Inauguraldissertation

zur
Erlangung der Doktorwürde
der
Gesamtfakultät für Mathematik, Ingenieur- und Naturwissenschaften
der
Ruprecht–Karls–Universität
Heidelberg

vorgelegt von
Maximilian Schmahl
aus
Mainz

Tag der mündlichen Prüfung:

Topics in Persistent Homology: From Morse Theory for Minimal Surfaces to Efficient Computation of Image Persistence

Betreuer:
Prof. Dr. Peter Albers
Prof. Dr. Ulrich Bauer

Abstract

We study some problems and develop some theory related to persistent homology, separated into two lines of investigation.

In the first part, we introduce lifespan functors, which are endofunctors on the category of persistence modules that filter out intervals from barcodes according to their boundedness properties. They can be used to classify injective and projective objects in the category of barcodes and the category of pointwise finite-dimensional persistence modules. They also naturally appear in duality results for absolute and relative versions of persistent (co)homology, generalizing previous results in terms of barcodes by de Silva, Morozov, and Vejdemo-Johansson. Due to their functoriality, we can apply these results to morphisms in persistent homology that are induced by morphisms between filtrations. This lays the groundwork for an efficient algorithm to compute barcodes of images and induced matchings of such morphisms, which performs computations in terms of relative cohomology and then translates to absolute homology via the aforementioned dualities. Our method is based on a previous algorithm by Cohen-Steiner, Edelsbrunner, Harer, and Morozov that did not make use of relative cohomology. Using it is crucial, however, because our algorithm applies the clearing optimization introduced by Chen and Kerber, which works particularly well in the context of relative cohomology. We provide an implementation of our algorithm for inclusions of filtrations of Vietoris–Rips complexes in the framework of the software Ripser by Ulrich Bauer.

In the second part, we introduce local connectedness conditions on a broad class of functionals that ensure that the persistent homology of their associated sublevel set filtration is q-tame, which, in particular, implies that they satisfy generalized Morse inequalities. We illustrate the applicability of these results by recasting the original proof of the unstable minimal surface theorem given by Morse and Tompkins in terms of persistent Čech homology in a modern and rigorous framework. Moreover, we show that the interleaving distance between the persistent singular homology and the persistent Čech homology of a filtration consisting of paracompact Hausdorff spaces is 0 if it satisfies a similar local connectedness condition to the one used to ensure q-tameness, generalizing a result by Mardešić for locally connected spaces to the setting of filtrations. In contrast to singular homology, the persistent Čech homology of a compact filtration is always upper semi-continuous, which has structural implications in the q-tame case: using a result by Chazal, Crawley-Boevey, and de Silva concerning radicals of persistence modules, we show that every lower semi-continuous q-tame persistence module can be decomposed as a direct sum of interval modules and that every upper semi-continuous q-tame persistence module can be decomposed as a product of interval modules.

Zusammenfassung

Wir betrachten einige Probleme und entwickeln etwas Theorie zu persistenter Homologie, unterteilt in zwei Forschungslinien.

Im ersten Teil führen wir Lifespan-Funktoren ein; also Endofunktoren auf der Kategorie der Persistenzmoduln, die Intervalle entsprechend ihrer Beschränktheitseigenschaften aus Barcodes herausfiltern. Mit diesen Funktoren können injektive und projektive Objekte in der Kategorie der Barcodes und der Kategorie der punktweise endlich dimensionalen Persistenzmoduln klassifiziert werden. Sie erscheinen ebenfalls auf natürliche Weise in Dualitätsresultaten für absolute und relative Versionen von persistenter (Ko)Homologie, welche Resultate bezüglich Barcodes von de Silva, Morozov und Vejdemo-Johansson verallgemeinern. Aufgrund ihrer Funktorialität können diese Resultate auf Morphismen in persistenter Homologie angewendet werden, die von Morphismen zwischen Filtrationen induziert werden. Dies legt den Grundstein für einen effizienten Algorithmus zur Berechnung von Barcodes von Bildern und induzierten partiellen Bijektionen solcher Morphismen, welcher Berechnungen in relativer Kohomologie ausführt und dann mittels der zuvor beschriebenen Dualität in absolute Homologie übersetzt. Diese Methode basiert auf einem vorangegangen Algorithmus von Cohen-Steiner, Edelsbrunner, Harer und Morozov, der relative Kohomologie nicht verwendet. Relative Kohomologie zu verwenden ist aber zentral, weil unser Algorithmus die sogenannte clearing Optimierung von Chen und Kerber benutzt, die im Zusammenspiel mit relativer Kohomologie besonders wirksam ist. Wir stellen eine Implementierung unseres Algorithmus für den Spezialfall von Vietoris–Rips Komplexen basierend auf der Software Ripser von Ulrich Bauer zur Verfügung.

Im zweiten Teil führen wir lokale Zusammenhangsbedingungen für die Subniveaumengenfiltrationen einer weiten Klasse von Funktionalen ein, die sicherstellen, dass die zugehörige persistente Homologie q-zahm ist, was insbesondere impliziert, dass diese Funktionale verallgemeinerte Morse-Ungleichungen erfüllen. Wir illustrieren die Anwendbarkeit dieser Resultate, indem wir den ursprünglichen Beweis des Satzes zu instabilen Minimalflächen von Morse und Tompkins mittels Čech Homologie auf moderne und präzise Weise aufbereiten. Darüber hinaus zeigen wir, dass die Interleaving-Distanz zwischen persistenter singulärer und persistenter Čech Homologie einer Filtration 0 ist, falls sie eine der vorherigen ähnliche lokale Zusammenhangsbedingung erfüllt, was ein Resultat von Mardešić für lokal zusammenhängende Räume auf Filtrationen verallgemeinert. Im Gegensatz zu singulärer Homologie hat die persistente Čech Homologie einer kompakten Filtration immer die Eigenschaft, oberhalbstetig zu sein, was im q-zahmen Fall strukturelle Auswirkungen hat: Mithilfe eines Resultats von Chazal, Crawley-Boevey und de Silva zu Radikalen von Persistenzmoduln zeigen wir, dass alle unterhalbstetigen q-zahmen Persistenzmoduln als direkte Summe von Intervallmoduln zerlegt werden können und alle oberhalbstetigen q-zahmen Persistenzmoduln als Produkt von Intervallmoduln zerlegt werden können.

Acknowledgments

First and foremost, big thanks to Peter Albers for being a great supervisor and for giving me the freedom to work on the problems I find most interesting, even if they are not symplectic at all. A big thank you goes to Peter and also Anna Wienhard for giving me the opportunity to be part of a vibrant working group and research environment with all the academic opportunities and support one could hope for. I also want to thank the German Research Foundation (DFG) for supporting my research through the Cluster of Excellence EXC-2181/1 *STRUCTURES*, and the Research Training Group RTG 2229 *Asymptotic Invariants and Limits of Groups and Spaces*.

I'm also very thankful for having had fantastic collaborators in Ulrich Bauer and Anibal Medina-Mardones who worked with me on various projects that became part of this thesis. My deepest gratitude especially goes to Uli for not only supervising me during my previous studies, but for also staying in touch to co-supervise me after I left Munich and for being my greatest source of mathematical inspiration and education. Apart from Uli and Anibal, this thesis benefitted a lot from the devoted proof-reading of different parts by Johanna, Lucas, and various anonymous referees, which I want to thank all of them for.

Thanks to everyone who played a part in creating a very pleasant atmosphere at the math department in Heidelberg, with special mentions going out to Alexandre, Arnaud, Kevin, and Valerio, for being great bouldering companions, workout partners, office mates, and overall friends. An important part of my experience studying and working in Heidelberg was the EP Math and Data team comprising of Daniel, Lukas, and Michael, as well as all the regular participants of our various seminars and workshops who gave me an outlet to talk about TDA in Heidelberg.

Thanks to all my friends from TUM, and especially Vincent, without whom I would have never gotten this far in my mathematical journey. Last but not least, thank you to all my friends and family, and in particular Chantal, for supporting me in all things not math.

Contents

Part I.

Introduction

1. A brief introduction to persistent homology

Before presenting the original contents of this thesis, we briefly recall some terminology from the existing literature that we will use freely throughout the rest of this work and that is necessary to state our results precisely. In Section 1.1 we remind the reader of persistent homology, persistence modules and their barcodes, as well as the stability theorem. Relevant to computational applications, we review the basic algorithm for determining barcodes of filtrations of simplicial complexes as well as the so-called clearing optimization in Section 1.2. We then present the framework of persistence diagrams as an alternative to barcodes and how they can be defined for q-tame persistence modules via the observable category in Section 1.3. To finish this primer, we discuss the less commonly used Čech homology theory in Section 1.4, which will be important for our considerations regarding Morse theory and minimal surfaces. A reader who is already familiar with these concepts and wants to skip this chapter should take away that for the rest of this thesis we fix a field $\mathbb{F}$, which, unless stated otherwise, will be used as the coefficient group for all homology groups and chain complexes and as the base field for all vector spaces. We also notationally fix an index set T for all persistence modules, which is always assumed to be totally ordered and sometimes specialized to the cases where T is finite or $T = \mathbb{R}$.

1.1. Persistence modules and barcodes

The main motivation for the development of persistence theory is the study of persistent homology, which is also the main concept of study in this thesis. We start by reviewing what persistent homology is and where it originates in § 1.1.1. In § 1.1.2, we then consider the structure theory of the underlying algebraic objects called persistence modules. In particular, we remind the reader of the basic concept of barcodes. We then review the standard notions of distance for persistence modules and barcodes, namely the interleaving distance in § 1.1.3, and the bottleneck distance in § 1.1.4. For completeness of the exposition, we also recall the stability theorem in § 1.1.5, which asserts that passing to barcodes is a 1-Lipschitz map.

1.1.1. Persistent homology and persistence modules

Persistent homology, the homology of a filtration of simplicial complexes, is a cornerstone in the foundations of topological data analysis. It has found numerous applications in a variety of disciplines, including for example computer vision, neuroscience, materials science, and evolutionary biology [Ble+21; Clo+20; Dab+12; Hu+19; Nak+15]. The most common setting studied in the topological data analysis literature is as follows: Given a finite metric space, to be thought of as a data set, one uses for example Rips, Čech, or

Delaunay complexes to construct a filtration of simplicial complexes

$$K_\bullet \colon \emptyset = K_0 \subseteq K_1 \subseteq \cdots \subseteq K_N = K,$$

where increasing the filtration parameter corresponds in some sense to increasing the distance at which points in the data set are allowed to be connected by simplices. By applying homology with coefficients in a field to each space and to each inclusion map one obtains a diagram of vector spaces

$$H_*(K_\bullet) \colon H_*(K_0) \to H_*(K_1) \to \cdots \to H_*(K).$$

Such a diagram is called a *persistence module*, and it decomposes into a direct sum of indecomposable diagrams, each supported on an interval [Bar94; Cra15; ZC05]. The collection of these intervals, called the *(persistence) barcode*, has proven to be a powerful invariant of the filtration. The passage from data sets to barcodes is often referred to as the *pipeline* of persistent homology and we refer to [Cha+16; EH10; Oud15; Pol+20] for a more detailed overview over each of its steps.

Category of persistence modules For our purposes, *the category of persistence modules* $\mathbf{Pers} = \mathbf{Vec}^{\mathbf{T}}$ is defined as the category of functors $\mathbf{T} \to \mathbf{Vec}$, where $\mathbf{T}$ is a totally ordered set $(T, \leq)$ considered as a poset category, and $\mathbf{Vec}$ is the category of vector spaces over some field $\mathbb{F}$, which we fix for the rest of this thesis. Spelled out, this means that a persistence module M consists of a vector space M_t for any $t \in T$, and for any pair $s, t \in T$ with $s \leq t$ a linear map $M_{s,t} \colon M_s \to M_t$ such that $M_{t,t}$ is the identity and

$$M_{s,t} \circ M_{r,s} = M_{r,t}$$

for any triple $r \leq s \leq t$. The maps $M_{s,t}$ are called *structure maps.* A morphism of persistence modules $\varphi \colon M \to N$ is a natural transformation, i.e., a collection of linear maps $\varphi_t \colon M_t \to N_t$ for all $t \in T$ such that

$$\varphi_t \circ M_{s,t} = N_{s,t} \circ \varphi_s$$

for all $s, t \in T$ with $s \leq t$. The category of persistence modules is an abelian category with kernels, cokernels, direct sums etc. given by their pointwise analogues.

Persistent homology The notion of persistence module above includes the persistent homology of filtered complexes coming from topological data analysis as described earlier. In this case, the index set T is a finite set. However, instead of starting with a data set or filtration of complexes, one may for example also start the pipeline with a not necessarily continuous function $f \colon X \to \mathbb{R}$ on some topological space X and consider its *sublevel set filtration*, which is defined by

$$f_{\leq t} = f^{-1}(-\infty, t].$$

Applying any vector space valued homology theory to each of the sublevel sets and the inclusion maps then yields a (graded) persistence module, this time indexed by $T = \mathbb{R}$. This persistence module is called the *persistent homology* of the sublevel set filtration, or sometimes just the persistent homology of the function. Slightly abusing terminology, we will sometimes say that the function f has a certain property if the persistent homology of its sublevel set filtration has that property. More generally, we call any diagram $\mathbf{T} \to \mathbf{Top}$ a *filtration* if its structure maps are all injective. Composing a diagram of topological spaces $\mathbf{T} \to \mathbf{Top}$ with a functor $\mathbf{Top} \to \mathbf{Vec}$ yields a persistence module, which, in case this functor is a homology, is called the *persistent homology* of the diagram. In the case where the functor $\mathbf{Top} \to \mathbf{Vec}$ applied to a filtration is contravariant, for example if one uses *co*homology instead of homology, the resulting persistence module will be indexed by the opposite category $\mathbf{T}^{\mathrm{op}}$, which is the poset category corresponding to $(T, \geq)$, i.e., T with the opposite of the original order. Purely algebraicly, one may also consider *persistent chain complexes*, i.e., diagrams of chain complexes $\mathbf{T} \to \mathbf{Ch}(\mathbf{Vec})$, and take their homology to obtain persistent homology. Note that persistent chain complexes may equivalently be thought of as chain complexes in the abelian category of persistence modules.

1.1.2. Barcodes and the structure of persistence modules

One of the most important questions in the theory of persistence is when a given persistence module can be decomposed into elementary building blocks: Recall that $I \subseteq T$ is called an *interval* if $I \neq \emptyset$ and if for all $s, u \in I$ and $t \in T$ with $s \leq t \leq u$ we have $t \in I$. For $I \subseteq T$ an interval, define a persistence module $C(I)$ via

$$C(I)_t = \begin{cases} \mathbb{F} & \text{if } t \in I, \\ 0 & \text{otherwise,} \end{cases}$$

with structure maps

$$C(I)_{s,t} = \begin{cases} \mathrm{id}_{\mathbb{F}} & \text{if } s, t \in I, \\ 0 & \text{otherwise.} \end{cases}$$

Such persistence modules are called *interval modules*. We say that a persistence module M has a *barcode* or is *interval-decomposable* if there exists an index set A and a collection of intervals $(I_\alpha)_{\alpha \in A}$ such that

$$M \cong \bigoplus_{\alpha \in A} C(I_\alpha).$$

A collection of intervals as above is called a *barcode* and a choice of such an isomorphism is called a *barcode decomposition* or *interval decomposition*. Barcode decompositions need not be unique, but by a version of the Krull–Remak–Schmidt–Azumaya theorem [Azu50] (see also [Cha+16, Theorem 2.7] and [Par70, Section 4.8]) the barcode itself is unique up to a choice of the index set if it exists. This justifies sometimes talking about *the* barcode of a persistence module M, often denoted by $B(M)$.

The most important existence result for barcodes is Crawley-Boevey's theorem [BC20; Cra15]. It states that every *pointwise finite-dimensional* (or PFD) persistence module, meaning every persistence module M for which M_t is a finite-dimensional vector space for all t, has a barcode. The full subcategory $\mathbf{vec}$ in $\mathbf{Vec}$ of finite-dimensional vector spaces is closed under taking kernels, cokernels, and finite direct sums, so the full subcategory of PFD persistence modules $\mathbf{pers} = \mathbf{vec}^{\mathbf{T}}$ is an abelian subcategory of $\mathbf{Pers} = \mathbf{Vec}^{\mathbf{T}}$.

Note that one can define a notion of morphism for barcodes called *overlap matching* such that barcodes form a category and the assignment $(I_\alpha)_{\alpha \in A} \mapsto \bigoplus_{\alpha \in A} C(I_\alpha)$ is a functor. An equivalent description of this category can be given in terms of so-called *matching diagrams* (see § 3.1.1).

1.1.3. Interleaving distance

We briefly review the definition of the interleaving distance for persistence modules indexed by $T = \mathbb{R}$ as introduced by Chazal et al. [Cha+09] and described in great detail and generality in [Les15]. Let M and N be two persistence modules. For $\delta \geq 0$, a *δ-interleaving* between M and N is given by two morphisms $M \to N(\delta)$ and $N \to M(\delta)$, where $M(\delta)$ is the δ-shift of M given by $M(\delta)_t = M_{t+\delta}$, such that the diagram

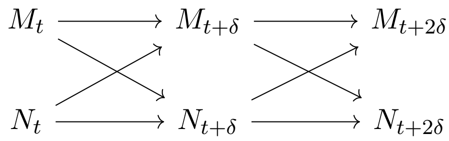

commutes for all $t \in \mathbb{R}$. Clearly, a 0-interleaving describes an isomorphism of persistence modules. A δ-interleaving for $\delta > 0$ may hence be thought of as an approximate isomorphism between persistence modules. The *interleaving distance* between M and N is then defined as

$$d_I(M, N) = \inf\{\delta \geq 0 \mid \text{there is a } \delta\text{-interleaving between } M \text{ and } N\}.$$

Note that this definition does not just make sense for the category of $\mathbb{R}$-indexed persistence modules, i.e., functors $\mathbf{R} \to \mathbf{Vec}$, but for any category of functors $\mathbf{R} \to \mathbf{C}$, where $\mathbf{C}$ is any category. The interleaving distance is easily verified to be an extended pseudo-metric on isomorphism classes of persistence modules. This means that it satisfies all the usual properties of a metric, except for the facts that two non-isomorphic persistence modules may be at interleaving distance 0 to each other and that two persistence modules may have infinite interleaving distance to each other. For example, for any pair of numbers $a, b \in \mathbb{R}$ with $a < b$, the interval modules $C((a, b))$, $C([a, b))$, $C((a, b])$, and $C([a, b])$ are all at distance 0 to each other, but none of them are isomorphic to each other. Moreover, the persistence modules $C([a, \infty))$ and $C([a, \infty)) \oplus C([a, \infty))$ have infinite interleaving distance to each other. In fact, one can easily check that the interleaving distance between two persistence modules with barcodes can only be finite if the number of unbounded intervals in their barcodes is the same.

1.1.4. Bottleneck distance

Corresponding to the interleaving distance for persistence modules, we review the bottleneck distance for barcodes consisting of intervals in $\mathbb{R}$ as introduced by Cohen-Steiner, Edelsbrunner, and Harer [CEH07]. Let $B = (I_\alpha)_{\alpha \in A}$ and $B' = (I'_{\alpha'})_{\alpha' \in A'}$ be barcodes consisting of intervals in $\mathbb{R}$ with index sets A and A', respectively. A *δ-matching* between B and B' is given by subsets $X \subseteq A$, $X' \subseteq A'$ and a bijection $f \colon X \to X'$ such that

$$|\sup I - \inf I| < 2\delta$$

whenever there is $\alpha \in A \setminus X$ with $I = I_\alpha$ or $\alpha' \in A' \setminus X'$ with $I = I'_{\alpha'}$, and

$$|\sup I'_{f(\alpha)} - \sup I_\alpha| < \delta,$$
$$|\inf I'_{f(\alpha)} - \inf I_\alpha| < \delta$$

for all $\alpha \in X$. The *bottleneck distance* between B and B' is then defined as

$$d_I(B, B') = \inf\{\delta \geq 0 \mid \text{there is a } \delta\text{-matching between } B \text{ and } B'\}.$$

The bottleneck distance defines an extended pseudo-metric on barcodes. As an example of different barcodes that have 0 bottleneck distance to each other, one may consider the single interval barcodes given by (a, b), $[a, b)$, $(a, b]$, and $[a, b]$ for $a < b \in \mathbb{R}$. The barcodes given by a single copy of $[a, \infty)$ and two copies of $[a, \infty)$ have infinite bottleneck distance to each other.

1.1.5. Stability theorem

We are now ready to review the *stability theorem* [BL15; Cha+09; Cha+16; CEH07], which has received widespread attention in the topological data analysis community and which provides one of the major justifications for the potential usefulness of persistent homology in applications. Roughly, the stability theorem asserts that passing from sufficiently well-behaved real-valued functions to the barcodes of the persistent homology of their sublevel sets is a 1-Lipschitz map with respect to the supremum norm and the bottleneck distance. In more detail, mapping from functions to barcodes can be seen as the composition of several 1-Lipschitz maps that make up the persistence pipeline: from functions (with the supremum norm) to filtrations by sublevel sets (with the interleaving distance) to persistence modules (with the interleaving distance) and finally to barcodes (with the bottleneck distance).

From functions to filtrations For the first map in the pipeline, recall that the definition of the interleaving distance does not only make sense for the category of persistence modules, but actually for any category of diagrams $\mathbf{R} \to \mathbf{C}$, where $\mathbf{C}$ is any category. In particular, we can define the interleaving distance for $\mathbb{R}$-indexed diagrams of topological spaces, so we can use it to compare sublevel set filtrations of real-valued functions. If X is some topological space and $f, g \colon X \to \mathbb{R}$ are functions, then, setting $e = \|f - g\|_\infty$, it is clear that for all $t \in \mathbb{R}$ we have sublevel set inclusions $f_{\leq t} \subseteq g_{\leq t+e}$ and $g_{\leq t} \subseteq f_{\leq t+e}$. The inclusion maps form an e-interleaving, showing that

$$f \mapsto f_{\leq \bullet}$$

is 1-Lipschitz with respect to the supremum norm and the interleaving distance for real-valued functions on a given topological space.

From filtrations to persistence modules As the next step in the pipeline, we note that for categories of diagrams $\mathbf{R} \to \mathbf{C}$ and $\mathbf{R} \to \mathbf{C}'$ and any functor $F \colon \mathbf{C} \to \mathbf{C}'$ it is clear that composing a δ-interleaving with F again yields a δ-interleaving. In particular, this means that the interleaving distance does not increase when passing from filtrations to persistence modules by composing with homology, so

$$X_\bullet \mapsto H(X_\bullet)$$

is 1-Lipschitz with respect to the interleaving distance for diagrams of topological spaces $X_\bullet$ and any choice of homology theory H.

From persistence modules to barcodes By far the most difficult step in proving the stability theorem is showing that passing from persistence modules that admit barcodes to their barcodes is 1-Lipschitz with respect to the interleaving and the bottleneck distance. This result is known as the *algebraic stability theorem* [BL15; Cha+09; Cha+16]. In fact, one can show that, at least for PFD persistence modules, passing from persistence modules to their barcodes

$$M \mapsto B(M)$$

is an isometry with respect to the interleaving and the bottleneck distance. Reviewing the proof of algebraic stability would go beyond the scope of this work. We mention that there are currently essentially two approaches, one via the non-constructive *interpolation lemma* [Cha+09; CEH07], and one by actually constructing matchings from interleavings via *induced matchings* [BL15], which will also play a motivating role for us later on.

Stability theorems Going through all of the previous steps, we obtain in total that

$$f \mapsto B(H(f_{\leq \bullet}))$$

is 1-Lipschitz for real-valued functions f and homology theories H for which $H_*(f_{\leq \bullet})$ admits a barcode. For completeness, we also mention that there are more computationally flavored versions of the stability theorem, for example stating that passing from finite metric spaces with the Gromov-Hausdorff distance to the barcodes of the persistent homology of their Rips- or Čech-filtrations with the bottleneck distance is 1-Lipschitz [CdSO14]. One may also define other metrics for barcodes than the bottleneck distance and there are corresponding stability results for these metrics [Coh+10].

1.2. Barcode computation

In order to apply persistent homology to data analysis problems, one needs to be able to compute barcodes for filtrations of simplicial complexes, and we now review the basic way to do so. We recall some terminology and the matrix reduction algorithm in § 1.2.1. Then, we explain how to use this algorithm on a matrix representing the simplicial boundary operator to obtain barcodes in § 1.2.2. In § 1.2.3, we then briefly describe the clearing optimization as a first speed-up for barcode computations.

1.2.1. Matrix reduction

We recall the basic matrix reduction algorithm that underlies the computation for barcodes of filtrations of simplicial complexes. If X is a matrix, we write x_i for the i-th column of X. For a non-zero column x_i, we define $\operatorname{pivot} x_i$ as the largest index where x_i has a non-zero entry and we write $\operatorname{pentry} x_i$ for this entry. We write $\operatorname{pivots} X$ for the set of all indices which occur as pivots of non-zero columns of X. A matrix is called *reduced* if no two non-zero columns have the same pivot. A *reduction* of a matrix X is a pair of matrices R and V such that R is reduced, V is full-rank and upper-triangluar, and we have $R = XV$. The following algorithm takes a matrix X with entries in a field and produces a reduction using left-to-right column additions. We write I_n for the identity matrix with $n \times n$ entries.

Input: Matrix X with entries in a field and n rows
Output: A reduction $R = XV$
  $R \leftarrow X$
  $V \leftarrow I_n$
  while $\exists i < j$ with $r_i \neq 0$ and $\operatorname{pivot} r_i = \operatorname{pivot} r_j$ **do**

$r_j \leftarrow r_j - \frac{\operatorname{pentry} r_j}{\operatorname{pentry} r_i} r_i$
$v_j \leftarrow v_j - \frac{\operatorname{pentry} r_j}{\operatorname{pentry} r_i} v_i$
end while
return R, V

1.2.2. Barcode computation via matrix reduction

Using matrix reduction as presented in § 1.2.1, we now describe how to compute barcodes for filtrations of simplicial complexes. A first method for this computation, not using the language of barcodes, yet, was given by Barannikov [Bar94]. In the context of topological data analysis, a first algorithm was given by Zomorodian and Carlsson [ZC05]. The formulation we present goes back to Cohen-Steiner, Edelsbrunner, and Morozov [CEM06].

Let

$$K_\bullet \colon \emptyset = K_{-\infty} = K_0 \subseteq K_1 \subseteq \cdots \subseteq K_N = K_\infty = K$$

be a filtration of finite simplicial complexes with finite index set. For simplicity of the notation, we assume that we are given a function $k \colon K \to \{0, \dots, N\}$ such that its sublevel set filtration is $K_\bullet$, i.e., such that $K_j = k^{-1}(\{0, \dots, j\})$. Let $\sigma_1, \dots, \sigma_n$ be the simplices of the full simplex K in *filtration order*, i.e., ordered such that $i \leq i'$ implies $k(\sigma_i) \leq k(\sigma_{i'})$. The simplices form an ordered basis for the simplicial chain complex $C_*(K)$. Hence, we can consider the corresponding *filtration boundary matrix* D, which is the matrix encoding the boundary operator $\partial \colon C_*(K) \to C_*(K)$ with respect to this ordered basis. Now, assume we have computed a reduction $R = DV$ of D. The row and column indices of these matrices correspond one-to-one to the simplices of K in filtration order. As such, the pivot of a column x corresponds to a simplex of K, denoted by $\operatorname{psimp} x$. Then $H_*(K_\bullet)$ has a barcode given by the multiset

$$\{[k(\operatorname{psimp} v_j), k(\operatorname{psimp} r_j)) \neq \emptyset \mid r_j \neq 0\} \cup \{[k(\operatorname{psimp} v_i), \infty) \mid r_i = 0 \text{ and } i \notin \operatorname{pivots} R\}.$$

Here, to deemphasize the dependence on the concrete index set $\{0, \dots, N\}$ we chose as notation, we write $[k(\operatorname{psimp} v_i), \infty)$ for the interval $[k(\operatorname{psimp} v_i), N]$ including the largest element of the index set.

1.2.3. Clearing

We recall the basic idea of the *clearing* optimization by Chen and Kerber [CK11] for the computation of barcodes of filtrations of simplicial complexes, also implicitly present in the work of de Silva, Morozov, and Vejdemo-Johansson [dSMV11]. Given a filtration boundary matrix D for simplices $\sigma_1, \dots, \sigma_n$ in filtration order for a filtration $K_\bullet$ as in § 1.2.2, we know that a reduction $R = DV$ of D can be used to determine the barcode of $H_*(K_\bullet)$. To compute such a reduction, we can make use of the homological grading on D to improve the basic matrix reduction algorithm from § 1.2.1. In short, this is because if we have $r_j \neq 0$, then we must have $r_i = 0$ for $i = \operatorname{pivot} r_j$, and we also have $\dim \sigma_i = \dim \sigma_j - 1$. This leads to the clearing procedure: Instead of simply reducing D by column operations from left to right, we reduce columns in decreasing order of their homological degree (increasing in the case of cohomology). Before using column operations to reduce the columns in dimension m, we set $r_j = 0$ for all indices j which appear as pivots of the already reduced columns in dimension $m + 1$.

Input: Filtration boundary matrix D with n rows
Output: A reduction $R = DV$
  $R \leftarrow D$
  $V \leftarrow I_n$
  for $m = \dim K, \dots, 1$ **do**
    while $\exists i < j$ with $r_i \neq 0$, $\operatorname{pivot} r_i = \operatorname{pivot} r_j$, and $\dim \sigma_i = m$ **do**
      $r_j \leftarrow r_j - \frac{\operatorname{pentry} r_j}{\operatorname{pentry} r_i} r_i$
      $v_j \leftarrow v_j - \frac{\operatorname{pentry} r_j}{\operatorname{pentry} r_i} v_i$
    end while
    for j with $\dim \sigma_j = m$ and $r_j \neq 0$ **do**
      $r_{\operatorname{pivot} r_j} \leftarrow 0$
      $v_{\operatorname{pivot} r_j} \leftarrow r_j$
    end for
  end for
  return R, V

1.3. The observable category

In some settings we consider, the standard notion of barcode is not quite suitable because some of the persistence modules that may appear do not admit one. To deal with this issue, we review the construction of the observable category of persistence modules in § 1.3.1. Via the radical construction, reviewed in § 1.3.2, the observable category leads to a well-defined notion of persistence diagram for q-tame persistence modules, which we discuss in § 1.3.3.

1.3.1. The observable category of persistence modules

By Crawley-Boevey's theorem [BC20; Cra15], every PFD persistence module has a barcode. This is no longer true for *q-tame* persistence modules M, which are persistence modules such that $\operatorname{rank} M_{s,t} < \infty$ for all $s < t$. For example, the $\mathbb{R}$-indexed persistence module $\prod_{n \in \mathbb{N}} C([0, n^{-1})$ does not decompose as a direct sum of interval modules. Still, there is a similar structure theory for q-tame persistence modules if one passes to the *observable category of persistence modules* [CCdS16]. For the construction of the observable category to work, we need to assume that the set T indexing our persistence modules is *dense*, meaning that for all $s, t \in T$ with $s < t$ there exists $u \in T$ with $s < u < t$.

Starting with some more terminology, a persistence module M is called *ephemeral* if $M_{s,t} = 0$ for all $s < t$. A morphism of persistence modules is called a *weak isomorphism* if its kernel and cokernel are ephemeral. Roughly, the idea for the observable category is to decompose persistence modules into interval modules only up to weak isomorphisms. However, since weak isomorphisms are not invertible, the relation of being weakly isomorphic is not symmetric. This can be dealt with by passing to a category in which weak isomorphism do become invertible.

Let M and N be persistence modules. An *observable morphism* $\varphi \colon M \dashrightarrow N$ is given by linear maps $\varphi_{t,u}\colon M_t \to N_u$ for $t < u$ such that the diagram

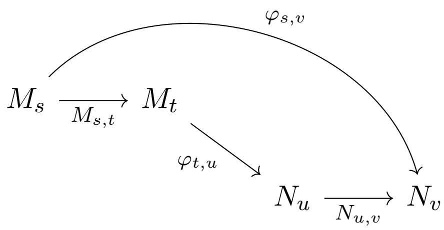

commutes for all $s \leq t < u \leq v$. The composition of two observable morphisms $\varphi\colon M \dashrightarrow N$ and $\psi\colon N \dashrightarrow P$ is defined by setting $(\psi \circ \varphi)_{s,u} = \psi_{t,u} \circ \varphi_{s,t}$ for some t with $s < t < u$, where such an index t can always be chosen because we assume T to be dense. This is easily seen to be well-defined by the above commutativity property. Observable identity morphisms are defined in the obvious way via the structure morphisms. We call the resulting category consisting of persistence modules with observable morphisms the *observable category of persistence modules* **Ob**. If persistence modules M and N are isomorphic in **Ob**, we write $M \simeq N$.

There is a functor $\pi\colon \mathbf{Pers} \to \mathbf{Ob}$ sending a persistence module to itself and a morphism $\varphi\colon M \to N$ to the observable morphism given by

$$\pi(\varphi)_{t,u} = N_{t,u} \circ \varphi_t = \varphi_u \circ M_{t,u}$$

for $t < u$. Importantly, π takes weak isomorphisms to isomorphisms by [CCdS16, Theorem 2.7].

1.3.2. Radicals

A major technical tool from [CCdS16] for working with q-tame persistence modules is the *radical* of a persistence module M, which is the persistence module $\operatorname{rad} M$ defined by

$$(\operatorname{rad} M)_t = \sum_{s<t} \operatorname{im} M_{s,t}$$

with structure maps given by the restrictions of those of M. The radical may also be abstractly described as the minimal submodule with an ephemeral quotient. Another possible definition is to view $(\operatorname{rad} M)_t$ as the image of the canonical morphism $\operatorname{colim}_{s<t} M_s \to M_t$.

An important property of radicals is that when passing to the observable category, a persistence module cannot be distinguished from its radical. More explicitly, we have that the inclusion $\operatorname{rad} M \hookrightarrow M$ is a weak isomorphism for every persistence module M, and hence $\operatorname{rad} M \simeq M$. This is helpful since, loosely speaking, the radical is usually more tame than the original module.

Observable barcode decompositions for q-tame persistence modules Now, assume M is a q-tame persistence module. While each vector space $\operatorname{im} M_{s,t}$ is finite-dimensional for $s < t$, the radical need still not be PFD. As an example, consider the $\mathbb{R}$-indexed persistence module $M = \bigoplus_{n \in \mathbb{N}} C((-n^{-1}, 0])$. Clearly, $\operatorname{rad} M = M$, but $(\operatorname{rad} M)_0 = M_0$ has infinite dimension. In any case, it turns out that the radical of a q-tame persistence module always has a barcode, given that the index set T is dense and every interval $I \subseteq T$ has a countable *coinitial* subset, i.e., a countable set $N \subseteq I \subseteq T$ such that for all $t \in I$ there exists $s \in N$ with $s \leq t$. These conditions are satisfied in the special case $T = \mathbb{R}$. That the radical has a barcode follows from a general version of Crawley-Boevey's theorem using chain conditions on kernel and images; see [CCdS16, Section 3.1 and Corollary 3.6] for more details. Via this barcode of the radical, we obtain an *observable barcode decomposition* for any q-tame persistence module, i.e., for any q-tame persistence module M there is a collection of intervals $(I_\alpha)_{\alpha \in A}$ such that

$$M \simeq \bigoplus_{\alpha \in A} C(I_\alpha).$$

However, such an observable barcode decomposition does not yield a well-defined barcode since different interval modules can be observably isomorphic. In fact, for any $a,\ b \in \mathbb{R}$ with $a < b$ the interval modules $C((a,b))$, $C([a,b))$, $C((a,b])$, and $C([a,b])$ are all observably isomorphic to each other. In order to get a well-defined invariant, we hence need to forget whether to include the endpoints of an interval or not, leading to the notion of persistence diagram.

1.3.3. Persistence diagrams

As an alternative to the barcode formalism, one may also describe the observable structure of persistence modules in terms of so-called persistence diagrams, which we describe here in the $\mathbb{R}$-indexed case. Starting with a barcode $B = (I_\alpha)_{\alpha \in A}$, we define its associated *persistence diagram* as the multiset given by the *multiplicity function* $\mathfrak{m}$ that associates to an element in

$$\mathcal{E} = \{(p,q) \mid p \in \mathbb{R} \cup \{-\infty\},\ q \in \mathbb{R} \cup \{+\infty\},\ p < q\}$$

the cardinality of the set $\{\alpha \in A \mid \inf I_\alpha = p,\ \sup I_\alpha = q\}$. Note that there may be foundational issues with this definition since there is no set containing all cardinals so that $\mathfrak{m}$ does not have a well-defined codomain. For the purposes of this thesis, this is not an issue since we will only be interested in persistence diagrams associated to q-tame persistence modules. For such persistence modules, the relevant cardinalities will always we finite so that $\mathfrak{m}$ can be considered as a function $\mathcal{E} \to \mathbb{N}$. The version of the persistence diagram defined above is also sometimes called the *undecorated* persistence diagram, because we do not record any information about whether the endpoints of an interval are included or not.

In the observable category, the persistence diagram is a complete invariant for interval-decomposable persistence modules. That is, we have $\bigoplus_{\alpha \in A} C(I_\alpha) \simeq \bigoplus_{\alpha' \in A'} C(I'_{\alpha'})$ if and only if the persistence diagrams associated to $(I_\alpha)_{\alpha \in A}$ and $(I'_{\alpha'})_{\alpha' \in A'}$ agree [CCdS16, Theorem 3.9]. This theorem follows from the fact that the observable category is a Grothendieck category and consequently a version of the Krull–Schmidt theorem holds; for details see the given reference. Finally, we can now define the *persistence diagram* of an $\mathbb{R}$-indexed q-tame persistence module as the persistence diagram associated to the barcode of its radical. This invariant then completely describes the isomorphism type of the persistence module in the observable category.

1.4. Čech and Vietoris homology

For parts of this thesis, we will use some less commonly studied homology theory, namely Čech homology. We review the standard definition of Čech homology and recall some of its important properties in § 1.4.1. We also review Vietoris homology in § 1.4.2, which is of historical interest in some of Morse's work that we consider in Chapter 7. These two homology theories turn out to be the same by a theorem of Dowker that we review in § 1.4.3.

1.4.1. Čech homology

Let us recall the definition of Čech homology as presented for example by Eilenberg and Steenrod [ES52, Section IX–X]. A *cover* of a topological space X is a set α consisting of open subsets of X such that $X = \bigcup_{U \in \alpha} U$. The set of all covers is denoted by $\mathrm{Cov}(X)$. It is a directed set with respect to the *refinement* relation, where a cover α of a space $A \subseteq X$ is said to *refine* a cover β of X if for all $U \in \alpha$ there exists $V \in \beta$ with $U \subseteq V$. Recall that for a cover $\alpha \in \mathrm{Cov}(X)$ its *nerve* $\mathrm{Nrv}(\alpha)$ is defined as the abstract simplicial complex

$$\mathrm{Nrv}(\alpha) = \{\beta \subseteq \alpha \mid \beta \text{ is finite and } \textstyle\bigcap_{U \in \beta} U \neq \emptyset\}.$$

For any abelian group G, the nerve construction composed with the functor of simplicial homology with coefficients in G defines a functor from $\mathrm{Cov}(X)$, regarded as a poset category, to the category of graded abelian groups. The *Čech homology with coefficients in G in dimension d of X* is defined as

$$\check{H}_d(X; G) = \lim_{\alpha \in \mathrm{Cov}(X)} H_d(\mathrm{Nrv}(\alpha); G).$$

Čech homology may also be defined for pairs of spaces (X, A) in a corresponding way, and there is also a natural definition of a boundary operator $\partial\colon \check{H}_*(X, A) \to \check{H}_{*-1}(A)$ induced by the boundary operators on nerves. For any coefficient group, Čech homology defines a homotopy invariant functor from the category of topological spaces to the category of abelian groups satisfying all of the Eilenberg–Steenrod axioms for homology, except for the fact that the long sequences for pairs of spaces need not be exact in general [ES52, Theorems IX.4.3, IX.4.4, IX.5.1, IX.6.1, and IX.7.6]. One does get long exact sequences, and thus a homology theory in the sense of Eilenberg–Steenrod, if one restricts to triangulable spaces [ES52, Corollary 9.4], or to vector space coefficients and pairs of compact Hausdorff spaces [Kel61].

Special properties of Čech homology One special feature of Čech homology compared to other homology theories such as singular homology is that it satisfies a strong version of the excision axiom for compact Hausdorff spaces [ES52, Theorem X.5.4], which, together with the fact that we have long exact sequences for compact Hausdorff spaces and field coefficients, implies by [tDie08, Theorem 10.7.2] that we get a natural *Mayer-Vietoris sequence*

$$\begin{array}{l} \cdots \longrightarrow \check{H}_{n+1}(X; \mathbb{F}) \\ \qquad\qquad \downarrow \\ \check{H}_n(X_1 \cap X_2; \mathbb{F}) \longrightarrow \check{H}_n(X_1; \mathbb{F}) \oplus \check{H}_n(X_2; \mathbb{F}) \longrightarrow \check{H}_n(X; \mathbb{F}) \\ \qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad \downarrow \\ \qquad\qquad\qquad\qquad\qquad\qquad \check{H}_{n-1}(X_1 \cap X_2; \mathbb{F}) \longrightarrow \cdots \end{array}$$

whenever $X_1,\ X_2 \subseteq X$ are all compact Hausdorff spaces with $X = X_1 \cup X_2$. Another advantage of Čech homology over other homology theories is provided by the fact that if $(X_i, A_i)_i$ is an inverse system of compact Hausdorff pairs, then the inverse limit $\lim_i(X_i, A_i)$ in the category of pairs of topological spaces is again a compact Hausdorff pair and for all d and G the natural map

$$\check{H}_d(\lim_i(X_i, A_i); G) \to \lim_i \check{H}_d(X_i, A_i; G)$$

is an isomorphism [ES52, Theorems VIII.3.6 and X.3.1].

1.4.2. Vietoris homology

We review the definition of Vietoris homology in terms of Vietoris complexes of covers, generalized from the original definition for metric spaces in terms of ball covers [Vie27]. For a cover α of a topological space X, we define its *Vietoris complex* to be the simplicial complex given by

$$\mathrm{Vtr}(\alpha) = \{\rho \subseteq X \mid \rho \text{ is finite and } \rho \subseteq U \text{ for some } U \in \alpha\}\,.$$

If α refines β, we get a simplicial inclusion $\pi\colon \mathrm{Vtr}(\alpha) \to \mathrm{Vtr}(\beta)$. We can compose the Vietoris complex construction with simplicial homology with coefficients in some abelian group G, which yields a functor from $\mathrm{Cov}(X)$ to the category of graded abelian groups. Then the *Vietoris homology with coefficients in G in dimension d of X* is defined as

$$H_d^V(X; G) = \lim_{\alpha \in \mathrm{Cov}(X)} H_d(\mathrm{Vtr}(\alpha); G).$$

Vietoris homology is also defined for pairs of spaces (X, A) in a corresponding way, and there is a natural boundary operator $\partial\colon H_*^V(X, A) \to H_{*-1}^V(A)$ induced by the boundary operators on nerves.

Metric Vietoris homology If X is a metric space, the construction above is not exactly the same as the construction of *metric* Vietoris homology, as originally defined by Vietoris [Vie27], which in modern notation can be expressed as the limit over all covers of X by metric δ-balls, i.e.,

$$\lim_{\alpha \in \mathrm{Balls}(X)} H_*(\mathrm{Vtr}(\alpha); G),$$

where

$$\mathrm{Balls}(X) = \{(B_\delta(x))_{x \in X} \mid \delta > 0\} \subseteq \mathrm{Cov}(X).$$

If the metric space X is compact, then for each cover α there exists $\lambda > 0$ by Lebesgue's number lemma [Mun00, Lemma 27.5] such that $(B_\lambda(x))_{x \in X}$ refines α. In other words, if the metric space X is compact, then $\mathrm{Balls}(X)$ is coinitial in $\mathrm{Cov}(X)$, that is to say, they yield the same limit and hence the two versions of Vietoris homology agree with each other.

1.4.3. Dowker's theorem

The Vietoris and nerve constructions are dual to each other in the sense of Dowker's theorem [Dow52], which asserts that the two complexes $\mathrm{Nrv}(\alpha)$ and $\mathrm{Vtr}(\alpha)$ are homotopy equivalent after geometric realization for any cover α of any topological space X. As a consequence, we have that $H_*(\mathrm{Nrv}(\alpha); G) \cong H_*(\mathrm{Vtr}(\alpha); G)$. This isomorphism is natural with respect to refinement of covers and maps between spaces, so we get a natural isomorphism

$$\check{H}_*(X; G) = \lim_{\alpha \in \mathrm{Cov}(X)} H_*(\mathrm{Nrv}(\alpha); G) \cong \lim_{\alpha \in \mathrm{Cov}(X)} H_*(\mathrm{Vtr}(\alpha); G) = H_*^V(X; G)$$

after passing to limits over $\mathrm{Cov}(X)$. This result naturally extends to the setting of pairs of spaces, and the isomorphism is also compatible with the boundary operators. Thus, we may use Čech and Vietoris homology completely interchangeably from now on. In particular, Vietoris homology also commutes with inverse limits for compact Hausdorff pairs, and is a homology theory in the sense of Eilenberg–Steenrod for such spaces if one restrict to field coefficients.

2. Results

We now introduce our results. Our contributions, and also this thesis, diverge into two more or less separate lines of research, with the first part being motivated by computational problems in topological data analysis, and the second part being inspired by the origins of persistence in Morse theory. Starting with Section 2.1, we present our findings in the first part. They are centered around the framework of lifespan functors, which is applied to classify injectives and projectives in certain categories as well as to duality results and the efficient computation of images in persistent homology. These results have previously appeared in two articles written in collaboration with Ulrich Bauer [BS21a; BS22]. In Section 2.2, we then discuss our findings in the second part, which are related to Morse theory and concern Čech homology, q-tameness, and the structure theory of persistence modules. These results have previoulsy appeared in three articles [BMS21; Sch22a; Sch22b], the first of which was written in collaboration with Ulrich Bauer and Anibal Medina-Mardones.

Recall that we fixed a totally ordered set T acting as the index set of our persistence modules, which will sometimes be specialized to T being a finite set or the real numbers. Moreover, recall that we fixed a field $\mathbb{F}$ underlying all vector spaces and used as a coefficient group for all (co)homology groups and (co)chain complexes unless stated otherwise.

2.1. Lifespan functors and applications

In this section we describe our results pertaining to what we call lifespan functors. To motivate their introduction, we describe the problem of computing images in persistent homology, and explain how lifespan functors naturally appear when talking about certain dualities in persistent homology in § 2.1.1. There, we also state the main theorem for lifespan functors of persistence modules, which describes their impact in terms of barcodes. In § 2.1.2 we state our main result concerning injectivity and projectivity of PFD persistence modules in terms of lifespan functors, and the previously mentioned dualities are generalized and described in more detail. Finally, § 2.1.3 contains our results on applying these dualities and other optimizations for computing images in persistent homology.

2.1.1. Introducing lifespan functors

Images in persistent homology Roughly, the famous (algebraic) stability theorem for persistence modules (reviewed in § 1.1.5) states that closely interleaved persistence modules admit matchings, i.e., partial bijections, between their barcodes that only match intervals whose endpoints are close to each other. One way to approach the stability theorem is via *induced matchings*, which were introduced by Bauer and Lesnick [BL15]. In the computational setting, given a morphism of simplicial filtrations $f_\bullet \colon L_\bullet \to K_\bullet$, i.e., a commutative diagram of simplicial maps

$$\begin{array}{ccccccccc}
\emptyset = L_{-\infty} = L_0 & \hookrightarrow & L_1 & \hookrightarrow & \dots & \hookrightarrow & L_{N-1} & \hookrightarrow & L_N = L_\infty = L \\
\downarrow f_0 & & \downarrow f_1 & & & & \downarrow f_{N-1} & & \downarrow f_N = f \\
\emptyset = K_{-\infty} = K_0 & \hookrightarrow & K_1 & \hookrightarrow & \dots & \hookrightarrow & K_{N-1} & \hookrightarrow & K_N = K_\infty = K,
\end{array}$$

the homology functor induces a morphism $H_*(f_\bullet)\colon H_*(L_\bullet) \to H_*(K_\bullet)$ of persistence modules. From this morphism the induced matching construction yields a matching between the barcodes of $H_*(L_\bullet)$ and $H_*(K_\bullet)$ that can be used to bound the bottleneck distance between these two barcodes from above. The induced matching is defined in terms of the *image persistence* of $f_\bullet$, i.e., the barcode of $\operatorname{im} H_*(f_\bullet)$, motivating the problem of computing this barcode. A first algorithm for this problem was given by Cohen-Steiner et al. [Coh+09] for the special case where $f_\bullet$ is of the form $L_\bullet = K_\bullet \cap L \hookrightarrow K_\bullet$. In addition to their algorithm for image persistence, Cohen-Steiner et al. [Coh+09] present algorithms for computing barcodes of the kernel and cokernel of the morphism $H_*(f_\bullet)$. All of their algorithms rely on the standard reduction of boundary matrices (reviewed in § 1.2.2).

The barcode of the image of a morphism of persistence modules has various applications besides the construction of the induced matching. Cohen-Steiner et al. [Coh+09] propose applications of the image barcode to recovering the persistent homology of a noisy function on a noisy domain; see also the related work by Chazal et al. [Cha+11]. More recently, Reani and Bobrowski [RB21] proposed a method that includes the computation of induced matchings in order to pair up common topological features in different data sets, with applications to statistical bootstrapping. Furthermore, the computation of image barcodes is used in a distributed algorithm for persistent homology computation based on the Mayer–Vietoris spectral sequence by Casas [Cas20]. Image persistence of endomorphisms such as Steenrod squares on the persistent (co)homology of a single filtration has also been proposed by Lupo, Medina-Mardones, and Tauzin [LMT22] as a tool to get more comprehensive invariants than the standard persistent (co)homology barcodes.

Despite the usefulness of image persistence, there are a few aspects that have prevented these techniques from being widely used in applications so far. Specifically, to the best of our knowledge, there was no publicly available implementation for the algorithm by Cohen-Steiner et al. [Coh+09]. Furthermore, computation using the known algorithms is slow in comparison to modern algorithms for a single filtration. Indeed, computing usual persistent homology for larger data sets arising in real-world applications only became feasible in recent years due to optimizations that exploit various structural properties and algebraic identities of the problem [Bau21; CK11; dSMV11]. The main motivation for the work summarized in this section is to develop a theory allowing for the adaptation of these speed-ups to the computation of images and induced matchings.

Dualities in persistent homology One of the most important improvements for barcode computations relies on the use of cohomology based algorithms. These were first studied by de Silva, Morozov, and Vejdemo-Johansson [dSMV11] and justified by certain duality results. In summary, these duality results provide correspondences between the barcodes for persistent homology and for *persistent cohomology*, i.e.,

$$H^*(K_\bullet)\colon H^*(K_0) \leftarrow H^*(K_1) \leftarrow \cdots \leftarrow H^*(K),$$

as well as the barcodes for *persistent relative homology*

$$H_*(K, K_\bullet)\colon H_*(K, K_0) \to H_*(K, K_1) \to \cdots \to H_*(K, K)$$

and *persistent relative cohomology*

$$H^*(K, K_\bullet)\colon H^*(K, K_0) \leftarrow H^*(K, K_1) \leftarrow \cdots \leftarrow H^*(K, K).$$

The homology persistence modules simply have the same barcode as their cohomology counterparts [dSMV11, Proposition 2.3]. For the absolute-relative correspondence [dSMV11, Proposition 2.4], it turns out that the bounded intervals in the barcodes of $H_{d-1}(K_\bullet)$ and $H_d(K, K_\bullet)$ are also exactly the same, and there is a one-to-one correspondence between intervals of the form $[a, \infty)$ in the barcode of $H_d(K_\bullet)$ and intervals of the form $(-\infty, a)$ in the barcode of $H_d(K, K_\bullet)$.

