal{X}$ and $\mathcal{Q} = \mathcal{K} = \emptyset$ always exists, but results in the same filtration and thus provides no savings in computational cost. Clearly, the desired goal is to choose an acyclic matching which minimizes the number of cells in \mathcal{A} , or equivalently maximizes the number of cells paired by $w: \mathcal{Q} \rightarrow \mathcal{K}$. It is known (see [25, Section 4.5]) that in general the problem of constructing an optimal acyclic matching is MAX-SNP hard.

5.2.1. Coreduction Pairs and Gradient Chains. Our approach to producing an acyclic matching is based on the coreduction homology algorithm of Mrozek and Batko [29] which has proven effective in computing homology of complexes [18, 19]. Heuristically, this algorithm

is based on the following idea. Let $\mathcal{X}' \subset \mathcal{X}$ be a sub-collection of cells. A pair of cells $\xi, \eta \in \mathcal{X}'$ form a *coreduction pair* in \mathcal{X}' if restricted to $C_*(\mathcal{X}')$

$$\partial \xi = u \cdot \eta$$

where $u \in \mathbf{R}$ is a unit. In this case we make the identifications $\xi \in \mathcal{K}, \eta \in \mathcal{Q}$ and $w(\eta) = \xi$, and remove both ξ and η from \mathcal{X}' .

From (4.1) it is clear that κ_μ – and hence the Morse boundary operator ∂_μ – is defined by summing over all connections between cells in \mathcal{A} . However, even enumerating all the connections between two such cells is a combinatorially explosive proposition. To circumvent this summation, we make use of the observation that the coreduction-based construction of the pairing $w: \mathcal{Q} \rightarrow \mathcal{K}$ is done by building paths in reverse order. Heuristically, we assign to each cell $\zeta \in \mathcal{X}$ a chain $g(\zeta) \in C_*(\mathcal{A})$ called the *gradient chain* such that if $s \in \mathcal{A}$, then $g(s) = \partial_\mu s$.

Initially we set $g(\zeta) := 0$. However, as the coreduction algorithm is used to construct the acyclic matching – that is, as the paths are constructed – the value of $g(\zeta)$ is suitably modified. Thus, the computation of ∂_μ can be performed during the construction of $w: \mathcal{Q} \rightarrow \mathcal{K}$ using the subroutine `UpdateGradientChain` presented below.

We make use of the following notation in the algorithms: given $\xi \in \mathcal{X}$, the *coboundary* cells of ξ are given by

$$\text{cb}(\xi) := \{\eta \in \mathcal{X} \mid \xi \prec \eta\}.$$

It is not necessary to impose any specific order on the cells in $\text{cb}(\xi)$.

To emphasize that we only need to store each cell once rather than save a copy for each subcomplex \mathcal{X}^m containing that cell, we partition the cells in \mathcal{X} by setting

$$\mathcal{N}^m = \mathcal{X}^m \setminus \mathcal{X}^{m-1}, \quad m = 1, \dots, M$$

where $\mathcal{X}^0 = \emptyset$. Note that each cell $\xi \in \mathcal{X}$ lies uniquely in $\mathcal{N}^{b(\xi)}$ where $b: \mathcal{X} \rightarrow \mathbb{Z}$ is as defined in (4.7). The partition $\{\mathcal{N}^m \mid k = 1, \dots, M\}$ defines the input to our algorithms; each cell ξ is eventually excised from $\mathcal{N}^{b(\xi)}$ either as an element of \mathcal{A} or in a coreduction pair. Given a cell $\xi \in \mathcal{N}^m$, we denote by $\text{cb}_{\mathcal{N}}(\xi)$ and $\partial^{\mathcal{N}}(\xi)$ the coboundary cells and the boundary chain when restricted to $\{\mathcal{N}^*\}$. Once a cell is removed from \mathcal{N}^m , it is also removed from the corresponding $\text{cb}_{\mathcal{N}}$ and boundary $\partial^{\mathcal{N}}$ of the remaining cells. Similarly, the cells of the output Morse complex $(\mathcal{A}, \kappa_\mu)$ are also partitioned via $\mathcal{N}_{\mathcal{A}}^m = \mathcal{A}^m \setminus \mathcal{A}^{m-1}$.

TABLE 1. **Algorithm:** `UpdateGradientChain`

	In: $\xi \in \mathcal{X}$;
	Out: Updates $g(\zeta)$ for each $\zeta \in \text{cb}_{\mathcal{N}}(\xi)$
01	for each $\zeta \in \text{cb}_{\mathcal{N}}(\xi)$
02	if $\xi \in \mathcal{N}_{\mathcal{A}}^m$ for some m
03	$g(\zeta) \leftarrow g(\zeta) + \kappa(\zeta, \xi) \cdot \xi$
04	else
05	$g(\zeta) \leftarrow g(\zeta) + \kappa(\zeta, \xi) \cdot g(\xi)$
06	end if
07	end for

5.2.2. Subroutines for Removing Cells. The next two subroutines perform tasks pertaining to removing cells. The first subroutine – called `MakeCritical` – chooses an arbitrary α' of minimal dimension in a non-empty \mathcal{N}^m and excises it as an element of \mathcal{A} .

TABLE 2. **Algorithm:** MakeCritical**In:** $m \in \{1, \dots, M\}$ so that $\mathcal{N}^m \neq \emptyset$ **Out:** $a' \in \mathcal{N}_{\mathcal{A}}^m$

01	choose $a' \in \mathcal{N}^m$ of min dimension
02	add a' to $\mathcal{N}_{\mathcal{A}}^m$;
03	updateGradientChain(a')
04	remove a' from \mathcal{N}^m
05	$\partial_{\mu} a' \leftarrow g(a')$
06	return a'

Thus, the output a' of MakeCritical becomes a generator of the free module $C_*(\mathcal{A}^m)$. The gradient chains of remaining coboundary cells $\text{cb}_{\mathcal{N}}(a')$ are then updated to reflect their incidence with a' . In this manner, the construction of gradient chains is from the “bottom-up”. Finally, the action of the Morse boundary operator ∂_{μ} on a' is recovered from the corresponding gradient chain $g(a')$.

The obvious operation of the second subroutine, called RemovePair, is to perform the reduction step from Section 4.2 with respect to a single coreduction pair k, q from the complex.

TABLE 3. **Algorithm:** RemovePair**In:** $k, q \in \mathcal{N}^{m*}$ with $\partial^{\mathcal{N}}k = u \cdot q$, Queue of cells Que, $n \in \mathbb{N}$ **Out:** Removes (k, q) as a cell pair from \mathcal{N}^{m*}

01	remove k from \mathcal{N}^{m*}
02	enqueue $\text{cb}_{\mathcal{N}}(q)$ in Que
03	if $\dim q = n$
04	$g(q) \leftarrow -\frac{g(k)}{u}$
05	updateGradientChain(q)
06	end if
07	remove q from \mathcal{N}^{m*}

Recall that on the theoretical level coreduction pairs are identified as w -paired cells and hence they define the paths of μ . Thus, before the coreduction pair can be removed two additional steps need to be performed involving the remaining coboundary cells $\text{cb}_{\mathcal{N}}(q)$ of q . First, we check if the removal of q has created new coreduction pairs. For this, it suffices to check cells in the coboundary of q and so we add those cells to a queue structure. Secondly, if the pair (k, q) potentially lies on a path between unpaired cells of adjacent dimension, the gradient chains of q and hence of its remaining coboundary cells are updated by a call to UpdateGradientChain.

5.2.3. MorseReduce. These subroutines are combined to form our main algorithm MorseReduce. The input to this algorithm is a filtration F of a complex (\mathcal{X}, κ) partitioned by $\{\mathcal{N}^m\}$ as described above; the incidence function κ represents knowledge of the boundary operator ∂ . The output is a new filtration F_{μ} of the Morse complex (\mathcal{A}, κ) – partitioned by $\{\mathcal{N}_{\mathcal{A}}^m\}$ – corresponding to a coreduction-based F -subordinate acyclic matching $\mu = (\mathcal{A}, w : \Omega \rightarrow \mathcal{K})$. The Morse incidence function κ_{μ} is recovered from the boundary operator ∂_{μ} .

