Technische Universität München

TUM School of Computation, Information and Technology

Geometric Complexes in Topological Data Analysis

Fabian Roll

Vollständiger Abdruck der von der TUM School of Computation, Information and Technology der Technischen Universität München zur Erlangung eines

Doktors der Naturwissenschaften (Dr. rer. nat)

genehmigten Dissertation.

Vorsitz:		Prof. Dr. Claudia Scheimbauer
Prüfende der Dissertation:	1.	Prof. Dr. Ulrich Bauer
	2.	Prof. Dr. Facundo Mémoli
	3.	Prof. Dr. Žiga Virk

Die Dissertation wurde am 16.02.2024 bei der Technischen Universität München eingereicht und durch die TUM School of Computation, Information and Technology am 02.07.2024 angenommen.

Acknowledgements

I want to thank Ulrich Bauer for giving me the opportunity to be part of his applied and computational topology group, and for advising me throughout my time at TUM. I am most grateful to him for giving me the freedom to develop and pursue my own research interests, which has helped me grow academically, but also personally. I also want to thank Facundo Mémoli for giving me the opportunity to do a research stay with his group at OSU, which I really enjoyed and that was eye-opening in many ways. This research has been supported by the TopMath program and the Deutsche Forschungsgemeinschaft (DFG – German Research Foundation) – Project-ID 195170736 – TRR109.

On a more personal level, I am grateful to have met so many interesting people over the last years. In particular, I want to thank Maria Elena Gonzalez Herrero, David Hien, Fabian Lenzen, Alexander Rolle, Jannik Steinmeier, and Nico Stucki for accompanying me during this journey, and for being more than just colleagues. Special thanks go to David Hien and Fabian Lenzen, with whom I had many interesting discussions and shared many memorable moments at conferences. Finally and most importantly, I want to thank my whole family for all the ways they are part of my life, in particular Laura for her love and support, especially during the last months.

Abstract

We study geometric complexes and their interaction with the persistent homology pipeline. We generalize a famous result of Rips on the contractibility of Vietoris–Rips complexes, with strong implications to the computation of persistent homology for tree-like metric data. We establish a close connection between discrete Morse theory and persistent homology with applications to shape reconstruction. We provide a variety of nerve theorems suitable for topological data analysis.

Contents

1. Introduction

Topological data analysis [Car09; EM13; Oud15; Ghr14] is a diverse area of research that aims to make concepts from the mathematical field of topology applicable to real-world data. By now, there is an extensive collection of papers available in the DONUT database [GLR22] that showcase the applicability of topological data analysis to other areas. For example, topological methods find application in biology [Tay+15; Mod+24], time series analysis [PH15; Bau+23a], image analysis [Car+08; Stu+23], and machine learning in general [HMR21; RS23]. The field of topological data analysis also stimulated other fields of mathematics [UZ16; Zav24].

We now describe the *persistent homology pipeline*, also called the *persistence pipeline*, which is the most prominent tool in topological data analysis. To capture the shape of a given point cloud, one typically constructs a combinatorial model that is suitable for computations. Such a combinatorial model can be obtained, for example, by using *geometric complexes*, which are constructed from the underlying metric structure of the point cloud and depend on one or more parameters. There are three fundamental geometric complexes, and variants thereof, that can be associated to a point cloud: The *Čech* and *Delaunay complex* (see Section 2.1.1), and the *Vietoris–Rips complex* (see Section 2.1.2), all of which depend on a single parameter. By increasing the parameter of any geometric complex from above, one obtains a filtration of simplicial complexes, whose *persistent homology* (see Section 2.2) captures the evolution of nonlinear topological structures and that can be encoded combinatorially in the associated *(persistence) barcode* (compare Fig. 1 and Section 2.2). Persistence barcodes can be computed, for example, by using any of the following software packages [The15; HG16; Bau21].

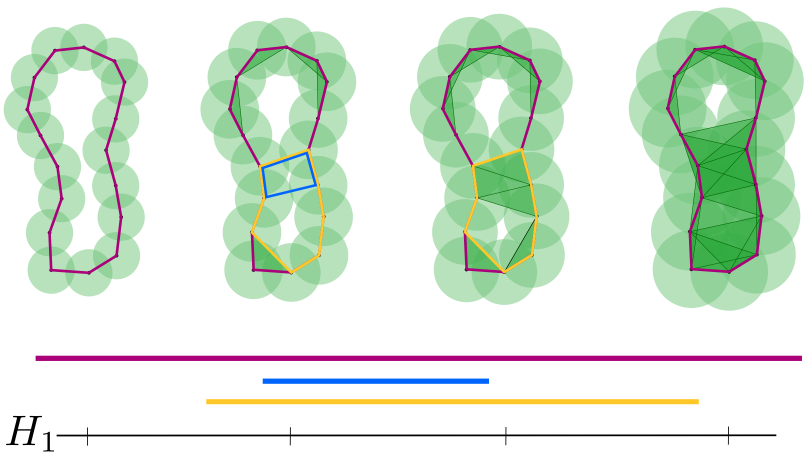

Figure 1: Čech filtration and its first persistence barcode associated to a point cloud.

This thesis aims to contribute to the study of geometric complexes and their interaction with the persistent homology pipeline. For more applications and recent results, see, for example, any of the following references and the references therein [Kah11; AA17; LV23; Rol22; BL22; BDK23; GM23; KM21; BOO22; Cha+23; BLL23; BS23].

Outline In Section 1.1 we briefly summarize the main contributions in this thesis. This thesis is based on and in large parts identical to the following three papers to which the author of this thesis contributed, in terms of the mathematical content and writing, in an essential way. The Sections 1.2 to 1.4 summarize the main statements Theorems A to C and E to H and Corollary D of these papers.

[Bau+23b] Ulrich Bauer, Michael Kerber, Fabian Roll, and Alexander Rolle. "A unified view on the functorial nerve theorem and its variations". In: *Expo. Math.* 41.4 (2023). DOI: 10.1016/j.exmath.2023.04.005.

[BR22] Ulrich Bauer and Fabian Roll. "Gromov Hyperbolicity, Geodesic Defect, and Apparent Pairs in Vietoris-Rips Filtrations". In: *38th International Symposium on Computational Geometry (SoCG 2022)*. Vol. 224. Leibniz International Proceedings in Informatics (LIPIcs). 2022. DOI: 10.4230/LIPIcs.SoCG.2022.15.

[BR24] Ulrich Bauer and Fabian Roll. "Wrapping Cycles in Delaunay Complexes: Bridging Persistent Homology and Discrete Morse Theory". In: *40th International Symposium on Computational Geometry (SoCG 2024)*. Vol. 293. Leibniz International Proceedings in Informatics (LIPIcs). 2024. DOI: 10.4230/LIPIcs.SoCG.2024.15.

We briefly discuss some aspects of the authors contributions to the papers above and the structure of this thesis. Sections 1.2, 2.1.2, 2.3.1 and 3 are essentially an updated and extended version of [BR22], and Sections 1.3, 2.2, 2.3.2, 2.3.4 and 4 are essentially the content of [BR24]. Both of these papers are joint work with Ulrich Bauer. Sections 1.4, 2.1.1, 2.3.3, 2.4 and 5 and Appendix A are essentially the content of [Bau+23b], which is joint work with Ulrich Bauer, Michael Kerber, and Alexander Rolle. We want to point out that the original idea for Section 5.3 is due to Alexander Rolle, with the main contributions of the author of this thesis being Lemmas 5.12 and 5.13, as well as the proof strategy for Theorem 5.14. We also want to point out, that Theorem 2.5, Theorem E, and Appendix A appeared in an earlier version already in the author's master's thesis [Rol20], where in explicit form the blowup complex and the bar construction, as in Sections 1.4, 2.4 and 5.1, were also used. Note that Figs. 9 to 11 and 14 resemble [Rol20, Fig. 3.1]. In the master's thesis [Rol20] the category of covered spaces, the nerve as a functor, and the work of Dugger and Isaksen [DI04] were also used, but Section 5.5 is only contained in [Bau+23b]. While the master's thesis [Rol20] is mainly focused on explicit proofs in different contexts, in contrast, a large part of the publication [Bau+23b], and therefore the present thesis, makes systematic use of the unified framework of model categories. The results in the present thesis that have previously appeared in the author's master's thesis [Rol20] are not essential parts of the dissertation, and are included only for the sake of readability and completeness.

Finally, the author of this thesis has coauthored the following paper, which is not part of this dissertation.

[Mod+24] Tejasvinee Atul Mody, Alexander Rolle, Nico Stucki, Fabian Roll, Ulrich Bauer, and Kay Schneitz. "Topological analysis of 3D digital ovules identifies cellular patterns associated with ovule shape diversity". In: *Development* 151.20 (2024). DOI: 10.1242/dev.202590.

1.1. Summary of Main Contributions

We first summarize the main research contributions in this thesis.

As discussed in more detail in Section 1.2, we generalize a famous result of Rips on the contractibility of Vietoris–Rips complexes, originally applied in geometric group theory, to metric spaces that are not necessarily geodesic metric spaces. Our generalization is constructive and also compatible with the Vietoris–Rips filtration (Theorem A). In particular, our result can be applied directly to point clouds. Moreover, our generalization has strong effects (Theorem B) on the computation of persistence barcodes for tree-like metric data and provides an explanation for the previously observed efficiency of Ripser [Bau21], a leading software for the computation of Vietoris–Rips persistence barcodes, on viral evolution data. More concretely, we previously observed that sorting the point cloud in a specific way brings down the computation time from a full day to about two minutes, without changing a single line of code (compare Remark 1.8).

As discussed in more detail in Section 1.3, there are two efficient shape reconstruction methods that are of practical use in industry, whose connection, up to now, has been poorly understood. One is Morse-theoretic, namely the Wrap complex (see Section 2.3.2) introduced by Edelsbrunner [Ede03] as a subcomplex of the Delaunay complex, and the other is homological in nature, namely the construction of lexicographically minimal homologous cycles (Definition 2.26) also considered by Cohen–Steiner, Lieutier, and Vuillamy [CLV22] in a similar setting. In Section 4, we establish a strong connection between the Morse-theoretic and the homological approach. Our main result Theorem C shows that the lexicographically minimal homologous cycles in a Delaunay complex are all supported on the corresponding Wrap complex, as illustrated in Fig. 2. As a consequence of our main result, we obtain Corollary D that establishes a connection between the persistence barcode computation of the Delaunay filtration and the Wrap complexes.

This thesis also contains main contributions that are mostly expository. As discussed in more detail in Section 1.4, the nerve theorem is a fundamental result for topological data analysis. It guarantees, for example, that a union of closed balls can be replaced with the associated Čech complex while keeping all homotopy theoretic information, such as the number of holes (compare Fig. 1). The literature on the nerve theorem is extensive, but unfortunately, it is hard to navigate. Moreover, some proofs of the nerve theorem use an outdated set of tools, making it difficult for non-experts to grasp the core ideas. Finally, to be applicable in the context of persistent homology, one needs, for example, a nerve theorem for closed covers that is functorial in an appropriate sense. These difficulties are reflected in the fact that even in published textbooks, variants of the nerve theorem are referenced that, in a strictly technical sense, do not apply to the intended use cases. We address this issue in Section 5 by providing a variety of functorial nerve theorems (compare Table 1), hoping that our treatment of the material will be helpful, especially for students and newcomers to the applied topology community.

1.2. Contractions in Vietoris–Rips Complexes

A famous result of Rips shows the contractibility of Vietoris–Rips complexes of geodesic metric spaces above a scale parameter depending on the hyperbolicity of the space (Lemma 1.1). We consider the notion of geodesic defect to generalize this result in Section 3 to general metric spaces in a way that establishes simplicial collapses and is compatible with the filtration (Theorem A). Motivated by computational aspects of persistent homology, we further show that for finite tree metrics the Vietoris–Rips complexes collapse to their corresponding subforests (Theorem B). These collapses are induced by the apparent pairs gradient, which is commonly used as an algorithmic optimization in persistent homology computations. Our results provide an explanation for the previously observed efficiency of this optimization on tree-like metric data (compare Remark 1.8).

Background The Vietoris–Rips complex (Definition 2.7) is a fundamental construction in algebraic, geometric, and applied topology. First introduced by Vietoris [Vie27] in order to make homology applicable to general compact metric spaces, it has also found important applications in geometric group theory [Gro87] and topological data analysis [SC04]. The role of the parameter t in these three application areas is notably different. The homology theory defined by Vietoris arises in the limit $t \to 0$, by forming an appropriate algebraic limit of the homologies of the Vietoris–Rips complexes. In contrast, the key applications in geometric group theory rely on the fact that the Vietoris–Rips complex of a hyperbolic geodesic space is contractible for a sufficiently large parameter. This observation, originally due to Rips and first published in Gromov's seminal paper on hyperbolic groups [Gro87], is a fundamental result about the topology of Vietoris–Rips complexes and plays a central role in the theory of hyperbolic groups.

Lemma 1.1 (Contractibility Lemma; Rips, Gromov [Gro87])**.** *Let X be a metric space that is δ-hyperbolic and star-geodesic with respect to some point $p \in X$. Then the complex $|\operatorname{Rips}_t(X)|$ is contractible for every $t > 0$ with $t \geq 4\delta$.*

Here, a metric space X is star-geodesic as in Definition 2.8. Moreover, it is called *δ-hyperbolic with respect to $p \in X$* for $\delta \geq 0$ (in the sense of Gromov [Gro87]) if for any three points $x, y, z \in X$ we have

$$d(y,z) + d(x,p) \leq \max\{d(y,x) + d(z,p), d(y,p) + d(z,x)\} + 2\delta, \tag{1.1}$$

and it is called *δ-hyperbolic* if it is δ-hyperbolic with respect to every point.

Remark 1.2. We briefly remark on the use of the Contractibility Lemma in geometric group theory, as described in more detail in [Hul; BH99]. Let G be a finitely generated group and S a finite generating set that is symmetric ($S = S^{-1}$). The associated *Cayley graph* has vertices given by the group elements and an edge between g and $g \cdot s$ for $g \in G$ and $s \in S$. The Cayley graph forms a metric space when equipped with the shortest path metric. Moreover, the group G acts on this metric space by group multiplication from the left, implying that G also acts on the associated Vietoris–Rips complex to any parameter. The group G is a *hyperbolic group* if the associated metric space is

δ-hyperbolic for some $\delta \geq 0$. For a hyperbolic group G, the Contractibility Lemma then implies that G acts on a contractible simplicial complex, namely the Vietoris–Rips complex to a large enough parameter, which in turn implies that G must admit an Eilenberg–MacLane space $K(G, 1)$ that is a CW-complex with finitely many cells in each dimension. In particular, the hyperbolic group G is finitely presented.

Finally, in applications to topological data analysis, one is typically interested in the persistent homology (see Section 2.2) of the entire Vietoris–Rips filtration. A notable difference to the classical applications is that the metric spaces under consideration are typically finite, and in particular not geodesic. This motivates our interest in a meaningful generalization of the Contractibility Lemma to general metric spaces. For example, in the context of evolutionary biology persistent Vietoris–Rips homology has been successfully applied to identify recombinations and recurrent mutations [CCR13; LRR20; Ble+21]. The metrics arising as genetic distances of aligned RNA or DNA sequences are typically very similar to trees, capturing the phylogeny of the evolution. This motivates our interest in the particular case of tree metrics. These metric spaces are known to have acyclic Vietoris–Rips homology in degree greater than 0, and so any homology is an indication of some evolutionarily relevant phenomenon.

Contributions Based on the notion of a discretely geodesic space defined by Lang [Lan13], which is a natural setting for hyperbolic groups, and motivated by techniques used in that paper, we consider the following quantitative geometric property that also appears in [BS11, p. 271].

Definition 1.3. A metric space X is *ν-almost geodesic with respect to $p \in X$* if for all $x \in X$ and $r, s \geq 0$ with $r + s = d(x, p)$ there exists a point $z \in X$ with

$$d(z, x) \leq r + \nu \text{ and } d(z, p) \leq s + \nu.$$

A metric space is *ν-almost geodesic* if it is ν-almost geodesic with respect to every point.

If X is ν-almost geodesic (with respect to p), then it is also ν'-almost geodesic (with respect to p) for every $\nu' \geq \nu$. With this in mind, it is natural to consider the infimum over all ν such that X is ν-almost geodesic with respect to p, which we call the *geodesic defect of X with respect to p*, denoted by $\operatorname{geod}_p(X)$. Moreover, we consider the *geodesic defect* of X, denoted by $\operatorname{geod}(X)$, defined as the infimum over all ν such that X is ν-almost geodesic, or equivalently, as the supremum $\operatorname{geod}(X) = \sup_{p \in X} \operatorname{geod}_p(X)$.

Our first main result, which follows directly from Theorems 3.14 and 3.16, is a generalization of the Contractibility Lemma that also applies to non-geodesic metric spaces using the notion of geodesic defect, and further produces collapses that are compatible with the Vietoris–Rips filtration above the collapsibility threshold. Recall that we write $D_r(p) = \{y \in X \mid d(p, y) \leq r\}$ and $B_r(p) = \{y \in X \mid d(p, y) < r\}$ for the closed metric ball and the open metric ball of radius r centered at p, respectively.

Theorem A. *Let X be a metric space that is δ-hyperbolic and ν-almost geodesic with respect to some point $p \in X$. Then for every $s > 2\nu$ with $s \geq 4\delta + 2\nu$ there exists a discrete gradient on the full simplicial complex $\mathrm{Cl}(X)$ that induces, for every $t > u \geq s$, the collapses*

$$\mathrm{Rips}_t(X) \searrow \mathrm{Rips}_t^<(X) \searrow \mathrm{Rips}_u(X) \searrow \{p\},$$

and for every $r > l \geq 0$, the collapses

$$\mathrm{Rips}_u(D_r(p)) \searrow \mathrm{Rips}_u(B_r(p)) \searrow \mathrm{Rips}_u(D_l(p)).$$

Remark 1.4. If $\delta > 0$, then $4\delta + 2\nu > 2\nu$ and one can choose $s = 4\delta + 2\nu$, simplifying the statement in Theorem A. Moreover, if X is finite, for every $s > 0$ there exists a sufficiently small $\epsilon > 0$ with $\mathrm{Rips}_s(X) = \mathrm{Rips}_{s+\epsilon}(X)$, implying that the assumption $s > 2\nu$ can also be dropped in this case for $\delta = 0$.

Remark 1.5. Variants of the Contractibility Lemma with weaker assumptions on the geodesicity of the space can be found, for example, in [Gro87, Remark 1.7.D] and [BH99, Proposition 3.23], whose proof relies on the use of homotopy groups. Moreover, the Contractibility Lemma was studied using a version of Bestvina–Brady discrete Morse theory [Zar22, p. 1199], with the associated function defined on the barycentric subdivision of the Vietoris–Rips complex. More concretely, it is argued that when X is the vertex set of the Cayley graph of a finitely generated hyperbolic group, then for every $t > u > 4\delta + 1$ the inclusion $|\mathrm{Rips}_u(X)| \hookrightarrow |\mathrm{Rips}_t(X)|$ is a homotopy equivalence, which implies that the inclusion $|\mathrm{Rips}_u(X)| \hookrightarrow |\mathrm{Cl}(X)|$ into the full simplicial complex on X is a homotopy equivalence, and so $|\mathrm{Rips}_u(X)|$ is contractible.

Related results about implications of the geometry of a metric space on the homotopy types of the associated Vietoris–Rips complexes can be found, for example, in [ALS13; ALS19; Lat01].

Example 1.6. An important special case is given by a finite *tree metric space* X, where X is the vertex set of a positively weighted tree $T = (X, E)$, and where the edge weights are taken as lengths. The metric d is the associated path length metric, i.e., for two points $x, y \in X$ their distance is the infimum total weight of any path starting in x and ending in y. The geodesic defect with respect to any point is $\mathrm{geod}(X) = \frac{1}{2} \max_{e \in E} l(e)$, where $l(e)$ is the length of the edge e. Moreover, X is 0-hyperbolic. In fact, a metric space is 0-hyperbolic if and only if it can be embedded isometrically into an $\mathbb{R}$-*tree* (see [Eva08, Theorems 3.38 and 3.40]).

Our second main result, which follows directly from Proposition 3.38, is a strengthened version of Theorem A for the special case of tree metric spaces that connects the collapses of the Vietoris–Rips complexes to the construction of apparent pairs (see Section 2.3.1), which play an important role as a computational shortcut in the software Ripser [Bau21] (compare Section 2.2). This result depends on a particular ordering of the vertices: we say that a total order of X is *compatible* with the tree T if it extends the unique tree partial order resulting from choosing some arbitrary root vertex as the minimal element.

Theorem B. *Let X be a finite tree metric space for a weighted tree $T = (X, E)$, whose vertices are totally ordered in a compatible way. Then the apparent pairs gradient for the lexicographically refined Vietoris–Rips filtration on the full simplicial complex* $\mathrm{Cl}(X)$ *induces the collapses*

$$\mathrm{Rips}_t(X) \searrow \mathrm{Rips}_t^{<}(X) \searrow \mathrm{Rips}_u(X) \searrow T_u$$

for every $t > u > 0$ such that no tree edge $e \in E$ has length $l(e) \in (u, t]$, where T_u is the subforest with vertices X and all edges of E with length at most u. In particular, the persistent homology of the Vietoris–Rips filtration is trivial in degree greater than 0*.*

In the special case of trees with unit edge length, the proofs in [Ada13, Proposition 2.2] and [Ada+20, Proposition 3] are similar in spirit to our proof of Theorem 3.31, which is based on discrete Morse theory.

Remark 1.7. For any metric space X and any point $p \in X$ one can choose a total order on X such that $x < y$ implies $d(x, p) \leq d(y, p)$. This total order is used in the proof of Theorems 3.14 and 3.16, and also, for example, by Kahle [Kah11] to study random Vietoris–Rips complexes in the supercritical regime. In the special case that X is a finite tree metric space, as in Example 1.6, this total order is compatible with the tree, as in Theorem B.

Remark 1.8. Given a vertex order $\leq$, the lexicographic order on simplices for the reverse vertex order $\geq$ coincides with the reverse colexicographic order for the original order $\leq$, which is used for computations in Ripser. As a consequence, when the input is a tree metric with the points ordered in reverse order of the distances to some arbitrarily chosen root, then Ripser will identify all non-tree simplices in apparent pairs, requiring not a single column operation to compute its trivial persistent homology. In practice, we observe that on data that is almost tree-like, such as genetic evolution distances, Ripser performs exceptionally well. The results of this section provide a partial geometric explanation for this behavior and yield a heuristic for preprocessing tree-like data by sorting the points to speed up the computation in such cases. In the application to the study of SARS-CoV-2 described in [Ble+21], ordering the genome sequences in reverse chronological order, as an approximation of the reverse tree order for the phylogenetic tree, lead to a huge performance improvement, bringing down the computation time for the persistence barcode from a full day to about 2 minutes.

Remark 1.9. In Section 3.4.1 we prove Theorem B in the special case when the tree metric is *generic*, meaning that the pairwise distances are distinct, by using the fact that the diameter function $\mathrm{diam}\colon \mathrm{Cl}(X) \to \mathbb{R}$ is a generalized discrete Morse function, defined on the full simplicial complex on X. The discrete Morse theory of other geometric complexes used in topological data analysis was studied, for example, by Bauer and Edelsbrunner [BE17], using the fact that for a finite subset of $\mathbb{R}^d$ in general position the Čech and Delaunay radius functions are generalized discrete Morse functions (Proposition 2.29).

1.3. Shape Reconstruction Methods

In Section 4 we study the connection between discrete Morse theory and persistent homology in the context of shape reconstruction methods. More specifically, we consider the construction of Wrap complexes (see Section 2.3.2), introduced by Edelsbrunner [Ede03] as a subcomplex of the Delaunay complex, and the construction of lexicographic optimal homologous cycles (Definition 2.26), also considered by Cohen–Steiner, Lieutier, and Vuillamy [CLV22] in a similar setting. We show that for any cycle in a Delaunay complex at a given radius parameter, the corresponding lexicographically optimal homologous cycle is supported on the Wrap complex to the same parameter (Theorem C), which establishes a close connection between the two methods. We obtain this result by establishing a fundamental connection between reductions of cycles in the computation of persistent homology and gradient flows in algebraic Morse theory.

Background Reconstructing shapes and submanifolds from point clouds is a classical topic in computational geometry. Starting in the 2000s, several key results have been achieved [Ede95; ACK01; Ame+02; Ede03; Dey07; CDR05; RS07; BG14; NSW08], leading to a method based on the Delaunay triangulation [CDR05] with theoretical homeomorphic reconstruction guarantees. The method is theoretical in nature, and there are some complexity and robustness issues that call into question its practical applicability. A major challenge is caused by *slivers* [Che+00], which are, in three dimensional Euclidean space, tetrahedra in a Delaunay triangulation with small volume but no short edges, and which have to be handled explicitly. In contrast, several related Delaunay-based methods have proven to be highly robust and successful in practice, in particular, Morse-theory based methods such as *Wrap* and related constructions [Ede03; Dey+05; RS07; Sad09; BE17; Por+22], and homological methods based on minimal cycles [CLV19; CLV22; CLV23; Vui21; AL22]. In particular, in three dimensional Euclidean space, these homological methods gracefully circumvent the issue of slivers, simply because 2-chains do not contain any tetrahedra. On the other hand, the Wrap complex (see Section 2.3.2) is always homotopy equivalent to a union of closed balls of a given radius (see Remark 2.38 and Section 1.4), but might contain critical sliver simplices. It is therefore desirable to identify a subcomplex that inherits some good properties of the Wrap complex and that is free from slivers.

Even though the Morse-theoretic and the homological method seem closely related in spirit, up to now, the connection between them has been poorly understood. While the development of the Wrap complex predated the introduction of Forman's discrete Morse theory [For98], it has subsequently been rephrased using this framework [BE17], thus making it possible to analyze the Wrap complex with a new set of tools. A connection to homology-based methods promises to further increase our understanding of the geometric and topological properties of both the Wrap complex and these algebraic constructions. Indeed, a synthesis of Wrap, discrete Morse theory, and persistent homology for surface reconstruction has been envisioned already in [Ede03].

Before discussing our contributions in detail, we remark that there are other methods to reconstruct a shape from a sample that are not directly based on the Delaunay

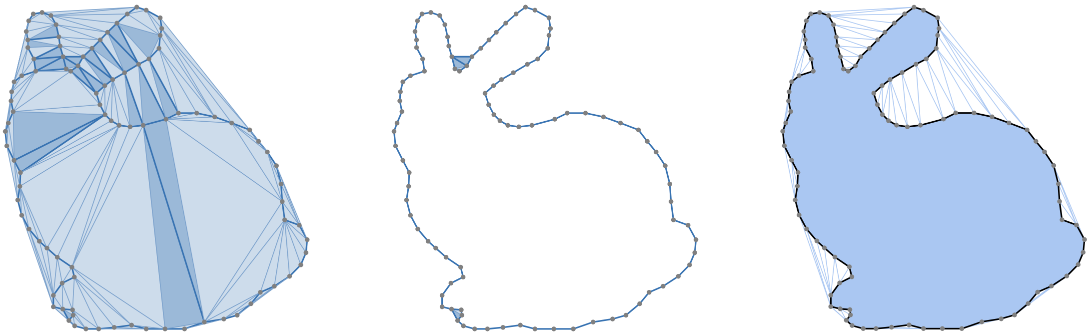

Figure 2: Left: Delaunay triangulation of a point cloud, with critical simplices highlighted. Middle: Wrap complex for a small radius parameter. Right: lexicographically minimal cycle for the most persistent feature (black contour), shown together with its bounding chain (shaded blue).

triangulation. For instance, the homotopy type of a shape can be reconstructed from the offset filtration of a sample [Att+23; NSW08], and therefore also by using Čech complexes. Similarly, the homotopy type of a shape can be reconstructed using variants of the Vietoris–Rips complex of a sample [LV23; Lat01]. Moreover, the homology of a shape can be inferred using methods based on persistent homology [CEH07; BCY18].

Contributions We establish a strong connection between the Morse-theoretic and the homological approaches mentioned above. Our main result, which follows from the more general Theorem 4.25 when applied to the Delaunay radius function (see Definition 2.3 and Proposition 2.29), establishes that the lexicographically minimal homologous cycles (Definition 2.26) in a Delaunay complex are all supported on the corresponding Wrap complex (see Section 2.3.2), as illustrated in Fig. 2. The lexicographic order on simplices is as in Section 2.3.1, and consider homology with coefficients in some field $\mathbb{F}$.

Theorem C. *Let $X \subset \mathbb{R}^d$ be a finite subset in general position, let $r \in \mathbb{R}$, and let $h \in H_*(\mathrm{Del}_r(X))$ be a homology class of the Delaunay complex $\mathrm{Del}_r(X)$. Then the lexicographically minimal cycle of h, with respect to the Delaunay-lexicographic order on the simplices, is supported on the Wrap complex $\mathrm{Wrap}_r(X)$.*

As a consequence of our main result, we obtain the following connection between the Wrap complex and the persistent homology (see Section 2.2) of the Delaunay filtration, by applying the more general Corollary 4.26 to the Delaunay radius function.

Corollary D. *Let $X \subset \mathbb{R}^d$ be a finite subset in general position and (σ, τ) a non-zero persistence pair of the lexicographically refined Delaunay filtration. Let $r = r_X(\sigma)$ and $s = r_X(\tau)$ be the radius of the smallest empty circumsphere of σ and of τ, respectively. Then the lexicographically minimal cycle of $[\partial\tau]$ in the Delaunay complex $\mathrm{Del}_s^<(X)$, given as the column R_τ of the totally reduced filtration boundary matrix, is supported on the Wrap complex $\mathrm{Wrap}_r(X)$.*

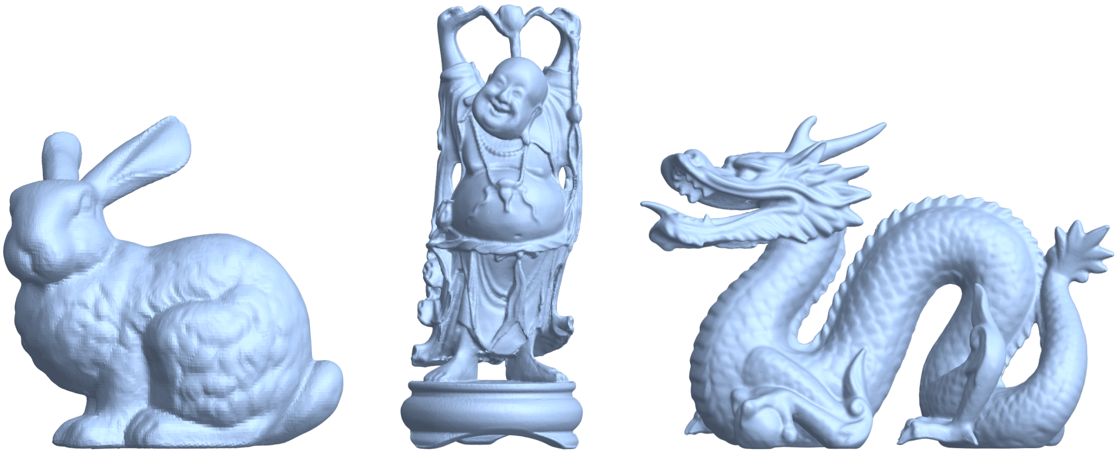

Figure 3: The lexicographically minimal cycle corresponding to the most persistent feature of the Delaunay filtration for three 3D scan point clouds [Sta] yields an accurate reconstruction of the surface.