The original proof for the absolute-relative correspondence uses a decomposition of filtered chain complexes. This strategy relies on a non-canonical choice, which does not extend to the functorial setting. We thus adopt a different point of view based on the long exact sequence of a pair in homology. Applying this functorial construction to a filtration $K_\bullet$, we obtain a long exact sequence of persistence modules

$$\cdots \longrightarrow \Delta H_d(K) \xrightarrow{\epsilon_d} H_d(K, K_\bullet) \xrightarrow{\partial} H_{d-1}(K_\bullet) \xrightarrow{\eta_{d-1}} \Delta H_{d-1}(K) \longrightarrow \cdots$$

where Δ denotes constant persistence modules. As it turns out, the first two of the short exact sequences

$$\begin{array}{ccccccccc} 0 & \longrightarrow & \operatorname{im}\partial & \longrightarrow & H_{d-1}(K_\bullet) & \longrightarrow & \operatorname{im}\eta_{d-1} & \longrightarrow & 0 \\ 0 & \longrightarrow & \operatorname{im}\epsilon_d & \longrightarrow & H_d(K, K_\bullet) & \longrightarrow & \operatorname{im}\partial & \longrightarrow & 0 \\ 0 & \longrightarrow & \operatorname{im}\eta_d & \longrightarrow & \Delta H_d(K) & \longrightarrow & \operatorname{im}\epsilon_d & \longrightarrow & 0 \end{array}$$

split (as a special case of Corollary 4.2.8), showing that $\operatorname{im}\partial \cong \ker \eta_{d-1} \cong \operatorname{coker}\epsilon_d$ is a summand of both $H_d(K, K_\bullet)$ and $H_{d-1}(K_\bullet)$. Its barcode consists of the bounded intervals of either persistence module. Moreover, the third short exact sequence has a constant persistence module $\Delta H_d(K)$ in the middle, implying that the persistence modules $\operatorname{im}\eta_d$ and $\operatorname{im}\epsilon_d$ determine each other. Together, this shows that the barcodes of $H_*(K_\bullet)$ and $H_*(K, K_\bullet)$ completely determine each other. For details see § 5.3.1.

By observing that $\Delta H_d(K) \cong \Delta \operatorname{colim} H_d(K) \cong \Delta \lim H_d(K, K_\bullet)$ and that ϵ and η are the counit and the unit of the adjunctions $\Delta \dashv \lim$ and $\operatorname{colim} \dashv \Delta$, respectively, we can generalize this discussion to arbitrary persistence modules. Taking images, kernels, and cokernels of the morphisms ϵ and η yields endofunctors on the category of persistence modules, which we call *lifespan functors*. In particular, the *mortal part* $(-)^\dagger = \ker \eta_{(-)}$ and the *immortal part* $(-)^\infty = \operatorname{im}\eta_{(-)}$ determine death in the persistence module, while the *nascent part* $(-)^* = \operatorname{coker}\epsilon_{(-)}$ and the *ancient part* $(-)^{-\infty} = \operatorname{im}\epsilon_{(-)}$ determine birth.

Effect of lifespan functors on barcodes The general definition of lifespan functors above also works in the category of *matching diagrams*, since these diagrams admit limits and colimits (Proposition 4.2.1). The category of matching diagrams is equivalent to the category of barcodes [BL20]. We will review the definitions of these categories in § 3.1.1 and the equivalence between them in § 3.1.2. The effect of the lifespan functors on persistence modules and matching diagrams is best described in terms of barcodes, as for the above example $\operatorname{im}\partial \cong \ker \eta_{d-1} \cong \operatorname{coker} \epsilon_d$, whose barcode corresponds to the bounded intervals; see also Figure 2.1 for an illustration.

Definition (Definition 4.2.3)**.** Write $\mathfrak{I} = \mathfrak{I}(T)$ for the set of all intervals in T. We define the following subsets of $\mathfrak{I}$.

$$\begin{aligned} \mathfrak{I}^* &= \{I \in \mathfrak{I} \mid I \text{ is strictly bounded below}\}, & \mathfrak{I}^{-\infty} &= \mathfrak{I} \setminus \mathfrak{I}^*, \\ \mathfrak{I}^\dagger &= \{I \in \mathfrak{I} \mid I \text{ is strictly bounded above}\}, & \mathfrak{I}^{\infty} &= \mathfrak{I} \setminus \mathfrak{I}^\dagger, \\ \mathfrak{I}^{\dagger,*} &= \mathfrak{I}^* \cap \mathfrak{I}^\dagger, & \mathfrak{I}^{-\infty,\infty} &= \mathfrak{I}^{-\infty} \cap \mathfrak{I}^{\infty}, \\ \mathfrak{I}^{-\infty,\dagger} &= \mathfrak{I}^{-\infty} \cap \mathfrak{I}^\dagger, & \mathfrak{I}^{*,\infty} &= \mathfrak{I}^* \cap \mathfrak{I}^{\infty}. \end{aligned}$$

If B is a barcode, we also define

$$B^\diamond = \{I \in B \mid I \in \mathfrak{I}^\diamond\}$$

for any $(-)^\diamond$ from the list above. These operations are referred to as *lifespan* operations.

With this terminology, we can now state the main theorem for lifespan functors of persistence modules. It tells us that the above manipulations of barcodes may be achieved on the level of persistence modules in a functorial way. This is particularly remarkable in light of the fact that there is no functor from **pers** to barcodes that assigns to each PFD persistence module its barcode [BL15, Proposition 5.10].

Theorem (Corollary 4.2.7)**.** *There are* lifespan functors $(-)^\diamond \colon \mathbf{Pers} \to \mathbf{Pers}$ *from the category of persistence modules to itself, constructed by forming images, kernels, and cokernels of the counit and the unit of the adjunctions* $\Delta \dashv \lim$ *and* $\operatorname{colim} \dashv \Delta$*, such that if* M *is a persistence module and* B *is a barcode of* M*, then* $B^\diamond$ *is a barcode for* $M^\diamond$*, where* $(-)^\diamond$ *is any lifespan functor.*

2.1.2. Some applications of lifespan functors

Injectives and projectives in categories of barcodes and persistence modules In some sense, the study of persistent homology may be seen as the study of homological algebra in the abelian category of persistence modules; see [BM21] for a detailed account of this perspective on persistence theory. As such, projective and injective resolutions, and hence projective and injective objects, in the category of persistence modules are of interest. Bubenik and Milićević [BM21] give classifications of interval modules that have these properties in **Pers**. Höppner [Höp83] and Höppner and Lenzing [HL81a; HL81b] give structural results for injective and projective objects in very general diagram categories including **Pers**, showing, among other things, that projective objects in **Pers** always have barcodes (long before the barcode terminology was introduced). Particularly in computational settings, however, one often only deals with PFD persistence modules and restricting from **Pers** to **pers** changes which persistence modules are injective and projective. Moreover, as argued by Bauer and Lesnick [BL20], it is also worthwhile to consider homological algebra in the categories of barcodes and matchings diagrams, which are not abelian but have the structure of a *p-exact* category (we will recall this in more detail in § 3.1.3). This structure also allows one to

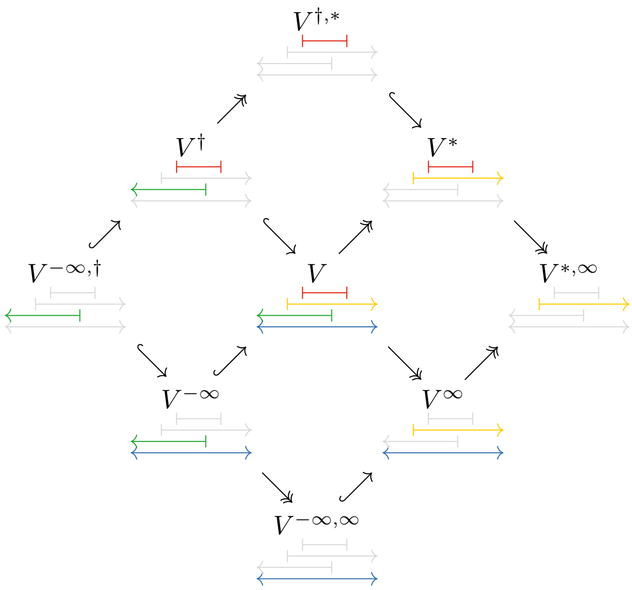

Figure 2.1.: Lifespan functors applied to a finite type $\mathbb{R}$-indexed persistence module V, visualized via their barcode according to Corollary 4.2.7.

develop large parts of the standard theory of homological algebra as presented for example in [Gra12].

As a first application of lifespan functors, we can now give a simple characterization of the projective and injective objects in the category of barcodes (or matching diagrams).

Theorem (Theorem 5.1.2)**.** *A barcode B is projective if and only if the mortal part of B satisfies $B^{\dagger} = 0$, and injective if and only if the nascent part of B satisfies $B^{*} = 0$.*

In the same vein, projective objects in **pers** can be characterized by their vanishing mortal parts and injective objects by their vanishing nascent part. In addition, these properties turn out to be equivalent to a persistence module only having injective or surjective structure maps, respectively, which is also true in the case of matching diagrams (Corollary 5.1.5).

Theorem (Theorem 5.2.3)**.** *Let M be a* PFD *persistence module.*

1. *The following are equivalent:*
   a) *All structure maps of M are monomorphisms.*
   b) *The mortal part of M satisfies $M^{\dagger} = 0$.*
   c) *M is projective in the category of* PFD *persistence modules* **pers**.
2. *The following are equivalent:*
   a) *All structure maps of M are epimorphisms.*
   b) *The nascent part of M satisfies $M^{*} = 0$.*
   c) *M is injective in the category of* PFD *persistence modules* **pers**.

Functorial dualities in persistent homology The lifespan functors allow us to succinctly express the previously described absolute-relative correspondence between the barcodes of $H_*(K_\bullet)$ and $H_*(K, K_\bullet)$ in terms of natural isomorphisms. The correspondence $[a, \infty) \leftrightarrow (-\infty, a)$ between the semi-infinite intervals is encapsulated via functors $(-)^\vartriangleright$ and $(-)^\vartriangleleft$, which are complementary to the lifespan functors and, under favorable circumstances, change the barcode of a persistence module by taking complements of semi-infinite intervals (Definition 4.1.7 and Propositions 4.2.9 and 4.2.10).

To state our result in its full generality, recall that the category of topological spaces has all colimits. In particular, if X is a diagram of topological spaces indexed by our index set T, it has a colimit. The natural map $X \to \Delta \operatorname{colim} X$ from X to the constant T-indexed diagram of its colimit $\Delta \operatorname{colim} X$ induces a map $C_*(X) \to C_*(\Delta \operatorname{colim} X)$ of persistent singular chain complexes. If X is a filtration, i.e., all of its structure maps are injective, then this map is a monomorphism. The cokernel of this map is again a persistent chain complex, and its homology is denoted by $H_*(\operatorname{colim} X, X)$ as a generalization of the persistent relative homology $H_*(K, K_\bullet)$ that we described in the computational setting of filtrations of simplicial complexes.

Theorem (Theorem 5.3.2)**.** *Let X be a filtration of topological spaces such that the natural maps* $\operatorname{colim} H_*(X) \to H_*(\operatorname{colim} X) \to \lim H_*(\operatorname{colim} X, X)$ *are isomorphisms. Then for all d we have the following isomorphisms, which are natural in X:*

$$\begin{aligned} H_{d-1}(X)^\dagger &\cong H_d(\operatorname{colim} X, X)^*, \\ H_d(X)^\vartriangleleft &\cong H_d(\operatorname{colim} X, X)^{-\infty}, \\ H_d(X)^\infty &\cong H_d(\operatorname{colim} X, X)^\vartriangleright. \end{aligned}$$

Dualities for image persistence In Theorem 5.3.2, naturality of the isomorphisms is inherited from the construction of the long exact homology sequence. In particular, from a morphism of filtrations of simplicial complexes $f_\bullet \colon L_\bullet \to K_\bullet$ inducing a map on colimits $f \colon L \to K$, we get an isomorphism $H_{d-1}(f_\bullet)^\dagger \cong H_d(f, f_\bullet)^*$ in the category of morphisms of persistence modules. We also get a morphism

$$\begin{array}{ccccccccccc} \cdots & \longrightarrow & \Delta H_d(K) & \xrightarrow{\epsilon_d} & H_d(K, K_\bullet) & \xrightarrow{\partial} & H_{d-1}(K_\bullet) & \xrightarrow{\eta_{d-1}} & \Delta H_{d-1}(K) & \longrightarrow & \cdots \\ & & \uparrow & & \uparrow & & \uparrow & & \uparrow & & \\ \cdots & \longrightarrow & \Delta H_d(L) & \xrightarrow{\epsilon_d} & H_d(L, L_\bullet) & \xrightarrow{\partial} & H_{d-1}(L_\bullet) & \xrightarrow{\eta_{d-1}} & \Delta H_{d-1}(L) & \longrightarrow & \cdots \end{array}$$

of long exact sequences, with vertical maps induced by $f_\bullet$. Note, however, that the induced sequences of kernels, images, and cokernels are no longer exact in general, so the rest of the proof of the absolute-relative correspondence for a single filtration does not carry over completely to the setting of image persistence described in § 2.1.1. In order to still obtain a useful absolute-relative correspondence for $\operatorname{im} H_*(f_\bullet)$, which we require for our algorithmic applications, we need to develop conditions for when the lifespan functors, which are in general not exact functors, commute with passing from morphisms to their images.

Theorem (Theorem 4.1.10)**.** *Let M and N be persistence modules and $\varphi \colon M \to N$ a morphism. If the induced map* $\operatorname{colim} \varphi \colon \operatorname{colim} M \to \operatorname{colim} N$ *is a monomorphism, we have canonical isomorphisms for the mortal and immortal parts*

$$\operatorname{im} \varphi^\dagger \cong (\operatorname{im} \varphi)^\dagger, \qquad \operatorname{im} \varphi^\infty \cong (\operatorname{im} \varphi)^\infty \cong M^\infty.$$

If the induced map $\lim \varphi \colon \lim M \to \lim N$ *is an epimorphism, we have canonical isomorphisms for the nascent and ancient parts*

$$\operatorname{im} \varphi^* \cong (\operatorname{im} \varphi)^*, \qquad \operatorname{im} \varphi^{-\infty} \cong (\operatorname{im} \varphi)^{-\infty} \cong N^{-\infty}.$$

Regarding the morphism $f_\bullet$ of simplicial filtrations from above, if the colimit morphism $f \colon L \to K$ is an isomorphism, then Theorems 4.1.10 and 5.3.2 together imply for example that

$$(\operatorname{im} H_{d-1}(f_\bullet))^\dagger \cong \operatorname{im} \left(H_{d-1}(f_\bullet)^\dagger\right) \cong \operatorname{im} \left(H_d(f, f_\bullet)^*\right) \cong (\operatorname{im} H_d(f, f_\bullet))^*.$$

This means that, as in the single filtration case, the bounded intervals in the barcodes of the images of the absolute and relative morphism agree.

To complete the picture of the known dualities for a single filtration in the image persistence setting, we will also state a functorial version of the correspondence between persistent homology and cohomology in terms of vector space duality (Proposition 5.3.4) and analyze how the lifespan functors behave with respect to dualization (§ 4.2.2). In summary, we then have the following correspondence for images of morphisms in persistent absolute homology and relative cohomology.

Corollary (Corollary 5.3.5). *Let* $f_\bullet \colon L_\bullet \to K_\bullet$ *be a morphism of filtrations of finite simplicial complexes inducing a map* $f \colon L \to K$ *on colimits such that* $H_*(f)$ *is an isomorphism. Assume that the index set* T *has a largest element and a smallest element* $t_{\min}$ *with* $L_{t_{\min}} = K_{t_{\min}} = \emptyset$*, so that no intervals in the barcodes of* $L_\bullet$ *and* $K_\bullet$ *contain* $t_{\min}$*. Then*

$$B(\operatorname{im} H_{d-1}(f_\bullet))^{\dagger,*} = B(\operatorname{im} H^d(f, f_\bullet))^{\dagger,*}$$

for all degrees d*, and the map* $I \mapsto T \setminus I$ *defines bijections*

$$\begin{aligned} B(\operatorname{im} H_*(f_\bullet))^\infty &\leftrightarrow B(H^*(L, L_\bullet))^{-\infty}, \\ B(\operatorname{im} H^*(f, f_\bullet))^{-\infty} &\leftrightarrow B(H_*(K_\bullet))^\infty. \end{aligned}$$

2.1.3. Efficient computation of barcodes for images in persistent homology

Clearing for image persistence We have reviewed in Section 1.2 that the basic algorithm for computing persistent homology is based on performing matrix reduction on a filtration boundary matrix, which can be improved by the clearing optimization. The basic algorithm for computing image persistence introduced by Cohen-Steiner et al. [Coh+09] not only requires the reduction of a filtration boundary matrix, but also the reduction of a permuted boundary matrix, to which clearing cannot be straightforwardly applied. We will remedy this by showing that before reducing it one can delete the columns in the permuted boundary matrix that were already reduced to 0 in the boundary matrix corresponding to the codomain filtration (Corollary 6.2.7).

The clearing optimization is particularly useful when applied in the setting of computing relative cohomology. This is due to the fact that there, the grading is *co*homological, so we can apply clearing in ascending rather than descending dimension. Thus, one can restrict to computing barcodes in small dimensions, which increases the feasibility of the computations and is usually sufficient for practical applications, see [Bau21] for a more detailed discussion.

To apply clearing in the relative cohomology setting for image persistence, we will reformulate the algorithm by Cohen-Steiner et al. [Coh+09] in the purely algebraic setting of filtered chain complexes of vector spaces. More precisely, we will consider two filtrations of (co)chain complexes $C_\bullet$ and $C'_\bullet$ and a monomorphism $\varphi_\bullet\colon C_\bullet \to C'_\bullet$ that induces an isomorphism on colimits. For our morphism $f_\bullet\colon L_\bullet \to K_\bullet$, his setup includes both the absolute homology case $C_*(L_\bullet) \hookrightarrow C_*(K_\bullet)$ and the relative cohomology case $C^*(K, K_\bullet) \hookrightarrow C^*(L, L_\bullet)$. The general idea for computing the image of $H_*(\varphi_\bullet)$ is then to first write it as a subquotient of $C'_\bullet$:

$$\operatorname{im} H_*(\varphi_\bullet) \cong \frac{\varphi_\bullet(Z_*(C_\bullet))}{\varphi_\bullet(Z_*(C_\bullet)) \cap B_*(C'_\bullet)},$$

where Z_* and B_* denote cycles and boundaries of the corresponding chain complexes, respectively, and where the intersection of persistence modules is to be interpreted indexwise, meaning that $(\varphi_\bullet(Z_*(C_\bullet)) \cap B_*(C'_\bullet))_t = \varphi_t(Z_*(C_t)) \cap B_*(C'_t)$. Performing matrix reductions that make use of the clearing optimization, we find a pair of inclusion-related *filtration compatible bases* for the filtrations appearing in the equation above. Filtration compatible bases provide a formal framework for many standard arguments for barcode computations via matrix reduction, and they can be interpreted as special instances of matching diagrams. Using the general theory of matching diagrams, the data we compute can then easily be shown to determine the barcode of $\operatorname{im} H_*(\varphi_\bullet)$ (Theorem 6.2.1).

The final algorithm We now come back to the computational setting and fix a monomorphism $f_\bullet\colon L_\bullet \to K_\bullet$ of filtrations of finite simplicial complexes inducing an isomorphism $f\colon L \to K$ on the colimits. If D^L is a filtration boundary matrix for $L_\bullet$ (see § 1.2.2), then its transpose along the anti-diagonal, denoted $(D^L)^\perp$, is a filtration *co*boundary matrix for the relative cochain complex $C^*(L, L_\bullet)$. Denoting by F the matrix encoding $C_*(f)$ with respect to the simplices of L and K in their respective filtration orders, we set $D^f = D^L F^{-1}$. Applying the general considerations from before in this setting and combining this with the translation between relative cohomology and absolute homology from § 2.1.2 yields an algorithm for computing the absolute homology image of $f_\bullet\colon L_\bullet \to K_\bullet$ by reducing the two coboundary matrices $(D^L)^\perp$ and $(D^f)^\perp$, which can be done with clearing.

Corollary (Corollary 6.2.8)**.** *Let $f_\bullet\colon L_\bullet \to K_\bullet$ be a monomorphism of filtrations of finite simplicial complexes inducing an isomorphism $f\colon L \to K$ on the colimits. Assume that the index set T has a largest element and a smallest element $t_{\min}$ with $L_{t_{\min}} = K_{t_{\min}} = \emptyset$. Let $L_\bullet$ and $K_\bullet$ be sublevel set filtrations of functions $l\colon L \to T$ and $k\colon K \to T$, respectively.*

Then the associated coboundary matrices $(D^L)^\perp$ and $(D^f)^\perp$ can be reduced with clearing, and given reductions $S = (D^f)^\perp W$ and $R = (D^L)^\perp V$ the barcode of $\operatorname{im} H_*(f_\bullet)$ can be determined as the multiset

$$\{[l(\operatorname{psimp} w_j), k(\operatorname{psimp} s_j)) \neq \emptyset \mid s_j \neq 0\} \cup \{[l(\operatorname{psimp} v_i), \infty) \mid r_i = 0 \textit{ and } i \notin \operatorname{pivots} R\},$$

where psimp c *is the simplex corresponding to the index* pivot c *for any column c.*

An implementation of this method based on Ripser [Bau21] is publicly available [BS21b]. Our software works under the assumption that $L_\bullet = \operatorname{Rips}_\bullet(X, d)$ and $K_\bullet = \operatorname{Rips}_\bullet(X, d')$ are filtrations of Vietoris–Rips complexes corresponding to two metrics d and d' on a finite set X that satisfy $d(x, y) \geq d'(x, y)$ for all $x, y \in X$ with the maps $f_t\colon L_t \to K_t$ being given by inclusion. Note that $L_t = \operatorname{Rips}_t(X, d)$ being a subcomplex of $K_t = \operatorname{Rips}_t(X, d')$ for all t is ensured by the inequality $d \geq d'$. The implementation also makes use of a version of the emergent and apparent pairs optimizations [Bau21], which we do not discuss here, referring to [BS22] instead.

2.2. Q-tameness, Čech homology, and Morse theory for minimal surfaces

In this section we describe our results concerning Čech homology, q-tameness, and the structure theory for q-tame semi-continuous persistence modules. These are inspired by Morse's work on minimal surfaces, which is part of a general framework that Morse termed functional topology. We start with a review of the unstable minimal surface theorem by Morse and Tompkins in § 2.2.1, focusing on a slightly reformulated version of the mountain pass theorem originally used to prove the existence of an unstable minimal surface. We also formulate Morse inequalities for q-tame persistence modules purely in terms of their persistence diagrams. Motivated by these developments, we then go on to see that filtrations of compact spaces satisfying a certain local connectedness condition have q-tame persistent homology in § 2.2.2, fixing an inaccuracy in Morse's work and providing a generalization of similar statements in the existing persistence literature. A similar local connectedness condition is sufficient for implying that the interleaving distance between the persistent singular and the persistent Čech homology of the filtration is 0, which we discuss in § 2.2.3. The persistent Čech homology of a filtration of compact Hausdorff spaces, also considered in Morse's functional topology, has a certain semi-continuity property, which appears in a structural result for q-tame persistence modules presented in § 2.2.4.

2.2.1. Morse inequalities in terms of persistence and minimal surfaces

Functional Topology Given that persistent homology can be used for analyzing sublevel sets of functions, it might not come as a surprise that it can be intimately linked to Morse theory. Nowadays, when thinking about the work of Marston Morse, our first thought probably involves a differentiable function on a closed smooth manifold, but more general settings should also be considered. Morse theory in the smooth context was presented in Milnor's famous book on the subject [Mil63], where he also gave a new proof of Bott's periodicity by applying Morse theory to the energy functional of paths in a Riemannian manifold, which notably goes beyond the compact setting. Another important example of the use of Morse's insights in an infinite-dimensional context is Floer's work on the Arnold conjecture and its many ramifications in symplectic topology, as surveyed for example in [Sal99]. Morse himself worked in a very general setting, publishing in the 1930s a pair of papers [Mor37; Mor40] and a monograph [Mor38] in which he established the key results of Morse theory in the broad context defined by semi-continuous functionals on metric spaces. He called the theory set forth in this body of work *functional topology* and used it to study questions about minimal surfaces motivated by Douglas' solution to Plateau's problem [Dou31]. In particular, Morse and Tompkins [MT39] used these techniques to prove a general *mountain pass theorem* – an existence result for saddle points – applying to functions that are not necessarily differentiable or not even continuous.

A mountain pass theorem for homotopically critical points To state the mountain pass theorem on the existence of saddle points in a non-differentiable, or even non-continuous setting, we need an appropriate notion of critical points.

Definition (Definition 7.2.1, [Str88, Definition II.6.1-II.6.2], [MT39, p. 445, 466])**.** Consider a real-valued function F on a metric space (M, d). A point $p \in M$ is called *homotopically* *regular* if there exists a neighborhood U of p in $F_{\leq F(p)}$ and a continuous map $\varphi\colon U \times [0,1] \to M$, which satisfies $\varphi(\cdot, 0) = \mathrm{id}_U$ and $\varphi(p,1) \neq p$, such that for every compact subset $V \subseteq U$ there exists a continuous *displacement function* $\delta\colon \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$. That is, a continuous function δ with $\delta(e) = 0$ if and only if $e = 0$ and

$$F(\varphi(x,s)) - F(\varphi(x,t)) \geq \delta(d(\varphi(x,s), \varphi(x,t)))$$

for all $x \in V$ and $0 \leq s \leq t \leq 1$. A point that is not homotopically regular is called *homotopically critical*. Function values of homotopically critical points will be called *critical values* and all other values will be called *regular values*. A *critical set* S is a closed and open subset of the subspace of all homotopically critical points with a given function value. It is said to be of *minimum type* if there exists a neighborhood N of the closure $\overline{S}$ of S, taken in M, such that the function values on $N \setminus S$ strictly exceed the function value on S.

Note that, in particular, an isolated local minimum constitutes a critical set of minimum type. Similarly, a critical submanifold of a Morse-Bott function on which the function values are locally minimized is also a critical set of minimum type.

Definition (Definition 7.2.2, [MT39, p. 445])**.** Let $F\colon M \to \mathbb{R}$ be a function on a metric space. We say that F is *weakly upper-reducible* if for all $p \in M$ and all $c > F(p)$ there exists a neighborhood U of p in $F_{<c}$, a positive constant $\eta > 0$, and a continuous map $\varphi\colon U \times [0,1] \to M$, which satisfies $\varphi(\cdot, 0) = \mathrm{id}_U$ and $\varphi(U, 1) \subseteq f_{\leq c-\eta}$, such that on every compact subset $V \subseteq U$ there exists a displacement function for φ as in Definition 7.2.1.

With this terminology in mind, we can now state the mountain pass theorem, formulated by Morse and Tompkins [MT39, Corollary 7.1] under slightly different assumptions, and with a slightly different conclusion. We provide more detailed comments regarding the differences between these assumptions and the conclusions of the mountain pass theorem in Subsections 7.2.2 and 8.2.1.

Theorem (Mountain pass theorem, Theorem 7.2.3, [MT39, Corollary 7.1])**.** *Let $F\colon M \to \mathbb{R}$ be a weakly upper-reducible function on a non-empty connected metric space with compact sublevel sets. Assume that the natural map* $\operatorname{colim} \check{H}_0(F_{\leq \bullet}) \to \check{H}_0(M)$ *is an isomorphism, and that $\check{H}_*(F_{\leq \bullet})$ is q-tame. If M contains two distinct critical sets of F of minimum type, then it also contains a critical set not of minimum type.*

It is worth noting that the details in [MT39] are incomplete, with some crucial theorems such as [MT39, Theorems 7.3 and 7.4, Corollary 7.1] being stated without proof, and with a citation to a paper in preparation that has never been published under the given name (we suppose that this paper is [Mor40]). Moreover, there is a gap in [Mor40], because [Mor40, Theorem 6.3], which establishes q-tameness, is incorrect as we will show in Corollary 8.2.3. We fix this by employing a similar but different condition that does indeed ensure q-tameness, which we discuss in § 2.2.2. This condition can be shown from the results of Morse and Tompkins [MT39] to be satisfied by the *Douglas functional*. From this the *unstable minimal surface theorem* can be deduced, showing the existence of a critical set not of minimum type for the Douglas functional (Theorem 7.2.14). In the intervening years, this result has been reproven and generalized in several directions using various techniques, and the problem class is still an active area of research [JS90; Jos91; MN21; Mon20; Str84].

Morse's work on functional topology did not have a long lasting impact on minimal surface theory or the calculus of variations in general; possibly in part because, as expressed by Struwe:

> The technical complexity and the use of a sophisticated topological machinery [...] tend to make Morse–Tompkins' original paper unreadable and inaccessible for the non-specialist. [Str88, p. 82]

A similar assessment was given by Bott, who writes in [Bot80, p. 934] that the papers [Mor37; Mor40] "are not easy reading" and constitute a "tour de force" by Morse. The intricacies of Morse's development notwithstanding, many of his ideas may be seen as early precursors to ideas that are now part of the standard framework of persistence theory.

Morse inequalities for cap numbers The most remarkable connections between functional topology and persistence theory come from Morse's paper [Mor40], where he developed the theory of *caps* and their *spans*. They capture much of the same information as the modern notion of persistence diagram, including concepts such as the persistence or birth and death of a homology class, although Morse's results still fall short of yielding global decompositions of persistence modules. Morse used his theory of caps to study functionals on a metric space by analyzing the evolution of the topology of their sublevel sets. A key tool to this end is a version of his eponymous inequalities for cap numbers, which expands their usual version in the compact and smooth setting. In this work, using persistence diagrams, we generalize the definition of these cap numbers to persistence modules and prove the existence of Morse inequalities for a large class of them.

Definition (Definition 7.1.1)**.** Let M be a graded q-tame persistence module with persistence diagram of the degree d part of M given by $\mathfrak{m}_d$. Whenever the sums below are well-defined, we define the (d,ϵ)*-cap numbers*

$$c_d^\epsilon = \alpha_d^\epsilon + \omega_{d-1}^\epsilon,$$

where

$$\alpha_d^\epsilon = \sum_{\substack{(p,q)\in\mathcal{E}\\ q-p>\epsilon\\ p\neq-\infty}} \mathfrak{m}_d(p,q), \qquad \omega_d^\epsilon = \sum_{\substack{(p,q)\in\mathcal{E}\\ q-p>\epsilon\\ q\neq\infty}} \mathfrak{m}_d(p,q)$$

are the *total number of births* and the *total number of deaths*, respectively, in degree d and with persistence greater than ϵ.

Comparing to the classical Morse inequalities, the cap numbers in dimension d act like the number of critical points with index d. As an analogue to the Betti numbers of the manifold appearing in the usual Morse inequalities, we define the *essential dimensions* p_d as

$$p_d = \sum_{p\in\mathbb{R}\cup\{-\infty\}} \mathfrak{m}_d(p,\infty),$$

which is also the dimension of the colimit of the degree d part of M.

Theorem (Theorem 7.1.2)**.** *Let $\epsilon > 0$, and let M be a graded q-tame persistence module with finite cap numbers c_d^ϵ and finite essential dimensions p_d for all d. If $\mathfrak{m}_d(-\infty,p) = 0$ for all $p \in \mathbb{R}\cup\{\infty\}$ and all d, then we have Morse inequalities*

$$\sum_{d=0}^{n} (-1)^{n-d}(c_d^\epsilon - p_d) \geq 0$$

for any dimension n.

2.2.2. Q-tameness for locally connected filtrations

Morse's local connectedness conditions Throughout his work on functional topology, in order to obtain q-tameness, Morse assumed slightly varying forms of local connectedness on the sublevel set filtrations of the functions he considered. In particular, Morse and Tompkins used the following condition from [Mor38; Mor40] in their applications to minimal surface theory:

> Let p be a point of M at which $F(p) = c$. The space M is said to be *locally F-connected* of order r at p if corresponding to each positive constant e there exists a positive constant δ such that each singular r-sphere on the δ-neighborhood of p and on $F_{c+\delta}$ bounds an $(r+1)$-cell of norm e on F_{c+e}. [Mor40, p. 431]

See also [Mor38, p. 25] and [MT39, p. 464], but note that the definitions given there contain evident typographical errors.

Morse then goes on to claim that the persistent Čech homology of the sublevel set filtration of a function F is q-tame, provided that F is bounded from below and satisfies the assumptions of local F-connectedness and compactness of sublevel sets. In the original (where the function is assumed to take values in $[0, 1)$) the claim reads:

> Let a and c be positive constants such that $a < c < 1$. The k^{th} connectivity $R^k(a, c)$ of F_a on F_c is finite. [Mor40, Theorem 6.3, p. 432]

Morse does not prove this statement in the given reference, but rather refers to [Mor38, Theorem 6.1]. Unfortunately, the above claim does not hold in general, as exemplified by the sublevel set filtration from Example 8.1.13, see also Corollary 8.2.3. The mountain pass theorem (Theorem 7.2.3) does, however, require q-tameness, so we want to develop a variation of the local connectedness condition above that is applicable in the minimal surface setting and indeed implies q-tameness.

Local connectedness revisited

Definition (Definition 8.1.2)**.** The sublevel set filtration of a function $f\colon X \to \mathbb{R}$ is called *locally homologically small* or LHS with respect to a homology theory H if for any $x \in X$, any neighborhood V of x, and any pair of indices s, t with $f(x) < s < t$ there is a neighborhood U of x with $U \subseteq V$ such that the inclusion $f_{\le s} \cap U \to f_{\le t} \cap V$ is *homologically small* or HS, i.e., has finite rank in every degree after applying H.

What we mean precisely by a *homology theory H* in this context is made more precise in § 8.1.1; importantly, our notion includes Čech homology with field coefficients on compact Hausdorff spaces, as well as any homology theory in the sense of the Eilenberg-Steenrod axioms.

Theorem (Theorem 8.1.4)**.** *If the sublevel set filtration of a function $F\colon X \to \mathbb{R}$ on a topological space X is* LHS *and consists of compact Hausdorff spaces, then its persistent homology is q-tame.*

We also introduce a weaker local-connectedness condition that can be used instead of LHS in the statement above if the filtration is defined by a continuous function (Corollary 8.1.12). The existence of a result of this kind has been suggested by Weinberger [Wei11], and a multiparameter version has been shown by Cagliari and Landi [CL11] with slightly stronger assumptions on the domain of the function.

To illustrate the suitability of our results in the context of minimal surface theory, we discuss why the Douglas functional has the LHS property for Čech homology (§ 8.2.2), so that its associated persistent Čech homology is indeed q-tame and the mountain pass theorem (Theorem 7.2.3) applies.

2.2.3. Comparison of singular and Čech homology in locally connected filtrations

Weak isomorphism from singular to Čech homology From the point of view of modern algebraic topology, the use of Čech homology in Morse's work on functional topology and in the mountain pass theorem (Theorem 7.2.3) might seem peculiar. We now discuss this in more detail, presenting an argument for why the use of Čech homology instead of the more common singular homology might not be necessary if one works with locally connected filtrations. As a counterpoint, we present an argument for why Čech homology might be useful after all in § 2.2.4. A general review of Čech homology can be found in Section 1.4.

Definition (Definition 9.3.3)**.** Let $f\colon X \to \mathbb{R}$ be a function on a topological space. We say that the sublevel set filtration $f_{\leq \bullet}$ induced by f is *homologically locally connected (*HLC*)* with respect to the homology theory H if for any $x \in X$, any neighborhood V of x, and any pair of indices s, t with $f(x) < s < t$ there is a neighborhood U of x with $U \subseteq V$ such that the inclusion $f_{\leq s} \cap U \to f_{\leq t} \cap V$ is taken to the trivial map by H.

Clearly, any filtration that is HLC is also LHS in the sense of Definition 8.1.2. By Theorem 8.1.4, we know that if the sublevel set filtration of f is LHS and all sublevel sets are compact Hausdorff, then the persistent homology of the sublevel set filtration is q-tame. There are also several results stating that the homology of a single compact Hausdorff space has finite rank if it is LHS (or HLC) (see [Bre97, Corollary II.17.7] for an example), where a single space X is called LHS (or HLC) if for any $x \in X$ and any neighborhood V of x there is a neighborhood U of x with $U \subseteq V$ such that the inclusion $U \to V$ induces a finite rank map (or the trivial map) in homology. Applying such finiteness results at every index, we see that the persistent homology of a filtration is PFD if all sublevel sets are locally connected in a suitable sense, and Theorem 8.1.4 may informally be seen as an observable (in the sense of [CCdS16]) version of such results.

Apart from finiteness results as the above one, local connectedness conditions play an important role in comparison results for different homology theories, see [Skl80] for a plethora of results in this direction. As a specific instance, Mardešić [Mar59] proved that paracompact Hausdorff spaces that are HLC with respect to singular homology have naturally isomorphic singular and Čech homology groups. Similarly to before, applying this to each sublevel set individually implies that the persistent singular and Čech homologies of a filtration are naturally isomorphic if the constituent sets are all HLC paracompact Hausdorff spaces. We now also provide an obserable version of the comparison result by Mardešić [Mar59].

Theorem (Theorem 9.3.4)**.** *If $f\colon X \to \mathbb{R}$ induces a sublevel set filtration of paracompact Hausdorff spaces that is* HLC *with respect to singular homology with coefficients in an abelian group G, then the natural map $\varphi\colon H_*(f_{\leq \bullet}; G) \to \check{H}_*(f_{\leq \bullet}; G)$ from its persistent singular to its persistent Čech homology is a weak isomorphism.*

The construction of the natural map φ from singular to Čech homology will be reviewed in § 9.1.1. Note that the persistent homologies appearing in the theorem need not be persistence modules in the sense that we have talked about before, because for general abelian coefficient groups G they will be functors from $\mathbb{R}$ to the category of G-modules and not necessarily to a category of vector spaces. However, the notion of weak isomorphism still makes sense in this more general setting and the result holds as stated. The interleaving distance between two persistence modules is 0 if there is a weak isomorphism from one to the other, so we also obtain the following corollary as a consequence of Theorem 9.3.4.

Corollary (Corollary 9.3.6). *If $f\colon X \to \mathbb{R}$ induces a sublevel set filtration of paracompact Hausdorff spaces that is locally connected with respect to singular homology with coefficients in an abelian group G, then*

$$d_I(H_*(f_{\leq \bullet}; G), \check{H}_*(f_{\leq \bullet}; G)) = 0.$$

Corollary 9.3.6 generalizes a recent result by Buhovsky et al. [Buh+22, Proposition 2.12], where f is assumed to be continuous and X is assumed to be a compact manifold. A persistence module that has 0 interleaving distance to a q-tame persistence module is again q-tame, so combining Theorem 9.3.4 with Theorem 8.1.4, we obtain that a compact filtration that is HLC with respect to singular homology is not only q-tame with respect to singular homology, but also with respect to Čech homology. Being q-tame, both persistence modules then admit a persistence diagram, and as a consequence of Corollary 9.3.6 the two persistence diagrams agree.

Local connectedness shift Our techniques will allow us to prove a stronger, quantitative version of Corollary 9.3.6. It bounds the interleaving distance between singular and Čech homology in terms of a measurement for how locally connected a filtration is.

Definition (Definition 9.1.4). Let $f\colon X \to \mathbb{R}$ be a function on a topological space. For $\delta \geq 0$, we say that the sublevel set filtration $f_{\leq \bullet}$ induced by f is *δ homologically locally connected ($\delta -$ HLC)* with respect to the homology theory H if for any $x \in X$, any neighborhood V of x, and any pair of indices s, t with $f(x) < s \leq s + \delta < t$ there is a neighborhood U of x with $U \subseteq V$ such that the inclusion $f_{\leq s} \cap U \to f_{\leq t} \cap V$ is taken to the trivial map by H. We define the *local connectedness shift of f with respect to H* as

$$\mathrm{lcs}_H(f) = \inf\{\delta > 0 \mid f \text{ is } \delta - \mathrm{HLC}\}.$$

Note that if the filtration induced by a function f is HLC, then its local connectedness shift is 0.

Theorem (Theorem 9.3.5). *If $f\colon X \to \mathbb{R}$ induces a sublevel set filtration of paracompact Hausdorff spaces and G is an abelian group, then*

$$d_I(H_{\leq d-1}(f_{\leq \bullet}; G), \check{H}_{\leq d-1}(f_{\leq \bullet}; G)) \leq d \cdot \mathrm{lcs}(f)$$

for all d, where $\mathrm{lcs}(f)$ is the local connectedness shift of f with respect to singular homology with coefficients in G.

Theorem 9.3.5 can be interpreted as providing a kind of stability theorem for when the function inducing the filtration is fixed but the homology theory varies from singular to Čech. Concerning stability with respect to variations of the function, we will also prove that the local connectedness shift lcs_H is 2-Lipschitz for any homology theory H (Proposition 9.3.7).

Our proofs of Theorems 9.3.4 and 9.3.5 closely follow the arguments presented in [Mar59] for a single space. Other approaches to proving our results could be to use an approximate nerve theorem [GS18] and analyze how the interleaving distance behaves with respect to inverse limits in the category of persistence modules, or to use an approach via cosheaves as in [Bre97], akin to the usual approach to comparison results between singular and Čech *co*homology via sheaf cohomology. In the present work, we do not investigate these possibilities further.

2.2.4. Structure of semi-continuous q-tame persistence modules

While Theorem 9.3.4 implies that singular and Čech homology have the same persistence diagrams for locally connected filtrations, Čech homology does have one advantage that is not visible in the persistence diagram: At least for filtrations consisting of compact Hausdorff spaces, persistent Čech homology is always upper semi-continuous.

Definition (Definition 10.1.1)**.** A persistence module M is called *upper semi-continuous (u.s.c.)* or *continuous from above* if the canonical map

$$M_t \to \lim_{s>t} M_s$$

is an isomorphism for all $t \in T$. It is called *lower semi-continuous (l.s.c.)* or *continuous from below* if the canonical map

$$\operatorname*{colim}_{s<t} M_s \to M_t$$

is an isomorphism for all $t \in T$.

Using the techniques developed in the context of the observable category by Chazal, Crawley-Boevey, and de Silva [CCdS16], we will show that, under some mild assumptions on the index set, semi-continuous q-tame persistence modules actually admit decompositions into interval modules up to isomorphism and not just weak isomorphism. While not explicitly stated by Chazal et al., the next result is an immediate corollary of [CCdS16, Corollary 3.6.]. The terms involving the index set were introduced in Section 1.3 and will also be recalled again in § 10.1.1. In the important special case $T = \mathbb{R}$, all assumptions are satisfied.

Theorem (Theorem 10.2.3)**.** *Let T be a dense totally ordered set such that every interval in T has a countable coinitial subset. Then every q-tame lower semi-continuous persistence module indexed by T has a barcode.*

With some additional work, we will prove the following novel result.

Theorem (Theorem 10.2.4)**.** *Let T be a dense totally ordered set such that every interval in T has a countable coinitial subset. Then for every q-tame upper semi-continuous persistence module M indexed by T there exists a collection of intervals $(I_\alpha)_{\alpha \in A}$, unique up to reordering, such that*

$$M \cong \prod_{\alpha \in A} C(I_\alpha).$$

In general, uniqueness statements for product decompositions are much harder to come by than in the case of direct sums where one has the Krull–Remak–Schmidt–Azumaya Theorem. We will also infer our uniqueness statement from this theorem, rather than from a general statement about products. In order to distinguish product and direct sum decompositions, when the distinction is not clear from the context, we suggest to called barcodes in the usual sense *additive barcodes*, while barcodes in the sense of Theorem 10.2.4 can be called *multiplicative barcodes*. In the PFD case, the two notions agree. In the q-tame case, however, there are persistence modules that have a multiplicative barcode, but no additive barcode and vice versa (Example 10.1.3).

2.3. Outline

The rest of this thesis will be structured like this chapter presenting our results was. Concretely, we have two more parts, with Part II corresponding to Section 2.1 and Part III corresponding to Section 2.2.

Part II will start with Chapter 3 recalling some more general persistence theory from the literature, in particular regarding barcodes and matching diagrams, expanding our treatment from Section 1.1. After this background chapter, we continue Part II with Chapter 4 corresponding to § 2.1.1, where the lifespan functors are introduced in more detail and the main theorems for the general theory are proven. We follow up with Chapter 5, where the injectivity and projectivity results, as well as the duality results stated in § 2.1.2 are proven and expanded on. We finish Part II in Chapter 6 by proving correctness of the method for efficiently computing image persistence outlined in § 2.1.3.

For Part III, we start with Chapter 7, where we present our persistence-flavored account of the Morse inequalities and the unstable minimal surface theorem previewed in § 2.2.1. In Chapter 8, we then discuss some shortcomings of Morse's original approach to functional topology and prove our theorem from § 2.2.2 that can be used to fix these shortcomings. Relating to the use of Čech homology in Morse's work, we then prove our comparison results with singular homology in Chapter 9, which completes the discussion from § 2.2.3. Finally, we prove our structural results for semi-continuous q-tame persistence modules from § 2.2.4 in Chapter 10.

2.4. Notes on collaborations and publications related to this thesis

As mentioned, most of the results in this thesis have previously appeared in the articles and preprints [BMS21; BS21a; BS22; Sch22a; Sch22b]. Large parts of this thesis also appear verbatim in these articles. We will now describe this in more detail.

The results in Section 2.1 and the work in Part II have previously appeared in two articles written in collaboration with Ulrich Bauer [BS21a; BS22], who initiated the projects and played an advisory role in them. In terms of content overlap, we have the following.

- Chapter 3 roughly corresponds to [BS21a, Section 2],
- Chapter 4 and § 2.1.1 roughly correspond to [BS21a, Sections 1, 3, and 4],
- Chapter 5 and § 2.1.2 roughly correspond to [BS21a, Sections 5 and 6], and

- Chapter 6 and § 2.1.3 roughly correspond to [BS22].

The results in Section 2.2 and the work in Part III have previoulsy appeared in three articles [BMS21; Sch22a; Sch22b], the first of which was written in collaboration with Ulrich Bauer and Anibal Medina-Mardones. These results are also partly a continuation of the author's master's thesis. Specifically, Theorem 7.1.5 has already appeared there, as have prototypical versions of Theorem 7.1.2 and Proposition 7.2.12. Those results that have previously appeared in the author's master's thesis are not essential parts of the present thesis and are mostly presented for completeness of the exposition and for the convenience of the reader. Ulrich Bauer supervised the author's master's thesis and again initiated and advised the collaborative project [BMS21]. In terms of content overlap, we have the following.

- Chapter 7 and § 2.2.1 roughly correspond to [BMS21, Sections 1, 3 and 5.1, Appendix B],
- Chapter 8 and § 2.2.2 roughly correspond to [BMS21, Sections 4 and 5.2],
- Chapter 9 and § 2.2.3 roughly correspond to [Sch22a], and
- Chapter 10 and § 2.2.4 roughly correspond to [Sch22b].

Part II.

Lifespan functors and applications

3. More background on persistence theory

In this chapter we augment the introduction to persistence theory from Chapter 1 with some more specific background knowledge that is only needed for the present part of the thesis. We start by defining the categories of barcodes and matching diagrams, reviewing the equivalence between them, and showing that they have a p-exact structure in Section 3.1. In Section 3.2, we then discuss the passage from these categories to the category of persistence modules.

For this chapter, we do not put any additional constraints on our totally ordered index set T.

3.1. Categories of barcodes and matching diagrams

We now review the interpretation and categorification of barcodes via matching diagrams following [BL15; BL20]. To start off, we define the categories of barcodes and matching diagrams in § 3.1.1. In § 3.1.2, we review the equivalence between these two categories. § 3.1.3 contains a brief recollection of what constitutes a p-exact category, and we recall that barcodes and matching diagrams have this structure.

3.1.1. Barcodes and matching diagrams

Barcodes revisited We have seen in § 1.1.2 that a persistence module M is said to have a barcode $(I_\alpha)_{\alpha\in A}$ if M is isomorphic to the direct sum over the interval modules $C(I_\alpha)$, and that the collection of intervals is unique up to a choice of the index set A. Making this more formal, we will now introduce a framework where the choice of A is encoded as part of the barcode.

Definition 3.1.1. We denote the set of all intervals in T as $\mathfrak{I}(T)$, or simply as $\mathfrak{I}$ when the index set is clear from the context. If A is an arbitrary set, we call any subset $B \subseteq \mathfrak{I} \times A$ a *barcode in* T.

The purpose of the set A in this definition is to distinguish multiple instances of the same interval, as in the standard construction of a disjoint union. If clear from the context, we sometimes suppress this index from the notation. If $B \subseteq \mathfrak{I} \times A$ is a barcode and I an interval in T, the cardinality of the set $\{a \in A \mid (I, a) \in B\}$ measures how many copies of I are in B, similar to the definition of persistence diagrams in § 1.3.3, but remembering the whole interval and not just its endpoints.

Overlap matchings

Definition 3.1.2. If A and B are sets, a subset $\sigma \subseteq A \times B$ is called a *matching* if for each $a \in A$ there is at most one $b \in B$ with $(a, b) \in \sigma$ and for each $b \in B$ there is at most one $a \in A$ with $(a, b) \in \sigma$.

When introducing the bottleneck distance between barcodes in § 1.1.4, we have already talked about certain matchings between barcodes, namely δ-matchings. Barcodes form a category in a way that is compatible with the passage to persistence modules, with the morphisms in the category of barcodes being certain other matchings.

Definition 3.1.3. Let I and J be intervals in T. We say that I *bounds* J *above* if for all $s \in J$ there exists $t \in I$ such that $s \leq t$. We say that I *bounds* J *below* if for all $u \in J$ there exists $t \in I$ such that $t \leq u$. We say that I *overlaps* J *above*, or that J *overlaps* I *below*, if their intersection is non-empty, I bounds J above, and J bounds I below.

With this terminology in mind, we can now define the category of barcodes.

Definition 3.1.4. For barcodes B and B', we call a matching $\sigma \subseteq B \times B'$ an *overlap matching* if for each $((I, a), (I', a')) \in \sigma$ the interval I overlaps the interval I' above. If $\sigma \subseteq B \times B'$ and $\tau \subseteq B' \times B''$ are overlap matchings, we define their *overlap composition* as

$$\begin{aligned} \tau \bullet \sigma = \{((I, a), (I'', a'')) \mid & \text{ there exists } (I', a') \in B' \text{ with} \\ & ((I, a), (I', a') \in \sigma, \\ & ((I', a'), (I'', a'') \in \tau, \text{ and} \\ & I \text{ overlaps } I'' \text{ above}\}. \end{aligned}$$

The resulting category with barcodes as objects, overlap matchings as morphisms and overlap composition will be denoted by $\mathbf{Barc}(\mathbf{T})$.

Note that two barcodes $B \subseteq \mathfrak{I} \times A$ and $B' \subseteq \mathfrak{I} \times A'$ are isomorphic if and only if there is a bijection $f \colon B \to B'$ such that for all $(I, a) \in B$ there is a' in A' with $f(I, a) = (I, a')$. In other words, B and B' are isomorphic if and only if the sets $\{a \in A \mid (I, a) \in B\}$ and $\{a' \in A' \mid (I, a') \in B'\}$ have the same cardinality for every interval $I \in \mathfrak{I}$.

Matching diagrams Apart from the requirement that I needs to overlap I'' above, the composition law from Definition 3.1.4 also makes sense for matchings between general sets and not just barcodes. This leads to the category of matching diagrams.

Definition 3.1.5. If A, B, C are sets and $\sigma \subseteq A \times B$, $\tau \subseteq B \times C$ are matchings, we define the composition $\tau \circ \sigma \subseteq A \times C$ as

$$\tau \circ \sigma = \{(a, c) \mid \text{ there exists } b \in B \text{ with } (a, b) \in \sigma \text{ and } (b, c) \in \tau\}.$$

The resulting category, with sets as objects, matchings as morphisms, and the above composition, will be denoted by $\mathbf{Mch}$. We define the category of *matching diagrams indexed by* $\mathbf{T}$ as the category $\mathbf{Mch}^{\mathbf{T}}$ of functors $\mathbf{T} \to \mathbf{Mch}$.

Note that the standard composition from Definition 3.1.5 and the overlap composition Definition 3.1.4 need not yield the same result when applied to overlap matchings between barcodes. This is due to the fact that if I overlaps I' above and I' overlaps I'', it need not be true that I overlaps I'' above. More specifically, the bounding conditions for I to overlap I'' above will be satisfied, but what may happen is that $I \cap I'' = \emptyset$. As an example where this occurs, one may consider the real intervals $I = [2, 3]$, $I' = [1, 2]$, and $I'' = [0, 1]$.

3.1.2. Equivalence of barcodes and matching diagrams

As we have mentioned before, the two categories $\mathbf{Mch}^{\mathbf{T}}$ and $\mathbf{Barc}(\mathbf{T})$ are equivalent. We will now review the construction of an explicit equivalence following [BL20].

From matching diagrams to barcodes

Definition 3.1.6. Let D be a matching diagram. We define its *components* as the set of equivalence classes

$$\mathcal{C}(D) = \left(\bigcup_{t \in T} \{t\} \times D_t \right) / \sim,$$

where the equivalence relation $\sim$ is defined as follows: for $t \leq u \in T$, $d \in D_t$, and $d' \in D_u$, we set $(t, d) \sim (u, d')$ if and only if $(d, d') \in D_{t,u}$. Note that each component $Q \in \mathcal{C}(D)$ can also be regarded as a matching diagram such that $Q_t \subseteq D_t$ has at most one element for each $t \in T$. For a component $Q \in \mathcal{C}(D)$, we define its *support* as the range of indices in T spanned by the component,

$$\operatorname{supp}(Q) = \{t \in T \mid (t, d) \in Q \text{ for some } d \in D_t\}.$$

Note that the component set construction may be regarded as a functor $\mathcal{C} \colon \mathbf{Mch}^{\mathbf{T}} \to \mathbf{Mch}$. Even more than that, it can not only be used to pass from matching diagrams to the matching category, but it can also be used to pass from matching diagrams to barcodes.