Note that we use a queue data structure Que which gets re-initialized once for each iteration of the outer **while** loop from line 02. We keep track of which cells are in Que so that no cell

TABLE 4. **Algorithm:** MorseReduce

	In: $\{\mathcal{N}_1^m\}^M$
	Out: $\{\mathcal{N}_A^m\}_1^M$
01	for each $m = 1, \dots, M$
02	while $\mathcal{N}^m \neq \emptyset$
03	$a' \leftarrow \text{MakeCritical}(m)$
04	Que := Empty Queue of Cells
05	enqueue $\text{cb}_{\mathcal{N}}(a')$ in Que
06	while Que $\neq \emptyset$
07	dequeue ξ from Que
08	if $\partial^{\mathcal{N}}\xi = 0$
09	enqueue $\text{cb}_{\mathcal{N}}(\xi)$ in Que
10	else if $\partial^{\mathcal{N}}\xi = u \cdot \eta$ for some $\eta \in \mathcal{N}^{b(\xi)}$ and unit $u \in \mathbf{R}$
11	RemovePair($\xi, \eta, \text{Que}, \dim a$)
12	end if
13	end while
14	end while
15	end for
16	return $\{\mathcal{N}_A^m\}_1^M$

is queued twice per such iteration. This can be achieved in practice either by storing an additional flag for each cell or by mirroring the queue in a separate data structure which has been optimized for search.

5.3. Verification of Correctness

We use Theorem 4.3.5 to confirm that the output filtration F_μ generated by the algorithm MorseReduce has the same persistent homology groups as those of the input filtration F .

THEOREM 5.3.1. *Let $F = \{\mathcal{X}^m \mid m = 1, \dots, M\}$ be a filtration of a complex (\mathcal{X}, κ) over a PID \mathbf{R} and define $\mathcal{N}^m := \mathcal{X}^m \setminus \mathcal{X}^{m-1}$ for each m . Then,*

- (1) MorseReduce terminates when applied to $\{\mathcal{N}_1^m\}^M$ and produces smaller collections of cells $\{\mathcal{N}_A^m\}_1^M$ from a Morse complex $(\mathcal{A}, \kappa_\mu)$ associated to a F -subordinate acyclic matching $\mu = (\mathcal{A}; w : \mathcal{Q} \rightarrow \mathcal{K})$ on \mathcal{X} .
- (2) The output $\{\mathcal{N}_A^m\}$ defines a filtration $F_\mu = \{\mathcal{A}^m\}_1^M$ of $(\mathcal{A}, \kappa_\mu)$ where each subcomplex \mathcal{A}^m is given by $\bigcup_{\ell=1}^m \mathcal{N}_A^\ell$ and the underlying incidence function κ_μ corresponds to the Morse boundary operator ∂_μ .
- (3) For each p, n and m , we have an isomorphism of the corresponding persistent homology group

$$H_n^p(\mathcal{X}^m) \cong H_n^p(\mathcal{A}^m).$$

PROOF. Each iteration of the outer **while** loop from line 02 permanently excises at least one cell a from \mathcal{N}^* via MakeCritical. In practice many more cells are also removed, since if (ξ, η) appear as a coreduction pair then they are excised during the RemovePair subroutine. Furthermore, observe that their appearance as a coreduction pair is equivalent to their being part of a path that terminates at a . The fact that no cell is queued twice during any iteration of the second **while** loop in line 07 guarantees the absence of infinite loops. Moreover, it is clear that the final size of each \mathcal{N}_A^m is smaller than the initial size of \mathcal{N}^m because MakeCritical

is only called once per iteration of the outer **while** loop and each call to `MakeCritical` results in a single cell from \mathcal{N}^m being removed and stored in the corresponding \mathcal{N}_A^m . Thus, $\mathcal{N}_A^m \subset \mathcal{N}^m$ for each m .

Observe from line 10 that if (ξ, η) is sent to `RemovePair`, then $b(\xi) = b(\eta)$ and $\kappa(\xi, \eta)$ equals some unit u in \mathbf{R} . Let $m_* = b(\xi)$, and note that defining $w^{m_*}(\eta) = \xi$ for each such pair constructs $w^{m_*}: \mathcal{Q}^{m_*} \rightarrow \mathcal{K}^{m_*}$. Combining this pairing information with the output of `MakeCritical` produces a partial matching $\mu = (\mathcal{A}, w: \mathcal{Q} \rightarrow \mathcal{K})$ which is subordinate to F .

To see that this partial matching is acyclic, observe from lines 10 and 11 that a pairing $w(\eta) = \xi$ is only made when η is the last remaining face of ξ , i.e., the unique cell in $\{\zeta \in \mathcal{N}^{b(\xi)} \mid \zeta \prec \xi\}$. Thus, all elements of \mathcal{Q} satisfying that could possibly come after ζ in a path must have already been excised before the pair (ξ, η) . Thus, no path of μ can be a cycle.

By Theorem 4.3.5, in order to show that the output determines a filtration F_μ with isomorphic persistent homology to F , it suffices to establish that F_μ is the Morse filtration associated to μ . Thus, we must ensure that the stored boundary ∂_μ of each cell $a \in \mathcal{A}$ built from the corresponding gradient chain $g(a)$ equals the boundary operator corresponding to the Morse incidence function κ_μ from (4.1). This is addressed by the subsequent proposition, which concludes the proof. \square

The proof of the following proposition employs the usual inner product $\langle \cdot, \cdot \rangle : C(\mathcal{X}) \times C(\mathcal{X}) \rightarrow \mathbf{R}$ on chains of the input complex (\mathcal{X}, κ) obtained by treating the cells in \mathcal{X} as an orthonormal basis.

PROPOSITION 5.3.2. *Assume the hypotheses and notation of Theorem 5.3.1. For cells a and a' in \mathcal{A} ,*

$$\langle g(a), a' \rangle = \kappa_\mu(a, a')$$

PROOF. We provide a brief summary of how gradient chains are constructed. Assume throughout that a' is removed via `MakeCritical`. Consider the following two cases.

- [a] Assume that ζ is an unremoved cell with $a' \prec \zeta$. Then, by line 02 of `MakeCritical` and the subsequent call to `UpdateGradientChain`, the gradient chain $g(\zeta)$ of ζ is incremented as follows:

$$g(\zeta) \leftarrow g(\zeta) + \kappa(\zeta, a') \cdot a'$$

Since this is the first instance of a' being added to gradient chains, we are guaranteed $\langle g(\zeta), a' \rangle = \kappa(\zeta, a')$ when `MakeCritical` returns a' .

- [q] Assume ζ is an arbitrary unremoved cell. Each cell q excised as an element of \mathcal{Q} via `RemovePair` inherits its gradient chain from the existing gradient chain of its paired cell $w(q)$ by the formula

$$g(q) = \frac{g(w(q))}{\kappa(w(q), q)}$$

This follows from line 04 of `RemovePair`. Since `UpdateGradientChain` is called in the next line, each remaining cell ζ satisfying $q \prec \zeta$ has its gradient chain incremented by $\kappa(\zeta, q) \cdot g(q)$. By the preceding formula for $g(q)$, we have

$$g(\zeta) \leftarrow g(\zeta) + \frac{\kappa(\zeta, q)}{-\kappa(w(q), q)} \cdot g(w(q))$$

Thus, there are two ways a critical cell a' appears with non-zero coefficient in the gradient chain $g(\zeta)$ of some hitherto unremoved cell ζ : either $a' \prec \zeta$ and we directly apply [a], or $\langle g(w(q)), a' \rangle \neq 0$ for some previously removed $q \in \mathcal{Q}$ with $q \prec \zeta$ and we apply [q]. Combining

these contributions, we have the following formula

$$\langle g(\zeta), a' \rangle = \kappa(\zeta, a') + \sum_{q \in \mathcal{Q}} \frac{\kappa(\zeta, q)}{-\kappa(w(q), q)} \langle g(w(q)), a' \rangle \quad (5.1)$$

Now assume that a cell a is eventually removed by `MakeCritical`. Recalling (4.1), we substitute $\zeta = a$ in (5.1) to get

$$\langle g(a), a' \rangle = \kappa(a, a') + \sum_{q \in \mathcal{Q}} \frac{\kappa(a, q)}{-\kappa(w(q), q)} \langle g(w(q)), a' \rangle$$

Applying (5.1) recursively to each $\langle g(w(q)), a' \rangle$ in the expression above completes the argument. \square

5.4. Complexity

Let (\mathcal{X}, κ) be a complex over a PID \mathbf{R} filtered by $\mathcal{F} = \{\mathcal{X}^k\}_1^K$ with face partial order \preceq generated by the usual relation:

$$\xi \prec \xi' \text{ if } \kappa(\xi', \xi) \neq 0 \in \mathbf{R}.$$

5.4.1. Parameters and Assumptions. We will describe the computational cost of using `MorseReduce` to construct an F-subordinate acyclic matching $(\mathcal{A}, w : \mathcal{Q} \rightarrow \mathcal{K})$ as well as the associated Morse complex $(\mathcal{A}, \kappa_\mu)$. These bounds on worst-case performance will be established in terms of the following basic *complexity parameters*.