The totally reduced filtration boundary matrix can be computed using Algorithm 2. For a sufficiently good sample of a compact d-submanifold of Euclidean space, the union of closed balls centered at the sample points deformation retracts onto the submanifold by a closest point projection [Att+23; NSW08]. As the Delaunay complex is naturally homotopy equivalent to the union of closed balls (see Section 1.4), this implies that the fundamental class of the manifold is captured in the d-dimensional persistent homology of the Delaunay filtration. Together with Corollary D, this suggests the following heuristic: Take the most persistent d-dimensional feature of the Delaunay filtration, i.e., the interval in the barcode with the largest death/birth ratio. Intuitively, this feature is born at a small scale and only gets filled in at a large scale. By Corollary D, the corresponding lexicographically minimal cycle is guaranteed to be supported on the Wrap complex for a small scale parameter. See Fig. 3 for an illustration, which can readily be reproduced using the code provided in [Rol23] by executing the following command on any machine with Docker installed and configured with sufficient memory (16GB recommended):

```
$ docker build -o output github.com/fabian-roll/wrappingcycles
```

We remark that Theorem C and Corollary D do not depend on a specific choice of coefficient field. In particular, when taking orientations of triangles into account, which are canonically oriented, as in Remark 2.22, according to a given total vertex order, Fig. 4 suggests that it is favorable in practice to use, for example, $\mathbb{F}_3$ coefficients to obtain a consistent orientation of the triangles in the cycle. It remains an interesting open problem to find suitable assumptions under which the lexicographically minimal cycle corresponding to a persistent feature of the Delaunay filtration can be guaranteed to reconstruct the sampled shape in a geometrically and topologically faithful way.

It is worth noting that the lexicographically optimal cycles considered in [CLV22; CLV19; CLV23; Vui21] are based on a slightly different total order on simplices, which

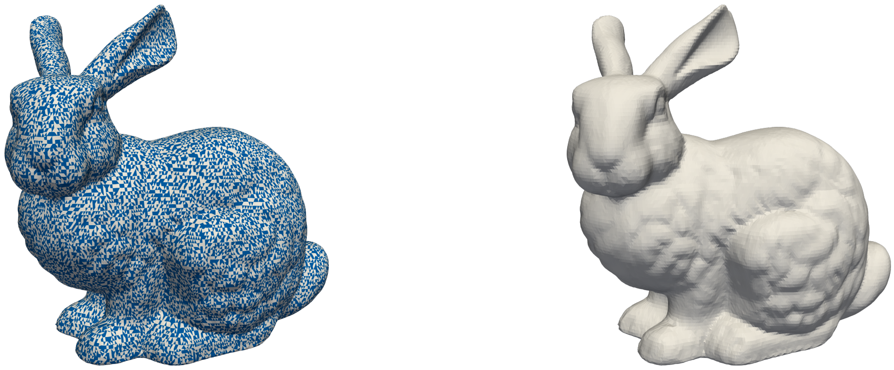

Figure 4: The lexicographically minimal cycle corresponding to the most persistent feature of the Delaunay filtration for the 3D scan point cloud of the Stanford bunny [Sta] with $\mathbb{F}_2$ coefficients (left) and $\mathbb{F}_3$ coefficients (right). Backfaces of triangles, oriented according to a given total vertex order, are colored in blue.

induces a simplexwise refinement of the sublevel set filtration of the Čech radius function (Definition 2.3). We now describe this total order in slightly more detail, following the exposition in [Vui21, Section 4.2.3]. The idea for this total order is based on a variational formulation of the Delaunay triangulation of the convex hull of a finite subset $P \subseteq \mathbb{R}^d$. To every simplex $\sigma = \{p_1, \dots, p_k\}$ in the full simplicial complex on P one can assign the weight

$$w_p(\sigma) = (\int_{|\sigma|} f_\sigma(x)^p dx)^{\frac{1}{p}}$$

given by the p-norm of the function $f_\sigma(x) = \sum_i \lambda_i(x) \|p_i\|_2^2 - \|x\|_2^2$, where the $\lambda_i(x)$ are the barycentric coordinates of x with respect to σ. Then, the Delaunay triangulation of the convex hull of P is the unique triangulation that minimizes the p-norm of the vector of weights of the simplices in the triangulation. Now, for two simplices σ, τ in the full simplicial complex on P one defines $\sigma \leq \tau$ if and only if for large enough p we have $w_p(\sigma) \leq w_p(\tau)$. Under some further assumptions, this induces a total order on simplices, and the Delaunay triangulation of the convex hull of P can be expressed in terms of a lexicographically optimal chain with respect to this total order.

Relating this particular choice of the total order on simplices to the results presented in this thesis, as well as to the results in [BE17], remains an interesting open problem.

1.4. Nerve Theorems in Applied Topology

The nerve theorem is a basic result of algebraic topology that plays a central role in computational and applied aspects of the subject. To be applicable to persistent homology (see Section 2.2), one needs a nerve theorem that is functorial in an appropriate sense, and furthermore, one often needs a nerve theorem for closed covers as well as for open covers. While the techniques for proving such functorial nerve theorems have long been available, there is, unfortunately, no general-purpose, explicit treatment of this topic in the literature. We address this by proving a variety of functorial nerve theorems in Section 5; also compare Table 1. In particular, we establish a “unified” nerve theorem (Theorem H) that subsumes many of the variants, using standard techniques from abstract homotopy theory.

Background If $\mathcal{U} = (U_i)_{i \in I}$ is a cover of a topological space X, then the *nerve* of $\mathcal{U}$, which dates back to Alexandroff [Ale28], is the simplicial complex $\mathrm{Nrv}(\mathcal{U})$ whose simplices are the finite subsets $J \subseteq I$ such that the intersection $\cap_{i \in J} U_i$ is non-empty. The nerve of a cover played an important role in the development of homology and cohomology theory. In particular, Čech (co)homology is given by the (co)limit of the (co)homology groups of the nerves of a directed system of open covers ordered by refinement. A historical exposition can be found in [EH80, Chapter 2].

The *nerve theorem*, whose early versions are due to Leray [Ler45], Borsuk [Bor48], and Weil [Wei52], is a basic result in algebraic and combinatorial topology. Roughly speaking, it says that if every non-empty finite intersection of cover elements is contractible, then, subject to some further tameness conditions on X and $\mathcal{U}$, the space X is homotopy equivalent to the nerve of $\mathcal{U}$.

Nowadays, the nerve theorem and the aspect of functoriality play a crucial role in topological data analysis. Nerves are the main way to replace a topological space, determined by the data points using geometric constructions, with a combinatorial model that is suitable for computations (see Section 2.1 and Fig. 1). Two prominent examples are the Čech complex and the Delaunay complex, which arise as nerves of a collection of closed balls and closed Voronoi balls, respectively. Another important example is the Vietoris–Rips complex, which is not usually defined as the nerve of a cover, though it is isomorphic to a nerve (Proposition 2.16). Note that, while one can choose whether to use open or closed sets when defining the Čech and Vietoris–Rips complexes, the only standard way to define the Delaunay complex uses closed sets.

These examples are typical, in that the topological spaces determined by data points usually depend on one or more parameters, leading to filtrations of topological spaces and covers. Now functoriality ensures that the corresponding nerves form a filtration as well. For example, if $X \subset \mathbb{R}^d$ is a finite set of points, the offset filtration O is the filtration of $\mathbb{R}^d$ with $O_r = \cup_{x \in X} D_r(x)$ for $r > 0$, where $D_r(x)$ is the closed ball about x of radius r. The nerve of the cover $\mathcal{U}_r = (D_r(x))_{x \in X}$ is the Čech complex (see Section 2.1.1 and Fig. 1), and as r varies, these complexes form a filtration as well. In this case, the nerve theorem says that O_r is homotopy equivalent to the nerve of $\mathcal{U}_r$.

Going further, one wants a nerve theorem to provide homotopy equivalences that are compatible with the inclusion maps in these two filtrations. This is necessary in particular if one is interested in persistent homology (see Section 2.2), which is an algebraic invariant of filtrations that encodes the homology of each step of the filtration, as well as the maps in homology induced by each inclusion. There are several ways in which the homotopy equivalences provided by a nerve theorem might be compatible with the inclusions in these two filtrations, as we will now explain.

In order to prove that the persistent homology of the offset filtration is isomorphic to the persistent homology of the associated Čech complex filtration, it suffices to have isomorphisms $H_n(O_r) \cong H_n(\mathrm{Nrv}(\mathcal{U}_r))$ such that all the squares of the following form commute:

$$\begin{array}{ccc} H_n(O_r) & \xrightarrow{\cong} & H_n(\mathrm{Nrv}(\mathcal{U}_r)) \\ \downarrow & & \downarrow \\ H_n(O_{r'}) & \xrightarrow[\cong]{} & H_n(\mathrm{Nrv}(\mathcal{U}_{r'})) \end{array} \tag{1.2}$$

By Theorem F below, such isomorphisms can be constructed from the induced maps of homotopy equivalences $|\,\mathrm{Nrv}(\mathcal{U}_r)| \to O_r$ between the nerves and the offsets such that all squares of the following form commute:

$$\begin{array}{ccc} O_r & \xleftarrow{\simeq} & |\,\mathrm{Nrv}(\mathcal{U}_r)| \\ \big\downarrow & & \big\downarrow \\ O_{r'} & \xleftarrow[\simeq]{} & |\,\mathrm{Nrv}(\mathcal{U}_{r'})| \end{array} \tag{1.3}$$

The construction of these compatible homotopy equivalences relies on the fact that the cover elements of the offset filtration are convex and that the inclusions $O_r \hookrightarrow O_{r'}$ are affine linear.

For a more general filtration $(X_r, \mathcal{A}_r)$, with $X_r = \bigcup \mathcal{A}_r$, a similar strategy does not necessarily produce commuting diagrams as in 1.3. However, if one is only interested in the filtration after applying homology or some other homotopy-invariant functor, then it suffices to have homotopy equivalences $X_r \to |\,\mathrm{Nrv}(\mathcal{A}_r)|$ such that all squares of the following form commute *up to homotopy:*

$$\begin{array}{ccc} X_r & \xrightarrow{\simeq} & |\,\mathrm{Nrv}(\mathcal{A}_r)| \\ \big\downarrow & \overset{H}{\nearrow\!\!\!\nearrow} & \big\downarrow \\ X_{r'} & \xrightarrow[\simeq]{} & |\,\mathrm{Nrv}(\mathcal{A}_{r'})| \end{array} \tag{1.4}$$

In the diagram, H is a homotopy from the bottom route around the square to the top route. Nerve theorems with this structure are often used in the study of persistent homology (for references, see the end of this subsection).

However, in some homotopy-theoretic approaches to topological data analysis, we need a nerve theorem that is compatible with the inclusions $X_r \hookrightarrow X_{r'}$ on the nose, and not just up to homotopy. In this section, we will prove nerve theorems that provide strictly commuting diagrams, at the cost of introducing an intermediary between the covered space and the nerve: we obtain a filtration Z_r and homotopy equivalences $Z_r \to X_r$ and $Z_r \to |\,\mathrm{Nrv}(\mathcal{A}_r)|$ such that all the diagrams of the following form commute:

$$\begin{array}{ccccc} X_r & \xleftarrow{\simeq} & Z_r & \xrightarrow{\simeq} & |\,\mathrm{Nrv}(\mathcal{A}_r)| \\ \big\downarrow & & \big\downarrow & & \big\downarrow \\ X_{r'} & \xleftarrow[\simeq]{} & Z_{r'} & \xrightarrow[\simeq]{} & |\,\mathrm{Nrv}(\mathcal{A}_{r'})| \end{array} \tag{1.5}$$

While one can avoid introducing intermediate objects in the special case of the offset filtration, this is not possible in general, as we explain below. Diagrams of the form

1.5 appear classically in the study of homotopy categories [GZ67]. More recently, they appear in Blumberg and Lesnick's work on the *homotopy interleaving distance* [BL17], a distance on diagrams of spaces that is universal among stable and homotopy-invariant distances. The idea is to define an equivalence relation on filtered spaces such that F and F' are related if they can be connected via an intermediate filtration, as above, with the horizontal arrows weak homotopy equivalences. Then, filtered spaces F_1 and F_2 are homotopy interleaved if F_1 is related to some F_1', F_2 is related to some F_2', and F_1' and F_2' are interleaved. An important motivation for nerve theorems that provide diagrams of the form 1.5 is that they can be used in frameworks like the one of Blumberg–Lesnick.

In this thesis, we establish a variety of nerve theorems that provide strictly commuting diagrams relating spaces and nerves. These nerve theorems are summarized in Table 1. Before we introduce the contents of this section in more detail, we highlight some aspects of our treatment of the material. The blowup complex is often used as an intermediate object for proving nerve theorems (serving as the space Z_r in Diagram 1.5). This space is closely related to the bar construction from abstract homotopy theory. We will discuss this fact in detail, explaining why it is advantageous to state nerve theorems using the blowup complex rather than the bar construction, but that one can still use the bar construction for the proofs. We prove a functorial nerve theorem for subsets of Euclidean space covered by closed, convex subsets using the blowup complex, and we also introduce a notion of pointed covers, which allows us to prove such a functorial nerve theorem that does not require an intermediate object at all. We consider simplicial complexes covered by subcomplexes and explain how one can use a bar construction in the category of posets to prove a functorial nerve theorem in this setting. Finally, after using standard model category arguments to prove more general nerve theorems, we give a series of examples demonstrating that most of the assumptions in these theorems are necessary.

Functorial Nerve Theorems In order to say precisely what we mean by a functorial nerve theorem, we need to explain how the nerve can be viewed as a functor. To this end, we will define the category of covered spaces.

To motivate our definition, we first briefly discuss a variant that is also common in the literature (see, e.g., [Bar02]). The objects in this category are of the form $(X, \mathcal{U})$, where X is a topological space and $\mathcal{U}$ is an unindexed cover of X. A map between covered spaces $f\colon (X, \mathcal{U}) \to (Y, \mathcal{V})$ is then given by a continuous map $f\colon X \to Y$ such that for any cover element $U \in \mathcal{U}$ there exists $V \in \mathcal{V}$ with $f(U) \subseteq V$. Choosing such a cover element $V \in \mathcal{V}$ for every element $U \in \mathcal{U}$ determines a simplicial map $\mathrm{Nrv}(\mathcal{U}) \to \mathrm{Nrv}(\mathcal{V})$ between the nerves. In general, different choices give different simplicial maps, but it will always be unique up to contiguity (see [Mun84, p. 67] for a definition). In particular, it follows that any two choices determine, up to homotopy, the same map on the geometric realization.

To avoid having to make choices, we work with indexed covers and record the choice of cover elements as above in a map between the indexing sets. This way, we circumvent the ambiguity of the induced map between the nerves up to homotopy.

Reference	O/C	X	$\|\mathcal{A}\|$	$\mathcal{A}$	Equivalence	Intermediate
Thm. 5.4	closed	$X \subset \mathbb{R}^d$	finite	convex	hom. eq.	Blowup($\mathcal{A}$)
Thm. 5.6	closed	$X \subset \mathbb{R}^d$	finite	pointed, convex	hom. eq.	none
Thm. 5.14	closed	simplicial complex	none	good cover by subcom-plexes	hom. eq.	Blowup($\mathcal{A}$)
Thm. 5.16	closed	compact, semi-algebraic	finite	good, semi-algebraic	hom. eq.	Blowup($\mathcal{A}$)
Thm. 5.20	closed	simplicial complex	loc. finite	$(k-t+1)$-good cover by subcom-plexes	k-connected	Blowup($\mathcal{A}$)
Thm. 5.25	open	none	none	good, numerable	hom. eq.	Blowup($\mathcal{A}$)
	open	none	none	weakly good	weak hom. eq.	Blowup($\mathcal{A}$)
	open	none	none	CG, hom'gy good	hom'gy iso.	Blowup($\mathcal{A}$)
	closed	CG	loc. finite, loc. finite-dim.	L-condition, good	hom. eq.	Blowup($\mathcal{A}$)
	closed	CG	loc. finite, loc. finite-dim.	L-condition, weakly good	weak hom. eq.	Blowup($\mathcal{A}$)
	closed	CG	loc. finite, loc. finite-dim.	L-condition, hom'gy good	hom'gy iso.	Blowup($\mathcal{A}$)

Table 1: A summary of the functorial nerve theorems in this thesis, for a cover $\mathcal{A}$ of a space X. Columns 2–5 summarize the assumptions: whether the cover elements are open or closed, assumptions on X, assumptions on the cardinality of $\mathcal{A}$, and additional assumptions on $\mathcal{A}$. Columns 6–7 summarize the conclusions: the type of equivalence established, and the intermediate object, if any. We use the abbreviations: homotopy equivalence (hom. eq.), homology isomorphism (hom'gy iso.), compactly generated (CG), homologically good (hom'gy good), Latching space condition (L-condition), locally finite (loc. finite), locally finite-dimensional (loc. finite-dim.).

Definition 1.10. Let X and Y be topological spaces, $(U_i)_{i\in I}$ a cover of X, and $(V_\ell)_{\ell\in L}$ a cover of Y. A *map of indexed covers* $\varphi\colon (U_i)_{i\in I} \to (V_\ell)_{\ell\in L}$ is specified formally by a map $\varphi\colon I \to L$ between the indexing sets, which we denote with the same symbol. We say that a continuous map $f\colon X \to Y$ is *carried by* φ if for all $i \in I$ we have $f(U_i) \subseteq V_{\varphi(i)}$.

If f is carried by φ and g is carried by ψ, then $g\circ f$ is carried by $\psi\circ\varphi$ if the compositions are defined. Hence, we get the following category.

Definition 1.11. The objects of the *category of covered spaces* Cov are pairs of the form $(X, (U_i)_{i\in I})$, where X is a topological space and $(U_i)_{i\in I}$ is a cover of X. A *morphism of covered spaces* $(f, \varphi)\colon (X, (U_i)_{i\in I}) \to (Y, (V_\ell)_{\ell\in L})$ consists of a continuous map $f\colon X \to Y$ and a map $\varphi\colon I \to L$ such that f is carried by the corresponding map of indexed covers $\varphi\colon (U_i)_{i\in I} \to (V_\ell)_{\ell\in L}$.

With this category in hand, we can define a functor $\mathrm{Spc}\colon \mathsf{Cov} \to \mathsf{Top}$ by forgetting the cover: Spc takes a pair $(X, (U_i)_{i\in I})$ to X. By taking the geometric realization of the nerve of a cover, we obtain another such functor. We denote by Simp the category of simplicial complexes.

Definition 1.12. Let X be a topological space, and let $\mathcal{U} = (U_i)_{i\in I}$ be a cover of X. For any $J \subseteq I$, we write $U_J = \bigcap_{i\in J} U_i$. The *nerve* of $\mathcal{U}$ is the simplicial complex $\mathrm{Nrv}(\mathcal{U})$ with simplices

$$\{J \subseteq I \mid |J| < \infty \text{ and } U_J \neq \emptyset\} .$$

A morphism of covered spaces $(f, \varphi)\colon (X, \mathcal{U}) \to (Y, \mathcal{V})$ induces a simplicial map between the nerves of the covers $\varphi_*\colon \mathrm{Nrv}(\mathcal{U}) \to \mathrm{Nrv}(\mathcal{V})$. Thus, the nerve can be seen to be a functor $\mathrm{Nrv}\colon \mathsf{Cov} \to \mathsf{Simp}$. By postcomposing this with the geometric realization functor $|\cdot|\colon \mathsf{Simp} \to \mathsf{Top}$, we get the functor $|\mathrm{Nrv}|\colon \mathsf{Cov} \to \mathsf{Top}$ that takes a pair $(X, \mathcal{U})$ to the geometric realization $|\mathrm{Nrv}(\mathcal{U})|$.

Remark 1.13. There is a variant of this definition, where the vertex set is the set of cover elements, in contrast to our definition, where it is the indexing set. While these definitions yield different simplicial complexes in general, as the same subset can appear multiple times in the indexed cover, they are always homotopy equivalent. More precisely, if $\mathcal{U} = (U_i)_{i\in I}$ is an indexed cover and $U_j \subseteq U_l$, with $j \neq l$, are cover elements, then the inclusion $|\mathrm{Nrv}(\mathcal{V})| \hookrightarrow |\mathrm{Nrv}(\mathcal{U})|$ is a homotopy equivalence, where $\mathcal{V} = (U_i)_{i\in I\setminus\{j\}}$: The link $\mathrm{Lk}(j) = \{\sigma \in \mathrm{Nrv}(\mathcal{U}) \mid j \notin \sigma,\ \sigma \cup \{j\} \in \mathrm{Nrv}(\mathcal{U})\}$ of the vertex j in $\mathrm{Nrv}(\mathcal{U})$ is a cone with apex l, i.e., for all $\sigma \in \mathrm{Lk}(j)$ we have $\sigma \cup \{l\} \in \mathrm{Nrv}(\mathcal{U})$. Therefore, there exists a collapse $\mathrm{Nrv}(\mathcal{U}) \searrow \mathrm{Nrv}(\mathcal{V})$ (see Section 2.3 for a definition; see also Example 2.31).

Now that we can understand the covered space and the nerve as functors, we can consider natural transformations that relate them. In general, if F_1 and F_2 are functors from some category $\mathcal{C}$ to Top, and $\sigma\colon F_1 \Rightarrow F_2$ is a natural transformation, one says that σ is a pointwise homotopy equivalence if the component $\sigma_C\colon F_1(C) \to F_2(C)$ is a homotopy equivalence for all objects C of $\mathcal{C}$. Similarly one can consider pointwise weak homotopy equivalences, pointwise homology isomorphisms, et cetera. This section is about nerve theorems that relate the covered space and the nerve through pointwise equivalences.

Most of these nerve theorems make use of a standard construction that is called the *blowup complex* by Zomorodian–Carlsson [ZC08], but goes back at least to Segal [Seg68]. It is a functor Blowup: $\mathsf{Cov} \to \mathsf{Top}$, along with natural transformations $\rho_S\colon \text{Blowup} \Rightarrow \text{Spc}$ and $\rho_N\colon \text{Blowup} \Rightarrow |\text{Nrv}|$. In particular, for any morphism of covered spaces $(f,\varphi)\colon (X,\mathcal{U}) \to (Y,\mathcal{V})$ there exists a commuting diagram of the following form:

$$\begin{array}{ccccc} X & \xleftarrow{\rho_S} & \text{Blowup}(\mathcal{U}) & \xrightarrow{\rho_N} & |\text{Nrv}(\mathcal{U})| \\ {\scriptstyle f}\downarrow & & \downarrow & & \downarrow{\scriptstyle |\varphi_*|} \\ Y & \xleftarrow[\rho_S]{} & \text{Blowup}(\mathcal{V}) & \xrightarrow[\rho_N]{} & |\text{Nrv}(\mathcal{V})| \end{array} \tag{1.6}$$

We write $[n]$ for the set $\{0,\ldots,n\}$. If $\mathcal{U} = (U_i)_{i\in[n]}$ is a finite cover of a space X, then the blowup complex is

$$\text{Blowup}(\mathcal{U}) = \bigcup_{J\in\text{Nrv}(\mathcal{U})} U_J \times \Delta^J \subseteq X \times \Delta^n ,$$

where Δ^n is the standard topological n-simplex and Δ^J is a face of Δ^n determined by the inclusion $J \subseteq [n]$. The idea is that each piece of X expands according to the number of cover elements that contain it. See Section 5.1 for the definition for arbitrary covers.

To begin, we give three functorial nerve theorems (Theorems E, F and G) whose proofs are relatively elementary and use techniques that are interesting in their own right. In Section 5.2, we prove the following functorial nerve theorem, which appeared in the author's master's thesis [Rol20] in a preliminary form:

Theorem E (Theorem 5.4)**.** *If $X \subset \mathbb{R}^d$, and $\mathcal{B} = (C_i)_{i\in[n]}$ is a cover by closed convex subsets, then the natural maps $\rho_S\colon \text{Blowup}(\mathcal{B}) \to X$ and $\rho_N\colon \text{Blowup}(\mathcal{B}) \to |\text{Nrv}(\mathcal{B})|$ are homotopy equivalences.*

The proof uses partitions of unity, and is similar to the strategy for open covers in Hatcher's textbook. In Section 5.2, we also prove a functorial nerve theorem for closed convex covers that does not require any intermediate object, subject to an additional assumption on the morphisms of covered spaces. Before we state this theorem, we elaborate shortly on why such a functorial nerve theorem cannot exist in general. The reason is simple: there are no natural transformations between Spc and $|\text{Nrv}|$ in either direction. Consider the covered spaces $(*,(*))$, where $*$ is the one-point space, and $(Y,(Y))$, where $Y \neq *$ is any space. For any point $p \in Y$ the inclusion $\iota_p\colon * \hookrightarrow Y$ gives rise to a morphism of covered space $(\iota_p, * \mapsto Y)$. If there existed a natural transformation $|\text{Nrv}| \Rightarrow \text{Spc}$, then this would already fix a single inclusion $\iota_q\colon * \hookrightarrow Y$ as part of such a morphism of covered spaces, implying that $Y = \{q\}$ is a single point, yielding a contradiction. Similarly, consider any covered space $(Z,(U,V))$ with $p \in U \cap V$ any point. Consider the two morphisms of covered spaces $(\iota_p, * \mapsto U), (\iota_p, * \mapsto V)\colon (*,(*)) \to (Z,(U,V))$. Then, these maps induce different simplicial maps on the nerves, implying that there exists no natural transformation $\text{Spc} \Rightarrow |\text{Nrv}|$.

Thus, in order to obtain a functorial nerve theorem that does not need an intermediate object, the map of indexed covers needs to have strong combinatorial control on the continuous map. To this end, we introduce the following notions.

Definition 1.14. A *pointed cover* $\mathcal{U}_* = (\mathcal{U} = (U_i)_{i\in I}, (u_\sigma)_{\sigma\in\mathrm{Nrv}(\mathcal{U})})$ of a topological space X is a cover $\mathcal{U}$ of X together with a point $u_\sigma \in U_\sigma$ for every $\sigma \in \mathrm{Nrv}(\mathcal{U})$.

The *category of pointed covered spaces* Cov_* has objects tuples of the form $(X, \mathcal{A}_*)$, where X is a topological space and $\mathcal{A}_* = (\mathcal{A}, (a_\sigma)_{\sigma\in\mathrm{Nrv}\,\mathcal{A}})$ is a pointed cover of X. A morphism $(f, \varphi)\colon (X, \mathcal{A}_*) \to (Y, \mathcal{B}_*)$ of pointed covered spaces is a morphism of covered spaces $(f, \varphi)\colon (X, \mathcal{A}) \to (Y, \mathcal{B})$ that respects the basepoints, i.e., such that for any $\sigma \in \mathrm{Nrv}(\mathcal{A})$ we have $f(a_\sigma) = b_{\varphi_*(\sigma)}$.

There is an obvious functor $\mathsf{Cov}_* \to \mathsf{Cov}$ that forgets the pointing, and hence we get, as for the category of covered space, the functors $\mathrm{Spc}\colon \mathsf{Cov}_* \to \mathsf{Top}$ and $|\,\mathrm{Sd}\,\mathrm{Nrv}\,|\colon \mathsf{Cov}_* \to \mathsf{Top}$, where $\mathrm{Sd}\,\mathrm{Nrv}$ is the subdivision of the nerve.

Now, we will describe a functorial nerve theorem that does not require an intermediate object. The subcategory ClConv_* of Cov_* consists of subsets of $\mathbb{R}^d$ that are covered by finitely many closed convex sets. Further, we restrict to morphisms of pointed covered spaces whose underlying continuous maps are affine linear on each cover element. Many covers of interest in topological data analysis are pointed.

Example 1.15. Let $\{x_0, \ldots, x_n\} \subseteq \mathbb{R}^d$ be a finite set of points. Then, we can point the cover $\mathcal{U}_r = (D_r(x_i))_{i\in[n]}$ of the union of closed balls $O_r = \bigcup_{i=0}^n D_r(x_i)$ in the following way: For each non-empty subset $\sigma \subseteq [n]$ there exists a smallest real number r_σ such that the intersection $(\mathcal{U}_{r_\sigma})_\sigma$ is non-empty. We define the point p_σ to be the unique point in this intersection. This gives the pointed cover $(\mathcal{U}_r, (p_\sigma)_{\sigma\in\mathrm{Nrv}(\mathcal{U}_r)})$ of O_r for each $r \in \mathbb{R}_{\geq 0}$. With this at hand, we see that the offset filtration is a functor $\mathbb{R}_{\geq 0} \to \mathsf{ClConv}_*$.

Theorem F (Theorem 5.6)**.** *For every pointed covered space $(X, \mathcal{A}_*) \in \mathsf{ClConv}_*$ there exists a homotopy equivalence*

$$\Gamma\colon |\,\mathrm{Sd}\,\mathrm{Nrv}(\mathcal{A})| \to X$$

that is natural with respect to the morphisms in ClConv_*.

One says that a cover is *good* if all non-empty finite intersections of cover elements are contractible. As we have already mentioned, nerve theorems usually assume that the covers involved are good. In Section 5.3, we again use the blowup complex to prove a functorial nerve theorem for simplicial complexes, which also appeared in the author's master's thesis [Rol20] in a preliminary form and with a different proof strategy:

Theorem G (Theorem 5.14)**.** *Let K be a simplicial complex and let $\mathcal{A} = (K_i \subseteq K)_{i\in I}$ be a good cover of K by subcomplexes. The natural maps $\rho_S\colon \mathrm{Blowup}(|\mathcal{A}|) \to |K|$ and $\rho_N\colon \mathrm{Blowup}(|\mathcal{A}|) \to |\,\mathrm{Nrv}(\mathcal{A})|$ are homotopy equivalences.*

The proof is related to work of Björner [Bjö03; Bjö81], and uses elementary methods from combinatorial homotopy theory for constructing homotopy equivalences between simplicial complexes, together with discrete Morse theory. Combining this result with a well-known theorem on triangulations of semi-algebraic sets (Lemma 5.15), we obtain a nerve theorem for compact semi-algebraic sets that are covered by finitely many closed, semi-algebraic subspaces.

Finally we use techniques from abstract homotopy theory to prove the following omnibus functorial nerve theorem. In particular, this result implies Theorems E and G. In parts 1(b) and 2(b) of the following Theorem H, we restrict attention to compactly-generated spaces. This is a standard hypothesis in algebraic topology, as these spaces form a "convenient" subcategory of topological spaces that is suitable for developing the machinery of homotopy theory. In part 1(b), the intersection A_T and the *latching space* $L(T) = \bigcup_{T \subsetneq J} A_J \subseteq A_T$ are assumed to satisfy the homotopy extension property; for example, CW-pairs satisfy the homotopy extension property (see Remark 2.66). These assumptions on the latching spaces together with the assumption that the cover is locally finite dimensional allow for inductive arguments analogous to arguments that employ induction over the skeleton of a CW-complex. In Sections 2.4.2 and 5.4 we introduce the notions used in the statement of the following theorem.