Definition 3.1.7. We define a functor $\mathcal{B} \colon \mathbf{Mch}^{\mathbf{T}} \to \mathbf{Barc}(\mathbf{T})$ by setting

$$\mathcal{B}(D) = \{(I, Q) \in \mathfrak{I} \times \mathcal{C}(D) \mid I = \operatorname{supp}(Q)\}$$

for any matching diagram D and

$$\mathcal{B}(\psi) = \{((I, Q), (I', R)) \in \mathcal{B}(D) \times \mathcal{B}(E) \mid Q_t \times R_t \subseteq \psi_t \text{ for all } t \in I \cap I'\}$$

for any morphism of matching diagrams $\psi \colon D \to E$.

As shown in [BL20], the support of a component is indeed an interval, a morphism of matching diagrams is mapped to an overlap matching by the above construction, and we indeed get a functor.

From barcodes to matching diagrams Conversely to the previous construction, we can also pass from barcodes to matching diagrams.

Definition 3.1.8. We define a functor $\mathcal{D} \colon \mathbf{Barc}(\mathbf{T}) \to \mathbf{Mch}^{\mathbf{T}}$ by setting $\mathcal{D}(B)$ for any barcode B to be the matching diagram given by

$$\mathcal{D}(B)_t = \{(I, a) \in B \mid t \in I\},$$
$$\mathcal{D}(B)_{t,u} = \{((I, a), (I', a')) \in \mathcal{D}(B)_t \times \mathcal{D}(B)_u \mid (I, a) = (I', a')\}.$$

For an overlap matching σ, we let $\mathcal{D}(\sigma)$ be the morphism of matching diagrams with

$$\mathcal{D}(\sigma)_t = \{((I, a), (I', a')) \in \sigma \mid t \in I \cap I'\}.$$

Again, we refer to [BL20] for the fact that $\mathcal{D}$ is a well-defined functor.

Theorem 3.1.9 ([BL20])**.** *The functors $\mathcal{B} \colon \mathbf{Mch}^{\mathbf{T}} \to \mathbf{Barc}(\mathbf{T})$ and $\mathcal{D} \colon \mathbf{Barc}(\mathbf{T}) \to \mathbf{Mch}^{\mathbf{T}}$ defined above are quasi-inverses. In particular, the categories $\mathbf{Mch}^{\mathbf{T}}$ and $\mathbf{Barc}(\mathbf{T})$ are equivalent.*

Note that in [BL20], the equivalences were denoted by E and F.

3.1.3. P-exactness for barcodes and matching diagrams

P-exact categories The categories of barcodes and matching diagrams have a certain property that allows one to do homological algebra in them.

Definition 3.1.10. A category is called *Puppe-exact* or *p-exact* if it has a zero object, it has all kernels and cokernels, every mono is a kernel and every epi is a cokernel, and every morphism has an epi-mono-factorization.

Put informally, a Puppe-exact category is an abelian category that need not have (co)products. For an in-depth look at p-exactness we refer to [Gra12]. Recall that in any category with kernels and cokernels, monos have vanishing kernels and epis have vanishing cokernels. While the converse is not true in general, it is true in p-exact categories.

Lemma 3.1.11 ([BP69, Korollar 2.4.4])**.** *A morphism in a p-exact category is mono if and only if its kernel vanishes and it is epi if and only if its cokernel vanishes.*

We will use Lemma 3.1.11 throughout without explicit reference.

Proposition 3.1.12 ([Gra12, Section 1.6.4])**.** **Mch** *is Puppe-exact. For a matching* $\sigma \subseteq A \times B$ *we have*

$$\begin{aligned} \ker \sigma &= \{a \in A \mid (a,b) \notin \sigma \textit{ for all } b \in B\}, \\ \operatorname{im} \sigma &= \{b \in B \mid (a,b) \in \sigma \textit{ for some } a \in A\}, \\ \operatorname{coker} \sigma &= \{b \in B \mid (a,b) \notin \sigma \textit{ for all } a \in A\}, \\ \operatorname{coim} \sigma &= \{a \in A \mid (a,b) \in \sigma \textit{ for some } b \in B\}. \end{aligned}$$

$\mathbf{Mch}^{\mathbf{T}}$ *is also Puppe-exact, with kernels, cokernels etc. given pointwise.*

Using the equivalence between $\mathbf{Mch}^{\mathbf{T}}$ and $\mathbf{Barc}(\mathbf{T})$, we can translate the constructions in Proposition 3.1.12 to describe the kernels, cokernels, and images of overlap matchings explicitly as barcodes.

Definition 3.1.13. For an overlap matching $\sigma \subseteq B \times B'$ and $(I,a) \in B$, $(I',a') \in B'$, we set

$$\ker(\sigma,(I,a)) = \begin{cases} (I \setminus I', a) & \text{if } ((I,a),(I',a')) \in \sigma, \\ (I,a) & \text{otherwise;} \end{cases}$$
$$\operatorname{coker}(\sigma,(I',a')) = \begin{cases} (I' \setminus I, a') & \text{if } ((I,a),(I',a')) \in \sigma, \\ (I',a') & \text{otherwise.} \end{cases}$$

Proposition 3.1.14 ([BL20])**.** *Let* $B \subseteq \mathfrak{I} \times A$ *and* $B' \subseteq \mathfrak{I} \times A'$ *be barcodes. Any overlap matching* $\sigma \subseteq B \times B'$ *has a kernel, coimage, image and cokernel in* $\mathbf{Barc}(\mathbf{T})$*, with*

$$\begin{aligned} \ker \sigma &= \{(J,a) \in \mathfrak{I} \times A \mid J = \ker(\sigma,(I,a)) \textit{ for } (I,a) \in B\} \\ \operatorname{coim} \sigma &= \{(J,a) \in \mathfrak{I} \times A \mid J = I \cap I' \textit{ for } ((I,a),(I',a')) \in \sigma\} \\ \operatorname{im} \sigma &= \{(J,a') \in \mathfrak{I} \times A' \mid J = I \cap I' \textit{ for } ((I,a),(I',a')) \in \sigma\} \\ \operatorname{coker} \sigma &= \{(J,a') \in \mathfrak{I} \times A' \mid J = \operatorname{coker}(\sigma,(I',a')) \textit{ for } (I',a') \in B'\}. \end{aligned}$$

3.2. From barcodes and matching diagrams to persistence modules

Of course, the main motivation for the introduction of the formal framework for barcodes and matching diagrams from Section 3.1 is their relationship to persistence modules. To elaborate on this, we briefly discuss the natural functors from the categories of barcodes and matching diagrams to the category of persistence modules in § 3.2.1. We then go on to review how dualization works in all of these categories in § 3.2.2.

3.2.1. Functors to persistence modules

Functor from barcodes to persistence modules With the overlap composition in $\mathbf{Barc}(\mathbf{T})$ being defined as in Definition 3.1.4, we can make the passage from barcodes to persistence modules into a functor.

Definition 3.2.1. If I and J are intervals such that I overlaps J above, there exists a canonical morphism on interval modules $\varphi(I,J)\colon C(I) \to C(J)$ defined by

$$\varphi(I,J)_t = \begin{cases} \mathrm{id}_{\mathbb{F}} & \text{if } t \in I \cap J, \\ 0 & \text{otherwise.} \end{cases}$$

We define the *barcode module* functor $\mathcal{M}\colon \mathbf{Barc}(\mathbf{T}) \to \mathbf{Pers}$ by sending a barcode B to the direct sum of interval modules $\bigoplus_{(I,a)\in B} C(I)$ and sending an overlap matching $\sigma \subseteq B \times B'$ to the direct sum of the morphisms $\varphi(I,I')\colon C(I) \to C(I')$ for all pairs $((I,a),(I',a')) \in \sigma$. If a persistence module M satisfies $\mathcal{M}(B) \cong M$ for some barcode $B \in \mathbf{Barc}(\mathbf{T})$, we say that B is a *barcode of* M.

Note that, in contrast to the usual language, we only talk about B being *a* barcode of M if $\mathcal{M}(B) \cong M$, because B is of course not unique. As discussed before, however, B is unique up to a choice of its index set, i.e., it is unique up to isomorphism in $\mathbf{Barc}(\mathbf{T})$.

Functor from matching diagrams to persistence modules The aforementioned functor from matching diagrams to persistence modules can now be defined as follows.

Definition 3.2.2. We define the functor $\mathcal{F}\colon \mathbf{Mch} \to \mathbf{Vec}$ by sending a set A to the free vector space generated by A and sending a matching $\sigma \subseteq A \times B$ to the linear extension of the map

$$a \mapsto \begin{cases} b & \text{if } (a,b) \in \sigma \\ 0 & \text{otherwise.} \end{cases}$$

By a slight abuse of notation, we also define the *matching module* functor $\mathcal{F}\colon \mathbf{Mch}^{\mathbf{T}} \to \mathbf{Pers}$ by applying $\mathcal{F}$ pointwise, i.e., $\mathcal{F}(D) = \mathcal{F} \circ D$.

From the definitions, one easily checks the following proposition.

Proposition 3.2.3. *There are natural isomorphisms* $\mathcal{F} \cong \mathcal{M} \circ \mathcal{B}$ *and* $\mathcal{M} \cong \mathcal{F} \circ \mathcal{D}$.

Importantly, the barcode and matching module functors translate the p-exact structure on barcodes and matching diagrams to p-exact structure on persistence modules, inherited by their abelian structure. That is, we have the following straightforward result.

Proposition 3.2.4. *The functor $\mathcal{F}$ preserves and reflects exactness, i.e., a sequence of matchings $V \to V' \to V''$ is exact if and only if the corresponding sequence of vector spaces $\mathcal{F}(V) \to \mathcal{F}(V') \to \mathcal{F}(V'')$ is exact. The same holds for $\mathcal{F}$ as a functor $\mathbf{Mch}^{\mathbf{T}} \to \mathbf{Pers}$ and for $\mathcal{M}\colon \mathbf{Barc}(\mathbf{T}) \to \mathbf{Pers}$.*

3.2.2. Dualization

If $\sigma \subseteq A \times B$ is a matching, we define its *opposite* matching

$$\sigma^\circ = \{(b, a) \mid (a, b) \in \sigma\} \subseteq B \times A.$$

This construction makes the category **Mch** self-dual, i.e., it yields an isomorphism between **Mch** and its opposite category, mapping every set to itself and every matching to its opposite matching. Taking opposite matchings at every index also yields a dualization construction for matching diagrams, and it also makes $\mathbf{Mch}^{\mathbf{T}}$ self-dual. As a consequence of the equivalence between $\mathbf{Mch}^{\mathbf{T}}$ and $\mathbf{Barc}(\mathbf{T})$, this also implies that $\mathbf{Barc}(\mathbf{T})$ is equivalent to its opposite category. Note, however, that in $\mathbf{Barc}(\mathbf{T})$, the opposite matching of an overlap matching need not be an overlap matching again in general, so the self-duality does not stem from the same construction.

Barcodes over the opposite poset and dualization of persistence modules One can not only compare $\mathbf{Barc}(\mathbf{T})$ to its opposite category $\mathbf{Barc}(\mathbf{T})^{\mathrm{op}}$, but also to the category of barcodes with respect to the opposite poset $(T, \geq)$, ie., to $\mathbf{Barc}(\mathbf{T}^{\mathrm{op}})$. To do so, note that a subset $I \subseteq T$ is an interval with respect to $\leq$ if and only if it is an interval with respect to $\geq$, and that an interval I overlaps and interval J above with respect to $\leq$ if and only if J overlaps I above with respect to $\geq$. This yields an obvious contravariant isomorphism between $\mathbf{Barc}(\mathbf{T})$ and $\mathbf{Barc}(\mathbf{T}^{\mathrm{op}})$ which maps each barcode to itself. Thus, we can compare barcodes of persistence modules indexed by $(T, \leq)$ with barcodes of persistence modules indexed by $(T, \geq)$. Algebraically, the following construction yields a method for comparing the two.

Definition 3.2.5. We define the contravariant *dualization* functor $(-)^\vee\colon \mathbf{Vec}^{\mathbf{T}} \to \mathbf{Vec}^{\mathbf{T}^{\mathrm{op}}}$ by applying vector space dualization pointwise, i.e., for a $\mathbf{T}$-indexed persistence module M, its *dual* $M^\vee$ is the $\mathbf{T}^{\mathrm{op}}$-indexed persistence module given by $M^\vee_t = \mathrm{Hom}_{\mathbf{Vec}}(M_t, \mathbb{F})$ for all $t \in T$.

We have the following well-known fact.

Lemma 3.2.6. *Let M be a* PFD *persistence module. Then B is a barcode for M if and only if it is a barcode for $M^\vee$.*

Note, however, that Lemma 3.2.6 does not extend beyong the PFD setting since no infinite-dimensional vector space is isomorphic to its dual.

Exactness of dualization Recall that $\mathbb{F}$ is injective as a module over itself, which means that the contravariant functor $\mathrm{Hom}_{\mathbf{Vec}}(-, \mathbb{F})\colon \mathbf{Vec} \to \mathbf{Vec}$ is exact. As pointwise application of an exact functor yields an exact functor of diagram categories, we get the following.

Lemma 3.2.7. *The dualization functor $(-)^\vee$ is exact. In particular, a morphism $\varphi\colon M \to N$ of persistence modules yields isomorphisms*

$$(\ker \varphi)^\vee \cong \operatorname{coker} \varphi^\vee, \qquad (\operatorname{im} \varphi)^\vee \cong \operatorname{im} \varphi^\vee, \qquad (\operatorname{coker} \varphi)^\vee \cong \ker \varphi^\vee.$$

4. Introducing lifespan functors

In this chapter, we will define the lifespan functors and prove the main result presented in § 2.1.1, as well as the result on lifespan functors and images described in § 2.1.2. We start with the definitions and the main properties of lifespan functors in a very general setting in Section 4.1. The special case of persistence modules and barcodes is then discussed in Section 4.2.

For this chapter, we do not put any additional constraints on our totally ordered index set T.

4.1. Lifespan in p-exact diagram categories

We will now start our introduction to lifespan functors. As the first step, we will define lifespan functors for very general diagrams in § 4.1.1. We then analyze some exactness properties of these functors, i.e., whether and when they commute with taking kernels and cokernels of morphisms in § 4.1.2.

4.1.1. Defining lifespan functors

Let $\mathbf{A}$ be any category with $\mathbf{T}$-shaped limits and colimits, so that we get functors $\lim\colon \mathbf{A}^{\mathbf{T}} \to \mathbf{A}$ and $\operatorname{colim}\colon \mathbf{A}^{\mathbf{T}} \to \mathbf{A}$, where $\mathbf{A}^{\mathbf{T}}$ denotes the category of functors $\mathbf{T} \to \mathbf{A}$. As for any functor category, we also have a diagonal functor $\Delta\colon \mathbf{A} \to \mathbf{A}^{\mathbf{T}}$, mapping each object to the corresponding constant diagram. Of course, this setting includes the case where $\mathbf{A} = \mathbf{Vec}$.

For each object V in $\mathbf{A}^{\mathbf{T}}$, the canonical maps $V_t \to \operatorname{colim} V$ for $t \in T$ form a natural transformation $\eta_V\colon V \to \Delta \operatorname{colim} V$. Recall that colim is left adjoint to the diagonal functor Δ, and the morphism η_V is the component at V for the unit $\eta\colon \mathrm{id}_{\mathbf{A}^{\mathbf{T}}} \to \Delta \circ \operatorname{colim}$ of the adjunction $\operatorname{colim} \dashv \Delta$. Similarly, the canonical maps $\lim V \to V_t$ give a natural transformation $\epsilon_V\colon \Delta \lim V \to V$, which is the counit $\epsilon\colon \Delta \circ \lim \to \mathrm{id}_{\mathbf{A}^{\mathbf{T}}}$ of the adjunction $\Delta \dashv \lim$. We thus get the diagram

$$\Delta \lim V \xrightarrow{\epsilon_V} V \xrightarrow{\eta_V} \Delta \operatorname{colim} V.$$

From now on, we assume that $\mathbf{A}$ is Puppe-exact, so that we can form kernels, cokernels, and images of morphisms in $\mathbf{A}$ and $\mathbf{A}^{\mathbf{T}}$. For a brief review of Puppe-exactness, see § 3.1.3 and for an in-depth review see the references therein.

Definition 4.1.1. We define the following functors $\mathbf{A}^{\mathbf{T}} \to \mathbf{A}^{\mathbf{T}}$.

1. The *mortal part* functor is defined as $(-)^\dagger = \ker \eta_{(-)}$.
2. The *immortal part* functor is defined as $(-)^\infty = \operatorname{im} \eta_{(-)}$.
3. The *nascent part* functor is defined as $(-)^* = \operatorname{coker} \epsilon_{(-)}$.
4. The *ancient part* functor is defined as $(-)^{-\infty} = \operatorname{im} \epsilon_{(-)}$.

By definition, for each object V in $\mathbf{A^T}$ we get a natural diagram

$$\begin{array}{ccccc} V^{\dagger} & & & & V^{*} \\ & \searrow & & \nearrow & \\ & & V & & \\ & \nearrow & & \searrow & \\ V^{-\infty} & & & & V^{\infty} \end{array}$$

with diagonal short exact sequences. We also get composite natural transformations

$$(-)^{\dagger} \to \mathrm{id}_{\mathbf{A^T}} \to (-)^{*} \qquad \text{and} \qquad (-)^{-\infty} \to \mathrm{id}_{\mathbf{A^T}} \to (-)^{\infty}$$

and can again form kernels, cokernels and images to get new functors.

Definition 4.1.2. We define the following functors $\mathbf{A^T} \to \mathbf{A^T}$.

1. The *finite part* functor is defined as $(-)^{\dagger,*} = \mathrm{im}((-)^{\dagger} \to (-)^{*})$.
2. The *constant part* functor is defined as $(-)^{-\infty,\infty} = \mathrm{im}((-)^{-\infty} \to (-)^{\infty})$.

Remark 4.1.3. The universal property of epi-mono-factorizations implies that we have a canonical isomorphism $V^{-\infty,\infty} \cong \mathrm{im}(\Delta \lim V \to \Delta \operatorname{colim} V)$ for all objects V in $\mathbf{A^T}$.

We will also form kernels and cokernels of the above composite morphisms. In the cases we are interested in, these turn out to coincide: by [Gra12, Lemma 2.2.4], pullbacks of monos and pushouts of epis exist in p-exact categories, and we have canonical isomorphisms

$$\ker(V^{\dagger} \to V^{*}) \cong V^{\dagger} \times_V V^{-\infty} \cong \ker(V^{-\infty} \to V^{\infty})$$
$$\mathrm{coker}(V^{\dagger} \to V^{*}) \cong V^{*} +_V V^{\infty} \cong \mathrm{coker}(V^{-\infty} \to V^{\infty})$$

for any V. Using this fact, we can make the following well-posed definition.

Definition 4.1.4. We define the following functors $\mathbf{A^T} \to \mathbf{A^T}$.

1. The *ancient mortal part* functor is defined as
$$(-)^{-\infty,\dagger} = \ker((-)^{\dagger} \to (-)^{*}) = \ker((-)^{-\infty} \to (-)^{\infty}).$$
2. The *immortal nascent part* functor is defined as
$$(-)^{*,\infty} = \mathrm{coker}((-)^{\dagger} \to (-)^{*}) = \mathrm{coker}((-)^{-\infty} \to (-)^{\infty}).$$

We give a common name to all the functors defined above.

Definition 4.1.5. For an object V in $\mathbf{A^T}$, we will call the diagram

$$\begin{array}{ccccccccc} & & & & V^{\dagger,*} & & & & \\ & & & \nearrow & & \searrow & & & \\ & & V^{\dagger} & & & & V^{*} & & \\ & \nearrow & & \searrow & & \nearrow & & \searrow & \\ V^{-\infty,\dagger} & & & & V & & & & V^{*,\infty} \\ & \searrow & & \nearrow & & \searrow & & \nearrow & \\ & & V^{-\infty} & & & & V^{\infty} & & \\ & & & \searrow & & \nearrow & & & \\ & & & & V^{-\infty,\infty} & & & & \end{array}$$

the *lifespan diagram of* V. We call the functors at the nodes of the diagram *lifespan functors* and the natural maps between them *lifespan transformations*. We use the notation $(-)^\diamond$ as a wildcard symbol for an arbitrary lifespan functors.

See Figure 2.1 for an example of a lifespan diagram of persistence modules.

Vanishing of lifespan functors Note that the lifespan diagram simplifies to a smaller diagram in many applications. For example, the short exact sequence $V^{-\infty,\dagger} \hookrightarrow V^{-\infty} \twoheadrightarrow V^{-\infty,\infty}$ on the bottom left vanishes if V is bounded below in the sense that there exists some $t \in T$ with $V_s = 0$ for all $s \leq t$. Similarly, the bottom right sequence vanishes if V is bounded above in the sense that there exists some $t \in T$ with $V_u = 0$ for all $u \geq t$. For the top left and the top right short exact sequences in the lifespan diagram, we have the following conditions.

Proposition 4.1.6. *Consider* $\lim, \operatorname{colim} \colon \mathbf{A}^\mathbf{T} \to \mathbf{A}$ *and an object* V *in* $\mathbf{A}^\mathbf{T}$.

1. *If* colim *is exact, then* $V^\dagger = 0$ *if and only if all structure maps of* V *are mono.*

2. *If* lim *is exact, then* $V^* = 0$ *if and only if all structure maps of* V *are epi.*

Proof. We only show the first statement since the second one is dual to it. So, assume that taking colimits is exact.

If $V^\dagger = 0$, then $V \to \Delta \operatorname{colim} V$ is mono, i.e., $V_t \to \operatorname{colim} V$ is mono for any $t \in T$. Now, for any structure map $V_t \to V_u$, we obtain that the composition $V_t \to V_u \to \operatorname{colim} V$ is mono since it is equal to the natural map $V_t \to \operatorname{colim} V$. This implies that $V_t \to V_u$ is mono.

Next, assume that all structure maps of V are mono and let $t \in T$. Define an object $\tilde{V}$ in $\mathbf{A}^\mathbf{T}$ by setting $\tilde{V}_s = V_s$ for any $s < t$ and $\tilde{V}_u = V_t$ for any $u \geq t$. There is an obvious map $\tilde{V} \to V$ consisting of structure maps of V and because we assume these structure maps to be mono the map $\tilde{V} \to V$ is mono, too. We assume that taking colimits is exact, so the induced map $\operatorname{colim} \tilde{V} \to \operatorname{colim} V$ is still mono. But $\operatorname{colim} \tilde{V} = V_t$ and the induced map is given by the natural map $V_t \to \operatorname{colim} V$. Hence, $V \to \Delta \operatorname{colim} V$ is mono, which implies $V^\dagger = 0$. ☐

Complementary functors The construction of the lifespan functors involves kernels, cokernels, and images of the natural transformations ϵ and η. Note, however, that we have not used $\ker \epsilon_{(-)}$ and $\operatorname{coker} \eta_{(-)}$ so far. These play a somewhat different role than the lifespan functors, as they do not yield subobjects or quotients of the object we start with. Still, their properties will be of similar importance, specifically, in Theorem 5.3.2.

Definition 4.1.7. We define the following functors $\mathbf{A}^\mathbf{T} \to \mathbf{A}^\mathbf{T}$.

1. The *ghost complement* functor is defined as $(-)^\triangleright = \ker \epsilon_{(-)}$.

2. The *unborn complement* functor is defined as $(-)^\triangleleft = \operatorname{coker} \eta_{(-)}$.

See Figure 4.1 for an illustration of the complement functors.

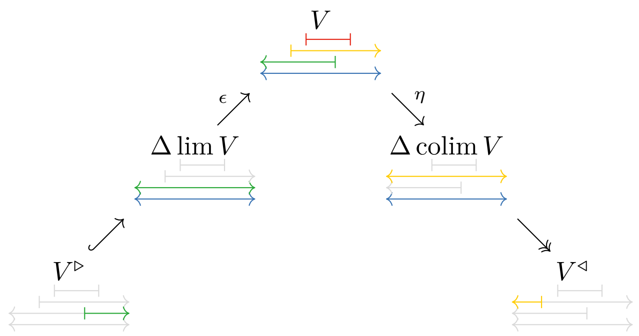

Figure 4.1.: Complement functors applied to a finite type $\mathbb{R}$-indexed persistence module V, visualized via their barcode according to Propositions 4.2.9 and 4.2.10.

4.1.2. Lifespan functors and images

One of our overall goals is to study images of morphisms in persistent homology. For that purpose, we want to study how the lifespan functors appearing in the statement of Theorem 5.3.2 behave with respect to this operation. The relevant theorems hold in the general setting, so, as before, let $\mathbf{A}$ be p-exact with $\mathbf{T}$-indexed limits and colimits.

Example 4.1.8. The following examples show that the nascent and mortal part do not preserve images. For both examples, let the index set be $\mathbb{Z}$.

1. Consider a morphism $\varphi\colon C([0,+\infty)) \to C([0,1])$ which has maximal rank everywhere, e.g., by taking φ_0 and φ_1 to be identities and all other maps 0. Clearly, φ is epi and in particular $(\operatorname{im}\varphi)^\dagger = C([0,1])^\dagger = C([0,1])$. However, we have $C([0,+\infty))^\dagger = 0$, so $\operatorname{im}\varphi^\dagger = 0$ and thus $\operatorname{im}\varphi^\dagger \neq (\operatorname{im}\varphi)^\dagger$.

2. Now let $\varphi\colon C([-1,0]) \to C((-\infty,0])$ be of maximal rank everywhere. By a similar argument, we get $\operatorname{im}\varphi^* = 0$ but $(\operatorname{im}\varphi)^* = C([-1,0])$.

While the preservation of images fails in general, there are classes of morphisms for which we get the desired result. We start with a lemma.

Lemma 4.1.9. *Let V and W be objects in $\mathbf{A}^{\mathbf{T}}$ and $\varphi\colon V \to W$ a morphism.*

1. *If φ is epi, then φ^* is epi; if φ is mono, then $\varphi^\dagger$ is mono.*

2. *If $\lim\varphi$ is epi, then $\varphi^{-\infty}$ is epi; if $\operatorname{colim}\varphi$ is mono, then φ^{∞} is mono.*

Proof. Assume φ is epi. Note that the canonical map $P \twoheadrightarrow P^*$ is also epi for any object P in $\mathbf{A}^{\mathbf{T}}$. We get a commutative diagram

$$\begin{array}{ccc} V & \xrightarrow{\varphi} & W \\ \downarrow & & \downarrow \\ V^* & \xrightarrow{\varphi^*} & W^* \end{array}$$

where the composition $V \to W^*$ is epi. Thus, φ^* must be epi, too. The other assertions can be shown analogously. $\square$

Theorem 4.1.10. *Let V and W be objects in $\mathbf{A}^{\mathbf{T}}$ and $\varphi\colon V \to W$ a morphism.*

1. *If* $\operatorname{colim}\varphi$ *is mono, we have canonical isomorphisms*
$$\operatorname{im}\varphi^\dagger \cong (\operatorname{im}\varphi)^\dagger, \qquad \operatorname{im}\varphi^\infty \cong (\operatorname{im}\varphi)^\infty \cong V^\infty.$$
2. *If* $\lim\varphi$ *is epi, we have canonical isomorphisms*
$$\operatorname{im}\varphi^* \cong (\operatorname{im}\varphi)^*, \qquad \operatorname{im}\varphi^{-\infty} \cong (\operatorname{im}\varphi)^{-\infty} \cong W^{-\infty}.$$

Proof. We only show the first part of the theorem, the second one being completely dual. So, assume that $\operatorname{colim}\varphi$ is mono. To start, consider the epi-mono-factorizations

$$V^\dagger \twoheadrightarrow \operatorname{im}\varphi^\dagger \hookrightarrow W^\dagger \qquad \text{and} \qquad V \overset{p}{\twoheadrightarrow} \operatorname{im}\varphi \overset{i}{\hookrightarrow} W$$

of $\varphi^\dagger$ and φ, respectively. Applying the mortal part functor to the second factorization and leaving the first one as is yields a commutative diagram

$$\begin{array}{ccccc} V^\dagger & \twoheadrightarrow & \operatorname{im}\varphi^\dagger & \hookrightarrow & W^\dagger \\ & {\scriptstyle p^\dagger}\searrow & & \nearrow{\scriptstyle i^\dagger} & \\ & & (\operatorname{im}\varphi)^\dagger & & \end{array}$$

By the universal property of epi-mono-factorizations, we get $\operatorname{im}\varphi^\dagger \cong (\operatorname{im}\varphi)^\dagger$ if $i^\dagger$ is mono and $p^\dagger$ is epi, which is what we will show now. Since i is mono, by Lemma 4.1.9 $i^\dagger$ is mono, too. To see that $p^\dagger$ is epi, consider the commutative diagram

$$\begin{array}{ccccccccc} 0 & \longrightarrow & V^\dagger & \longrightarrow & V & \longrightarrow & V^\infty & \longrightarrow & 0 \\ & & \downarrow{\scriptstyle p^\dagger} & & \downarrow{\scriptstyle p} & & \downarrow{\scriptstyle p^\infty} & & \\ 0 & \longrightarrow & (\operatorname{im}\varphi)^\dagger & \longrightarrow & \operatorname{im}\varphi & \longrightarrow & (\operatorname{im}\varphi)^\infty & \longrightarrow & 0 \end{array}$$

with short exact rows. Applying the snake lemma (which holds in p-exact categories, see [Gra12, Lemma 6.2.8]), we get an exact sequence

$$\ldots \longrightarrow \ker p^\infty \longrightarrow \operatorname{coker} p^\dagger \longrightarrow \operatorname{coker} p \longrightarrow \ldots.$$

By construction, p is epi so that $\operatorname{coker} p = 0$. By assumption, $\operatorname{colim}\varphi = \operatorname{colim} i \circ \operatorname{colim} p$ is mono, so $\operatorname{colim} p$ is mono. Using the second part of Lemma 4.1.9, we get that p^∞ is mono as well. Thus, we get $\ker p^\infty$. The exact sequence above now reads

$$\ldots \longrightarrow 0 \longrightarrow \operatorname{coker} p^\dagger \longrightarrow 0 \longrightarrow \ldots,$$

so that $\operatorname{coker} p^\dagger$ and hence $p^\dagger$ is epi as needed. This finishes the proof of $\operatorname{im}\varphi^\dagger \cong (\operatorname{im}\varphi)^\dagger$.

Next, recall that colim preserves epis, so $\operatorname{colim} p$ is not only mono as noted before but also epi. Hence, $\operatorname{colim} p$ is an isomorphism because monos and epis in p-exact categories are always normal, see [Gra12, §1.5.4]. Thus, we get a commutative diagram

$$\begin{array}{ccccc} V & \twoheadrightarrow & \operatorname{im}\eta_V & \hookrightarrow & \Delta\operatorname{colim}(V) \\ \downarrow{\scriptstyle p} & & & & \uparrow{\scriptstyle (\Delta\operatorname{colim} p)^{-1}} \\ \operatorname{im}\varphi & \twoheadrightarrow & \operatorname{im}\eta_{\operatorname{im}\varphi} & \hookrightarrow & \Delta\operatorname{colim}(\operatorname{im}\varphi) \end{array}$$

with the epi-mono-factorizations of η_V and $\eta_{\operatorname{im}\varphi}$ in the rows. Uniqueness of the epi-mono-factorization implies that the middle terms have to agree, so we obtain $V^\infty = \operatorname{im}\eta_V \cong \operatorname{im}\eta_{\operatorname{im}\varphi} = (\operatorname{im}\varphi)^\infty$. With our assumption that $\operatorname{colim}\varphi$ is mono, the first part of Lemma 4.1.9 yields that φ^∞ is also mono. Thus, we have $\operatorname{im}\varphi^\infty = V^\infty$. This finishes the proof. □

A kernel-cokernel correspondence In our computational applications, we will often need to assume that maps on the limit and colimit are not only mono or epi, but actually isomorphisms. In this case, we have the following result, which we will not need later on and record for its own sake. It might be of use in devising efficient algorithms for the computation of kernels and cokernels of maps in persistent homology, which we do not investigate further in this thesis. For this purpose, one can also show results similar to Theorem 4.1.10 with kernels and cokernels replacing images.

Proposition 4.1.11. *Let V and W be objects in $\mathbf{A}^{\mathbf{T}}$ and $\varphi\colon V \to W$ a morphism.*

1. *If* $\operatorname{colim}\varphi$ *is an isomorphism, we have a canonical isomorphism*

$$\operatorname{coker}\varphi^\infty \cong \ker\varphi^{\triangleleft}.$$

2. *If* $\lim\varphi$ *is an isomorphism, we have a canonical isomorphism*

$$\ker\varphi^{-\infty} \cong \operatorname{coker}\varphi^{\triangleright}.$$

Proof. We only show the first assertion, the proof of the second one is completely analogous. Consider the diagram

$$\begin{array}{ccccccccc} 0 & \longrightarrow & V^\infty & \longrightarrow & \Delta\operatorname{colim}V & \longrightarrow & V^{\triangleleft} & \longrightarrow & 0 \\ & & \downarrow{\scriptstyle\varphi^{-\infty}} & & \downarrow{\scriptstyle\Delta\operatorname{colim}\varphi} & & \downarrow{\scriptstyle\varphi^{\triangleleft}} & & \\ 0 & \longrightarrow & W^\infty & \longrightarrow & \Delta\operatorname{colim}W & \longrightarrow & W^{\triangleleft} & \longrightarrow & 0 \end{array}$$

Its rows are exact, so we can apply the snake lemma to get an exact sequence

$$\ldots \longrightarrow \ker(\Delta\operatorname{colim}\varphi) \longrightarrow \ker\varphi^{\triangleleft} \longrightarrow \operatorname{coker}\varphi^\infty \longrightarrow \operatorname{coker}(\Delta\operatorname{colim}\varphi) \longrightarrow \ldots$$

By assumption, the terms on the left and the right vanish, which yields the claim. □

4.2. Lifespan of persistence modules

We will now specialize the discussion of the lifespan functors from general p-exact diagram categories to those categories we are chiefly interested in, namely persistence modules and barcodes. More precisely, we analyze how lifespan functors and their complementary functors change the barcodes of persistence modules in § 4.2.1. To finish the chapter, we discuss how lifespan functors behave with respect to dualization of persistence modules and the corresponding change from $\mathbf{T}$ to $\mathbf{T}^{\mathrm{op}}$ as the index category in § 4.2.2.

4.2.1. Lifespan functors and barcodes

In order to give an explicit description of how our lifespan functors change the barcode of an interval-decomposable persistence module, we will take a detour via matching diagrams as reviewed in § 3.1.1. We can apply the theory of lifespan functors to category of matching diagrams $\mathbf{Mch}^{\mathbf{T}}$ because $\mathbf{Mch}$ is p-exact and every object in $\mathbf{Mch}^{\mathbf{T}}$ has a limit and a colimit, as we will show using the component set from Definition 3.1.6.

Proposition 4.2.1. *Every matching diagram D in $\mathbf{Mch}^{\mathbf{T}}$ has a limit and a colimit.*

Proof. The limit is given by

$$\lim D = \{Q \in \mathfrak{C}(D) \mid \operatorname{supp}(Q) \text{ is not strictly bounded below}\},$$

with natural maps $\lim D \to D_t$ matching a class Q to its representative in D_t if there is one. We can also explicitly construct the colimit of D as

$$\operatorname{colim} D = \{Q \in \mathcal{C}(D) \mid \operatorname{supp}(Q) \text{ is not strictly bounded above}\}.$$

Here, the natural maps $D_t \to \operatorname{colim} D$ match an element to its equivalence class if this class is contained in the set above. We omit the straightforward verification that these construction satisfy the universal properties of limits and colimits. □

Remark 4.2.2. The construction above can be adapted to show that $\mathbf{Mch}$ not only has totally ordered limits and colimits, but all cofiltered limits and filtered colimits.

Lifespan functors for barcodes We will now look at how the lifespan functors behave when being transported to barcodes via the equivalences $\mathcal{B}$ and $\mathcal{D}$ between barcodes and matching diagrams reviewed in Theorem 3.1.9. Recall that $\mathfrak{I} = \mathfrak{I}(T)$ denotes the set of all intervals in T. We introduce some more notation.

Definition 4.2.3. We define the following subsets of the intervals $\mathfrak{I}$ in T.

$$\begin{aligned} \mathfrak{I}^* &= \{I \in \mathfrak{I} \mid I \text{ is strictly bounded below}\}, & \mathfrak{I}^{-\infty} &= \mathfrak{I} \setminus \mathfrak{I}^*, \\ \mathfrak{I}^\dagger &= \{I \in \mathfrak{I} \mid I \text{ is strictly bounded above}\}, & \mathfrak{I}^{\infty} &= \mathfrak{I} \setminus \mathfrak{I}^\dagger, \\ \mathfrak{I}^{\dagger,*} &= \mathfrak{I}^* \cap \mathfrak{I}^\dagger, & \mathfrak{I}^{-\infty,\infty} &= \mathfrak{I}^{-\infty} \cap \mathfrak{I}^{\infty}, \\ \mathfrak{I}^{-\infty,\dagger} &= \mathfrak{I}^{-\infty} \cap \mathfrak{I}^\dagger, & \mathfrak{I}^{*,\infty} &= \mathfrak{I}^* \cap \mathfrak{I}^{\infty}. \end{aligned}$$

If B is a barcode, we also define

$$B^\diamond = \{(I, a) \in B \mid I \in \mathfrak{I}^\diamond\}$$

for any lifespan functor $(-)^\diamond$.

Theorem 4.2.4. *Let B be a barcode. We have*

$$\mathcal{B}(\mathcal{D}(B)^\diamond) \cong B^\diamond$$

for all lifespan functors $(-)^\diamond$. Moreover, under these isomorphisms, all lifespan transformations correspond to the respective inclusions and coinclusions.

Proof. From the definitions of $\mathcal{B}$ and $\mathcal{D}$ as well as the explicit constructions of limits and colimits for matching diagrams in the proof of Proposition 4.2.1, we obtain

$$\begin{aligned}\mathcal{B}(\Delta \lim \mathcal{D}(B)) &\cong \{(T,(I,a)) \in \mathfrak{I} \times B \mid I \in \mathfrak{I}^{-\infty}\},\\ \mathcal{B}(\Delta \operatorname{colim} \mathcal{D}(B)) &\cong \{(T,(I,a)) \in \mathfrak{I} \times B \mid I \in \mathfrak{I}^{\infty}\}.\end{aligned}$$

The overlap matching $\mathcal{B}(\epsilon_{\mathcal{D}(B)})\colon \mathcal{B}(\Delta \lim \mathcal{D}(B)) \to \mathcal{B}(\mathcal{D}(B)) \cong B$ matches every interval $(T,(I,a))$ with $I \in \mathfrak{I}^{-\infty}$ to (I,a). Similarly, $\mathcal{B}(\eta_{\mathcal{D}(B)})$ matches every element (I,a) with $I \in \mathfrak{I}^{\infty}$ to $(T,(I,a))$. All lifespan functors are given on the level of barcodes by first forming kernels, cokernels, and images of $\mathcal{B}(\epsilon_{\mathcal{D}(B)})$ and $\mathcal{B}(\eta_{\mathcal{D}(B)})$, and then kernels, cokernels, and images of the resulting composite lifespan transformations. Hence, the claim follows by applying the formulas for kernels, cokernels, and images of overlap matchings from Proposition 3.1.14 several times. □

Lifespan functors and the passage to persistence modules Next, we want to show that all the lifespan functors are compatible with the matching module functor $\mathcal{F}$. Since $\mathcal{F}$ is exact, a straightforward proof strategy would be to show that $\mathcal{F}$ also commutes with lim and colim and then use the fact that all lifespan functors are obtained from lim and colim by forming kernels, cokernels, and images. For colimits, this works out.

Lemma 4.2.5. *The functor $\mathcal{F}\colon \mathbf{Mch} \to \mathbf{Vec}$ commutes with $\mathbf{T}$-indexed colimits.*

Proof. Recall that in the proof of Proposition 4.2.1 we constructed the colimit of a matching diagram D as the set of components $Q \in \mathcal{C}(D)$ whose support is in $\mathfrak{I}^{\infty}$. Further, recall from the definition of the component set that each component can be regarded as a matching diagram. As such, D is canonically isomorphic to the disjoint union (which is not the coproduct, but rather a *butterfly product* in **Mch**, see [Gra12, Section 2.1.7]) of all its components. Clearly, $\mathcal{F}$ takes disjoint unions to direct sums. Moreover, for each component $Q \in \mathcal{C}(D)$ the colimit of $\mathcal{F}(Q)$ is one-dimensional if $\operatorname{supp} Q \in \mathfrak{I}^{\infty}$ and trivial else. Altogether, we obtain a natural isomorphism

$$\operatorname{colim} \mathcal{F}(D) \cong \operatorname{colim} \bigoplus_{Q \in \mathcal{C}(D)} \mathcal{F}(Q) \cong \bigoplus_{Q \in \mathcal{C}(D)} \operatorname{colim} \mathcal{F}(Q) \cong \bigoplus_{\substack{Q \in \mathcal{C}(D)\\ \operatorname{supp}(Q) \in \mathfrak{I}^{\infty}(T)}} \mathbb{F} \cong \mathcal{F}(\operatorname{colim} D),$$

proving the claim. □

In contrast to colimits, $\mathcal{F}$ generally does not commute with $\mathbf{T}$-indexed limits: Consider the matching diagram D indexed by the negative integers and given by $D_{-n} = \{1, \dots, n\}$ with structure maps matching each number to itself. Then $\mathcal{F}(\lim D) = \bigoplus_{n \in \mathbb{N}} \mathbb{F}$, but $\lim \mathcal{F}(D) = \prod_{n \in \mathbb{N}} \mathbb{F}$.

Instead, we will use a more explicit argument to show that $\mathcal{F}$ commutes with the ancient part, which, together with the colimit, can also be used as a starting point to construct the other lifespan functors by forming kernels, cokernels, and images.

Theorem 4.2.6. *Let D be a matching diagram. We have canonical isomorphisms*

$$\mathcal{F}(D)^{\diamond} \cong \mathcal{F}(D^{\diamond})$$

for all lifespan functors $(-)^{\diamond}$, which commute with the lifespan transformations.

Proof. We start by showing that $\mathcal{F}$ commutes with the ancient part. For this, consider the epi-mono-factorizations $\Delta \lim D \twoheadrightarrow D^{-\infty} \hookrightarrow D$ and $\Delta \lim \mathcal{F}(D) \twoheadrightarrow \mathcal{F}(D)^{-\infty} \hookrightarrow \mathcal{F}(D)$. Recall that $\mathcal{F}$ preserves exactness and hence also monos and epis. Thus, by applying $\mathcal{F}$ to the first diagram, we get another epi-mono-factorization. The universal property of the limit also induces a unique morphism $\mathcal{F}(\Delta \lim D) \to \Delta \lim \mathcal{F}(D)$ through which the cone morphism $\mathcal{F}(\Delta \lim D) \to \mathcal{F}(D)$ factors. We obtain a commutative diagram

$$\begin{array}{ccccc} \mathcal{F}(\Delta \lim D) & \twoheadrightarrow & \mathcal{F}(D^{-\infty}) & \hookrightarrow & \mathcal{F}(D) \\ \downarrow & & & & \uparrow \mathrm{id} \\ \Delta \lim \mathcal{F}(D) & \twoheadrightarrow & \mathcal{F}(D)^{-\infty} & \hookrightarrow & \mathcal{F}(D) \end{array}$$

Since epi-mono factorizations are unique up to unique isomorphism, we only need to show that the composite morphism $\mathcal{F}(\Delta \lim D) \to \mathcal{F}(D)^{-\infty}$ is epi in order to obtain our claim. So let $t_0 \in T$ and $m \in \mathcal{F}(D)^{-\infty}_{t_0}$. Because m is in the ancient part, i.e., the image of the natural map $\lim \mathcal{F}(D) \to \mathcal{F}(D)_{t_0}$, there exists a family $(m_t)_t$ with $m_t \in \mathcal{F}(D)_t$, $m_{t_0} = m$ and such that $\mathcal{F}(D)_{s,t}(m_s) = m_t$ whenever $s \leq t$. Now, choose finite index sets A_t for $t \in T$ such that we can write each m_t as a linear combinations $m_t = \sum_{\alpha \in A_t} \lambda_{\alpha,t} d_{\alpha,t}$ with $d_{\alpha,t} \in D_t$ and $\lambda_{\alpha,t} \neq 0$ for all $t \in T$. Because $\mathcal{F}(D)_{s,t_0}(m_s) = m_{t_0}$ holds for any $s \leq t_0$, we obtain that for any $\alpha \in A_{t_0}$ and $s \leq t_0$ there exists $\beta = \beta(\alpha, s) \in A_s$ with $(d_{\beta,s}, d_{\alpha,t_0}) \in D_{s,t_0}$. In particular, the component $Q_\alpha \in \mathcal{C}(D)$ represented by d_{α,t_0} has support in $\mathfrak{I}^{-\infty}$ for any $\alpha \in A_{t_0}$. Thus, $m = m_{t_0}$ is the image of $\sum_{\alpha \in A_{t_0}} \lambda_{\alpha,t_0} Q_\alpha \in \mathcal{F}(\lim D)$ under the composite morphism $\mathcal{F}(\lim D) \to \mathcal{F}(D)^{-\infty}_{t_0}$. Hence, this map is epi as we needed to show.

The claimed isomorphisms for the other lifespan functors can now be deduced from the isomorphisms we have shown already: Consider the commutative squares

$$\begin{array}{ccc} \mathcal{F}(D) & \xrightarrow{\eta_{\mathcal{F}(D)}} & \Delta \operatorname{colim} \mathcal{F}(D) \\ \downarrow \mathrm{id} & & \downarrow \\ \mathcal{F}(D) & \xrightarrow{\mathcal{F}(\eta_D)} & \mathcal{F}(\Delta \operatorname{colim} D) \end{array} \quad \text{and} \quad \begin{array}{ccc} \mathcal{F}(D)^{-\infty} & \xrightarrow{\alpha_{\mathcal{F}(D)}} & \mathcal{F}(D) \\ \downarrow & & \downarrow \mathrm{id} \\ \mathcal{F}(D^{-\infty}) & \xrightarrow{\mathcal{F}(\alpha_D)} & \mathcal{F}(D) \end{array}$$

We have just shown that the vertical maps in the square on the right are isomorphisms. The vertical maps in the square on the left are isomorphisms because $\mathcal{F}$ and colim commute by Lemma 4.2.5. Thus, we obtain

$$\mathcal{F}(D)^\dagger = \ker \eta_{\mathcal{F}(D)} \cong \mathcal{F}(\ker \eta_D) = \mathcal{F}(D^\dagger)$$
$$\mathcal{F}(D)^\infty = \operatorname{im} \eta_{\mathcal{F}(D)} \cong \mathcal{F}(\operatorname{im} \eta_D) = \mathcal{F}(D^\infty)$$
$$\mathcal{F}(D)^* = \operatorname{coker} \alpha_{\mathcal{F}(D)} \cong \mathcal{F}(\operatorname{coker} \alpha_D) = \mathcal{F}(D^*).$$

By these isomorphisms, the vertical maps in the commutative squares

$$\begin{array}{ccc} \mathcal{F}(D)^\dagger & \xrightarrow{\beta_{\mathcal{F}(D)}} & \mathcal{F}(D)^* \\ \downarrow & & \downarrow \\ \mathcal{F}(D^\dagger) & \xrightarrow{\mathcal{F}(\beta_D)} & \mathcal{F}(D^*) \end{array} \quad \text{and} \quad \begin{array}{ccc} \mathcal{F}(D)^{-\infty} & \xrightarrow{\gamma_{\mathcal{F}(D)}} & \mathcal{F}(D)^\infty \\ \downarrow & & \downarrow \\ \mathcal{F}(D^{-\infty}) & \xrightarrow{\mathcal{F}(\gamma_D)} & \mathcal{F}(D^\infty) \end{array}$$

are isomorphisms, too. This yields

$$\begin{aligned}
\mathcal{F}(D)^{-\infty,\dagger} &= \ker \beta_{\mathcal{F}(D)} \cong \mathcal{F}(\ker \beta_D) = \mathcal{F}(D^{-\infty,\dagger}) \\
\mathcal{F}(D)^{\dagger,*} &= \operatorname{im} \beta_{\mathcal{F}(D)} \cong \mathcal{F}(\operatorname{im} \beta_D) = \mathcal{F}(D^{\dagger,*}) \\
\mathcal{F}(D)^{*,\infty} &= \operatorname{coker} \gamma_{\mathcal{F}(D)} \cong \mathcal{F}(\operatorname{coker} \gamma_D) = \mathcal{F}(D^{*,\infty}) \\
\mathcal{F}(D)^{-\infty,\infty} &= \operatorname{im} \gamma_{\mathcal{F}(D)} \cong \mathcal{F}(\operatorname{im} \gamma_D) = \mathcal{F}(D^{-\infty,\infty}).
\end{aligned}$$ □

Lifespan functors for interval-decomposable persistence modules Finally, combining the fact that the matching module functor $\mathcal{F}$ and the lifespan functors commute by Theorem 4.2.6 with the fact that $\mathcal{F}$ and the barcode module functor $\mathcal{M}$ are compatible with each other by Proposition 3.2.3, we can use the formulas for the effect of lifespan functors on barcodes from Theorem 4.2.4 to describe how the lifespan functors change the barcodes of persistence modules.

Corollary 4.2.7. *Let M be a persistence module. If B is a barcode of M, then*

$$B^{\diamond} = \{(I, a) \in B \mid I \in \mathfrak{I}^{\diamond}\}$$

is a barcode for $M^{\diamond}$, where $(-)^{\diamond}$ is any lifespan functor.

From Corollary 4.2.7 we obtain that all the short exact sequences in the lifespan diagram of an interval-decomposable persistence module are obtained up to isomorphism by applying the barcode module functor $\mathcal{M}$ to a short exact sequence of barcodes of the form $B' \hookrightarrow B' \sqcup B'' \twoheadrightarrow B''$. The inclusion and coinclusion into and out of the disjoint union only match bars with identical underlying intervals, so they admit one-sided inverses by Lemma 5.1.1. Applying $\mathcal{M}$ to the sequence of barcodes above preserves these one-sided inverses, so we obtain the following corollary.

Corollary 4.2.8. *All the short exact sequences in the lifespan diagram of an interval-decomposable persistence module split.*

Complementary functors for interval-decomposable persistence modules For the unborn complement, the expected formula for the effect on barcodes and the compatibility with $\mathcal{F}$ hold as for the lifespan functors. We summarize the results and omit the analogous proofs.

Proposition 4.2.9. *The unborn complement satisfies*

$$\mathcal{B}(\mathcal{D}(B)^{\triangleleft}) \cong B^{\triangleleft} := \{(T \setminus I, a) \mid (I, a) \in B, I \neq T \text{ and } I \in \mathfrak{I}^{\infty}\}$$

for any barcode B. Moreover, the unborn complement commutes with the matching module functor up to natural transformation. In particular, if M is a persistence module with barcode B, then $B^{\triangleleft}$ is a barcode of $M^{\triangleleft}$.

For the ghost complement, however, not all of the corresponding statements hold in general: It does not commute with the matching module functor and thus does not change the barcode of a persistence module as we would like it to. This is closely related to the fact that $\mathcal{F}$ does not commute with limits as mentioned before Theorem 4.2.6, which can even happen for PFD persistence modules indexed by the real line as the example $\bigoplus_{n \in \mathbb{N}} C((-\infty, -n))$

shows. The problems with the ghost complement disappear for classes of persistence modules where limits do commute with $\mathcal{F}$, e.g., those of *finite type*, which are persistence modules with a finite barcode. Similarly, everything works out as desired if the index set has a minimal element $t_{\min}$, because then we have $\lim \mathcal{F}(D) \cong \mathcal{F}(D_{t_{\min}}) \cong \mathcal{F}(\lim D)$.

Proposition 4.2.10. *Let B be a barcode satisfying $\lim \mathcal{F}(\mathcal{D}(B)) \cong \mathcal{F}(\lim \mathcal{D}(B))$. Then we have*

$$\mathcal{B}(\mathcal{D}(B)^{\triangleright}) \cong B^{\triangleright} := \{(T \setminus I, a) \mid (I, a) \in B, I \neq T \text{ and } I \in \mathfrak{I}^{-\infty}\}.$$

Moreover, on the full subcategory of these barcodes the ghost complement commutes with the matching module functor up to natural transformation. In particular, if M is a persistence module with barcode B satisfying the requirement above, then $B^{\triangleright}$ is a barcode of $M^{\triangleright}$.