- (1) The *input size* – denoted by n – is the number of cells in \mathcal{X} .
- (2) The *output size* is the number of cells in the filtered Morse complex \mathcal{A} which we denote by m . Note that m is partitioned by $m = m_0 + \dots + m_D$ where m_d is the cardinality of d -dimensional cells in \mathcal{A} . Constructing an optimal acyclic matching – that is, a matching that minimizes m – is MAX-SNP hard [25]. Providing sharp bounds on optimal m values relative to n for arbitrary complexes would require major breakthroughs in algebraic topology as well as graph theory. Therefore, we leave m as a parameter.
- (3) The *coboundary mass* p of \mathcal{X} is defined as

$$p = \sup_{\xi \in \mathcal{X}} \#\{\eta \in \mathcal{X} \mid \kappa(\eta, \xi) \neq 0\},$$

where $\#$ denotes cardinality. Thus, the coboundary mass bounds the number of cells $\eta \in \mathcal{X}$ which satisfy $\xi \prec \eta$ for a given cell $\xi \in \mathcal{X}$. Even though p may safely be bounded by n , in most situations this is a gross over-estimate. For example, the coboundary mass of a d -dimensional cubical grid is bounded above by $2d$ independent of the total number of cubes present.

For the purposes of complexity analysis, we also make these two *simplifying assumptions*:

- (1) we assume that adding, removing or locating a cell $\xi \in \mathcal{X}$ incurs a constant cost, and
- (2) we assume that ring operations in \mathbf{R} may be performed in constant time so that the cost of adding and scaling gradient chains is linear in the length of the chains involved.

5.4.2. Evaluating Complexity. We begin by evaluating the complexity of a single iteration of the outer **while** loop from Line 02 of `MorseReduce`. Assume that in this iteration the call to `MakeCritical` via Line 03 has returned a cell A' of dimension d . Since in each iteration of this **while** loop we add a cell to `Que` at most once, the maximum size attainable by `Que` is n . Moreover, each `Que` insertion involves testing the coboundary of a cell which requires at most p operations. In light of these bounds, we will just assume that the total cost of managing the

Que data structure within a single **while** iteration depends linearly on $n \cdot p$ and we will not separately tabulate the cost of each Que operation.

We also require the following observations regarding the cost of the three subroutines in terms of the complexity parameters defined previously.

- (1) The cost of calling `UpdateGradientChain` on a d -dimensional cell equals $O(p \cdot m_d)$. This follows from the fact that we must iterate over each cell ζ in the remaining coboundary of ξ and update the gradient chain $g(\zeta)$ which consists of d -dimensional cells in \mathcal{A} .
- (2) A call to `MakeCritical` in Line 03 also costs $O(p \cdot m_d)$, since the only non-trivial operation is the call to `UpdateGradientChain` in Line 03.
- (3) In the worst case, the **if** statement from Line 03 of `RemovePair` always evaluates positively and hence `UpdateGradientChain` is called. Thus, each call to `RemovePair` also incurs a worst case cost of $O(p \cdot m_d)$ since all other non-trivial operations only involve Que insertion.

Since the inner **while** loop from Line 06 of `MorseReduce` depends only on the size of Que, it may run at most n times. Thus, the cost of iterating the outer **while** loop from Line 02 reduces to a single call to `MakeCritical`, the management of the Que structure, and at most n calls to `RemovePair`. Adding these respective contributions, the total cost of a single iteration of this outer **while** loop equals

$$O(p \cdot m_d) + O(n \cdot p) + O(n \cdot p \cdot m_d)$$

The third quantity clearly dominates the first two, so the desired complexity estimate of the outer **while** loop when A' has dimension d is $O(n \cdot p \cdot m_d)$.

It now suffices to estimate how many iterations of the outer **while** loop are actually executed in a single instance of `MorseReduce`. But this is straightforward: each such iteration corresponds to exactly one cell $A' \in \mathcal{A}$ as returned by `MakeCritical`, so this **while** loop executes precisely m times. Partitioning $m = m_0 + \dots + m_D$ by dimension as usual, we estimate the following total cost of running `MorseReduce` in terms of our complexity parameters:

$$O\left(n \cdot p \cdot \sum_{d=0}^D m_d^2\right)$$

In light of this expression, it is convenient to define the number $\tilde{m} \leq m^2$ by

$$\tilde{m} = \sum_{d=0}^D m_d^2.$$

Thus, we have the following result regarding the computational complexity of `MorseReduce`.

PROPOSITION 5.4.1. *Assume that `MorseReduce` is executed on a filtered complex \mathcal{X} of top dimension D , size n and coboundary mass p . If the resulting Morse complex \mathcal{A} has size $m = m_0 + \dots + m_D$, then the worst-case complexity is bounded by $O(n \cdot p \cdot \tilde{m})$, where $\tilde{m} = m_0^2 + \dots + m_D^2$.*

Thus, the cost of computing the maps induced on homology by inclusions $\mathcal{X}^k \subset \mathcal{X}^{k+1}$ in the filtered complex \mathcal{X} over any PID \mathbf{R} reduces from $O(n^3)$ without `MorseReduce` to $O(n \cdot p \cdot \tilde{m} + m^3)$ after `MorseReduce`. When m is much smaller than n , the first term is dominant and one observes essentially linear cost in terms of the input size n . In the special case when \mathbf{R} is a field, recall that the standard persistence algorithm of [40] has cubic complexity in the size of the input filtered complex. Therefore, the total cost of computing the persistence intervals of \mathcal{X} via the algorithm of [40] after applying `MorseReduce` to \mathcal{X} also equals $O(n \cdot p \cdot \tilde{m} + m^3)$.

REMARK 5.4.2. The efficiency of our approach depends crucially on m being much smaller than n . In the worst case, no cells get paired and we are left with $m = n$. Examples of filtered complexes which realize these *pathologies* may be easily constructed in two ways:

- (1) consider a complex \mathcal{X} such that any non-zero incidence $\kappa(\xi, \xi') \in \mathbf{R}$ is never a unit for any pair of cells $\xi, \xi' \in \mathcal{X}$. Alternately,
- (2) consider a complex \mathcal{X} so that whenever $\kappa(\xi, \xi') \neq 0$ for cells ξ, ξ' we have $b(\xi) \neq b(\xi')$. Since matched cells are required to have the same b -values by Definition 4.3.1, no non-trivial matching is possible in this case.

It is easy to test the input complex for both pathologies in $O(n \cdot p)$ time by checking each pair of cells $\xi, \xi' \in \mathcal{X}$ with non-trivial incidence $\kappa(\xi, \xi') \neq 0$. Moreover, these pathologies are extremely rare in practical situations such as those involving simplicial or cubical complexes arising from typical experimental data, since

- (1) for both cubical and simplicial complexes all non-zero incidences are units ± 1 in any PID \mathbf{R} , so the first pathology is avoided,
- (2) often, the b -values are only prescribed on top-dimensional cells (such as grayscale pixel or voxel values for image data). In these situations, each lower dimensional cell recursively inherits its b -value as the minimum b -value encountered among its co-boundary cells. This guarantees the existence of at least some cells $\xi \prec \xi'$ with $b(\xi) = b(\xi')$ and avoids the second pathology,
- (3) in other cases, the b -values are inherited from lower dimensional cells. A prime example is the Vietoris-Rips complex built around point cloud data. Here each simplex inherits the maximum b -value encountered in its 1-skeleton. Again, this process ensures the existence of dimensionally adjacent cells which share b -values and hence avoids the second pathology.

As we demonstrate in Section 5.5, Morse theoretic pre-processing is remarkably effective for computing persistent homology of several types of filtered complexes arising from experimental data.

We believe that a more nuanced approach to analyzing the effectiveness of combinatorial Morse theory would require imposing reasonable probability measures on the set of all complexes and proving statements regarding the expected fraction of cells reduced. We leave such considerations to future work.