Theorem H (Unified Nerve Theorem 5.25)**.** *Let X be a topological space and let $\mathcal{A} = (A_i)_{i \in I}$ be a cover of X.*

1. *Consider the natural map $\rho_S \colon \mathrm{Blowup}(\mathcal{A}) \to X$.*
   - *(a) If $\mathcal{A}$ is an open cover, then ρ_S is a weak homotopy equivalence. If furthermore X is a paracompact Hausdorff space, or, more generally, if $\mathcal{A}$ is numerable, then ρ_S is a homotopy equivalence.*
   - *(b) Assume that X is compactly generated and that $\mathcal{A}$ is a closed cover that is locally finite and locally finite dimensional. If for any $T \in \mathrm{Nrv}(\mathcal{A})$ the latching space $L(T) \subseteq A_T$ is a closed subset and the pair $(A_T, L(T))$ satisfies the homotopy extension property, then ρ_S is a homotopy equivalence.*
2. *Consider the natural map $\rho_N \colon \mathrm{Blowup}(\mathcal{A}) \to |\,\mathrm{Nrv}(\mathcal{A})|$.*
   - *(a) If $\mathcal{A}$ is (weakly) good, then ρ_N is a (weak) homotopy equivalence.*
   - *(b) If for all $J \in \mathrm{Nrv}(\mathcal{A})$ the space A_J is compactly generated and $\mathcal{A}$ is homologically good with respect to a coefficient ring R, then ρ_N is an R-homology isomorphism.*

We now summarize the ingredients that go into the proof. First of all, the blowup complex is closely related to the bar construction (Definition 2.57). While the blowup complex has a natural map to the nerve of a cover, the bar construction has instead a natural map to the subdivision of the nerve, which is why we use the blowup complex in the statements of the theorems. However, the bar construction is in some ways easier to work with; in Section 5.1 we explain how we can work with either the blowup complex or the bar construction, whichever is more convenient.

The statement in 1(a) about weak homotopy equivalences follows from work of Dugger–Isaksen [DI04]; we give a short proof in Section 5.5. The statement in 1(a) about homotopy equivalences is proved in Hatcher's textbook [Hat02, Proposition 4G.2]. Statement 1(b) follows from a standard argument using Reedy model structures, which is similar to the proof of [Dug08, Corollary 14.17], for example. Both parts of 2(a) follow from the fact that the bar construction is homotopical, both for homotopy equivalences and weak homotopy equivalences. In the case of weak homotopy equivalences, this again uses Dugger–Isaksen [DI04]; see Proposition 5.43. We also give a short proof of this in Section 5.5. Finally, 2(b) is proved using the bar construction in the setting of simplicial R-modules.

In summary, the proof of Theorem H is straightforward, given some powerful tools for studying diagrams of spaces from abstract homotopy theory. In Sections 2.4 and 5.4 we provide an introduction to these tools.

The Literature on Nerve Theorems We now summarize the literature on the nerve theorem, with a particular focus on results that address functoriality.

The original work of Alexandroff [Ale28] on nerves, as well as the early nerve theorems of Leray [Ler45, Théorème 12] [Wu62] and Borsuk [Bor48, Corollary 3], considered closed covers, motivated in part by covers of polytopes by simplices. Open covers were considered by Weil [Wei52, Section 5], McCord [McC67], and Segal [Seg68]. There is renewed interest in the case of closed covers in applied topology, motivated by geometric constructions such as alpha shapes [EKS83].

A common way to relate the nerve of an open cover $\mathcal{U}$ with the covered space X is by a partition of unity subordinate to the cover. Such a partition of unity defines a map from X to the nerve of $\mathcal{U}$ in a straightforward way, which is a homotopy equivalence if $\mathcal{U}$ is good. This idea appears in the work of Weil and Segal, and the textbook proofs of the nerve theorem by Hatcher [Hat02, Corollary 4G.3] and Kozlov [Koz08, Theorem 15.21]. Moreover, up to homotopy this map is independent of the choice of partition of unity (see [ES52, Ch. X.11] for a discussion in a slightly different setting), and this map commutes up to homotopy with the maps on spaces and nerves induced by a morphism of covered spaces. This was observed by Chazal–Oudot [CO08] (for certain inclusions of covered spaces) and by Bauer–Edelsbrunner–Jabłoński–Mrozek [Bau+20] (for general morphisms); for similar results, see Lim–Mémoli–Okutan [LMO22, Theorem 6] and Virk [Vir21, Lemma 5.1].

In the case of open covers of paracompact Hausdorff spaces, Botnan–Spreemann [BS15], following the approach that goes back at least to Segal, observe that the blowup complex provides a zigzag of natural transformations relating covered spaces and nerves.

Ferry–Mischaikow–Nanda [FMN14] consider covers by open and closed balls in Euclidean space; they also use the blowup complex as an intermediate object relating spaces and nerves, but use the Vietoris–Smale theorem on proper maps with contractible fibers to obtain a homotopy equivalence from nerve to space with control over the image of each simplex.

Bendich–Cohen-Steiner–Edelsbrunner–Harer–Morozov [Ben+07; Mor08] give a nerve theorem for certain closed convex covers in Euclidean space; they define a map from the subdivision of the nerve of $\mathcal{U}$ to the space by choosing a point in each non-empty intersection $\mathcal{U}_J$, mapping the vertex J to this point, and extending by piecewise linear interpolation. They show this map commutes up to homotopy with maps induced by certain morphisms of covered spaces.

The references in the previous four paragraphs also give examples of applications in which functoriality of the nerve theorem is important. For more, see, e.g., work on approximate nerve theorems [GS18; CS18] and a comparison of persistent singular and Čech homology [Sch22].

Borel–Serre [BS73, Theorem 8.2.1] prove a nerve theorem for locally finite and closed covers whose nerve is finite dimensional and such that all finite non-empty intersections of cover elements are absolute retracts for metric spaces. Using similar techniques, Nagórko [Nag07] proves a nerve theorem for locally finite, locally finite dimensional, star-countable closed covers of normal spaces such that all non-empty intersections of cover elements are absolute extensors for metric spaces.

Björner [Bjö03] gives a proof of an n-connectivity version of the nerve theorem, which we discuss in Section 5.3.3. Given a good cover of a finite simplicial complex by subcomplexes, Barmak [Bar11] proves a related result, showing that the simplicial complex and the nerve have the same simple homotopy type.

2. Preliminaries

We recall some notions and statements that are used in the subsequent sections. Some of those have novel aspects, most notably Propositions 2.34, 2.37, 2.43 and 2.55.

2.1. Geometric Complexes

In this section, we define three fundamental geometric complexes that can be associated to a point cloud, namely the *Čech* and *Delaunay complex* (see Section 2.1.1), and the *Vietoris–Rips complex* (see Section 2.1.2), all of which give rise to a filtration of simplicial complexes. Moreover, we recall the nerve theorem and concepts from metric geometry.

We first recall some basic notions from combinatorial topology [Mun84; EH10; Koz08]. A *simplicial complex* K is a collection of non-empty finite sets such that for any set $\sigma \in K$ and any non-empty subset $\rho \subseteq \sigma$ one has $\rho \in K$. The *vertex set* of K is the union $\operatorname{Vert} K = \bigcup_{\sigma \in K} \sigma$. A set $\sigma \in K$ is called a *simplex*, and $\dim \sigma = \operatorname{card} \sigma - 1$ is its *dimension*. A simplex $\rho \subseteq \sigma$ is said to be a *face* of σ and σ a *coface* of ρ. If additionally $\dim \rho = \dim \sigma - 1$, then we call ρ a *facet* of σ, σ a *cofacet* of ρ, and (ρ, σ) a *facet pair*. The *star* of σ, $\operatorname{St}_K(\sigma)$, is the set of cofaces of σ in K, and the *closure* of σ, $\operatorname{Cl}(\sigma)$, is the set of its faces. For a subset $E \subseteq K$, we write $\operatorname{St}_K(E) = \bigcup_{\sigma \in E} \operatorname{St}_K(\sigma)$ and $\operatorname{Cl}(E) = \bigcup_{\sigma \in E} \operatorname{Cl}(\sigma)$. Moreover, for a set X, we write $\operatorname{Cl}(X) = \{\sigma \subseteq X \mid \sigma \text{ non-empty and finite}\}$ for the full simplicial complex on X. The *closed star* of σ, $\operatorname{clst}_K(\sigma)$, is the closure of the star and the *link* of σ, $\operatorname{Lk}_K(\sigma)$, is the set of simplices in the closed star that are disjoint from σ. A *simplicial map* $f \colon K \to L$ between two simplicial complexes is given by a vertex map $f \colon \operatorname{Vert} K \to \operatorname{Vert} L$ such that for any $\sigma \in K$ we have $f(\sigma) \in L$. We denote the category of simplicial complexes together with simplicial maps by Simp. Moreover, we denote by $|K|$ the *geometric realization* of K; it gives rise to a functor $|\cdot| \colon \mathsf{Simp} \to \mathsf{Top}$ into the category of topological spaces. Write Po for the category of posets. Let $\operatorname{Pos} \colon \mathsf{Simp} \to \mathsf{Po}$ be the functor that takes a simplicial complex to its poset of simplices (ordered by inclusion), and let $\operatorname{Flag} \colon \mathsf{Po} \to \mathsf{Simp}$ be the functor that takes a poset P to the simplicial complex whose vertices are the elements of P and whose n-simplices are the chains $x_0 < \cdots < x_n$ of elements in P. The *barycentric subdivision* of a simplicial complex K is $\operatorname{Sd}(K) = \operatorname{Flag}(\operatorname{Pos}(K))$. There is an affine linear homeomorphism $\alpha_K \colon |\operatorname{Sd}(K)| \to |K|$ defined by the vertex map that sends a vertex σ of $\operatorname{Sd}(K)$ to the barycenter of $|\sigma|$ in $|K|$. Note that, while the homeomorphism α_K is natural with respect to inclusions of simplicial complexes, it is not natural with respect to general simplicial maps. For any poset P we denote by P^{op} the *opposite poset*, i.e., the poset with the same underlying set and $x \leq_{P^{\mathrm{op}}} y$ if and only if $x \geq_P y$. Let P be a poset, viewed as a category with objects the elements of the poset and with a unique morphism $p \to q$ if and only if $p \leq q$. A *filtration* of simplicial complexes is a functor $K_\bullet \colon P \to \mathsf{Simp}$ such that for any $p \leq q$ the map $K_p \to K_q$ is injective. A function $f \colon K \to \mathbb{R}$ defined on a simplicial complex is *monotonic* if for any $\sigma, \tau \in K$ with $\sigma \subseteq \tau$ we have $f(\sigma) \leq f(\tau)$. For a monotonic function $f \colon K \to \mathbb{R}$ we write $S_r(f) = f^{-1}(-\infty, r] \subseteq K$ for the sublevel set and $S_r^<(f) = f^{-1}(-\infty, r) \subseteq K$ for the open sublevel set of f at scale $r \in \mathbb{R}$.

2.1.1. Čech and Delaunay Complexes

We now define the Čech and Delaunay complex of a subset of a metric space; see Fig. 5 for an illustration. We also explain how these complexes can alternatively be described, in the special case of subsets of Euclidean space, in terms of smallest enclosing spheres and smallest empty circumspheres, respectively. Finally, we present a nerve theorem that applies to the Čech and Delaunay complexes of finite subsets of Euclidean space.

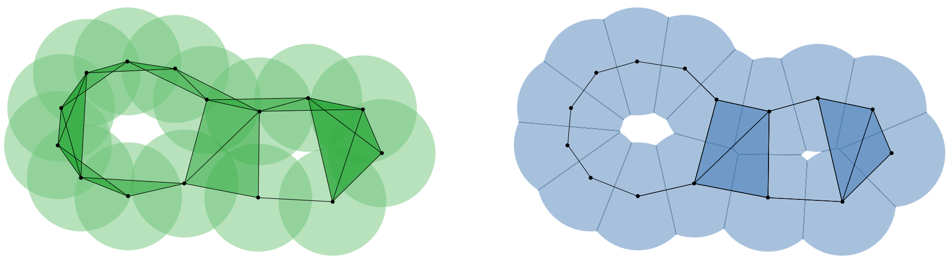

Figure 5: A cover by closed balls (left) and closed Voronoi balls (right) together with the corresponding Čech and Delaunay complex of a finite subset of the plane.

Let Y be a metric space, let $p \in Y$ be a point, and let $r \geq 0$ be a real number. We write $D_r(p) = \{y \in Y \mid d(p, y) \leq r\}$ and $B_r(p) = \{y \in Y \mid d(p, y) < r\}$ for the closed metric ball and the open metric ball of radius r centered at p, respectively. Moreover, we write $S_r(p) = \{y \in Y \mid d(p, y) = r\}$ for the sphere of radius r centered at p. We first define the Čech complex.

Definition 2.1. Let $X \subseteq Y$ be a subset of a metric space and $r \geq 0$ a real number. The *Čech complex* $\check{\mathrm{C}}\mathrm{ech}_r(X, Y)$ is the nerve of the cover given by closed balls in Y of radius r centered at the points in X:

$$\check{\mathrm{C}}\mathrm{ech}_r(X, Y) = \{\emptyset \neq \sigma \subseteq X \mid \sigma \text{ finite}, \bigcap_{x \in \sigma} D_r(x) \neq \emptyset\}.$$

The *Čech complex* $\check{\mathrm{C}}\mathrm{ech}_r^{<}(X, Y)$ is the nerve of the cover given by open balls in Y of radius r centered at the points in X:

$$\check{\mathrm{C}}\mathrm{ech}_r^{<}(X, Y) = \{\emptyset \neq \sigma \subseteq X \mid \sigma \text{ finite}, \bigcap_{x \in \sigma} B_r(x) \neq \emptyset\}.$$

To simplify the notation, we write $\check{\mathrm{C}}\mathrm{ech}(X, Y)$ for $\check{\mathrm{C}}\mathrm{ech}_\infty(X, Y)$. If the ambient metric space Y is unambiguous from the context, we also write $\check{\mathrm{C}}\mathrm{ech}_r(X)$ for $\check{\mathrm{C}}\mathrm{ech}_r(X, Y)$, and $\check{\mathrm{C}}\mathrm{ech}_r^{<}(X)$ for $\check{\mathrm{C}}\mathrm{ech}_r^{<}(X, Y)$.

We now define *Voronoi balls*, which are used in the definition of the Delaunay complex below. Let $X \subseteq Y$ be a subset of a metric space, $p \in X$ a point, and $r \geq 0$ a real number. The *Voronoi domain* $\mathrm{Vor}(p, X)$ of p with respect to X is given by

$$\mathrm{Vor}(p, X) = \{y \in Y \mid d(y, p) \leq d(y, x) \text{ for all } x \in X\}.$$

The *(closed) Voronoi ball* $\mathrm{Vor}_r(p, X)$ and the *(strict) Voronoi ball* $\mathrm{Vor}_r^<(p, X)$ of radius $r \geq 0$ centered at $p \in X$ are given by

$$\mathrm{Vor}_r(p, X) = D_r(p) \cap \mathrm{Vor}(p, X),$$
$$\mathrm{Vor}_r^<(p, X) = B_r(p) \cap \mathrm{Vor}(p, X).$$

We now define the Delaunay complex, sometimes also referred to as *alpha* complex.

Definition 2.2. Let $X \subseteq Y$ be a subset of a metric space and $r \geq 0$ a real number. The *Delaunay complex* $\mathrm{Del}_r(X, Y)$ is the nerve of the cover given by closed Voronoi balls in Y of radius r centered at the points in X:

$$\mathrm{Del}_r(X, Y) = \{\emptyset \neq \sigma \subseteq X \mid \sigma \text{ finite}, \bigcap_{x \in \sigma} \mathrm{Vor}_r(x, X) \neq \emptyset\}.$$

The *Delaunay complex* $\mathrm{Del}_r^<(X, Y)$ is the nerve of the cover given by strict Voronoi balls in Y of radius r centered at the points in X:

$$\mathrm{Del}_r^<(X, Y) = \{\emptyset \neq \sigma \subseteq X \mid \sigma \text{ finite}, \bigcap_{x \in \sigma} \mathrm{Vor}_r^<(x, X) \neq \emptyset\},$$

To simplify the notation, we write $\mathrm{Del}(X, Y)$ for $\mathrm{Del}_\infty(X, Y)$. If the ambient metric space Y is unambiguous from the context, we also write $\mathrm{Del}_r(X)$ for $\mathrm{Del}_r(X, Y)$, and $\mathrm{Del}_r^<(X)$ for $\mathrm{Del}_r^<(X, Y)$.

Minimal Enclosing Spheres and Circumspheres Assume that Y is the Euclidean space $\mathbb{R}^d$ equipped with the Euclidean metric $d(x, y) = \|x - y\|_2$. We now describe, following [BE17], how the Čech and Delaunay complex can alternatively be described in terms of smallest enclosing spheres and smallest empty circumspheres, respectively.

Let $X \subseteq \mathbb{R}^d$ be any subset, and let $\sigma \subseteq X$. We say that a sphere $S = S_r(p)$ is a *circumsphere* of σ if all points of σ lie on S,

$$\sigma \subseteq S = \{y \in \mathbb{R}^d \mid d(p, y) = r\},$$

and it is an *empty circumsphere* of σ if additionally no point of X lies inside S,

$$X \cap B_r(p) = X \cap \{y \in \mathbb{R}^d \mid d(p, y) < r\} = \emptyset.$$

Moreover, we say that S is an *enclosing sphere* of σ if all points of σ lie on or inside S,

$$\sigma \subseteq D_r(p) = \{y \in \mathbb{R}^d \mid d(p, y) \leq r\}.$$

We denote by $\mathrm{On}\, S = X \cap S$ the set of points in X that lie on the sphere, and we denote by $\mathrm{Incl}\, S = X \cap D_r(p)$ the set of points in X that are enclosed by the sphere. It is not difficult to see that if σ has an enclosing sphere, then it also has a smallest enclosing sphere, i.e., an enclosing sphere with smallest possible radius, which we denote by $S(\sigma, \emptyset)$. Similarly, if σ has a circumsphere, then it also has a smallest circumsphere.

Further, if σ has an empty circumsphere, then it also has a smallest empty circumsphere, which we denote by $S(\sigma, X)$. Note that such a smallest empty circumsphere is also the smallest circumsphere of some point set, namely $\operatorname{On} S$. The smallest circumsphere of σ, if it exists, is given by the unique circumsphere whose center lies in the *affine hull*

$$\operatorname{aff} \sigma = \{\sum_{x \in \sigma} \lambda_x \cdot x \mid \sum_{x \in \sigma} \lambda_x = 1\}$$

of σ. This can be deduced, for example, from the *Karush–Kuhn–Tucker conditions*, which give optimality conditions for (nonlinear) optimization problems with inequality constraints, as explained in [BE17]. A sufficient condition for the existence of a circumsphere of σ is that σ is *affinely independent*, meaning that the affine hull of σ is a $\dim \sigma$-dimensional affine subspace, or equivalently, the coefficients of any affine combination are unique in the sense that if

$$\sum_{x \in \sigma} \lambda_x \cdot x = \sum_{x \in \sigma} \mu_x \cdot x \quad \text{with} \quad \sum_{x \in \sigma} \lambda_x = \sum_{x \in \sigma} \mu_x = 1,$$

then $\lambda_x = \mu_x$ for all $x \in \sigma$. In this case, we denote by $\operatorname{Front} S = \{x \in \operatorname{On} S \mid \lambda_x > 0\}$ the set of points $x \in \operatorname{On} S$ that contribute positively to the unique affine combination of the center $z = \sum_{x \in \operatorname{On} S} \lambda_x \cdot x$, with $\sum_{x \in \operatorname{On} S} \lambda_x = 1$, of the smallest circumsphere of σ.

Čech and Delaunay complexes are related to spheres in the following way. Observe that for a non-empty and finite subset $\sigma \subseteq X$, we have $\sigma \in \check{\mathrm{C}}\mathrm{ech}_r(X)$ and $\sigma \in \check{\mathrm{C}}\mathrm{ech}_r^<(X)$ if and only if there exists an enclosing sphere S of σ with radius at most r and strictly smaller than r, respectively. Moreover, $\sigma \in \mathrm{Del}_r(X)$ and $\sigma \in \mathrm{Del}_r^<(X)$ if and only if there exists an empty circumsphere S of σ with radius at most r and strictly smaller than r, respectively. This leads to the following definition.

Definition 2.3. Let $X \subseteq \mathbb{R}^d$ be any subset. The *Čech radius function* is the monotonic function

$$r_\emptyset \colon \check{\mathrm{C}}\mathrm{ech}(X) \to \mathbb{R}$$

that assigns to a simplex the radius of its smallest enclosing sphere. Moreover, the *Delaunay radius function* is the monotonic function

$$r_X \colon \mathrm{Del}(X) \to \mathbb{R}$$

that assigns to a simplex the radius of its smallest empty circumsphere.

Note that for any real number $r \geq 0$, the sublevel sets of the Čech radius function and the Delaunay radius function at scale r are the Čech complex $\check{\mathrm{C}}\mathrm{ech}_r(X)$ and the Delaunay complex $\mathrm{Del}_r(X)$, respectively. Similarly, their open sublevel sets at scale r are the Čech complex $\check{\mathrm{C}}\mathrm{ech}_r^<(X)$ and the Delaunay complex $\mathrm{Del}_r^<(X)$, respectively.

In [BE17] the authors show that the Čech radius function and the Delaunay radius function are both generalized discrete Morse functions (Proposition 2.29), if X is a generic subset of $\mathbb{R}^d$ in the following sense.

Definition 2.4. A finite subset $X \subseteq \mathbb{R}^d$ is in *general position* if for every $\sigma \subseteq X$ of at most $d+1$ points

- σ is affinely independent, and
- no point of $X \setminus \sigma$ lies on the smallest circumsphere of σ.

We make use of the fact above in Section 2.3.2, where we extend the definition of the *Wrap complex* [BE17], associated to the Delaunay radius function, to a more general setting. This is then used in Section 4, where we study the connection between discrete Morse theory and persistent homology in the context of shape reconstruction methods.

The Nerve Theorem for Closed Convex Covers Recall that both the Čech complex $\check{\mathrm{C}}\mathrm{ech}_r(X)$ and the Delaunay complex $\mathrm{Del}_r(X)$ are nerves of closed covers, namely the covers given by closed balls and closed Voronoi balls, respectively. We now present a nerve theorem that allows us to replace the union of balls, which equals the union of Voronoi balls to the same parameter, with the Čech or Delaunay complex without loosing any homotopy theoretic information.

Motivated by the Čech and Delaunay complexes of finite point sets, we consider nerves of finite closed and convex covers of subsets of $\mathbb{R}^d$. We write $[n]$ for the set $\{0, \dots, n\}$. Let $\mathcal{C} = (C_i)_{i \in [n]}$ be a collection of closed convex subsets of $\mathbb{R}^d$, and let X be their union. We now explain the construction of a continuous map $\Gamma\colon |\operatorname{Sd}\operatorname{Nrv}(\mathcal{C})| \to X$ that is a homotopy equivalence, establishing a nerve theorem for this setting. In Section 5 we extend this result to prove the two functorial versions Theorem E and Theorem F. Each vertex $J \in \operatorname{Sd}\operatorname{Nrv}(\mathcal{C})$ represents a simplex in the nerve $\operatorname{Nrv}(\mathcal{C})$, and hence we can choose a point p_J from the non-empty intersection $C_J = \bigcap_{j \in J} C_j$. By convexity of the cover elements in $\mathcal{C}$, this choice extends uniquely to a map $\Gamma\colon |\operatorname{Sd}\operatorname{Nrv}(\mathcal{C})| \to X$ that is affine linear on each simplex of the barycentric subdivision; see Fig. 6 for an illustration. Similar constructions can be found in the literature [BT82, Theorem 13.4] [Hau95, p. 179] [Ben+07, p. 544].

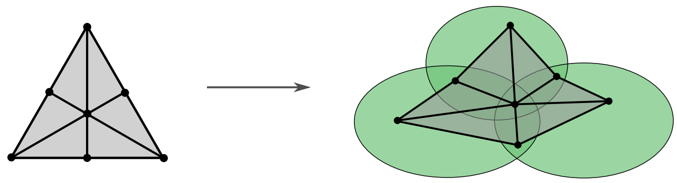

Figure 6: Illustration of the map Γ.

The following theorem already appeared, with the additional assumption on the cover elements to be compact, in the author's master's thesis [Rol20]. For a proof of this theorem, see Appendix A, where a homotopy inverse Ψ to Γ is constructed by using a partition of unity subordinate to an open thickening of the cover elements, similar as in the familiar proof of the nerve theorem for open covers [Hat02, Proposition 4G.2].

Theorem 2.5. *The map Γ is a homotopy equivalence. In particular, the nerve $|\operatorname{Nrv}(\mathcal{C})|$ is homotopy equivalent to the space X.*

In Section 1.4 we give a thorough overview of all the (functorial) nerve theorems in this thesis, some of which also apply to the case of open covers. Compare Table 1.

Remark 2.6. The proof strategy of Theorem 2.5 can be generalized to the following setting. Let $\mathcal{A} = (A_i)_{i \in I}$ be a finite closed cover of a metric space X. Assume there exists a continuous map

$$\sigma \colon \bigcup_{i \in I} (A_i \times A_i \times [0,1]) \to X,$$

that restricts for every $i \in I$ to a bicombing $A_i \times A_i \times [0,1] \to A_i$ (see Definition 2.10). In particular, this implies that every non-empty intersection of finitely many cover elements in $\mathcal{A}$ is contractible (compare Remark 2.11). Thus, similarly to the above, there exists a continuous map $\Gamma \colon |\operatorname{Sd}\operatorname{Nrv}(\mathcal{A})| \to X$ that satisfies the analog of Proposition A.6 in this setting. The rest of the proof also transfers directly to this setting, where the straight line homotopy between the identity id_X and the composition $\Gamma \circ \Psi$ is replaced with the homotopy $H \colon X \times [0,1] \to X$ given by

$$H(x,t) = \sigma(x, \Gamma \circ \Psi(x), t).$$

We remark that such a map σ is also a crucial ingredient in the proof of the nerve theorem due to Weil [Wei52, Section 5]. The reasoning above establishes a nerve theorem, which applies, for example, to the cover of the circle by three closed arcs as in Example 2.61, where the map σ is obtained by normalizing the straight line homotopy in the plane.

2.1.2. Vietoris–Rips Complexes and Metric Geometry

We now define the Vietoris–Rips complex of a metric space (see Fig. 7 for an illustration) and recall some concepts from metric geometry. Moreover, we explain how any metric space can be isometrically embedded into an associated universal metric space, such that its Vietoris–Rips complex is isomorphic to a Čech complex inside this universal metric space. In particular, Vietoris–Rips complexes are also nerves of covers.

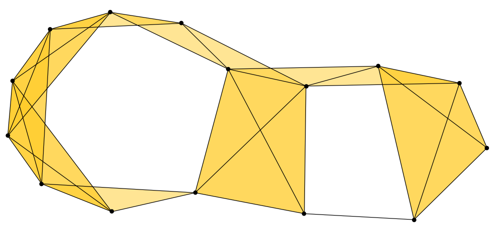

Figure 7: The Vietoris–Rips complex of a finite subset of the plane.

Definition 2.7. Let X be a metric space and $t \geq 0$ a real number. The *Vietoris–Rips complex* $\operatorname{Rips}_t(X)$ is the simplicial complex consisting of non-empty and finite subsets of X with diameter at most t:

$$\operatorname{Rips}_t(X) = \{\emptyset \neq \sigma \subseteq X \mid \sigma \text{ finite}, \ \operatorname{diam} \sigma \leq t\}.$$

The *Vietoris–Rips complex* $\mathrm{Rips}_t^<(X)$ is the simplicial complex consisting of non-empty and finite subsets of X with diameter strictly smaller than t:

$$\mathrm{Rips}_t^<(X) = \{\emptyset \neq \sigma \subseteq X \mid \sigma \text{ finite},\ \operatorname{diam} \sigma < t\}.$$

Note that the Vietoris–Rips complex is a *clique complex*, meaning that a non-empty and finite subset $\sigma \subseteq X$ is contained in $\mathrm{Rips}_t(X)$ if and only if every two element subset of σ is contained in $\mathrm{Rips}_t(X)$.

Geodesic Metric Spaces and Bicombings We recall some concepts from metric geometry, assuming some familiarity with the basic notions (for example, see [BH99]). Let X be a metric space. A *geodesic* between two points x and y in X is an isometric map $[0, d(x, y)] \to X$ such that $0 \mapsto x$ and $d(x, y) \mapsto y$. Equivalently, it is a path $\gamma\colon [0, 1] \to X$ between x and y such that $d(\gamma(t), \gamma(s)) = |t - s| d(x, y)$ for every $t, s \in [0, 1]$.

Definition 2.8. A metric space X is *star-geodesic with respect to* $p \in X$ if for every point $x \in X$ there exists a geodesic between p and x. Further, the metric space X is *geodesic* if it is star-geodesic with respect to every point in X.

Remark 2.9. Geodesics, if they exist, do not need to be unique nor depend continuously on their endpoints. Consider, for example, the unit circle $S^1 \subseteq \mathbb{R}^2$ equipped with the intrinsic metric, i.e., the distance between two points is given by the angle between those two vectors in radians. This metric space is not uniquely geodesic, as between any two antipodal points there always exist two geodesics, one running clockwise and the other counterclockwise. An example of a metric space that has unique geodesics, but which do not depend continuously on the endpoints is given in [BH99, Exercise 3.14].

To overcome the issues mentioned in Remark 2.9, one can consider spaces equipped with a distinguished family of geodesics. This idea dates back at least to Busemann [BP87].

Definition 2.10. Let X be a metric space. A continuous map $\sigma\colon X \times X \times [0, 1] \to X$ is a *bicombing* if for any two points $x, y \in X$ the map σ restricts to a path

$$\sigma_{x,y} = \sigma(x, y, \cdot)\colon [0, 1] \to X$$

between x and y, i.e., we have $\sigma_{x,y}(0) = x$ and $\sigma_{x,y}(1) = y$. It is a *geodesic bicombing* if for any two points $x, y \in X$ the path $\sigma_{x,y}$ is a geodesic between x and y.

Remark 2.11. Note that by the tensor-hom adjunction, also known as currying, the continuous maps $X \times X \times [0, 1] \to X$ correspond bijectively to the continuous maps $X \times X \to P(X)$, where $P(X) = \{\eta\colon [0, 1] \to X \text{ continuous}\}$ is the path space equipped with the compact-open topology. Under this identification, a bicombing corresponds to a section of the evaluation map $\mathrm{ev}\colon P(X) \to X \times X,\ \eta \mapsto (\eta(0), \eta(1))$ and vice versa.