Remark 4.2.11. For persistence modules, some of the lifespan functors admit more explicit descriptions. In particular, the mortal part of a persistence module $M = ((M_t)_t, (m_{s,t})_{s,t})$ is the submodule given by the subspaces $M_t^{\dagger} = \bigcup_u \ker m_{t,u} \subseteq M_t$. In this form, the construction has been considered before by Höppner and Lenzing in [HL81a]. They describe it as analogous to taking the submodule of all torsion elements of a module over some integral domain. Certain categories of persistence modules can be shown to be equivalent to categories of modules over some ring (see [CK18; Les15; ZC05]), and under these equivalences, the mortal part indeed corresponds to the torsion submodule. Furthermore, the immortal part has also been considered before in applications of barcodes to symplectic geometry. For a recent example, see [Dah21].

Note that, on the other hand, while the ancient part is always a submodule of the intersection of images, $M_t^{-\infty} \subseteq \bigcap_s \operatorname{im} m_{s,t}$, in general the two need not be isomorphic. The persistence module M_3 described in § 5.2.2 provides a counterexample.

4.2.2. Lifespan and dualization

When passing from homology to cohomology, we will see later in Proposition 5.3.4 that what happens on the level of persistence modules is dualization. When passing from absolute to relative persistent homology, our correspondence result Theorem 5.3.2 will involve the lifespan functors. So, in order to get the full picture involving all four persistence modules associated to a diagram of spaces, we now also have to analyze whether dualization is compatible with lifespan functors.

Because we dualize, we will not only consider persistence modules indexed by $(T, \leq)$, but also ones indexed by $(T, \geq)$. When interpreting lifespan in terms of barcodes, it is important to note that the direction of the order on the index set changes the meaning of the different classes of intervals we consider, i.e., Definition 4.2.3 depends on whether we use the usual or the opposite order. For example, we have $\mathfrak{I}^*(T, \leq) = \mathfrak{I}^{\dagger}(T, \geq)$ and $\mathfrak{I}^{-\infty}(T, \leq) = \mathfrak{I}^{\infty}(T, \geq)$. Thus, one should expect duals of mortal parts to correspond to nascent parts of duals and so on. To avoid confusion, we introduce some notation.

Notation 4.2.12. In the context of the indexing category $\mathbf{T}^{\mathrm{op}}$, we will write

$$(-)_{\dagger} := (-)^{*}, \quad (-)_{\infty} := (-)^{-\infty}, \quad (-)_{*} := (-)^{\dagger}, \quad (-)_{-\infty} := (-)^{\infty},$$
$$(-)_{\dagger,*} := (-)^{\dagger,*}, \quad (-)_{-\infty,\infty} := (-)^{-\infty,\infty}, \quad (-)_{*,\infty} := (-)^{-\infty,\dagger}, \quad (-)_{-\infty,\dagger} := (-)^{*,\infty},$$
$$(-)_{\triangleleft} := (-)^{\triangleright}, \quad (-)_{\triangleright} := (-)^{\triangleleft}.$$

The above convention now yields $\mathfrak{I}_\diamond(T, \geq) = \mathfrak{I}^\diamond(T, \leq)$ for any lifespan functor $(-)^\diamond$.

Proposition 4.2.13. *Let M be a persistence module. We have canonical isomorphisms*

$$(M^\dagger)^\vee \cong (M^\vee)_\dagger, \qquad (M^\infty)^\vee \cong (M^\vee)_\infty, \qquad (M^\triangleleft)^\vee \cong (M^\vee)_\triangleleft$$

Proof. The functor $\operatorname{Hom}(-, \mathbb{F})$ takes colimits to limits, so we have a canonical isomorphism $(\Delta \operatorname{colim} M)^\vee \cong \Delta \lim M^\vee$. Together with the kernel, cokernel, and image descriptions for dual maps from Lemma 3.2.7, this yields the claim. □

The limit functor for persistence modules commonly exhibits less desirable properties than the colimit functor. For example, the limit functor does not preserve exactness and does not commute with the functor $\mathcal{F}$ while the colimit functor does. A similar phenomenon arises with dualization of persistence modules, preventing the previous proposition from holding for all lifespan functors: In general, we do not have an isomorphism between $(\Delta \lim M)^\vee$ and $\Delta \operatorname{colim} M^\vee$, because the vector spaces $\operatorname{Hom}(\lim M, \mathbb{F})$ and $\operatorname{colim} M^\vee$ need not be isomorphic. However, if $(T, \leq)$ has a smallest element $t_{\min}$, then we have $(\Delta \lim M)^\vee \cong \Delta \operatorname{Hom}(M_{t_{\min}}, \mathbb{F}) \cong \Delta \operatorname{colim} M^\vee$. Thus, we get the following.

Proposition 4.2.14. *Assume that $(T, \leq)$ has a smallest element and let M be a $\mathbf{T}$-indexed persistence module. Then we have canonical isomorphisms $(M^\diamond)^\vee \cong (M^\vee)_\diamond$ for any lifespan functor or complementary functor $(-)^\diamond$.*

For later use, we also record the following completely equivalent reformulation in terms of persistence modules indexed by the opposite order.

Proposition 4.2.15. *Assume that $(T, \leq)$ has a largest element and let M be a $\mathbf{T}^{\mathrm{op}}$-indexed persistence module. Then we have canonical isomorphisms $(M_\diamond)^\vee \cong (M^\vee)^\diamond$ for any lifespan functor of complementary functor $(-)^\diamond$.*

Note that above, we use the notation $(-)^\vee$ as a functor from $\mathbf{T}$-indexed persistence modules to $\mathbf{T}^{\mathrm{op}}$-indexed persistence modules and also vice versa.

Furthermore, in the PFD case, applying any lifespan functor $(-)^\diamond$ to a persistence module M has the same effect on barcodes as the corresponding functor $(-)_\diamond$ applied to the dual persistence module $M^\vee$.

Proposition 4.2.16. *Let M be a* PFD *persistence module. Then $M^\diamond$ and $(M^\vee)_\diamond$ have the same barcodes for any lifespan functor $(-)^\diamond$.*

Proof. By Lemma 3.2.6 we know that PFD persistence modules have the same barcode as their duals, so the claim follows immediately from the explicit formula in Corollary 4.2.7 for the effect of lifespan functors on barcodes. □

5. Some applications of lifespan functors

In this chapter, we prove the results stated in § 2.1.2. We start by discussing injective and projective objects in the equivalent categories of barcodes and matching diagrams in Section 5.1. Afterwards, we classify injective and projective objects in **pers** in Section 5.2, and discuss which parts of this classification also hold in **Pers**. As our last application, we discuss natural dualities in persistent homology in terms of lifespan functors in Section 5.3.

For this chapter, we do not put any general constraints on our totally ordered index set T. In Section 5.2, we will make some special considerations in the case where $T = \mathbb{R}$, and in Section 5.3 we will at some point consider the case where T is finite.

5.1. Injective and projective barcodes and matching diagrams

As our first application, we classify injective and projective objects in the p-exact categories of barcodes and matching diagrams. We start by characterizing injective and projective barcodes in terms of vanishing lifespan functors in § 5.1.1. These characterizations are then translated to the category of matching diagrams and combined with previous results on lifespan functors in § 5.1.2.

5.1.1. Injective and projective barcodes

As a first application, we will now use our lifespan functors to characterize projective and injective objects in the categories of barcodes. The following characterization of split mono and epi overlap matchings will be important in what follows.

Lemma 5.1.1. *Let σ be an overlap matching and assume that σ is mono or epi. Then σ is split if and only if $((I, a), (I', a')) \in \sigma$ implies $I = I'$.*

Proof. If $((I, a), (I', a')) \in \sigma$ implies $I = I'$ for some overlap matching σ, then its opposite matching σ° is again an overlap matching. If σ is epi, this yields a right inverse and if σ is mono, this yields a left inverse.

If on the other hand σ is split mono or split epi, there needs to be an overlap matching τ with $((I', a'), (I, a)) \in \tau$ whenever $((I, a), (I', a')) \in \sigma$. Since both σ and τ are overlap matchings, I and I' overlap each other above, so we have $I = I'$ whenever $((I, a), (I', a')) \in \sigma$. □

In other words, Lemma 5.1.1 states that an overlap matching is split if and only if its opposite matching is again an overlap matching.

Theorem 5.1.2. *A barcode B is projective if and only if $B^\dagger = 0$, and injective if and only if $B^* = 0$.*

Proof. We will only show the first statement, the second one can be shown analogously. First, assume that $B^\dagger = 0$. In order to show that B is projective, we consider some overlap matching $\sigma\colon B \to B'$ and need to show that it factors through an arbitrary epi $\tau\colon B'' \to B'$. Consider σ and τ as ordinary matchings and set $\rho = \tau^\circ \circ \sigma$, where τ° is the opposite matching of τ (see Definition 3.1.2). We show that ρ is in fact an overlap matching, i.e., that for any $((I,a),(I'',a'')) \in \rho$ we have that I overlaps I'' above:

Since $B^\dagger = 0$, we have $I \in \mathfrak{I}^\infty$, so that I bounds any other interval, and in particular I'', above. What is left to check is that I'' bounds I below and that the two intervals have non-empty intersection. If $((I,a),(I'',a'')) \in \rho$, then by definition of ρ there is some $(I',a') \in B'$ such that $((I,a),(I',a')) \in \sigma$ and $((I'',a''),(I',a')) \in \tau$. Since τ is epi, its cokernel vanishes and we obtain $I' \subseteq I''$ from the explicit cokernel formula in Proposition 3.1.14. Moreover, I overlaps I' above, so we know that I' bounds I below and that $I \cap I' \neq \emptyset$. Together with $I' \subseteq I''$, this implies that I'' bounds I below and that $I \cap I'' \neq \emptyset$. In total, I overlaps I'' above and ρ is an overlap matching.

Now, an easy calculation verifies that we have $\tau \bullet \rho = \sigma$, i.e., that when considering ρ as an overlap matching, its overlap composition with τ recovers σ. Hence, we have shown that σ factors through τ, so B is projective.

Next, assume that $B^\dagger \neq 0$. We want to show that in this case B is not projective by constructing a barcode B' and an epi $\sigma\colon B' \to B$ that does not split. To do so, choose $(I,a) \in B$ such that $I \in \mathfrak{I}^\dagger$, which is possible by our assumption $B^\dagger \neq 0$. Define

$$J = \{t \in T \mid \text{there exists } s \in I \text{ with } s \leq t\}.$$

Clearly, J is an interval in T and it overlaps I above. We define

$$B' = (B \setminus \{(I,a)\}) \cup \{(J,a)\}$$

and $\sigma\colon B' \to B$ by matching each element of $B \setminus \{(I,a)\}$ to itself and matching (J,a) to (I,a). This matching σ has trivial cokernel since $I \subseteq J$, so σ is epi as desired. But, we have $I \neq J$ since $I \in \mathfrak{I}^\dagger$. Thus, σ matches non-identical intervals and consequently does not split by Lemma 5.1.1, so B cannot be projective. □

5.1.2. Injective and projective matching diagrams

Translating Theorem 5.1.2 via the equivalence of barcodes and matching diagrams, we also obtain that a matching diagram is projective if and only if its mortal part vanishes and injective if and only if nascent part vanishes.

Corollary 5.1.3. *A matching diagram D is projective if and only if $D^\dagger = 0$, and injective if and only if $D^* = 0$.*

By Proposition 4.1.6 we know that vanishing mortal and nascent part can equivalently be described in terms of the structure maps of a diagram, given that taking limits and colimits of diagrams is exact. It is therefore interesting to check whether taking limits and colimits of matching diagrams is exact.

Proposition 5.1.4. *The functors* $\operatorname{colim}, \lim\colon \mathbf{Mch}^{\mathbf{T}} \to \mathbf{Mch}$ *are exact.*

Proof. Let $D \to D' \to D''$ be an exact sequence of matching diagrams. By Proposition 3.2.4 the functor $\mathcal{F}$ preserves exactness, so the sequence $\mathcal{F}(D) \to \mathcal{F}(D') \to \mathcal{F}(D'')$ remains exact. It is well-known that taking colimits of persistence modules is exact. Thus, the sequence $\operatorname{colim} \mathcal{F}(D) \to \operatorname{colim} \mathcal{F}(D') \to \operatorname{colim} \mathcal{F}(D'')$ is still exact. Using that $\mathcal{F}$ commutes with taking $\mathbf{T}$-indexed colimits by Lemma 4.2.5, we get that $\mathcal{F}(\operatorname{colim} D) \to \mathcal{F}(\operatorname{colim} D') \to \mathcal{F}(\operatorname{colim} D'')$ is also exact. By Proposition 3.2.4 the functor $\mathcal{F}$ reflects exactness, so $\operatorname{colim} D \to \operatorname{colim} D' \to \operatorname{colim} D''$ is exact, proving that taking colimits of matching diagrams is exact. Self-duality of the category **Mch** implies that taking limits then has to be exact, too. □

Knowing that taking limits and colimits of matching diagrams is exact, we can now combine the equivalent conditions for vanishing mortal and nascent parts from Proposition 4.1.6 and Corollary 5.1.3 to obtain the following.

Corollary 5.1.5. *A matching diagram D is projective if and only if all of its structure maps are mono, and injective if and only if all of its structure maps are epi.*

A natural question to ask is whether statements analogous to the above Corollaries 5.1.3 and 5.1.5 also hold for persistence modules instead of matching diagrams: can we characterize projectivity/injectivity or structure maps being mono/epi by vanishing mortal/nascent parts?

5.2. Injective and projective persistence modules

After having classified injectives and projectives in the categories of barcodes and matching diagrams, we will now check which of the characterizations given in Section 5.1 also hold for persistence modules. Specifically, in § 5.2.1, we show that the characterization of injective and projective matching diagrams from before also holds in the case of PFD persistence modules. We finish by discussing which parts of the characterization also hold for general persistence modules in § 5.2.2.

5.2.1. Injective and projective objects in pers

Before talking about PFD persistence modules, we prove one more result about lifespan functors of persistence modules in general.

Proposition 5.2.1. *For any persistence module M, we have $M^\dagger = 0$ if and only if all structure maps of M are mono. Moreover, if $M^* = 0$, then all structure maps of M are epi.*

Proof. The first part of the proposition is just a special case of the first part of Proposition 4.1.6, noting that taking colimits of persistence modules is exact.

For the second part, we repeat parts of the dual version of the proof of Proposition 4.1.6: If $M^* = 0$, then $\Delta \lim M \to M$ is epi, i.e., $\lim M \to M_t$ is epi for all $t \in T$. This implies in particular that for any structure map $M_s \to M_t$, the composition $\lim M \to M_s \to M_t$ is epi since it is equal to $\lim M \to M_t$. As a consequence, $M_s \to M_t$ needs to be epi, finishing the proof. □

As we will see in § 5.2.2, the converse to the second part of the proposition does not hold in general. In the category **pers** of PFD persistence modules, however, we can indeed characterize projectives and injectives in a way analogous to matching diagrams $\mathbf{Mch}^{\mathbf{T}}$. We prove a lemma before stating and proving the main theorem.

Lemma 5.2.2. *Let $B \subseteq \mathfrak{I} \times A$ be a barcode such that $\mathcal{M}(B)$ is* PFD *and let $I \in \mathfrak{I}^\infty$ be an interval. If there exists an epimorphism $\varphi\mathcal{M}(B) \to C(I)$, then there exists $(J, a) \in B$ with $I = J$ and φ induces the identity morphism $C(J) \to C(I)$.*

Proof. □

Theorem 5.2.3. *Let M be a* PFD *persistence module.*

1. *The following are equivalent:*
   a) *All structure maps of M are mono.*
   b) *$M^\dagger = 0$.*
   c) *M is projective in* **pers**.
2. *The following are equivalent:*
   a) *All structure maps of M are epi.*
   b) *$M^* = 0$.*
   c) *M is injective in* **pers**.

Proof. Starting with the first part of the theorem, we first note that we have already shown that $M^\dagger = 0$ is equivalent to M having mono structure maps for any persistence module M in Proposition 5.2.1. Thus, what is left to show for the first part is that $M^\dagger = 0$ is equivalent to M being projective in the PFD category. To do so, we fix a barcode decomposition $M \cong \bigoplus_\alpha C(I_\alpha)$, which is possible by Crawley-Boevey's theorem [BC20; Cra15] since M is PFD.

Now, assume that $M^\dagger = 0$, or equivalently $M = M^\infty$. We want to show that M is projective in **pers**. A direct sum of projectives is projective, so it suffices to check that the interval modules $C(I_\alpha)$ in the decomposition of M are projective in **pers**. Recall that **pers** is abelian, so in order to show that $C(I_\alpha)$ is projective in **pers** we only need to show now that any epimorphism $\varphi \colon N \to C(I_\alpha)$ with N PFD splits:

Choosing a barcode decomposition $N \cong \bigoplus_\beta C(J_\beta)$ induces maps $\varphi_\beta \colon C(J_\beta) \to C(I_\alpha)$ for each β. According to the dual version of [BL15, Lemma 4.3] (which Bauer and Lesnick state in the $\mathbb{R}$-indexed case, but whose proof generalizes to our T-indexed setting), there has to be some β_0 such that $I_\alpha \subseteq J_{\beta_0}$ and that simultaneously J_{β_0} has to overlap I_α above. Since we assume $M^\dagger = 0$, $I_\alpha \in \mathfrak{I}^\infty$ holds by Corollary 4.2.7, so we obtain $I_\alpha = J_{\beta_0}$. This yields that φ_{β_0} is an isomorphism. We can thus define $\psi \colon C(I_\alpha) \to N$ as the composition

$$C(I_\alpha) \xrightarrow{\varphi_{\beta_0}^{-1}} C(J_{\beta_0}) \hookrightarrow \bigoplus_\beta C(J_\beta) \cong N.$$

By construction, we have $\varphi \circ \psi = \varphi_{\beta_0}^{-1} \circ \varphi_{\beta_0}$, which is the identity on $C(I_\alpha)$, so φ splits. Thus, we have shown that $C(I_\alpha)$, and consequently M, is projective.

Next, we assume that $M^\dagger \neq 0$ and show that M is not projective in **pers**. Because the mortal part of M does not vanish, there now has to be some α_0 with $I_{\alpha_0} \in \mathfrak{I}^\dagger$. We proceed as in the proof of Theorem 5.1.2 and define

$$J = \{t \in T \mid \text{there exists } s \in I_{\alpha_0} \text{ with } s \leq t\}.$$

Clearly, J is an interval in T and it overlaps I_{α_0} above. The canonical map $C(J) \to C(I_{\alpha_0})$ is an epi, which we can use to obtain an epi

$$\bigoplus_{\alpha \neq \alpha_0} C(I_\alpha) \oplus C(J) \to \bigoplus_{\alpha \neq \alpha_0} C(I_\alpha) \oplus C(I_{\alpha_0}) \cong M$$

in **pers**, which is an isomorphism on all summands except for $C(J)$. If this map would split, the splitting would induce a morphism $C(I_{\alpha_0}) \to C(J)$, which cannot exist since I_{α_0} by construction does not overlap J above. Thus, the epi we constructed does not split and M is not projective in **pers**. This finishes the proof of the first part.

For the second part, we first observe that in the PFD setting, barcode decompositions can not only be interpreted as direct sums but even as biproducts. Note that a PFD persistence module may have a barcode consisting of infinitely many intervals, so this assertion is not guaranteed by **pers** being abelian and thus having finite biproducts. However, the observation is still true due to the fact that direct sums and products of persistence modules are given pointwise, and they coincide for finite-dimensional vector spaces, so they also coincide for PFD persistence modules. Thus, since we have biproduct decompositions, one can now show that $M^* = 0$ is equivalent to M being injective in **pers** by dualizing the previous argument and exploiting the fact that products of injectives are again injectives.

That $M^* = 0$ implies M having epi structure maps has been shown for all persistence modules M before in Proposition 5.2.1, so what remains to be checked is that $M^* = 0$ if M is PFD and its structure maps are epi. To see that this is the case, one can use the fact that the functor $\lim\colon \mathbf{pers} \to \mathbf{Vec}$ is exact (because derived inverse limits of PFD persistence modules vanish [Jen70, Proposition 1.1], [Roo62, Théorème 2] and reuse the argument in the proof of Proposition 4.1.6 to show that $\Delta \lim M \to M$ is epi if the structure maps of M are epi, which implies that $M^* = 0$. □

Any PFD persistence module has a barcode and the lifespan functors are compatible with the passage to barcodes, so another way of phrasing the previous theorem is that a PFD persistence module is projective or injective in **pers** if and only if its barcode has the corresponding property in $\mathbf{Barc}(\mathbf{T})$.

5.2.2. Injective and projective objects in Pers

When considering persistence modules beyond the PFD category, some of the equivalences established in Theorem 5.2.3 do not hold anymore in general. We give a few examples.

First consider the case $T = \mathbb{R}$, which is the most important for persistent homology. In this case, if the structure maps of M are epi, then $M^* = 0$, providing a converse to the second part of Proposition 5.2.1. The assertion can be shown by a simple argument as in the proof of [Zel51, Lemma 7]. Moreover, any injective object in $\mathbf{Vec}^{\mathbf{R}}$ has epi structure maps (see [Höp83; HL81a] for classification results for injectives in **Pers**). The converse does not hold: the real-indexed persistence module $M_1 = C(-\infty, 0)$ satisfies $M_1^* = 0$ and has epi structure maps, but it is not injective in $\mathbf{Vec}^{\mathbf{R}}$ because the obvious mono $M_1 = C(-\infty, 0) \to \prod_{n \in \mathbb{N}_{>0}} C\left(-\infty, -\frac{1}{n}\right)$ does not split. Similarly, any projective object in $\mathbf{Vec}^{\mathbf{R}}$ has vanishing mortal part and mono structure maps (see [HL81b] for a classification of projectives in **Pers**). Again, the converse does not hold: the real-indexed interval module $M_2 = C(0, \infty)$ satisfies $M_1^\dagger = 0$ and has mono structure maps, but it is not projective in $\mathbf{Vec}^{\mathbf{R}}$ because the obvious epi $\bigoplus_{n \in \mathbb{N}_{>0}} C\left(\frac{1}{n}, \infty\right) \to C(0, \infty) = M_2$ does not split.

For general totally ordered indexing sets T, we still have the implications that any injective object in **Pers** has epi structure maps and that any projective object in **Pers** has mono structure maps and vanishing mortal part. However, for persistence modules, having epi structure maps does not always imply vanishing nascent part: there is a non-zero persistence module M_3 indexed by the opposite poset of the first uncountable ordinal ω_1 whose structure maps are all epi, but which satisfies $\lim M_3 = 0$ ([HS54, Section 3]), so that $M_3^* = M_3 \neq 0$. We are presently unable to determine whether injective persistence modules necessarily have vanishing nascent part.

5.3. Dualities in terms of lifespan functors

As we have seen in § 2.1.1, the main motivation for the introduction of lifespan functors is that they naturally appear when discussing certain duality results in persistent homology. In § 5.3.1, we prove these natural duality results in terms of lifespan functors and dualization of persistence modules. These dualities are then combined in § 5.3.2 with previous results on when lifespan functors commute with taking images of morphisms from § 4.1.2 in order to obtain correspondence results between barcodes of images of maps in absolute and relative persistent (co)homology.

All chain complexes and (co)homology groups in this section are understood to be singular (co)homology with coefficients in our fixed field $\mathbb{F}$.

5.3.1. Persistent homology dualities

Absolute-relative correspondence We will now prove a generalization of the absolute-relative correspondence [dSMV11, Proposition 2.4] involving our lifespan functors. In order for this to work nicely, we only consider filtrations that satisfy the following condition.

Definition 5.3.1. Let X be a $\mathbf{T}$-indexed diagram of topological spaces. Recall that X is called a filtration if all of its structure maps are injective. We say that X is *colimit proper* if the natural maps $\operatorname{colim} H_d(X) \to H_d(\operatorname{colim} X)$ and $H_d(\operatorname{colim} X) \to \lim H_d(\operatorname{colim} X, X)$ are isomorphisms for all d.

Note that colimit properness is always satisfied if the diagram X is initially empty and eventually constant. In particular, if the index set has a largest element $t_{\max}$ and a minimal element $t_{\min}$ then every X with $X_{t_{\min}} = \emptyset$ is colimit proper. These properties are usually given in the computational setting for persistent homology.

Theorem 5.3.2. *Let X be a colimit proper filtration of topological spaces. For all d, we have the following isomorphisms, which are natural in X:*

$$\begin{aligned} H_{d-1}(X)^{\dagger} &\cong H_d(\operatorname{colim} X, X)^*, \\ H_d(X)^{\triangleleft} &\cong H_d(\operatorname{colim} X, X)^{-\infty}, \\ H_d(X)^{\infty} &\cong H_d(\operatorname{colim} X, X)^{\triangleright}. \end{aligned}$$

Proof. To shorten notation, we write A for $\operatorname{colim} X$. Since X is a filtration, the natural map $C_*(X) \to C_*(\Delta A)$ is mono. Denoting the cokernel of this map by $C_*(A, X)$, we have a short exact sequence

$$0 \longrightarrow C_*(X) \longrightarrow C_*(\Delta A) \longrightarrow C_*(A, X) \longrightarrow 0$$

of persistent chain complexes. This induces a long exact sequence of persistence modules

$$\cdots \longrightarrow \Delta H_d(A) \xrightarrow{\epsilon_d} H_d(A,X) \xrightarrow{\partial} H_{d-1}(X) \xrightarrow{\eta_{d-1}} \Delta H_{d-1}(A) \longrightarrow \cdots .$$

Since we assume X to be colimit proper, the map ϵ_d can be identified with the counit

$$\epsilon_{H_d(A,X)} \colon \Delta \lim H_d(A,X) \to H_d(A,X)$$

of the adjunction $\Delta \dashv \lim$. Similarly, the map η_{d-1} may be identified with the unit

$$\eta_{H_{d-1}(X)} \colon H_{d-1}(X) \to \Delta \operatorname{colim} H_{d-1}(X)$$

of the adjunction $\operatorname{colim} \dashv \Delta$. Applying the definition of the lifespan functors, the claimed isomorphisms are now simply given by exactness of the above sequence :

$$\begin{aligned} H_{d-1}(X)^{\dagger} &\cong \ker \eta_{d-1} \cong \operatorname{coker} \epsilon_d \cong H_d(A,X)^{*}, \\ H_d(X)^{\triangleleft} &\cong \operatorname{coker} \eta_d \cong \operatorname{im} \epsilon_d \cong H_d(A,X)^{-\infty}, \\ H_d(X)^{\infty} &\cong \operatorname{im} \eta_d \cong \ker \epsilon_d \cong H_d(A,X)^{\triangleright}. \end{aligned}$$

These isomorphisms are natural in X as a direct consequence of the fact that the construction of the long exact sequence is natural in X. □

Using the barcode formulas for lifespan functors and complements in Corollary 4.2.7 and Propositions 4.2.9 and 4.2.10, one can easily recover the original duality result by de Silva et al. [dSMV11, Proposition 2.4]. Moreover, naturality in the filtration variable implies that for a morphism $f \colon X \to Y$ between colimit proper filtrations with $\phi = \operatorname{colim} f$ we also get isomorphisms

$$H_{d-1}(f)^{\dagger} \cong H_d(\phi, f)^{*}, \qquad H_d(f)^{\triangleleft} \cong H_d(\phi, f)^{-\infty}, \qquad H_d(f)^{\infty} \cong H_d(\phi, f)^{\triangleright}$$

in the category of morphisms of persistence modules. These also translate to isomorphisms between the corresponding images, kernels, and cokernels.

Note that the isomorphism between the mortal part of the absolute persistent homology and the nascent part of the relative persistent homology in the proof of Theorem 5.3.2 is induced by the boundary operator. This means that if an interval in the nascent part of the barcode of the relative persistent homology is represented by some relative cycle, the boundary of this cycle represents the same interval in the absolute persistent homology in one dimension lower, as observed in [dSMV11].

Remark 5.3.3. While the above result is stated for persistent homology of filtrations of spaces, a similar statement holds in the purely algebraic setting. Given a filtered chain complex C, we can consider the short exact sequence

$$0 \longrightarrow C \longrightarrow \Delta \operatorname{colim} C \longrightarrow C^{\triangleleft} \longrightarrow 0.$$

We can then continue as in the proof above to get natural isomorphisms

$$H_{d-1}(C)^{\dagger} \cong H_d(C^{\triangleleft})^{*} \qquad H_d(C)^{\triangleleft} \cong H_d(C^{\triangleleft})^{-\infty} \qquad H_d(C)^{\infty} \cong H_d(C^{\triangleleft})^{\triangleright}.$$

Homology-cohomology correspondence For completeness, we also record a functorial version of the correspondence between persistent homology and persistent cohomology [dSMV11, Proposition 2.3], which follows immediately from the universal coefficient theorem.

Proposition 5.3.4. *Let X be a $\mathbf{T}$-indexed diagram of topological spaces. For all d, we have the following isomorphisms, which are natural in X:*

$$H_d(X)^\vee \cong H^d(X),$$
$$H_d(\operatorname{colim} X, X)^\vee \cong H^d(\operatorname{colim} X, X).$$

While the correspondence in [dSMV11, Proposition 2.3] is stated on the level of barcodes, the natural isomorphism asserted in Proposition 5.3.4 appears in its proof, which essentially combines the previous statement with the fact that PFD persistence modules have the same barcode as their duals (Lemma 3.2.6).

As in the absolute-relative correspondence, naturality in the variable X yields corresponding isomorphisms in the category of morphisms of persistence modules for maps $f\colon X \to Y$.

5.3.2. Absolute homology images from relative cohomology images

As a concrete application, we want to explain how to use our previous results for the efficient computation of barcodes for images of morphisms in persistent homology. Note that similar considerations also apply for kernels and cokernels of such morphisms.

As mentioned before, and as is explained, e.g., in [Bau21], one of the most efficient ways currently known to compute the barcode of the persistent homology of a filtration of finite simplicial complexes is to actually compute the barcode of the persistent relative cohomology with the clearing optimization, and to then translate this to persistent homology via the two duality results from de Silva, Morozov, and Vejdemo-Johansson [dSMV11].

Our generalizations of these duality results now allow us to proceed similarly for the image of a map $f\colon X \to Y$. Since we are talking about computational speed-ups, X and Y are assumed to be filtrations with PFD persistent homology indexed by a totally ordered set $\mathbf{T}$ with a smallest element $t_{\min}$ and a largest element $t_{\max}$. We also assume that $X_{t_{\min}} = Y_{t_{\min}} = \emptyset$, so that both filtrations are colimit proper, and we assume that $\operatorname{colim} H_d(f) = \lim H_d(\operatorname{colim} f, f) = H_d(f_{t_{\max}})$ is an isomorphism.

In order to obtain the barcode for $\operatorname{im} H_d(f)$ from barcodes computed with relative cohomology, we start by applying the (non-natural) decomposition

$$\operatorname{im} H_d(f) \cong (\operatorname{im} H_d(f))^\dagger \oplus (\operatorname{im} H_d(f))^\infty$$

from Corollary 4.2.8. We consider both summands separately, making use of the fact that taking barcodes is compatible with direct sums.

Recovering the mortal part Starting with the first summand, we observe that because $\operatorname{colim} H_d(f) = H_d(f_{t_{\max}})$ is an isomorphism, and in particular a monomorphism, we have

$$(\operatorname{im} H_d(f))^\dagger \cong \operatorname{im}(H_d(f)^\dagger)$$

using the first part of Theorem 4.1.10. The natural duality Theorem 5.3.2, which we can apply since X and Y are colimit proper, provides an isomorphism

$$\operatorname{im}(H_d(f)^\dagger) \cong \operatorname{im}(H_{d+1}(\operatorname{colim} f, f)^*).$$

An application of the second part of Theorem 4.1.10 yields the isomorphism

$$\operatorname{im}(H_{d+1}(\operatorname{colim} f, f)^*) \cong (\operatorname{im} H_{d+1}(\operatorname{colim} f, f))^*$$

using that $\lim H_d(\operatorname{colim} f, f) = H_d(f_{t_{\max}})$ is epi. Finally, the duality of homology and cohomology from Proposition 5.3.4 yields an isomorphism

$$(\operatorname{im} H_{d+1}(\operatorname{colim} f, f))^* \cong ((\operatorname{im} H^{d+1}(\operatorname{colim} f, f))^\vee)^*,$$

where we also make use of the fact that applying dualization twice yields the identity on PFD persistence modules. Finally, because our index set has a largest element, Proposition 4.2.15 gives

$$((\operatorname{im} H^{d+1}(\operatorname{colim} f, f))^\vee)^* \cong ((\operatorname{im} H^{d+1}(\operatorname{colim} f, f))_*)^\vee.$$

In total, the above implies that $(\operatorname{im} H_d(f))^\dagger$ and $(\operatorname{im} H^{d+1}(\operatorname{colim} f, f))_*$ have the same barcode by Lemma 3.2.6 because we are in the PFD setting, so we can obtain the mortal part of the absolute homology barcode from some relative cohomology persistence module.

Recovering the immortal part For the second term in the mortal-immortal decomposition of $\operatorname{im} H_d(f)$, we have

$$(\operatorname{im} H_d(f))^\infty \cong H_d(X)^\infty$$

by Theorem 4.1.10. We proceed again with our natural absolute-relative duality from Theorem 5.3.2 to obtain

$$H_d(X)^\infty \cong H_d(\operatorname{colim} X, X)^\triangleright.$$

Since all modules are PFD, passing to cohomology with Proposition 5.3.4 yields

$$(H_d(\operatorname{colim} X, X))^\triangleright \cong ((H^d(\operatorname{colim} X, X))^\vee)^\triangleright.$$

Proposition 4.2.15 finally yields

$$((H^d(\operatorname{colim} X, X))^\vee)^\triangleright \cong (H^d(\operatorname{colim} X, X)_\triangleright))^\vee.$$

Thus, we can also obtain the immortal part of the absolute homology barcode from some relative cohomology persistence module.

Summary in the computational setting Using the results on lifespan functors, complementary functors, and barcodes Corollary 4.2.7 and Propositions 4.2.9 and 4.2.10, we can translate the previous isomorphisms to barcode formulas. We record them in the following corollary, formulated with notation as in the computational setting of filtrations of finite simplicial complexes for later use.

Corollary 5.3.5. *Let $f_\bullet \colon L_\bullet \to K_\bullet$ be a morphism of filtrations of finite simplicial complexes inducing a map $f \colon L \to K$ on colimits such that $H_*(f)$ is an isomorphism. Assume that the index set T has a smallest element $t_{\min}$ and a largest element, and assume that $L_{t_{\min}} = K_{t_{\min}} = \emptyset$, so that no intervals in the barcodes of $L_\bullet$ and $K_\bullet$ contain $t_{\min}$. Then*

$$B(\operatorname{im} H_{d-1}(f_\bullet))^{\dagger,*} = B(\operatorname{im} H^d(f, f_\bullet))_{\dagger,*}$$

for all degrees d, and the map $I \mapsto T \setminus I$ defines bijections

$$B(\operatorname{im} H_*(f_\bullet))^\infty \leftrightarrow B(H^*(L, L_\bullet))_{-\infty},$$
$$B(\operatorname{im} H^*(f, f_\bullet))_{-\infty} \leftrightarrow B(H_*(K_\bullet))^\infty.$$

Note that our formulation of the above result in the introduction was slightly ambiguous because there we did not distinguish the lifespan operations depending on whether a barcode comes from a $(T, \leq)$-indexed or a $(T, \geq)$-indexed persistence module. That is, we interpreted all barcodes there as living in $\mathbf{Barc}(\mathbf{T})$ even though, technically, some of them would live in $\mathbf{Barc}(\mathbf{T}^{\mathrm{op}})$. In the present formulation, all barcodes are assumed to live in the 'correct' category coming from the index set of the underlying persistence modules.

6. Efficient computation of barcodes for images in persistent homology

In this chapter, we develop the algorithm for computing images in persistent homology described in § 2.1.3, making use of the duality results established in Section 5.3. We start by describing some general theory for bases of filtrations of vector spaces in Section 6.1. This theory is then applied in the context of filtrations of simplicial complexes to compute barcodes for images of maps in persistent homology in Section 6.2, where we also describe the speed-ups for the basic algorithm.

In order to emphasize the difference between persistent and non-persistent objects, we will throughout the chapter use the convention of denoting persistence modules and filtrations with a $(-)_\bullet$ subscript, as we have done before when talking about computational applications. Since we are dealing with the computational setting, the reader may also think of all vector spaces in this chapter being finite-dimensional and the index set T having finitely many elements, although many of the statements also hold more generally.

6.1. Filtration compatible bases

Before talking about our computational results, we introduce the framework of filtration compatible bases and show some results in this general framework that will be helpful in the computational setting later on. In § 6.1.1, we define the notion of filtration compatible basis and state an important result for obtaining barcodes of quotients of filtrations that is an immediate consequence of the theory of matching diagrams. We then prove some basic but helpful results on filtration compatible bases whose proofs essentially boil down to linear algebra in § 6.1.2. Similar statements to the ones we make in this section may be seen as folklore results in the topological data analysis community, but to the best of our knowledge there is no comprehensive and systematic treatment as the one we present in the literature, yet.

6.1.1. Filtration compatible bases for barcode computations

Definition 6.1.1. We say that a persistence module $M_\bullet$ is a *filtration* of the vector space $M = \operatorname{colim} M_\bullet$ if $M_t \subseteq M_u$ for all $t \leq u$ and the structure maps $M_{t,u}$ are given by the subspace inclusions. For any $m \in M$, we define its *support* in $M_\bullet$ as $\operatorname{supp}_{M_\bullet}(m) = \{t \in T \mid m \in M_t)\}$. A basis $\mathfrak{M}$ of M will be called *filtration compatible* if $\mathfrak{M}_t = \mathfrak{M} \cap M_t$ is a basis for M_t for all $t \in T$. If $(\mathfrak{M}, \leq)$ is an ordered basis for M, we say that it is a *filtration compatible ordered basis* if it is filtration compatible and $m \leq m' \in \mathfrak{M}$ implies $\operatorname{supp} m' \subseteq \operatorname{supp} m$.

If $M_\bullet$ and $M'_\bullet$ are filtrations of vector spaces, we write $M_\bullet \subseteq M'_\bullet$ if $M_t \subseteq M'_t$ for all $t \in T$. We write $M'_\bullet / M_\bullet$ for the persistence module given by $(M'_\bullet / M_\bullet)_t = M'_t / M_t$, i.e., for the cokernel of the inclusion-induced morphism $M_\bullet \hookrightarrow M'_\bullet$. Similarly, if $M''_\bullet$ is another filtration with $M''_\bullet \subseteq M'_\bullet$, we write $M_\bullet \cap M''_\bullet$ for the persistence module given by $(M_\bullet \cap M''_\bullet)_t = M_t \cap M''_t$.

Observe that if $M_\bullet$ is a filtration of vector spaces and $\mathfrak{M}$ is a filtration compatible basis, then $(\operatorname{supp}(m))_{m\in\mathfrak{M}}$ is a barcode of $M_\bullet$. By interpreting $\mathfrak{M}$ as a matching diagram in the obvious way, this may be seen as a special case of the general equivalence of matching diagrams and barcodes as reviewed in § 3.1.2. With this in mind, the following result is an immediate consequence of Propositions 3.1.14 and 3.2.4.

Proposition 6.1.2. *Let $M_\bullet \subseteq M'_\bullet$ be filtrations of vector spaces with respective filtration compatible bases $\mathfrak{M}$ and $\mathfrak{M}'$ related by an inclusion $\mathfrak{M} \subseteq \mathfrak{M}'$. Then $M'_\bullet/M_\bullet$ has a barcode given by*

$$(\operatorname{supp}_{M'_\bullet}(m) \setminus \operatorname{supp}_{M_\bullet}(m))_{m\in\mathfrak{M}} \cup (\operatorname{supp}_{M'_\bullet}(m))_{m\in\mathfrak{M}'\setminus\mathfrak{M}}.$$

6.1.2. Some linear algebra for filtrations

Basis exchange lemma for filtrations We now develop some helpful facts about filtration compatible bases. We start with a lemma relating supports of basis elements with filtration compatibility.

Lemma 6.1.3. *Let $M_\bullet$ be a filtration of the vector space M with filtration compatible basis $\mathfrak{M}$. Let $\mathfrak{M}'$ be another basis for M such that there exists a bijection $g\colon \mathfrak{M} \to \mathfrak{M}'$ with $\operatorname{supp}_{M_\bullet}(m) = \operatorname{supp}_{M_\bullet}(g(m))$ for all $m \in \mathfrak{M}$. Then $\mathfrak{M}'$ is a filtration compatible basis for $M_\bullet$.*

Proof. We need to show that $\mathfrak{M}'_t$ is a basis for M_t for all t. Since $\mathfrak{M}'$ is a basis, $\mathfrak{M}'_t = \mathfrak{M} \cap M_t$ is linearly independent. We assume $\mathfrak{M}_t$ to be a basis for M_t, so it suffices to show that $\mathfrak{M}_t$ and $\mathfrak{M}'_t$ have the same cardinality.

To see this, note that if $m \in \mathfrak{M}_t = \mathfrak{M} \cap M_t$, we must have $t \in \operatorname{supp}_{M_\bullet}(m) = \operatorname{supp}_{M_\bullet}(g(m))$, so that $g(m) \in \mathfrak{M}'_t = \mathfrak{M}' \cap M_t$ holds. Thus, g restricts to a map $\mathfrak{M}_t \to \mathfrak{M}'_t$. Similarly, the restriction of g^{-1} to $\mathfrak{M}'_t$ yields a map $\mathfrak{M}'_t \to \mathfrak{M}_t$. As the restrictions of g and g^{-1} are inverse to each other, $\mathfrak{M}_t$ and $\mathfrak{M}'_t$ indeed have the same cardinality. ☐

Intersections of filtrations Next, we prove a version of a standard fact about intersections of vector spaces for filtrations.

Lemma 6.1.4. *Let $M'_\bullet, M''_\bullet \subseteq M_\bullet$ be filtrations of vector spaces and let $\mathfrak{M}'$ and $\mathfrak{M}''$ be filtration compatible bases for $M'_\bullet$ and $M''_\bullet$, respectively, such that $\mathfrak{M}' \cup \mathfrak{M}''$ is linearly independent. Then $\mathfrak{M}' \cap \mathfrak{M}''$ is a filtration compatible basis for $M'_\bullet \cap M''_\bullet$. Moreover, for all $m \in \mathfrak{M}' \cap \mathfrak{M}''$ we have*

$$\operatorname{supp}_{M'_\bullet \cap M''_\bullet}(m) = \operatorname{supp}_{M'_\bullet}(m) \cap \operatorname{supp}_{M''_\bullet}(m).$$

Proof. We want to show that $\mathfrak{M}' \cap \mathfrak{M}''$ is a filtration compatible basis for $M'_\bullet \cap M''_\bullet$, i.e., that

$$(\mathfrak{M}' \cap \mathfrak{M}'')_t = \mathfrak{M}' \cap \mathfrak{M}'' \cap M'_t \cap M''_t = \mathfrak{M}'_t \cap \mathfrak{M}''_t$$

is a basis for $(M'_\bullet \cap M''_\bullet)_t = M'_t \cap M''_t$ for all $t \in T$. By standard linear algebra, for subspaces $V', V'' \subseteq V$ of some vector space V with respective bases $\mathfrak{V}'$ and $\mathfrak{V}''$, the intersection $\mathfrak{V}' \cap \mathfrak{V}''$ of the bases is a basis for the intersection $V' \cap V''$ of spaces if the union $\mathfrak{V}' \cup \mathfrak{V}''$ of the bases is linearly independent. In particular, $\mathfrak{M}'_t \cap \mathfrak{M}''_t$ is a basis for $M'_t \cap M''_t$ since the union $\mathfrak{M}'_t \cup \mathfrak{M}''_t \subseteq \mathfrak{M}' \cup \mathfrak{M}''$ is linearly independent by assumption. Thus, $\mathfrak{M}' \cap \mathfrak{M}''$ is a filtration compatible basis for $M' \cap M''$.

For the supports, we have

$$\begin{aligned}\operatorname{supp}_{M'_\bullet \cap M''_\bullet}(m) &= \{t \in T \mid m \in M'_t \cap M''_t\}\\ &= \{t \in T \mid m \in M'_t\} \cap \{t \in T \mid m \in M''_t\}\\ &= \operatorname{supp}_{M'_\bullet}(m) \cap \operatorname{supp}_{M''_\bullet}(m)\end{aligned}$$

for all $m \in \mathfrak{M}' \cap \mathfrak{M}''$, proving the claim. □

Note that the union $\mathfrak{M}' \cup \mathfrak{M}''$ in the statement of Lemma 6.1.4 is to be interpreted as the set-theoretic union of the two bases and not as the concatenation of families of vectors. In this second interpretation, the union could only be linearly independent if the intersection is empty, in which case the statement of the lemma is not meaningful. We will use the special case of Lemma 6.1.4 where $M'_\bullet$ is included in $M''_\bullet$ at the last filtration step (but not necessarily before):

Corollary 6.1.5. *Let $M'_\bullet, M''_\bullet \subseteq M_\bullet$ be filtrations of vector spaces $M' \subseteq M'' \subseteq M$, respectively. Moerover, let $\mathfrak{M}' \subseteq \mathfrak{M}''$ be filtration compatible bases for $M'_\bullet$ and $M''_\bullet$, respectively. Then $\mathfrak{M}'$ is a filtration compatible basis for $M'_\bullet \cap M''_\bullet$. Moreover, for all $m \in \mathfrak{M}'$ we have*

$$\operatorname{supp}_{M'_\bullet \cap M''_\bullet}(m) = \operatorname{supp}_{M'_\bullet}(m) \cap \operatorname{supp}_{M''_\bullet}(m).$$

Rank-nullity for filtrations To finish the general discussion about filtration compatible bases, we state a version of the rank-nullity-theorem for filtrations.

Lemma 6.1.6. *Let $\phi_\bullet \colon M_\bullet \to P_\bullet$ be a morphism of filtrations of vector spaces and consider the linear map $\phi = \phi_N \colon M \to P$. Let $\mathfrak{M}$ be a filtration compatible basis for $M_\bullet$, let $\mathfrak{M}' = \mathfrak{M} \cap \ker \phi$, and assume that $\mathfrak{M}'' = (\phi(m))_{m \in \mathfrak{M} \setminus \mathfrak{M}'}$ is a linearly independent family of vectors. Then $\mathfrak{M}'$ is a filtration compatible basis for $\ker \phi_\bullet$, and $\mathfrak{M}''$ is a filtration compatible basis for $\operatorname{im} \phi_\bullet$. Moreover,*

$$\operatorname{supp}_{\ker \phi_\bullet}(m') = \operatorname{supp}_{M_\bullet}(m') \quad \text{and} \quad \operatorname{supp}_{\operatorname{im} \phi_\bullet}(\phi(m)) = \operatorname{supp}_{M_\bullet}(m)$$

for all $m' \in \mathfrak{M}'$ and for all $m \in \mathfrak{M} \setminus \mathfrak{M}'$.

Proof. We show that $\mathfrak{M}' = \mathfrak{M} \cap \ker \phi$ is a filtration compatible basis for $\ker \phi_\bullet$. Because $M_\bullet$ is a filtration, i.e., its structure maps are all injective, we have $\ker \phi_t = M_t \cap \ker \phi$ for all $t \in T$. Denoting the constant filtration of $\ker \phi$ by $\Delta \ker \phi$, we thus have $\ker \phi_\bullet = M_\bullet \cap \Delta \ker \phi$. By standard linear algebra, $\mathfrak{M}'' = (\phi(m))_{m \in \mathfrak{M} \setminus \mathfrak{M}'}$ being linearly independent implies that $\mathfrak{M}'$ is a basis for $\ker \phi$. Hence, $\mathfrak{M}'$ is a filtration compatible basis for the constant filtration $\Delta \ker \phi$. We can apply Corollary 6.1.5 to obtain that $\mathfrak{M}'$ is a filtration compatible basis for the intersection $M_\bullet \cap \Delta \ker \phi = \ker \phi_\bullet$ satisfying

$$\operatorname{supp}_{\ker \phi_\bullet}(m') = \operatorname{supp}_{\Delta \ker \phi}(m') \cap \operatorname{supp}_{M_\bullet}(m') = T \cap \operatorname{supp}_{M_\bullet}(m') = \operatorname{supp}_{M_\bullet}(m')$$

for all $m' \in \mathfrak{M}'$.

To show that $\mathfrak{M}''$ is a filtration compatible basis for $\operatorname{im} \phi_\bullet$, we need to show that $\mathfrak{M}''_t = \mathfrak{M}'' \cap \operatorname{im} \phi_t$ is a basis for $\operatorname{im} \phi_t$ for all $t \in T$. Since we assume $\mathfrak{M}''$ to be linearly independent, $\mathfrak{M}''_t$ is linearly independent as well. It remains to be shown that $\mathfrak{M}''_t$ is also a generating set for $\operatorname{im}\phi_t$. Because $\mathfrak{M}_t$ is a basis for M_t by assumption, the family $(\phi_t(m))_{m\in\mathfrak{M}_t\setminus\ker\phi_t}$ generates $\operatorname{im}\phi_t$. Hence, it suffices to show that

$$(\phi_t(m))_{m\in\mathfrak{M}_t\setminus\ker\phi_t} \subseteq (\phi(m))_{m\in\mathfrak{M}\setminus\mathfrak{M}'} \cap \operatorname{im}\phi_t = \mathfrak{M}''_t.$$

So let $m \in \mathfrak{M}_t \setminus \ker\phi_t$. We have $\phi_t(m) \neq 0 \in P_t$, so since $P_\bullet$ is a filtration we also get $\phi(m) \neq 0 \in P$. Thus, $m \in \mathfrak{M}_t \setminus \ker\phi_t \subseteq \mathfrak{M} \setminus \ker\phi = \mathfrak{M} \setminus \mathfrak{M}'$. This implies $\phi_t(m) = \phi(m) \in \mathfrak{M}''_t$ as needed, so $\mathfrak{M}''_t$ generates $\operatorname{im}\phi_t$ and hence forms a basis for $\operatorname{im}\phi_t$. Bases are minimal generating sets, so we obtain an equality $(\phi_t(m))_{m\in\mathfrak{M}_t\setminus\ker\phi_t} = (\phi(m))_{m\in\mathfrak{M}\setminus\mathfrak{M}'} \cap \operatorname{im}\phi_t$ from the previously shown inclusion. Thus, for $m \in \mathfrak{M} \setminus \mathfrak{M}'$ we have $\phi(m) \in \operatorname{im}\phi_t$ if and only if $m \in \mathfrak{M}_t = \mathfrak{M} \cap M_t$ if and only if $m \in M_t$. This yields $\operatorname{supp}_{\operatorname{im}\phi_\bullet}(\phi(m)) = \operatorname{supp}_{M_\bullet}(m)$, which finishes the proof. □

In contrast to the considerations about linear independent unions regarding Lemma 6.1.4, we now really need the family $(\phi(m))_{m\in\mathfrak{M}\setminus\mathfrak{M}'}$ to be linearly independent (meaning in particular that no vectors appear more than once) and not just the set $\{\phi(m)) \mid m \in \mathfrak{M}\setminus\mathfrak{M}'\}$. Otherwise, the conclusion of Lemma 6.1.6 may be false. Moreover, note that if one drops the assumption of the above lemma that $P_\bullet$ – and hence the image $\operatorname{im}\phi_\bullet$ – is a filtration, then it may happen that $\mathfrak{M}' = \mathfrak{M} \cap \ker\phi$ is a basis for the vector space $\ker\phi$ but not a filtration compatible basis for the filtration $\ker\phi_\bullet$.

6.2. Computing images in persistent homology

We will now present our results on the efficient computation of image persistence. In § 6.2.1, we describe the basic algorithm for computing barcodes of images of morphisms in persistent homology induced by monomorphisms of filtrations of chain complexes. We then describe how to use clearing in this setting in § 6.2.2. Finally, we formulate the full algorithm for obtaining images in persistent homology in § 6.2.3, which works by applying the previously described algorithm to relative cohomology and then translating via the absolute-relative correspondence.

Recall that throughout the chapter we assume all vector spaces to be finite-dimensional, thus all persistence modules to be PFD, and our index set T to have finitely many elements. As usual, all (co)homology groups in this section are understood to be singular or simplicial (co)homology with coefficients in our fixed field $\mathbb{F}$.