5.5. Experimental Results

The popularity of persistent homology as a tool for understanding large datasets has led to a variety of highly efficient implementations and preprocessing algorithms (see [33, 17, 38] for instance). These approaches rely heavily on the efficient storability of cubical datasets of low dimension over \mathbb{Z}_2 coefficients, and there appears to be little hope of similar techniques succeeding on other types of complexes. Since MorseReduce only requires face relations on the input complex as encoded by the underlying incidence function, it applies to filtered complexes independent of coefficient ring and dimensionality. The coreduction-based strategy of [30] has similarly wide applicability but it only pairs those cells which have gradient paths descending to unpaired cells of dimension zero, and therefore results in the reduction of much fewer cells when compared to MorseReduce.

Note that since the output of MorseReduce is a filtration in its own right, it is possible to iterate the algorithm until the number of reductions performed becomes essentially negligible. Thus, the cells output by an iteration of MorseReduce get further partitioned by the subsequent iteration and may get paired by the associated acyclic matching.

We demonstrate the results of `MorseReduce` on cubical grids, simplicial complexes, Vietoris-Rips complexes and movies. The cubical complexes come from sub-level sets of finite element Cahn-Hilliard simulations and the simplicial complexes arise from brain imaging data. The Vietoris-Rips complexes come from point clouds of experimental granular flow data. The largest datasets by far, courtesy of M. Schatz, are two black and white movies obtained by segmenting Rayleigh-Bénard convection data, each successive frame consisting of about 155,000 three dimensional cubes.

The implementation of `MorseReduce` benchmarked here was coded in C++ using the standard template library and compiled using the GNU C++ compiler with optimization level O3. All computations were performed on an Intel Core i5 machine with 32 GB of available RAM and virtual memory disabled. The source code for our implementation is available at [3].

The comparison is with our implementation of the standard algorithm for computing persistent homology as found in [40] which we will denote by SP. While this algorithm may also be found in various flavors and as part of the software package `jPlex` [2] or from the `Dionysus` project [1], the present comparison is fair because the same data structures are used in both cases. The SP results simply provide the time taken when no discrete Morse theoretic pre-processing is performed while holding all other implementation-specific factors constant. Thus, if more efficient implementations of SP exist, then preprocessing with `MorseReduce` will vastly improve the performance of those implementations as well.

While the results of Theorem 5.3.1 apply to input filtrations over any PID \mathbf{R} , the usual computation of persistence intervals via SP requires \mathbf{R} to be a field. In the experimental results that follow, we have performed all reductions over \mathbb{Z} , but we assume $\mathbf{R} = \mathbb{Z}_2$ throughout when applying SP to the reduced filtration output by `MorseReduce`. The following table demonstrates the performance comparison of computing persistence with and without pre-processing by `MorseReduce`.

Type	Dim	# Frames	# Cells	Red. # Cells	SP	MR + SP
C	2	16	0.26 M	2.31 K	2.4	0.9
C	3	25	1.24 M	3.35 K	12.7	4.2
C	3	2400	5.25 M	9.50 K	195.8	46.6
S	5	20	12.86 M	27.13 K	73.9	8.8
S	5	5000	0.86 M	99.8 K	1951.6	303.1
VR	2	100	2.34 M	86.33 K	1277.0	37.7
VR	3	50	9.50 M	18.31 K	286.5	47.2
VR	3	250	1.29 M	53.95 K	551.1	125.2
M	3	209	259.21 M	1.25 K	DNF	7,213
M	3	215	266.67 M	2.10 K	DNF	7,416

TABLE 5. Experimental Results.

The table is arranged as follows: the first column indicates the type of complexes in the filtration (Cubical, Simplicial, Vietoris-Rips or Movie) while the second column contains the maximum dimension of the cells present in the filtration. The third column contains the length M of each input filtration $\mathcal{F} = \{\mathcal{X}^m\}_1^M$. The next two columns provide the size (in number of cells) of the filtration before and after Morse reduction. The penultimate column provides the time taken by our implementation of SP to compute persistence intervals over \mathbb{Z}_2 of the filtration, whereas the last column provides the total time taken to first apply `MorseReduce` and then compute the persistence intervals of the reduced filtration with SP. DNF indicates that the given algorithm failed to terminate because it ran out of memory. All times are in seconds.

A final note to illustrate the power and scalability of the Morse theoretic approach: the movie datasets were far too large to be held in memory all at once. Our approach involved storing about 30 frames at a time and removing paired cells from all but the last frame. This freed up considerable memory which we used to input the remaining portions of the movies in pieces, each comprising 30 consecutive frames. At each stage we left the last frame unreduced so that the next piece of the movie could be attached to it, and so on. In this way, extremely large and complicated persistence computations may be brought within the scope of commodity hardware. To the best of our knowledge, there is no other publically available technique which yields persistence intervals of a large filtration from such local computations without ever holding all the cells in memory at once.

CHAPTER 6

Applications to Exact Sequences

*The infernal Serpent; he it was whose guile,
stirred up with envy and revenge, deceived
the mother of mankind, what time his pride
had cast him out from Heaven, with all his host;*

–John Milton, *Paradise Lost*

6.1. Introduction

Exact sequences of chain complexes are basic objects of fundamental interest in algebraic topology and homological algebra. Given a short exact sequence S of free, finitely generated chain complexes \mathcal{C}, \mathcal{D} and \mathcal{E} over a principal ideal domain \mathbf{R} ,

$$0 \rightarrow \mathcal{C} \xrightarrow{\alpha} \mathcal{D} \xrightarrow{\beta} \mathcal{E} \rightarrow 0$$

there exist well-defined *connecting homomorphisms* $\Delta_n : H_n(\mathcal{E}) \rightarrow H_{n-1}(\mathcal{C})$ of homology modules so that the following sequence $\mathcal{L}(S)$ is exact

$$\dots \xrightarrow{\Delta_{n+1}} H_n(\mathcal{C}) \xrightarrow{\alpha_n^*} H_n(\mathcal{D}) \xrightarrow{\beta_n^*} H_n(\mathcal{E}) \xrightarrow{\Delta_n} H_{n-1}(\mathcal{C}) \xrightarrow{\alpha_{n-1}^*} \dots$$

Here α^* and β^* are the induced maps on homology corresponding to the chain maps α and β respectively.

The explicit construction of these connecting homomorphisms Δ – and hence, of $\mathcal{L}(S)$ – has various important applications in computational topology and dynamics. For instance, the Mayer-Vietoris long exact sequence [20, Ch. 2.2] allows one to compute the homology modules of a cell complex in terms of the homology modules of its subcomplexes and their intersections. Similarly, the long exact sequence of a triple [20, Ch. 2.1] is used to compute Conley indices [11, 14] in computational dynamics. While our focus is on constructing these two long exact sequences, the methods outlined here are rather general and apply to a significantly larger class of such constructions.

The construction of Δ typically involves bringing matrix representatives of the chain maps and/or the boundary operators into normal form. These operations are at least of cubical complexity in the size of the bases of the underlying chain modules. In many practical cases, these bases are large enough to render matrix algebra untenable in terms of both time taken and memory consumed.

Our strategy utilizes discrete Morse theory from Chapter 4 to perform efficient reductions on the bases of the chain modules while preserving the underlying homology modules. This generates a new short exact sequence S' which is weakly equivalent to S but has significantly smaller chain modules. Consequently,

- (1) $\mathcal{L}(S')$ is isomorphic to $\mathcal{L}(S)$, and
- (2) the matrices which need to be normalized for constructing $\mathcal{L}(S')$ are typically much smaller than those needed to construct $\mathcal{L}(S)$.

Thus, the connecting homomorphisms may be constructed for $\mathcal{L}(S')$ via the usual matrix operations on these smaller matrices and then transported isomorphically to $\mathcal{L}(S)$.

We briefly discuss the complexity of constructing long exact sequences in Section 6.2. Discrete Morse theory is then applied to simplify constructions of the long exact sequence of a triple and the Mayer Vietoris sequence in Sections 6.3 and 6.4 respectively.

6.2. The Zigzag Lemma

The Zigzag Lemma is a standard tool in homological algebra which describes the construction of long exact sequences in homology as a consequence of the *Snake Lemma*. This result describes a functorial construction from the category of short exact sequences of chain complexes to the category of long exact sequences.

Its statement and detailed proof may be found in most standard presentations of algebraic topology [20, Theorem 2.13] and homological algebra [39, Lemma 1.3.2]. We provide a brief outline of the argument in order to discuss the complexity of various algebraic operations involved therein.