Such a section is also known as a *continuous motion planning*, and it is not too difficult to see that it exists if and only if X is contractible. Motivated by this, Farber [Far03] introduced the *topological complexity* $\mathrm{TC}(X)$ of X as the minimal number k such that there is a finite open cover $(U_i)_{i \in \{1,\ldots,k\}}$ of $X \times X$ and local sections $s_i\colon U_i \to P(X)$ of the evaluation map, meaning that $\mathrm{ev} \circ s_i = \mathrm{id}_{U_i}$. For example, the topological complexity of the circle is $\mathrm{TC}(S^1) = 2$.

Hyperconvex Metric Spaces and the Injective Hull The following property ensures that the Čech complex of a subset of a metric space agrees with its Vietoris–Rips complex. See [EK01] for a thorough introduction to hyperconvex metric spaces.

Definition 2.12. A metric space is *hyperconvex* if it is geodesic and if any collection of closed balls has the Helly property, meaning that if any two of these balls have a non-empty intersection, then all these balls have a non-empty intersection.

The following lemma is a direct consequence of this definition and the fact that Vietoris–Rips complexes are clique complexes; see also [LMO22, Proposition 2.2].

Lemma 2.13. *If Y is a hyperconvex metric space and $X \subseteq Y$ is a subspace, then for any $t \geq 0$ we have*

$$\mathrm{Rips}_{2t}(X) = \check{\mathrm{C}}\mathrm{ech}_t(X, Y) \quad \textit{and} \quad \mathrm{Rips}_{2t}^{<}(X) = \check{\mathrm{C}}\mathrm{ech}_t^{<}(X, Y).$$

Alternatively, hyperconvex metric spaces can be characterized by the following lifting property [Lan13; Isb64].

Proposition 2.14. *A metric space Y is hyperconvex if and only if it is* injective*, meaning that for any isometric embedding $e\colon A \to B$ and any 1-Lipschitz map $f\colon A \to Y$, there exists a 1-Lipschitz map $F\colon B \to Y$ such that $F \circ e = f$, i.e., such that the following diagram commutes:*

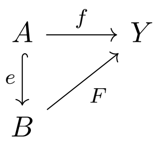

Every metric space embeds into an injective metric space: Let X be a metric space. We now describe its *injective hull* $E(X)$, following Lang [Lan13]. A function $f\colon X \to \mathbb{R}$ with $f(x) + f(y) \geq d(x, y)$ for all $x, y \in X$ is *extremal* if $f(x) = \sup_{y \in X}(d(x, y) - f(y))$ for every $x \in X$. The difference between any two extremal functions turns out to be bounded, and so we can equip the set $E(X)$ of extremal functions with the metric induced by the supremum norm, i.e., $d(f, g) = \sup_{x \in X} |f(x) - g(x)|$. We define an isometric embedding $e\colon X \to E(X)$ by $y \mapsto d_y$, where $d_y(x) = d(y, x)$.

By the following result, as can be found in [Lan13; Isb64], the injective hull $E(X)$ of X is the smallest injective metric space containing X, justifying the naming.

Proposition 2.15. *For any metric space X, the metric space $E(X)$ is injective, and therefore hyperconvex. Moreover, if $i\colon X \to Y$ is an isometric embedding into some other injective metric space Y, then there exists an isometric embedding $j\colon E(X) \to Y$ such that $j \circ e = i$, i.e., such that the following diagram commutes:*

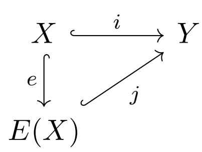

As $e\colon X \to E(X)$ is an isometric embedding of X into its injective hull, the following is a direct consequence of Lemma 2.13 and Proposition 2.15. It shows that Vietoris–Rips complexes are also nerves of covers.

Proposition 2.16. *Let X be any metric space. Then for any $t \geq 0$ we have*

$$\mathrm{Rips}_{2t}(X) \cong \check{\mathrm{C}}\mathrm{ech}_t(e(X), E(X)) \quad \textit{and} \quad \mathrm{Rips}_{2t}^{<}(X) \cong \check{\mathrm{C}}\mathrm{ech}_t^{<}(e(X), E(X)).$$

In Section 3.2 we demonstrate how this can be used to show that the Vietoris–Rips complex of a δ*-hyperbolic* and ν*-almost geodesic* metric space is contractible for a sufficiently large parameter, by using the nerve theorem. By the following result, as can be found in [Lan13; Isb64; LMO22], we can continuously interpolate between any two points in an injective metric space using geodesics, and closed as well as open balls are geodesically convex. In particular, covers by closed or open balls are good covers in the sense of Section 1.4.

Proposition 2.17. *Let Y be an injective metric space. There exists a geodesic bicombing $\sigma\colon Y \times Y \times [0,1] \to Y$ such that for all $x, y, z, p \in Y$ and $t \in [0,1]$ we have*

$$d(\sigma_{x,y}(t), \sigma_{z,p}(t)) \leq (1-t) \cdot d(x,z) + t \cdot d(y,p).$$

In particular, Y is contractible. Moreover, non-empty intersections of closed balls are contractible, and non-empty intersections of open balls are also contractible.

Gromov–Hausdorff Distance The similarity between two metric spaces can be expressed in terms of the *Gromov–Hausdorff distance* [BBI01; BH99], which measures how far two metric spaces are from being isometric.

A *correspondence* C between two metric spaces X and Y is a subset $C \subseteq X \times Y$ such that $\mathrm{pr}_X(C) = X$ and $\mathrm{pr}_Y(C) = Y$, where $\mathrm{pr}_X\colon X \times Y \to X$ and $\mathrm{pr}_Y\colon X \times Y \to Y$ are the projections to X and Y, respectively.

Definition 2.18. Let X and Y be metric spaces. The *Gromov–Hausdorff distance* between X and Y is given by

$$d_{\mathrm{GH}}(X,Y) = \frac{1}{2} \inf\{\mathrm{dis}(C) \mid C \text{ correspondence between } X \text{ and } Y\},$$

where $\mathrm{dis}(C) = \sup\{|d_X(x,x') - d_Y(y,y')| \ \mid (x,y),(x',y') \in C\}$ is the *distortion* of C.

We remark that the Gromov–Hausdorff distance defines a metric on the space of isometry classes of compact metric spaces [BBI01], and that determining the exact value of the Gromov–Hausdorff distance between two metric spaces is a challenging task, in general, even for spheres [LMS23]. We use the Gromov–Hausdorff distance in the statements of the stability results Propositions 2.21, 3.1 and 3.7.

2.2. Persistent Homology

We briefly recall some aspects of *persistent homology* [ELZ02; Rob99; ZC05; EH10], which, in short, is the homology of a filtration. In particular, we introduce the *barcode* of a *persistence module*, explain how it can be computed in the special case of a finitely filtered based chain complex from a reduction of the filtration boundary matrix, and discuss the *apparent pairs* shortcut that can be used to speed up the computations. Finally, we introduce *lexicographically minimal chains*, which are used in Section 4.

Let $K_\bullet \colon \mathbb{R}_{\geq 0} \to \mathsf{Simp}$ be any of the filtrations from Section 2.1 that can be associated to a finite metric space. By applying the simplicial homology functor $H_d \colon \mathsf{Simp} \to \mathsf{vec}_\mathbb{F}$ to the filtration $K_\bullet$, where $\mathsf{vec}_\mathbb{F}$ is the category of finite dimensional vector spaces over a field $\mathbb{F}$, we obtain the *persistent homology* $H_d(K_\bullet) \colon \mathbb{R}_{\geq 0} \to \mathsf{vec}_\mathbb{F}$ of the filtration $K_\bullet$.

This motivates the following.

Definition 2.19. An ($\mathbb{R}_{\geq 0}$-indexed) *persistence module* is a functor $M \colon \mathbb{R}_{\geq 0} \to \mathsf{vec}_\mathbb{F}$, meaning that M assigns to every $r \in \mathbb{R}_{\geq 0}$ a finite dimensional $\mathbb{F}$-vector space M_r, and to every pair $r \leq t$ of non-negative real numbers a linear map $M_{t,r} \colon M_r \to M_t$, such that $M_{r,r} = \mathrm{id}_{M_r}$ and if $s \in \mathbb{R}_{\geq 0}$ with $r \leq s \leq t$, then $M_{t,r} = M_{t,s} \circ M_{s,r}$.

Persistence modules decompose into elementary summands [Cha+16; BC20].

Proposition 2.20. *Let $M \colon \mathbb{R}_{\geq 0} \to \mathsf{vec}_\mathbb{F}$ be a persistence module. There exists a unique multiset* $\mathrm{Barc}(M)$ *of intervals in $\mathbb{R}_{\geq 0}$, called the* barcode *of M, such that M decomposes as*

$$M \cong \bigoplus_{I \in \mathrm{Barc}(M)} \mathbb{F}(I),$$

where $\mathbb{F}(I) \colon \mathbb{R}_{\geq 0} \to \mathsf{vec}_\mathbb{F}$ is the interval module *given by*

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

and such that for $r, t \in I$ with $r \leq t$ we have $\mathbb{F}(I)_{t,r} = \mathrm{id}_\mathbb{F}$.

We show further below, where we also introduce some more terminology, how the barcode of the persistent homology of an elementwise filtered based chain complex can be obtained from a reduction of the filtration boundary matrix. The barcode of an arbitrary filtration of a based chain complex can then be obtained through the barcode of any elementwise refinement; compare [Bau21, Proposition 2.1].

One central result for persistent homology is the *stability theorem* [CEH07], which states that similar persistence modules have similar barcodes. There are many variants of the stability theorem in the literature [Cha+16; Cha+09; CSO14; BL20]. We present the following result from [CSO14] as a representative, noting that there is an analogous result for the Čech filtration. Here, $d_B(\mathrm{Barc}(M), \mathrm{Barc}(N))$ is the *bottleneck distance* between two barcodes, which measures, in analogy to the Gromov–Hausdorff distance (Definition 2.18), how far two barcodes are from being equal.

Proposition 2.21. *Let X and Y be any finite metric spaces. For any homology degree d we have*

$$d_B(\mathrm{Barc}(M), \mathrm{Barc}(N)) \leq 2d_{\mathrm{GH}}(X, Y),$$

where $M = H_d(\mathrm{Rips}_\bullet(X))$ and $N = H_d(\mathrm{Rips}_\bullet(Y))$.

Based Chain Complexes and Filtrations We assume the reader to be familiar with the basics of homological algebra (see, e.g., [Mun84; Wei94]). By a *based chain complex* (C_*, Σ_*) (sometimes also called a *Lefschetz complex* [Lef42]) we mean a bounded chain complex $C_* = (C_n, \partial)_{n\in\mathbb{N}}$ of finite dimensional vector spaces over a field $\mathbb{F}$ together with a basis Σ_n for each C_n. Consider the canonical bilinear form $\langle\cdot,\cdot\rangle$ on C_* for the given basis Σ_*, i.e., for $a, b \in \Sigma_*$ we have $\langle a, b\rangle = 0$ if $a \neq b$ and $\langle a, a\rangle = 1$. Given two basis elements $c \in \Sigma_n$ and $e \in \Sigma_{n+1}$ such that $\langle\partial e, c\rangle \neq 0$, we call c a *facet* of e and e a *cofacet* of c, and we call (c, e) a *facet pair*.

A *filtration* of (C_*, Σ_*) is a collection of based chain complexes $(C^i_*, \Sigma^i_*)_{i\in I}$, where I is a totally ordered indexing set, such that $\Sigma^i_* \subseteq \Sigma_*$ spans the subcomplex C^i_* of C_* for all $i \in I$, and $i \leq j$ implies $\Sigma^i_* \subseteq \Sigma^j_*$. We call the filtration an *elementwise filtration* if for any j with immediate predecessor i we have that $\Sigma^j_* \setminus \Sigma^i_*$ contains exactly one basis element σ_j. Thus, elementwise filtrations of (C_*, Σ_*) correspond bijectively to total orders $<$ on Σ_* such that prefixes $\downarrow\sigma_j = \{\sigma_i \mid \sigma_i \leq \sigma_j\}$ span subcomplexes.

Remark 2.22. Our main example for a based chain complex is the simplicial chain complex $C_*(K)$ of a finite simplicial complex K with coefficients in a field. If the vertices of K are totally ordered, then there is a canonical basis of $C_n(K)$ consisting of the n-dimensional simplices of K oriented according to the given vertex order, and a simplexwise filtration of K induces a canonical elementwise filtration of $C_*(K)$.

Matrix Reduction For a based chain complex $(C_*, \Sigma_* = \sigma_1 < \cdots < \sigma_l)$ with an elementwise filtration, we often identify an element of C_* with its coordinate vector in $\mathbb{F}^l$. The *filtration boundary matrix* D of an elementwise filtration is the matrix that represents the boundary map ∂ with respect to the total order on the basis elements induced by the filtration.

For a matrix R, we denote by R_j the jth column of R, and by $R_{i,j}$ the entry of R in row i and column j. The *pivot* of a column R_j, denoted by $\mathrm{PivInd}\, R_j$, is the maximal row index i with $R_{i,j} \neq 0$, taken to be 0 if all entries are 0. Otherwise, the non-zero entry is called the *pivot entry*, denoted by $\mathrm{PivEnt}\, R_i$. We denote the collection of all non-zero column pivots by $\mathrm{PivInds}\, R = \{i \mid i = \mathrm{PivInd}\, R_j \neq 0\}$. Moreover, we call a column R_j *reduced* if its pivot cannot be decreased by adding a linear combination of the columns R_i with $i < j$, and we call the matrix R reduced if all its columns are reduced. Finally, we call the matrix R *totally reduced* if for each $i < j$ we have $R_{s,j} = 0$, where $s = \mathrm{PivInd}\, R_i$.

We call a matrix S a *reduction matrix* if it is a full rank upper triangular matrix such that $R = D \cdot S$ is reduced and S is *homogeneous*, meaning that respects the degrees in the chain complex. Any such reduction $R = D \cdot S$ of the filtration boundary matrix

Input: $D = \partial$ an $l \times l$ filtration boundary matrix
Result: $R = D \cdot S$ with R reduced and S full rank upper triangular
$R = D$; $S = \mathrm{Id}$;
for $j = 1$ **to** l **do**
  while *there exists* $i < j$ *with* $\operatorname{PivInd} R_i = \operatorname{PivInd} R_j > 0$ **do**
    $\mu = -\operatorname{PivEnt} R_j / \operatorname{PivEnt} R_i$;
    $R_j = R_j + \mu \cdot R_i$;
    $S_j = S_j + \mu \cdot S_i$;
return R, S

Algorithm 1: Standard matrix reduction

induces a direct sum decomposition (see, e.g., [Bau21; SMV11]) of C_* into elementary chain complexes in the following way: If $R_j \neq 0$, then we have the summand

$$\cdots \to 0 \to \langle S_j \rangle \xrightarrow{\partial} \langle R_j \rangle \to 0 \to \dots \ ,$$

in which case we call j a *death index*, $i = \operatorname{PivInd} R_j$ a *birth index*, and (i, j) an *index persistence pair*. If $R_i = 0$ and $i \notin \operatorname{PivInds} R$, then we have the summand

$$\cdots \to 0 \to \langle S_i \rangle \to 0 \to \dots \ ,$$

in which case we call i an *essential index*. Moreover, we call the element σ_i a *birth*, *death*, or *essential element*, if its index is a birth, death, or essential index. Similarly, we call a pair of elements (σ_i, σ_j) a *persistence pair*, if the pair (i, j) is an index persistence pair. Note that this is independent of the specific reduction of the filtration boundary matrix. By taking the intersection with the filtration, one obtains elementary filtered chain complexes, in which R_j is a cycle appearing in the filtration at index $i = \operatorname{PivInd} R_j$ and becoming a boundary when S_j enters the filtration at index j, and in which an essential cycle S_i enters the filtration at index i. Thus, the barcode of the persistent homology of the elementwise filtration is given by the collection of intervals

$$\{[i, j) \mid (i, j) \text{ index persistence pair}\} \cup \{[i, \infty) \mid i \text{ essential index}\}.$$

Such a reduction $R = D \cdot S$ of the filtration boundary matrix can be computed by a variant of Gaussian elimination [CEM06], as in Algorithm 1. A slight modification is Algorithm 2, which computes a totally reduced filtration boundary matrix, as used in Corollary D. This is also known as *exhaustive reduction*, and appears in various forms in the literature [EÖ20; CLV22; EZ03; ZC05].

Apparent Pairs Many optimization schemes have been developed in order to speed up the computation of persistent homology. One of them is based on *apparent pairs* [Bau21], a concept which lies at the interface of persistence and discrete Morse theory.

Definition 2.23. Let $(C_*, \Sigma_* = \sigma_1 < \cdots < \sigma_l)$ be a based chain complex with an elementwise filtration. We call a pair of basis elements (σ_i, σ_j) an *apparent pair* if σ_i is the maximal facet of σ_j and σ_j is the minimal cofacet of σ_i.

Input: $D = \partial$ an $l \times l$ filtration boundary matrix
Result: $R = D \cdot S$ with R totally reduced and S full rank upper triangular
$R = D$; $S = \mathrm{Id}$;
for $j = 1$ **to** l **do**
  while *there exist* $s, i < j$ *with* $\operatorname{PivInd} R_i = s$ *and* $R_{s,j} \neq 0$ **do**
    $\mu = -R_{s,j}/R_{s,i}$;
    $R_j = R_j + \mu \cdot R_i$;
    $S_j = S_j + \mu \cdot S_i$;
return R, S

Algorithm 2: Exhaustive matrix reduction

In the context of persistence, the interest in apparent pairs stems from the following observation [Bau21], which is immediate from the definitions.

Lemma 2.24. *For any apparent pair (σ, τ) of an elementwise filtration, the column of τ in the filtration boundary matrix is reduced, and (σ, τ) is a persistence pair.*

Lexicographic Optimality We now introduce the lexicographic order on chains for a based chain complex $(C_*, \Sigma_* = \sigma_1 < \cdots < \sigma_l)$ with an elementwise filtration, extending the definitions in [CLV22]. For any chain $c = \sum_i \lambda_i \sigma_i \in C_n$, we define its *support* $\operatorname{supp}_{\Sigma_*} c$ to be the set of basis elements $\sigma_i \in \Sigma_n$ with $\lambda_i \neq 0$. Note that this is not to be confused with the supporting subcomplex.

Given a totally ordered set $(X, \leq)$, we consider the lexicographic order $\preceq$ on the power set 2^X given by identifying any subset $A \subseteq X$ with its characteristic function and considering the lexicographic order on the set of characteristic functions. Explicitly, for $A, B \subseteq X$ we have $A \preceq B$ if and only if $A = B$ or the maximal element of the symmetric difference $(A \setminus B) \cup (B \setminus A)$ is contained in B.

Definition 2.25. The *lexicographic preorder* $\sqsubseteq$ on the collection of chains C_n is given by $c_1 \sqsubseteq c_2$ if and only if $\operatorname{supp}_{\Sigma_*} c_1 \preceq \operatorname{supp}_{\Sigma_*} c_2$ in the lexicographic order on subsets of Σ_n. We write $\sqsubset$ for the corresponding strict preorder.

If we consider a simplicial chain complex with coefficients in $\mathbb{Z}/2\mathbb{Z}$, then this preorder is a total order, and it coincides with the one considered in [CLV22, Definition 2.1].

Definition 2.26. We call a chain $c \in C_n$ *lexicographically minimal*, or *irreducible*, if there exists no strictly smaller homologous chain $(c + \partial e) \sqsubset c$ in the lexicographic preorder, where $e \in C_{n+1}$. Otherwise, we call the chain c *reducible*.

Remark 2.27. It follows directly from Proposition 4.18 that each homology class of the chain complex C_* has a unique lexicographically minimal representative cycle, even for coefficient fields different from $\mathbb{Z}/2\mathbb{Z}$.

2.3. Discrete Morse Theory

In this section, we recall some basic notions from combinatorial topology and discrete Morse theory, which was introduced by Forman [For98; For02] and that is a central tool to study geometric complexes. It is also closely related to Whitehead's simple homotopy theory [Whi50; Coh73]. Some of the presented results have novel aspects, most notably Propositions 2.34, 2.37, 2.43 and 2.55.

We use the following generalization of a discrete Morse function, originally due to Forman [For98; BE17; Fre09].

Definition 2.28. Let K be a finite simplicial complex. A function $f\colon K \to \mathbb{R}$ is a *generalized discrete Morse function* if

- f is monotonic, i.e., for any $\sigma, \tau \in K$ with $\sigma \subseteq \tau$ we have $f(\sigma) \leq f(\tau)$, and
- there exists a (unique) partition $\hat{V}$ of K into intervals $[\rho, \phi] = \{\psi \in K \mid \rho \subseteq \psi \subseteq \phi\}$ in the face poset such that any pair of simplices $\sigma \subseteq \tau$ satisfies $f(\sigma) = f(\tau)$ if and only if σ and τ belong to a common interval in the partition.

The collection of *regular* intervals, $[\rho, \phi]$ with $\rho \neq \phi$, is the *discrete gradient* V of f on K, and any singleton interval $[\sigma, \sigma] = \{\sigma\}$, as well as the corresponding simplex σ, is *critical*.

If W is another discrete gradient on K, then we say that V is a *refinement* of W if each interval in the gradient partition $\hat{W}$ is a disjoint union of intervals in $\hat{V}$. If the refinement preserves the set of critical simplices, we call it a *regular refinement*. Moreover, if each regular interval in V only consists of a pair of simplices, we simply call f a *discrete Morse function*. We often refer to a discrete gradient without explicit mention of the function f, noting that different functions can have the same gradient.

Recall that for a monotonic function $g\colon K \to \mathbb{R}$ we write $S_r(g) = g^{-1}(-\infty, r] \subseteq K$ for the sublevel set and $S_r^{<}(g) = g^{-1}(-\infty, r) \subseteq K$ for the open sublevel set of g at scale $r \in \mathbb{R}$. By the following result from [BE17], the Čech as well as the Delaunay radius function (Definition 2.3), for any finite subset of $\mathbb{R}^d$ in general position (Definition 2.4), are both generalized discrete Morse functions.

Proposition 2.29 (Bauer, Edelsbrunner [BE17])**.** *Let $X \subset \mathbb{R}^d$ be a finite subset in general position. The Čech radius function $r_\emptyset\colon$ Čech$(X) \to \mathbb{R}$ is a generalized discrete Morse function with discrete gradient given by the regular intervals in*

$$\{[\operatorname{On} S, \operatorname{Incl} S] \mid \sigma \in \text{Čech}(X),\ S = S(\sigma, \emptyset)\},$$

where $S(\sigma, \emptyset)$ is the smallest enclosing sphere of σ. Moreover, the Delaunay radius function $r_X\colon \operatorname{Del}(X) \to \mathbb{R}$ is also a generalized discrete Morse function with discrete gradient given by the regular intervals in

$$\{[\operatorname{Front} S, \operatorname{On} S] \mid \sigma \in \operatorname{Del}(X),\ S = S(\sigma, X)\},$$

where $S(\sigma, X)$ is the smallest empty circumsphere of σ. In both cases, the critical simplices are the centered Delaunay simplices*, i.e., the simplices σ such that the center of the smallest empty circumsphere of σ is contained in the convex hull of σ.*

We show in Section 3.4.1 that for generic tree metrics, the diameter function, whose sublevel sets correspond to the Vietoris–Rips complexes of the metric space, is also a generalized discrete Morse function.

An *elementary collapse* $K \searrow K \setminus \{\sigma, \tau\}$ is the removal of a pair of simplices (σ, τ), where σ is a *free facet* of τ, i.e., τ the unique proper coface of σ. A *collapse* $K \searrow L$ onto a subcomplex L is a finite sequence of elementary collapses starting in K and ending in L; see Section 2.3.3 for an extension to infinite complexes. An elementary collapse can be realized continuously by a strong deformation retraction and therefore collapses preserve the homotopy type. A discrete gradient can induce a collapse.

Proposition 2.30 (Forman [For98]; see also [Koz20, Theorem 10.9])**.** *Let K be a finite simplicial complex and let $L \subseteq K$ be a subcomplex. Assume that V is a discrete gradient on K such that the complement $K \setminus L$ is the union of intervals in V. Then there exists a collapse $K \searrow L$. In particular, the inclusion $|L| \hookrightarrow |K|$ is a homotopy equivalence.*

Example 2.31. Let K be a finite simplicial complex and let $p \notin \operatorname{Vert} K$. The *cone on K with apex p* is the simplicial complex

$$\operatorname{Cone}_p K = K \cup \{\tau \cup \{p\} \mid \tau \in K\} \cup \{\{p\}\}.$$

The monotonic function

$$f\colon \operatorname{Cone}_p K \to \mathbb{R}, \ \sigma \mapsto \dim(\sigma \cup \{p\})$$

is a discrete Morse function with discrete gradient $V = \{(\tau, \tau \cup \{p\}) \mid \tau \in K\}$. The only critical simplex is $\{p\}$ and thus, by Proposition 2.30, there exists a collapse $\operatorname{Cone}_p K \searrow \{p\}$.

2.3.1. Discrete Morse Theory and Apparent Pairs

We explain the connection between discrete Morse theory and apparent pairs (Definition 2.23). Let $f\colon K \to \mathbb{R}$ be a monotonic function defined on a finite simplicial complex, and assume that the vertices of K are totally ordered. The *f-lexicographic order* is the total order $\leq_f$ on K given by ordering the simplices

- by their value under f,
- then by dimension,
- then by the lexicographic order induced by the total vertex order.

Note that the f-lexicographic order induces a simplexwise filtration of K, and, recalling Remark 2.22, every simplexwise filtration of K induces a canonical elementwise filtration of the simplicial chain complex $C_*(K)$. A persistence pair (σ, τ) of this simplexwise filtration is a *zero persistence pair* if $f(\sigma) = f(\tau)$. The definition of apparent pairs (Definition 2.23) specializes to the simplicial setting as follows: an *apparent pair* of a simplexwise filtration of K is a pair of simplices (σ, τ) in K, such that σ is the maximal facet of τ and τ is the minimal cofacet of σ, with respect to the simplexwise filtration. The following [Bau21, Lemma 3.5] is essentially a consequence of the definitions.

Lemma 2.32. *The collection of apparent pairs of a simplexwise filtration is a discrete gradient, the* apparent pairs gradient*.*

Remark 2.33. Variants of the observations from Example 2.31 are used in the proofs of Theorems 3.14 and 3.16 and Theorem B. If the vertices of $\operatorname{Cone}_p K$ are totally ordered such that $p < v$ for every $v \in \operatorname{Vert} K$ and $g\colon \operatorname{Cone}_p K \to \mathbb{R}$ is any monotonic function with $g(\sigma) = g(\sigma \cup \{p\})$, then the collection of zero persistence apparent pairs of the g-lexicographic order on $\operatorname{Cone}_p K$ is precisely the discrete gradient V from Example 2.31. In particular, if X is a finite metric space and for some parameters $t > u > 0$ there is a point $p \in X$ with $d(p, x) \leq u$ for every $x \in X$, then the apparent pairs gradient of the diam-lexicographic order on the Vietoris–Rips complex $\operatorname{Rips}_t(X)$ induces the collapses $\operatorname{Rips}_t(X) \searrow \operatorname{Rips}_u(X) \searrow \{p\}$ if, for example, we choose a total order on X such that $x < y$ implies $d(x, p) \leq d(y, p)$.

There is a further connection between apparent pairs and discrete Morse functions: Let $f\colon K \to \mathbb{R}$ be a generalized discrete Morse function, defined on a finite simplicial complex, with discrete gradient V: We can refine every regular interval $[\rho, \phi] \in V$ towards an arbitrary vertex $x \in \phi \setminus \rho$ by partitioning $[\rho, \phi]$ into the pairs $(\psi \setminus \{x\}, \psi \cup \{x\})$ for all $\psi \in [\rho, \phi]$. It is easy to see that this yields a discrete gradient, called a *vertex refinement* of V, with the same set of critical simplices [BE17; Fre09], meaning that it is a regular refinement of V. If the vertices are chosen minimally according to the total order on $\operatorname{Vert} K$, we call the discrete gradient

$$\widetilde{V} = \{(\psi \setminus \{v\}, \psi \cup \{v\}) \mid \psi \in [\rho, \phi] \in V,\ v = \min(\phi \setminus \rho)\}$$

the *minimal vertex refinement of V*.

By the following proposition, this regular refinement is induced by the apparent pairs gradient for the simplexwise filtration determined by the f-lexicographic order. In particular, the zero persistence pairs of this simplexwise filtration are precisely the zero persistence apparent pairs, and the V-critical simplices of K are precisely the simplices that are either essential or contained in a non-zero persistence pair.

Proposition 2.34. *The zero persistence apparent pairs with respect to the f-lexicographic order are precisely the gradient pairs of the minimal vertex refinement $\widetilde{V}$.*

Proof. Let (σ, τ) be a zero persistence apparent pair. Then $f(\sigma) = f(\tau)$, and σ and τ are contained in the same regular interval $I = [\rho, \phi]$ of V. Let v be the minimal vertex in $\phi \setminus \rho$. By assumption, σ is the maximal facet of τ, and τ is the minimal cofacet of σ. Hence, σ is lexicographically maximal among all facets of τ in I, and τ is lexicographically minimal under all cofacets of σ in I. By the assumption that (σ, τ) forms an apparent pair, we cannot have $v \in \sigma$, as otherwise $\tau \setminus \{v\}$ would be a larger facet of τ than σ. Similarly, we cannot have $v \notin \tau$, as otherwise $\sigma \cup \{v\}$ would be a smaller cofacet of σ than τ. This means that $\tau = \sigma \cup \{v\}$ and therefore $(\sigma, \tau) \in \widetilde{V}$.

Conversely, assume that $(\sigma, \tau) \in \widetilde{V}$ holds. Consider the interval $I = [\rho, \phi]$ of V with $(\sigma, \tau) \subseteq I$ and let v be the minimal vertex in $\phi \setminus \rho$. By construction of $\widetilde{V}$, $\sigma = \tau \setminus \{v\}$ is the lexicographically maximal facet of τ in I and $\tau = \sigma \cup \{v\}$ is the lexicographically minimal cofacet of σ in I. Therefore, (σ, τ) is a zero persistence apparent pair. □

2.3.2. Descending Complexes and Gradient Refinements

Motivated by Proposition 2.34 and with a view towards Section 4, we extend the definition of the *Wrap complex* from [BE17] to an arbitrary subset C of the set of critical simplices of a discrete gradient V on a finite simplicial complex K, and study its behavior under gradient refinements. Moreover, we also extend it to monotonic functions $g\colon K \to \mathbb{R}$ that are *compatible* with V, i.e., such that for any pair of simplices $\sigma \subseteq \tau$ that belong to a common interval $I \in \hat{V}$ we have $g(\sigma) = g(\tau) =: g(I) \in \mathbb{R}$.

The gradient partition $\hat{V}$ has a canonical poset structure $\leq_{\hat{V}}$ given by the transitive closure of the relation

$$I \sim J \text{ if and only if there exists a face } \sigma \in I \text{ of a simplex } \tau \in J.$$

The *lower set* of a subset $A \subseteq \hat{V}$ is the set of intervals $\downarrow A = \{I \in \hat{V} \mid \exists J \in A : I \leq_{\hat{V}} J\}$, and for $r \in \mathbb{R}$ we denote the discrete gradient V restricted to the sublevel set $S_r(g)$ by

$$V_r = \{I \in V \mid g(I) \leq r\}.$$

Note that if $I \leq_{\hat{V}} J$, then $g(I) \leq g(J)$, and hence for a subset $A \subseteq \hat{V_r} \subseteq \hat{V}$ the lower sets with respect to the canonical poset structure on $\hat{V_r}$ and with respect to the canonical poset structure on $\hat{V}$ coincide.