6.2.1. Image barcodes via matrix reduction

We now turn to the setting of image persistence, initially formulated in the simplex-free setting of filtrations of chain complexes. A chain complex of persistence modules $C_\bullet$ with differential $\partial_\bullet$ is called a filtration of a chain complex of vector spaces C with differential ∂ if $C_\bullet$ is a filtration of C as a vector space and the map induced by $\partial_\bullet$ on the colimit of $C_\bullet$ is the same as ∂. Let $C_\bullet$ and $C'_\bullet$ be filtrations of the chain complexes C and C' with corresponding filtration compatible ordered bases $\mathfrak{C}$ and $\mathfrak{C}'$. Let D and D' be the corresponding filtration boundary matrices, i.e., the matrices representing ∂ and ∂' with respect to $\mathfrak{C}$ and $\mathfrak{C}'$. Assume that we are given a monomorphism of filtrations of chain complexes $\varphi_\bullet\colon C_\bullet \to C'_\bullet$ such that the map $\varphi\colon C \to C'$ on the final filtration step is an isomorphism. Let F be the matrix representing φ with respect to $\mathfrak{C}$ and $\mathfrak{C}'$ and define the *mixed basis boundary matrix* $D^\varphi = DF^{-1} = F^{-1}D'$.

Our goal is to determine a barcode for $\operatorname{im} H_*(\varphi_\bullet)$ via matrix reduction, so assume we have $R = DV$ and $R^\varphi = D^\varphi V^\varphi$ reduced with V and V^φ full-rank and upper-triangular. The columns of R, D, V, R^φ and D^φ should be interpreted as coordinate vectors with respect to $\mathfrak{C}$ and the columns of V^φ as coordinate vectors with respect to $\mathfrak{C}'$. Recall that if X is a matrix, we denote its jth column by x_j. The main result of this section can then be stated as follows.

Theorem 6.2.1. *The image of $H_*(\varphi_\bullet)$ has a barcode given by the multiset*

$$\left\{\operatorname{supp}_{C_\bullet}(r_j^\varphi) \setminus \operatorname{supp}_{C'_\bullet}(v_j^\varphi) \neq \emptyset \mid r_j^\varphi \neq 0\right\} \cup \left\{\operatorname{supp}_{C_\bullet}(v_i) \mid r_i = 0 \text{ and } i \notin \operatorname{pivots} R\right\}.$$

Special cases of the general theorem The first algorithm to compute images in persistent homology given by Cohen-Steiner et al. [Coh+09] works essentially in the same way as the algorithm implied by Theorem 6.2.1. That is, Cohen-Steiner et al. [Coh+09] construct D and D^φ as above and obtain the same formula for the barcode of $\operatorname{im} H_*(\varphi)$. However, they prove correctness of this method only in the special case where $C_\bullet = C_*(L_\bullet)$ and $C'_\bullet = C_*(K_\bullet)$ are the simplicial chain complexes of filtrations of simplicial complexes $L_\bullet \hookrightarrow K_\bullet$ such that $L_t = K_t \cap L$ for some fixed complex L, which implies that the filtration orders on the simplices induced by $L_\bullet$ and $K_\bullet$ are the same. Theorem 6.2.1 shows that this condition is not necessary.

If we assume that $C_\bullet = C'_\bullet$ and that φ is the identity then Theorem 6.2.1 yields a formula for the barcode of $H_*(C_\bullet)$. Note that the intervals $\operatorname{supp}_{C_\bullet}(v_i)$ only depend on $C_\bullet$, so we obtain that for any choice of φ as above, the intervals in the barcode of $\operatorname{im} H_*(\varphi_\bullet)$ that are not bounded above are precisely the same as those in the barcode of $H_*(C_\bullet)$. This can also be inferred from the first part of Theorem 4.1.10.

In the case where $C_\bullet = C'_\bullet = C_*(K_\bullet)$ is the simplicial chain complex of a filtration of simplicial complexes $K_\bullet$, the statement above is exactly the same as the classical barcode formula reviewed in § 1.2.2. As noted by de Silva, Morozov, and Vejdemo-Johansson [dSMV11], this formula may also be applied to the persistent relative cohomology of $K_\bullet$, where $C_\bullet = C'_\bullet = C^*(K, K_\bullet)$ and K denotes the last complex in the filtration. They also note that if D is a filtration boundary matrix for $C_*(K_\bullet)$, then $D^\perp$ is a filtration (co)boundary matrix for $C^*(K, K_\bullet)$. Here, $(-)^\perp$ denotes transposing a matrix along the anti-diagonal. Similarly, we note that if we have a non-trivial morphism $f_\bullet \colon L_\bullet \to K_\bullet$ inducing maps $C_*(f_\bullet)$ and $C^*(f, f_\bullet)$, then if D^f is a mixed basis boundary matrix in the absolute homology case, $(D^f)^\perp$ is a mixed basis (co)boundary matrix in the relative cohomology case.

Proving the general theorem The proof of Theorem 6.2.1 will be based on a sequence of intermediate results. As alluded to before, the general idea is to write

$$\operatorname{im} H_*(\varphi_\bullet) \cong \frac{\varphi(Z_*(C_\bullet))}{\varphi(Z_*(C_\bullet)) \cap B_*(C'_\bullet)},$$

where Z_* and B_* denote cycles and boundaries of the corresponding chain complexes, respectively. We will find filtration compatible bases $\mathfrak{Z}$ and $\mathfrak{B}$ for $\varphi(Z_*(C_\bullet))$ and $\varphi(Z_*(C_\bullet)) \cap B_*(C'_\bullet)$, respectively, such that $\mathfrak{B} \subseteq \mathfrak{Z}$ holds and we can apply Proposition 6.1.2.

If X is a matrix, we will write $\operatorname{cols} X$ for the family of all its non-zero column vectors.

Lemma 6.2.2. *The familiy* $\operatorname{cols} V^\varphi$ *is a filtration compatible basis for* $C'_\bullet$,

$$\mathfrak{B} = \operatorname{cols} FR^\varphi$$

is a filtration compatible basis for $B_*(C'_\bullet)$*, and for all* j *with* $r_j^\varphi \neq 0$ *we have*

$$\operatorname{supp}_{B_*(C'_\bullet)}(Fr_j^\varphi) = \operatorname{supp}_{C'_\bullet}(v_j^\varphi).$$

Proof. We start by showing that $\operatorname{cols} V^\varphi$ is a filtration compatible basis for $C'_\bullet$: We have $\operatorname{pivot} v_j^\varphi = j$ since V^φ is full-rank and upper-triangular. It follows that v_j^φ has the same support in $C'_\bullet$ as the jth element of $\mathfrak{C}'$. Thus, $\operatorname{cols} V^\varphi$ is a filtration compatible basis for $C'_\bullet$ by Lemma 6.1.3.

Next, note that $(\partial(v))_{v \in \operatorname{cols} V^\varphi \setminus \ker \partial} = \operatorname{cols} FR^\varphi$ is linearly independent since R^φ is reduced and F has full rank. Thus, we can apply Lemma 6.1.6 to the map of filtrations $\partial_\bullet \colon C'_\bullet \to C'_\bullet$ and the filtration compatible basis $\operatorname{cols} V^\varphi$ to obtain that $\operatorname{cols} FR^\varphi$ is a filtration compatible basis for $B_*(C'_\bullet) = \operatorname{im} \partial_\bullet$. The assertion on the supports also follows from the support formula in Lemma 6.1.6. □

Now that we have a filtration compatible basis for $B_*(C'_\bullet)$, we want to extend it to a filtration compatible basis for $\varphi_\bullet(Z_*(C_\bullet))$.

Lemma 6.2.3. *Let*

$$\mathfrak{X} = \operatorname{cols} R^\varphi \cup \{v_j \mid j \notin \operatorname{pivots} R^\varphi\} \quad \textit{and} \quad \mathfrak{X}' = \mathfrak{X} \cap \ker \partial.$$

Then $\mathfrak{X}$ *is a filtration compatible basis for* $C_\bullet$, $\mathfrak{Z} = F\mathfrak{X}'$ *is a filtration compatible basis for* $\varphi_\bullet(Z_*(C_\bullet))$*, and for all* $x \in \mathfrak{X}'$ *we have*

$$\operatorname{supp}_{\varphi_\bullet(Z_*(C_\bullet))}(Fx) = \operatorname{supp}_{C_\bullet}(x).$$

Proof. We start by showing that $\mathfrak{X}$ is a filtration compatible basis for $C_\bullet$. The same argument as in the beginning of the proof of Lemma 6.2.2 yields that $\operatorname{cols} V$ is a filtration compatible basis for $C_\bullet$. Next, note that $\mathfrak{X}$ is linearly independent since all elements have unique pivots: R^φ is reduced and we only consider those v_j with $\operatorname{pivot} v_j = j \notin \operatorname{pivots} R^\varphi$. Moreover, we have a bijection $\mathfrak{X} \to \operatorname{cols} V$ given by mapping v_j to itself and mapping r_j^φ to $v_{\operatorname{pivot} r_j^\varphi}$. Now, $\operatorname{pivot} r_j^\varphi = \operatorname{pivot} v_{\operatorname{pivot} r_j^\varphi}$ implies

$$\operatorname{supp}_{C_\bullet}\left(r_j^\varphi\right) = \operatorname{supp}_{C_\bullet}\left(v_{\operatorname{pivot} r_j^\varphi}\right).$$

Since $\operatorname{cols} V$ is a filtration compatible basis for $C_\bullet$, Lemma 6.1.3 now implies that $\mathfrak{X}$ is also a filtration compatible basis for $C_\bullet$.

Next, we can apply Lemma 6.1.6 to the boundary operator $\partial_\bullet \colon C_\bullet \to C_\bullet$ and the filtration compatible basis $\mathfrak{X}$ since $(\partial(v))_{v \in \mathfrak{X} \setminus \mathfrak{X}'} \subseteq \operatorname{cols} R$ is linearly independent by reducedness of R. We obtain that $\mathfrak{X}' = F^{-1}\mathfrak{Z}$ is a filtration compatible basis for $\ker \partial_\bullet = Z_*(C_\bullet)$ with $\operatorname{supp}_{Z_*(C_\bullet)}(x) = \operatorname{supp}_{C_\bullet}(x)$ for all $x \in \mathfrak{X}'$. The claim now follows from the fact that $\varphi_\bullet$ is mono, so that its restriction is an isomorphism $Z_*(C_\bullet) \to \varphi_\bullet(Z_*(C_\bullet))$ represented by F. □

Having extended the filtration compatible basis for $B_*(C'_\bullet)$ to one for $\varphi_\bullet(Z_*(C_\bullet))$, we also obtain one for $\varphi_\bullet(Z_*(C_\bullet)) \cap B_*(C'_\bullet)$.

Lemma 6.2.4. *The family $\mathfrak{B}$ is a filtration compatible basis for $\varphi_\bullet(Z_*(C_\bullet)) \cap B_*(C'_\bullet)$, and for all j with $r^\varphi_j \neq 0$ we have*

$$\operatorname{supp}_{\varphi_\bullet(Z_*(C_\bullet))\cap B_*(C'_\bullet)}(Fr^\varphi_j) = \operatorname{supp}_{C_\bullet}(r^\varphi_j) \cap \operatorname{supp}_{C'_\bullet}(v^\varphi_j).$$

Proof. Recall that $\mathfrak{B}$ is a filtration compatible basis for $B_*(C'_\bullet)$, and $\mathfrak{Z}$ extends $\mathfrak{B}$ to one for $\varphi_\bullet(Z_*(C_\bullet))$. Now Corollary 6.1.5 together with the support equalities from Lemmas 6.2.2 and 6.2.3 yield the claim. □

We prove a final lemma.

Lemma 6.2.5. $\operatorname{pivots} R = \operatorname{pivots} R^\varphi$.

Proof. The matrices D and $D^\varphi = DF^{-1}$ have the same column space. Matrix reduction does not change column spaces, so R and R^φ also have the same column space. In particular, every non-zero column of R is a non-trivial linear combination of non-zero columns of R^φ and vice versa. The pivot of a linear combination of a reduced set of column vectors must be the same as the pivot of one of these vectors, so we indeed obtain $\operatorname{pivots} R = \operatorname{pivots} R^\varphi$. □

We are now prepared to prove the main result of this section.

Proof of Theorem 6.2.1. By definition of the induced map in homology, we have

$$\operatorname{im} H_*(\varphi_\bullet) \cong \frac{\varphi_\bullet(Z_*(C_\bullet))}{\varphi_\bullet(Z_*(C_\bullet)) \cap B_*(C'_\bullet)}.$$

The claim follows by applying Proposition 6.1.2 to the inclusion $\varphi_\bullet(Z_*(C_\bullet)) \cap B_*(C'_\bullet) \subseteq \varphi_\bullet(Z_*(C_\bullet))$ with the filtration compatible bases $\mathfrak{B} \subseteq \mathfrak{Z}$, with supports as previously determined in Lemmas 6.2.3 and 6.2.4. Note that in the basis $\mathfrak{Z}$ we choose columns Fv_i with $i \notin \operatorname{pivots} R^\varphi$, while the formula in Theorem 6.2.1 requires $i \notin \operatorname{pivots} R$. But these conditions are equivalent by Lemma 6.2.5. □

6.2.2. Clearing for image persistence

Before discussing clearing in the image setting, we first note that the clearing procedure as reviewed in § 1.2.3 not only works for simplicial complexes, but also for the algebraic setting of chain complexes we are considering presently. That is, if D is a filtration boundary matrix for $C_\bullet$ and $R = DV$ is a reduction, then if $r_j \neq 0$ we must have $r_i = 0$ for $i = \operatorname{pivot} r_j$. Now, assume that the matrix D represents the boundary operator with respect to a filtration compatible basis $\mathfrak{C}$ that is compatible with the homological grading of $C_\bullet$ in the sense that the restriction of this basis to each grading summand is again a basis of that summand. Then we can again perform the reduction of D degree-wise and clear columns in lower degrees using pivot information from higher degrees. Note that these considerations also work for *co*homological grading, where the reduction is then performed in ascending rather than descending degrees. Thus, we can also use clearing in the setting of relative cohomology of filtrations of simplicial complexes.

Returning to the algebraic image setting, we now assume that the bases $\mathfrak{C}$ and $\mathfrak{C}'$ as well as the map $\varphi_\bullet \colon C_\bullet \to C'_\bullet$ are compatible with the grading in the sense as above. Here, there is no direct analogue to the procedure outlined above: the mixed basis boundary matrix D^φ fails to admit the property described above, i.e., it may happen that $r^\varphi_j \neq 0$ but $r^\varphi_i \neq 0$

for $i = \operatorname{pivot} r_j^\varphi$. In order to obtain a useful condition for columns of R^φ to be zero, we need to additionally consider the boundary matrix $D' = FD^\varphi$, which represents the boundary operator on $C'_\bullet$ with respect to $\mathfrak{C}'$, and assume we have a reduction $R' = D'V'$.

Proposition 6.2.6. *Let $R' = D'V'$ and $R^\varphi = D^\varphi V^\varphi$ be reduced. For all indices j we have $r_j^\varphi = 0$ if and only if $r'_j = 0$.*

Proof. First, note that $r_j^\varphi = 0$ if and only if $Fr_j^\varphi = 0$ because F is invertible. Moreover, FR^φ and R' have the same column space, since $FR^\varphi = R'(V')^{-1}V^\varphi$. Thus, the number of zero columns of R^φ is the same as the number of zero columns of R' since their ranks are equal and their non-zero columns are linearly independent. Now, it suffices to show that $r_j^\varphi = 0$ implies $r'_j = 0$, so assume $r_j^\varphi = 0$. Then $Fr_j^\varphi = 0$, but Fr_j^φ is also the same as the j-th column of $R'(V')^{-1}V^\varphi$. This is a linear combination of columns of R' with non-zero coefficient for r'_j since $(V')^{-1}V^\varphi$ is full-rank and upper-triangular. Non-zero columns of R' are linearly independent, so this linear combination can only be zero if $r'_j = 0$. □

In order to apply clearing to the reduction of D^φ, one can now reduce D' with clearing as usual, and clear the columns with the same indices in D^φ. Even more than that, one can not only clear the columns of D^φ whose index appears as a pivot in R', but rather *every* column with the same index as a 0 column in R', meaning also those that have been reduced to 0 via column operations on D'. With this optimization, there are no more columns that have to be reduced to 0 in the reduction of D^φ even before any column operations on D^φ have been performed. The only reduction steps that are left to be performed are those that make pivots unique among the non-zero columns.

Corollary 6.2.7. *If D' has already been reduced to R', one can set $r_j^\varphi = 0$ for all j with $r'_j = 0$ before reducing D^φ, and no further columns of R^φ will be reduced to* 0.

6.2.3. Assembling barcodes from (co)homology computations

Given a morphism of filtrations of chain complexes $\varphi_\bullet \colon C_\bullet \to C'_\bullet$ as before, we have now seen how to compute the barcode of $\operatorname{im} H_*(\varphi)$ via reductions of the filtration boundary matrix D and the mixed basis boundary matrix D^φ in Theorem 6.2.1. The matrix D can be reduced using the standard clearing procedure as we recalled it in § 1.2.3. The matrix D^φ can not be straightforwardly reduced with clearing, but we can clear columns in it according to Corollary 6.2.7 if we simultaneously also reduce the filtration boundary matrix D' corresponding to the domain $C'_\bullet$ of $\varphi_\bullet$.

Recalling our concrete setting of simplicial complexes and their persistent homology, assume that we are given filtrations $L_\bullet$ and $K_\bullet$ of two isomorphic simplicial complexes $L \cong K$ and a monomorphism $f_\bullet \colon L_\bullet \to K_\bullet$ inducing an isomorphism $f \colon L \to K$. Let D^L and D^K be filtration boundary matrices for the respective filtrations and let D^f be the corresponding mixed basis boundary matrix. Applying the previous results with $\varphi_\bullet = C_*(f_\bullet)$, we see that the barcode of $\operatorname{im} H_*(f_\bullet)$ can be determined via reductions of D^L and D^f and that the reduction of D^f may be performed with clearing if D^K has previously been reduced.

From relative cohomology to absolute homology As is also known from the single filtration case, clearing needs to be initialized by performing a full persistence computation in the first homological degree for which persistence is computed. Because persistence computations are often only feasible in low dimensions and practitioners are often only interested in barcodes in low degrees, it is much more powerful to apply clearing for *co*homological grading, allowing for the initialization to be performed in degree 0; see also [Bau21] for a more detailed discussion on why clearing and cohomology work particularly well together. Thus, our goal is to perform cohomological computations and still recover the image $\operatorname{im} H_*(f_\bullet)$ in homology. However, the persistent cochain complex giving rise to persistent cohomology is not a filtration, so the basic matrix reduction algorithm does not directly apply there. Thus, we instead perform computations in the relative cohomology setting given by the map $H^*(f, f_\bullet)$, coming from the filtrations $C^*(K, K_\bullet)$ and $C^*(L, L_\bullet)$. The image of the relative cohomology map no longer has the same barcode as $\operatorname{im} H_*(f_\bullet)$, but we can translate between the two barcodes with Corollary 5.3.5. In particular, Corollary 5.3.5 implies that in order to determine the barcode of $\operatorname{im} H_*(f_\bullet)$, it suffices to compute $B(H^*(L, L_\bullet))_{-\infty}$ and $B(\operatorname{im} H^*(f, f_\bullet))_{\dagger,*}$. Following the discussion in § 6.2.1 on special cases of our general theorem, we observe that $B(H^*(L, L_\bullet))_{-\infty}$ may be determined from a reduction of the coboundary matrix $(D^L)^\perp$ and that $B(\operatorname{im} H^*(f, f_\bullet))_{\dagger,*}$ may be determined from a reduction of the coboundary matrix $(D^f)^\perp$. In the relative cohomology setting, the matrices $(D^L)^\perp$ and $(D^f)^\perp$ play the roles of D' and D^φ in the general setting, so by Corollary 6.2.7 we can simultaneously reduce these matrices with clearing. In total, we can hence obtain the barcode of $\operatorname{im} H_*(f_\bullet)$ via two boundary matrix reductions, both performed with clearing in cohomological grading, i.e., ascending dimension of simplices.

The final algorithm We summarize the discussion in the following theorem. To simplify notation, we will assume that we are given funtions k and l on $K \cong L$ that induce the filtrations $K_\bullet$ and $L_\bullet$, respectively, via their sublevel set filtrations. For example, the functions l and k would be given by the diameter functions if $K_\bullet$ and $L_\bullet$ are Vietoris–Rips filtrations for different metrics on the same set of points. Moreover, recall that if A is a matrix, we write a_j for its jth column. To determine barcodes from reductions of boundary matrices, recall that the column and row indices of the matrices D^L, $(D^L)^\perp$, D^f, $(D^f)^\perp$, etc., correspond to the simplices of $K \cong L$ (in different orders). In particular, the pivot index of a column vector c will in this context always correspond to a unique simplex, which we denote by $\operatorname{psimp} c$. The support of the chain represented by c in the respective filtration $C_*(K_\bullet)$ or $C_*(L_\bullet$ (depending on with respect to which filtration order compatible basis c is a coordinate vector) can then be determined as $\operatorname{supp}(c) = [a(\operatorname{psimp} c), \infty)$, where a is a placeholder for either k or l (with the same dependence as before) and where $[t, \infty)$, as usual, denotes the interval extending from t to the largest index in our index set set T. Combining Theorem 6.2.1 and Corollaries 5.3.5 and 6.2.7, we now get the following.

Corollary 6.2.8. *The matrices $(D^L)^\perp$ and $(D^f)^\perp$ can be reduced with clearing, and given reductions $S = (D^f)^\perp W$ and $R = (D^L)^\perp V$, the barcode of $\operatorname{im} H_*(f_\bullet)$ can be determined as the multiset*

$$\{[l(\operatorname{psimp} w_j), k(\operatorname{psimp} s_j)) \neq \emptyset \mid s_j \neq 0\} \cup \{[l(\operatorname{psimp} v_i), \infty) \mid r_i = 0 \textit{ and } i \notin \operatorname{pivots} R\}\,.$$

Recall that the column and row indices of the coboundary matrices indicated by $(-)^\perp$ correspond to the simplices of $K \cong L$ in reverse filtration order. Hence, the pivot simplex of a column vector appearing in the theorem will be the *first* simplex appearing in the simplicial filtration among those that correspond to a non-zero entry of the column, while for the usual boundary matrices D^L, D^K, etc., the pivot simplex of a column would be the one that appears *last* in the filtration among those simplices that correspond to a non-zero entry.

For the convenience of the reader, we summarize the algorithm resulting from Corollary 6.2.8 in pseudo-code in §6.2.3, keeping the notation from before. As mentioned in §2.1.3, an implementation of this method based on Ripser [Bau21] is publicly available [BS21b]. Our software works under the assumption that $L_\bullet = \mathrm{Rips}_\bullet(X, d)$ and $K_\bullet = \mathrm{Rips}_\bullet(X, d')$ are filtrations of Vietoris–Rips complexes corresponding to two metrics d and d' on a finite set X that satisfy $d(x, y) \geq d'(x, y)$ for all $x, y \in X$ with the maps $f_t \colon L_t \to K_t$ being given by inclusion. Note that $L_t = \mathrm{Rips}_t(X, d)$ being a subcomplex of $K_t = \mathrm{Rips}_t(X, d')$ for all t is ensured by the inequality $d \geq d'$. The implementation also makes use of a version of the emergent and apparent pairs optimizations [Bau21], which we do not discuss here, referring to [BS22] instead. It also uses some more technical optimizations such as sparse matrix representations and implicit storage of certain matrices. In particular, neither the boundary nor coboundary matrices or their reductions are explicitly stored, but rather (re-)computed as needed from the sparse matrices V, W and the distance matrices for d and d'. The computation is also not initialized by performing matrix reduction in dimension 0, but rather by performing an efficient algorithm for finding minimum spanning trees. These optimizations do not differ from the case of a single filtration and are discussed in [Bau21].

Algorithm 6.2.9 Algorithm to compute image persistence via two matrix reductions with clearing in cohomological grading

Input: Filtration boundary matrix D^L with n columns, mixed basis boundary matrix D^f
Output: Barcode of $\operatorname{im} H_*(f)$

$R \leftarrow (D^L)^\perp$
$V \leftarrow I_n$
$S \leftarrow (D^f)^\perp$
$W \leftarrow I_n$
$B \leftarrow \emptyset$
for $m = 0, \dots, \dim L - 1$ **do**
  while $\exists i < j$ with $r_i \neq 0$, $\operatorname{pivot} r_i = \operatorname{pivot} r_j$, and $\dim \operatorname{psimp} r_i = m + 1$ **do**
    $r_j \leftarrow r_j - \frac{\operatorname{pentry} r_j}{\operatorname{pentry} r_i} r_i$
    $v_j \leftarrow v_j - \frac{\operatorname{pentry} r_j}{\operatorname{pentry} r_i} v_i$
  end while
  for j with $r_j \neq 0$ and $\dim \operatorname{psimp} r_j = m + 1$ **do**
    $r_{\operatorname{pivot} r_j} \leftarrow 0$
    $v_{\operatorname{pivot} r_j} \leftarrow r_j$
  end for
end for
for j with $r_j = 0$ **do**
  $s_j \leftarrow 0$
  $w_j \leftarrow v_j$
  if $j \notin \operatorname{pivots} R$ **then**
    $B \leftarrow B \sqcup \{[l(\operatorname{psimp} v_j), \infty)\}$
  end if
end for
while $\exists i < j$ with $s_i \neq 0$ and $\operatorname{pivot} s_i = \operatorname{pivot} s_j$ **do**
  $s_j \leftarrow s_j - \frac{\operatorname{pentry} s_j}{\operatorname{pentry} s_i} s_i$
  $w_j \leftarrow w_j - \frac{\operatorname{pentry} s_j}{\operatorname{pentry} s_i} w_i$
end while
for j with $s_j \neq 0$ **do**
  if $l(\operatorname{psimp} w_j) < k(\operatorname{psimp} s_j)$ **then**
    $B \leftarrow B \sqcup \{[l(\operatorname{psimp} w_j), k(\operatorname{psimp} s_j))\}$
  end if
end for
return B

Part III.

Q-tameness, Čech homology, and Morse theory for minimal surfaces

7. Morse inequalities in terms of persistence and minimal surfaces

In this chapter, we prove the results related to Morse's theory of functional topology presented in § 2.2.1. We start with our approach to Morse inequalities via persistence diagrams in Section 7.1. Afterwards, we discuss the proof of the unstable minimal surface theorem by Morse and Tompkins in Section 7.2.

In this chapter, we only consider persistent homology of sublevel set filtrations of real-valued functions, so we consider persistence modules indexed by $T = \mathbb{R}$. As usual, all homology groups are understood to be taken with field coefficients.

7.1. Morse inequalities in terms of persistence diagrams

We will now present our approach to Morse inequalities for persistence modules. In § 7.1.1, we consider the cap numbers of a q-tame persistence module and show that they satisfy Morse inequalities. We then discuss a condition for the finiteness of these numbers in § 7.1.2.

7.1.1. Morse inequalities for cap numbers

Classical Morse theory Recall that for a Morse function f on a closed smooth manifold X, the classical Morse inequalities [Mor96] state that for any non-negative integer n the following holds:

$$\sum_{d=0}^{n} (-1)^{n-d}\big(c_d(f) - \beta_d(X)\big) \;\geq\; 0, \tag{7.1.1}$$

where $c(f)$ is the number of critical points of f with index d and $\beta_d(X)$ is the d^{th} Betti number of X.

If no two critical points of f have the same value, the sets of critical points and values are naturally in one-to-one correspondence, which, in turn, are in one-to-one correspondence with the homological changes in the sublevel set filtration of f, i.e., the endpoints of the intervals appearing in the barcode of the persistent homology of $f_{\leq \bullet}$. More precisely, an index d critical point may either kill an existing homology class, in which case it corresponds to the right endpoint of an interval in the barcode of $H_{d-1}(f_{\leq \bullet})$, or it may give rise to a new homology class, in which case it corresponds to the left endpoint of an interval in the barcode of $H_d(f_{\leq \bullet})$.

The Betti numbers of X may also be expressed in terms of barcodes because they agree with the number of intervals that extend to $+\infty$. Thus, the Morse inequalities above can be expressed entirely in terms of the barcode (or persistence diagram) of the persistent homology of the sublevel set filtration of the function, which encodes the homological changes in the filtration.

Birth, death, and cap numbers The approach of counting homological changes instead of critical points described above is also what Morse used in the non-smooth setting of functional topology. To keep track of the number of d-dimensional homological events at filtration value t that persist for at least time $\epsilon > 0$ but not indefinitely, Morse [Mor40] defined the (d, t, ϵ)-*cap numbers* of a filtration. The definition given by Morse is specific to Vietoris homology and implicitly relies on the fact that the resulting persistence module is continuous from above. In terms of persistence barcodes, the (d, t, ϵ)-cap number correspond to the number of bars in the d^{th} barcode with left endpoint t and length greater than ϵ, plus the number of bars in the $(d-1)^{\text{th}}$ barcode with right endpoint t and length greater than ϵ. In the compact smooth setting, for sufficiently small ϵ, the (d, t, ϵ)-cap number equals the number of critical points of index d and value t, which either create homology in degree d or destroy homology in degree $(d-1)$. In [Mor40, Corollary 12.3], Morse proves a version of his eponymous inequalities using cap numbers as a replacement for numbers of critical points with the stated goal of making the inequalities applicable in settings where the function may not be smooth or the number of critical points may not be finite.

We now take a more general persistence-based approach that allows us to go beyond the setting of Vietoris homology. Working entirely in the algebraic setting, we will fix a graded q-tame persistence module M. Of course, one may think of M as the persistent homology of a q-tame filtration for any choice of homology theory, but M could for example also arise as the filtered Floer homology of some Hamiltonian on a symplectic manifold. Since M is q-tame, it has a persistence diagram in every degree d with multiplicity function denoted $\mathfrak{m}_d \colon \mathcal{E} \to \mathbb{N}$. In analogy to Morse's definitions, we may define, for an integer d and real numbers t and $\epsilon > 0$, the (d, t, ϵ)-*cap number* of our graded q-tame persistence module M in terms of its persistence diagram as

$$c_d^\epsilon(t) \ = \alpha_d^\epsilon(t) + \omega_{d-1}^\epsilon(t),$$

where

$$\alpha_d^\epsilon(t) = \sum_{\substack{q \in \mathbb{R} \cup \{\infty\} \\ q - t > \epsilon}} \mathfrak{m}_d(t, q), \qquad \omega_d^\epsilon(t) = \sum_{\substack{p \in \mathbb{R} \cup \{-\infty\} \\ t - p > \epsilon}} \mathfrak{m}_d(p, t)$$

are the *number of births* and the *number of deaths*, respectively, in degree d, at parameter t, and with persistence greater than ϵ. Note that finiteness of these quantities is ensured by the q-tameness of M and the use of a non-zero ϵ bounding below the persistence of the considered features. To see the necessity of this second condition, consider the q-tame persistence module given by the infinite product $\prod_{n \in \mathbb{N}_{>0}} C([0, 1/n))$ whose cap numbers $c^\epsilon(0)$ tend to ∞ as ϵ tends to 0.

Definition 7.1.1. Let M be a graded q-tame persistence module with persistence diagram given by $\mathfrak{m}$. Whenever the sums below are well-defined, we define the (d, ϵ)-*cap numbers*

$$c_d^\epsilon = \sum_t c_d^\epsilon(t) = \alpha_d^\epsilon + \omega_{d-1}^\epsilon,$$

where

$$\alpha_d^\epsilon = \sum_{t \in \mathbb{R}} \alpha_d^\epsilon(t) = \sum_{\substack{(p,q) \in \mathcal{E} \\ q - p > \epsilon \\ p \neq -\infty}} \mathfrak{m}_d(p, q), \qquad \omega_d^\epsilon = \sum_{t \in \mathbb{R}} \omega_d^\epsilon(t) = \sum_{\substack{(p,q) \in \mathcal{E} \\ q - p > \epsilon \\ q \neq \infty}} \mathfrak{m}_d(p, q)$$

are the *total number of births* and the *total number of deaths*, respectively, in degree d and with persistence greater than ϵ.

In the language of lifespan functors from Chapter 4, α_d^ϵ is the number of bars longer than ϵ in the barcode of the nascent part of the radical of the degree d part of M. Similarly, ω_d^ϵ is the number of bars longer than ϵ in the barcode of the mortal part of the radical of the degree d part of M.

Morse inequalities for cap numbers Comparing to the classical Morse inequalities, the cap numbers in dimension d act like the number of critical points with index d. As an analogue to the Betti numbers of the manifold appearing in the usual Morse inequalities, Morse defines quantities p_d referred to as *essential dimensions*, which, under the same assumptions as before, can be expressed in the language of persistence diagrams as

$$p_d \;=\; \sum_{p \in \mathbb{R} \cup \{-\infty\}} \mathfrak{m}_d(p, \infty),$$

which is also the dimension of the colimit of the degree d part of M.

Theorem 7.1.2. *Let $\epsilon > 0$, and let M be a graded q-tame persistence module with finite cap numbers c_d^ϵ and finite essential dimensions p_d for all d. If $\mathfrak{m}_d(-\infty, p) = 0$ for all $p \in \mathbb{R} \cup \{\infty\}$ and all d, then we have Morse inequalities*

$$\sum_{d=0}^{n} (-1)^{n-d}(c_d^\epsilon - p_d) \;\geq\; 0 \tag{7.1.2}$$

for any dimension n.

Proof. Recall that the dth ϵ-cap number is defined as

$$c_d^\epsilon = \alpha_d^\epsilon + \omega_{d-1}^\epsilon.$$

Since we assume $\mathfrak{m}_d(-\infty, p) = 0$ for all p, we have

$$p_d \;= \alpha_d^\epsilon - \omega_d^\epsilon$$

The difference of the two numbers, which appears in the Morse inequalities to be shown, is thus

$$c_d^\epsilon - p_d = \omega_{d-1}^\epsilon + \omega_d^\epsilon,$$

and their sum is

$$\sum_{d=0}^{n} (-1)^{n-d}(c_d^\epsilon - p_d) = \omega_n^\epsilon \geq\; 0$$

as claimed. □

Given the assertion that the persistence module M has a persistence diagram, one can interpret the previous proof as saying that the Morse inequalities simply express the (trivial) fact that the number of right endpoints in the persistence diagram is nonnegative. The simplicity of this observation illustrates the usefulness of interpreting fundamental facts in Morse theory through the lens of persistence theory.

As we will show in Theorem 7.1.5, the finiteness assumptions are satisfied if M is initially and eventually constant. Hence, as a special case, the theorem yields generalized Morse inequalities for any bounded real-valued function whose sublevel set filtration has q-tame persistent homology, including classical Morse functions f on closed smooth manifolds X. As outlined in our motivation for the definition of cap numbers, in this setting our inequalities (7.1.2) agree with the classical inequalities (7.1.1), as $\beta_d(X) = p_d$ and $c_d^\epsilon = \#\,\mathrm{crit}_d(f)$ for $\epsilon > 0$ smaller than the minimal difference between any two critical values of f. Morse inequalities for unbounded functions can still be obtained by restricting the function to an arbitrary sublevel set. Using Čech homology and considering bounded q-tame functions, our cap numbers and essential dimensions agree with the corresponding historical notions from [Mor40], so in this case our inequalities (7.1.2) also agree with the inequalities [Mor40, Corollary 12.3].

To apply our inequalities, one needs q-tameness, and we will give topological conditions that ensure q-tameness (Theorem 8.1.4). These conditions are in particular satisfied by the Douglas functional (Proposition 8.2.5), which we will review later on, and which motivated the developments in [Mor40].

Remark 7.1.3. In addition to the formulation of the inequalities in terms of cap numbers, Morse also proposed a generalized version of critical points, which he called homotopically critical, and which formalizes the idea of criticality of a point in terms of topological changes in the sublevel set filtration. This notion was employed by Morse and Tompkins in their study of minimal surfaces [MT39], which we will partly review in Section 7.2 and for which a thorough historical account may be found in [Str88, Section II.6].

The usefulness of this notion might however be limited in some cases of interest. In Floer theory, for example, critical points of the action functional corresponding to a Hamiltonian usually do not have a finite index and thus do not lead to a change in the homotopy type of sublevel sets. In this setting, our approach of formulating the inequalities purely in algebraic terms might be more suitable: while these critical points are not topologically visible, they do correspond to features in the persistence diagram of the filtered Floer homology.

7.1.2. Finiteness of cap numbers

The results in this subsection have previously appeared in the authors master's thesis. We repeat them for completeness, but do not claim any novelty.

An important setting where all cap numbers are well-defined is when M is the persistent homology of the sublevel set filtration of a bounded function. A more general statement can be made using the following notion.

Definition 7.1.4. A persistence module M is said to be *initially constant* if there is $s \in \mathbb{R}$ such that $M_{r,s}$ is an isomorphism for all $r \leq s$. Similarly, it is said to be *eventually constant* if there is $u \in \mathbb{R}$ such that $M_{u,v}$ is an isomorphism for all $u \leq v$.

Theorem 7.1.5. *Let M be a q-tame persistence module that is both initially and eventually constant. If $\mathfrak{m}$ is the multiplicity function of the persistence diagram of N, then for each $\epsilon > 0$,*

$$\sum_{\substack{(p,q)\in\mathcal{E} \\ q-p>\epsilon}} \mathfrak{m}(p,q) < \infty.$$

Proof. Let $s, u \in \mathbb{R}$ be as in Definition 7.1.4. We split the sum whose finiteness we want to show in two parts

$$\sum_{\substack{(p,q)\in\mathcal{E}\\ q-p>\epsilon}} \mathfrak{m}(p,q) \;=\; \sum_{\substack{(p,q)\in\mathcal{E}\setminus T\\ q-p>\epsilon}} \mathfrak{m}(p,q) \;+\; \sum_{\substack{(p,q)\in T\\ q-p>\epsilon}} \mathfrak{m}(p,q)$$

where T denotes the triangle

$$T = \{(p,q) \in \mathcal{E} \mid s \leq p < q \leq u\}.$$

For the first summand, observe that, because M is constant below s and constant above u, we have $\mathfrak{m}(p,q) = 0$ whenever one of $-\infty < p < s$ or $q < s$ or $p > u$ or $u < q < \infty$ holds. This implies

$$\sum_{\substack{(p,q)\in\mathcal{E}\setminus T\\ q-p>\epsilon}} \mathfrak{m}(p,q) \;=\; \sum_{s<q<u} \mathfrak{m}(-\infty,q) \;+\; \sum_{s<p<u} \mathfrak{m}(p,\infty)$$

which is clearly finite because M is q-tame.

For the second summand, note that

$$\sum_{\substack{(p,q)\in T\\ q-p>\epsilon}} \mathfrak{m}(p,q) \;\leq\; \sum_{(p,q)\in T^\epsilon} \mathfrak{m}(p,q),$$

where T^ϵ is the smaller triangle

$$T^\epsilon = \{(p,q) \in T \mid q - p \geq \epsilon\}.$$

Thus, in order to prove the theorem, it suffices to show that we have

$$\sum_{(p,q)\in T^\epsilon} \mathfrak{m}(p,q) \;<\; \infty.$$

To do this, we consider open quadrants

$$Q(x,y) = \{(p,q) \in \mathbb{R}^2 \mid p < x \text{ and } y < q\}.$$

Covering the compact set T^ϵ by the open quadrants $Q\left(x, x + \frac{\epsilon}{2}\right)$ for $x \in \mathbb{R}$, we may choose a finite subcover given by, say, $x_1, \ldots, x_n$. We obtain

$$\sum_{(p,q)\in T^\epsilon} \mathfrak{m}(p,q) \;\leq\; \sum_{i=1}^{n} \sum_{\substack{(p,q)\in\\ Q(x_i, x_i+\frac{\epsilon}{2})}} \mathfrak{m}(p,q).$$

Each of the sums $\sum_{(p,q)\in Q\left(x_i, x_i+\frac{\epsilon}{2}\right)} \mathfrak{m}(p,q)$ over the quadrants $Q\left(x_i, x_i + \frac{\epsilon}{2}\right)$ is finite since M is q-tame (which is where the name q-tame or *quadrant*-tame comes from, see [Cha+16, Section 3.8]). □

Of course, the persistent homology of a function, and even a bounded one, will generally not be bounded as a persistence module. For bounded functions, the persistent homology of the sublevel set filtration will be eventually constant with the homology of the domain of the function as value. If this homology of the domain is finite-dimensional, Theorem 7.1.5 can still be applied to obtain finiteness of cap numbers: one can split the persistent homology into its mortal and immortal part (see Chapter 4), apply Theorem 7.1.5 to the mortal part, and observe that the cap numbers of the immortal part are finite because the homology of the domain is finite-dimensional.

7.2. Persistence for minimal surfaces

Next, we present our approach to minimal surfaces via the persistent homology of the Douglas functional. Concretely, we will prove a mountain pass theorem for homotopically critical points of functionals in § 7.2.1. Our version of this theorem has slightly different assumptions and conclusions than the historical version by Morse and Tompkins that they used to prove the unstable minimal surface theorem, and we will discuss some of these differences in § 7.2.2. To finish, we deduce the unstable minimal surface theorem from our mountain pass theorem in § 7.2.3, barring the fact that the Douglas functional has q-tame persistent Čech homology, which we discuss in § 8.2.2.

7.2.1. A mountain pass theorem for homotopically critical points

Homotopically critical points As alluded to in Remark 7.1.3, Morse and Tompkins consider a homotopical notion of critical point for a general function $F\colon M \to \mathbb{R}$ on a metric space [MT39, p. 445], see also [Mor43]. Their goal was to apply this general theory to the Douglas functional to study minimal surfaces. As mentioned, we will discuss this in more detail Section 7.2.

Definition 7.2.1 ([Str88, Definition II.6.1-II.6.2], [MT39, p. 445, 466])**.** Consider a real-valued function F on a metric space (M, d). A point $p \in M$ is called *homotopically regular* if there exists a neighborhood U of p in $F_{\leq F(p)}$ and a continuous map $\varphi\colon U \times [0,1] \to M$, which satisfies $\varphi(\cdot, 0) = \mathrm{id}_U$ and $\varphi(p, 1) \neq p$, such that for every compact subset $V \subseteq U$ there exists a continuous *displacement function* $\delta\colon \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$. That is, a continuous function δ with $\delta(e) = 0$ if and only if $e = 0$ and

$$F(\varphi(x, s)) - F(\varphi(x, t)) \geq \delta(d(\varphi(x, s), \varphi(x, t)))$$

for all $x \in V$ and $0 \leq s \leq t \leq 1$. A point that is not homotopically regular is called *homotopically critical*. Function values of homotopically critical points will be called *critical values* and all other values will be called *regular values*. A *critical set* S is a closed and open subset of the subspace of all homotopically critical points with a given function value. It is said to be of *minimum type* if there exists a neighborhood N of the closure $\overline{S}$ of S, taken in M, such that the function values on $N \setminus S$ strictly exceed the function value on S.

Note that, in particular, an isolated local minimum constitutes a critical set of minimum type. Similarly, a critical submanifold of a Morse-Bott function on which the function values are locally minimized is also a critical set of minimum type.

A mountain pass theorem

Definition 7.2.2 ([MT39, p. 445])**.** Let $F\colon M \to \mathbb{R}$ be a function on a metric space. We say that the sublevel set filtration of F is *compact* if all sublevel sets of F are compact. We say that F is *weakly upper-reducible* if for all $p \in M$ and all $c > F(p)$ there exists a neighborhood U of p in $F_{\leq c}$, a positive constant $\eta > 0$, and a continuous map $\varphi\colon U \times [0,1] \to M$, which satisfies $\varphi(\cdot, 0) = \mathrm{id}_U$ and $\varphi(U, 1) \subseteq F_{\leq c-\eta}$, such that on every compact subset $V \subseteq U \cap F_{\geq c-\eta}$ there exists a displacement function for φ as in Definition 7.2.1.

In [Mor38, p. 36], the stronger notion of *upper-reducibility*, without the prefix *weakly*, is defined analogously to Definition 7.2.2 with the slight difference that the existence of a displacement function is not only demanded for compact $V \subseteq U \cap F_{\geq c-\eta}$, but for all compact $V \subseteq U$.

Morse and Tompkins state the following mountain pass theorem under slightly different assumptions, and with a slightly different conclusion.

Theorem 7.2.3 (Mountain pass theorem, [MT39, Corollary 7.1])**.** *Let $F\colon M \to \mathbb{R}$ be a weakly upper-reducible function on a non-empty connected metric space with compact sublevel set filtration. Assume that the natural map* $\operatorname{colim} \check{H}_0(F_{\leq \bullet}) \to \check{H}_0(M)$ *is an isomorphism, and that the sublevel set filtration of F is* LHS *with respect to Čech homology. If M contains two distinct critical sets of F of minimum type, then it also contains a critical set not of minimum type.*

What it means for the sublevel set filtration of F to be LHS will be defined in Definition 8.1.2. This condition on F is used to ensure that $\check{H}_0(F_{\leq \bullet})$ is q-tame (Theorem 8.1.4).

We note that, under the conditions on F from Theorem 7.2.3, the Morse inequalities (7.1.2) apply to $\check{H}_*(F_{\leq \bullet})$, which can be checked with the help of Theorems 7.1.5 and 8.1.4 and Lemma 7.2.10 and the fact that F must be bounded below if its sublevel sets are compact. Using Morse inequalities is also the approach of Morse and Tompkins to the mountain pass theorem: they relate homotopically critical points to cap numbers and then interpret the mountain pass theorem as a version of the Morse inequalities (7.1.2) for $n = 1$. Our approach will be similar, but in the formulation of Theorem 7.2.3, the mountain pass theorem is actually a stronger statement than the Morse inequalities for $n = 1$. We will comment on this in more detail at the end of § 7.2.2.

It is also worth noting that the details in [MT39] are incomplete, with some crucial theorems such as [MT39, Theorems 7.3 and 7.4, Corollary 7.1] being stated without proof, and with a citation to a paper in preparation that has never been published under the given name (we suppose that this paper is [Mor40]). Moreover, there is a gap in [Mor40], because [Mor40, Theorem 6.3], which establishes q-tameness in the setting of Morse–Tompkins, is incorrect as we will show in Corollary 8.2.3. The assumptions we choose for our version of the mountain pass theorem are adapted from the original assumptions to the modern language of persistence theory, and they fix the problem with q-tameness. Still, our assumptions can be established for the Douglas functional from the results of Morse and Tompkins [MT39]. We provide more detailed comments regarding the differences between these assumptions and the conclusions of the mountain pass theorem in Subsections 7.2.2 and 8.2.1.

Remark 7.2.4. We mention that more general homology theories can be considered in Theorem 7.2.3. The precise hypotheses on the homology theory H are that it is additive, taking non-zero values on non-empty sets in dimension 0, and such that $H(F_{\leq \bullet})$ is continuous from above, i.e., $H(F_{\leq t}) \to \lim_{u>t} H(F_{\leq u})$ is an isomorphism for all $t \in \mathbb{R}$. One may also replace the topological conditions of M being connected, the natural map $\operatorname{colim} H_0(F_{\leq \bullet}) \to H_0(M)$ being an isomorphism, and F being LHS by the algebraic conditions that $p_0 = 1$ and $H(F_{\leq \bullet})$ be q-tame.

Towards a proof of the mountain pass theorem The primary application that Morse and Tompkins had in mind for the mountain pass theorem is the unstable minimal surface theorem that we review in § 7.2.3. While more efficient and more general proofs for the existence of an unstable minimal surface have subsequently been established [DHS10; Str88], including less restrictive assumptions on the boundary curve, the original approach of Morse and Tompkins is notable for its close connection to persistent homology and the generality in which it is applicable in principle. As an illustration, we will now give a proof of Theorem 7.2.3 following Morse and Tompkins [MT39] using the previously developed language of persistence, filling in some gaps left by Morse and Tompkins. As a first step towards proving Theorem 7.2.3, we give general formulas for *numbers of births* $\alpha(t) = \sum_{q\in(t,\infty]} \mathfrak{m}(t,q)$ and *numbers of deaths* $\omega(t) = \sum_{p\in[-\infty,t)} \mathfrak{m}(p,t)$ in a q-tame persistence module M with persistence diagrams given by the multiplicity function $\mathfrak{m}$.

Lemma 7.2.5. *Let M be a q-tame persistence module. For each $t \in \mathbb{R}$ we either have*

$$\alpha(t) = \dim \operatorname{coker} \left(\operatorname*{colim}_{s<t} M_s \to \lim_{u>t} M_u \right)$$

or both quantities are infinite. Similarly, we either have

$$\omega(t) = \dim \ker \left(\operatorname*{colim}_{s<t} M_s \to \lim_{u>t} M_u \right)$$

or both quantities are infinite.

Proof. The internal colimits $\operatorname{colim}_{s<t} M_s$ and limits $\lim_{u>t} M_u$ are invariant under weak isomorphisms by Lemmas 10.1.9 and 10.1.10, and the same is true for the numbers of births and deaths $\alpha(t)$ and $\omega(t)$ because they are defined through persistence diagrams, which themselves are invariant under weak isomorphisms. Hence, using the fact that the inclusion $\operatorname{rad} M \hookrightarrow M$ is a weak isomorphism, we conclude that none of the quantities appearing in the statement of the lemma change if we replace M by its radical. This radical admits a barcode decomposition and the claims follow from an explicit computation on the corresponding barcode module $M' := \bigoplus_{\lambda\in\Lambda} C(I_\lambda)$:

First, note that we have

$$\dim \lim_{u>t} M'_u = \sharp \{\lambda \in \Lambda \mid \inf I_\lambda \leq t,\ \sup I_\lambda > t\}$$

if one side of the equation is finite, and we also have

$$\dim \operatorname*{colim}_{s<t} M'_s = \sharp \{\lambda \in \Lambda \mid \inf I_\lambda < t,\ \sup I_\lambda \geq t\},$$

as well as

$$\dim \operatorname{im} \left(\operatorname*{colim}_{s<t} M'_s \to \lim_{u>t} M'_u \right) = \sharp \{\lambda \in \Lambda \mid \inf I_\lambda < t,\ \sup I_\lambda > t\}.$$

From this, if again one side of the equation is finite, we obtain

$$\begin{aligned}
\alpha(t) &= \sum_{q\in(t,\infty]} \mathfrak{m}(t,q) \\
&= \sharp \{\lambda \in \Lambda \mid \inf I_\lambda = t,\ \sup I_\lambda > t\} \\
&= \sharp \{\lambda \in \Lambda \mid \inf I_\lambda \leq t,\ \sup I_\lambda > t\} - \sharp \{\lambda \in \Lambda \mid \inf I_\lambda < t,\ \sup I_\lambda > t\} \\
&= \dim \lim_{u>t} M'_u - \dim \operatorname{im} \left(\operatorname*{colim}_{s<t} M'_s \to \lim_{u>t} M'_u \right) \\
&= \dim \operatorname{coker} \left(\operatorname*{colim}_{s<t} M'_s \to \lim_{u>t} M'_u \right)
\end{aligned}$$

and we also obtain

$$\begin{aligned}
\omega(t) &= \sum_{p\in[-\infty,t)} \mathfrak{m}(p,t) \\
&= \sharp\{\lambda \in \Lambda \mid \inf I_\lambda < t,\ \sup I_\lambda = t\} \\
&= \sharp\{\lambda \in \Lambda \mid \inf I_\lambda < t,\ \sup I_\lambda \geq t\} - \sharp\{\lambda \in \Lambda \mid \inf I_\lambda < t,\ \sup I_\lambda > t\} \\
&= \dim \operatorname*{colim}_{s<t} M'_s - \dim \operatorname{im}\left(\operatorname*{colim}_{s<t} M'_s \to \lim_{u>t} M'_u\right) \\
&= \dim \ker\left(\operatorname*{colim}_{s<t} M'_s \to \lim_{u>t} M'_u\right). \qquad \square
\end{aligned}$$

If M is continuous from below or above (see Definition 10.1.1), then the colimits and limits appearing in the formulas above may simply be replaced by the constituent vector spaces of M. As a special case, this yields the following corollary.

Corollary 7.2.6. *Let M be a q-tame persistence module. If M is continuous from above and below at t, we have*

$$\alpha(t) = \omega(t) = 0.$$

Topological evolution at regular values We now return to the setting of functions F on metric spaces and prove some more lemmas. To emphasize which persistence-theoretic notions are relevant, we will work with a general homology theory H such that $H(F_{\leq \bullet})$ has certain properties, stated in the lemmas. In fact, beyond these properties, H need not actually be a homology theory in the usual sense of the word, but can be any graded vector space valued homotopy invariant functor. In all cases, Čech homology as used in Theorem 7.2.3 satisfies the necessary conditions. We start with a first result stating that, as in the case of smooth Morse theory, the homotopy type of sublevel sets does not change leading up to regular values. A weaker version of Lemma 7.2.7 is stated in [Mor38, Lemma 8.1], under the slightly stronger assumption of upper-reducibility. In [MT39], Morse and Tompkins introduce and use *weak* upper-reducibility, noting that the results and arguments from [Mor38, Sections 7 and 8] still apply. For convenience of the reader, we revisit the relevant arguments using our setting, loosely following the exposition by Struwe [Str88, Remark II.6.3] (who is working under yet another slightly different set of assumptions).