LEMMA 6.2.1 (Zigzag Lemma). *Consider a short exact sequence S of chain complexes given by*

$$0 \rightarrow \mathcal{C} \xrightarrow{\alpha} \mathcal{D} \xrightarrow{\beta} \mathcal{E} \rightarrow 0.$$

There exist canonical connecting morphisms $\Delta_n : H_n(\mathcal{E}) \rightarrow H_{n-1}(\mathcal{C})$ of homology modules such that the following sequence $\mathcal{L}(S)$ is exact

$$\dots \xrightarrow{\Delta_{n+1}} H_n(\mathcal{C}) \xrightarrow{\alpha_n^*} H_n(\mathcal{D}) \xrightarrow{\beta_n^*} H_n(\mathcal{E}) \xrightarrow{\Delta_n} H_{n-1}(\mathcal{C}) \xrightarrow{\alpha_{n-1}^*} \dots$$

PROOF. Consider the following commutative diagram induced by the boundary operators from S_n to S_{n-1} :

$$\begin{array}{ccccccc} C_n & \xrightarrow{\alpha_n} & D_n & \xrightarrow{\beta_n} & E_n & \longrightarrow & 0 \\ \downarrow \partial_n & & \downarrow \delta_n & & \downarrow d_n & & \\ 0 & \longrightarrow & C_{n-1} & \xrightarrow{\alpha_{n-1}} & D_{n-1} & \xrightarrow{\beta_{n-1}} & E_{n-1} \end{array}$$

Given a cycle $\mathbf{x} \in E_n$, by the surjectivity of β there exists some $\mathbf{y} \in D_n$ with $\beta(\mathbf{y}) = \mathbf{x}$. Let $\mathbf{z} \in D_{n-1}$ be given by $\mathbf{z} = \delta_n(\mathbf{y})$. It is easy to see that $\mathbf{z} \in \ker \beta$, since β is a chain map:

$$\beta(\mathbf{z}) = \beta \circ \delta_n(\mathbf{y}) = d_n \circ \beta(\mathbf{y}) = d_n(\mathbf{x}) = 0.$$

Here the last equality follows from the assumption that \mathbf{x} is a cycle. By exactness of the sequence at the bottom, $\mathbf{z} \in \text{img } \alpha$, so there exists some $\mathbf{w} \in C_{n-1}$ with $\alpha(\mathbf{w}) = \mathbf{z}$. Since α is a chain map, we have

$$\alpha \circ \partial_{n-1}(\mathbf{w}) = \delta_{n-1} \circ \alpha(\mathbf{w}) = \delta_{n-1}(\mathbf{z}) = \delta_{n-1} \circ \delta_n(\mathbf{y}) = 0$$

By the injectivity of α , we have $\partial_{n-1}(\mathbf{w}) = 0$ and therefore \mathbf{w} is a cycle. It is easy to check that defining Δ_n as the map which sends the homology class of \mathbf{x} to the homology class of \mathbf{w} yields the desired long exact sequence in a well-defined manner. \square

In order to make the complexity of constructing Δ transparent, consider the simpler situation where \mathbf{R} is a principal ideal domain and each \mathbf{R} -module in \mathcal{C} , \mathcal{D} and \mathcal{E} is finitely generated and free; assume also that the morphisms α and β have been expressed as matrices with \mathbf{R} coefficients relative to some fixed bases of these free modules. Here are the four essential steps in the proof outlined above:

- (1) Computing the homology modules of \mathcal{E} so that we may choose representative cycles \mathbf{x} ,
- (2) Solving the matrix equation $\beta(\mathbf{y}) = \mathbf{x}$ for \mathbf{y} ,
- (3) Solving the matrix equation $\alpha(\mathbf{w}) = \mathbf{z}$ for \mathbf{w} , where $\mathbf{z} = \delta_n(\mathbf{y})$, and finally
- (4) Computing the homology modules of \mathcal{C} so that we may express \mathbf{w} as a \mathbf{R} -linear combination of representative cycles.

Each step reduces to performing row and column operations on matrices with \mathbf{R} coefficients. For instance, computing the homology modules of \mathcal{C} and \mathcal{E} is a matter of bringing the corresponding matrix representations of the boundary operators ∂ and d into their Smith normal forms. Similarly, solving the matrix equation $\beta(\mathbf{y}) = \mathbf{x}$ for \mathbf{y} involves bringing the matrix representation of β into reduced echelon form.

In the general spirit of the rest of our basic approach, reducing the sizes of the bases of the chain modules C_n , D_n and E_n in an efficient manner while preserving homology modules would also reduce the sizes of all matrices in sight, thereby providing a considerable computational advantage. The following classical result [31, Theorem 24.2] establishes functoriality of the construction outlined above.

PROPOSITION 6.2.2. *Given a morphism $\Omega : S \rightarrow S'$ of short exact sequences of chain complexes*

$$\begin{array}{ccccccccc} 0 & \longrightarrow & \mathcal{C} & \xrightarrow{\alpha} & \mathcal{D} & \xrightarrow{\beta} & \mathcal{E} & \longrightarrow & 0 \\ & & \downarrow \omega^{\mathcal{C}} & & \downarrow \omega^{\mathcal{D}} & & \downarrow \omega^{\mathcal{E}} & & \\ 0 & \longrightarrow & \mathcal{C}' & \xrightarrow{\alpha'} & \mathcal{D}' & \xrightarrow{\beta'} & \mathcal{E}' & \longrightarrow & 0 \end{array}$$

there is an induced morphism $\mathcal{L}(\Omega) : \mathcal{L}(S) \rightarrow \mathcal{L}(S')$ of the associated long exact sequences. That is, the following diagram commutes.

$$\begin{array}{ccccccccccc} \dots & \xrightarrow{\Delta_{n+1}} & H_n(\mathcal{C}) & \xrightarrow{\alpha_n^*} & H_n(\mathcal{D}) & \xrightarrow{\beta_n^*} & H_n(\mathcal{E}) & \xrightarrow{\Delta_n} & H_{n-1}(\mathcal{C}) & \xrightarrow{\alpha_{n-1}^*} & \dots \\ & & \downarrow \omega_n^{\mathcal{C}*} & & \downarrow \omega_n^{\mathcal{D}*} & & \downarrow \omega_n^{\mathcal{E}*} & & \downarrow \omega_{n-1}^{\mathcal{C}*} & & \\ \dots & \xrightarrow{\Delta'_{n+1}} & H_n(\mathcal{C}') & \xrightarrow{\alpha_n'^*} & H_n(\mathcal{D}') & \xrightarrow{\beta_n'^*} & H_n(\mathcal{E}') & \xrightarrow{\Delta'_n} & H_{n-1}(\mathcal{C}') & \xrightarrow{\alpha_{n-1}'^*} & \dots \end{array}$$

This functoriality enables us to bypass some of the algebraic complexity of constructing Δ in the following way. We use algebraic Morse theory to perform homology-preserving reductions on the bases of the underlying chain modules of \mathcal{C} , \mathcal{D} and \mathcal{E} . This produces a short exact sequence S' of chain complexes whose chain modules have smaller bases than those of S along with an explicit weak equivalence $\Omega : S \rightarrow S'$. Then, the induced map $\mathcal{L}(\Omega) : \mathcal{L}(S) \rightarrow \mathcal{L}(S')$ is an isomorphism of long exact sequences because all the vertical maps are invertible. Thus, constructing Δ' via the Zigzag Lemma for S' allows us to recover Δ for S via the simple formula

$$\Delta_n \equiv (\omega_{n-1}^{\mathcal{C}*})^{-1} \circ \Delta'_n \circ (\omega_n^{\mathcal{E}*})$$

We remark here that evaluating this composition reduces to *multiplying* matrices rather than bringing them into normal form; in this case, it is well known that sparseness actually leads to significant computational gains.

6.3. The Long Exact Sequence of a Triple

Let (X, κ) be a cell complex in the sense of Definition 2.3.1 over a principal ideal domain \mathbf{R} . Consider a filtration F of X given by $\mathcal{U} \subset \mathcal{V} \subset X$ where \mathcal{U} and \mathcal{V} are subcomplexes and let $F(X)$ be the corresponding chain filtration:

$$C_*(\mathcal{U}) \xrightarrow{i} C_*(\mathcal{V}) \xrightarrow{j} C_*(X)$$

where i and j are chain maps induced by inclusion. The *long exact sequence of the triple* \mathcal{U}, \mathcal{V} and \mathcal{X} is defined to be the long exact sequence $\mathcal{L}(S)$ associated to the following short exact sequence S of relative chain complexes

$$0 \rightarrow C_*(\mathcal{V}, \mathcal{U}) \xrightarrow{\alpha} C_*(\mathcal{X}, \mathcal{U}) \xrightarrow{\beta} C_*(\mathcal{X}, \mathcal{V}) \rightarrow 0 \quad (6.1)$$

where α is the inclusion of relative chains and β is the projection of relative chains. That is,

$$\alpha(\mathbf{v} + C_*(\mathcal{U})) = j(\mathbf{v}) + C_*(\mathcal{U}) \text{ and } \beta(\mathbf{x} + C_*(\mathcal{U})) = \mathbf{x} + C_*(\mathcal{V})$$

for $\mathbf{v} \in C_*(\mathcal{V})$ and $\mathbf{x} \in C_*(\mathcal{X})$.