Definition 2.35. Let V be a discrete gradient on a finite simplicial complex K. The *descending complex* $D(V)$ is the subcomplex

$$D(V) = \bigcup \downarrow \{\{\sigma\} \mid \sigma \in K \text{ critical}\}$$

of K given by the union of intervals in the lower sets of the critical intervals. More generally, if C is a subset of the set of critical simplices, the *descending complex* $D(V,C)$ is the subcomplex

$$D(V,C) = \bigcup \downarrow \{\{\sigma\} \mid \sigma \in C\}$$

of K. Moreover, for a monotonic function $g\colon K \to \mathbb{R}$ that is compatible with V, the *descending complex* $D_r(V,g)$ at scale $r \in \mathbb{R}$ is the subcomplex

$$D_r(V,g) = D(V_r) = D(V, \mathrm{Crit}_r(V,g)) = \bigcup \downarrow \{\{\sigma\} \mid \sigma \in K \text{ critical}, \, g(\sigma) \leq r\}$$

of $S_r(g)$, where $\mathrm{Crit}_r(V,g) = \{\sigma \in K \text{ critical} \mid g(\sigma) \leq r\}$. If V is the discrete gradient of a generalized discrete Morse function f, we simply write $D_r(f)$ for $D_r(V,f)$.

See Fig. 8 and Example 2.36. The descending complex $D(V)$ in the context of discrete Morse theory is motivated by the concept of a *descending* or *stable manifold* of a critical point from smooth Morse theory, which is central in the original definition of the Wrap complex [Ede03]. Note that the descending complex $D_r(r_X) \subseteq \mathrm{Del}_r(X)$ of the Delaunay radius function r_X (see Definition 2.3 and Proposition 2.29) is precisely the *Wrap complex*, $\mathrm{Wrap}_r(X)$, from [BE17]; see Fig. 2 for an illustration. Finally, it has been shown that a variant of the Wrap complex can be used to reconstruct a submanifold up to homotopy type or even homeomorphism by choosing a suitable subset of critical simplices [Dey+05; RS07; Sad09], which also motivates our definition of $D(V,C)$.

Example 2.36. For a generic tree metric space X, as in Section 3.4.1, the diameter function diam: $\mathrm{Cl}(X) \to \mathbb{R}$ is a generalized discrete Morse function. The critical simplices are the vertices X and the tree edges E. Thus, for any $r \geq 0$, the descending complex $D_r(\mathrm{diam})$ is the subforest T_r with vertices X and all edges of E with length at most r.

We now study the behavior of the descending complex under gradient refinements.

Proposition 2.37. *Let V be a discrete gradient on a finite simplicial complex K, let C be a subset of V-critical simplices, and let W be a refinement of V. Then the descending complex $D(W, C)$ is a subcomplex of the descending complex $D(V, C)$.*

Proof. Note first that every V-critical simplex is also W-critical, as W is a refinement of V. By the same reason, there exists a set map $\varphi\colon \hat{W} \to \hat{V}$ between the gradient partitions such that for every $B \in \hat{W}$ we have $B \subseteq \varphi(B)$. It follows straightforwardly from the definition of the canonical poset structures on $\hat{W}$ and $\hat{V}$ that φ is a poset map. Thus, for every W-critical simplex $\sigma \in C$ and interval $A \in \hat{W}$ with $A \leq_{\hat{W}} \{\sigma\}$, we have $\varphi(A) \leq_{\hat{V}} \varphi(\{\sigma\}) = \{\sigma\} \in \hat{V}$, as σ is also V-critical. It now follows directly from the construction of the descending complexes, that $A \subseteq \varphi(A) \subseteq D(V, C)$ and $D(W, C) \subseteq D(V, C)$. □

Remark 2.38. Let V be a discrete gradient on a finite simplicial complex K. Recall from Proposition 2.30, that if L is a subcomplex of K, and the complement $K \setminus L$ is the disjoint union of regular intervals in V, then V induces a collapse $K \searrow L$. It follows directly from this and the construction of $D(V)$, that $D(V)$ is the smallest subcomplex of K such that V induces a collapse $K \searrow D(V)$. Moreover, if W is a regular refinement of V, which implies $D(W) \subseteq D(V)$ by Proposition 2.37, the complement $D(V) \setminus D(W)$ is the disjoint union of regular intervals in W. Similar to before, it follows that W induces a collapse $D(V) \searrow D(W)$. In particular, the inclusion $|D(W)| \hookrightarrow |D(V)|$ is a homotopy equivalence. See Fig. 8 for an example.

Figure 8: Left: Generalized discrete gradient (blue) with corresponding descending complex (green). Right: lexicographic gradient refinement (blue) with corresponding descending complex (green).

Recalling Definition 2.35, and as regular refinements preserve the set of critical simplices, Proposition 2.37 directly implies the following.

Corollary 2.39. *Let $g\colon K \to \mathbb{R}$ be a monotonic function, defined on a finite simplicial complex, that is compatible with the discrete gradients V and W on K, where V is a regular refinement of W. Then for any $r \in \mathbb{R}$ the descending complex $D_r(V, g)$ is a subcomplex of the descending complex $D_r(W, g)$.*

2.3.3. Discrete Morse Theory for Infinite Complexes

For the proofs of Theorems 3.14 and 3.16, which imply Theorem A, and Lemma 5.13, which is used in the proof of Theorem 5.14, we use a variant of Proposition 2.30 for infinite complexes that we describe now. For an extension to infinite cell complexes that is similar to the results presented here, see [Bat02].

First, we extend the notion of a discrete gradient to infinite complexes. For finite simplicial complexes, the following definitions are equivalent to the ones above that are given in terms of (non-generalized) discrete Morse functions [For02]. Let K be a (not necessarily finite) simplicial complex. By a *discrete vector field* on K we mean a partition $\hat{V}$ of K into singletons $\{\sigma\}$, σ is then called a *critical simplex*, and pairs (σ, τ) with σ a facet of τ. Equivalently, a discrete vector field is given by a matching in the Hasse diagram $\mathcal{H}(K)$ of the face poset of K, i.e., the directed graph whose nodes are the simplices and whose arcs are the pairs (σ, τ) with σ a facet of τ. We call the collection V of *regular* pairs, $(\sigma, \tau) \in \hat{V}$ with $\sigma \neq \tau$, a *discrete gradient* if the corresponding matching in the Hasse diagram $\mathcal{H}(K)$ is acyclic, i.e., the directed graph $\mathcal{H}(K, V)$, that is obtained from $\mathcal{H}(K)$ by reversing all the arcs (σ, τ) for which $(\sigma, \tau) \in V$, is acyclic. Note that it suffices to check that there are no non-trivial closed *gradient paths*, also called *V-paths*, i.e., that there are no undirected paths in $\mathcal{H}(K)$ that are of the form $\sigma_0 \to \tau_0 \leftarrow \cdots \to \tau_{n-1} \leftarrow \sigma_n$ with $(\sigma_i, \tau_i) \in V$, $\sigma_i \neq \sigma_{i+1}$ and $\sigma_0 = \sigma_n$.

For any such discrete gradient V on a simplicial complex K, there is, as in Section 2.3.2, a canonical poset structure on the gradient partition $\hat{V}$: For two elements $A, B \in \hat{V}$ we have $A \leq_{\hat{V}} B$ if and only if there exists a sequence $A = C_0, C_1, \ldots, C_n = B$ in $\hat{V}$ such that for every i there exist elements $x_{i-1} \in C_{i-1}, x_i \in C_i$ with x_{i-1} a face of x_i.

Moreover, for any element $A \in \hat{V}$ we define its *height* to be

$$\mathrm{ht}(A) = \sup\{n \in \mathbb{N} \mid \exists\ A = B_0 > \cdots > B_n \text{ in } \hat{V}\}.$$

The following lemma is useful in practice.

Lemma 2.40. *The height* $\mathrm{ht}(A)$ *is finite for every* $A \in \hat{V}$ *if and only if for every simplex* $\sigma \in K$ *its V-path height*

$$\mathrm{ht}_V(\sigma) = \sup\{n \in \mathbb{N} \mid \exists\ V\text{-path } \sigma = \sigma_0 \to \tau_0 \leftarrow \cdots \to \tau_{n-1} \leftarrow \sigma_n\}.$$

is finite.

Proof. Every V-path $\sigma_0 \to \tau_0 \leftarrow \cdots \to \tau_{n-1} \leftarrow \sigma_n$, with $(\sigma_i, \tau_i) \in V$, induces a descending chain $(\sigma_0, \tau_0) > \cdots > (\sigma_{n-1}, \tau_{n-1})$. Hence, if the height is finite for every $A \in \hat{V}$ this implies that the V-path height is finite for every simplex in K.

For the converse, we employ induction over the dimension $\dim A = \dim \min A$ of an element $A \in \hat{V}$. If $\dim A = 0$, then $\mathrm{ht}(A) = \mathrm{ht}_V(\min A)$ and this is finite by assumption. For the induction step, consider the set $F_{\min A}$ of all V-paths starting in $\min A$. For such a gradient path $\gamma = \sigma_0 \to \tau_0 \leftarrow \cdots \to \tau_{n-1} \leftarrow \sigma_n$ write $\mathrm{end}\, \gamma = \sigma_n$ and $\mathrm{length}\, \gamma = n$. Then we can bound $\mathrm{ht}(A)$ from above as follows:

$$\mathrm{ht}(A) \leq \max\{\mathrm{length}\, \gamma + 1 + \mathrm{ht}(B) \mid \gamma \in F_{\min A},\ \sigma \subsetneq \mathrm{end}\, \gamma,\ \sigma \in B \in \hat{V}\}.$$

To complete the induction step, note that for every B as above we have $\dim B < \dim A$ and hence $\mathrm{ht}(B)$ is finite by the induction assumption. Thus, it suffices to show that the set $F_{\min A}$ is finite. This can be seen as follows: Given a gradient path γ as above, then the path ends in σ_i or there are $\dim \tau_i = \dim \sigma_0 + 1$ choices for σ_{i+1} once σ_i is fixed. Hence, the cardinality of $F_{\min A}$ is bounded from above by $(\dim(\min A)+2)^{\mathrm{ht}_V(\min A)}$. □

Remark 2.41. The last part of the proof of Lemma 2.40 shows that if the V-path height of a simplex $\sigma \in K$ is finite, then there are only finitely many gradient paths starting in σ. In particular, if the V-path height of every simplex is finite, then there are only finitely many gradient paths between any two simplices.

Remark 2.42. If the V-path height of a simplex $\sigma \in K$ is finite, then there does not exist an infinite gradient path starting in σ, i.e., there does not exist for every $i \in \mathbb{N}$ a gradient pair $(\sigma_i, \tau_i) \in V$ such that for all $n \in \mathbb{N}$ those assemble to a gradient path $\sigma = \sigma_0 \to \tau_0 \leftarrow \cdots \to \tau_{n-1} \leftarrow \sigma_n$. The converse is also true: We show the contraposition of the statement, i.e., we show that if $\mathrm{ht}_V(\sigma) = \infty$, then there exists an infinite gradient path starting in σ. First note that, given any gradient path $\sigma_0 \to \tau_0 \leftarrow \cdots \to \tau_{n-1} \leftarrow \sigma_n$, then the path ends in σ_i or there are $\dim \tau_i = \dim \sigma_0 + 1$ choices for σ_{i+1} once σ_i is fixed. As by assumption $\mathrm{ht}_V(\sigma) = \infty$, there exists a sequence $(\gamma_k)_k$ of gradient paths starting in σ, such that the length of γ_k is at least k. Since there are only finitely many choices in each step of any such γ_k, there exists, by a variant of the pigeonhole principle, for every $i \in \mathbb{N}$ a gradient pair $(\sigma_i, \tau_i) \in V$, with $\sigma_0 = \sigma$, that assemble for every $n \in \mathbb{N}$ to a gradient path $\sigma = \sigma_0 \to \tau_0 \leftarrow \cdots \to \tau_{n-1} \leftarrow \sigma_n$. Thus, there exists an infinite gradient path starting in σ, proving the claim.

The essential ideas for the proof of the following proposition, which we provide further below and that extends Proposition 2.30 to infinite complexes, can be found already in the proof of [Bro92, Proposition 1], which predates Forman's papers.

Proposition 2.43. *Let $L \subseteq K$ be a pair of simplicial complexes and let V be a discrete gradient on K such that for every element $A \in V$ its height $\mathrm{ht}(A)$ is finite. Moreover, assume that $K \setminus L$ is the union of pairs in V. Then the inclusion $|L| \hookrightarrow |K|$ is a homotopy equivalence.*

In light of Proposition 2.30, we extend the terminology to say that the discrete gradient in Proposition 2.43 induces a *collapse* $K \searrow L$.

Remark 2.44. Note that the homotopy equivalence $|L| \hookrightarrow |K|$ is not necessarily induced by a finite sequence of elementary collapses. However, a close inspection of the proof shows that it is a (possibly infinite) concatenation of strong deformation retracts that are obtained by simultaneously executing a family of elementary collapses induced by the discrete gradient V. This justifies the extension of terminology above.

Remark 2.45. The finite height assumption in Proposition 2.43 is crucial. Consider the following infinite simplicial complex K with the indicated discrete gradient V:

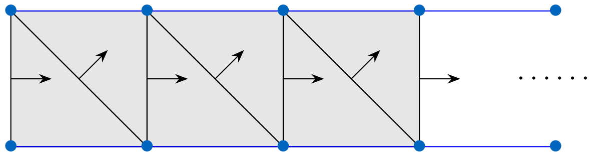

Let L be the subcomplex consisting of the simplices indicated in blue, i.e., all horizontal 1-simplices and their vertices. The complement $K \setminus L$ is the union of pairs in V, but the inclusion $|L| \hookrightarrow |K|$ is not a homotopy equivalence. Note that every gradient facet has infinite V-path height, and thus, no element of V has finite height.

Before we prove Proposition 2.43 we need one small lemma.

Lemma 2.46. *Let K be a simplicial complex and $K_0 \subseteq K_1 \subseteq \cdots \subseteq K$ a filtration of subcomplexes whose union is K such that for all $i \in \mathbb{N}$ the inclusion $|K_i| \hookrightarrow |K_{i+1}|$ is a homotopy equivalence. Then the inclusion $|K_0| \hookrightarrow |K|$ is also a homotopy equivalence.*

Proof. By Whitehead's theorem [Hat02, Theorem 4.5], it suffices to show that for all $n \in \mathbb{N}$ the induced morphism on homotopy groups $g_n \colon \pi_n(|K_0|) \to \pi_n(|K|)$ is an isomorphism. For any map $f \colon S^n \to |K|$ its image is compact and hence contained in some $|K_i|$. As $|K_0| \to |K_i|$ is a homotopy equivalence, it follows that the homotopy class $[f]$ is in the image of the composite $\pi_n(|K_0|) \to \pi_n(|K_i|) \to \pi_n(|K|)$ and hence also in the image of g_n. This shows surjectivity. A similar argument applied to any homotopy $h \colon S^n \times [0,1] \to |K|$ shows that g_n is injective. This proves the lemma. □

Proof of Proposition 2.43. Without loss of generality, we can assume that L is the union of critical simplices. Consider the filtration $L = K_0 \subseteq K_1 \subseteq \cdots \subseteq K$ of K, where K_i is the subcomplex

$$K_i = L \cup \bigcup_{A \in V,\ \mathrm{ht}(A) \leq i} A.$$

We show that for every $i \in \mathbb{N}$ the inclusion $|K_i| \hookrightarrow |K_{i+1}|$ is a homotopy equivalence and the proposition then follows from Lemma 2.46: Let $(\sigma, \tau) = A \in V$ be any element, with σ a facet of τ, such that $\mathrm{ht}(A) = i + 1$. Then, σ is a free facet of τ in K_{i+1}, as otherwise there would exist a pair $B \in V$ with $\mathrm{ht}(B) \leq i + 1$ and $B > A$. But this cannot be true, because then the last property implies that the height of B satisfies $\mathrm{ht}(B) \geq \mathrm{ht}(A) + 1 = i + 2$, contradicting the construction. A similar argument shows that τ is not properly contained in any simplex of K_{i+1}. Therefore, the complement $K_{i+1} \setminus K_i$ is partitioned by pairs in V of height $i + 1$ and the corresponding simplices to different pairs can only possibly intersect in the subcomplex K_i. Thus, executing the elementary collapses that are encoded by those pairs simultaneously induces a strong deformation retract

$$|K_{i+1}| \to |K_{i+1} \setminus \bigcup_{A \in V,\ \mathrm{ht}(A) = i+1} A| = |K_i|$$

and hence the inclusion $|K_i| \hookrightarrow |K_{i+1}|$ is a homotopy equivalence. □

We now describe how to construct discrete gradients from others. The following proposition follows directly from Kozlov [Koz20, Theorem 16.8]; variants of this proposition also appear in Hersh [Her05, Lemma 4.1] and Johnsson [Jon08, Lemma 4.2].

Proposition 2.47. *Let P be the face poset of a simplicial complex K, and let Q be any poset. Let $\varphi\colon P \to Q$ be a map of posets, and for all $q \in Q$ let V_q be an acyclic matching in the Hasse diagram of $\varphi^{-1}(q)$. Then the union $V = \bigcup_q V_q$ of these matchings is itself an acyclic matching in the Hasse diagram of P, i.e., a discrete gradient on K.*

Remark 2.48. Recall that a (non-generalized) discrete Morse function $f\colon K \to \mathbb{R}$ is a poset map from the face poset of K to the poset of real numbers, such that every fiber is isomorphic to a disjoint union of posets with one element, which are of the form $\{0\}$, and posets with two elements, which are of the form $\{0 < 1\}$. Similarly, every fiber of a generalized discrete Morse function is isomorphic to a disjoint union of intervals in the face poset of K, and acyclic matchings in the Hasse diagrams of those correspond to gradient refinements. Thus, in principle, we can think of the poset map in Proposition 2.47 as being a vast generalization of a discrete Morse function.

2.3.4. Algebraic Morse Theory and Gradient Flows

We now recall some aspects of algebraic Morse theory [Skö06; JW09; Koz08], also called algebraic discrete Morse theory, with a view towards its use in Section 4.1. Let (C_*, Σ_*) be a based chain complex, as defined in Section 2.2.

Definition 2.49. A function $f\colon \Sigma_* \to \mathbb{R}$ is an *algebraic Morse function* if

- f is monotonic, i.e., for any facet $\sigma \in \Sigma_n$ of $\tau \in \Sigma_{n+1}$ we have $f(\sigma) \leq f(\tau)$, and
- there exists a (unique) disjoint collection V of facet pairs such that every facet pair (μ, η) satisfies $f(\mu) = f(\eta)$ if and only if $(\mu, \eta) \in V$.

We call V the *algebraic gradient* of f on Σ_*, and we call a basis element *critical* if it is not contained in any pair of V. Moreover, for $(\sigma, \tau) \in V$ we call σ a *gradient facet* and τ a *gradient cofacet*.

We often refer to an algebraic gradient without explicit mention of the associated algebraic Morse function.

Remark 2.50. The definitions of algebraic Morse function and algebraic gradient generalize those from discrete Morse theory. Let $f\colon K \to \mathbb{R}$ be a (non-generalized) discrete Morse function, defined on a finite simplicial complex, with discrete gradient V. Recall that the simplicial chain complex $C_*(K)$ has a basis Σ_* consisting of the simplices of K with some chosen orientation. We can now interpret f as an algebraic Morse function on this basis Σ_* and the discrete gradient V as an algebraic gradient.

Gradient Flows We now introduce the flow determined by an algebraic gradient, and study its behavior under gradient containment. In Section 4 we then analyze its relation to lexicographically minimal cycles (Proposition 4.16), as well as the descending complex (Proposition 4.24). While Section 4 is mainly focused on cycles, in this section we present results that hold more generally for chains. Moreover, in Section 4.1.4 we demonstrate how the gradient flow on a cycle can be interpreted as a variant of Gaussian elimination, tying it closely to the exhaustive reduction (Algorithm 2).

Let (C_*, Σ_*) be a based chain complex and V an algebraic gradient on Σ_*. The following definition, originally for discrete gradients [For98], carries over naturally to the algebraic setting.

Definition 2.51. The *flow* $\Phi\colon C_* \to C_*$ determined by V is the chain map given by

$$\Phi(c) = c + \partial\,\mathrm{F}(c) + \mathrm{F}(\partial c),$$

where $\mathrm{F}\colon C_* \to C_{*+1}$ is the unique linear map defined on the basis elements $\sigma \in \Sigma_*$ as

$$\mathrm{F}(\sigma) = \begin{cases} -\frac{1}{\langle \partial\tau, \sigma\rangle} \cdot \tau & \text{if } \sigma \text{ is contained in a pair } (\sigma,\tau) \in V, \\ 0 & \text{otherwise.} \end{cases}$$

Note that, by construction, the map F is a chain homotopy between the identity and the flow Φ. In particular, if c is a cycle, then the flow reduces to $\Phi(c) = c + \partial\,\mathrm{F}(c)$ and therefore acts on each homology class of the chain complex by a change of representative cycle. Moreover, if $f\colon \Sigma_* \to \mathbb{R}$ is an algebraic Morse function with algebraic gradient V, then, by construction, the associated flow Φ acts for any $r \in \mathbb{R}$ on the subcomplex of C_* spanned by the sublevel set $S_r(f) = f^{-1}(-\infty, r]$.

Forman [For98] proved that the sequence $\Phi, \Phi^2, \ldots$ stabilizes in the case that the chain complex is the cellular chain complex of a finite CW-complex. This generalizes straightforwardly to our setting of finite chain complexes, and we denote the *stabilized flow* by $\Phi^\infty = \Phi^r$, where r is a large enough natural number:

Proposition 2.52. *Let (C_*, Σ_*) be a based chain complex and V an algebraic gradient on Σ_* with associated flow Φ. There exists an $r \in \mathbb{N}$ such that for all $s \geq r$ we have $\Phi^s = \Phi^r$.*

The proof of this proposition uses the following lemma, which is also a straightforward generalization of the corresponding statement in [For98], and is of independent interest.

Lemma 2.53. *Let $\sigma \in \Sigma_n$ be critical. Then for all $r \in \mathbb{N}$ the iterated flow $\Phi^{r+1}(\sigma)$ is given by*

$$\Phi^{r+1}(\sigma) = \sigma + w + \Phi(w) + \cdots + \Phi^r(w),$$

where $w = \mathrm{F}(\partial\sigma)$. Moreover, we have $\Phi^s(w) \in \operatorname{im}\mathrm{F}$ for all $s \in \mathbb{N}$, and $\mathrm{F}(\Phi^{r+1}(\sigma)) = 0$.

Proof. We prove the first claim by induction over $r \in \mathbb{N}$. To establish the base case, let $r = 0$. As σ is critical, we have $\mathrm{F}(\sigma) = 0$. By definition of the gradient flow, we thus have $\Phi(\sigma) = \sigma + \mathrm{F}(\partial\sigma) = \sigma + w$, establishing the base case. For the induction step, assume the claim holds for some $r \in \mathbb{N}$, and we prove it for $r+1$. Note that

$$\begin{aligned}\Phi^{r+2}(\sigma) = \Phi(\Phi^{r+1}(\sigma)) &= \Phi(\sigma) + \Phi(w) + \Phi^2(w) + \cdots + \Phi^{r+1}(w) \\ &= \sigma + w + \Phi(w) + \cdots + \Phi^{r+1}(w),\end{aligned}$$

where we used the induction hypothesis and the base case. This proves the first claim.

We now show the second claim. By definition, we have $\mathrm{F} \circ \mathrm{F} = 0$, implying

$$\Phi \circ \mathrm{F} = (1 + \partial \circ \mathrm{F} + \mathrm{F} \circ \partial) \circ \mathrm{F} = \mathrm{F} \circ (1 + \partial \circ \mathrm{F}),$$

and therefore $\operatorname{im}(\Phi \circ \mathrm{F}) \subseteq \operatorname{im} \mathrm{F}$. As $w = \mathrm{F}(\partial\sigma) \in \operatorname{im} \mathrm{F}$, the claim $\Phi^s(w) \in \operatorname{im} \mathrm{F}$ now follows directly from the previous observation. Thus, the chain $\Phi^{r+1}(\sigma)$ is the sum of the critical basis element σ and elements in the image of F. Using $\mathrm{F}(\sigma) = 0$ and $\mathrm{F} \circ \mathrm{F} = 0$, this implies $\mathrm{F}(\Phi^{r+1}(\sigma)) = 0$ as claimed. □

Proof of Proposition 2.52. Let $f\colon \Sigma_* \to \mathbb{R}$ be an algebraic Morse function with algebraic gradient V and associated flow Φ. Recall that the linear map F is such that $\mathrm{F}(\sigma) = -\langle \partial\tau, \sigma\rangle^{-1} \cdot \tau$ if $(\sigma, \tau) \in V$ and 0 on all other basis elements of Σ_*. Thus, if $(\sigma, \tau) \in V$, the unique basis element in the support of $\mathrm{F}(\sigma)$, namely τ, has the same function value as σ. As Σ_* is finite and hence the image of the map f is also finite, we can therefore prove the statement inductively over the function values of the basis elements.

To establish the base case of the induction, let $\sigma \in \Sigma_*$ be a basis element with minimal function value. Then, σ is necessarily a critical element contained in Σ_0, and therefore, we have $\Phi(\sigma) = \sigma$, by definition of the gradient flow. This establishes the base case.

For the induction step, let $\sigma \in \Sigma_*$ be a basis element and assume the claim holds for all other basis elements with strictly smaller function value than σ. If σ is a gradient facet, or a gradient cofacet, then the support of $\Phi(\sigma) = \sigma + \partial\,\mathrm{F}(\sigma) + \mathrm{F}(\partial\sigma)$ does not contain σ, by definition of the flow Φ and the linear map F. Moreover, every element in the support of $\Phi(\sigma)$ has strictly smaller function value than σ, by definition of the algebraic Morse function f. Thus, the claim follows in this case from the induction assumption. If σ is a critical element, by Lemma 2.53, it suffices to prove that $\Phi^r(w)$, where $w = \mathrm{F}(\partial\sigma)$, is equal to zero for large enough r. As $\Phi^s(w) \in \operatorname{im} \mathrm{F}$ for all $s \in \mathbb{N}$, it thus suffices to prove that for every chain c in the image of F, the flow Φ strictly decreases the maximal function value of the basis elements in the support of c. As the image of f is finite, this then implies the desired claim. So, let τ be any basis element in the support of the chain $c \in \operatorname{im} \mathrm{F}$, meaning that τ is a gradient cofacet with respect to the algebraic gradient V. Analogously to the first case above, the chain $\Phi(\tau) = \tau + \mathrm{F}(\partial\tau)$ does not contain τ, by definition of the flow Φ and the linear map F. Moreover, every element in the support of $\Phi(\tau)$ has strictly smaller function value than τ, by definition of the algebraic Morse function f. Thus, the claim also follows in this case. □

Based on the concept of discrete flow, Forman [For98] considered the subcomplex of the cellular chain complex given by the flow invariant chains and proved the following proposition in this special case, which also generalizes to our setting. Consider the subcomplex of C_* consisting of the Φ-invariant chains,

$$C_*^\Phi = \{c \in C_* \mid \Phi(c) = c\}.$$

Proposition 2.54. *The Φ-invariant chains are spanned by the image of the critical basis elements under the stabilized flow $C_n^\Phi = \operatorname{span}\{\Phi^\infty(\sigma) \mid \sigma \in \Sigma_n$ critical$\}$.*

Proof. Let $f\colon \Sigma_* \to \mathbb{R}$ be an algebraic Morse function with algebraic gradient V and associated flow Φ. Let $c \in C_n$ be a Φ-invariant chain; it suffices to prove that c is a linear combination of chains of the form $\Phi^\infty(\sigma)$ with $\sigma \in \Sigma_n$ critical. Let $\sigma \in \Sigma_n$ be a basis element in the support of c with maximal function value. It follows, through similar arguments as in the induction step in the proof of Proposition 2.52, that σ must be a critical element. Then, the chain

$$c - \langle c, \sigma\rangle \cdot \Phi^\infty(\sigma)$$

is also Φ-invariant, with one critical element less in its support than c. As there are only finitely many critical elements in Σ_*, by repeating the argument above, there must exist critical elements $\sigma_1, \dots, \sigma_k$ in Σ_n such that the chain

$$c - \langle c, \sigma_1\rangle \cdot \Phi^\infty(\sigma_1) - \cdots - \langle c, \sigma_k\rangle \cdot \Phi^\infty(\sigma_k) \tag{2.1}$$

is also Φ-invariant and that contains no critical elements in its support. Since, as before, elements in the support of a Φ-invariant chain with maximal function value are always critical, the chain in Eq. (2.1) must be zero, proving the claim. □

We relate the flow invariant chains of an algebraic gradient to those of its subgradients.

Proposition 2.55. *Let (C_*, Σ_*) be a based chain complex and $W \subseteq V$ two algebraic gradients on Σ_*. Consider the flows $\Psi, \Phi\colon C_* \to C_*$ determined by W and V, respectively. Then any Φ-invariant chain is also Ψ-invariant, i.e., we have $C_*^\Phi \subseteq C_*^\Psi$.*

Proof. Let $c \in C_n$ be a Φ-invariant chain. Recall from Definition 2.51 that Φ is given by $\Phi(c) = c + \partial\,\mathrm{F}(c) + \mathrm{F}(\partial c)$, where $\mathrm{F}\colon C_* \to C_{*+1}$ is the linear map with $\mathrm{F}(\sigma) = -\langle \partial\tau, \sigma\rangle^{-1} \cdot \tau$ if $(\sigma, \tau) \in V$ and 0 on all other basis elements of Σ_*. By Proposition 2.54 and linearity, we can assume without loss of generality that c is of the form $\Phi^\infty(\eta)$ for a V-critical basis element $\eta \in \Sigma_n$. By definition, $\Phi^\infty = \Phi^{r+1}$ for a large enough $r \in \mathbb{N}$, and by Lemma 2.53, the chain $c = \Phi^{r+1}(\eta)$ satisfies $\mathrm{F}(c) = \mathrm{F}(\Phi^{r+1}(\eta)) = 0$. Hence, we have

$$c = \Phi(c) = c + \mathrm{F}(\partial c) + \partial\,\mathrm{F}(c) = c + \mathrm{F}(\partial c)$$

from which we conclude that $\mathrm{F}(\partial c) = 0$ holds. By construction of F, this implies that c and ∂c are the sums of critical elements and gradient cofacets of V.