Lemma 7.2.7. *Let $F\colon M \to \mathbb{R}$ be a weakly upper-reducible function a metric space with compact sublevel set filtration. If $t \in \mathbb{R}$ is a regular value of F, then there exists $\epsilon > 0$ such that the inclusion $F_{\leq s} \hookrightarrow F_{\leq t}$ is a homotopy equivalence for all $s \in (t-\epsilon, t]$.*

Proof. Let $t \in \mathbb{R}$ be a regular value of F. For each $p \in F_{\leq t}$ we can find a number $\delta_p > 0$ and a continuous map $\varphi_p\colon B_{\delta_p}(p) \times [0,1] \to M$, where $B_{\delta_p}(p)$ is the open metric ball around p in $F_{\leq t}$ with radius δ_p, such that

1. $\varphi_p(\cdot, 0) = \mathrm{id}$,
2. for all compact $V \subseteq B_{\delta_p}(p)$ there exists $\epsilon(p,V) > 0$ with
    a) $\varphi_p(V \cap F_{\leq s} \times [0,1]) \subseteq F_{\leq s}$ for all $s \in (t-\epsilon(p,V), t]$, and
    b) $\varphi(V,1) \subseteq F_{\leq t-\epsilon(p,V)}$.

For p with $F(p) < t$ the existence of such δ_p and φ_p is guaranteed by the assumption that F is weakly upper-reducible. For p with $F(p) = t$ it follows from the assumption that t is a regular value and hence p is homotopically regular, by the following argument.

If p is homotopically regular, there exists a neighborhood U of p in $F_{\leq t}$ and a continuous map $\varphi\colon U \times [0,1] \to M$ such that $\varphi(\cdot,0) = \mathrm{id}_U$, $\varphi(p,1) \neq p$, and such that on every compact set $V \subseteq U$ there exists a displacement function for φ and F. It is clear that $\varphi(V \cap F_{\leq s} \times [0,1]) \subseteq F_{\leq s}$ for all $s \leq t$ and $V \subseteq U$ compact since there exists a displacement function and hence the values of F can not increase along trajectories of φ. It remains to be checked that we can shrink U enough such that for all compact $V \subseteq U$ there is $\epsilon(p,V)$ with $\varphi(V,1) \subseteq F_{\leq t-\epsilon(p,V)}$. Assume for a contradiction that for all neighborhoods $W \subseteq U$ there exists a compact $V \subseteq W$ such that for all $\epsilon > 0$ there is some $q \in V$ with $t \geq F(q) \geq F(\varphi(q,1)) > t - \epsilon$, where we note that $F(q) \geq F(\varphi(q,1))$ is a consequence of the existence of a displacement function on V. It follows that there exists a sequence $(q_n)_n$ in $F_{\leq t}$ with $F(q_n) \to t$ and $q_n \to p$ for $n \to \infty$. Since φ is continuous we have that $\varphi(q_n,1)$ converges to $\varphi(p,1)$, which is different from p by the choice of φ, so there can be no infinite subsequence of $(q_n)_n$ consisting of fixed points for $\varphi(\cdot,1)$. Without loss of generality, we can thus assume that $\varphi(q_n,1) \neq q_n$ for all n. Now consider the compact set $K = \{q_n \mid n \in \mathbb{N}\} \cup \{p\}$. By assumption, we can choose a continuous displacement function δ on K. Since K is compact, φ is continuous, δ is continuous, and $\varphi(\cdot,1)$ has no fixed points in K, the function $x \mapsto \delta(d(x,\varphi(x,1)))$ attains a minimum $\epsilon > 0$ on K. We get $F(x) - t \geq F(x) - F(\varphi(x,1)) \geq \epsilon > 0$ for all $x \in K$, which contradicts the fact that $F(q_n) \to t$ for $n \to \infty$. This proves the existence of the required δ_p and φ_p for p homotopically regular with $F(p) = t$.

Because $F_{\leq t}$ is compact, we can now choose finitely many points $p_1, \dots, p_n$ such that $F_{\leq t}$ is covered by $(B_{\frac{\delta_{p_i}}{3}}(p_i))_i$, and we set

$$\epsilon := \max_{i=1,\dots,n} \epsilon\left(p_i, \overline{B}_{\frac{\delta_{p_i}}{3}}(p_i)\right).$$

We extend the corresponding homotopies φ_{p_i} from $B_{\delta_{p_i}}(p_i)$ to all of $F_{\leq t}$, not changing them on $B_{\frac{\delta_{p_i}}{3}}(p_i)$, and being the identity on $F_{\leq t} \setminus B_{\delta_{p_i}}(p_i)$: Let $\psi\colon [0,1] \to [0,1]$ be the continous map given by

$$\psi(x) = \begin{cases} 1 & \text{if } x \in [0, \frac{1}{3}), \\ -3x + 2 & \text{if } x \in [\frac{1}{3}, \frac{2}{3}), \\ 0 & \text{if } x \in [\frac{2}{3}, 1]. \end{cases}$$

We define $\tilde{\varphi}_i\colon F_{\leq t} \times [0,1] \to F_{\leq t}$ by

$$\tilde{\varphi}_i(x,r) = \begin{cases} \varphi_{p_i}\left(x, \psi\left(\frac{d(x,p_i)}{\delta_{p_i}}\right) \cdot r\right) & \text{if } x \in B_{\delta_{p_i}}(p_i), \\ x & \text{otherwise.} \end{cases}$$

Finally, we define $\varphi\colon F_{\leq t} \times [0,1] \to F_{\leq t}$ as the concatenation of the maps $\tilde{\varphi}_i$, i.e., $\varphi(x,0) = x$ and $\varphi(x,r) = \varphi_i\left(x, n \cdot r - i\right)$ for $r \in \left(\frac{i}{n}, \frac{i+1}{n}\right]$. From the properties of the maps φ_{p_i} chosen in the beginning and the construction of the maps $\tilde{\varphi}_i$, we obtain that

1. $\varphi(\cdot,0) = \mathrm{id}$,
2. $\varphi(F_{\leq s} \times [0,1]) \subseteq F_{\leq s}$ for all $s \in (t-\epsilon, t]$, and

3. $\varphi(F_{\leq t}, 1) \subseteq F_{\leq t-\epsilon}$.

This shows that, indeed, for $s \in (t-\epsilon, t]$ the inclusion $F_{\leq s} \hookrightarrow F_{\leq t}$ is a homotopy equivalence, with homotopy inverse given by φ. □

Topological evolution at critical values We proceed by showing that function values with non-vanishing *cap numbers* $c_d(t) \coloneqq \alpha_d(t) + \omega_{d-1}(t)$, for any d, are indeed critical values. This corresponds to [Mor38, Theorem 8.1]. Note that the definition of cap numbers used by Morse is phrased in terms of relative homology, differing from our definition in terms of absolute homology. The equivalence of both definitions is shown in Proposition 7.2.12.

Lemma 7.2.8. *Let $F\colon M \to \mathbb{R}$ be a weakly upper-reducible function on a metric space with compact sublevel set filtration. Assume that $H(F_{\leq \bullet})$ is q-tame and continuous from above, and consider $t \in \mathbb{R}$ and its birth and death numbers $\alpha_d(t)$ and $\omega_{d-1}(t)$. If $\alpha_d(t) > 0$ or $\omega_{d-1}(t) > 0$ for some degree d, then t is a critical value of F.*

Proof. Following Lemma 7.2.7, we know that if t is a regular value, there exists $\epsilon > 0$ such that the inclusion $F_{\leq s} \hookrightarrow F_{\leq t}$ is a homotopy equivalence for all $s \in (t-\epsilon, t]$. Thus, $H(F_{\leq \bullet})$ is continuous from below at every regular value. However, we assume $H(F_{\leq \bullet})$ to be also continuous from above at every value, and in particular at regular values. Hence, Corollary 7.2.6 implies that $\alpha_d(t) = \omega_{d-1}(t) = 0$ for all d whenever t is regular, which proves the claim. □

Next, we will analyze how the homology of sublevel sets changes at function values of critical sets of minimum type.

Lemma 7.2.9. *Let $F\colon M \to \mathbb{R}$ be a weakly upper-reducible function on a metric space with compact sublevel set filtration and let S be a critical set of minimum type with value t. Assume that H is additive and that $H(F_{\leq \bullet})$ is q-tame and continuous from above.*

1. *The number of births at t satisfies $\alpha_d(t) \geq \dim H_d(S)$ for all d.*

2. *If there are no homotopically critical points with value t outside S, then the number of deaths at t satisfies $\omega_d(t) = 0$ for all d.*

Proof. We start by showing that S is a topological summand of $F_{\leq t}$ in the sense that $F_{\leq t}$ is homeomorphic to the disjoint union $S \sqcup (F_{\leq t} \setminus S)$. It suffices to show that S is open and closed in $F_{\leq t}$. By definition, there exists a neighborhood N of $\overline{S}$ in M such that the function values of F on $N \setminus S$ exceed t. In particular, we have $F_{\leq t} \cap N = S$, showing that S is open in $F_{\leq t}$. Because N contains $\overline{S}$, we also obtain $F_{\leq t} \cap \overline{S} = S$, showing that S is closed in $F_{\leq t}$.

Using additivity of H and Lemma 7.2.5, we now obtain

$$\begin{aligned}
\alpha_d(t) &= \dim \operatorname{coker}\big(\operatorname*{colim}_{s<t} H_d(F_{\leq s}) \to \lim_{u>t} H_d(F_{\leq u})\big) \\
&= \dim \operatorname{coker}\big(\operatorname*{colim}_{s<t} H_d(F_{\leq s}) \to H_d(F_{\leq t})\big) \\
&= \dim \operatorname{coker}\big(\operatorname*{colim}_{s<t} H_d(F_{\leq s}) \to H_d(F_{\leq t} \setminus S) \oplus H_d(S)\big) \\
&\geq \dim H_d(S),
\end{aligned}$$

where we have used the assumption that $H(F_{\leq \bullet})$ is continuous from above for the second equality and the fact that $F_{\leq s} \subseteq (F_{\leq t} \setminus S)$ for all $s < t$ for the final inequality.

Now, assume that S is the set of all homotopically critical points of F with value t and consider the restriction of F to $F_{\leq t} \setminus S$, denoted by G. The set $F_{\leq t} \setminus S$ is compact because $F_{\leq t}$ is compact and S is open in $F_{\leq t}$, so the sublevel set filtration of G is compact. G is also weakly upper-reducible because the same is true for F. Moreover, any homotopically critical point of G is a homotopically critical point of F, so our assumption that S is the set of all homotopically critical points with value t of F implies that t is a regular value for G. By Lemma 7.2.7, this implies that there exists $\epsilon > 0$ such that the inclusion $F_{\leq s} = G_{\leq s} \hookrightarrow G_{\leq t} = (F_{\leq t} \setminus S)$ is a homotopy equivalence for for all $s \in (t - \epsilon, t]$ Hence, again using additivity of H, continuity from above, and Lemma 7.2.5, we obtain that

$$\begin{aligned}
\omega_d(t) &= \dim\ker\big(\operatorname*{colim}_{s<t} H_d(F_{\leq s}) \to \lim_{u>t} H_d(F_{\leq u})\big) \\
&= \dim\ker\big(\operatorname*{colim}_{s<t} H_d(F_{\leq s}) \to H_d(F_{\leq t})\big) \\
&= \dim\ker\big(\operatorname*{colim}_{s<t} H_d(F_{\leq s}) \to \operatorname*{colim}_{s<t} H_d(F_{\leq s}) \oplus H_d(S)\big) \\
&= 0
\end{aligned}$$

as claimed. □

Final arguments As the last preparatory result, we will show that connectedness of M together with the condition on the colimit of the persistent homology of F in the setting of Theorem 7.2.3 imply that the 0-th essential dimension is trivial for Čech homology.

Lemma 7.2.10. *Let $F\colon M \to \mathbb{R}$ be a function on a connected non-empty metric space with compact sublevel set filtration such that the natural map* $\operatorname{colim} \check{H}_0(F_{\leq \bullet}) \to \check{H}_0(M)$ *is an isomorphism and $\check{H}_0(F_{\leq \bullet})$ is q-tame. Then the essential dimension of $\check{H}_0(F_{\leq \bullet})$ satisfies $\check{p}_0 = 1$.*

Proof. Since M is connected and non-empty, we have $\dim H_0(M) = 1$ for singular homology, and since M is non-empty we have $\dim \check{H}_0(M) \geq 1$. Now the natural map from singular to Čech homology is always surjective for compact metric spaces in dimension 0 [EK00], so the morphism $H_0(F_{\leq \bullet}) \to \check{H}_0(F_{\leq \bullet})$ is epi. Taking direct limits of vector spaces is exact, so the natural map $\operatorname{colim} H_0(F_{\leq \bullet}) \to \operatorname{colim} \check{H}_0(F_{\leq \bullet})$ is still epi. We assume $\operatorname{colim} \check{H}_0(F_{\leq \bullet}) \to \check{H}_0(M)$ to be an isomorphism, so the composition $\operatorname{colim} H_0(F_{\leq \bullet}) \to \check{H}_0(M)$ is still epi. But this map factors through $H_0(M)$, which has dimension 1, so we get $\dim \check{H}_0(M) = 1$. Together with our assumption on the colimit, we obtain $\check{p}_0 = \dim \operatorname{colim} \check{H}_0(F_{\leq \bullet}) = \dim \check{H}_0(M) = 1$. □

Remark 7.2.11. Note that $\check{p}_0 = 1$ does not already follow from just M being connected, i.e., the assumption $\operatorname{colim} \check{H}_0(F_{\leq \bullet}) \to \check{H}_0(M)$ is non-vacuous: For the lower semi-continuous function $F\colon [0,1] \to \mathbb{R}$ defined as $F(0) = 0$ and $F(t) = \frac{1}{t}$ for $t > 0$, we have for $t > 0$ that $F_{\leq t} = \{0\} \cup [\frac{1}{t}, 1]$. Hence, we get $\check{p}_0 = \dim \operatorname{colim} \check{H}_0(F_{\leq \bullet}) = 2$ despite M being contractible and the sublevel sets of F being compact.

We are now ready to give a proof of the mountain pass theorem.

Proof of Theorem 7.2.3. First, we note that $\check{H}_*(F_{\leq\bullet})$ is q-tame by Theorem 8.1.4 because we assume that $F_{\leq\bullet}$ is LHS for Čech homology. As a consequence, $\check{H}_*(F_{\leq\bullet})$ has a well-defined persistence diagram, and thus we can consider cap numbers, births, deaths, etc. We write $\check{c}_d$, $\check{\alpha}_d$, $\check{\omega}_d$, and $\check{p}_d$ for the cap numbers, births, deaths, and essential dimensions of $\check{H}_d(F_{\leq\bullet})$, respectively. The sublevel set filtration of $F_{\leq\bullet}$ is also assumed to be compact, M is assumed to be connected and non-empty, $\operatorname{colim} \check{H}_0(F_{\leq\bullet}) \to \check{H}_0(M)$ is assumed to be an isomorphism, and F is assumed to be weakly upper-reducible, so the previous lemmas are all applicable.

Now assume that F has two distinct critical sets S_1 and S_2 of minimum type with values t_1 and t_2, respectively. Since both critical sets are non-empty, we have $\check{H}_0(S_i) \neq 0$, and thus the first assertion of Lemma 7.2.9 implies $\check{\alpha}_0(t_i) \geq \dim \check{H}_0(S_i) \geq 1$ for $i = 1, 2$, indicating the existence of at least one feature for each $i = 1, 2$ with birth t_i and some positive persistence $\epsilon_i > 0$ in the persistence diagram of $\check{H}_0(F_{\leq\bullet})$. Choosing $0 < \epsilon < \min_i \epsilon_i$, this implies that $\check{\alpha}_0^\epsilon(t_i) \geq 1$ for $i = 1, 2$, and thus $\check{c}_0^\epsilon \geq \check{\alpha}_0^\epsilon(t_1) + \check{\alpha}_0^\epsilon(t_2) \geq 1 + 1 = 2$. Now, from Lemma 7.2.10 we obtain that $\check{p}_0 = 1$, which yields $\check{\omega}_0^\epsilon = \check{c}_0^\epsilon - \check{p}_0 \geq 2 - 1 = 1$. Thus, there must be some $t \in \mathbb{R}$ with $\check{\omega}_0(t) \geq \check{\omega}_0^\epsilon)(t) > 0$, so that we may apply Lemma 7.2.8 to obtain that the set S of homotopically critical points at value t is non-empty. If S were of minimum type, then we would have $\check{\omega}_0(t) = 0$ by the second assertion of Lemma 7.2.9, contradicting the choice of t. Hence, S cannot be of minimum type, which finishes the proof. □

7.2.2. Some historical comparisons

Equivalence of definitions of cap numbers We will now discuss some differences and equivalences between our approach and that of Morse and Tompkins. As a starting point for the comparison of the notion of cap number, we observe that the basic use of the homotopy lemma Lemma 7.2.7 is to show that if t is an endpoint of a feature in the persistence diagram associated to our function F, i.e., if $c(t) > 0$, then there must be a homotopically critical point with value t. To achieve this, we assume that F is weakly upper-reducible and its sublevel set filtration is compact. A very similar statement appears in [Mor38, Theorem 8.1], which says that if t is a *cap limit*, then there must be a homotopically critical point with value t. There are two slight differences to our approach: First, in [Mor38, Theorem 8.1] it is assumed that F is not only weakly upper-reducible but upper-reducible. Second, t is not required to satisfy $c(t) > 0$, but required to be a *cap limit*, meaning that $\operatorname{colim}_{s:s<t} \check{H}(F_{\leq t}, F_{\leq s}) \neq 0$ [Mor38, p. 12]. We have already discussed the first point when introducing Lemma 7.2.7. Regarding the second point, we now show that our cap numbers are non-zero for some t if and only if t is a cap limit.

Proposition 7.2.12. *Let $F \colon M \to \mathbb{R}$ be a weakly upper-reducible function on a metric space with compact sublevel set filtration and fix $t \in \mathbb{R}$. Assume that H has long exact sequences for pairs of compact spaces, and that $H_*(F_{\leq\bullet})$ is q-tame and continuous from above. Then $\dim \operatorname{colim}_{s:s<t} H_d(F_{\leq t}, F_{\leq s}) = c_d(t)$ or both quantities are infinite.*

Proof. First, note that we have a splitting

$$\operatorname*{colim}_{s:s<t} H_d(F_{\leq t}, F_{\leq s}) \cong \operatorname{im}\left(H_d(F_{\leq t}) \to \operatorname*{colim}_{s:s<t} H_*(F_{\leq t}, F_{\leq s})\right) \oplus \operatorname{coker}\left(H_d(F_{\leq t}) \to \operatorname*{colim}_{s:s<t} H_d(F_{\leq t}, F_{\leq s})\right).$$

Computing with the long exact sequences of pairs in homology, we have

$$\begin{aligned}
\alpha_d(t) &= \dim \operatorname{coker}\left(\operatorname*{colim}_{s:s<t} H_d(F_{\leq s} \hookrightarrow F_{\leq t})\right) \\
&= \dim \operatorname*{colim}_{s:s<t} \operatorname{coker}\left(H_d(F_{\leq s} \hookrightarrow F_{\leq t})\right) \\
&= \dim \operatorname*{colim}_{s:s<t} \operatorname{im}\left(H_d(F_{\leq t}) \to H_d(F_{\leq t}, F_{\leq s})\right) \\
&= \dim \operatorname{im}\left(H_d(F_{\leq t}) \to \operatorname*{colim}_{s:s<t} H_d(F_{\leq t}, F_{\leq s})\right),
\end{aligned}$$

and

$$\begin{aligned}
\omega_{d-1}(t) &= \dim \ker\left(\operatorname*{colim}_{s:s<t} H_{d-1}(F_{\leq s} \hookrightarrow F_{\leq t})\right) \\
&= \dim \operatorname*{colim}_{s:s<t} \ker\left(H_{d-1}(F_{\leq s} \hookrightarrow F_{\leq t})\right) \\
&= \dim \operatorname*{colim}_{s:s<t} \operatorname{coker}\left(H_d(F_{\leq t}) \to H_d(F_{\leq t}, F_{\leq s})\right) \\
&= \dim \operatorname{coker}\left(H_d(F_{\leq t}) \to \operatorname*{colim}_{s:s<t} H_d(F_{\leq t}, F_{\leq s})\right),
\end{aligned}$$

where we have in both cases used the fact that taking filtered colimits of vector spaces is exact, and where the first equalities hold in case both sides of the equation are finite because of Lemma 7.2.5 and the fact that $H(F_{\leq \bullet})$ is continuous from above. Since $c_d(t) = \alpha_d(t) + \omega_{d-1}(t)$, this finishes the proof. □

Remark 7.2.13. Proposition 7.2.12 only references cap limits, i.e., levels t for which the *cap space* $\operatorname{colim}_{s:s<t} H_d(F_{\leq t}, F_{\leq s})$ is non-zero. However, Morse not only considers these cap spaces on their own, but he also defines the *span* of a cap, and considers the space of caps with span greater than some $\epsilon > 0$, see for example [Mor40, Section 11]. Using slight generalizations of the arguments proving Proposition 7.2.12 and Lemma 7.2.5, one can show that the dimensions of these spaces agree with $c_d^\epsilon(t)$. We already mentioned this when talking about Morse inequalities, which Morse originally formulated for dimensions of such cap spaces.

Note that the assumptions of Proposition 7.2.12 are in particular satisfied for Čech homology with field coefficients and a function F whose sublevel set filtration is LHS with respect to this homology theory. To continue the comparison between our approach to the mountain pass theorem and the one by Morse and Tompkins, we will now discuss several differences between the requisite definitions and the respective assumptions and conclusions.

Continuity of displacement functions Our Definition 7.2.1 of homotopically regular point differs slightly from that of Morse [Mor38, p. 30] and Morse and Tompkins [MT39, p. 444–445] as their displacement functions are not explicitly assumed to be continuous. As used in the proof of Lemma 7.2.7, we require this property to establish that a homotopically regular point with value t admits a neighborhood U and a homotopy φ as in Definition 7.2.1 such that for any compact $V \subseteq U$ there is $\epsilon > 0$ with $\varphi(V, 1) \subseteq F_{\leq t-\epsilon}$. That such neighborhoods exist is claimed by Morse in the proof of [Mor38, Lemma 8.1], where continuity of the displacement functions is not mentioned explicitly. It is not clear to us whether the claim actually holds without continuity, or whether the assumption of continuity was made implicitly. Regardless, all displacement functions considered by Morse and Tompkins [MT39] and used for the purposes of this work are indeed continuous.

Critical sets of minimum type Our Definition 7.2.1 of a critical set S of *minimum type* also differs slightly from that of Morse and Tompkins [MT39, p. 466], who do not require the neighborhood N of S on which the function values exceed those on S to contain the closure of S. Without this additional assumption, however, Theorem 7.2.3 does not hold, as shown by the example $F\colon [0,1] \to \mathbb{R}$ with $F(0) = F(1) = 0$ and $F(t) = 1$ for $0 < t < 1$: F has the four critical sets $\{0\}$, $\{1\}$, $\{0\} \cup \{1\}$ and $(0,1)$, which all satisfy the minimum type condition if the neighborhood N need not contain their closure, but then there is no critical set that is not of minimum type.

Assumptions of the mountain pass theorem In [MT39, Corollary 7.1], the assumptions that Morse and Tompkins use for F and M are that the sublevel set filtration is compact, weakly upper-reducible, "regular at infinity", and that M is "locally F-connected". Compactness is also required for our version, as is weak upper-reducibility. Local F-connectedness is replaced by the LHS condition (this point will be discussed in more detail in § 8.2.1). Regularity at infinity roughly corresponds to our assumption that M is connected and the natural map $\operatorname{colim} \check{H}_0(F_{\leq \bullet}) \to \check{H}_0(M)$ is an isomorphism (see [MT39, p. 444] and our proof of Theorem 7.2.14). This assumption may also be replaced by the assumption that $\check{p}_0 = 1$, which we deduce from our assumptions in Lemma 7.2.10.

Conclusions of the mountain pass theorem The precise statement of [MT39, Corollary 7.1] postulates the existence of a homotopic critical point such that every critical set containing it has a positive first *type number*. This corresponds to our statement in so far as the existence of a critical set not of minimum type implies that there are ϵ and t such that $c_1^\epsilon(t) > 0$ (see Lemma 7.2.8 and the proof of Theorem 7.2.3). Note, however, that the existence of a positive cap number $c_1^\epsilon(t) > 0$ does not conversely imply the existence of a critical set not of minimum type: Consider the parabola $P = \{(x, x^2) \in \mathbb{R}^2 \mid x \in \mathbb{R}\}$, let $M = S^1 \times P$ and let $F\colon M \to \mathbb{R}$ be the Morse–Bott function given by the projection to the second coordinate in P. There is a single critical set $C = S^1 \times \{(0,0)\}$, which is of minimum type, but we still have $c_1^\epsilon(0) > 0$ for any $\epsilon > 0$ if H is a homology theory such that $H_1(C) \neq 0$, e.g., singular or Čech homology. To establish the existence of a critical set not of minimum type in the proof of Theorem 7.2.3, we show that not only $c_1^\epsilon(t) > 0$, but that we indeed have $\omega_0^\epsilon(t) > 0$, i.e., that the positive cap number is created by the death of a feature. This also illustrates that our version of the mountain pass theorem is stronger than the Morse inequalities (7.1.2) for $n = 1$, which merely imply $c_1^\epsilon(t) > 0$. In particular, we have mentioned before that the Morse inequalities are equivalent to the fact that numbers of deaths are non-negative, but we need one of them to be positive for our mountain pass theorem to hold.

7.2.3. The unstable minimal surface theorem

The Douglas functional Morse and Tompkins [MT39] considered the following setting introduced by Douglas. Let $g\colon \mathbb{R} \to \mathbb{R}^n$ be a 2π-periodic function representing a simple closed curve such that g is differentiable with Lipschitz derivative. Let $\widetilde{\Omega}$ be the space of continuous non-decreasing functions $\varphi\colon \mathbb{R} \to \mathbb{R}$ with $\varphi(t + 2\pi) = \varphi(t) + 2\pi$ for all t and $\varphi(\alpha_i) = \alpha_i$ for three fixed distinct points $\alpha_i \in [0, 2\pi)$. The *Douglas functional* on $\widetilde{\Omega}$ associated to the curve g is defined [Dou31] as

$$A_g(\varphi) = \frac{1}{16\pi} \int_0^{2\pi} \int_0^{2\pi} \left\| \frac{g(\varphi(\alpha)) - g(\varphi(\beta))}{\sin \frac{\alpha - \beta}{2}} \right\|_2^2 \mathrm{d}\alpha \, \mathrm{d}\beta.$$

The Douglas functional evaluated at φ is equal to the Dirichlet energy of the unique harmonic extension of the reparametrized curve $g \circ \varphi$ to a parametrized surface. The Dirichlet energy is an upper bound for the area, with equality if the parametrization is conformal. Let $\Omega_g = \{\varphi \in \widetilde{\Omega} \mid A_g(\varphi) < \infty\}$, equipped with the C^0 metric. The set Ω_g is non-empty because g is continuously differentiable and hence rectifiable, which implies $A_g(\mathrm{id}_{\mathbb{R}}) < \infty$ [Dou31, p. 267-268]. Moreover, the sublevel sets of A_g are compact [MT39, p. 448]. Since A_g is bounded below by 0, this implies that A_g attains a global minimum. The corresponding surface is then a solution of *Plateau's problem*, which asks for a surface homeomorphic to a disk with boundary g and minimum area.

The unstable minimal surface theorem Morse and Tompkins show [MT39, Theorem 6.2] that each homotopically critical point of the Douglas functional A_g indeed corresponds to a *minimal surface* – a surface with vanishing mean curvature – and using this correspondence the following result, also reviewed in [Str88, Theorem II.6.10], can be deduced from Theorem 7.2.3.

Theorem 7.2.14 (Unstable minimal surface theorem [MT39, p. 472])**.** *If the space Ω_g contains two minimal surfaces contained in distinct critical sets of minimum type of the functional A_g, then it also contains an* unstable *minimal surface, i.e., a minimal surface contained in a critical set that is not of minimum type.*

Proof. The Douglas functional A_g is weakly upper-reducible by [MT39, Theorem 5.1], its sublevel set filtration is compact according to [MT39, p. 448], and as previously mentioned Ω_g is non-empty by [Dou31, p. 267-268] because g is rectifiable. Moreover, Ω_g is contractible and hence connected by [MT39, Theorem 4.3]. The sublevel set filtration of A_g is also LHS for Čech homology according to Proposition 8.2.5. Finally, we have that colim $\check{H}_0((A_g)_\bullet) \to \check{H}_0(\Omega_g)$ is an isomorphism according to [MT39, p. 444] because A_g satisfies the *regularity at infinity* condition by [MT39, Theorem 4.3]. In total, Theorem 7.2.3 applies to A_g. Since any homotopically critical point of A_g corresponds to a minimal surface spanned by g [MT39, Theorem 6.2], this implies the claim. □

In [MT39, Section 8], Morse and Tompkins provide an example of a curve g for which Ω_g indeed contains two distinct critical sets of minimum type, so that g then also spans an unstable minimal surface.

What is left to complete our discussion of the mountain pass theorem and the unstable minimal surface theorem is why the LHS condition implies q-tameness (Theorem 8.1.4, why the Douglas functional is q-tame for Čech homology (Proposition 8.2.5), and what the issue with Morse's approach to q-tameness in [Mor40, Theorem 6.3] is (Corollary 8.2.3).

8. Q-tameness for locally connected filtrations

In this chapter, we prove the main result on q-tameness for locally connected filtrations from § 2.2.2 and discuss the issues with Morse's local connectedness conditions. The proof of our result that a locally homologically small filtration is q-tame is presented in Section 8.1. We then discuss the shortcomings of Morse's approach in Section 8.2, where we also comment on the q-tameness of the Douglas functional, which is the last missing ingredient in the proof of the unstable minimal surface theorem discussed in Section 7.2.

All results are formulated in terms of sublevel set filtrations of real-valued functions, so we consider persistence modules indexed by $T = \mathbb{R}$. However, all statements in § 8.1.1, and in particular Theorem 8.1.4 naturally generalize to filtrations indexed by arbitrary totally ordered sets with identical proofs. As usual, all homology groups are understood to be taken with field coefficients, but Theorem 8.1.4 can also be generalized to more general coefficient groups. Note that, the homology theories appearing this chapter need not strictly be homology theories in the usual sense; what we mean precisely by homology in the present context is made precise in § 8.1.1. For some parts, we will also consider only Čech homology.

8.1. Persistence diagrams for locally homologically small filtrations

In this section, we present our findings on topological conditions guaranteeing q-tameness. We start by discussing that locally homologically small filtrations are q-tame in § 8.1.1. For sublevel set filtrations induced by continuous maps, a slightly weaker condition suffices, which we discuss in § 8.1.2.

8.1.1. Local connectedness and q-tameness

We have seen that q-tame functions admit persistence diagrams, which can be used to formulate Morse inequalities and to prove a mountain pass theorem as discussed in Chapter 7. With q-tameness being a rather abstract algebraic property, we now establish concrete topological conditions that ensure the q-tameness of a function. Our definitions are motivated by similar conditions considered by Morse in his work on functional topology. We present a historical account in Section 8.2.

Mayer-Vietoris property Whether a function is q-tame or not depends on the functor that is used to pass from the sublevel set filtration of the function to a persistence module. The general idea that we want to employ is to deduce global finiteness properties like q-tameness from local ones. Thus, the functors we consider should have a certain property allowing us to do so.

Definition 8.1.1. Let $H = (H_d)_{d\in\mathbb{Z}}\colon \mathbf{Top} \to \mathbf{Vect}$ be a fixed graded homotopy invariant functor, which we call a *homology theory.* A triple of spaces $X_1, X_2 \subseteq X$ is said to have a *Mayer–Vietoris sequence* for H if the inclusion-induced maps can be completed to a long exact sequence

$$\begin{array}{ccccccc} \cdots \longrightarrow H_{n+1}(X) & & & & & & \\ \downarrow & & & & & & \\ H_n(X_1 \cap X_2) & \longrightarrow & H_n(X_1) \oplus H_n(X_2) & \longrightarrow & H_n(X) & & \\ & & & & \downarrow & & \\ & & & & H_{n-1}(X_1 \cap X_2) & \longrightarrow & \cdots \ . \end{array}$$

We say that H has the *open* (resp. *compact*) *Mayer–Vietoris property* if there are natural Mayer–Vietoris sequences for all triples $X_1, X_2 \subseteq X$ with $X = X_1 \cup X_2$ and $X_i \subseteq X$ open (resp. compact Hausdorff).

For the rest of this section, we will assume that H is a homology theory that has either the open or the compact Mayer–Vietoris property and for which there is n_0 such that H_n is 0 for all $n \leq n_0$. We will also assume that H_n takes finite-dimensional values on one-point spaces. Note that this includes singular homology with field coefficients, which has the open Mayer–Vietoris property (like any homology theory in the sense of the Eilenberg–Steenrod axioms [ES52, Section I]), and it also includes Čech homology with field coefficients, which has the compact Mayer–Vietoris property (see § 1.4.1).

Locally homologically small filtrations

Definition 8.1.2. The sublevel set filtration of a function $f\colon X \to \mathbb{R}$ is called *locally homologically small* or LHS with respect to a homology theory H if for any $x \in X$, any neighborhood V of x, and any pair of indices s, t with $f(x) < s < t$ there is a neighborhood U of x with $U \subseteq V$ such that the inclusion $f_{\leq s} \cap U \to f_{\leq t} \cap V$ is *homologically small* or HS, i.e., has finite rank in every degree after applying H.

Definition 8.1.3. We say that a sublevel set filtration is *compact* if all sublevel sets are compact Hausdorff spaces.

If $f_{\leq t}$ is compact for all t, then the function f is necessarily lower-semicontinuous and bounded from below (see [MT39, p. 444] or [Str88, Theorem 3.1]).

Our main result, proven in the remainder of this subsection, is that for compact sublevel set filtrations the LHS condition implies q-tameness, and consequently also the existence of a persistence diagram.

Theorem 8.1.4. *If the sublevel set filtration of a function $f\colon X \to \mathbb{R}$ is compact and* LHS*, then it is also q-tame. In particular, f has a persistence diagram.*

The general proof strategy is inspired by the proof of Wilder's Finiteness Theorem [Wil49, p. 325] as presented by Bredon [Bre97, Section II.17]. We collect the main ideas in several lemmas.

Lemma 8.1.5. *Given a commutative diagram of modules over a principal ideal domain*

$$\begin{array}{ccccc}
A_{1,1} & \longrightarrow & A_{1,2} & & \\
\uparrow & & \uparrow & & \\
A_{2,1} & \longrightarrow & A_{2,2} & \longrightarrow & A_{2,3} \\
 & & \uparrow & & \uparrow \\
 & & A_{3,2} & \longrightarrow & A_{3,3}
\end{array}$$

where the middle row is exact and both $A_{2,1} \to A_{1,1}$ and $A_{3,3} \to A_{2,3}$ have finitely generated images, then so does $A_{3,2} \to A_{1,2}$.

Proof. This is proven via a straightforward diagram chase. For more details see [Bre97, Lemma II.17.3]. □

Lemma 8.1.6. *Let X be a locally compact Hausdorff space. For any compact subset K and open set U with $K \subseteq U$ there exists a compact set K' such that*

$$K \subseteq \operatorname{int}(K') \subseteq K' \subseteq U.$$

Proof. For any $x \in K$ choose a compact neighborhood $C(x) \subseteq U$. We have

$$K \subseteq \bigcup_{x \in K} \operatorname{int}(C(x)).$$

Since K is compact, there is a finite subset $\{x_1, \ldots, x_m\}$ of elements in K so that

$$K \subseteq \bigcup_{i=1}^{m} \operatorname{int}(C(x_i)) \subseteq \operatorname{int}\left(\bigcup_{i=1}^{m} C(x_i)\right) \subseteq \bigcup_{i=1}^{m} C(x_i) \subseteq U$$

Defining $K' = \bigcup_{i=1}^{m} C(x_i)$ finishes the proof. □

We want to use Lemma 8.1.6 on the domain of the function whose sublevel set filtration we consider. However, we do not want to assume the domain to be locally compact for Theorem 8.1.4. To circumvent this, we will work in one of the sublevel sets, which are assumed to be compact Hausdorff and hence locally compact. This requires the use of a slight weakening of the LHS condition.

Definition 8.1.7. For $u \in \mathbb{R}$, the sublevel set filtration of a function $f\colon X \to \mathbb{R}$ is called LHS *below* u if, for any $x \in X$, any neighborhood V of x, and any pair of indices s, t with $f(x) < s < t < u$, there is a neighborhood U of x with $U \subseteq V$ such that the inclusion $f_{\leq s} \cap U \hookrightarrow f_{\leq t} \cap V$ is HS.

Lemma 8.1.8. *Let $f\colon X \to \mathbb{R}$ be a function whose sublevel set filtration is* LHS*. Fix $u \in \mathbb{R}$ and let $g\colon Y \to \mathbb{R}$ be the restriction of f to the sublevel set $Y = f_{\leq u}$. Then the sublevel set filtration defined by g is* LHS *below u.*

Proof. Let $x \in Y$, let V be a neighborhood of x in Y, and consider indices s, t with $g(x) < s < t < u$. We need to find a neighborhood $U \subseteq V$ of x such that the inclusion $g_{\leq s} \cap U \hookrightarrow g_{\leq t} \cap V$ is HS.

Since $Y \subseteq X$ carries the subspace topology, we may choose a neighborhood V' of x in X such that $V = V' \cap Y$. The sublevel set filtration of f is assumed to be LHS, so there is a neighborhood $U' \subseteq V'$ of x in X such that the inclusion $f_{\leq s} \cap U' \hookrightarrow f_{\leq t} \cap V'$ is HS. We set $U = U' \cap Y$, which defines a neighborhood of x in Y.

Now $s < t < u$ implies that $g_{\leq s} = f_{\leq s}$ and $g_{\leq t} = f_{\leq t}$. Moreover, we have $f_{\leq s} \cap Y = f_{\leq s} \cap f_{\leq u} = f_{\leq s}$ and $f_{\leq t} \cap Y = f_{\leq t} \cap f_{\leq u} = f_{\leq t}$. Thus, we obtain $g_{\leq s} \cap U = f_{\leq s} \cap Y \cap U' = f_{\leq s} \cap U'$ and $g_{\leq t} \cap V = f_{\leq t} \cap Y \cap V' = f_{\leq t} \cap V'$. This implies that the inclusion $g_{\leq s} \cap U \hookrightarrow g_{\leq t} \cap V$ is HS because it agrees with the inclusion $f_{\leq s} \cap U' \hookrightarrow f_{\leq t} \cap V'$, which is HS by assumption. This finishes the proof. □

Lemma 8.1.9. *Let $f \colon Y \to \mathbb{R}$ be a function on a locally compact Hausdorff space Y whose sublevel set filtration is compact and* LHS *below $u \in \mathbb{R}$, and consider subsets $C \subseteq L \subseteq Y$ with C compact and L open. For any $s < t < u$ the inclusion $C \cap f_{\leq s} \hookrightarrow L \cap f_{\leq t}$ is* HS.

Proof. Recall our assumption that the underlying homology theory H has either the open or the compact Mayer–Vietoris property and that there is some n_0 such that H_n is zero for all $n \leq n_0$. The statement of the lemma holds for HS_n in place of HS for any $n \leq n_0$ since H_n induces the zero map. We will proceed by induction on $n \geq n_0$ assuming the statement for HS_{n-1}.

We define $\Sigma_{s,t}$ to be the collection of all open subsets $V \subseteq Y$ whose closure $\overline{V}$ is compact, contained in L, and has an open neighborhood U with $\overline{V} \subseteq U \subseteq L$ for which there exists $s' \in (s, t)$ such that the inclusion $U \cap f_{\leq s'} \hookrightarrow L \cap f_{\leq t}$ is HS_n. We will show that $\Sigma_{s,t}$ has the following two properties:

1. Any point $x \in L \cap f_{\leq s}$ has a neighborhood $V_x \in \Sigma_{s,t}$.

2. If $V_1, V_2 \in \Sigma_{s,t}$ then $V_1 \cup V_2 \in \Sigma_{s,t}$.

Assuming them for the moment, the first property allows us to cover $C \cap f_{\leq s}$ by sets $V_x \in \Sigma_{s,t}$, $x \in C \cap f_{\leq s}$. Because both C and $f_{\leq s}$ are compact, $C \cap f_{\leq s}$ is again compact, and hence the cover can be chosen finite, represented by say $x_1, \dots, x_m$. By the second property, we have

$$V \coloneqq \bigcup_{i=1}^{m} V_{x_i} \in \Sigma_{s,t}.$$

Thus, there exists an open neighborhood U of $\overline{V}$ in L and $s' \in (s, t)$ such that the inclusion $U \cap f_{\leq s'} \hookrightarrow L \cap f_{\leq t}$ is HS_n. The inclusion $C \cap f_{\leq s} \hookrightarrow L \cap f_{\leq t}$ factors through the previous one, so it is HS_n as well. What is left to do is to show that $\Sigma_{s,t}$ has the two claimed properties.

Next, we will show using the LHS property that $\Sigma_{s,t}$ has the first required property, i.e., that any point $x \in L \cap f_{\leq s}$ has a neighborhood in $\Sigma_{s,t}$. Choose an arbitrary $s' \in (s, t)$. Since the sublevel set filtration of f is LHS below u and we have $f(x) \leq s < s' < t < u$, there is an open neighborhood $U_x \subseteq L$ such that the inclusion $U_x \cap f_{\leq s'} \hookrightarrow L \cap f_{\leq t}$ is HS, so in particular HS_n. By local compactness of Y we can choose a compact neighborhood K_x of x contained in U_x. Now $V_x = \operatorname{int}(K_x)$ is a neighborhood of x with $V_x \in \Sigma_{s,t}$.

Finally, using the Mayer–Vietoris property and the induction hypothesis we will show that $\Sigma_{s,t}$ has the second required property, i.e., that it is closed under finite unions. So for

$i \in \{1,2\}$ let $V_i \in \Sigma_{s,t}$ with U_i and $s'_i \in (s,\, t)$ such that $\overline{V_i} \subseteq U_i \subseteq L$ and $U_i \cap f_{\leq s'_i} \hookrightarrow L \cap f_{\leq t}$ is HS_n. Writing $K_i = \overline{V_i}$, we use Lemma 8.1.6 to construct compact sets K'_i such that

$$V_i \subseteq K_i \subseteq V'_i \subseteq K'_i \subseteq U_i \subseteq L$$

where $V'_i = \mathrm{int}(K'_i)$. The union $V_1 \cup V_2 \subseteq L$ is open, its closure $\overline{V_1 \cup V_2}$ is compact, and $\overline{V_1 \cup V_2} \subseteq V'_1 \cup V'_2 \subseteq L$. Thus, we obtain $V_1 \cup V_2 \in \Sigma_{s,t}$ if we can show that there is an $s' \in (s,\, t)$ such that the inclusion $(V'_1 \cup V'_2) \cap f_{\leq s'} \hookrightarrow L \cap f_{\leq t}$ is HS_n. To do so, we set $s'' = \min_i s'_i$ and choose $s' \in (s,\, s'')$. For proving that $(V'_1 \cup V'_2) \cap f_{\leq s'} \hookrightarrow L \cap f_{\leq t}$ is HS_n we now distinguish the two cases where H has either the open or the compact Mayer–Vietoris property.

For the open Mayer–Vietoris property, note that for both $i \in \{1,2\}$ the inclusions $U_i \cap f_{\leq s''} \hookrightarrow L \cap f_{\leq t}$ are HS_n. Moreover, the inclusion $V'_1 \cap V'_2 \cap f_{\leq s'} \hookrightarrow U_1 \cap U_2 \cap f_{\leq s''}$ is HS_{n-1} because it factors through the inclusion $K'_1 \cap K'_2 \cap f_{\leq s'} \hookrightarrow U_1 \cap U_2 \cap f_{\leq s''}$, which is HS_{n-1} by the induction hypothesis. Because the V_i and V'_i are open and because H has the open Mayer–Vietoris property, we obtain the following commutative diagram satisfying the assumptions of Lemma 8.1.5:

$$\begin{array}{ccccc}
H_n(L \cap f_{\leq t}) \oplus H_n(L \cap f_{\leq t}) & \longrightarrow & H_n(L \cap f_{\leq t}) & & \\
\uparrow & & \uparrow & & \\
H_n(U_1 \cap f_{\leq s''}) \oplus H_n(U_2 \cap f_{\leq s''}) & \longrightarrow & H_n((U_1 \cup U_2) \cap f_{\leq s''}) & \longrightarrow & H_{n-1}(U_1 \cap U_2 \cap f_{\leq s''}) \\
 & & \uparrow & & \uparrow \\
 & & H_n((V'_1 \cup V'_2) \cap f_{\leq s'}) & \longrightarrow & H_{n-1}(V'_1 \cap V'_2 \cap f_{\leq s'}).
\end{array}$$

We conclude that the inclusion $(V'_1 \cup V'_2) \cap f_{\leq s'} \hookrightarrow L \cap f_{\leq t}$ is HS_n, which finishes this part of the proof.

For the compact Mayer–Vietoris property, we apply Lemma 8.1.6 once more to obtain compact sets K''_i such that

$$V_i \subseteq K_i \subseteq V'_i \subseteq K'_i \subseteq V''_i \subseteq K''_i \subseteq U_i \subseteq L$$

where $V''_i = \mathrm{int}(K''_i)$. The rest of the proof is then analogous to the previous case: We have that for both $i \in \{1,2\}$ the inclusion $K''_i \cap f_{\leq s''} \hookrightarrow L \cap f_{\leq t}$ is HS_n because it factors through $U_i \cap f_{\leq s''} \hookrightarrow L \cap f_{\leq t}$. Moreover, the inclusion $K'_1 \cap K'_2 \cap f_{\leq s'} \hookrightarrow K''_1 \cap K''_2 \cap f_{\leq s''}$ is HS_{n-1} because it factors through the inclusion $K'_1 \cap K'_2 \cap f_{\leq s'} \hookrightarrow V''_1 \cap V''_2 \cap f_{\leq s''}$, which is HS_{n-1} by the induction hypothesis. Because the K'_i and K''_i as well as the sublevel sets of f are all compact and because H has the compact Mayer–Vietoris property, we obtain the following commutative diagram satisfying the assumptions of Lemma 8.1.5:

$$\begin{array}{ccccc}
H_n(L \cap f_{\leq t}) \oplus H_n(L \cap f_{\leq t}) & \longrightarrow & H_n(L \cap f_{\leq t}) & & \\
\uparrow & & \uparrow & & \\
H_n(K''_1 \cap f_{\leq s''}) \oplus H_n(K''_2 \cap f_{\leq s''}) & \rightarrow & H_n((K''_1 \cup K''_2) \cap f_{\leq s''}) & \rightarrow & H_{n-1}(K''_1 \cap K''_2 \cap f_{\leq s''}) \\
 & & \uparrow & & \uparrow \\
 & & H_n((K'_1 \cup K'_2) \cap f_{\leq s'}) & \rightarrow & H_{n-1}(K'_1 \cap K'_2 \cap f_{\leq s'}).
\end{array}$$

We conclude that the inclusion $(K'_1 \cup K'_2) \cap f_{\leq s'} \hookrightarrow L \cap f_{\leq t}$ is HS_n, and so the same is true for the inclusion $(V'_1 \cup V'_2) \cap f_{\leq s'} \hookrightarrow L \cap f_{\leq t}$ as it factors through the previous one. $\square$

We can now complete the proof of the claim stating that for compact sublevel set filtrations, LHS implies q-tameness.

Proof of Theorem 8.1.4. By definition, the sublevel set filtration of f is q-tame if and only if the inclusion $f_{\leq s} \hookrightarrow f_{\leq t}$ is HS for all pairs $s < t$. Choose $u \in \mathbb{R}$ with $u > t$ and let $g\colon Y \to \mathbb{R}$ be the restriction of f to the sublevel set $Y = f_{\leq u}$. Since we assume f to induce a LHS sublevel set filtration, by Lemma 8.1.8 the sublevel set filtration of g is LHS below u. Clearly, the sublevel set filtration of g is also compact, and its domain Y is locally compact being a compact Hausdorff space by assumption. Thus, we can apply Lemma 8.1.9 to the filtration $g_{\leq \bullet}$ with $C = L = Y$ to obtain that the inclusion $f_{\leq s} = C \cap g_{\leq s} \hookrightarrow L \cap g_{\leq t} = f_{\leq t}$ is HS. □

8.1.2. The case of a continuous function

We now describe a weaker version of the LHS property that implies q-tameness for sublevel set filtration induced by continuous functions.

Definition 8.1.10. The sublevel set filtration of a function $f\colon X \to \mathbb{R}$ is said to be *weakly locally homologically small* or *weakly* LHS if for any $x \in X$, any neighborhood V of x, and any index $t > f(x)$, there is an index s with $f(x) < s < t$ and a neighborhood U of x with $U \subseteq V$ such that the inclusion $f_{\leq s} \cap U \to f_{\leq t} \cap V$ is HS.

Clearly, any LHS sublevel set filtration is also weakly LHS: while the weak LHS property merely requires the existence of an index $s \in (f(x), t)$ satisfying the HS condition, the LHS property requires the HS condition to hold for any $s \in (f(x), t)$. If the filtration is induced by a continuous function, the converse also holds as the following theorem shows.

Lemma 8.1.11. *If the sublevel set filtration of a continuous function $f\colon X \to \mathbb{R}$ is weakly* LHS*, then it is also* LHS*.*

Proof. Fix $x \in X$, a neighborhood V of x and indices $f(x) < s < t$. We need to show that there is a neighborhood $U \subseteq V$ of x such that the inclusion $f_{\leq s} \cap U \to f_{\leq t} \cap V$ is HS.

To do so, we start by using the weak LHS property to choose a neighborhood $U' \subseteq V$ of x and an index $s' \in (f(x),\, t)$ such that the inclusion $f_{\leq s'} \cap U' \to f_{\leq t} \cap V$ is HS. Now, we choose $U = f_{<s'} \cap U'$, where $f_{<s'} = f^{-1}(-\infty, s')$. Note that this choice of U still defines a neighborhood of x because f is assumed to be continuous, so that $f_{<s'}$ is an open subset of X.

We obtain that $f_{\leq s} \cap U \subseteq f_{\leq s'} \cap U'$, so that the inclusion $f_{\leq s} \cap U \to f_{\leq t} \cap V$ factors through the inclusion $f_{\leq s'} \cap U' \to f_{\leq t} \cap V$. This second map is HS by construction. Any map that factors through an HS map is also HS, so the proof is complete. □

The following result is deduced directly from Lemma 8.1.11 and Theorem 8.1.4. The existence of a result of this kind has been suggested by Weinberger [Wei11], and a multiparameter version has been shown by Cagliari and Landi [CL11] with slightly stronger assumptions on the domain of the function.

Corollary 8.1.12. *If the sublevel set filtration of a continuous function $f\colon X \to \mathbb{R}$ is compact and weakly* LHS*, then it is also q-tame.*

As the following example illustrates, the continuity assumption in the above theorem is crucial.

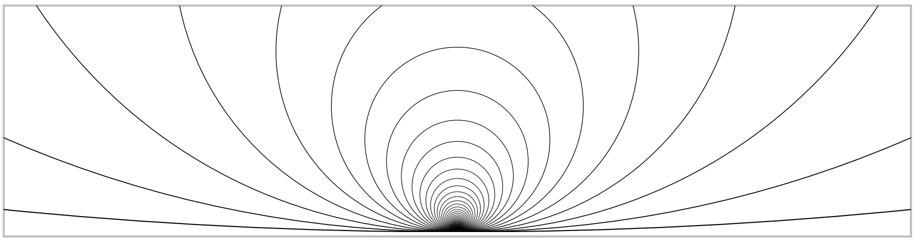

Figure 8.1.: A closeup of the Hawaiian earring $\mathbb{H}^1$.

Example 8.1.13. Consider the *d-dimensional Hawaiian earring*

$$\mathbb{H}^d = \bigcup_{n\in\mathbb{N}} \left\{ (x_0, \dots, x_d) \in \mathbb{R}^{d+1} \;\middle|\; \left(x_0 - \frac{1}{n}\right)^2 + x_1^2 + \cdots + x_d^2 = \left(\frac{1}{n}\right)^2 \right\},$$

which is a compact subspace of $\mathbb{R}^{d+1}$. The function $f\colon \mathbb{H}^d \to \mathbb{R}$ whose value at the origin is 0 and is 1 everywhere else defines a compact and weakly LHS sublevel set filtration that is not q-tame with respect to H if $H_n(\mathbb{H}^d)$ is infinite-dimensional for some n.

To verify that f has compact sublevel sets we notice that all sublevel sets are either the empty set, the singleton containing the origin, or $\mathbb{H}^d$ itself, all compact Hausdorff spaces.