If $\mathcal{U} = \emptyset$, then S reduces to the familiar *short exact sequence of the pair* $(\mathcal{X}, \mathcal{V})$ given by

$$0 \rightarrow C_*(\mathcal{V}) \rightarrow C_*(\mathcal{X}) \rightarrow C_*(\mathcal{X}, \mathcal{V}) \rightarrow 0$$

Let $\mathcal{Y} \in \{\mathcal{U}, \mathcal{V}, \mathcal{X}\}$. Assume the existence of acyclic matchings

$$\mu_{\mathcal{Y}} = (\mathcal{A}_{\mathcal{Y}}; \omega_{\mathcal{Y}} : \mathcal{Q}_{\mathcal{Y}} \rightarrow \mathcal{K}_{\mathcal{Y}}) \text{ on } \mathcal{Y} \quad (6.2)$$

subordinate to the filtration $\mathcal{U} \subset \mathcal{V} \subset \mathcal{X}$. In practice, one may construct these matchings via the algorithms of Chapter 5. By Theorem 4.3.5, note that $\mathcal{A} := \mathcal{A}_{\mathcal{X}}$ is the Morse complex associated to $\mu_{\mathcal{X}}$ and $\mathcal{A}_{\mathcal{U}} \subset \mathcal{A}_{\mathcal{V}} \subset \mathcal{A}$ is a filtration of that complex. Here is the associated chain filtration $F_{\mu}(\mathcal{A})$:

$$C_*(\mathcal{A}_{\mathcal{U}}) \xrightarrow{i_{\mu}} C_*(\mathcal{A}_{\mathcal{V}}) \xrightarrow{j_{\mu}} C_*(\mathcal{A}_{\mathcal{X}})$$

Let $\psi_{\mathcal{Y}} : C_*(\mathcal{Y}) \rightarrow C_*(\mathcal{A}_{\mathcal{Y}})$ and $\phi_{\mathcal{Y}} : C_*(\mathcal{A}_{\mathcal{Y}}) \rightarrow C_*(\mathcal{A}_{\mathcal{Y}})$ be the chain equivalences from the proof of Theorem 4.2.2 induced by $\mu_{\mathcal{Y}}$. Define Ψ_{μ} to be the sequence of maps $\psi_{\mathcal{Y}}$ for $\mathcal{Y} = \mathcal{U}, \mathcal{V}, \mathcal{X}$ and similarly define Φ_{μ} . By Theorem 4.3.5 and Remark 4.3.7, $\Psi_{\mu} : F(\mathcal{X}) \rightarrow F_{\mu}(\mathcal{A})$ and $\Phi_{\mu} : F_{\mu}(\mathcal{A}) \rightarrow F(\mathcal{X})$ are strong equivalences of chain filtrations.

THEOREM 6.3.1. *Let S be the short exact sequence (6.1) of a triple of complexes $\mathcal{U} \subset \mathcal{V} \subset \mathcal{X}$ and let $\mu_{\mathcal{Y}}$ for $\mathcal{Y} \in \{\mathcal{U}, \mathcal{V}, \mathcal{X}\}$ be subordinate matchings as defined in (6.2). Then, the following sequence S^{μ} is short exact*

$$0 \rightarrow C_*(\mathcal{A}_{\mathcal{V}}, \mathcal{A}_{\mathcal{U}}) \xrightarrow{\alpha'} C_*(\mathcal{A}_{\mathcal{X}}, \mathcal{A}_{\mathcal{U}}) \xrightarrow{\beta'} C_*(\mathcal{A}_{\mathcal{X}}, \mathcal{A}_{\mathcal{V}}) \rightarrow 0$$

Here α' and β' are induced by inclusion and projection of relative chains respectively. Moreover, S^{μ} is weakly equivalent to S .

PROOF. Since $\mathcal{A}_{\mathcal{U}} \subset \mathcal{A}_{\mathcal{V}} \subset \mathcal{A}_{\mathcal{X}} = \mathcal{A}$ is a filtration by Theorem 4.3.5, it is immediately clear that S^{μ} is a short exact sequence and it suffices to construct a weak equivalence $\Omega : S \rightarrow S^{\mu}$. It follows from Remarks 2.2.2 and 4.3.7 that for $\mathcal{Y}, \mathcal{Z} \in \{\mathcal{U}, \mathcal{V}, \mathcal{X}\}$ such that $\mathcal{Y} \subset \mathcal{Z}$, the chain equivalences $\psi_{\mathcal{Z}}$ and $\phi_{\mathcal{Z}}$ induce relative chain equivalences

$$\psi_{\mathcal{Z}, \mathcal{Y}} : C(\mathcal{Z}, \mathcal{Y}) \rightarrow C(\mathcal{A}_{\mathcal{Z}}, \mathcal{A}_{\mathcal{Y}}) \text{ and } \phi_{\mathcal{Z}, \mathcal{Y}} : C(\mathcal{A}_{\mathcal{Z}}, \mathcal{A}_{\mathcal{Y}}) \rightarrow C(\mathcal{Z}, \mathcal{Y})$$

Now, we have the following diagram between S and S^{μ} :

$$\begin{array}{ccccccccc} 0 & \longrightarrow & C(\mathcal{V}, \mathcal{U}) & \xrightarrow{\alpha} & C(\mathcal{X}, \mathcal{U}) & \xrightarrow{\beta} & C(\mathcal{X}, \mathcal{V}) & \longrightarrow & 0 \\ \downarrow & & \downarrow \psi_{\mathcal{V}, \mathcal{U}} & & \downarrow \psi_{\mathcal{X}, \mathcal{U}} & & \downarrow \psi_{\mathcal{X}, \mathcal{V}} & & \downarrow \\ 0 & \longrightarrow & C(\mathcal{A}_{\mathcal{V}}, \mathcal{A}_{\mathcal{U}}) & \xrightarrow{\alpha'} & C(\mathcal{A}_{\mathcal{X}}, \mathcal{A}_{\mathcal{U}}) & \xrightarrow{\beta'} & C(\mathcal{A}_{\mathcal{X}}, \mathcal{A}_{\mathcal{V}}) & \longrightarrow & 0 \end{array}$$

Routine computations using the fact that $\Psi_\mu : F(\mathcal{X}) \rightarrow F_\mu(\mathcal{A})$ is a morphism of chain filtrations confirm that this diagram commutes. For instance, consider the first square:

$$\begin{aligned} \psi_{\mathcal{X},\mathcal{U}} \circ \alpha(\mathbf{v} + C(\mathcal{U})) &= \psi_{\mathcal{X},\mathcal{U}}(j(\mathbf{v}) + C(\mathcal{U})) \\ &= \psi_{\mathcal{X}} \circ j(\mathbf{v}) + C(\mathcal{A}_{\mathcal{U}}) \\ &= j_\mu \circ \psi_{\mathcal{V}}(\mathbf{v}) + C(\mathcal{A}_{\mathcal{U}}) \\ &= \alpha'(\psi_{\mathcal{V}}(\mathbf{v}) + C(\mathcal{A}_{\mathcal{U}})) \\ &= \alpha' \circ \psi_{\mathcal{V},\mathcal{U}}(\mathbf{v} + C(\mathcal{U})) \end{aligned}$$

Let $\Omega : S \rightarrow S^\mu$ be the morphism consisting of the maps $\psi_{z,y}$. Note that a morphism $\Upsilon : S^\mu \rightarrow S$ may be constructed analogously via the maps $\phi_{z,y}$. Clearly, Υ is a weak inverse of Ω as desired. \square

6.4. The Mayer-Vietoris Sequence

Let (\mathcal{X}, κ) be a cell complex and consider subcomplexes $\mathcal{U}, \mathcal{V} \subset \mathcal{X}$ such that $\mathcal{X} = \mathcal{U} \cup \mathcal{V}$. Note that the intersection $\mathcal{Y} = \mathcal{U} \cap \mathcal{V}$ is also a (possibly empty) subcomplex of \mathcal{X} . Consider the following diamond of inclusion maps of chains

$$\begin{array}{ccc} & C(\mathcal{X}) & \\ i \nearrow & & \nwarrow j \\ C(\mathcal{U}) & & C(\mathcal{V}) \\ k \nwarrow & & \nearrow l \\ & C(\mathcal{Y}) & \end{array}$$