Similarly to before, Ψ is given by $\Psi(c) = c + \partial E(c) + E(\partial c)$, where $E\colon C_* \to C_{*+1}$ is the linear map with $E(\sigma) = -\langle \partial\tau, \sigma\rangle^{-1} \cdot \tau$ if $(\sigma, \tau) \in W$ and 0 on all other basis elements of Σ_*. As $W \subseteq V$, it holds true that c and ∂c are the sums of critical elements and gradient cofacets of W, as well. By construction of E, this implies $E(c) = E(\partial c) = 0$, and hence

$$\Psi(c) = c + \partial E(c) + E(\partial c) = c + 0 + 0 = c,$$

meaning that c is Ψ-invariant, proving the claim. □

Note that the flow Φ can be written as a sum of flows

$$\Phi = \sum_{(a,b)\in V} \Phi^{(a,b)} - (\operatorname{card} V - 1) \cdot \mathrm{id},$$

where $\Phi^{(a,b)}$ is the flow determined by the algebraic gradient $\{(a, b)\}$ on Σ_*. Together with Proposition 2.55, this proves the following.

Corollary 2.56. *Let V be an algebraic gradient on Σ_* with associated flow $\Phi\colon C_* \to C_*$. Then a chain is Φ-invariant if and only if it is $\Phi^{(a,b)}$-invariant for every pair $(a, b) \in V$.*

2.4. Abstract Homotopy Theory

To be able to prove functorial nerve theorems in Section 5, we now recall some notions from homotopy theory, such as simplicial model categories. Most prominently, we will make use of the *bar construction*, which is a standard model for the homotopy colimit: like the colimit, the homotopy colimit can be defined via a universal property, but since this universal property is phrased in terms of derived categories, it takes some work to define it precisely. A full discussion of the homotopy colimit is beyond the scope of this section (see [Dug08] for a nice introduction to the topic, or [Rie14, Part I] for a more abstract approach). However, in order to explain the properties of the bar construction that we will use, we will at least describe the problem that the homotopy colimit addresses. So, in this section we will introduce a basic problem with the colimit of a diagram of topological spaces, give an idea of how the homotopy colimit addresses this problem, define the bar construction and explain some properties of the bar construction that can be used to prove functorial nerve theorems in Section 5.

2.4.1. Homotopy Colimits and the Bar Construction

While colimits are used everywhere in topology to construct new spaces, the colimit operation fails to respect homotopy equivalences, in the following sense. Take $\mathcal{A}$ to be the category that looks like this:

$$\bullet \leftarrow \bullet \rightarrow \bullet$$

and consider the commutative diagram:

$$\begin{array}{ccccc} D^n & \hookleftarrow & S^{n-1} & \hookrightarrow & D^n \\ \downarrow & & \downarrow \mathrm{id} & & \downarrow \\ * & \longleftarrow & S^{n-1} & \longrightarrow & * \end{array}$$

Here, the top maps are the boundary inclusions. We think of the rows as $\mathcal{A}$-shaped diagrams, and the vertical maps define a natural transformation between these two $\mathcal{A}$-shaped diagrams. Every component of this natural transformation is a homotopy equivalence, but the colimit of the top row is the sphere S^n, while the colimit of the bottom row is a one-point space $*$, so the induced map between the colimits cannot be a homotopy equivalence.

More generally, let $\mathcal{C}$ be a small category, and write $\mathsf{Top}^{\mathcal{C}}$ for the category of functors $\mathcal{C} \to \mathsf{Top}$. One says that a functor $\Omega\colon \mathsf{Top}^{\mathcal{C}} \to \mathsf{Top}$ is *homotopical* if, given a natural transformation $\lambda\colon F \Rightarrow F'$ between $\mathcal{C}$-shaped diagrams F and F' that is a pointwise homotopy equivalence, the induced map $\Omega(F) \to \Omega(F')$ is also a homotopy equivalence. For any small category $\mathcal{C}$, the colimit defines a functor $\operatorname{colim}\colon \mathsf{Top}^{\mathcal{C}} \to \mathsf{Top}$, and the previous example shows that this functor is not homotopical in general.

A homotopy colimit is a homotopical functor $\operatorname{hocolim}\colon \mathsf{Top}^{\mathcal{C}} \to \mathsf{Top}$, together with a natural transformation hocolim $\Rightarrow$ colim that makes hocolim, in some sense, the best possible homotopical approximation of the colimit functor. We now show how to construct a particular model for the homotopy colimit, called the bar construction, and we will see that it can be thought of as a "thickened" version of the colimit; see Example 2.61 for an illustration.

We write Δ^n for the standard topological n-simplex, and for $0 \le i \le n$, we write $d^i\colon \Delta^{n-1} \hookrightarrow \Delta^n$ for the inclusion of the face opposite the i^{th} vertex.

Definition 2.57. Let P be a poset, and let $F\colon P \to \mathsf{Top}$ be a diagram of topological spaces. The *bar construction* $\operatorname{Bar}(F)$ of F is the quotient space

$$\operatorname{Bar}(F) = \left(\coprod_{\sigma=(v_0<\cdots<v_n)} F(v_0) \times \Delta^n \right) \Big/ \sim$$

where the disjoint union is taken over all chains in P, and the equivalence relation $\sim$ is generated as follows. For a chain $\sigma = (v_0 < \cdots < v_n)$ and $0 \le i \le n$, we write

$$\tau_i = (v_0 < \cdots < \hat{v}_i < \cdots < v_n) = (w_0 < \cdots < w_{n-1})$$

for the subchain with v_i left out, noting that if $i > 0$, then $w_0 = v_0$, and if $i = 0$, then $w_0 = v_1$. Now for any $x \in F(v_0)$ and $\alpha \in \Delta^{n-1}$, we identify $(x, d^i(\alpha))$ in the copy of $F(v_0) \times \Delta^n$ indexed by σ with $(F(v_0 \le w_0)(x), \alpha)$ in the copy of $F(w_0) \times \Delta^{n-1}$ indexed by τ_i.

Example 2.58. Let $P = \{0 < 1\}$. Then a diagram $F\colon P \to \mathsf{Top}$ is just a map $F(0) \to F(1)$, and the bar construction $\operatorname{Bar}(F)$ is the mapping cylinder of this map.

Recall the functor $\operatorname{Pos}\colon \mathsf{Simp} \to \mathsf{Po}$ from Section 2.1 that takes a simplicial complex to its poset of simplices (ordered by inclusion).

Definition 2.59. Let X be a topological space and $\mathcal{U} = (U_i)_{i\in I}$ a cover of X. Writing $P_{\mathcal{U}} = \operatorname{Pos}(\operatorname{Nrv}(\mathcal{U}))^{\mathrm{op}}$, the *nerve diagram* of the cover $\mathcal{U}$ is the functor $\mathcal{D}_{\mathcal{U}}\colon P_{\mathcal{U}} \to \mathsf{Top}$ with $\mathcal{D}_{\mathcal{U}}(J) = U_J$.

Remark 2.60. In many cases, the colimit of the diagram $\mathcal{D}_{\mathcal{U}}$ simply gives us back X: the inclusions $U_J \subseteq X$ induce a continuous map $\operatorname{colim} \mathcal{D}_{\mathcal{U}} \to X$, which is in fact a bijection. If $\mathcal{U}$ is an open cover, or if it is a closed cover that is *locally finite* (i.e., every point of X has an open neighborhood that intersects only finitely many cover elements), then this bijection is a homeomorphism.

Example 2.61. In Section 5, we will mainly consider bar constructions of diagrams associated to a cover. For example, consider the following cover $\mathcal{U}$ of the circle S^1:

Figure 9: A cover by three arcs (left) and the intersections of those (right).

If we label the three arcs $\{b, r, g\}$, the poset $P_{\mathcal{U}}$ associated to this cover has the following form:

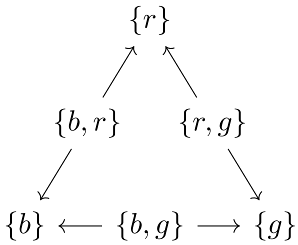

By definition, the bar construction $\operatorname{Bar}(\mathcal{D}_{\mathcal{U}})$ of the nerve diagram $\mathcal{D}_{\mathcal{U}} \colon P_{\mathcal{U}} \to \mathsf{Top}$ associated to the cover $\mathcal{U}$ is built from pieces indexed by chains $v_0 < \cdots < v_n$ in $P_{\mathcal{U}}$ and are of the form $\mathcal{D}_{\mathcal{U}}(v_0) \times \Delta^n$. More concretely, the bar construction in our example is built from the following pieces:

Figure 10: Pieces indexed by chains in $P_{\mathcal{U}}$ of length zero (left) and of length one (right).

After making all identifications, the bar construction $\operatorname{Bar}(\mathcal{D}_{\mathcal{U}})$ is the following "thickened" version of $\operatorname{colim} \mathcal{D}_{\mathcal{U}} \cong S^1$:

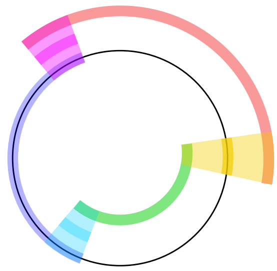

Figure 11: The bar construction of the nerve diagram.

By $\mathrm{Diag}_{\mathsf{Po}}(\mathsf{Top})$ we denote the category of diagrams over a poset: the objects in this category are tuples (P, F), where P is a poset and $F\colon P \to \mathsf{Top}$ is a functor. A morphism $(g, \lambda)\colon (P, F) \to (R, G)$ consists of a poset map $g\colon P \to R$ and a natural transformation $\lambda\colon F \Rightarrow G \circ g$. Then the bar construction defines a functor $\mathrm{Bar}\colon \mathrm{Diag}_{\mathsf{Po}}(\mathsf{Top}) \to \mathsf{Top}$: a morphism (g, λ) induces a continuous map $\mathrm{Bar}(F) \to \mathrm{Bar}(G)$ defined by the maps

$$\lambda(v_0) \times |\operatorname{Flag}(g)|\colon F(v_0) \times \Delta^n \to G(g(v_0)) \times \Delta^m ,$$

where $|\operatorname{Flag}(g)|\colon \Delta^n \to \Delta^m$ is the affine map that sends the vertex v_i to $g(v_i)$. Moreover, the projection maps $F(v_0) \times \Delta^n \to F(v_0)$ define a natural map $\mathrm{Bar}(F) \to \operatorname{colim} F$.

There are analogues of this quotient space construction in other settings, which are also called bar constructions. We will encounter some of these in Section 5.4. For a very general discussion of the bar construction, including a proof that it is a model of the homotopy colimit, see [Rie14, Chapters 4–5]. The bar construction for topological spaces is homotopical (see [Koz08, Theorem 15.12] or [Hat02, Proposition 4G.1]):

Proposition 2.62. *Let P be a poset, $F, G\colon P \to \mathsf{Top}$ diagrams of topological spaces, and let $\lambda\colon F \Rightarrow G$ be a natural transformation. If the component $\lambda(v)$ is a homotopy equivalence for all $v \in P$, then so is the induced map* $\mathrm{Bar}(F) \to \mathrm{Bar}(G)$.

2.4.2. Simplicial Model Categories

In order to prove the unified nerve theorem (Theorem 5.25), we will need a generalization of the bar construction (Definition 2.57) to other settings than the category of topological spaces. To make sense of the homotopy invariance property (Proposition 2.62) in other settings, we will need a general framework for studying analogues of homotopy equivalences in other contexts. There are many choices for such frameworks: we will work with *model categories*, which have been a standard tool of abstract homotopy theory since they were introduced by Quillen in the 1960s. A thorough introduction to model categories is beyond the scope of this section (see, e.g., [DS95] for a friendly introduction), but we will briefly introduce the aspects of model categories that are most relevant to Section 5. We first recall some notions from topology.

To avoid pathological behavior in the category Top of all topological spaces, results in algebraic topology are often restricted to certain full subcategories that include all the spaces of primary interest (such as CW complexes) and that have better categorical properties. For example, it is often convenient to work in a category of topological spaces that is *cartesian closed*: roughly speaking, this means that for any spaces X and Y in the category, we have a "mapping space" Y^X in the category such that for a fixed space Z, the set of maps $Z \to Y^X$ is in bijection with the set of maps $X \times Z \to Y$, and this bijection is natural in Y and Z. Letting $Z = *$, we see that the points of Y^X are in bijection with continuous maps $X \to Y$. Such mapping spaces play an important role in algebraic topology, because they encode homotopy-theoretic information. For example, a path $\gamma\colon [0,1] \to Y^X$ in the mapping space corresponds to a homotopy $H\colon X \times [0,1] \to Y$. There is more than one standard choice for a cartesian closed subcategory. We will consider the following one.

Definition 2.63. A topological space X is *weak Hausdorff* if $g(K)$ is closed in X for every continuous map $g\colon K \to X$ with K compact Hausdorff. A subspace A of X is *compactly closed* if $g^{-1}(A)$ is closed in K for every continuous map $g\colon K \to X$ with K compact Hausdorff. A space X is a *k-space* if every compactly closed subspace of X is closed. A space X is *compactly generated* if it is a weak Hausdorff k-space. The full subcategory of Top of compactly generated spaces is denoted by CGSpc.

A note of warning: there is conflicting terminology in the literature surrounding compactly generated spaces. See [May99, Chapter 5] or [Str09] for basic facts about these spaces. For example, there exist inclusions and adjoint functors

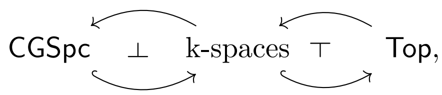

where k-spaces is the full subcategory of Top consisting of k-spaces.

Example 2.64. Many spaces are compactly generated:

- Every closed subspace of a compactly generated space is compactly generated.
- Every CW-complex is compactly generated.
- Every locally compact Hausdorff space is compactly generated [Str09, Proposition 1.7]. In particular, $\mathbb{R}^d$ is compactly generated.

Finally, let us recall the homotopy extension property.

Definition 2.65. Let X be topological spaces and let A be a subset. We say that the pair (X, A) satisfies the *homotopy extension property* if for every commutative diagram of the following shape the dotted arrow exists

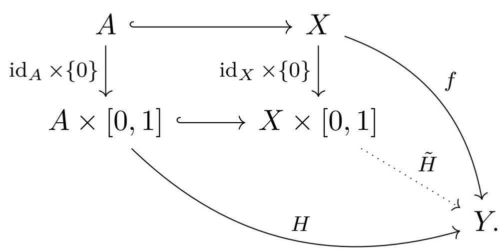

In words, the pair (X, A) has the homotopy extension property if for any map f, every homotopy H of f on A can be extended to a homotopy $\tilde{H}$ of f defined on all of X.

Remark 2.66. A large class of pairs has the homotopy extension property. For example, if X is a CW-complex and A a subcomplex, then (X, A) satisfies the homotopy extension property ([Hat02, Proposition 0.16] or [Koz08, Proposition 7.10]). We say more about the homotopy extension property in Section 5.4.1.

Model Categories A model category is a category together with three distinguished classes of morphisms, the *weak equivalences*, *fibrations*, and *cofibrations*, which are required to satisfy certain axioms. An admissible choice of these classes is called a *model structure* on the underlying category. The distinguished classes of morphisms also determine two distinguished classes of objects: an object X is *fibrant* if the unique map from X to the terminal object is a fibration, and it is *cofibrant* if the unique map from the initial object is a cofibration. Before we give the axioms in Definition 2.69, it is useful to have in mind some basic examples.

Example 2.67. There are several important model categories whose objects are topological spaces. As discussed earlier in this section, in order to avoid pathological behavior, one often considers some subcategory of Top; we choose the subcategory of compactly generated spaces. There is a model structure on the category of compactly generated spaces for which the weak equivalences are the homotopy equivalences and the cofibrations are the *Hurewicz cofibrations*, which are the maps $i \colon A \to X$ that satisfy the homotopy extension property (see Definition 2.65, and replace the inclusion $A \subset X$ with i). This is called the *Hurewicz model structure*. It was originally shown to be a model structure (on the category of all topological spaces) by Strøm [Str72]; see [MP12, Theorem 17.1.1] for an account in the setting of compactly generated spaces. Every space is both fibrant and cofibrant in the Hurewicz model structure, which is quite rare.

There is another model structure on the category of compactly generated spaces, called the *Quillen model structure*, for which the weak equivalences are the weak homotopy equivalences, i.e., the maps that induce a bijection on path components and an isomorphism on homotopy groups for all choices of base point. This was first studied by Quillen in his original work on model categories [Qui67]; see [MP12, Theorem 17.2.2] for an account in our setting. Every space is fibrant in the Quillen model structure, and every CW complex is cofibrant.

Example 2.68. Model categories can be used to study homological algebra. Let R be a commutative ring. There is a model structure on the category of non-negatively graded chain complexes of R-modules, for which the weak equivalences are the quasi-isomorphisms, and the cofibrations are those monomorphisms that have a degreewise-projective cokernel. In particular, the cofibrant objects are the degreewise-projective chain complexes. This is another of the original examples from [Qui67].

Definition 2.69. A *model category* $\mathcal{M}$ is a category which is equipped with three subcategories of morphisms called *weak equivalences*, *fibrations* and *cofibrations* such that the following axioms hold:

1. The category $\mathcal{M}$ has all small limits and colimits.
2. (2-of-3) If f and g are maps of $\mathcal{M}$ such that $g \circ f$ is defined and two of the maps f, g, $g \circ f$ are weak equivalences, then so is the third.
3. If f is a retract of g and g is a weak equivalence, fibration, or cofibration, then so is f.
4. Given a commutative square
$$\begin{array}{ccc} A & \xrightarrow{f} & X \\ {\scriptstyle i}\downarrow & & \downarrow{\scriptstyle p} \\ B & \xrightarrow[g]{} & Y, \end{array}$$
where i is a cofibration and p is a fibration, then there is a map $h\colon B \to X$ such that $f = h \circ i$ and $g = p \circ h$ if either of i or p is a weak equivalence.
5. Any map f can be factored as (i) $f = p \circ i$, where i is a cofibration and p is a fibration and a weak equivalence, and (ii) $f = p' \circ i'$ where i' is a cofibration and a weak equivalence and p' is a fibration.

Remark 2.70. The definition of model category has evolved since it was first introduced. For example, we require all small limits and colimits, while Quillen originally required only all finite limits and colimits. For a discussion, see [Hov99, Chapter 1]

Remark 2.71. In a model category the weak equivalences together with the fibrations or the cofibrations determine the third subcategory; see, e.g., [Hov99, Lemma 1.1.10].

Many algebraic topologists prefer to work with certain kinds of combinatorial models of spaces, rather than with topological spaces themselves. These combinatorial models are called *simplicial sets*, and they are somewhat similar to simplicial complexes. While they may appear more complicated than simplicial complexes – for example, every simplicial set has simplices in every dimension, even the simplicial set that models the one-point space – they have better categorical properties. For example, there is a geometric realization functor $|-|$ from the category of simplicial sets to the category of compactly generated topological spaces, and this functor preserves all small colimits and all finite limits (we define this construction in Section 5.4.2). So, one can take limits and colimits in the category of simplicial sets, and these will model the corresponding limits and colimits of topological spaces. See [Fri12] for a friendly introduction to this topic.

Definition 2.72. The *simplex category*, denoted by Δ, has as objects the finite ordinals $\{[n] = \{0, 1, \ldots, n\} \mid n \geq 0\}$, with the morphisms being the order preserving maps.

Definition 2.73. A *simplicial set* is a functor $X\colon \Delta^{\mathrm{op}} \to \mathsf{Set}$, and a *morphism of simplicial sets* is a natural transformation. The set $X_n = X([n])$ is the set of *n-simplices* of X. The category of simplicial sets is denoted by sSet. More generally, if $\mathcal{C}$ is any category, a *simplicial object in* $\mathcal{C}$ is a functor $\Delta^{\mathrm{op}} \to \mathcal{C}$, and the category of simplicial objects in $\mathcal{C}$ is denoted by $\mathsf{s}(\mathcal{C})$.

Example 2.74. The *Yoneda embedding* $Y\colon \Delta \hookrightarrow \mathrm{sSet}$, $[n] \mapsto \mathrm{Hom}_\Delta(-,[n])$ gives rise to a simplicial set for each $n \in \mathbb{N}$. We denote $Y([n])$ by Δ^n and call it the *standard n-simplex*.

Example 2.75. Let X be a topological space. The *singular simplicial set of X* is the simplicial set $\mathrm{Sing}(X)$ with

$$\mathrm{Sing}(X)([n]) = \mathrm{hom}(|\Delta^n|, X),$$

where $|\Delta^n|$ is the standard topological n-simplex, and $\mathrm{hom}(-,-)$ denotes the set of continuous maps.

A fundamental fact about the relationship between simplicial sets and topological spaces is that the functor $\mathrm{Sing}\colon \mathsf{CGSpc} \to \mathsf{sSet}$ is right adjoint to the geometric realization $|-|\colon \mathsf{sSet} \to \mathsf{CGSpc}$ mentioned above. This adjunction is what allows us to use simplicial sets as a model for spaces; we say more about this below.

Example 2.76. Let $\mathcal{C}$ be a category. The *(categorical) nerve of $\mathcal{C}$* is the simplicial set $N(\mathcal{C})$ such that

$$N(\mathcal{C})([n]) = \{v_0 \to v_1 \to \cdots \to v_n \mid \text{string of composable morphisms in } \mathcal{C}\}.$$

If $\mathcal{A} = (A_i)_{i\in I}$ is a cover of a topological space, then the finite non-empty intersections of cover elements form a category with morphisms given by inclusion $A_J \hookrightarrow A_{J'}$ if $J' \subseteq J$. The nerve of this category and the nerve of the cover have homeomorphic geometric realizations, explaining the common name for the two constructions.

We can now introduce two more examples of model categories, both of which play a role in the proof of Theorem 5.25.

Example 2.77. The category of simplicial sets can be equipped with a model structure, called the *Quillen model structure*, where the the cofibrations are the monomorphisms and the weak equivalences are those maps that are mapped to weak homotopy equivalences when applying the geometric realization functor; see, for example, [GJ09, Chapter I]. From the definition of cofibration, it follows that every simplicial set is cofibrant.

Example 2.78. An alternative to the chain complexes of Example 2.68 is the category of simplicial R-modules. Here, R is a commutative ring as before, and a simplicial R-module is a simplicial object in the category R-Mod of R-modules, i.e., a functor $\Delta^{\mathrm{op}} \to R$-$\mathsf{Mod}$. Denoting the category of simplicial R-modules by $\mathsf{s}(R\text{-}\mathsf{Mod})$, we let $\mathsf{R}\colon \mathsf{sSet} \to \mathsf{s}(R\text{-}\mathsf{Mod})$ denote the functor that is induced by the free R-module functor $\mathsf{R}\colon \mathsf{Set} \to R\text{-}\mathsf{Mod}$; the forgetful functor $U\colon \mathsf{s}(R\text{-}\mathsf{Mod}) \to \mathsf{sSet}$ is right adjoint to R. Then there is a model structure on $\mathsf{s}(R\text{-}\mathsf{Mod})$ such that the weak equivalences and fibrations are exactly those morphisms whose underlying map of simplicial sets is a weak equivalence and fibration, respectively [GS07, Proposition 4.2]. Moreover, a continuous map $X \to Y$ is an R-homology isomorphism if and only if the induced map $\mathsf{R}(\mathrm{Sing}(X)) \to \mathsf{R}(\mathrm{Sing}(Y))$ is a weak equivalence of simplicial R-modules [Wei94, Dold-Kan Theorem 8.4.1].

We have now encountered two important adjunctions connecting model categories: the adjunction $(|-|, \mathrm{Sing})$ relating spaces with simplicial sets, and (R, U) relating simplicial sets with simplicial R-modules. In general, a *Quillen adjunction* between model categories is an adjunction such that the left adjoint preserves cofibrations and trivial cofibrations, or equivalently, the right adjoint preserves fibrations and trivial fibrations. These adjunctions are the main way to relate model categories; both of the adjunctions just mentioned are Quillen adjunctions.

A Quillen adjunction (F, G), where the left adjoint F is a functor $\mathcal{M} \to \mathcal{N}$, is a *Quillen equivalence* if, for all cofibrant X in $\mathcal{M}$ and all fibrant Y in $\mathcal{D}$, a map $FX \to Y$ is a weak equivalence in $\mathcal{N}$ if and only if the corresponding map $X \to GY$ is a weak equivalence in $\mathcal{M}$. The adjunction $(|-|, \mathrm{Sing})$ is a Quillen equivalence when CGSpc is given the Quillen model structure; see for example [Hov99, Theorem 2.4.25, Theorem 3.6.7]. This simple definition has powerful consequences, and we now describe one that plays a role in the proof of Theorem 5.25. Since (F, G) is an adjunction, there is a natural map $\eta\colon X \to GFX$ called the *unit*, and another $\varepsilon\colon FGY \to Y$ called the *counit*. If (F, G) is a Quillen equivalence, then these maps are weak equivalences, subject to additional fibrancy and cofibrancy assumptions. See [Hov99, Proposition 1.3.13] for details. Once we know that $(|-|, \mathrm{Sing})$ is a Quillen equivalence, then it follows immediately that the unit $K \to \mathrm{Sing}(|K|)$ is a weak equivalence for every simplicial set K, and the counit $|\mathrm{Sing}(Y)| \to Y$ is a weak equivalence for every compactly-generated space Y. The additional fibrancy and cofibrancy assumptions are vacuous in this case, as every simplicial set is cofibrant and every space is fibrant.

Simplicial Model Categories At the beginning of this section, we discussed the importance of mapping spaces. A simplicial model category $\mathcal{M}$ is a model category equipped with additional structure that generalizes this feature of algebraic topology; see [Rie14, Definition 11.4.4] for a precise definition. For any two objects X and Y of a simplicial model category $\mathcal{M}$, we have a simplicial set $\mathbf{Hom}_{\mathcal{M}}(X, Y)$ that encodes homotopy-theoretic information about X and Y. Formally, one requires that the model category $\mathcal{M}$ be enriched in simplicial sets, and tensored and cotensored. One then imposes an additional axiom that relates this structure to the model structure. We will omit the formal definitions, since we will not use most of the structure explicitly. Rather, for the proof of the unified nerve theorem, we will need a few facts about simplicial model categories, principally Proposition 5.42. However, we will use the *tensoring* explicitly in order to define the bar construction in a simplicial model category (Definition 5.38), and so we introduce it now. If $\mathcal{M}$ is a simplicial model category, then for any object X of $\mathcal{M}$ and any simplicial set K there is an object $X \otimes K$ of $\mathcal{M}$, and this construction gives a functor $\mathcal{M} \times \mathsf{sSet} \to \mathcal{M}$. Furthermore, for any X, the functor $\mathbf{Hom}_{\mathcal{M}}(X, -)\colon \mathcal{M} \to \mathsf{sSet}$ has a left adjoint given by $X \otimes -\colon \mathsf{sSet} \to \mathcal{M}$. This adjunction in particular yields an isomorphism

$$\hom_{\mathcal{M}}(X \otimes K, Y) \cong \hom_{\mathsf{sSet}}(K, \mathbf{Hom}_{\mathcal{M}}(X, Y))$$

for any X, Y in $\mathcal{M}$ and any simplicial set K. The motivation for the terminology "tensoring", and the notation $\otimes$, comes from this adjunction, which is analogous to the tensor-hom adjunction from linear algebra.

Example 2.79. The category CGSpc of compactly-generated spaces is enriched in simplicial sets, and tensored and cotensored. This makes CGSpc a simplicial model category with either the Hurewicz or Quillen model structures. If X and Y are compactly generated spaces, $\mathbf{Hom}(X, Y)$ is the simplicial set with

$$\mathbf{Hom}(X, Y)_n = \hom(X \times |\Delta^n|, Y).$$

So, the zero-simplices of $\mathbf{Hom}(X, Y)$ are maps from X to Y, the one-simplices are homotopies, the two-simplices are "homotopies between homotopies", and so on. The operation $\otimes$ is characterized by $X \otimes \Delta^n = X \times |\Delta^n|$.

Example 2.80. The Quillen model structure on sSet gives a simplicial model category, and the operation $\otimes$ is the cartesian product.

Example 2.81. The category $\mathsf{s}(R\text{-}\mathsf{Mod})$ is a simplicial model category, with the model structure described in Example 2.78. If K is a simplicial set and M is a simplicial R-module, then $M \otimes K$ is the simplicial R-module with $(M \otimes K)_n = M_n \otimes_R \mathsf{R}(K_n)$.

3. Vietoris–Rips Filtrations of Hyperbolic and Almost Geodesic Spaces

In this section, we prove the main results summarized in Section 1.2, namely Theorems A and B. More concretely, in Section 3.1, we establish some independent facts about the hyperbolicity and the geodesic defect of metric spaces. In Section 3.2, we slightly modify the proof strategy of the Contractibility Lemma from [LMO22], which uses the injective hull of a metric space, to extend the result to non-geodesic spaces. In Section 3.3, we provide an alternative proof strategy for the Contractibility Lemma for non-geodesic spaces that is based on discrete Morse theory and that provides collapses compatible with the Vietoris–Rips filtration. Finally, in Section 3.4, we argue that this result has strong implications to the computation of persistent homology for tree-like metric data, by showing that for tree metrics these collapses are induced by the apparent pairs gradient (see also Remark 3.24), which is closely related to the persistence computation (Lemma 2.24).

3.1. Hyperbolicity and Geodesic Defect of Metric Spaces

We establish some basic facts about the hyperbolicity and the geodesic defect of metric spaces, such as their stability with respect to the Gromov–Hausdorff distance. The results shown here are independent of the results in the following sections.

Recall the definitions from Section 1.2. If X is δ-hyperbolic, then it is also δ'-hyperbolic for every $\delta' \geq \delta$. With this in mind, it is natural to consider the infimum over all δ such that X is δ-hyperbolic, which is called the *hyperbolicity* $\operatorname{hyp}(X)$ of X. It follows from Eq. (1.1) that we have

$$\operatorname{hyp}(X) = \sup_{w,x,y,z \in X} \frac{d(w,x) + d(y,z) - \max\{d(w,y) + d(x,z), d(w,z) + d(x,y)\}}{2}, \quad (3.1)$$

and from this equivalent description of the hyperbolicity it can be seen that X is indeed $\operatorname{hyp}(X)$-hyperbolic. Moreover, note that every subspace of a δ-hyperbolic space is also δ-hyperbolic and that if X is δ-hyperbolic with respect to some point $p \in X$, then X is 2δ-hyperbolic [Gro87, Corollary 1.1.B].

Similar spaces have similar hyperbolicity.

Proposition 3.1 (Mémoli et al. [MOW21])**.** *Let X and Y be any metric spaces and $s = d_{\mathrm{GH}}(X, Y)$. If X is δ-hyperbolic, then Y is $(\delta + 4s)$-hyperbolic. Hence,*

$$|\operatorname{hyp}(X) - \operatorname{hyp}(Y)| \leq 4 d_{\mathrm{GH}}(X, Y).$$

Remark 3.2. By definition, if X is δ-hyperbolic it is also δ-hyperbolic with respect to every point $p \in X$, and if X is δ-hyperbolic with respect to some point $p \in X$ we just stated that it is also 2δ-hyperbolic. Both implications are tight in the following sense. If $p \in X$ is a point such that for every other point $q \in X$ there is an isometry of X sending q to p, e.g., take the vertices of a square in the plane, then X is δ-hyperbolic with respect to q if and only if it is δ-hyperbolic with respect to p. In particular, X is δ-hyperbolic if and only if it is δ-hyperbolic with respect to p. For the other implication, consider the following graph with vertex set $\{a, b, c, d, e\}$, whose edges all have length two:

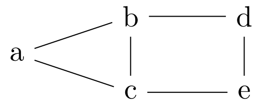

For this graph metric, the smallest hyperbolicity constant with respect to a is equal to 1 and for every other point it is equal to 2. In particular, the hyperbolicity of this space is $2 = 2 \cdot 1$.