In order to verify that the sublevel set filtration of f is weakly LHS, we consider $x \in \mathbb{H}^d$, V a neighborhood of x in $\mathbb{H}^d$ and $t > f(x)$. We need to find a neighborhood $U \subseteq V$ of x and $s \in (f(x), t)$ such that the inclusion $f_{\leq s} \cap U \to f_{\leq t} \cap V$ is HS. Since we assume that H_n is homotopy invariant and takes singletons to finite-dimensional spaces, it suffices to find U as above such that $f_{\leq s} \cap U \to f_{\leq t} \cap V$ is homotopic to a constant map.

If x is the origin, we have $f(x) = 0$ and choose $s \in (0, \min\{t, 1\})$. Then $f_{\leq s} = \{x\}$, so with $U = V$ the inclusion $f_{\leq s} \cap U \to f_{\leq t} \cap V$ is the inclusion of $\{x\}$ into $f_{\leq t} \cap V$, which is a constant map, so the weak LHS condition is trivially satisfied.

For x not the origin we have $f(x) = 1$ and choose $s \in (1, t)$ arbitrarily, so that $f_{\leq s} = f_{\leq t} = \mathbb{H}^d$. Note that since x is not the origin, there is a unique d-sphere in $\mathbb{H}^d$ that contains x. Clearly, we may choose $\delta > 0$ so small that $B_\delta(x) = \{y \in \mathbb{R}^{d+1} \mid \|x - y\| < \delta\} \cap \mathbb{H}^d$ is a topological ball contained in this sphere and contained in V. The ball $B_\delta(x)$ can be contracted to $\{x\}$ in V, so choosing $U = B_\delta(x)$, we obtain that the inclusion $f_{\leq s} \cap U \to f_{\leq t} \cap V$ is homotopic to the constant map with value x.

It remains to be shown that $f_{\leq \bullet}$ is not q-tame for H. This follows directly from our assumption that $H_n(\mathbb{H}^d)$ is not finite-dimensional for some n, as $f_{\leq t}$ is constant with value $\mathbb{H}^d$ for $t \geq 1$.

8.2. Local connectedness in functional topology

After having established our own results on q-tameness, we will now come back to the motivating problems in Morse's functional topology and the applications to minimal surfaces. We will compare the local connectedness conditions that Morse historically wanted to use to ensure q-tameness to our own, and highlight a gap in his treatment in § 8.2.1. To finish this chapter, we discuss that our previous results can be used to circumvent this gap in the main application of functional topology to the Douglas functional in § 8.2.2.

8.2.1. Morse's local connectedness conditions

We have mentioned in § 2.2.2 that Morse used different local connectedness conditions in his work on functional topology to ensure q-tameness, for example the local F-connectedness condition from [Mor38; Mor40] that we quoted in § 2.2.2. Using similar language to the one used for our LHS condition, local F-connectedness is equivalent to the following notion applicable to general topological spaces.

Definition 8.2.1. The sublevel set filtration of a function $f\colon X \to \mathbb{R}$ is said to be *weakly locally connected of all orders*, or *weakly* πLC, if for any $x \in X$, V a neighborhood of x, and any index $t > f(x)$, there is an index s with $f(x) < s < t$ and a neighborhood U of x with $U \subseteq V$ such that the inclusion $f_{\leq s} \cap U \to f_{\leq t} \cap V$ induces trivial maps on homotopy groups.

Morse claims in [Mor38, Theorem 6.1] and [Mor40, Theorem 6.3, p. 432] that the persistent Čech homology of this sublevel set filtration is q-tame, provided that F is bounded from below and satisfies the assumptions of local F-connectedness and compactness of sublevel sets. Unfortunately, this claim does not hold in general, as exemplified by the sublevel set filtration from Example 8.1.13. To elaborate on this, we consider a stronger version of weak local connectedness.

Definition 8.2.2. The sublevel set filtration of a function $f\colon X \to \mathbb{R}$ is said to be *weakly locally contractible*, or *weakly* LC, if for any $x \in X$, V a neighborhood of x, and any index $t > f(x)$, there is an index s with $f(x) < s < t$ and a neighborhood U of x with $U \subseteq V$ such that the inclusion $f_{\leq s} \cap U \to f_{\leq t} \cap V$ is homotopic to a constant map.

Clearly, being weakly LC implies being weakly πLC and, if the homology H takes finite-dimensional values on one-point spaces, also weakly LHS. Observe that Example 8.1.13 actually establishes that the filtration given there is weakly LC, so not even the weak LC condition is sufficient to ensure the q-tameness of compact sublevel set filtrations that are induced by non-continuous functions in general. In particular, our construction invalidates Morse's claim quoted before that weak πLC implies q-tameness for Čech homology because Čech homology satisfies the assumptions on the homology theory made in Example 8.1.13.

Specifically, using the fact that Čech homology of compact Hausdorff spaces commutes with inverse limits, it is straightforward to verify that the Čech homology in degree d of the d-dimensional Hawaiian earring $\mathbb{H}^d$ is isomorphic to $\prod_{n \in \mathbb{N}} \mathbb{F}$, which is infinite-dimensional over $\mathbb{F}$. Moreover, the singular homology of $\mathbb{H}^d$ is also infinite-dimensional, as proven in [BM62]. In summary, we have the following result contradicting [Mor40, Theorem 6.3, p. 432].

Corollary 8.2.3. *The function $f\colon \mathbb{H}^d \to \mathbb{R}$ with value 0 at the origin and 1 elsewhere defines a weakly* LC *compact sublevel set filtration that is not q-tame with respect to either singular or Čech homology.*

8.2.2. Local connectedness of the Douglas functional

Corollary 8.2.3 highlights a gap in the argument of Morse and Tompkins on minimal surfaces, as the sublevel set filtration of the Douglas functional A_g is not actually shown to be q-tame. Luckily, this gap can be readily fixed by applying Theorem 8.1.4. This is because the proof given in [MT39, Theorem 7.2, p.464] for the local connectedness of the sublevel set filtration induced by A_g can actually be seen to establish a stronger property, described next.

Definition 8.2.4. The sublevel set filtration of a function $f\colon X \to \mathbb{R}$ is said to be *locally contractible* or LC if for any $x \in X$, any neighborhood V of x and any pair of indices $f(x) < s < t$ there is a neighborhood $U \subseteq V$ of x such that the map $f_{\leq s} \cap U \to f_{\leq t} \cap V$ is homotopic to a constant map.

Proposition 8.2.5 ([MT39, p.464])**.** *The Douglas functional $A_g\colon \Omega_g \to \mathbb{R}$ induces an* LC *sublevel set filtration. In particular, it is* LHS *and hence q-tame for Čech homology.*

Morse introduced another condition that he also called local F-connectedness in an earlier article. It roughly corresponds to being πLC with a certain added uniformity property. In the original it reads:

> The space M will be said to be locally F-connected for the order n if corresponding to n, an arbitrary point p on M, and an arbitrary positive constant e, there exists a positive constant δ with the following property. For $c \geq F(p)$ any singular n-sphere on $F \leq c$ (the continuous image on $F \leq c$ of an ordinary n-sphere) on the δ-neighborhood p_δ of p is the boundary of a singular $(n+1)$-cell on $F \leq c + e$ and on p_e. [Mor37, p.421–422]

Morse also claims in the given reference that this condition is sufficient for q-tameness, but without providing a proof. Whether this statement is true or not is not covered by our analysis, because the πLC and LHS conditions generally do not imply each other. We expect the quoted claim to be true, but do not investigate it further.

9. Comparison of singular and Čech homology in locally connected filtrations

In this chapter, we prove the comparison results for singular and Čech homology presented in § 2.2.3. We start the proofs by constructing maps from singular to Čech homology and back in Section 9.1. We then construct certain chain homotopies on singular and Čech complexes in Section 9.2, showing that the previously defined maps are in some sense inverses to each other. Finally, we put these results together to prove our main theorems in Section 9.3, including the short proof that the local connectedness shift of a function is 2-Lipschitz.

Throughout the chapter, we fix an abelian group G and use G as a coefficient group for all chain complexes and homology groups that appear while dropping it from the notation most of the time. Note that, in contrast to the rest of this thesis, this coefficient group does not need to be a field here. This means that the persistent homologies appearing in this chapter are not persistence modules in the sense that we have considered before, but rather diagrams of abelian groups indexed by our index set T. In this context, there is no structure theory of persistence modules as in the case of vector spaces, but the concepts of interleaving distance and weak isomorphisms still make sense. This is all that we need in the present context. We consider persistent homology of sublevel set filtrations of real-valued functions here, so we have $T = \mathbb{R}$.

9.1. Constructing maps from singular to Čech homology and back

We start our discussion of the comparison between singular and Čech homology by constructing some chain maps. In § 9.1.1, we review the standard construction of the map from singular to Čech homology, or, more precisely, to Vietoris homology. We then construct approximate inverses to this map for certain filtrations in § 9.1.2.

9.1.1. The map from singular to Čech homology

We recall the natural map from singular to Čech homology, defined in terms of Vietoris complexes. For a review of the definition of Vietoris complexes and the isomorphism between Čech and Vietoris homology see Section 1.4. To do so, we need a description of singular homology depending on covers of spaces. Denoting by $C_*(X)$ the singular chain complex of a topological space X, we define $C_*(\alpha)$ for some cover α of X to be the subcomplex of $C_*(X)$ generated by those singular simplices whose image is contained in some set $U \in \alpha$. The union of all images of singular simplices making up a singular chain is called the *support* of the chain. With this terminology, $C_*(\alpha)$ is the subcomplex of $C_*(X)$ consisting of all chains whose support is contained in some $U \in \alpha$. If α refines β there is an inclusion of chain complexes $\eta\colon C_*(\alpha) \to C_*(\beta)$. The inclusion $C_*(\alpha) \to C_*(\{X\}) = C_*(X)$ can be shown to induce an isomorphism in homology for all $\alpha \in \mathrm{Cov}(X)$ [ES52, Theorem VI.8.2], with inverse given by a subdivision construction [Mar58]. In particular, we get a description of the singular homology of X, denoted $H_*(X)$ throughout the chapter, as the inverse limit

$$H(X) \cong \lim_{\alpha \in \mathrm{Cov}(X)} H(C_*(\alpha)).$$

We write $V_*(\alpha)$ for the simplicial chain complex of $\mathrm{Vtr}(\alpha)$ and $V_*^d(\alpha)$ for the simplicial chain complex of the d-skeleton $\mathrm{Vtr}^d(\alpha)$ of $\mathrm{Vtr}(\alpha)$. Simplices in $\mathrm{Vtr}(\alpha)$ given by distinct points $v_0, \dots, v_n \in U \in \alpha$ are denoted by $\{v_0, \dots, v_n\}$, and *oriented* simplices in $V_*(\alpha)$ by $[v_0, \dots, v_n]$. Now, let $e_0, \dots, e_n$ be the vertices of the standard n-simplex. For any $\alpha \in \mathrm{Cov}(X)$, we obtain a chain map $\mu\colon C_*(\alpha) \to V_*(\alpha)$ by sending a singular n-simplex σ to the oriented Vietoris simplex $[\sigma(e_0), \dots, \sigma(e_n)]$ if the $\sigma(e_i)$ are all distinct and to 0 else. These chain maps are natural with respect to refinement, so they give rise to a map

$$\varphi = \lim H(\mu)\colon H_*(X) \cong \lim_{\alpha \in \mathrm{Cov}(X)} H(C_*(\alpha)) \to \lim_{\alpha \in \mathrm{Cov}(X)} H(V_*(\alpha)) \cong \check{H}_*(X),$$

which is natural in the argument X. Recall that the last isomorphism is provided by Dowker's theorem as reviewed in § 1.4.3. Thus, we obtain a map

$$\varphi\colon H_*(f_{\leq \bullet}) \to \check{H}_*(f_{\leq \bullet})$$

from the persistent singular to the persistent Čech homology of the sublevel set filtration of a real-valued function f.

9.1.2. A map from Čech to singular homology

Admissible covers and chain maps We will now show how covers satisfying certain local triviality assumptions can be used to construct maps from Čech homology to singular homology in filtrations of spaces.

If α is a cover of a topological space X and $U \in \alpha$, we define the *star* of U with respect to α as

$$\mathrm{St}_\alpha U = \bigcup_{\substack{U' \in \alpha \\ U \cap U' \neq \emptyset}} U' \subseteq X.$$

A cover α is called a *star refinement* of another cover β if for all $U \in \alpha$ there exists $V \in \beta$ with $\mathrm{St}_\alpha U \subseteq V$. We say that the star refinement is HLC with respect to a homology theory $\tilde{H}$ if V can be chosen such that the inclusion $\mathrm{St}_\alpha U \to V$ becomes trivial after applying $\tilde{H}$.

Definition 9.1.1. Let $A_0 \subseteq A_1 \subseteq \cdots \subseteq A_d \subseteq X$ be topological spaces. A sequence of covers α_i of A_i is called *admissible* if α_i is an HLC star refinement of α_{i+1} with respect to reduced singular homology for all i. Given an admissible sequence of subspace covers, a chain map $\lambda\colon V_*^d(\alpha_0) \to C_*(\alpha_d)$ is called *admissible* if

- for any 0-simplex $[v]$, $\lambda([v])$ is a singular 0-simplex taking the value v and if
- for any oriented n-simplex ρ in $V_*^d(\alpha_0)$ there exists $V \in \alpha_n$ such that the support of the singular n-chain $\lambda(\rho)$ and all vertices of ρ are contained in V.

For our later application to the sublevel set filtration of a function $f\colon X \to \mathbb{R}$, one may think of the sets A_i in the definition above as sublevel sets of f.

Lemma 9.1.2. *Let $A_0 \subseteq A_1 \subseteq \cdots \subseteq A_d \subseteq X$ be topological spaces and let $(\alpha_i)_i$ be an admissible sequence of covers. Then there exists an associated admissible chain map $\lambda\colon V_*^d(\alpha_0) \to C_*(\alpha_d)$.*

Proof. We define the desired map λ on oriented Vietoris simplices and extend linearly. We also proceed inductively in the dimension of the simplices.

For a 0-simplex $[v]$, we set $\lambda([v])$ to be the singular simplex with value v. For an oriented 1-simplex $[v_0, v_1] \in V_*^d(\alpha_0)$, we choose $U \in \alpha_0$ with $v_0, v_1 \in U$. The cover α_0 is an HLC star refinement of α_1, so we may choose $V \in \alpha_1$ such that the inclusion $U \subseteq V$ is trivial in singular homology. This means that the augmented singular 0-cycle $c = \lambda(\partial([v_0, v_1])) = \lambda([v_0]) - \lambda([v_1]) \in C_*(U)$ becomes a boundary when considered as a chain in V, so there exists a 1-chain $c' \in C_*(V) \subseteq C_*(\alpha_1)$ such that $\partial(c') = c$. We set $\lambda([v_0, v_1]) = c'$.

Now, inductively, assume that λ has been defined for oriented $(n-1)$-simplices and consider an oriented n-simplex $\rho = [v_0, \ldots, v_n]$, where $2 \leq n \leq d$. Let $\rho_i = [v_0, \ldots, \hat{v}_i, \ldots, v_n]$ be the i-th boundary component of ρ. By our inductive assumption, we have singular $(n-1)$-chains $\lambda(\rho_i) \in C_*(\alpha_d)$, and using admissibility we may choose sets $U_i \in \alpha_{n-1}$ such that the support of $\lambda(\rho_i)$, as well as the points v_j for $j \neq i$ are contained in U_i. Because we assume $n \geq 2$ we can for all i choose $j \notin \{0, i\}$ to obtain $v_j \in U_0 \cap U_i$, so we must have $U_0 \cap U_i \neq \emptyset$ for all i. It follows that $U_i \subseteq \mathrm{St}_{\alpha_{n-1}} U_0$ so that the support of $\lambda(\rho_i)$ is contained in $\mathrm{St}_{\alpha_{n-1}} U_0$ for all i. Since α_{n-1} is an HLC star refinement of α_n we may choose $V \in \alpha_n$ such that $\mathrm{St}_{\alpha_{n-1}} U_0 \subseteq V$ and the corresponding inclusion map is trivial in singular homology. Hence, the $(n-1)$-cycle $c = \lambda(\partial(\rho)) = \sum_i (-1)^i \lambda(\rho_i) \in C_*(\mathrm{St}_{\alpha_{n-1}} U_0)$ becomes a boundary when considered as a chain in V, i.e., there exists a singular n-chain $c' \in C_*(V) \subseteq C_*(\alpha_n)$ such that $\partial(c') = c$. Setting $\lambda(\rho) = c'$ yields a chain map as desired. □

Independence of choices for admissible covers and chain maps Given a filtration of spaces as in Lemma 9.1.2 and Definition 9.1.1, there might be many choices of admissible subspace covers, and for each such choice there might again be many choices of admissible chain maps. We will now show that the result in homology is in some sense independent of these choices. To state the result, recall that η and π are the maps induced by refinement of covers on singular and Vietoris complexes, respectively.

Lemma 9.1.3. *Let $A_0 \subseteq A_1 \subseteq \cdots \subseteq A_d \subseteq X$ and $A'_0 \subseteq A'_1 \subseteq \cdots \subseteq A'_d \subseteq X$ be topological spaces with $A_i \subseteq A'_i$ for all i and assume we are given admissible sequences of subspace covers α_i of A_i and α'_i of A'_i in X such that α_i refines α'_i for all i. Further, assume we are given admissible chain maps λ and λ' associated to the sequences α_i and α'_i, respectively. Then the diagram*

$$\begin{array}{ccc} H_{\leq d-1}(C_*(\alpha_d)) & \xrightarrow{H(\eta)} & H_{\leq d-1}(C_*(\alpha'_d)) \\ {\scriptstyle H(\lambda)}\uparrow & & {\scriptstyle H(\lambda')}\uparrow \\ H_{\leq d-1}(V_*^d(\alpha_0)) & \xrightarrow{H(\pi)} & H_{\leq d-1}(V_*^d(\alpha'_0)) \end{array}$$

commutes.

Proof. We prove the claim by constructing a suitable chain homotopy for the diagram above, i.e., a map $D\colon V_*^{d-1}(\alpha_0) \to C_*(\alpha'_d)$ such that $D \circ \partial + \partial \circ D = \eta \circ \lambda - \lambda' \circ \pi$. Note that it suffices to construct D on the $(d-1)$-skeleton because we only want to show commutativity of the diagram in homology up to dimension $d-1$. In particular, the existence of D on the $(d-1)$-skeleton already implies that $\eta \circ \lambda - \lambda' \circ \pi$ maps $(d-1)$-cycles to $(d-1)$-boundaries.

As for the construction of admissible chain maps in the proof of Lemma 9.1.2, we will perform the construction of D only on oriented simplices and proceed inductively on their dimension. As part of our induction hypothesis, and similarly to the admissibility condition for chain maps from Definition 9.1.1, we will require that for any oriented n-simplex ρ of $V_*^{d-1}(\alpha_0)$ there exists $U \in \alpha'_{n+1}$ such that the support of the singular $(n+1)$-chain $D(\rho)$, as well as the vertices of ρ are contained in U.

For a 0-simplex $[v]$, we set $D([v]) = 0$, which clearly satisfies the required conditions because $\lambda([v])$ and $\lambda'([v])$ are both singular 0-simplices that take value v. Now consider an oriented 1-simplex $[v_0, v_1] \in V_*^{d-1}(\alpha_0)$. Because λ and λ' are admissible, there exist $U \in \alpha_1$ and $U' \in \alpha'_1$ such that U contains the support of $\eta(\lambda([v_0, v_1]))$, U' contains the support of $\lambda'(\pi([v_0, v_1]))$, and $v_0,\ v_1 \in U \cap U'$. Since α_1 refines α'_1, we can choose $U'' \in \alpha'_1$ such that $U \subseteq U''$. In particular, we then have $U'' \cap U' \neq \emptyset$, so $U'' \cup U' \subseteq \operatorname{St}_{\alpha'_1} U'$. By assumption, α'_1 is an HLC star refinement of α'_2, so we may choose $V \in \alpha'_2$ such that $\operatorname{St}_{\alpha'_1} U' \subseteq V$ and the inclusion map is trivial in singular homology. Since $c = \eta(\lambda([v_0, v_1])) - \lambda'(\pi([v_0, v_1]))$ is a singular 1-cycle supported in $U'' \cup U' \subseteq \operatorname{St}_{\alpha'_1} U'$, this implies that there is a singular 2-chain $c' \in C_*(V) \subseteq C_*(\alpha'_2)$ such that $\partial(c') = c$. Setting $D([v_0, v_1]) = c'$ then satisfies the above requirements. In particular, we can calculate

$$\partial(D([v_0, v_1])) + D(\partial([v_0, v_1])) = \partial(c') + 0 = c = \eta(\lambda([v_0, v_1])) - \lambda'(\pi([v_0, v_1])).$$

Next, assume that D has been defined for oriented $(n-1)$-simplices and consider an oriented n-simplex $\rho = [v_0, \dots, v_n] \in V_*^{d-1}(\alpha_0)$, where $2 \leq n$, with boundary components ρ_i. As before, we may use the admissibility of λ and λ' to choose $U \in \alpha_n$ and $U' \in \alpha'_n$ such that U contains the support of $\eta(\lambda(\rho))$, U' contains the support of $\lambda'(\pi(\rho))$, and both U and U' contain the vertices of ρ. Again, we choose some $U'' \in \alpha'_n$ such that $U \subseteq U''$. Using the induction hypothesis, we may also choose $U_i \in \alpha'_n$ such that the support of $D(\rho_i)$ and all of the vertices of ρ_i are contained in U_i for all i. Because we assume $n \geq 2$ we can for all i choose $j \notin \{0, i\}$ to obtain $v_j \in U_0 \cap U_i$, so we must have $U_0 \cap U_i \neq \emptyset$ for all i. It follows that $U_i \subseteq \operatorname{St}_{\alpha'_n} U_0$. We also have that $v_n \in U'' \cap U_0$ and $v_n \in U' \cap U_0$, so we get $U'', U' \subseteq \operatorname{St}_{\alpha'_n} U_0$, too. In total, we obtain that the singular n-cycle $c = \eta(\lambda(\rho)) - \lambda'(\pi(\rho)) - D(\partial(\rho))$ has support in $U'' \cup U' \cup \bigcup_i U_i \operatorname{St}_{\alpha'_n} U_0$. Since α'_n is an HLC star refinement of α'_{n+1} we may now choose $V \in \alpha'_{n+1}$ such that $\operatorname{St}_{\alpha'_n} U_0 \subseteq V$ and the corresponding inclusion map is trivial in singular homology. Thus, there is some $(n+1)$-chain $c' \in C_*(V) \subseteq C_*(\alpha'_{n+1})$ with $\partial(c') = c$. Setting $D(\rho) = c'$, we calculate

$$\partial(D(\rho)) + D(\partial(\rho)) = \partial(c') + D(\partial(\rho)) = c + \eta(\lambda(\rho)) - \lambda'(\pi(\rho)) - c = \eta(\lambda(\rho)) - \lambda'(\pi(\rho)),$$

which finishes the proof. □

Note that Lemma 9.1.3 can in particular be applied in the case where $A_i = A'_i$ and $\alpha_i = \alpha'_i$, which implies that the maps η and π are identities and hence two choices of associated chain maps λ and λ' for the same sequence of covers yield identical maps in homology up to dimension $d-1$.

Admissible maps in locally connected sublevel set filtrations Coming back to our setting of functions $f\colon X \to \mathbb{R}$ and their sublevel set filtrations, we define the following.

Definition 9.1.4. Let $f\colon X \to \mathbb{R}$ be a function on a topological space. For $\delta \geq 0$, we say that the sublevel set filtration $f_{\leq \bullet}$ induced by f is δ *homologically locally connected* ($\delta -$ HLC*)* with respect to the homology theory H if for any $x \in X$, any neighborhood V of x, and any pair of indices s, t with $f(x) < s \leq s + \delta < t$ there is a neighborhood U of x with $U \subseteq V$ such that the inclusion $f_{\leq s} \cap U \to f_{\leq t} \cap V$ is taken to the trivial map by H. We define the *local connectedness shift of f with respect to H* as

$$\mathrm{lcs}_H(f) = \inf\{\delta > 0 \mid f \text{ is } \delta - \mathrm{HLC}\}.$$

We now consider a dimension d and indices $s < t$ satisfying $t - s > d \cdot \mathrm{lcs}_{\mathrm{singular}}(f)$ to define

$$\begin{aligned}\mathcal{I}_{s,t,d} = \{(t_0, \dots, t_d, \alpha_0, \dots, \alpha_d) \mid & t_i \in \mathbb{R}, t_0 = s, t_d = t,\\ & t_{i+1} - t_i > \mathrm{lcs}_{\mathrm{singular}}(f) \text{ for all } i,\\ & (\alpha_i)_i \text{ admissible covers for } f_{\leq t_0} \subseteq \cdots \subseteq f_{\leq t_d} \subseteq X\}.\end{aligned}$$

We define an order $\leq$ on $\mathcal{I}_{s,t,d}$ by saying that $(t_0, \dots, \alpha_d) \leq (t'_0, \dots, \alpha'_d)$ if and only if $t_i \leq t'_i$ – so that in particular $f_{\leq t_i} \subseteq f_{\leq t'_i}$ – and α_i refines α'_i for all i.

Lemma 9.1.5. *Let $f\colon X \to \mathbb{R}$ be a function whose sublevel sets are paracompact Hausdorff spaces, and consider a dimension d and indices $s < t$ satisfying $t - s > d \cdot \mathrm{lcs}_{\mathrm{singular}}(f)$. Then $(\mathcal{I}_{s,t,d}, \leq)$ is a non-empty directed set.*

Proof. It is clear that the relation $\leq$ defines a preorder because the refinement order for covers of a single space defines a preorder.

For non-emptyness, we proceed inductively. Start with an arbitrary choice of t_i with $t_0 = s$, $t_d = t$, $t_{i+1} - t_i > \mathrm{lcs}_{\mathrm{singular}}(f)$, and an arbitrary choice of $\alpha_d \in \mathrm{Cov}(f_{\leq t})$. Next, given α_{i+1}, choose α'_i as a cover of $f_{\leq t_i}$ such that for every $U' \in \alpha'_i$ there exists $V \in \alpha_{i+1}$ with $U' \subseteq V$ with the inclusion $f_{\leq t_i} \cap U' \to f_{\leq t_{i+1}} \cap V$ being trivial for $\tilde{H}$. This choice is possible since $t_{i+1} - t_i > \mathrm{lcs}_{\mathrm{singular}}(f)$. Next, choose a star refinement α_i of α'_i, which is possible because $f_{\leq t_i}$ is assumed to be a paracompact Hausdorff space [Eng89, Theorem 5.1.12]. Clearly, α_i is then an HLC star refinement of α_{i+1}.

To show directedness, we consider elements $(t_0, \dots, \alpha_d)$ and $(t'_0, \dots, \alpha'_d)$ in $\mathcal{I}_{s,t,d}$ and construct a common lower bound $(t''_0, \dots, \alpha''_d)$ for them as follows. First, we set $t''_i = \min\{t_i, t'_i\}$, which satisfies $t''_{i+1} - t''_i > \mathrm{lcs}_{\mathrm{singular}}(f)$ so that f is $(t_{i+1} - t_i) -$ HLC for all i. Next, we choose $\alpha''_d \in \mathrm{Cov}(f_{\leq t''_d})$ as an arbitrary common refinement of α_d and α'_d. Finally, we inductively define $\alpha''_i \in \mathrm{Cov}(f_{\leq t''_i})$ by constructing an HLC star refinement β_i of α''_{i+1} as above, and then choosing α''_i as a common refinement of β_i, α_i, and α'_i. ☐

We can now define an inverse system of maps indexed by the non-empty directed set $\mathcal{I}_{s,t,d}$ by mapping $(t_0, \dots, t_d, \alpha_0, \dots, \alpha_d)$ to

$$H_{\leq d-1}(V^d_*(\alpha_0)) \xrightarrow{H(\lambda)} H_{\leq d-1}(C_*(\alpha_d)),$$

where λ is some choice of admissible chain map associated to the covers α_i. The connecting maps corresponding to the relations on $\mathcal{I}_{s,t,d}$ are given by pairs $(H(\pi), H(\eta))$, which is well-defined by Lemma 9.1.3.

Next, we want to define a map $\psi_{s,t}\colon \check{H}_{\leq d-1}(f_{\leq s}) \to H_{\leq d-1}(f_{\leq t})$ playing the role of an approximate inverse to φ as the inverse limit of the above maps $H(\lambda)$ over the set $\mathcal{I}_{s,t,d}$. To do so, what remains to be shown is that the limit of the domains of the $H(\lambda)$ is $\check{H}_{\leq d-1}(f_{\leq s})$ and that the limit of the codomains is $H_{\leq d-1}(f_{\leq t})$. We define

$$\begin{aligned}\mathcal{C}_{s,t,d} &= \{\alpha \in \operatorname{Cov}(f_{\leq s}) \mid \text{ there exists } (t_0,\dots,t_d,\alpha_0,\dots,\alpha_d) \in \mathcal{I}_{s,t,d} \text{ with } \alpha_0 = \alpha\},\\ \mathcal{C}'_{s,t,d} &= \{\alpha \in \operatorname{Cov}(f_{\leq t}) \mid \text{ there exists } (t_0,\dots,t_d,\alpha_0,\dots,\alpha_d) \in \mathcal{I}_{s,t,d} \text{ with } \alpha_d = \alpha\}.\end{aligned}$$

Both sets will be considered as preordered sets with the refinement relation.

Lemma 9.1.6. *Let $f\colon X \to \mathbb{R}$ be a function whose sublevel sets are paracompact Hausdorff spaces, and consider a dimension d and indices $s < t$ satisfying $t - s > d \cdot \operatorname{lcs}_{\text{singular}}(f)$. Then the subsets $\mathcal{C}_{s,t,d} \subseteq \operatorname{Cov}(f_{\leq s})$ and $\mathcal{C}'_{s,t,d} \subseteq \operatorname{Cov}(f_{\leq t})$ are coinitial.*

Proof. First, note that the part of the proof of Lemma 9.1.5 where non-emptyness of $\mathcal{I}_{s,t,d}$ is shown actually establishes that α_d can be chosen arbitrarily, so that $\mathcal{C}'_{s,t,d} = \operatorname{Cov}(f_{\leq t})$. For the other assertion, let $\alpha \in \operatorname{Cov}(f_{\leq s})$ be an arbitrary cover. We choose some element $(t_0,\dots,t_d,\alpha_0,\dots,\alpha_d) \in \mathcal{I}_{s,t,d}$, which is possible because the set is non-empty. If α'_0 is a common refinement of α and α_0, then α'_0 is clearly still an HLC star refinement of α_1, so that $(t_0,\dots,t_d,\alpha'_0,\alpha_1\dots,\alpha_d)$ is an element of $\mathcal{I}_{s,t,d}$. Hence, we have $\alpha'_0 \in \mathcal{C}_{s,t,d}$ and α'_0 refines α, so $\mathcal{C}_{s,t,d} \subseteq \operatorname{Cov}(f_{\leq s})$ is indeed coinitial. □

As a consequence of Lemma 9.1.6, we get that the domain of $\lim_{\mathcal{I}_{s,t,d}} H(\lambda)$ is

$$\begin{aligned}\lim_{(t_0,\dots,t_d,\alpha_0,\dots,\alpha_d)\in\mathcal{I}_{s,t,d}} H_{\leq d-1}(V^d_*(\alpha_0)) &\cong \lim_{\alpha\in\mathcal{C}_{s,t,d}} H_{\leq d-1}(V^d_*(\alpha))\\ &\cong \lim_{\alpha\in\operatorname{Cov}(f_{\leq s})} H_{\leq d-1}(V_*(\alpha))\\ &\cong \check{H}_{\leq d-1}(f_{\leq s}),\end{aligned}$$

where we have made use of the fact that for any simplicial complex K, its d-skeleton determines its homology up to degree $d-1$, i.e., $H_{\leq d-1}(K) \cong H_{\leq d-1}(K^d)$. Similarly, we obtain that the codomain of $\lim_{\mathcal{I}_{s,t,d}} H(\lambda)$ is

$$\begin{aligned}\lim_{(t_0,\dots,t_d,\alpha_0,\dots,\alpha_d)\in\mathcal{I}_{s,t,d}} H_{\leq d-1}(C_*(\alpha_d)) &\cong \lim_{\alpha\in\mathcal{C}'_{s,t,d}} H_{\leq d-1}(C_*(\alpha))\\ &\cong \lim_{\alpha\in\operatorname{Cov}(f_{\leq t})} H_{\leq d-1}(C_*(\alpha))\\ &\cong H_{\leq d-1}(f_{\leq t}).\end{aligned}$$

Hence, we indeed get a well-defined map

$$\psi_{s,t} = \lim_{\mathcal{I}_{s,t,d}} H(\lambda)\colon \check{H}_{\leq d-1}(f_{\leq s}) \to H_{\leq d-1}(f_{\leq t})$$

whenever $t - s > d \cdot \operatorname{lcs}_{\text{singular}}(f)$. From this, we obtain a diagram

$$\begin{array}{ccc} H_{\leq d-1}(f_{\leq s};G) & \longrightarrow & H_{\leq d-1}(f_{\leq t};G)\\ \downarrow{\scriptstyle\varphi_s} & \overset{\psi_{s,t}}{\nearrow} & \downarrow{\scriptstyle\varphi_t}\\ \check{H}_{\leq d-1}(f_{\leq s};G) & \longrightarrow & \check{H}_{\leq d-1}(f_{\leq t};G)\end{array} \tag{9.1.1}$$

whose commutativity we want to show next. We will also show that the maps $\psi_{s,t}$ are compatible across different choices of s and t. Together with the commutativity of Diagram (9.1.1), this will imply our main comparison results.

9.2. Constructing homotopies on singular and Čech complexes

The goal of this section is to show commutativity of Diagram (9.1.1). For this purpose, we construct chain homotopies on singular complexes in § 9.2.1, and chain homotopies on Čech complexes in § 9.2.2.

9.2.1. A homotopy on singular chain complexes

We now prove that the triangle in Diagram (9.1.1) involving the inclusion-induced map in singular homology commutes.

Proposition 9.2.1. *Let $f\colon X \to \mathbb{R}$ be a function whose sublevel sets are paracompact Hausdorff spaces and consider a dimension d, as well as $\delta > d \cdot \mathrm{lcs}_{singular}(f)$. Then for all $s \in \mathbb{R}$ the diagram*

$$\begin{array}{ccc} H_{\leq d-1}(f_{\leq s}) & \longrightarrow & H_{\leq d-1}(f_{\leq s+\delta}) \\ {\scriptstyle \varphi_s}\big\downarrow & \nearrow_{\psi_{s,s+\delta}} & \\ \check{H}_{\leq d-1}(f_{\leq s}) & & \end{array} \tag{9.2.1}$$

commutes.

We start with a lemma. It is analogous to [Mar59, Lemma 8].

Lemma 9.2.2. *Let $A_0 \subseteq A_1 \subseteq \cdots \subseteq A_d \subseteq X$ be topological spaces with an admissible sequence of covers α_i and an admissible chain map λ. Then the diagram*

$$\begin{array}{ccc} H_{\leq d-1}(C_*(\alpha_0)) & \xrightarrow{H(\eta)} & H_{\leq d-1}(C_*(\alpha_d)) \\ {\scriptstyle H(\mu)}\big\downarrow & \nearrow_{H(\lambda)} & \\ H_{\leq d-1}(V^d_*(\alpha_0)) & & \end{array} \tag{9.2.2}$$

commutes.

Proof. We follow the general outline of the proof of Lemma 9.1.3 and construct a chain homotopy $D\colon C_*(\alpha_0) \to C_*(\alpha_d)$ for the above diagram, i.e., a map such that $D \circ \partial + \partial \circ D = \lambda \circ \mu - \eta$. Again, it suffices to define D for simplices up to dimension $d-1$, and to do so we proceed inductively on the dimension. As part of the induction hypothesis, we will again assume that for any n-simplex σ of $C_*(\alpha_0)$ with $n \leq d-1$ there exists $U \in \alpha'_{n+1}$ such that the singular $(n+1)$-chain $D(\sigma)$ and the singular simplex σ are supported in U.

On 0-simplices, we set D to be 0. For a 1-simplex σ, we distinguish two cases: either $\mu(\sigma) \neq 0$ or $\mu(\sigma) = 0$. If $\mu(\sigma)$ is not 0, it is an oriented 1-simplex in $V_*(\alpha_0)$. Because λ is admissible, we may hence choose $U \in \alpha_1$ such that the support of $\lambda(\mu(\sigma))$ and the vertices of $\mu(\sigma)$ are contained in U. We have $\sigma \in C_*(\alpha_0)$ and α_0 refines α_1, so we may also choose $U' \in \alpha_1$ such that U' contains the support of σ. Note that the support of σ contains the vertices of $\mu(\sigma)$, so we have $U \cap U' \neq \emptyset$. Hence, the 1-cycle $\lambda(\mu(\sigma)) - \eta(\sigma)$ is supported in $\mathrm{St}_{\alpha_1} U$. Since the covers α_i are admissible, we may now choose $V \in \alpha_2$ such that $\mathrm{St}_{\alpha_1} U \subseteq V$ and the inclusion map is trivial in singular homology. It follows that there exists a 2-chain $c' \in C_*(V) \subseteq C_*(\alpha_2)$ with $\partial(c') = c$. We set $D(\sigma) = c'$.

If $\mu(\sigma) = 0$, we choose $U \in \alpha_0$ such that the image of σ is contained in U. Since α_0 is an HLC star refinement of α_2, we may pick $V \in \alpha_2$ such that $U \subseteq V$ and the inclusion map is trivial in singular homology. Note that $\mu(\sigma) = 0$ implies that the two vertices of σ are mapped to the same point in U. In particular, this implies that $c = \sigma$ is a 1-cycle in $C_*(U)$, so there exists a 2-chain $c' \in C_*(V) \subseteq C_*(\alpha_2)$ with $\partial(c') = c$. We set $D(\sigma) = c'$.

Now, assume that D has been defined as required for simplices up to dimension $n-1$ and let σ be an n-simplex in $C_*(\alpha_0)$, where $2 \leq n \leq d-1$. Again, we distinguish between the cases where either $\mu(\sigma) \neq 0$ or $\mu(\sigma) = 0$. If $\mu(\sigma)$ is not 0, it is an oriented n-simplex of $V_*(\alpha_0)$. Using admissibility of λ, we may consequently choose $U \in \alpha_n$ such that the support of $\lambda(\mu(\sigma))$ and the vertices of $\mu(\sigma)$ are contained in U. We may also choose $U' \in \alpha_n$ such that the support of σ is contained in U'. With the induction hypothesis on D, it is also possible to choose $U_i \in \alpha_n$ for all boundary simplices σ_i of σ such that U_i contains the supports of $D(\sigma_i)$ and σ_i. Writing $v_0, \dots, v_n$ for the vertices of $\mu(\sigma)$, we obtain $v_n \in U \cap U' \cap U_0$ and $v_j \in U_i \cap U_0$ for $j \notin \{0, i\}$. Hence, the singular n-cycle $c = \lambda(\mu(\sigma)) - \eta(\sigma) - D(\partial(\sigma))$ is supported in $\mathrm{St}_{\alpha_n}(U_0)$. Using admissibility of the α_i, we can now choose $V \in \alpha_{n+1}$ such that $\mathrm{St}_{\alpha_n}(U_0) \subseteq V$ and the inclusion map is homologically trivial. Thus, there exists $c' \in C_*(V) \subseteq C_*(\alpha_{n+1})$ such that $\partial(c') = c$ and we set $D(\sigma) = c'$.

If $\mu(\sigma) = 0$, we start by choosing $U \in \alpha_n$ such that contains the support of σ, which is possible because $\sigma \in C_*(\alpha_0)$ and α_0 refines α_n. For every boundary simplex σ_i of σ we use the induction hypothesis to choose $U_i \in \alpha_n$ such that U_i contains the supports of σ_i and $D(\sigma_i)$. A routine argument again implies that $U \cap U_0 \neq \emptyset$ and $U_i \cap U_0 \neq \emptyset$ for all i, so the n-cycle $c = \lambda(\mu(\sigma)) - \eta(\sigma) - D(\partial(\sigma)) = -\eta(\sigma) - D(\partial(\sigma))$ is supported in $\mathrm{St}_{\alpha_n}(U_0)$. Since the α_i are admissible we may choose $V \in \alpha_{n+1}$ such that $\mathrm{St}_{\alpha_n}(U_0)$ includes homologically trivially into V. This implies that there is an $(n+1)$-chain $c' \in C_*(V) \subseteq C_*(\alpha_{n+1}$ such that $\partial(c') = c$. Setting $D(\sigma) = c'$ finishes the construction. We omit the straightforward verification that D indeed has all required properties. □

Proof of Proposition 9.2.1. Using Lemma 9.1.3, we may consider an inverse system indexed by $\mathcal{I}_{s,s+\delta,d}$, mapping $(t_0, \dots, \alpha_d)$ to the Diagram (9.2.2), which commutes by Lemma 9.2.2. The connecting maps in this inverse system corresponding to relations in $\mathcal{I}_{s,s+\delta,d}$ are given by $(H(\eta), H(\eta), H(\pi))$. Taking the inverse limit of this system yields the commutative Diagram (9.2.1) as claimed. □

9.2.2. A homotopy on Čech complexes

We now prove that the triangle in Diagram (9.1.1) involving the inclusion-induced map in Čech homology commutes.

Proposition 9.2.3. *Let $f\colon X \to \mathbb{R}$ be a function whose sublevel sets are paracompact Hausdorff spaces and consider a dimension d, as well as $\delta > d \cdot \mathrm{lcs}_{singular}(f)$. Then for all $s \in \mathbb{R}$ the diagram*

$$\begin{array}{ccc} & & H_{\leq d-1}(f_{\leq s+\delta}) \\ & \overset{\psi_{s,s+\delta}}{\nearrow} & \downarrow \varphi_{s+\delta} \\ \check{H}_{\leq d-1}(f_{\leq s}) & \longrightarrow & \check{H}_{\leq d-1}(f_{\leq s+\delta}) \end{array} \tag{9.2.3}$$

commutes.

We start with some terminology and intermediate results.

Definition 9.2.4. Let X be a topological space and $\alpha \in \mathrm{Cov}(X)$. If $\rho = [v_0, \dots, v_n] \in V_*(\alpha)$ is an oriented Vietoris n-simplex with $U \in \alpha$ such that $v_i \in U$ for all i, and $x \in U$ is any point, we define the *join* $x \vee \rho$ as the oriented Vietoris $(n+1)$-simplex $[x, v_0, \dots, v_n]$. By convention, we understand this to mean that $x \vee \rho = 0$ if there is some j with $x = v_j$. If $\sum_i a_i \rho_i \in V_*(\alpha)$ is a Vietoris n-chain with simplices ρ_i, coefficients a_i, sets $U_i \in \alpha$ such that the vertices of ρ_i are in U_i, and $x \in \bigcap_i U_i$ is any point, we define the *join* $x \vee \sum_i a_i \rho_i$ as the $(n+1)$-chain $\sum_i a_i (x \vee \rho_i)$.

Lemma 9.2.5. *Let X be a topological space and $\alpha \in \mathrm{Cov}(X)$. Let $c = \sum_i a_i \rho_i \in V_*(\alpha)$ be an n-chain with $n \geq 1$ and choose sets $U_i \in \alpha$ such that the vertices of ρ_i are contained in U_i. Let $x \in \bigcap_i U_i$ be any point. Then $\partial(x \vee c) + x \vee \partial(c) = c$.*

Proof. Writing $\rho_i = [v_{i,0}, \dots, v_{i,n}]$ for the simplices of c, we can write the boundary of ρ_i as $\partial(\rho_i) = \sum_{k=0}^n (-1)^k [v_{i,0}, \dots, \widehat{v_{i,k}}, \dots, v_{i,n}]$, where $\widehat{v_{i,k}}$ means that $v_{i,k}$ is excluded from the given list. Note that this requires our assumption that $n \geq 1$. Denoting $c' = \partial(x \vee c) + x \vee \partial(c)$, we calculate

$$\begin{aligned}
c' &= \sum_i a_i \left(\partial(x \vee \rho_i) + (x \vee \partial(\rho_i))\right) \\
&= \sum_i a_i \left(\partial([x, v_{i,0}, \dots, v_{i,n}]) + \left(x \vee \sum_{k=0}^n (-1)^k [v_{i,0}, \dots, \widehat{v_{i,k}}, \dots, v_{i,n}]\right)\right) \\
&= \sum_i a_i \left([v_{i,0}, \dots, v_{i,n}] + \sum_{k=0}^n \left(\left((-1)^{k+1} + (-1)^k\right) [x, v_{i,0}, \dots, \widehat{v_{i,k}}, \dots, v_{i,n}]\right)\right) \\
&= \sum_i a_i [v_{i,0}, \dots, v_{i,n}] \\
&= c.
\end{aligned}$$

□

The following lemma is analogous to [Mar59, Lemma 9].

Lemma 9.2.6. *Let $A_0 \subseteq A_1 \subseteq \cdots \subseteq A_d \subseteq X$ be topological spaces with an admissible sequence of covers α_i and an admissible chain map λ. Then the diagram*

$$\begin{array}{ccc}
 & & H_{\leq d-1}(C_*(\alpha_d)) \\
 & \overset{H(\lambda)}{\nearrow} & \downarrow H(\mu) \\
H_{\leq d-1}(V_*^d(\alpha_0)) & \xrightarrow[H(\pi)]{} & H_{\leq d-1}(V_*^d(\alpha_d))
\end{array} \tag{9.2.4}$$

commutes.

Proof. We prove the claim by constructing a suitable chain homotopy for the diagram above, i.e., a map $D \colon V_*^d(\alpha_0) \to V_*^d(\alpha_d)$ such that $D \circ \partial + \partial \circ D = \mu \circ \lambda - \pi$. It suffices to define D on simplices and extend it linearly. For 0-simplices $\rho = [v]$, we set $D(\rho) = 0$, which satisfies the above equation because we have $\mu(\lambda(\rho)) = [v] = \pi(\rho)$ since $\lambda(\rho)$ takes value v. If ρ is an oriented n-simplex with $n \geq 1$, we choose $U \in \alpha_n$ such that the support of $\lambda(\rho)$ and the vertices of ρ are contained in U, which is possible since λ is admissible. Next, we choose an arbitrary point $x \in U$ and define $D(\rho) = x \vee (\mu(\lambda(\rho)) - \pi(\rho))$. We have

$n \geq 1$, so it follows from the Lemma 9.2.5 that

$$\begin{aligned} D(\partial(\rho)) + \partial(D(\rho)) &= x \vee (\mu(\lambda(\partial(\rho))) - \pi(\partial(\rho))) + \partial\,(x \vee (\mu(\lambda(\rho) - \pi(\rho))) \\ &= \partial(x \vee (\mu(\lambda(\rho)) - \pi(\rho))) + \partial\,(x \vee (\mu(\lambda(\rho) - \pi(\rho))) \\ &= \mu(\lambda(\rho)) - \pi(\rho). \end{aligned}$$ □

Proof of Proposition 9.2.3. Using Lemma 9.1.3, we may consider an inverse system indexed by $\mathcal{I}_{s,s+\delta,d}$, mapping $(t_0, \ldots, \alpha_d)$ to the Diagram (9.2.4), which commutes by Lemma 9.2.6. The connecting maps in this inverse system corresponding to relations in $\mathcal{I}_{s,s+\delta,d}$ are given by $(H(\eta), H(\pi), H(\pi))$. Taking the inverse limit of this system yields the commutative Diagram (9.2.3) as claimed. □

9.3. Final proofs

We are now ready to give proofs of the main theorems stated in § 2.2.3, which we do in § 9.3.1. To finish the chapter, we show that the local connectedness shift is 2-Lipschitz for the supremum norm in § 9.3.2.

9.3.1. Proofs of the singular to Čech comparison results

As a last preparatory result, we show that the construction of the maps $\psi_{s,t}$ is in some sense consistent among different choices for the indices s and t.

Proposition 9.3.1. *Let $f\colon X \to \mathbb{R}$ be a function whose sublevel sets are paracompact Hausdorff spaces, and consider a dimension d, an index $s \in \mathbb{R}$, and $\delta > d \cdot \mathrm{lcs}_{\mathrm{singular}}(f)$. Then the diagram*

$$\begin{array}{ccccc} & & H_{\leq d-1}(f_{\leq s+\delta}) & \longrightarrow & H_{\leq d-1}(f_{\leq s+2\delta}) \\ & \overset{\psi_{s,s+\delta}}{\nearrow} & & \overset{\psi_{s+\delta,s+2\delta}}{\nearrow} & \\ \check{H}_{\leq d-1}(f_{\leq s}) & \longrightarrow & \check{H}_{\leq d-1}(f_{\leq s+\delta}) & & \end{array} \tag{9.3.1}$$

commutes.

Proof. We write $s' = s+\delta$ and $s'' = s+2\delta$. Now, define a map $\tau\colon \mathcal{I}_{s',s'',d} \to \mathcal{I}_{s,s',d}$ as follows. Starting with an element $(t'_0, \ldots, t'_d, \alpha'_0, \ldots, \alpha'_d) \in \mathcal{I}_{s',s'',d}$, we set $t_i = t'_i - \delta$. We also choose $\alpha_d \in \mathrm{Cov}(f_{\leq t_d})$ as an arbitrary refinement of α'_d. Then, similar to the construction of common refinements in the proof of Lemma 9.1.5, we inductively choose $\beta_i \in \mathrm{Cov}(f_{\leq t_i})$ as an HLC star refinement of $\alpha_{i+1} \in \mathrm{Cov}(f_{\leq t_{i+1}})$, and then choose α_i as a common refinement of β_i and α'_i. This process is possible because $t_{i+1} - t_i = t'_{i+1} - t'_i > \mathrm{lcs}_{\mathrm{singular}}(f)$. By construction, $(t_0, \ldots, \alpha_d)$ is an element of $\mathcal{I}_{s,s',d}$, so we may set $\tau(t'_0, \ldots, t'_d, \alpha'_0, \ldots, \alpha'_d) = (\tau(t_0), \ldots, \tau(t_d), \tau(\alpha_0), \ldots, \tau(\alpha_d)) = (t_0, \ldots, t_d, \alpha_0, \ldots, \alpha_d)$.

Since we have $f_{\tau(t'_i)} \subseteq f_{t'_i}$ and $\tau(\alpha'_i)$ refines α'_i for all i, we can apply Lemma 9.1.3 to obtain a commutative diagram

$$\begin{array}{ccc} H_{\leq d-1}(C_*(\tau(\alpha'_d))) & \xrightarrow{H(\eta)} & H_{\leq d-1}(C_*(\alpha'_d)) \\ {\scriptstyle H(\lambda)}\big\uparrow & & \big\uparrow{\scriptstyle H(\lambda')} \\ H_{\leq d-1}(V^d_*(\tau(\alpha'_0))) & \xrightarrow{H(\pi)} & H_{\leq d-1}(V^d_*(\alpha'_0)) \end{array}$$

for any element $(t'_0, \ldots, t'_d, \alpha'_0, \ldots, \alpha'_d) \in \mathcal{I}_{s',s'',d}$, where λ and λ' are choices of admissible chain maps. In other words, we obtain a morphism between the inverse systems of admissible maps $H(\lambda)$ and $H(\lambda')$, which by a standard procedure gives rise to a map between their limits. That is, we obtain the claimed commutative Diagram (9.3.1). That the limits of the maps $H(\pi)$ and $H(\eta)$ yield the inclusion-induced maps is an immediate consequence of their definitions. ☐

We summarize the results from Propositions 9.2.1, 9.2.3 and 9.3.1 in the following corollary.

Corollary 9.3.2. *Let $f\colon X \to \mathbb{R}$ be a function whose sublevel sets are paracompact Hausdorff spaces and consider a dimension d, as well as $\delta > d \cdot \mathrm{lcs}_{\mathrm{singular}}(f)$. Then for all $s \in \mathbb{R}$ there are maps $\psi_{s,s+\delta}\colon \check{H}_{d-1}(f_{\leq s}) \to H_*(f_{\leq s+\delta})$ such that the diagram*

$$\begin{array}{ccccc}
H_{\leq d-1}(f_{\leq s}) & \longrightarrow & H_{\leq d-1}(f_{\leq s+\delta}) & \longrightarrow & H_{\leq d-1}(f_{\leq s+2\delta}) \\
\Big\downarrow \varphi_s & \nearrow \psi_{s,s+\delta} & \Big\downarrow \varphi_{s+\delta} & \nearrow \psi_{s+\delta,s+2\delta} & \Big\downarrow \varphi_{s+2\delta} \\
\check{H}_{\leq d-1}(f_{\leq s}) & \longrightarrow & \check{H}_{\leq d-1}(f_{\leq s+\delta}) & \longrightarrow & \check{H}_{\leq d-1}(f_{\leq s+2\delta})
\end{array} \tag{9.3.2}$$

commutes.

To state the first of our two main theorems, we emphasize the notion of being 0-HLC with some shorthand terminology.