Define the chain maps $\alpha : C(\mathcal{Y}) \rightarrow C(\mathcal{U}) \oplus C(\mathcal{V})$ and $\beta : C(\mathcal{U}) \oplus C(\mathcal{V}) \rightarrow C(\mathcal{X})$ as follows.

$$\alpha(\mathbf{y}) = (k(\mathbf{y}), l(\mathbf{y})) \text{ and } \beta(\mathbf{u}, \mathbf{v}) = i(\mathbf{u}) - j(\mathbf{v}).$$

Now, the following sequence S of chain complexes is short exact:

$$0 \rightarrow C(\mathcal{Y}) \xrightarrow{\alpha} C(\mathcal{U}) \oplus C(\mathcal{V}) \xrightarrow{\beta} C(\mathcal{X}) \rightarrow 0 \quad (6.3)$$

and the corresponding long exact sequence $\mathcal{L}(S)$ is called the *Mayer Vietoris sequence* associated to \mathcal{U} and \mathcal{V} .

There are two filtrations of interest in the inclusion diagram above: $\mathcal{Y} \subset \mathcal{U} \subset \mathcal{X}$ and $\mathcal{Y} \subset \mathcal{V} \subset \mathcal{X}$. Assume the existence of acyclic matchings $\mu_z = (\mathcal{A}_z; w_z : \mathcal{Q}_z \rightarrow \mathcal{K}_z)$ for $z \in \{\mathcal{Y}, \mathcal{U}, \mathcal{V}\}$ such that the compatibility requirement $\mu_{\mathcal{U}} \supset \mu_{\mathcal{Y}} \subset \mu_{\mathcal{V}}$ is satisfied. We call $\mu_{\mathcal{Y}}$ the *intersection* of $\mu_{\mathcal{U}}$ and $\mu_{\mathcal{V}}$ in this case.

DEFINITION 6.4.1. Given acyclic matchings $\mu_{\mathcal{U}} \supset \mu_{\mathcal{Y}} \subset \mu_{\mathcal{V}}$ as above, set $\mathcal{A}_{\mathcal{X}} = \mathcal{A}_{\mathcal{U}} \cup \mathcal{A}_{\mathcal{V}}$, $\mathcal{K}_{\mathcal{X}} = \mathcal{K}_{\mathcal{U}} \cup \mathcal{K}_{\mathcal{V}}$ and $\mathcal{Q}_{\mathcal{X}} = \mathcal{Q}_{\mathcal{U}} \cup \mathcal{Q}_{\mathcal{V}}$. For each $q \in \mathcal{Q}$, define

$$w_{\mathcal{X}}(q) = \begin{cases} w_{\mathcal{U}}(q), & \text{if } q \in \mathcal{Q}_{\mathcal{U}} \\ w_{\mathcal{V}}(q), & \text{if } q \in \mathcal{Q}_{\mathcal{V}} \end{cases}$$

noting that if $q \in \mathcal{Q}_{\mathcal{U}} \cap \mathcal{Q}_{\mathcal{V}} = \mathcal{Q}_{\mathcal{Y}}$ then both expressions equal $w_{\mathcal{Y}}(q)$.

It is straightforward to check that $\mu_{\mathcal{X}} = (\mathcal{A}_{\mathcal{X}}; w_{\mathcal{X}} : \mathcal{Q}_{\mathcal{X}} \rightarrow \mathcal{K}_{\mathcal{X}})$ is a partial matching on \mathcal{X} . To see that $\mu_{\mathcal{X}}$ is indeed acyclic, note that any path ρ of $\mu_{\mathcal{X}}$ must start with an element

$q \in \mathcal{Q}_X = \mathcal{Q}_U \cup \mathcal{Q}_V$. If $q \in \mathcal{Q}_U$ then ρ is also a path of μ_U and must consequently be acyclic. Similarly, if $q \in \mathcal{Q}_V$ then ρ is also a path of μ_V .

By this construction, we obtain the compatibility $\mu_U \subset \mu_X \supset \mu_V$. We call μ_X the *union* of μ_U and μ_V , and emphasize that it is only defined when the *intersection* μ_Y exists.

REMARK 6.4.2. Given subcomplexes U, V of X with $Y = U \cap V$, it is possible to use the Algorithm MorseReduce from Chapter 5 to construct acyclic matchings μ_Z for $Z \in \{Y, U, V, X\}$ so that μ_Y and μ_X are the intersection and union of μ_U and μ_V respectively. To see this, note that we may use MorseReduce to construct a matching on X subordinate to the following filtration:

$$Y \subset U \subset X.$$

Set $\mu_Z = (\mathcal{A}_Z; w_Z : \mathcal{Q}_Z \rightarrow \mathcal{K}_Z)$ for $Z \in \{Y, U, X\}$ and note that we have $\mu_Y \subset \mu_U \subset \mu_X$ by Theorem 4.3.5. Finally, define $\mu_V = (\mathcal{A}_V; w_V : \mathcal{Q}_V \rightarrow \mathcal{K}_V)$ as follows: $\mathcal{D}_V = \mathcal{D}_Y \sqcup (\mathcal{D}_X \setminus \mathcal{D}_U)$ for $\mathcal{D} \in \{\mathcal{A}, \mathcal{K}, \mathcal{Q}\}$ and

$$w_V(q) = \begin{cases} w_X(q) & q \in \mathcal{Q}_X \setminus \mathcal{Q}_U \\ w_Y(q) & q \in \mathcal{Q}_Y \end{cases}$$

It is straightforward to check that $\mu_Y \subset \mu_V \subset \mu_X$ as desired.

Letting $\psi_Z : C(Z) \rightarrow C(\mathcal{A}_Z)$ and $\phi_Z : C(\mathcal{A}_Z) \rightarrow C(Z)$ be the chain equivalences associated to μ_Z for $Z \in \{Y, U, V, X\}$, we obtain the following commutative diagram of diamonds by Theorem 4.3.5.

$$\begin{array}{ccccc}
 & & C_*(\mathcal{A}_X) & & \\
 & & \uparrow \psi_X & & \\
 & & C_*(X) & & \\
 & i_\mu \nearrow & & \nwarrow j_\mu & \\
 C_*(\mathcal{A}_U) & \xleftarrow{\psi_U} & C_*(U) & & C_*(V) \xrightarrow{\psi_V} C_*(\mathcal{A}_V) \\
 & \nwarrow k_\mu & \uparrow i & \downarrow j & \\
 & & C_*(Y) & & \\
 & \swarrow k_\mu & \downarrow \psi_Y & \swarrow l_\mu & \\
 & & C_*(\mathcal{A}_Y) & &
 \end{array}$$

where the maps on the outer diamond are induced by inclusions of critical chains. A similar diagram exists for the maps ϕ_Z with $Z \in \{Y, U, V, X\}$.

THEOREM 6.4.3. Let (X, κ) be a cell complex with subcomplexes U and V such that $X = U \cup V$ and set $Y = U \cap V$. Let S be the short exact sequence (6.3). Let $\mu_Z = (\mathcal{A}_Z; w_Z : \mathcal{Q}_Z \rightarrow \mathcal{K}_Z)$ be acyclic matchings on $Z = \{Y, U, V, X\}$ so that μ_Y and μ_X are the intersection and union respectively of μ_U and μ_V . Then, the following sequence S^μ is short exact

$$0 \rightarrow C(\mathcal{A}_Y) \xrightarrow{\alpha'} C(\mathcal{A}_U) \oplus C(\mathcal{A}_V) \xrightarrow{\beta'} C(\mathcal{A}_X)$$

where α' and β' are defined analogously to α and β from S . Moreover, S is chain homotopy equivalent to S^μ .