Remark 3.3. In contrast, for the geodesic defect the relationship between the local and the global quantities can differ arbitrarily. Note that if X is ν-almost geodesic with respect to some point $p \in X$, then a straightforward estimate involving the triangle inequality shows that for any $q \in X$ it is also $(\nu + \frac{3}{2}d(q,p))$-almost geodesic with respect to q. However, the space X does not need to be ν-almost geodesic for any ν. For example, consider the subspace X of the plane $\mathbb{R}^2$ given by the coordinate axes equipped with the Euclidean metric. Then X is 0-almost geodesic with respect to the origin, but it is not ν-almost geodesic for any ν: Let $t > 0$ be arbitrary and take $x = (0,t), p = (t,0)$, and $r = s = \frac{\sqrt{2}}{2}t$. A straightforward calculation yields the lower bound $\operatorname{geod}(X) \geq \operatorname{geod}_p(X) \geq (1 - \frac{\sqrt{2}}{2})t$, which directly implies the claim.

We now establish lower and upper bounds on the geodesic defect, relating it to other geometric quantities. For a metric space X consider its *separation* $\operatorname{sep}(X) = \inf_{x \neq y} d(x,y)$.

Proposition 3.4. *For any metric space X we have* $\operatorname{geod}(X) \geq \frac{1}{2}\operatorname{sep}(X)$.

Proof. If X is not ν-almost geodesic for any ν, then $\operatorname{geod}(X) = \infty$ and there is nothing to prove. Thus, assume that X is ν-almost geodesic. Let $\epsilon > 0$ be arbitrary and let $u, w \in X$ be any two points with $u \neq w$ and $d(u,w) - \epsilon \leq I := \inf_{x \neq y} d(x,y)$. Then any other point has distance at least $d(u,w) - \epsilon$ to u and w. As X is ν-almost geodesic, there exists a point $z \in X$ with $d(u,z) \leq \frac{1}{2}d(u,w) + \nu$ and $d(w,z) \leq \frac{1}{2}d(u,w) + \nu$. If $z = w$, then the first inequality implies $\frac{1}{2}I \leq \frac{1}{2}d(u,w) \leq \nu$ and hence $\frac{1}{2}I \leq \operatorname{geod}(X)$, because $\operatorname{geod}(X)$ is the infimum over all ν such that X is ν-almost geodesic. If $z \neq w$, then

$$d(u,w) - \epsilon \leq I \leq d(w,z) \leq \frac{1}{2}d(u,w) + \nu$$

and therefore $\frac{1}{2}I - \epsilon \leq \frac{1}{2}d(u,w) - \epsilon \leq \nu$. Letting ϵ tend to zero implies $\frac{1}{2}I \leq \nu$ and hence $\frac{1}{2}I \leq \operatorname{geod}(X)$, because $\operatorname{geod}(X)$ is the infimum over all ν such that X is ν-almost geodesic. □

A metric subset $X \subseteq Y$ is *r-dense* if for every $y \in Y$ there exists $x \in X$ with $d(y,x) \leq r$. The following proposition proves an upper bound on the geodesic defect for r-dense subsets of a geodesic space. A partial converse for δ-hyperbolic spaces is given by Proposition 3.10.

Proposition 3.5. *Let X be an r-dense subset of a geodesic metric space Y. Then X is r-almost geodesic. In particular,* $\operatorname{geod}(X) \leq r$.

Proof. Let $x, y \in X$ be any two points and let $t, s \geq 0$ be with $t + s = d(x, y)$. Choose an isometric embedding $\gamma\colon [0, d(x, y)] \to Y$ with $\gamma(0) = x$ and $\gamma(d(x, y)) = y$. As X is r-dense, there exists a point $z \in X$ with $d(\gamma(t), z) \leq r$. By the triangle inequality, we get $d(x, z) \leq d(\gamma(0), \gamma(t)) + d(\gamma(t), z) \leq t + r$ and similarly $d(y, z) \leq s + r$. This proves that X is r-almost geodesic. □

A 0-almost geodesic space is also called *metrically convex* [EK01]. If X is a geodesic space, then its geodesic defect is zero, but the converse is not always true. In fact, for any *length space* [BBI01], i.e., a metric space where the distance between two points is the infimum of lengths of paths connecting them, the geodesic defect is zero. It is worth noting that a complete metric space is a length space if and only if its geodesic defect is zero, and it is a geodesic space if and only if it is 0-almost geodesic [BBI01, Section 2.4]. The punctured unit disk in $\mathbb{R}^2$ is an example for a space that has geodesic defect 0 but that is not 0-almost geodesic. However, we have the following.

Proposition 3.6. *Let X be a proper metric space, i.e., assume that every closed ball is compact. Then X is* $\operatorname{geod}(X)$*-almost geodesic.*

Proof. Let $x, y \in X$ be two points and $r, s \geq 0$ with $r + s = d(x, y)$. For every natural number $n \in \mathbb{N}$ the space X is ν_n-almost geodesic, where $\nu_n = \operatorname{geod}(X) + \frac{1}{n}$, and hence there exists a point $z_n \in X$ with $d(x, z_n) \leq r + \nu_n$ and $d(y, z_n) \leq s + \nu_n$. This sequence is contained in the closed ball of radius $r + \operatorname{geod}(X) + 1$ centered at x, which is compact by assumption. Hence, there exists a convergent subsequence $z_{n_k} \to z$. The limit point $z \in X$ satisfies $d(x, z) \leq r + \operatorname{geod}(X)$ and $d(y, z) \leq s + \operatorname{geod}(X)$. Therefore, X is $\operatorname{geod}(X)$-almost geodesic. □

We now show that the geodesic defect, like the hyperbolicity, is a Gromov–Hausdorff stable quantity.

Proposition 3.7. *Let X and Y be any metric spaces and $s > d_{\mathrm{GH}}(X, Y)$. If X is ν-almost geodesic, then Y is $(\nu + 3s)$-almost geodesic. Hence,*

$$|\operatorname{geod}(X) - \operatorname{geod}(Y)| \leq 3 d_{\mathrm{GH}}(X, Y).$$

Proof. Let C be a correspondence between X and Y with $\operatorname{dis} C \leq 2s$. Furthermore, let $y, y' \in Y$ be any two points, and let $r, t \geq 0$ be such that $r + t = d(y, y')$. Choose two corresponding points $x, x' \in X$, with $(x, y), (x', y') \in C$. For $u = r + \operatorname{dis} C/2$ and $w = t + \operatorname{dis} C/2$, we have $u + w \geq d(x, x')$. As X is ν-almost geodesic, there exists a point $z \in X$ such that $d(x, z) \leq u + \nu$ and $d(x', z) \leq w + \nu$. We can choose a corresponding point $p \in Y$, with $(z, p) \in C$. For this point, we get

$$d(y, p) \leq d(x, z) + \operatorname{dis} C \leq (u + \nu) + 2s = (r + \operatorname{dis} C/2) + \nu + 2s \leq r + \nu + 3s$$

and similarly $d(y', p) \leq t + \nu + 3s$. Thus Y is $(\nu + 3s)$-almost geodesic. □

If X is an r-dense subset of a geodesic space Y, then $d_{\mathrm{GH}}(X, Y) \leq r$, and the above proposition implies that X is $(3s)$-almost geodesic for every $s > r$. In particular, $\operatorname{geod}(X) \leq 3r$. Note however that in this case Proposition 3.5 gives the stronger bound $\operatorname{geod}(X) \leq r$.

For length spaces, it is known that all structure maps in the first persistent homology of the Vietoris–Rips filtration are surjective [CSO14, Corollary 6.2]. This statement generalizes to arbitrary metric spaces using the geodesic defect as follows.

Proposition 3.8. *Let X be a ν-almost geodesic metric space. For every $t > u > 2\nu$ the canonical map $H_1(\mathrm{Rips}_u(X)) \to H_1(\mathrm{Rips}_t(X))$ is surjective.*

Proof. Let $\{x, y\} \in \mathrm{Rips}_t(X)$ be an edge with length $d(x, y) = t$. As X is ν-almost geodesic, there exists a point $z \in X$ with $d(x, z) \leq \frac{1}{2}t + \nu$ and $d(y, z) \leq \frac{1}{2}t + \nu$. By assumption, we have $\frac{1}{2}t + \nu = \frac{1}{2}t + \frac{1}{2}u - (\frac{1}{2}u - \nu) < t - l$, where $l = (\frac{1}{2}u - \nu)$. Hence, the simplex $\{z, x, y\}$ is contained in $\mathrm{Rips}_t(X)$ and the simplicial chain $[x, y]$ is homologous to $[z, y] - [z, x] \in C_1(\mathrm{Rips}_{t-l}(X))$. As every simplicial chain in $C_1(\mathrm{Rips}_t(X))$ is a finite sum of edges and $l > 0$ is a constant, it follows that finitely many reapplications of the argument above yields that this chain is homologous to a chain in $C_1(\mathrm{Rips}_u(X))$, proving the claim. ☐

Remark 3.9. We briefly remark on how the hyperbolicity and geodesic defect can be computed for a finite metric space $X = \{x_1, \ldots, x_n\}$ with n points. By an expression similar to Eq. (3.1), the hyperbolicity of X with respect to some point $p \in X$ can be computed by brute force in $\Theta(n^3)$ time. More efficiently [FIV15; Dua14]: The hyperbolicity of X with respect to p is given by the largest entry of the matrix $A \odot A - A$, where A is the matrix with $A_{i,j}$ equal to the *Gromov product* of x_i and x_j,

$$A_{i,j} = \frac{1}{2}(d(x_i, p) + d(x_j, p) - d(x_i, x_j)) = (x_i, x_j)_p,$$

and $A \odot A$ is the $(\max, \min)$-*product* given by $(A \odot A)_{i,j} = \max_k \min\{A_{i,k}, A_{k,j}\}$. The $(\max, \min)$-product can be computed in $O(n^{(3+\omega)/2})$ time [DP09] with ω such that two arbitrary real $m \times m$ square matrices can be multiplied in $O(m^\omega)$ time, which is known to be possible with $2 \leq \omega < 3$.

We now describe a brute force way to compute the geodesic defect of X with respect to some point $p \in X$. First, fix any point $x \in X$, let $r, s \geq 0$ be any real numbers such that $r + s = d(x, p)$, and let $\nu \geq 0$ be arbitrary. If there exists a point $z \in X$ such that $d(z, x) \leq r + \nu$ and $d(z, p) \leq s + \nu$, then we have

$$\max\{d(z, x) - r, d(z, p) - d(x, p) + r\} \leq \nu.$$

Thus, the optimal value for ν, such that a point $z \in X$ as in the inequality above exists for any $r, s \geq 0$ with $r + s = d(x, p)$, is given by

$$\nu^x = \max_{r \in [0, d(x,p)]} \min_{z \in X} f_z^x(r), \tag{3.2}$$

where

$$f_z^x(r) = \max\{d(z,x) - r, d(z,p) - d(x,p) + r\} \tag{3.3}$$

is a piecewise linear function $\mathbb{R} \to \mathbb{R}$ in the variable r. The function f_z^x attains its minimum at

$$r_z^x = \frac{1}{2}(d(z,x) + d(x,p) - d(z,p)),$$

which is the Gromov product $(z,p)_x$. Note that $\min_{z \in X} f_z^x(r)$ is also a piecewise linear function in the variable r. For any two points $y, z \in X$ the corresponding functions f_y^x and f_z^x, as in Eq. (3.3), can potentially intersect at

$$r_{y,z}^x = \frac{1}{2}(d(y,x) + d(x,p) - d(z,p)),$$

or at $r_{z,y}^x$. Thus, to compute the maximum of the function $\min_{z \in X} f_z^x$, as in Eq. (3.2), we discretize the interval $[0, d(x,p)]$ by only taking the values

$$D^x = \{r_{y,z}^x \mid y, z \in X\} \cap [0, d(x,p)]$$

into account. For each discretization point $l \in D^x$, we iterate over all points in X, to obtain the minimum $M_l = \min_{z \in X} f_z^x(l)$. The value ν^x is then given by $\nu^x = \max_{l \in D^x} M_l$, and the geodesic defect of X with respect to p is given by $\mathrm{geod}_p(X) = \max_{x \in X} \nu^x$. We remark that in each step it suffices to only take the points of X into account that are contained in the *lune* of x and p, given by $\{y \in X \mid d(y,x) \leq d(x,p),\ d(y,p) \leq d(x,p)\}$. Thus, the geodesic defect of X with respect to p can be computed in $O(n^4)$ time. To the best of the author's knowledge, the geodesic defect has not been study from a computational point of view before. It remains an interesting open problem to find better algorithms for computing the geodesic defect of a finite metric space.

3.2. Rips' Contractibility Lemma via the Injective Hull

We adapt some known facts about embeddings of metric spaces into their injective hull (see Section 2.1.2) using the geodesic defect, to prove a version of the Contractibility Lemma for δ-hyperbolic ν-almost geodesic metric spaces, following [LMO22].

The following is essentially due to Lang [Lan13]. Originally, it has been stated for a special case, but the proof applies verbatim to the below statement involving the geodesic defect, which indeed provided the motivation for our definition. Note that the definition of δ-hyperbolic used in [Lan13] differs from the one used here by a factor of 2.

Proposition 3.10 (Lang [Lan13, Proposition 1.3])**.** *Let X be a δ-hyperbolic ν-almost geodesic metric space. Then the injective hull $E(X)$ is δ-hyperbolic, and every point in $E(X)$ has distance at most $2\delta + \nu$ to $e(X)$.*

This result yields a generalization of the Contractibility Lemma using the injective hull analogously to the proof for geodesic spaces in [LMO22, Corollary 8.1].

Theorem 3.11. *Let X be a δ-hyperbolic ν-almost geodesic metric space. Then the complex $|\mathrm{Rips}_t^<(X)|$ is contractible for every $t > 0$ with $t > 4\delta + 2\nu$.*

Proof. By Proposition 3.10, we know that the collection of open balls with radius $\frac{t}{2}$ centered at the points in $e(X)$ covers $E(X)$. The nerve of this cover is isomorphic to

$$\check{\mathrm{C}}\mathrm{ech}^{<}_{\frac{t}{2}}(e(X), E(X)).$$

Thus, Proposition 2.16, Propositions 2.15 and 2.17, and the nerve theorem (Theorem H) imply

$$|\,\mathrm{Rips}^{<}_{t}(X)| = |\,\check{\mathrm{C}}\mathrm{ech}^{<}_{\frac{t}{2}}(e(X), E(X))| \simeq E(X) \simeq *. \qquad \square$$

Remark 3.12. Proposition 2.17 implies that for any cover of an injective metric Y by closed balls, there exists a geodesic bicombing $\sigma \colon Y \times Y \times [0,1] \to Y$ that restricts to a map as assumed in Remark 2.6. Thus, if the metric space X in Theorem 3.11 is finite, then, in the proof of said statement, we can replace the nerve theorem (Theorem H) by the nerve theorem in Remark 2.6, and conclude similarly that $|\,\mathrm{Rips}_t(X)|$ is contractible for every $t > 0$ with $t \geq 4\delta + 2\nu$.

3.3. Filtered Collapsibility of Vietoris–Rips Complexes

We revisit the original proof of the Contractibility Lemma in [Gro87], adapted to discrete Morse theory (see Section 2.3). We extend the statement beyond geodesic spaces using the geodesic defect, strengthen the assertion of contractibility to collapsibility, and further extend the result to become compatible with the Vietoris–Rips filtration.

We first establish a key fact that is essential in the construction of the discrete gradients in the proofs of Theorems 3.14 and 3.16.

Lemma 3.13. *Let X be a metric space that is δ-hyperbolic and ν-almost geodesic with respect to some point $p \in X$. Then for any point $x \in X$ with $d(x,p) \geq 2\delta + \nu$ there is a point $z_x \in X$ such that*

- *$d(z_x,p) \leq d(x,p) - 2\delta$ and $d(z_x,x) \leq 2\delta + 2\nu$,*
- *for any $t \geq 4\delta + 2\nu$ and $y \in X$ with $d(y,p) \leq d(x,p)$ and $d(y,x) \leq t$, we have $d(y,z_x) \leq t$,*
- *for any $t > 4\delta + 2\nu$ and $y \in X$ with $d(y,p) \leq d(x,p)$ and $d(y,x) < t$, we have $d(y,z_x) < t$.*

Proof. For $r = 2\delta + \nu$ and $s = d(x,p) - 2\delta - \nu$ we have $r + s = d(x,p)$, and therefore, by the assumption that X is a ν-almost geodesic space with respect to p, there is a point $z_x \in X$ with

$$d(z_x,p) \leq s + \nu = d(x,p) - 2\delta,$$

and

$$d(z_x,x) \leq r + \nu = 2\delta + 2\nu = (4\delta + 2\nu) - 2\delta. \tag{3.4}$$

This proves the first claim. We illustrate the situation as follows:

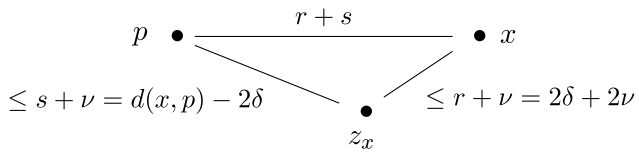

We now prove the second and third claim. Note that if $t \geq 4\delta + 2\nu$, then Eq. (3.4) implies

$$d(z_x, x) \leq t - 2\delta,$$

and if $t > 4\delta + 2\nu$, then Eq. (3.4) implies

$$d(z_x, x) < t - 2\delta.$$

By assumption, we have

$$d(y, p) \leq d(x, p),$$

and we either have

$$d(y, x) \leq t \quad \text{or} \quad d(y, x) < t.$$

Recall that the four-point condition (1.1) states

$$d(y, z_x) \leq \max\{d(y, x) + d(z_x, p),\ d(y, p) + d(z_x, x)\} + 2\delta - d(x, p).$$

For the non-strict case, the four-point condition together with the previous inequalities yields

$$d(y, z_x) \leq \max\{t + (d(x, p) - 2\delta), d(x, p) + (t - 2\delta)\} + 2\delta - d(x, p) = t,$$

which proves the second claim. We illustrate the situation as follows:

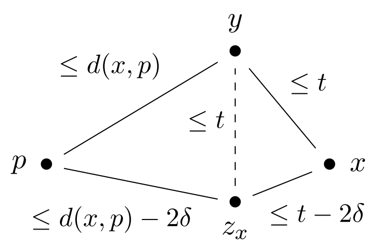

For the strict case, the four-point condition together with the previous inequalities yields

$$d(y, z_x) \leq \max\{\underbrace{d(y, x)}_{<t} + \underbrace{d(z_x, p) - d(x, p)}_{\leq -2\delta},\ \underbrace{d(y, p) - d(x, p)}_{\leq 0} + \underbrace{d(z_x, x)}_{<t-2\delta}\} + 2\delta < t,$$

which proves the third claim. $\square$

Let X be a metric space and $p \in X$ a point. Consider the sublevel set filtration of the distance function $d(-,p)\colon X \to \mathbb{R}$, which induces, for every $t \geq 0$, a sublevel set filtration of the Vietoris–Rips complex by the function

$$\tilde{d}_p\colon \operatorname{Rips}_t(X) \to \mathbb{R},\ \sigma \mapsto \max_{y\in\sigma} d(y,p).$$

Recall that we write $D_r(p) = \{y \in X \mid d(p,y) \leq r\}$ and $B_r(p) = \{y \in X \mid d(p,y) < r\}$ for the closed metric ball and the open metric ball of radius r centered at p, respectively.

Theorem 3.14. *Let X be a metric space that is δ-hyperbolic and ν-almost geodesic with respect to some point $p \in X$. Then for every $t > 2\nu$ with $t \geq 4\delta + 2\nu$ there exists a discrete gradient on the Vietoris–Rips complex* $\operatorname{Rips}_t(X)$

- *that is compatible with the sublevel set filtration of $\tilde{d}_p$,*
- *whose only critical simplex is $\{p\}$, and*
- *such that every simplex has finite height.*

In particular, the discrete gradient induces a collapse

$$\operatorname{Rips}_t(X) \searrow \{p\},$$

and for every $r > l \geq 0$, the collapses

$$\operatorname{Rips}_t(D_r(p)) \searrow \operatorname{Rips}_t(B_r(p)) \searrow \operatorname{Rips}_t(D_l(p)).$$

Proof. Without loss of generality, assume that $\delta > 0$; if X is 0-hyperbolic with respect to p, then it is also δ-hyperbolic with respect to p for any $\delta > 0$. As $t > 2\nu$, we may simply choose $\delta = \frac{t-2\nu}{4}$, which still satisfies $t \geq 4\delta + 2\nu$.

Extend the preorder on X induced by the sublevel set filtration of $d(-,p)$ to a total order $<$, so that $x < y$ implies $d(x,p) \leq d(y,p)$. Every simplex $\sigma \in \operatorname{Rips}_t(X)$ has a maximal vertex with respect to this total order, $\max\sigma$. For every point $x \in X \setminus \{p\}$ we construct a discrete gradient V_x that is compatible with the sublevel set filtration of $\tilde{d}_p$ and induces a collapse $K_{\leq x} \searrow K_{<x}$, where

$$K_{\leq x} = \{\sigma \in \operatorname{Rips}_t(X) \mid \max\sigma \leq x\} \quad \text{and} \quad K_{<x} = \{\sigma \in \operatorname{Rips}_t(X) \mid \max\sigma < x\}.$$

Note that the complement $K_{\leq x} \setminus K_{<x}$ consists of all simplices of $\operatorname{Rips}_t(X)$ that contain x as the maximal vertex.

First assume $d(x,p) < t$. Then for any vertex z of $K_{\leq x}$ we have $d(z,p) \leq d(x,p) < t$, so $\{z,p\}$ is a 1-simplex in $K_{\leq x}$. This implies that for every simplex $\sigma \in K_{\leq x}$ we also have $\sigma \cup \{p\} \in K_{\leq x}$, meaning that $K_{\leq x}$ is a simplicial cone with apex p. Pairing the simplices containing p with those not containing p, we obtain a discrete gradient inducing a collapse $K_{\leq x} \searrow K_{<x}$:

$$V_x = \{(\sigma \setminus \{p\}, \sigma \cup \{p\}) \mid \sigma \in K_{\leq x} \setminus K_{<x}\}. \tag{3.5}$$

Now assume $d(x,p) \geq t \geq 4\delta + 2\nu$. By Lemma 3.13, there is a point $z_x \in X$ with $d(z_x,p) \leq d(x,p) - 2\delta$, implying $z_x < x$. We show that for every simplex $\sigma \in K_{\leq x} \setminus K_{<x}$ the union $\sigma \cup \{z_x\}$ is also a simplex in $K_{\leq x} \setminus K_{<x}$: Note first that for any $y \in \sigma$ we have $y \leq \max \sigma = x$ and hence $d(y,p) \leq d(x,p)$ by construction of the total order on X. As $x \in \sigma$ by assumption, we also have $d(y,x) \leq \operatorname{diam} \sigma \leq t$. Since $t \geq 4\delta + 2\nu$, it therefore follows from Lemma 3.13 that any vertex y of σ has distance $d(y,z_x) \leq t$ to z_x, showing that $\sigma \cup \{z_x\}$ is also a simplex in $K_{\leq x} \setminus K_{<x}$. Similarly to the above, pairing the simplices containing z_x with those not containing z_x yields a discrete gradient inducing a collapse $K_{\leq x} \searrow K_{<x}$:

$$V_x = \{(\sigma \setminus \{z_x\}, \sigma \cup \{z_x\}) \mid \sigma \in K_{\leq x} \setminus K_{<x}\}. \tag{3.6}$$

By Proposition 2.47, the union $V = \bigcup_x V_x$ is a discrete gradient on $\operatorname{Rips}_t(X)$ whose only critical simplex is $\{p\}$. By Lemma 2.40 and Proposition 2.43, the discrete gradient V induces the collapses in the statement of Theorem 3.14 if every simplex $\sigma \in \operatorname{Rips}_t(X)$ has finite V-path height, which we now show. Let

$$\sigma = \sigma_0 \to \tau_0 \leftarrow \cdots \to \tau_{n-1} \leftarrow \sigma_n$$

be any V-path in $\operatorname{Rips}_t(X)$ with $n \geq 1$, implying that $\sigma \neq \{p\}$.

Before providing an upper bound on the length of the gradient path, we consider a single step in the gradient path more closely. Assume first that σ_i is such that $d(\max \sigma_i, p) < t$. By construction of the discrete gradient in Eq. (3.5), we know that σ_i is a gradient cofacet if $p \in \sigma_i$, and hence in this case the gradient path must end in σ_i, implying $n = i$. If $p \notin \sigma_i$, we must have that σ_i is paired with

$$\tau_i = \sigma_i \cup \{p\}$$

and the facet σ_{i+1} of τ_i, which is different from σ_i, contains p. Hence, the simplex σ_{i+1} is a gradient cofacet of V_x from Eq. (3.5), since $d(\max \sigma_{i+1}, p) \leq d(\max \sigma_i, p) < t$ and $p \in \sigma_{i+1}$. Thus, the gradient path must end in σ_{i+1}, implying $n = i+1$.

Now assume that σ_i is such that $d(\max \sigma_i, p) \geq t$. Recall that $z_{\max \sigma_i}$ is the point from Lemma 3.13 that satisfies $d(z_{\max \sigma_i}, p) \leq d(\max \sigma_i, p) - 2\delta$ and $z_{\max \sigma_i} < \max \sigma_i$. By construction of the discrete gradient in Eq. (3.6), we know that σ_i is a gradient cofacet if $z_{\max \sigma_i} \in \sigma_i$, and hence in this case the gradient path must end in σ_i, implying $n = i$. If $z_{\max \sigma_i} \notin \sigma_i$, we must have that σ_i is paired with

$$\tau_i = \sigma_i \cup \{z_{\max \sigma_i}\}.$$

Now, either the gradient path ends in σ_{i+1}, implying $n = i+1$, or σ_{i+1} is the unique facet of τ_i different from σ_i that is paired with a gradient cofacet τ_{i+1} in V, namely the unique other facet not simultaneously containing $z_{\max \sigma_i}$ and $\max \sigma_i$. It is given by

$$\sigma_{i+1} = \tau_i \setminus \{\max \sigma_i\}.$$

In particular, as $z_{\max \sigma_i} < \max \sigma_i$ and $\max \sigma_i \notin \sigma_{i+1}$, we get $\max \sigma_{i+1} < \max \sigma_i$. This paragraph can be summarized as follows: In each step of the gradient path, where σ_i is such that $d(\max\sigma_i, p) \geq t$, the furthest point to p, namely $\max\sigma_i$, is replaced with the point $z_{\max\sigma_i}$ that satisfies $d(z_{\max\sigma_i}, p) \leq d(\max\sigma_i, p) - 2\delta$.

We now provide an upper bound on the length of the gradient path. Assume without loss of generality that $n > \operatorname{card}\sigma$. We know that $d(\max\sigma_i, p) \geq t$ for all $0 \leq i < n-1$, as otherwise $n \leq i+1 < n$ by the first case in the previous discussion. Note also that for all $y \in \sigma$ we have $y \leq \max\sigma$ and hence $d(y,p) \leq d(\max\sigma, p)$, or equivalently, $d(y,p) - 2\delta \leq d(\max\sigma, p) - 2\delta$. It now follows from the second case in the previous discussion, that after at most $\operatorname{card}\sigma$ steps in the gradient path each element $y \in \sigma$ with $d(y,p) > d(\max\sigma, p) - 2\delta$ got replaced with a point $z_y \in X$ such that

$$d(z_y, p) \leq d(y,p) - 2\delta \leq d(\max\sigma, p) - 2\delta.$$

In particular, this implies

$$0 \leq d(\max\sigma_{\operatorname{card}\sigma}, p) \leq d(\max\sigma, p) - 2\delta. \tag{3.7}$$

Moreover, for the largest natural number m with $m \cdot \operatorname{card}\sigma - 1 < n - 1$ we get the estimate

$$0 \leq d(\max\sigma_{m\cdot\operatorname{card}\sigma}, p) \leq d(\max\sigma_{(m-1)\cdot\operatorname{card}\sigma}, p) - 2\delta \leq \cdots \leq d(\max\sigma_0, p) - m\cdot 2\delta,$$

where we used that $\operatorname{card}\sigma_i = \operatorname{card}\sigma_0 = \operatorname{card}\sigma$ for all i, and applied Eq. (3.7) to the truncated gradient path $\sigma_{(m-j)\cdot\operatorname{card}\sigma} \to \tau_{(m-j)\cdot\operatorname{card}\sigma} \leftarrow \cdots \to \tau_{n-1} \leftarrow \sigma_n$ for all j with $m \geq j \geq 1$. Since $\sigma_0 = \sigma$, we equivalently have

$$m \leq \frac{d(\max\sigma, p)}{2\delta},$$

and thus we get the estimate

$$n - 1 \leq (m+1)\cdot\operatorname{card}\sigma - 1 \leq (\frac{d(\max\sigma, p)}{2\delta} + 1)\cdot\operatorname{card}\sigma - 1. \tag{3.8}$$

This shows that the maximal length of a gradient path starting in σ is bounded from above by a constant only dependent on σ and δ, implying that the V-path height of σ is finite. □

Remark 3.15. For a finite simplicial complex K, Barmak and Minian [BM12] define a particular type of simplicial collapse, called *elementary strong collapse*, which is a collapse from K to $K \setminus \operatorname{St}_K(v)$ such that the link of the vertex v is a simplicial cone. If X is finite, the proof of Theorem 3.14 also shows that for $t \geq 4\delta + 2\nu$ there exists a sequence of elementary strong collapses from $\operatorname{Rips}_t(X)$ to $\{*\}$.

We now extend the proof strategy of Theorem 3.14 to obtain a filtration-compatible strengthening of the Contractibility Lemma. Let X be a metric space and $p \in X$ a point. Consider the two-parameter sublevel set filtration induced by

$$(\operatorname{diam}, \tilde{d}_p)\colon \operatorname{Cl}(X) \to \mathbb{R}^2,\ \sigma \mapsto (\operatorname{diam}\sigma, \max_{y\in\sigma} d(y,p))$$

on the full simplicial complex on X, where $\mathbb{R}^2 = \mathbb{R}\times\mathbb{R}$ carries the product partial order.