Definition 9.3.3. Let $f\colon X \to \mathbb{R}$ be a function on a topological space. We say that the sublevel set filtration $f_{\leq \bullet}$ induced by f is *homologically locally connected (*HLC*)* with respect to the homology theory H if for any $x \in X$, any neighborhood V of x, and any pair of indices s, t with $f(x) < s < t$ there is a neighborhood U of x with $U \subseteq V$ such that the inclusion $f_{\leq s} \cap U \to f_{\leq t} \cap V$ is taken to the trivial map by H.

Clearly, being HLC is equivalent to being 0-HLC in the sense of Definition 9.1.4. Moreover, any filtration that is HLC is also LHS in the sense of Definition 8.1.2. As our main result for HLC filtrations, we have the following.

Theorem 9.3.4. *If $f\colon X \to \mathbb{R}$ induces a filtration of paracompact Hausdorff spaces that is* HLC *with respect to singular homology with coefficients in an abelian group G, then the natural map $\varphi\colon H_*(f_{\leq\bullet}; G) \to \check{H}_*(f_{\leq\bullet}; G)$ from its persistent singular to its persistent Čech homology is a weak isomorphism.*

Proof. We fix a dimension d and indices $s < t$ and have to show that for $s < t$ the natural maps $\ker \varphi_s \to \ker \varphi_t$ and $\operatorname{coker} \varphi_s \to \operatorname{coker} \varphi_t$ are 0. Note that $\mathrm{lcs}_{\mathrm{singular}}(f) = 0$ because we assume the filtration of f to be HLC. Thus, for $\delta = t - s$, we have $\delta > 0 = (d+1) \cdot \mathrm{lcs}_{\mathrm{singular}}(f)$. So by Corollary 9.3.2 there exists a map $\psi_{s,t}\colon \check{H}_{\leq d}(f_{\leq s}) \to H_{\leq d}(f_{\leq t})$ such that Diagram (9.3.2) commutes. Let $i_{s,t}$ denote the inclusion $f_{\leq s} \to f_{\leq t}$. Given the above diagram, we see that for any $h \in \ker \varphi_s$, we have

$$H_{\leq d}(i_{s,t})(h) = \psi_{s,t}(\varphi_s(h)) = 0,$$

so that the natural map $\ker \varphi_s \to \ker \varphi_t$ must be 0. Similarly, we obtain that for any $h \in \check{H}_{\leq d}(f_{\leq s})$, we have

$$\check{H}_{\leq d}(i_{s,t})(h) = \varphi_t(\psi_{s,t}(h)) \in \operatorname{im} \varphi_t,$$

so that the natural map $\operatorname{coker} \varphi_s \to \operatorname{coker} \varphi_t$ must be 0. ☐

Note that the compatibility part of Corollary 9.3.2 is not necessary for the proof above, i.e., for the proof above it suffices to have established Propositions 9.2.1 and 9.2.3, and Proposition 9.3.1 is not needed. It is needed, however, for the following quantitative result.

Theorem 9.3.5. *If $f\colon X \to \mathbb{R}$ induces a filtration of paracompact Hausdorff spaces and G is an abelian group, then*

$$d_I\big(H_{\leq d-1}(f_{\leq \bullet}; G), \check{H}_{\leq d-1}(f_{\leq \bullet}; G)\big) \leq d \cdot \operatorname{lcs}(f)$$

for all d, where $\operatorname{lcs}(f)$ is the local connectedness shift of f with respect to singular homology with coefficients in G.

Proof. We fix a dimension d. To show that the claimed inequality holds, it suffices to show that $H_{\leq d-1}(f_{\leq \bullet})$ and $\check{H}_{\leq d-1}(f_{\leq \bullet})$ are δ-interleaved for any $\delta > d \cdot \operatorname{lcs}_{\mathrm{singular}}(f)$. For such δ, we get maps $\psi_{s,s+\delta}\colon \check{H}_{\leq d}(f_{\leq s}) \to H_{\leq d}(f_{\leq s+\delta})$ for every $s \in \mathbb{R}$ from Corollary 9.3.2 such that Diagram (9.3.2) commutes. It follows that the maps $\psi_{s,s+\delta}$ form a δ-interleaving together with the maps $\tilde{\varphi}_{s,s+\delta}\colon H_{\leq d}(f_{\leq s}) \to \check{H}_{\leq d}(f_{\leq s+\delta})$ given by composing φ_s with the structure map $\check{H}_{\leq d}(f_{\leq s}) \to \check{H}_{\leq d}(f_{\leq s+\delta})$. This finishes the proof. □

As an immediate corollary, we obtain the following.

Corollary 9.3.6. *If $f\colon X \to \mathbb{R}$ induces a filtration of paracompact Hausdorff spaces that is locally connected with respect to singular homology with coefficients in an abelian group G, then*

$$d_I\big(H_*(f_{\leq \bullet}; G), \check{H}_*(f_{\leq \bullet}; G)\big) = 0.$$

9.3.2. Lipschitz-continuity of the local connectedness shift

We finish by proving that the local connectedness shift is a Lipschitz map with respect to the supremum norm.

Proposition 9.3.7. *Let X be a topological space and H a functor from topological spaces to a category with a 0 object. Then for any functions $f, g\colon X \to \mathbb{R}$ we have*

$$|\operatorname{lcs}_H(f) - \operatorname{lcs}_H(g)| \leq 2 \cdot \|f - g\|_\infty.$$

Proof. We set $e = \|f - g\|_\infty$. It suffices to show that if $f_{\leq \bullet}$ is δ-HLC for some $\delta \geq 0$, then $g_{\leq \bullet}$ is $(\delta + 2e)$-HLC. So let $x \in X$ and consider indices s, t with $f(x) < s \leq s + \delta + 2e < t$, as well as a neighborhood V of x in X. If f is δ-HLC, then we can choose a neighborhood U of x such that $U \subseteq V$ and the inclusion $f_{\leq s+e} \cap U \to f_{\leq s+e+\delta} \cap V$ is taken to 0 by H. Now, note that we have $g_{\leq s} \subseteq f_{\leq s+e}$ and $f_{\leq s+e+\delta} \subseteq g_{\leq s+2e+\delta}$ because $e = \|f - g\|_\infty$. It follows that we have a commutative diagram

$$\begin{array}{ccccc} & & f_{\leq s+e} \cap U & \longrightarrow & f_{\leq s+e+\delta} \cap V & & \\ & \nearrow & & & & \searrow & \\ g_{\leq s} \cap U & & \longrightarrow & & & & g_{\leq s+2e+\delta} \cap V \end{array}$$

consisting of inclusion maps. In particular, $g_{\leq s} \cap U \to g_{\leq s+2e+\delta} \cap V$ factors through $f_{\leq s+e} \cap U \to f_{\leq s+e+\delta} \cap V$. The second map is taken to 0 by H, so the same is true for the first one. Hence, the sublevel set filtration of g is indeed $(\delta + 2e)$-HLC if the sublevel set filtratiton of f is δ-HLC, which proves the claim. □

10. Structure of semi-continuous q-tame persistence modules

In this chapter, we prove the structural results for persistence modules presented § 2.2.4. We start with some preparatory lemmas on semi-continuous interval modules and internal limits and colimits in Section 10.1. In Section 10.2, we then prove our main structure theorems for semi-continuous q-tame persistence modules.

At various points in the chapter, we will need to make some assumptions on the index set T, which are always explicitly stated in the respective results.

10.1. Basic results on semi-continuity

We start our discussion of semi-continuity by proving some preliminary results. First, we recall some terminology and consider some examples of semi-continuous persistence modules in § 10.1.1. In § 10.1.2, we then characterize semi-continuous interval modules. Finally, we show that weak isomorphisms are taken to isomorphisms by internal limits and colimits in § 10.1.3.

10.1.1. Semi-continuous persistence modules

Let us begin by introducing and recalling some terminology, partly taken from [CCdS16]. Recall that our index set T is *dense* if for all $s, u \in T$ with $s < u$ there exists $t \in T$ with $s < t < u$. If $N \subseteq I \subseteq T$ are subsets, N is said to be *coinitial* in I if for all $t \in I$ there exists $s \in N$ with $s \leq t$. N is said to be *cofinal* in I if for all $t \in I$ there exists $s \in N$ with $t \leq s$.

Definition 10.1.1. If M is a persistence module, we define a persistence module $\underline{M}$ by

$$\underline{M}_t = \lim_{s>t} M_s$$

with the obvious structure maps. The canonical maps $M_t \to \lim_{s>t} M_s$ form a morphism $M \to \underline{M}$, and we say that M is *upper semi-continuous (u.s.c.)* or *continuous from above* if this morphism is an isomorphism. We also define a persistence module $\overline{M}$ by

$$\overline{M}_t = \operatorname*{colim}_{s<t} M_s,$$

again with the obvious structure maps. The canonical maps $\operatorname{colim}_{s<t} M_s \to M_t$ form a morphism $\overline{M} \to M$, and we say that M is *lower semi-continuous (l.s.c.)* or *continuous from below* if this morphism is an isomorphism.

The constructions defined in Definition 10.1.1 clearly extend to endofunctors on the category of persistence modules. We have already mentioned in § 1.3.2 that the *radical* of a persistence module M may be seen as the image of the canonical map $\overline{M} \to M$, and that under some assumptions on the index set, the radical of a q-tame persistence module has a barcode. We record.

Theorem 10.1.2. *[CCdS16, Corollary 3.6.] Let T be a dense totally ordered set such that every interval in T has a countable coinitial subset. If M is a q-tame persistence module indexed by T, its radical* $\operatorname{rad} M$ *has a barcode.*

Semi-continuous persistence modules appear naturally in many different contexts. Some authors, especially within symplectic topology, even go as far as to consider almost exclusively semi-continuous persistence modules, e.g., Buhovsky et al. [Buh+22] and Polterovich and Shelukhin [PS16].

Example 10.1.3.

- One of the standard examples of a q-tame persistence module indexed by $\mathbb{R}$ that does not have a barcode in the usual sense is

$$\prod_{n\in\mathbb{N}} C\left(\left[0, n^{-1}\right)\right).$$

  It is upper semi-continuous by the previous two lemmas, so it has a multiplicative barcode by Theorem 10.2.4. Clearly, this is given by $([0, n^{-1}))_{n\in\mathbb{N}}$. In particular, this persistence module has a multiplicative barcode but no additive barcode.

- Consider the $\mathbb{R}$-indexed persistence module

$$\bigoplus_{n\in\mathbb{N}} C\left(\left(-n^{-1}, 0\right]\right).$$

  It is lower semi-continuous by the previous two lemmas and also q-tame. It has an additive barcode but no multiplicative barcode.

- Let $X\colon \mathbf{T} \to \mathbf{Top}$ be a diagram of topological spaces. If X_t is a compact Hausdorff space for all $t \in T$ and X is upper semi-continuous, i.e., $X_t \to \lim_{s>t} X_s$ is an isomorphism for all $t \in T$, then the persistent Čech homology $\check{H}_*(X)$ is upper semi-continuous by [ES52, Theorem X.3.1.]. This is also what we have considered before in the functional topology setting, where X was the sublevel set filtration of some function.

- Let X be a topological space and $f\colon X \to \mathbb{R}$ a continuous map. Write $f_{<t}$ for the open sublevel set of f at t. Since f is continuous, we have $f_{<t} = \operatorname{colim}_{s<t} f_{<s}$ for all t, where the colimit is taken in the category of topological spaces. Using the fact that the interval $(-\infty, t)$ has a countable cofinal subset, the main theorem in [May99, Section 14.6.] implies that the sublevel set persistence $H(f_{<\bullet})$ is lower semi-continuous. Here, H is any generalized homology theory with values in $\mathbf{Vec}_{\mathbb{F}}$.

- For any persistence module M indexed by $\mathbf{T}$, we get a dual persistence module $({}^{\vee}M)$ indexed by $\mathbf{T}^{\mathrm{op}}$ defined by composing the functor $M\colon \mathbf{T} \to \mathbf{Vec}$ with the contravariant functor $\operatorname{Hom}_{\mathbf{Vec}}(-, \mathbb{F})\colon \mathbf{Vec} \to \mathbf{Vec}$. We have reviewed this in more detail in § 3.2.2.

If M is lower semi-continuous, then $M^\vee$ is upper semi-continuous:

$$\begin{aligned}\underline{M^\vee}_t &= \lim_{s<t} \operatorname{Hom}(M_s, \mathbb{F}) \\ &= \operatorname{Hom}(\operatorname*{colim}_{s<t} M_s, \mathbb{F}) \\ &= \operatorname{Hom}(M_t, \mathbb{F}) \\ &= M_t^\vee,\end{aligned}$$

where equality should be interpreted as 'canonically isomorphic'.

However, if M is upper semi-continuous, $M^\vee$ need *not* be lower semi-continuous: Consider $M = \prod_{n\in\mathbb{N}} C\left([0, n^{-1})\right)$ as in the first example. An easy calculation shows that

$$\overline{M^\vee}_0 \cong \bigoplus_{n\in\mathbb{N}} \mathbb{F},$$

but we also have

$$M_0^\vee = \operatorname{Hom}\left(\prod_{n\in\mathbb{N}} \mathbb{F}, \mathbb{F}\right),$$

which is isomorphic to the double dual space of $\bigoplus_{n\in\mathbb{N}} \mathbb{F}$. Since no infinite-dimensional vector space is isomorphic to its double dual, we obtain that $M^\vee$ is not lower semi-continuous at 0.

10.1.2. Semi-continuous interval modules

As the next step, we analyze how the functors defined in Definition 10.1.1 behave for interval modules.

Definition 10.1.4. For $t \in T$, we write

$$\begin{aligned}\uparrow t &= \{s \in T \mid s > t\}, \\ \downarrow t &= \{s \in T \mid s < t\},\end{aligned}$$

for the strict upset and the strict downset of t. If $I \subseteq T$ is an interval, we define

$$\begin{aligned}\underline{I} &= \{t \in T \mid I \cap \uparrow t \text{ is non-empty and coinitial in } \uparrow t\}, \\ \overline{I} &= \{t \in T \mid I \cap \downarrow t \text{ is non-empty and cofinal in } \downarrow t\}, \\ \operatorname{rad} I &= \underline{I} \cap \overline{I}.\end{aligned}$$

Lemma 10.1.5. *Let $I \subseteq T$ be an interval. Then the sets $\underline{I}$, $\overline{I}$ and $\operatorname{rad} I$ are again intervals in T if they are non-empty.*

Proof. We only show the claim for $\underline{I}$, the other ones can be shown similarly. Let $s, u \in \underline{I}$ and $t \in T$ with $s < t < u$. We need to show that $t \in \underline{I}$, i.e., for $a \in \uparrow t$ we need to find $b \in I \cap \uparrow t$ with $b \leq a$.

First, we show that $u \in I$: Since $u \in \underline{I}$, there exists some $v \in I \cap \uparrow u$. Since $u \in \uparrow s$ and $s \in \underline{I}$, there exists $c \in I \cap \uparrow s$ with $c \leq u$. We have $c \leq u < v$ and $c, v \in I$. Since I is an interval, we get $u \in I$. In particular, we have $u \in I \cap \uparrow t$, so this set is non-empty.

Now consider $a \in \uparrow t$ again. We have $u \in I \cap \uparrow t$, so if $u \leq a$, we can set $b = u$ and are done. If $a < u$, pick $c \in I \cap \uparrow s$ with $c \leq a$. This is possible since $a \in \uparrow t \subseteq \uparrow s$ and $s \in \underline{I}$. Then, we have $c \leq a < u$ and $c, u \in I$, which implies $a \in I$. So in this case, we can simply set $b = a$ and the proof is finished. □

Recall that for an interval I, we denote the corresponding interval module as defined in §1.1.2 by $C(I)$. While we do not consider the empty set to be an interval, we set $C(\emptyset) = 0$. Then, the lemma below still holds true if the involved sets are empty.

Lemma 10.1.6. *For any interval $I \subseteq T$ we have*

$$\begin{aligned} \underline{C(I)} &\cong C(\underline{I}), \\ \overline{C(I)} &\cong C(\overline{I}), \\ \operatorname{rad} C(I) &\cong C(\operatorname{rad} I). \end{aligned}$$

Proof. Again, we only show the first isomorphism and the others can be shown analogously. For all $t \in \underline{I}$, we have

$$\underline{C(I)}_t = \lim_{s \in \uparrow t} C(I)_s = \lim_{s \in I \cap \uparrow t} C(I)_s = \lim_{s \in I \cap \uparrow t} \mathbb{F} = \mathbb{F}.$$

For $t \notin \underline{I}$, we have that $I \cap \uparrow t$ is empty or that there exists $t_0 \in \uparrow t$ such that there is no $s \in I \cap \uparrow t$ with $s \leq t_0$. In the first case, we have

$$\underline{C(I)}_t = \lim_{s \in \uparrow t} C(I)_s = \lim_{s \in \uparrow t} 0 = 0.$$

In the second case, we have

$$\underline{C(I)}_t = \lim_{s \in \uparrow t} C(I)_s = \lim_{t < s \leq t_0} C(I)_s = \lim_{t < s \leq t_0} 0 = 0.$$

Thus, $\underline{C(I)}$ and $C(\underline{I})$ agree pointwise. Clearly, their structure maps also agree and we obtain the claim. □

Semi-continuity is now easy to characterize for interval modules.

Lemma 10.1.7. *Let $I \subseteq T$ be an interval.*

1. *$C(I)$ is l.s.c. if and only if $I = \overline{I}$.*
2. *$C(I)$ is u.s.c. if and only if $I = \underline{I}$.*

Proof. Both claims follow immediately from Lemma 10.1.6 and the fact that for any two interval modules $C(J)$ and $C(J')$ we have $C(J) \cong C(J')$ if and only if $J = J'$. □

As an illustration, observe that a real interval $I \subseteq \mathbb{R}$ satisfies $I = \overline{I}$ if and only if $I = \mathbb{R}$ or $I = [a, b)$ for some $a \in \mathbb{R}$ and $b \in \mathbb{R} \cup \{\infty\}$. Analogously, we have $I = \underline{I}$ if and only if $I = \mathbb{R}$ or $I = (a, b]$ for some $a \in \mathbb{R} \cup \{-\infty\}$ and $b \in \mathbb{R}$.

For the proof of our main theorem, we will need the following.

Lemma 10.1.8. *Let T be a dense totally ordered set and $I \subseteq T$ an interval. If $\underline{I} = I$, then $\operatorname{rad} I$ is non-empty and $\underline{\operatorname{rad} I} = I$.*

Proof. First, note that $\underline{I} = I$ implies that $\underline{I}$ is non-empty. In other words, there exists $t \in T$ such that $I \cap \uparrow t \neq \emptyset$ is coinitial in $\uparrow t$. We will show that $I \cap \uparrow t = \operatorname{rad} I \cap \uparrow t$. This immediately implies that $\operatorname{rad} I$ is non-empty. It also shows that $\underline{I} \subseteq \underline{\operatorname{rad} I}$. The other inclusion obviously also holds, so in total we get $\underline{\operatorname{rad} I} = \underline{I} = I$ as claimed.

It is clear that $I \cap \uparrow t \supseteq \operatorname{rad} I \cap \uparrow t$. To see the other inclusion, consider $s \in I \cap \uparrow t$. We need to show that $s \in \operatorname{rad} I = I \cap \overline{I}$, so it is enough to check that $s \in \overline{I}$. So let $a \in \downarrow s$. We need to find $b \in I \cap \downarrow s$ with $b \geq a$.

If $a > t$, we have $a \in I$: Since $I \cap \uparrow t$ is coinitial in $\uparrow t$, we may choose $s' \in I \cap \uparrow t$ with $s' \leq a$. Now $s' \leq a < s$ and $s, s' \in I$, so $a \in I$ because I is an interval. In this case, we can set $b = a$ and are done.

If $a \leq t$, we use the fact that T is dense to choose $c \in T$ with $t < c < s$. By the same argument as before, we get $c \in I$ and can set $b = c$. This finishes the proof. ☐

10.1.3. Internal limits and colimits

As the next step, we will analyze how the internal limit and colimit functors in Definition 10.1.1 behave with respect to weak isomorphisms.

Lemma 10.1.9. *Let $\varphi\colon M \to N$ be a weak isomorphism of persistence modules. Then φ induces an isomorphism*

$$\overline{\varphi}\colon \overline{M} \to \overline{N}.$$

Proof. Since taking direct limits of vector spaces is exact, the same is true for the functor $\overline{(-)}$. Thus, this functor commutes with kernels and cokernels, so we get

$$\ker \overline{\varphi} \cong \overline{\ker \varphi}$$

and

$$\operatorname{coker} \overline{\varphi} \cong \overline{\operatorname{coker} \varphi}$$

Since $\ker \varphi$ and $\operatorname{coker} \varphi$ are ephemeral by assumption, we get that in both cases the right-hand side vanishes. So $\overline{\varphi}$ has trivial kernel and cokernel, which proves the claim. ☐

Lemma 10.1.10. *Assume that every interval in T has a countable coinitial subset. Let $\varphi\colon M \to N$ be a weak isomorphism of persistence modules. Then φ induces an isomorphism*

$$\underline{\varphi}\colon \underline{M} \to \underline{N}.$$

Proof. Consider the epi-mono-factorization of φ as

$$M \xrightarrow{p} \operatorname{im} \varphi \xhookrightarrow{i} N.$$

In order to show that $\underline{\varphi}$ is an isomorphism, it suffices to prove that $\underline{p}$ and $\underline{i}$ are isomorphisms. First, consider the short exact sequence

$$0 \longrightarrow \operatorname{im} \varphi \xrightarrow{i} N \longrightarrow \operatorname{coker} \varphi \longrightarrow 0.$$

Since taking inverse limits of vector spaces is left-exact, the functor $\underline{(-)}$ is also left-exact. Thus, we get an exact sequence

$$0 \longrightarrow \underline{\operatorname{im} \varphi} \xrightarrow{\underline{i}} \underline{N} \longrightarrow \underline{\operatorname{coker} \varphi}.$$

By assumption, $\operatorname{coker} \varphi$ is ephemeral, so we have $\underline{\operatorname{coker} \varphi} = 0$, which implies that $\underline{i}$ is an isomorphism. Next, consider the short exact sequence

$$0 \longrightarrow \ker \varphi \longrightarrow M \xrightarrow{p} \operatorname{im} \varphi \longrightarrow 0.$$

For each $t \in T$, the interval $\{s \in T \mid s > t\}$ has a countable coinitial subset by assumption. Since $\ker \varphi$ is ephemeral, the inverse system $(\ker \varphi_s)_{s>t}$ satisfies the Mittag-Leffler property for all $t \in T$. Thus, by [Gro61, Proposition 13.2.2.] the sequence

$$0 \longrightarrow \lim_{s>t} \ker \varphi_s \longrightarrow \lim_{s>t} M_s \longrightarrow \lim_{s>t} \operatorname{im} \varphi_s \longrightarrow 0$$

is exact for all $t \in T$. Consequently, the sequence

$$0 \longrightarrow \underline{\ker \varphi} \longrightarrow \underline{M} \xrightarrow{\underline{p}} \underline{\operatorname{im} \varphi} \longrightarrow 0$$

is also exact. We have $\underline{\ker \varphi} = 0$ since $\ker \varphi$ is assumed to be ephemeral. Hence, $\underline{p}$ is an isomorphism and the proof is finished. □

Remark 10.1.11. The previous lemma also holds if we replace the assumption on T by the assumption that T be a dense order. In this case, the lemma is a consequence of the fact that $\underline{(-)}$ defines a functor on the observable category of persistence modules and that weak isomorphisms turn to isomorphisms when mapped to the observable category ([CCdS16, Remark 2.12., Theorem 2.9.]).

10.2. Decompositions of semi-continuous persistence modules

We are now ready to give proofs of our main structure theorems. We provide some results on q-tame and semi-continuous direct sums and products of persistence modules in § 10.2.1. These are applied, together with the results from Section 10.1, to prove our structural results on semi-continuous q-tame persistence modules in § 10.2.2.

10.2.1. Sums and products of persistence modules

Before proceeding to our main theorems about additive and multiplicative barcodes of persistence modules, i.e., decompositions into direct sums and products of persistence modules, we record two more facts about these constructions.

Semi-continuity for sums and products

Lemma 10.2.1. *Let $(M_\alpha)_{\alpha \in A}$ be a collection of persistence modules.*

1. *$\bigoplus_{\alpha \in A} M_\alpha$ is l.s.c. if and only if all M_α are l.s.c.*
2. *$\prod_{\alpha \in A} M_\alpha$ is u.s.c. if and only if all M_α are u.s.c.*

Proof. It is easy to check that taking direct sums of persistence modules is conservative, so the canonical map $\bigoplus_{\alpha \in A} \overline{M_\alpha} \to \bigoplus_{\alpha \in A} M_\alpha$ is an isomorphism if and only if all M_α are l.s.c. Colimits commute with each other, so we also have a canonical isomorphism

$$\overline{\bigoplus_{\alpha \in A} M_\alpha} \cong \bigoplus_{\alpha \in A} \overline{M_\alpha}.$$

This implies the first claim. The second claim follows analogously because taking products of persistence modules is also conservative and limits commute with each other. □

In particular, we obtain that if an l.s.c. persistence module has an additive barcode, then all intervals I appearing in this barcode satisfy $I = \overline{I}$. Similarly, if a u.s.c. persistence module has a multiplicative barcode, then all intervals I in this barcode satisfy $I = \underline{I}$.

Weak isomorphism from sum to product As the last preparatory results, we have the following.

Proposition 10.2.2. *Let $(M_\alpha)_{\alpha\in A}$ be a collection of persistence modules such that their product $\prod_{\alpha\in A} M_\alpha$ is q-tame. Then the canonical map*

$$\bigoplus_{\alpha\in A} M_\alpha \to \prod_{\alpha\in A} M_\alpha$$

is a weak isomorphism.

Proof. Denote the map above by φ. Clearly, φ has trivial kernel. Thus, it suffices to show that $\operatorname{coker}\varphi$ is ephemeral. So let $s,t\in T$ with $s<t$ and consider the diagram

$$\begin{array}{ccccc}
\left(\bigoplus_{\alpha\in A} M_\alpha\right)_t & \xhookrightarrow{\varphi_t} & \left(\prod_{\alpha\in A} M_\alpha\right)_t & \xrightarrow{p_t} & \operatorname{coker}\varphi_t \\
\sigma_{s,t}\uparrow & & \pi_{s,t}\uparrow & & \uparrow\gamma_{s,t} \\
\left(\bigoplus_{\alpha\in A} M_\alpha\right)_s & \xhookrightarrow{\varphi_s} & \left(\prod_{\alpha\in A} M_\alpha\right)_s & \xrightarrow{p_s} & \operatorname{coker}\varphi_s,
\end{array}$$

where we added some shorthand notation for the structure maps of the persistence modules we consider. We need to check that $\gamma_{s,t}=0$. Since p_s is epi, it is enough to show $\gamma_{s,t}\circ p_s=0$. Commutativity of the above diagram implies that

$$\gamma_{s,t}\circ p_s = p_t\circ\pi_{s,t}.$$

Note that $p_t\circ\varphi_t=0$, so we are done if we can show that $\pi_{s,t}$ factors through φ_t. To see that this is the case, we factor $\sigma_{s,t}$ and $\pi_{s,t}$ through their images to obtain a diagram

$$\begin{array}{ccc}
\left(\bigoplus_{\alpha\in A} M_\alpha\right)_t & \xhookrightarrow{\varphi_t} & \left(\prod_{\alpha\in A} M_\alpha\right)_t \\
\uparrow & & \uparrow \\
\operatorname{im}\sigma_{s,t} & \xrightarrow{\psi_{s,t}} & \operatorname{im}\pi_{s,t} \\
\uparrow & & \uparrow \\
\left(\bigoplus_{\alpha\in A} M_\alpha\right)_s & \xhookrightarrow{\varphi_s} & \left(\prod_{\alpha\in A} M_\alpha\right)_s
\end{array}.$$

We can canonically identify

$$\operatorname{im}\sigma_{s,t}\cong\bigoplus_{\alpha\in A}\operatorname{im}(M_\alpha)_{s,t}$$

and

$$\operatorname{im}\pi_{s,t}\cong\prod_{\alpha\in A}\operatorname{im}(M_\alpha)_{s,t}.$$

From commutativity of the previous diagram, it is easy to see that under this identification $\psi_{s,t}$ is simply the canonical inclusion of the direct sum into the product. But here, this map is an isomorphism since $\operatorname{im}\pi_{s,t}$ is finite-dimensional by our q-tameness assumption. Thus, we can invert $\psi_{s,t}$, yielding a factorization of $\pi_{s,t}$ as

$$\left(\prod_{\alpha\in A} M_\alpha\right)_s \twoheadrightarrow \operatorname{im}\pi_{s,t} \xrightarrow{\psi_{s,t}^{-1}} \operatorname{im}\sigma_{s,t} \hookrightarrow \left(\bigoplus_{\alpha\in A} M_\alpha\right)_t \xhookrightarrow{\varphi_t} \left(\prod_{\alpha\in A} M_\alpha\right)_t.$$

As explained above, this finishes the proof. □

10.2.2. Structural results for semi-continuous q-tame persistence modules

We are now ready to prove the main results of this chapter. While not explicitly stated by Chazal et al., the next result is an immediate corollary of the fact that radicals of q-tame persistence modules have barcodes under appropriate assumptions on the index set. [CCdS16, Corollary 3.6.].

Theorem 10.2.3. *Let T be a dense totally ordered set such that every interval in T has a countable coinitial subset. Then every q-tame lower semi-continuous persistence module indexed by T has a barcode.*

Proof. By definition, a persistence module M is lower semi-continuous if the canonical morphism $\overline{M} \to M$ is an isomorphism. In particular, a lower semi-continuous persistence module is isomorphic to its radical. Thus, the claim is an immediate consequence of Theorem 10.1.2. □

Our main contribution of this chapter is the following.

Theorem 10.2.4. *Let T be a dense totally ordered set such that every interval in T has a countable coinitial subset. Then for every q-tame upper semi-continuous persistence module M indexed by T there exists a collection of intervals $(I_\alpha)_{\alpha\in A}$, unique up to reordering, such that*

$$M \cong \prod_{\alpha\in A} C(I_\alpha).$$

Proof. Under our assumptions, $\operatorname{rad} M$ has a barcode by Theorem 10.1.2, say $(I_\alpha)_{\alpha\in A}$. We claim that M is isomorphic to the product over the interval modules $C(\underline{I_\alpha})$.

First, we have

$$M \cong \underline{M}$$

since we assume M to be u.s.c. Since the canonical map $\operatorname{rad} M \to M$ is a weak isomorphism (as a consequence of [CCdS16, Proposition 2.11.]), Lemma 10.1.10 implies

$$\underline{M} \cong \underline{\operatorname{rad} M}.$$

Recall that $\underline{(-)}$ is a functor, so

$$\underline{\operatorname{rad} M} \cong \underline{\bigoplus_{\alpha\in A} C(I_\alpha)}$$

because the barcode of $\operatorname{rad} M$ is given by the I_α. The inclusion of the direct sum into the product is a weak isomorphism in the q-tame case (Proposition 10.2.2), so Lemma 10.1.10 implies

$$\underline{\bigoplus_{\alpha\in A} C(I_\alpha)} \cong \underline{\prod_{\alpha\in A} C(I_\alpha)}.$$

Since limits commute with products we also get

$$\underline{\prod_{\alpha \in A} C(I_\alpha)} \cong \prod_{\alpha \in A} \underline{C(I_\alpha)}.$$

We have $\underline{C(I_\alpha)} \cong C(\underline{I_\alpha})$ by Lemma 10.1.6, so that

$$\prod_{\alpha \in A} \underline{C(I_\alpha)} \cong \prod_{\alpha \in A} C(\underline{I_\alpha}).$$

Putting everything together yields that M is indeed isomorphic to the product over the $C(\underline{I_\alpha})$.

The uniqueness part of the statement essentially follows by reversing the above argument. Suppose $(J_\beta)_{\beta \in B}$ are also intervals such that

$$M \cong \prod_{\beta \in B} C(J_\beta).$$

We want to prove that $(J_\beta)_{\beta \in B}$ and $(\underline{I_\alpha})_{\alpha \in A}$ agree up to reordering. Note that this in particular implies that each $\underline{I_\alpha}$ is non-empty. Since $M \cong \prod_\beta C(J_\beta)$ is u.s.c. each factor $C(J_\beta)$ must be u.s.c. as well by Lemma 10.2.1. Together with Lemma 10.1.7 this yields

$$\underline{J_\beta} = J_\beta.$$

Thus, by Lemma 10.1.8 we get that $\operatorname{rad} J_\beta$ is non-empty and consequently an interval for all β. Next, we will show that $(\operatorname{rad} J_\beta)_{\beta \in B}$ is a barcode for $\operatorname{rad} M$: Consider

$$\operatorname{rad} M \cong \operatorname{rad} \prod_{\beta \in B} C(J_\beta) = \operatorname{im}\left(\overline{\prod_{\beta \in B} C(J_\beta)} \to \prod_{\beta \in B} C(J_\beta) \right).$$

Recall that the inclusion of the direct sum into the product is a weak isomorphism in our case, so together with Lemma 10.1.9 we obtain that

$$\operatorname{im}\left(\overline{\prod_{\beta \in B} C(J_\beta)} \to \prod_{\beta \in B} C(J_\beta) \right) \cong \operatorname{im}\left(\overline{\bigoplus_{\beta \in B} C(J_\beta)} \to \prod_{\beta \in B} C(J_\beta) \right),$$

where the map on the right is equal to the composition of the natural map

$$\overline{\bigoplus_{\beta \in B} C(J_\beta)} \to \bigoplus_{\beta \in B} C(J_\beta)$$

and the inclusion $\bigoplus_{\beta \in B} C(J_\beta) \to \prod_{\beta \in B} C(J_\beta)$. Since this inclusion is mono, we get

$$\operatorname{im}\left(\overline{\bigoplus_{\beta \in B} C(J_\beta)} \to \prod_{\beta \in B} C(J_\beta) \right) \cong \operatorname{im}\left(\overline{\bigoplus_{\beta \in B} C(J_\beta)} \to \bigoplus_{\beta \in B} C(J_\beta) \right).$$

Direct sums and the functor $\overline{(-)}$ commute. The same is true for direct sums and images, so we get

$$\operatorname{im}\left(\overline{\bigoplus_{\beta \in B} C(J_\beta)} \to \bigoplus_{\beta \in B} C(J_\beta) \right) \cong \bigoplus_{\beta \in B} \operatorname{im}(\overline{C(J_\beta)} \to C(J_\beta)) = \bigoplus_{\beta \in B} \operatorname{rad} C(J_\beta).$$

We have $\operatorname{rad} C(J_\beta) \cong C(\operatorname{rad} J_\beta)$ by Lemma 10.1.6, so we get

$$\bigoplus_{\beta \in B} \operatorname{rad} C(J_\beta) \cong \bigoplus_{\beta \in B} C(\operatorname{rad} J_\beta).$$

In total, we have shown that $(\operatorname{rad} J_\beta)_{\beta \in B}$ is indeed a barcode for $\operatorname{rad} M$.

Using the Krull–Remak–Schmidt–Azumaya Theorem, we obtain that the two barcodes $(I_\alpha)_{\alpha \in A}$ and $(\operatorname{rad} J_\beta)_{\beta \in B}$ agree up to reordering. This implies that also $(\underline{I_\alpha})_{\alpha \in A}$ and $(\underline{\operatorname{rad} J_\beta})_{\beta \in B}$ agree up to reordering. Now recall that we have $\underline{J_\beta} = \overline{J_\beta}$ for all β because M is u.s.c. By Lemma 10.1.8, we get that $(\underline{\operatorname{rad} J_\beta})_{\beta \in B} = (\overline{J_\beta})_\beta$. Thus, $(\underline{I_\alpha})_{\alpha \in A}$ and $(J_\beta)_{\beta \in B}$ agree up to reordering. This finishes the proof. □

Bibliography

[Azu50] Gorô Azumaya. “Corrections and supplementaries to my paper concerning Krull–Remak–Schmidt’s theorem”. *Nagoya Math. J.* 1 (1950) (cit. on p. 5).

[Bar94] S. A. Barannikov. “The framed Morse complex and its invariants”. *Singularities and bifurcations*. Vol. 21. Adv. Soviet Math. Amer. Math. Soc., 1994 (cit. on pp. 4, 9).

[BM62] M. G. Barratt and John Milnor. “An example of anomalous singular homology”. *Proc. Amer. Math. Soc.* 13 (1962) (cit. on p. 102).

[Bau21] Ulrich Bauer. “Ripser: efficient computation of Vietoris-Rips persistence barcodes”. *J. Appl. Comput. Topol.* 5.3 (2021) (cit. on pp. 18, 23, 24, 62, 73, 74).

[BL15] Ulrich Bauer and Michael Lesnick. “Induced matchings and the algebraic stability of persistence barcodes”. *J. Comput. Geom.* 6.2 (2015) (cit. on pp. 7, 8, 17, 20, 37, 58).

[BL20] Ulrich Bauer and Michael Lesnick. “Persistence Diagrams as Diagrams: A Categorification of the Stability Theorem”. *Topological Data Analysis*. Vol. 15. Abel Symposia. Springer, 2020 (cit. on pp. 19, 20, 37, 39, 40).

[BMS21] Ulrich Bauer, Anibal M. Medina-Mardones, and Maximilian Schmahl. “Persistent homology for functionals”. Preprint, arXiv:2107.14247. 2021 (cit. on pp. 17, 32, 33).

[BS21a] Ulrich Bauer and Maximilian Schmahl. “Lifespan Functors and Natural Dualities in Persistent Homology”. Preprint, arXiv:2012.12881. 2021 (cit. on pp. 17, 32).

[BS21b] Ulrich Bauer and Maximilian Schmahl. *Ripser for image persistence*. GitHub. 2021 (cit. on pp. 24, 74).

[BS22] Ulrich Bauer and Maximilian Schmahl. “Efficient Computation of Image Persistence”. Preprint, arXiv:2201.04170. 2022 (cit. on pp. 17, 24, 32, 33, 74).

[Ble+21] Michael Bleher et al. “Topology identifies emerging adaptive mutations in SARS-CoV-2”. Preprint, arXiv:2106.07292. 2021 (cit. on p. 3).

[BC20] Magnus Bakke Botnan and William Crawley-Boevey. “Decomposition of persistence modules”. *Proc. Amer. Math. Soc.* 148.11 (2020) (cit. on pp. 5, 10, 58).

[Bot80] Raoul Bott. “Marston Morse and his mathematical works”. *Bull. Amer. Math. Soc. (N.S.)* 3.3 (1980) (cit. on p. 27).

[Bre97] Glen E. Bredon. *Sheaf theory*. Springer, 1997 (cit. on pp. 29, 31, 96, 97).

[BP69] Hans-Berndt Brinkmann and Dieter Puppe. *Abelsche und exakte Kategorien, Korrespondenzen*. Springer, 1969 (cit. on p. 40).

[BM21] Peter Bubenik and Nikola Milićević. "Homological algebra for persistence modules". *Found. Comput. Math.* 21.5 (2021) (cit. on p. 20).

[Buh+22] Lev Buhovsky et al. "Coarse nodal count and topological persistence". Preprint, arXiv:2206.06347. 2022 (cit. on pp. 30, 118).

[CL11] Francesca Cagliari and Claudia Landi. "Finiteness of rank invariants of multi-dimensional persistent homology groups". *Applied Mathematics Letters* 24.4 (2011) (cit. on pp. 29, 100).

[Cas20] Álvaro Torras Casas. "Distributing Persistent Homology via Spectral Sequences". Preprint, arXiv:1907.05228. 2020 (cit. on p. 18).

[Cha+09] Frédéric Chazal, David Cohen-Steiner, Marc Glisse, Leonidas J. Guibas, and Steve Y. Oudot. "Proximity of Persistence Modules and Their Diagrams". *Proceedings of the Twenty-Fifth Annual Symposium on Computational Geometry.* 2009 (cit. on pp. 6–8).

[CCdS16] Frédéric Chazal, William Crawley-Boevey, and Vin de Silva. "The observable structure of persistence modules". *Homology Homotopy Appl.* 18.2 (2016) (cit. on pp. 10–12, 29, 31, 117, 118, 122, 124).

[Cha+16] Frédéric Chazal, Vin de Silva, Marc Glisse, and Steve Oudot. *The structure and stability of persistence modules*. Springer, 2016 (cit. on pp. 4, 5, 7, 83).

[CdSO14] Frédéric Chazal, Vin de Silva, and Steve Oudot. "Persistence stability for geometric complexes". *Geom. Dedicata* 173 (2014) (cit. on p. 8).

[Cha+11] Frédéric Chazal, Leonidas J. Guibas, Steve Y. Oudot, and Primoz Skraba. "Scalar field analysis over point cloud data". *Discrete Comput. Geom.* 46.4 (2011) (cit. on p. 18).

[CK11] Chao Chen and Michael Kerber. "Persistent homology computation with a twist". *Proceedings of the 27th European Workshop on Computational Geometry.* 2011 (cit. on pp. 9, 18).

[Clo+20] James Clough et al. "A Topological Loss Function for Deep-Learning based Image Segmentation using Persistent Homology". *IEEE transactions on pattern analysis and machine intelligence* (2020) (cit. on p. 3).

[CEH07] David Cohen-Steiner, Herbert Edelsbrunner, and John Harer. "Stability of persistence diagrams". *Discrete Comput. Geom.* 37.1 (2007) (cit. on pp. 6–8).

[Coh+10] David Cohen-Steiner, Herbert Edelsbrunner, John Harer, and Yuriy Mileyko. "Lipschitz functions have L_p-stable persistence". *Found. Comput. Math.* 10.2 (2010) (cit. on p. 8).

[Coh+09] David Cohen-Steiner, Herbert Edelsbrunner, John Harer, and Dmitriy Morozov. "Persistent homology for kernels, images, and cokernels". *Proceedings of the Twentieth Annual ACM-SIAM Symposium on Discrete Algorithms.* 2009 (cit. on pp. 18, 23, 69).

[CEM06] David Cohen-Steiner, Herbert Edelsbrunner, and Dmitriy Morozov. "Vines and vineyards by updating persistence in linear time". *Computational geometry.* 2006 (cit. on p. 9).

[CK18] René Corbet and Michael Kerber. “The representation theorem of persistence revisited and generalized”. *J. Appl. Comput. Topol.* 2.1-2 (2018) (cit. on p. 53).

[Cra15] William Crawley-Boevey. “Decomposition of pointwise finite-dimensional persistence modules”. *J. Algebra Appl.* 14.5 (2015) (cit. on pp. 4, 5, 10, 58).

[Dab+12] Y. Dabaghian, F. Mémoli, L. Frank, and G. Carlsson. “A Topological Paradigm for Hippocampal Spatial Map Formation Using Persistent Homology”. *PLOS Computational Biology* 8.8 (2012) (cit. on p. 3).

[Dah21] Lucas Dahinden. “C^0-stability of topological entropy for contactomorphisms”. *Commun. Contemp. Math.* 23.6 (2021) (cit. on p. 53).

[dSMV11] Vin de Silva, Dmitriy Morozov, and Mikael Vejdemo-Johansson. “Dualities in persistent (co)homology”. *Inverse Problems* 27.12 (2011) (cit. on pp. 9, 18, 19, 60–62, 69).

[DHS10] Ulrich Dierkes, Stefan Hildebrandt, and Friedrich Sauvigny. *Minimal surfaces.* 2nd ed. Springer, 2010 (cit. on p. 86).

[Dou31] Jesse Douglas. “Solution of the problem of Plateau”. *Trans. Amer. Math. Soc.* 33.1 (1931) (cit. on pp. 25, 94).

[Dow52] C. H. Dowker. “Homology groups of relations”. *Ann. of Math. (2)* 56 (1952) (cit. on p. 14).

[EK00] Katsuya Eda and Kazuhiro Kawamura. “The surjectivity of the canonical homomorphism from singular homology to Čech homology”. *Proc. Amer. Math. Soc.* 128.5 (2000) (cit. on p. 90).

[EH10] Herbert Edelsbrunner and John L. Harer. *Computational topology: An introduction.* American Mathematical Society, 2010 (cit. on p. 4).

[ES52] Samuel Eilenberg and Norman Steenrod. *Foundations of algebraic topology.* Princeton University Press, 1952 (cit. on pp. 13, 14, 96, 106, 118).

[Eng89] Ryszard Engelking. *General topology.* Translated from the Polish by the author. Heldermann, 1989 (cit. on p. 109).

[GS18] Dejan Govc and Primoz Skraba. “An approximate nerve theorem”. *Found. Comput. Math.* 18.5 (2018) (cit. on p. 31).

[Gra12] Marco Grandis. *Homological algebra. The interplay of homology with distributive lattices and orthodox semigroups.* World Scientific Publishing, 2012 (cit. on pp. 21, 40, 44, 47, 50).

[Gro61] A. Grothendieck. “Éléments de géométrie algébrique. III. Étude cohomologique des faisceaux cohérents.” *Inst. Hautes Études Sci. Publ. Math.* 11 (1961) (cit. on p. 122).

[HS54] Graham Higman and A. H. Stone. “On inverse systems with trivial limits”. *J. London Math. Soc.* 29 (1954) (cit. on p. 60).

[Höp83] Michael Höppner. “A note on the structure of injective diagrams”. *Manuscripta Math.* 44.1-3 (1983) (cit. on pp. 20, 59).

[HL81a] Michael Höppner and Helmut Lenzing. “Diagrams over ordered sets: A simple model of abelian group theory”. *Abelian group theory (Oberwolfach, 1981).* Vol. 874. Lecture Notes in Math. Springer, 1981 (cit. on pp. 20, 53, 59).

[HL81b] Michael Höppner and Helmut Lenzing. “Projective diagrams over partially ordered sets are free”. *J. Pure Appl. Algebra* 20.1 (1981) (cit. on pp. 20, 59).

[Hu+19] Xiaoling Hu, Fuxin Li, Dimitris Samaras, and Chao Chen. “Topology-Preserving Deep Image Segmentation”. *Advances in Neural Information Processing Systems*. Vol. 32. 2019 (cit. on p. 3).

[Jen70] C. U. Jensen. “On the vanishing of $\varprojlim^{(i)}$”. *J. Algebra* 15 (1970) (cit. on p. 59).

[JS90] J. Jost and M. Struwe. “Morse-Conley theory for minimal surfaces of varying topological type”. *Invent. Math.* 102.3 (1990) (cit. on p. 26).

[Jos91] Jürgen Jost. *Two-dimensional geometric variational problems*. John Wiley & Sons, Ltd., 1991 (cit. on p. 26).

[Kel61] G. M. Kelly. “The exactness of Čech homology over a vector space”. *Proc. Cambridge Philos. Soc.* 57 (1961) (cit. on p. 13).

[Les15] Michael Lesnick. “The theory of the interleaving distance on multidimensional persistence modules”. *Found. Comput. Math.* 15.3 (2015) (cit. on pp. 6, 53).

[LMT22] Umberto Lupo, Anibal M. Medina-Mardones, and Guillaume Tauzin. “Persistence Steenrod modules”. *J. Appl. Comput. Topol.* (2022) (cit. on p. 18).

[Mar58] S. Mardešić. “Equivalence of singular and Čech homology for ANR-s. Application to unicoherence”. *Fund. Math.* 46 (1958) (cit. on p. 106).

[Mar59] Sibe Mardešić. “Comparison of singular and Čech homology in locally connected spaces”. *Michigan Math. J.* 6 (1959) (cit. on pp. 29, 31, 111, 113).

[MN21] Fernando C. Marques and André Neves. “Morse theory for the area functional”. *São Paulo J. Math. Sci.* 15.1 (2021) (cit. on p. 26).

[May99] J. P. May. *A concise course in algebraic topology*. University of Chicago Press, 1999 (cit. on p. 118).

[Mil63] J. Milnor. *Morse theory*. Based on lecture notes by M. Spivak and R. Wells. Princeton University Press, 1963 (cit. on p. 25).

[Mon20] Rafael Montezuma. “A mountain pass theorem for minimal hypersurfaces with fixed boundary”. *Calc. Var. Partial Differential Equations* 59.6 (2020) (cit. on p. 26).

[Mor37] Marston Morse. “Functional topology and abstract variational theory”. *Ann. of Math. (2)* 38.2 (1937) (cit. on pp. 25, 27, 103).

[Mor38] Marston Morse. *Functional topology and abstract variational theory*. Gauthier-Villars, 1938 (cit. on pp. 25, 28, 85, 87, 89, 91–93, 102).

[Mor40] Marston Morse. “Rank and span in functional topology”. *Ann. of Math. (2)* 41 (1940) (cit. on pp. 25–28, 80, 82, 85, 92, 94, 102).

[Mor43] Marston Morse. “Functional topology”. *Bull. Amer. Math. Soc.* 49 (1943) (cit. on p. 84).

[Mor96] Marston Morse. *The calculus of variations in the large*. Reprint of the 1932 original. American Mathematical Society, 1996 (cit. on p. 79).

[MT39] Marston Morse and C. Tompkins. "The existence of minimal surfaces of general critical types". *Ann. of Math. (2)* 40.2 (1939) (cit. on pp. 25, 26, 28, 82, 84–87, 92–94, 96, 102, 103).

[Mun00] James R. Munkres. *Topology*. Prentice Hall, 2000 (cit. on p. 14).

[Nak+15] Takenobu Nakamura, Yasuaki Hiraoka, Akihiko Hirata, Emerson G Escolar, and Yasumasa Nishiura. "Persistent homology and many-body atomic structure for medium-range order in the glass". *Nanotechnology* 26.30 (2015) (cit. on p. 3).

[Oud15] Steve Y. Oudot. *Persistence theory: from quiver representations to data analysis*. American Mathematical Society, 2015 (cit. on p. 4).

[Par70] Bodo Pareigis. *Categories and functors*. Translated from the German. Academic Press, 1970 (cit. on p. 5).

[Pol+20] L. Polterovich, D. Rosen, K. Samvelyan, and J. Zhang. *Topological Persistence in Geometry and Analysis*. American Mathematical Society, 2020 (cit. on p. 4).

[PS16] Leonid Polterovich and Egor Shelukhin. "Autonomous Hamiltonian flows, Hofer's geometry and persistence modules". *Selecta Math. (N.S.)* 22.1 (2016) (cit. on p. 118).

[RB21] Yohai Reani and Omer Bobrowski. "Cycle Registration in Persistent Homology with Applications in Topological Bootstrap". Preprint, arXiv:2101.00698. 2021 (cit. on p. 18).

[Roo62] Jan-Erik Roos. "Bidualité et structure des foncteurs dérivés de $\varinjlim$ dans la catégorie des modules sur un anneau régulier". *C. R. Acad. Sci. Paris* 254 (1962) (cit. on p. 59).

[Sal99] Dietmar Salamon. "Lectures on Floer homology". *Symplectic geometry and topology*. Vol. 7. IAS/Park City Math. Ser. Amer. Math. Soc., 1999 (cit. on p. 25).

[Sch22a] Maximilian Schmahl. "Comparison of persistent singular and Čech homology for locally connected filtrations". Preprint, arXiv:2207.00848. 2022 (cit. on pp. 17, 32, 33).

[Sch22b] Maximilian Schmahl. "Structure of semi-continuous q-tame persistence modules". *Homology Homotopy Appl.* 24.1 (2022) (cit. on pp. 17, 32, 33).

[Skl80] E. G. Skljarenko. "On homologically locally connected spaces". *Izv. Akad. Nauk SSSR Ser. Mat.* 44.6 (1980) (cit. on p. 29).

[Str84] Michael Struwe. "On a critical point theory for minimal surfaces spanning a wire in $\mathbf{R}^n$". *J. Reine Angew. Math.* 349 (1984) (cit. on p. 26).

[Str88] Michael Struwe. *Plateau's problem and the calculus of variations*. Princeton University Press, 1988 (cit. on pp. 25, 27, 82, 84, 86, 87, 94, 96).

[tDie08] Tammo tom Dieck. *Algebraic topology*. European Mathematical Society, 2008 (cit. on p. 13).

[Vie27] L. Vietoris. "Über den höheren Zusammenhang kompakter Räume und eine Klasse von zusammenhangstreuen Abbildungen". *Math. Ann.* 97.1 (1927) (cit. on p. 14).

[Wei11] Shmuel Weinberger. “What is. . .persistent homology?” *Notices Amer. Math. Soc.* 58.1 (2011) (cit. on pp. 28, 100).

[Wil49] Raymond Louis Wilder. *Topology of Manifolds*. American Mathematical Society, 1949 (cit. on p. 96).

[Zel51] Daniel Zelinsky. “Rings with ideal nuclei”. *Duke Math. J.* 18 (1951) (cit. on p. 59).

[ZC05] Afra Zomorodian and Gunnar Carlsson. “Computing persistent homology”. *Discrete Comput. Geom.* 33.2 (2005) (cit. on pp. 4, 9, 53).