PROOF. By the existence of the outer diamond in the preceding diagram, it is clear that S^μ is a short exact sequence, so we focus on constructing a weak equivalence $\Omega : S \rightarrow S'$. Consider

the following diagram between S and S^μ :

$$\begin{array}{ccccccccc}
 0 & \longrightarrow & C(\mathcal{Y}) & \xrightarrow{\alpha} & C(\mathcal{U}) \oplus C(\mathcal{V}) & \xrightarrow{\beta} & C(\mathcal{X}) & \longrightarrow & 0 \\
 \downarrow & & \downarrow \psi_{\mathcal{Y}} & & \downarrow (\psi_{\mathcal{U}}, \psi_{\mathcal{V}}) & & \downarrow \psi_{\mathcal{X}} & & \downarrow \\
 0 & \longrightarrow & C(\mathcal{Y}) & \xrightarrow{\alpha'} & C(\mathcal{U}) \oplus C(\mathcal{V}) & \xrightarrow{\beta'} & C(\mathcal{X}) & \longrightarrow & 0
 \end{array}$$

From the diagram between diamonds it can be checked that this diagram of short exact sequences commutes. Consider the third square for example and observe that for $(\mathbf{u}, \mathbf{v}) \in C(\mathcal{U}) \oplus C(\mathcal{V})$ we have

$$\begin{aligned}
 \beta' \circ (\psi_{\mathcal{U}}, \psi_{\mathcal{V}})(\mathbf{u}, \mathbf{v}) &= \beta'(\psi_{\mathcal{U}}(\mathbf{u}), \psi_{\mathcal{V}}(\mathbf{v})) \\
 &= i_{\mu} \circ \psi_{\mathcal{U}}(\mathbf{u}) - j_{\mu} \circ \psi_{\mathcal{V}}(\mathbf{v}) \\
 &= \psi_{\mathcal{X}} \circ i(\mathbf{u}) - \psi_{\mathcal{X}} \circ j(\mathbf{v}) \\
 &= \psi_{\mathcal{X}}(i(\mathbf{u}) - j(\mathbf{v})) \\
 &= \psi_{\mathcal{X}} \circ \beta(\mathbf{u}, \mathbf{v})
 \end{aligned}$$

Thus, a morphism $\Omega : S \rightarrow S^\mu$ may be defined via the sequence of vertical maps in the diagram above. To see that Ω is in fact a weak equivalence, note that we can analogously define a weak inverse $\Upsilon : S^\mu \rightarrow S$ using the maps $\phi_{\mathcal{Z}} : C(\mathcal{A}_{\mathcal{Z}}) \rightarrow C(\mathcal{Z})$. \square

Bibliography

- [1] **Dionysus**, <http://hg.mrzv.org/dionysus>.
- [2] **jPlex**, <http://comptop.stanford.edu/u/programs/jplex>.
- [3] **Perseus**, <http://www.math.rutgers.edu/~vidit/perseus.html>.
- [4] E. Batzies and V. Welker. Discrete morse theory for cellular resolutions. *J. Reine Angew. Math*, 543:147–168, 2000.
- [5] Mladen Bestvina. PL Morse theory. (English). *Math. Commun.*, 13(2):149–162, 2008.
- [6] Raoul Bott. Morse theory indomitable. *Publications Mathématiques de L’IHÉS*, 68(1):99–114, 1988.
- [7] Gunnar Carlsson. Topology and data. *Bull. Amer. Math. Soc. (N.S.)*, 46(2):255–308, 2009.
- [8] Gunnar Carlsson and Vin De Silva. Topological approximation by small simplicial complexes. *Preprint*, 2003.
- [9] Manoj K. Chari. On discrete Morse functions and combinatorial decompositions. *Discrete Math.*, 217(1-3):101–113, 2000. Formal power series and algebraic combinatorics (Vienna, 1997).
- [10] David Cohen-Steiner, Herbert Edelsbrunner, and John Harer. Stability of persistence diagrams. *Discrete Comput. Geom.*, 37(1):103–120, 2007.
- [11] Charles Conley. *Isolated invariant sets and the Morse index*, volume 38 of *CBMS Regional Conference Series in Mathematics*. American Mathematical Society, 1978.
- [12] H. Edelsbrunner and J. Harer. Persistent homology – a survey. In *Twenty Years After*, eds. J. E. Goodman, J. Pach and R. Pollack, *AMS.*, 2007.
- [13] Robin Forman. Morse theory for cell complexes. *Advances in Mathematics*, 134:90–145, 1998.
- [14] Robert Franzosa. Index filtrations and the homology index braid for partially ordered Morse decompositions. *Trans. Amer. Math. Soc.*, 298:193–213, 1986.
- [15] Robert Ghrist. Barcodes: the persistent topology of data. *Bull. Amer. Math. Soc. (N.S.)*, 45(1):61–75, 2008.
- [16] Robert Ghrist. Three examples of applied and computational homology. *Nieuw Arch. Wiskd. (5)*, 9(2):122–125, 2008.
- [17] D. Gunther, J. Reininghaus, H. Wagner, and I. Hotz. Memory efficient computation of persistent homology for 3d image data using discrete Morse theory. *Proc. Conference on Graphics, Patterns and Images*, 24, to appear.
- [18] Shaun Harker, Konstantin Mischaikow, Marian Mrozek, and Vidit Nanda. Discrete Morse theoretic algorithms for computing homology of complexes and maps. *Preprint*, 2008.
- [19] Shaun Harker, Konstantin Mischaikow, Marian Mrozek, Vidit Nanda, Hubert Wagner, Mateusz Juda, and Pawel Dlotko. The efficiency of a homology algorithm based on discrete Morse theory and coreductions. In *Proceedings of the 3rd International Workshop on Computational Topology in Image Context*, volume 1 of *Image A*, pages 41–47. 2010.
- [20] Allen Hatcher. *Algebraic Topology*. Cambridge University Press, 2002.
- [21] Tomasz Kaczynski, Konstantin Mischaikow, and Marian Mrozek. *Computational homology*, volume 157 of *Applied Mathematical Sciences*. Springer-Verlag, New York, 2004.
- [22] Dmitry Kozlov. Discrete Morse theory for free chain complexes. *Comptes Rendus Mathématique*, 340:867–872, 2005.
- [23] Dmitry Kozlov. *Combinatorial Algebraic Topology*, volume 21 of *Algorithms and Computation in Mathematics*. Springer, 2008.
- [24] Solomon Lefschetz. *Algebraic topology*. American Mathematical Society Colloquium Publications, v. 27. American Mathematical Society, New York, 1942.
- [25] Thomas Lewiner, Helio Lopes, and Geovan Tavares. Toward optimality in discrete Morse theory. *Experiment. Math*, 12(3):271–286, 2003.
- [26] J. Milnor. *Morse theory*. Princeton University Press, 1963. Based on lecture notes by M. Spivak and R. Wells.
- [27] Konstantin Mischaikow and Vidit Nanda. Morse theory for filtrations and efficient computation of persistent homology. *Under review*, 2011.
- [28] Konstantin Mischaikow and Vidit Nanda. Reconstructing the induced map on homology from images of random samples. *Under review*, 2012.

- [29] Marian Mrozek and Bogdan Batko. The coreduction homology algorithm. *Discrete and Computational Geometry*, 41:96–118, 2009.
- [30] Marian Mrozek and Thomas Wanner. Coreduction homology algorithm for inclusions and persistent homology. *Comput. Math. Appl.*, 60(10):2812–2833, 2010.
- [31] J.R. Munkres. *Elements of algebraic topology*. The Benjamin/Cummings Publishing Company, Inc., 1984.
- [32] P. Niyogi, S. Smale, and S. Weinberger. Finding the homology of submanifolds with high confidence from random samples. *Discrete and Computational Geometry*, 39:419–441, 2008.
- [33] V. Robins, P.J. Wood, and A.Ā. Sheppard. Discrete Morse theory for grayscale digital images. *IEEE Transactions on Pattern Analysis and Machine Learning*, pages 1–14, 2010.
- [34] Herbert Simon. Designing organizations for an information-rich world. In Martin Greenberger, editor, *Computers, Communications and the Public Interest*, pages 37–72. The Johns Hopkins Press, 1971.
- [35] Emil Sköldberg. Morse theory from an algebraic viewpoint. *Transactions of the AMS*, 358(1):115–129, 2006.
- [36] Edwin H. Spanier. *Algebraic topology*. McGraw-Hill Book Co., New York, 1966.
- [37] A. W. Tucker. Cell spaces. *Ann. of Math. (2)*, 37(1):92–100, 1936.
- [38] Hubert Wagner, Chao Chen, and Erald Vucini. Efficient computation of persistent homology for cubical data. *Proceedings of Workshop on Topology-based Methods in Data Analysis and Visualization*, to appear.
- [39] Charles Weibel. *An introduction to homological algebra*. Cambridge Univ. Press, 1994.
- [40] Afra Zomorodian and Gunnar Carlsson. Computing persistent homology. *Discrete Comput. Geom.*, 33(2):249–274, 2005.