Theorem 3.16. *Let X be a metric space that is δ-hyperbolic and ν-almost geodesic with respect to some point $p \in X$. Then for every $s > 2\nu$ with $s \geq 4\delta + 2\nu$ there exists a discrete gradient on the full simplicial complex* $\mathrm{Cl}(X)$

- *that is compatible with the two-parameter sublevel set filtration of* $(\mathrm{diam}, \tilde{d}_p)$*,*
- *whose only critical simplices are* $\mathrm{Rips}_s(X)$*, and*
- *such that every simplex has finite height.*

In particular, the discrete gradient induces, for every $t > u \geq s$, the collapses

$$\mathrm{Rips}_t(X) \searrow \mathrm{Rips}_t^<(X) \searrow \mathrm{Rips}_u(X),$$

and, for every $r > l \geq 0$, the collapses

$$\mathrm{Rips}_t(D_r(p)) \searrow \mathrm{Rips}_t(B_r(p)) \cup \mathrm{Rips}_s(D_r(p)) \searrow \mathrm{Rips}_t(D_l(p)) \cup \mathrm{Rips}_s(D_r(p)).$$

Proof. Without loss of generality, assume that $\delta > 0$; if X is 0-hyperbolic with respect to p, then it is also δ-hyperbolic with respect to p for any $\delta > 0$. As $s > 2\nu$, we may simply choose $\delta = \frac{s-2\nu}{4}$, which still satisfies $s \geq 4\delta + 2\nu$.

Extend the preorder on X induced by the sublevel set filtration of $d(-,p)$ to a total order $<$, so that $x < y$ implies $d(x,p) \leq d(y,p)$. Every simplex $\sigma \in \mathrm{Cl}(X)$ has a maximal vertex with respect to this total order, $\max \sigma$. For every $t > s$, we construct a discrete gradient W_t that is compatible with the sublevel set filtration of $(\mathrm{diam}, \tilde{d}_p)$. Further, we show that the union $W = \bigcup_{s<t} W_t$ is a discrete gradient on $\mathrm{Cl}(X)$ that induces the claimed collapses. To this end, for every point $x \in X \setminus \{p\}$, we construct a discrete gradient V_x^t on $K_{\leq x}^t$ that induces a collapse $K_{\leq x}^t \searrow K_{<x}^t$, where

$$K_{\leq x}^t = \{\sigma \in \mathrm{Rips}_t(X) \mid \max \sigma \leq x\} \quad \text{and} \quad K_{<x}^t = \{\sigma \in K_{\leq x}^t \mid \mathrm{diam}\, \sigma < t \text{ or } \max \sigma < x\}.$$

Note that $K_{\leq x}^t \setminus K_{<x}^t$ consists of all simplices of $\mathrm{Rips}_t(X)$ with diameter t that contain x as the maximal vertex.

First assume $d(x,p) < t$. Let $\sigma \in K_{\leq x}^t \setminus K_{<x}^t$. As x is the maximal vertex of σ, we have $d(v,p) \leq d(x,p) < t$ for all $v \in \sigma$. Since σ has diameter t, this implies that $\sigma \cup \{p\}$ also has diameter $t$and contains an edge $e \subseteq \sigma \setminus \{p\} \subseteq \sigma$ not containing p with $\mathrm{diam}\, e = t$. Therefore, $\sigma \setminus \{p\}$ also has diameter t. As $p < x$, both simplices $\sigma \setminus \{p\}$ and $\sigma \cup \{p\}$ contain x as the maximal vertex and are thus contained in $K_{\leq x}^t \setminus K_{<x}^t$. Pairing the simplices containing p with those not containing p, we obtain a discrete gradient inducing a collapse $K_{\leq x}^t \searrow K_{<x}^t$:

$$V_x^t = \{(\sigma \setminus \{p\}, \sigma \cup \{p\}) \mid \sigma \in K_{\leq x}^t \setminus K_{<x}^t\}. \tag{3.9}$$

Now assume $d(x,p) \geq t > s \geq 4\delta + 2\nu$. By Lemma 3.13, there is a point $z_x \in X$ with $d(z_x,p) \leq d(x,p) - 2\delta$, implying $z_x < x$. We show that for every simplex $\sigma \in K_{\leq x}^t \setminus K_{<x}^t$, the simplices $\sigma \setminus \{z_x\}$ and $\sigma \cup \{z_x\}$ are also contained in $K_{\leq x}^t \setminus K_{<x}^t$: To this end, we show first that any vertex y of σ has distance $d(y, z_x) \leq t$ to z_x. Note that for any $y \in \sigma$ we

have $y \leq \max \sigma = x$ and hence $d(y,p) \leq d(x,p)$ by construction of the total order on X. As $x \in \sigma$ by assumption, we also have $d(y,x) \leq \operatorname{diam} \sigma = t$. Since $t > s \geq 4\delta + 2\nu$, it therefore follows from Lemma 3.13 that any vertex y of σ has distance $d(y,z_x) \leq t$ to z_x. Moreover, it follows that if $d(y,x) < t$, then $d(y,z_x) < t$. Hence, $\operatorname{diam}(\sigma \cup \{z_x\}) = t$, and $\operatorname{diam} \sigma = t$ implies $\operatorname{diam} \sigma \setminus \{z_x\} = t$, by an argument similar to the above in the case that $d(x,p) < t$. As $z_x < x$, both simplices $\sigma \setminus \{z_x\}$ and $\sigma \cup \{z_x\}$ contain x as the maximal vertex and are thus contained in $K^t_{\leq x} \setminus K^t_{<x}$. Pairing the simplices containing z_x with those not containing z_x, we obtain a discrete gradient inducing a collapse $K^t_{\leq x} \searrow K^t_{<x}$:

$$V^t_x = \{(\sigma \setminus \{z_x\}, \sigma \cup \{z_x\}) \mid \sigma \in K^t_{\leq x} \setminus K^t_{<x}\}. \tag{3.10}$$

By Proposition 2.47, for every $t > s$, the union $W_t = \bigcup_x V^t_x$ is a discrete gradient on $\operatorname{Rips}_t(X)$, and similarly, the union $W = \bigcup_{s<t} W_t$ is a discrete gradient on $\operatorname{Cl}(X)$ whose only critical simplices are $\operatorname{Rips}_s(X)$. By Lemma 2.40 and Proposition 2.43, the discrete gradient W induces the collapses in the statement of Theorem 3.16 if every simplex $\sigma \in \operatorname{Cl}(X)$ has finite W-path height, which we now show. Let

$$\sigma = \sigma_0 \to \tau_0 \leftarrow \cdots \to \tau_{n-1} \leftarrow \sigma_n$$

be any W-path in $\operatorname{Cl}(X)$ with $n \geq 1$, implying that $\operatorname{diam} \sigma > s$.

We provide an upper bound on the length of the gradient path, by showing that it decomposes into two gradient paths, each of which is a valid gradient path in the proof of Theorem 3.14. For this, we make use of the fact that the gradients in Eqs. (3.5) and (3.6) and Eqs. (3.9) and (3.10), respectively, are constructed in essentially the same way. The lengths of these two gradient paths are bounded from above by

$$B := \left(\frac{d(\max \sigma, p)}{2\delta} + 1\right) \cdot \operatorname{card} \sigma,$$

as in Eq. (3.8). This shows that $n \leq 2B$ and the maximal length of a W-path starting in σ is bounded from above by a constant only dependent on σ and δ, implying that the W-path height of σ is finite. Note first that along the W-path above the sequence of real numbers $(\operatorname{diam} \sigma_i)_i$ is monotonously decreasing. If for all $i < n$ we have $d(\max \sigma_i, p) \geq \operatorname{diam} \sigma_i$, then the entire gradient path is a valid gradient path in the proof of Theorem 3.14, with $t = \operatorname{diam} \sigma_{n-1}$, and we are done. Otherwise, let $i < n$ be the smallest integer such that $d(\max \sigma_i, p) < \operatorname{diam} \sigma_i$. We now decompose the gradient path into two parts. The first part is the truncated gradient path

$$\sigma_0 \to \tau_0 \leftarrow \cdots \to \tau_i \leftarrow \sigma_{i+1},$$

where $\tau_i = \sigma_i \cup \{p\}$, $p \in \sigma_{i+1}$, and which is a valid gradient path in the proof of Theorem 3.14 with $t = \operatorname{diam} \sigma_i$. The second part is the truncated gradient path

$$\sigma_{i+1} \to \cdots \leftarrow \sigma_n, \tag{3.11}$$

which possibly only consists of a single simplex; we therefore assume without loss of generality that $i + 1 < n$. We argue that this second gradient path is also valid. Note that we must have $p \in \sigma_j$ for all j with $i+1 \leq j \leq n$, by construction of the discrete gradients in Eqs. (3.9) and (3.10). As $i+1 < n$, we have $d(\max \sigma_{i+1}, p) \geq \operatorname{diam} \sigma_{i+1}$, since otherwise the simplex σ_{i+1} is either critical or a gradient cofacet in $V_x^t \subseteq W$ from Eq. (3.9), where $t = \operatorname{diam} \sigma_{i+1}$ and $x = \max \sigma_{i+1}$, which implies either way that the gradient path must end in σ_{i+1} and yields the contradiction $n = i+1 < n$. Similar to before, if for all $i+1 \leq k \leq n$ we have $d(\max \sigma_k, p) \geq \operatorname{diam} \sigma_k$, then the gradient path in Eq. (3.11) is a valid gradient path in the proof of Theorem 3.14, with $t = \operatorname{diam} \sigma_{n-1}$, and we are done. Otherwise, let $i+1 < k \leq n$ be the smallest integer such that $d(\max \sigma_k, p) < \operatorname{diam} \sigma_k$. As $p \in \sigma_k$, the simplex σ_k is either critical or a gradient cofacet in $V_x^t \subseteq W$ from Eq. (3.9), where $t = \operatorname{diam} \sigma_k$ and $x = \max \sigma_k$. Either way, the gradient path must end in σ_k, implying that $n = k$. In particular, the gradient path in Eq. (3.11) is a valid gradient path in the proof of Theorem 3.14, with $t = \operatorname{diam} \sigma_{n-1}$, proving the claim. ☐

Remark 3.17. We are not aware of an example of a space with positive hyperbolicity showing that this bound is tight. However, it is tight for every finite tree metric, as can be deduced from Example 1.6.

3.4. Collapsing Vietoris–Rips Complexes of Trees by Apparent Pairs

In this section, we analyze the Vietoris–Rips filtration of a tree metric space X for a positively weighted finite tree $T = (X, E)$, with the goal of proving the collapses in Theorem B using the apparent pairs gradient (see Section 2.3.1). To this end, we introduce two other discrete gradients: the *canonical gradient*, which is independent of any choices, and the *perturbed gradient*, which coarsens the canonical gradient and can be interpreted as a gradient that arises through a symbolic perturbation of the edge lengths. We then show that the intervals in the perturbed gradient are refined by apparent pairs of the lexicographically refined Vietoris–Rips filtration, with respect to a particular total order on the vertices.

Recall that we write $D_r(p) = \{y \in X \mid d(p,y) \leq r\}$ and $S_r(p) = \{y \in X \mid d(p,y) = r\}$ for the closed metric ball and the sphere of radius r centered at p, respectively.

Lemma 3.18. *Let $x, y \in X$ be two distinct points at distance $d(x,y) = r$. Then we have* $\operatorname{diam} D_r(x) \cap D_r(y) = r$*. Furthermore, if $a, b \in D_r(x) \cap D_r(y)$ are points with $d(a,b) = r$, then these points are contained in the union $S_r(x) \cup S_r(y)$.*

Proof. We start by showing the first claim. Let $a, b \in D_r(x) \cap D_r(y)$ be any two points. We show that $d(a,b) \leq r$ holds, implying $\operatorname{diam} D_r(x) \cap D_r(y) \leq r$. Because $x, y \in D_r(x) \cap D_r(y)$ we also have $\operatorname{diam} D_r(x) \cap D_r(y) \geq r$, proving equality.

Write $[n] = \{1, \ldots, n\}$ and let $\gamma \colon ([n], \{\{i, i+1\} \mid i \in [n-1]\}) \to T$ be the unique shortest path $x \rightsquigarrow y$. Moreover, let Ψ_a and Ψ_b be the unique shortest paths $x \rightsquigarrow a$ and $x \rightsquigarrow b$, respectively. Consider the largest numbers $t_a, t_b \in [n]$ with $\gamma(t_a) = \Psi_a(t_a)$ and $\gamma(t_b) = \Psi_b(t_b)$ and assume without loss of generality $t_a \leq t_b$. Note that the unique shortest path $a \rightsquigarrow b$ is then given by the concatenation $a \rightsquigarrow \gamma(t_a) \rightsquigarrow \gamma(t_b) \rightsquigarrow b$, where

$\gamma(t_a) \rightsquigarrow \gamma(t_b)$ is the restricted path $\gamma_{|[t_a,t_b]}$. By assumption, we have $d(a, y) \leq r$ and this implies the inequality

$$d(a, \gamma(t_a)) + d(\gamma(t_a), y) = d(a, y) \leq r = d(x, y) = d(x, \gamma(t_a)) + d(\gamma(t_a), y),$$

which is equivalent to $d(a, \gamma(t_a)) \leq d(x, \gamma(t_a))$. Similarly, the assumption $d(x, b) \leq r$ implies $d(\gamma(t_b), b) \leq d(\gamma(t_b), y)$. Thus, the distance $d(a, b)$ satisfies

$$\begin{aligned} d(a, b) &= d(a, \gamma(t_a)) + d(\gamma(t_a), \gamma(t_b)) + d(\gamma(t_b), b) \\ &\leq d(x, \gamma(t_a)) + d(\gamma(t_a), \gamma(t_b)) + d(\gamma(t_b), y) = d(x, y) = r, \end{aligned} \tag{3.12}$$

which finishes the proof of the first claim.

We now show the second claim; assume $d(a, b) = r$. From the inequalities (3.12) and $d(a, \gamma(t_a)) \leq d(x, \gamma(t_a))$, $d(\gamma(t_b), b) \leq d(\gamma(t_b), y)$ together with the assumption $d(a, b) = r$, we deduce the equalities $d(a, \gamma(t_a)) = d(x, \gamma(t_a))$ and $d(\gamma(t_b), b) = d(\gamma(t_b), y)$. Hence,

$$d(a, y) = d(a, \gamma(t_a)) + d(\gamma(t_a), y) = d(x, \gamma(t_a)) + d(\gamma(t_a), y) = d(x, y) = r$$

and similarly $d(x, b) = r$, proving the second claim. □

Lemma 3.19. *Each edge $e \in \mathrm{Cl}(X)$ has a unique maximal coface Δ_e with* $\operatorname{diam} \Delta_e = \operatorname{diam} e$. *Moreover, $\Delta_e = e$ if and only if e is a tree edge.*

Proof. By definition, e corresponds to two points $x, y \in X$ at distance $d(x, y) = r$. If e is contained in the simplex $\Delta \in \mathrm{Rips}_r(X)$, then the points in Δ lie in the intersection $D_r(x) \cap D_r(y)$, which has diameter r by Lemma 3.18. Hence, the maximal simplex Δ_e is spanned by all the points in $D_r(x) \cap D_r(y)$.

If e is a tree edge of length r, then this intersection only contains x and y, and hence $\Delta_e = e$. Conversely, if e is not a tree edge, then this intersection contains at least one vertex different from x and y that lies on the unique shortest path $x \rightsquigarrow y$. This implies $e \subsetneq \Delta_e$. □

3.4.1. Generic Tree Metrics

Before dealing with the general case, let us focus on the special case where the finite tree metric space X is *generic*, meaning that the pairwise distances are distinct. In this case, Lemma 3.19 has the following implication.

Corollary 3.20. *If X is a finite and generic tree metric space, then the diameter function*

$$\operatorname{diam}\colon \mathrm{Cl}(X) \to \mathbb{R}, \ \sigma \mapsto \operatorname{diam} \sigma$$

is a generalized discrete Morse function, defined on the full simplicial complex on X, with discrete gradient

$$\{[e, \Delta_e] \mid \textit{non-tree edge } e \subseteq \mathrm{Cl}(X)\},$$

and the critical simplices are the vertices X and the tree edges E.

We call this gradient the *generic diameter gradient*.

Example 3.21. Consider the following weighted tree with vertex set $\{a, b, c, d\}$:

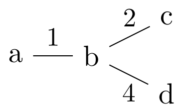

The generic diameter gradient is given by $\{[\{a, c\}, \{a, b, c\}], [\{a, d\}, \{a, b, d\}], [\{c, d\}, \{a, b, c, d\}]\}$. These intervals are the preimages under the diameter function of the non-tree distances 3, 5, and 6, respectively.

Together with Proposition 2.30, this yields the following theorem.

Theorem 3.22. *If X is a finite and generic tree metric space, then the generic diameter gradient induces, for every $r > 0$, the collapses*

$$\mathrm{Rips}_r(X) \searrow (\mathrm{Rips}_r^<(X) \cup T_r) \searrow T_r.$$

Moreover, it follows from Proposition 2.34 that for the Vietoris–Rips filtration, refined lexicographically with respect to an arbitrary total order on the vertices, the zero persistence apparent pairs refine the generic diameter gradient, and therefore also induce the above collapses.

Theorem 3.23. *If the finite tree metric space X is generic, then the apparent pairs gradient induces, for every $r > 0$, the collapses*

$$\mathrm{Rips}_r(X) \searrow (\mathrm{Rips}_r^<(X) \cup T_r) \searrow T_r.$$

Remark 3.24. There is a connection between the proof strategy of Theorem 3.16 and the proof strategy of Theorem B, in the special case of a generic tree metric X with root vertex p. As pointed out in Remark 1.7, the proof of Theorem 3.16 depends on a total order on X such that $x < y$ implies $d(x, p) \leq d(y, p)$. Moreover, the proof partitions the points into two parts, near and far, and the simplices are then partitioned by their farthest point, into simplices with only near points and simplices with some far point. More precisely, the first part deals with the simplices σ such that with $x = \max \sigma$ and $r = \operatorname{diam} \sigma$ we have $d(x, p) < r$. The discrete gradient in Eq. (3.9) consists of apparent pairs with respect to the lexicographically refined Vietoris–Rips filtration, as discussed in Remark 2.33. The second part deals with the simplices such that $d(x, p) \geq r$. Assume that $e = \{y, x\} \in \mathrm{Rips}_r(X)$ is the unique edge with $\operatorname{diam} e = r$. Recalling the proof of Lemma 3.13, we illustrate the situation as follows, where $z \in X$ is constructed from the fact that X is $(\frac{1}{2} \max_{e \in E} l(e))$-almost geodesic by Example 1.6 and the fact that the tree metric space is δ-hyperbolic for any $\delta > 0$:

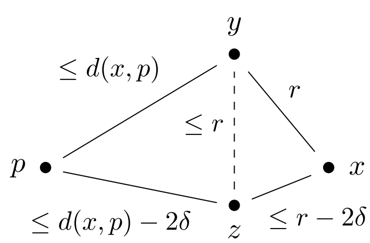

Note that $\delta > 0$ and $\operatorname{diam} e = r$ imply that z is neither equal to y nor x and thus $d(z, y) < r$. Therefore, $z \in \Delta_e \setminus e$, and the gradient in Eq. (3.10) is a vertex refinement of the interval $[e, \Delta_e]$ in the generic diameter gradient towards the vertex z.

Conversely, if $\delta > 0$ is chosen small enough (smaller than the *discretization radius* $\operatorname{discrad}(X) = \frac{1}{2} \inf\{|a-b| \mid a, b \in D_X,\ a \neq b\}$, where $D_X = \{d(x, y) \mid x, y \in X\}$ is the *distance set*), it is not too difficult to see that any $z \in \Delta_e \setminus e$ satisfies the inequalities above and can therefore be used in the proof of Theorem 3.16. In particular, the apparent pairs gradient used in Theorem 3.23, which yields a minimal vertex refinement by Proposition 2.34, is a possible choice for the gradient used in that proof.

3.4.2. Arbitrary Tree Metrics

We now turn to the general case, where Proposition 2.34 is not directly applicable anymore, as the diameter function is not necessarily a generalized discrete Morse function. Nevertheless, we show that Theorem 3.23 is still true without the genericity assumption, if the vertices X are ordered in a compatible way. To this end, we consider for every $r > 0$ the subcomplex

$$K_r := \operatorname{Rips}_r^{<}(X) \cup T_r$$

of $\operatorname{Rips}_r(X)$ and show that the complement

$$C_r := \operatorname{Rips}_r(X) \setminus K_r$$

is the set of all cofaces of non-tree edges of length r. We further show that it is partitioned into regular intervals in the face poset, and that this constitutes a discrete gradient.

Let $\Delta \in C_r \subseteq \operatorname{Rips}_r(X) =: R_r$ be a maximal simplex. We write

$$E_\Delta = \{e \in C_r \mid \dim e = 1,\ \Delta_e = \Delta\}$$

for the set of edges $e \in C_r$ with $\Delta_e = \Delta$, which is equivalently described as the set of non-tree edges of length r contained in Δ.

Lemma 3.25. *We have $\operatorname{St}_{R_r}(E_\Delta) = C_r \cap \operatorname{Cl}(\Delta)$. In particular, C_r is the set of all cofaces of non-tree edges of length r.*

Proof. The inclusion $\operatorname{St}_{R_r}(E_\Delta) \subseteq C_r \cap \operatorname{Cl}(\Delta)$ holds by definition of E_Δ. To show the inclusion $\operatorname{St}_{R_r}(E_\Delta) \supseteq C_r \cap \operatorname{Cl}(\Delta)$, let $\sigma \in C_r \cap \operatorname{Cl}(\Delta)$ be any simplex. As the Vietoris–Rips complex is a clique complex, there exists an edge $e \subseteq \sigma \subseteq \Delta$ with $\operatorname{diam} e = r$. By Lemma 3.19, this edge can not be a tree edge, as otherwise $\Delta = \Delta_e = e \notin C_r$ contradicting the assumption $\Delta \in C_r$, and hence $e \in C_r$. Therefore, by definition, $e \in E_\Delta$ and $\sigma \in \operatorname{St}_{R_r}(e) \subseteq \operatorname{St}_{R_r}(E_\Delta)$. □

Lemma 3.26. *If two distinct maximal simplices $\Delta, \Delta' \in C_r = \operatorname{Rips}_r(X) \setminus K_r$ intersect in a common face $\Delta \cap \Delta'$, then this face is contained in K_r.*

Proof. Assume for a contradiction that $\emptyset \neq \Delta \cap \Delta' \notin K_r$, implying $\Delta \cap \Delta' \in C_r$. By Lemma 3.25, there exists an edge $e \in E_\Delta \subseteq C_r$ with $e \subseteq \Delta \cap \Delta'$. By uniqueness of the maximal simplex containing e (Lemma 3.19), this implies $\Delta = \Delta'$ and contradicting the assumption that Δ, Δ' are distinct. □

We denote by L_Δ the set of all vertices of Δ that are not contained in any edge in E_Δ.

Lemma 3.27. *Let $e = \{u, w\} \in E_\Delta$ be an edge. Then any point $x \in X \setminus \{u, w\}$ on the unique shortest path $u \rightsquigarrow w$ of length r in T is contained in L_Δ. In particular, L_Δ is non-empty.*

Proof. By assumption, we have $d(u, x) < r$, $d(w, x) < r$ and $d(u, w) = r$. Therefore, $\operatorname{diam}\{u, w, x\} = r$ and $x \in \{u, w, x\} \subseteq \Delta_e = \Delta$. Assume for a contradiction that x is contained in an edge in E_Δ. Then it follows from Lemma 3.18 that we have $d(u, x) = r$ or $d(w, x) = r$, contradicting the above. We conclude that $x \in L_\Delta$. □

The Canonical Gradient We now describe a discrete gradient that is compatible with the diameter function and induces the same collapses as in Theorem 3.22 even if the tree metric is not generic. This construction is *canonical* in the sense that it does not depend on the choice of an order on the vertices, in contrast to the subsequent constructions.

Lemma 3.28. *For any two edges $f, e \in E_\Delta$ and any vertex $v \in f$ there exists a vertex $z \in e$ such that $\{v, z\} \in E_\Delta$ is an edge in E_Δ.*

Proof. Let $f = \{v, w\}, e = \{x, y\}$; note that $d(v, w) = d(x, y) = r$. Since f and e are both contained in the maximal simplex Δ, we have $v, w \in D_r(x) \cap D_r(y)$. Both $\{v, x\}$ and $\{v, y\}$ are contained in $\{v, x, y\} \subseteq \Delta$ and Lemma 3.18 implies that at least one of these two edges has length r; call this edge e_v. It follows from Lemma 3.19 that e_v is not a tree edge, and therefore $e_v \in E_\Delta$. □

Lemma 3.29. *The set $\operatorname{St}_{R_r}(E_\Delta) = C_r \cap \operatorname{Cl}(\Delta)$ is partitioned by the intervals*

$$W_\Delta = \{[\cup S, (\cup S) \cup L_\Delta] \mid \emptyset \neq S \subseteq E_\Delta\}, \tag{3.13}$$

and these form a discrete gradient on $\operatorname{Cl}(\Delta)$ inducing a collapse $\operatorname{Cl}(\Delta) \searrow (K_r \cap \operatorname{Cl}(\Delta))$.

Proof. The intervals in W_Δ are disjoint and contained in $\operatorname{St}_{R_r}(E_\Delta)$ by construction. They are regular, because L_Δ is non-empty (by Lemma 3.27). By Proposition 2.30, it remains to show that the intervals in W_Δ partition $\operatorname{St}_{R_r}(E_\Delta) = \operatorname{Cl}(\Delta) \setminus (K_r \cap \operatorname{Cl}(\Delta))$ and that W_Δ is a discrete gradient.

To show the first claim, it suffices to prove that any simplex $\sigma \in \operatorname{St}_{R_r}(E_\Delta)$ is contained in a regular interval of W_Δ. Consider the simplex $\tau = \sigma \setminus L_\Delta \subseteq \sigma$. As $\sigma \in \operatorname{St}_{R_r}(E_\Delta)$, there exists an edge $e \in E_\Delta$ with $e \subseteq \sigma$. By the definition of L_Δ, we have $e \subseteq \sigma \setminus L_\Delta = \tau$. Any other vertex $v \in \tau \setminus e$ is also contained in one of the edges E_Δ. By Lemma 3.28, there exists an edge $e_v = \{v, w\} \in E_\Delta$, where $w \in e$. Then $\tau = e \cup \bigcup_{v \in \tau \setminus e} e_v$ and $\sigma \in [\tau, \tau \cup L_\Delta] \in W_\Delta$.

The second claim now follows from the observation that the function

$$\sigma \mapsto \begin{cases} \dim(\sigma \cup L_\Delta) & \sigma \in \operatorname{St}_{R_r}(E_\Delta) \\ \dim \sigma & \sigma \notin \operatorname{St}_{R_r}(E_\Delta) \end{cases}$$

is a generalized discrete Morse function with discrete gradient W_Δ. □

Example 3.30. Consider the following tree with vertex set $X = \{a, b, c, d\}$, whose edges all have length one:

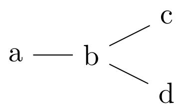

The complex $\mathrm{Rips}_2(X)$ is the full simplicial complex $\mathrm{Cl}(X)$ with maximal simplex $\Delta = X$. Clearly, the diameter function is not a generalized discrete Morse function, as Δ has three different minimal faces of same diameter, $E_\Delta = \{\{a, c\}, \{a, d\}, \{c, d\}\}$. The set L_Δ only contains the vertex $\{b\}$. Therefore, we get

$$W_\Delta = \{(\{a, c\}, \{a, b, c\}), (\{a, d\}, \{a, b, d\}), (\{c, d\}, \{b, c, d\}), (\{a, c, d\}, \{a, b, c, d\})\}.$$

We now consider the union $W_r = \bigcup_\Delta W_\Delta$, where Δ runs over all maximal simplices in C_r and W_Δ is as in (3.13). We call $W = \bigcup_r W_r$ the *canonical gradient.*

Theorem 3.31. *The canonical gradient is a discrete gradient on* $\mathrm{Cl}(X)$*. For every* $r > 0$*, it induces the collapses*

$$\mathrm{Rips}_r(X) \searrow \mathrm{Rips}_r^<(X) \cup T_r \searrow T_r.$$

Proof. Let Δ be a maximal simplex in $\Delta \in C_r = \mathrm{Rips}_r(X) \setminus K_r$, where $K_r = \mathrm{Rips}_r^<(X) \cup T_r$. It follows from Lemma 3.29 that the set W_Δ is a discrete gradient on the full subcomplex $\mathrm{Cl}(\Delta) \subseteq \mathrm{Rips}_r(X)$ that partitions $\mathrm{St}_{R_r}(E_\Delta) = \mathrm{Cl}(\Delta) \setminus (K_r \cap \mathrm{Cl}(\Delta))$ and that induces a collapse $\mathrm{Cl}(\Delta) \searrow (K_r \cap \mathrm{Cl}(\Delta))$.

It follows directly from Lemma 3.26 and Proposition 2.47 that the union $W_r = \bigcup_\Delta W_\Delta$ is a discrete gradient on $\mathrm{Rips}_r(X)$. Again by Proposition 2.47, the union $W = \bigcup_r W_r$ is a discrete gradient on $\mathrm{Cl}(X)$.

By construction of the W_Δ, the union W_r partitions the complement $\mathrm{Rips}_r(X) \setminus K_r$. Hence, by Proposition 2.30, it induces a collapse $\mathrm{Rips}_r(X) \searrow K_r = \mathrm{Rips}_r^<(X) \cup T_r$. Since only the vertices and the tree edges are critical for W, this also yields the collapse to T_r. ∎

The Perturbed Gradient Assume that X is totally ordered. We construct a coarsening of the canonical gradient to the *perturbed gradient*, such that under a specific total order of X the perturbed gradient is refined by the zero persistence apparent pairs of the diam-lexicographic order $<$ on simplices.

Every simplex $\sigma \in \mathrm{Cl}(X)$ contains a unique maximal edge $e_\sigma \in \mathrm{Cl}(X)$ with respect to $<$, which satisfies $\operatorname{diam} e_\sigma = \operatorname{diam} \sigma$. The following is an extension of Lemma 3.19.

Lemma 3.32. *Each edge* $e \in \mathrm{Cl}(X)$ *has a unique maximal coface* Σ_e *with* $\operatorname{diam} \Sigma_e = \operatorname{diam} e$ *and* $e_{\Sigma_e} = e$*, meaning that the maximal edge contained in* Σ_e *is equal to* e*.*

Proof. Let $r = \operatorname{diam} e$ and consider the union $\Sigma_e := \bigcup_{\sigma : e_\sigma = e} \sigma \subseteq \mathrm{Cl}(X)$. By Lemma 3.19, we have $\Sigma_e \subseteq \Delta_e$, implying that $\operatorname{diam} e = \operatorname{diam} \Sigma_e = \operatorname{diam} \Delta_e$.

It remains to show that the maximal edge contained in Σ_e is equal to e. To this end, let $\tilde{e} \subseteq \Sigma_e$ be any edge with $\operatorname{diam}\tilde{e} = r$. Write $e = \{x, y\}$ with $x < y$ and $\tilde{e} = \{a, b\}$ with $a < b$. By construction of Σ_e, there exist simplices $\sigma_a, \sigma_b \subseteq \mathrm{Cl}(X)$ with $a \in \sigma_a, b \in \sigma_b$ and $e_{\sigma_a} = e_{\sigma_b} = e$. Note that $\{x, y, a\} \subseteq \sigma_a$ and $\{x, y, b\} \subseteq \sigma_b$, as well as $\operatorname{diam}\sigma_a = \operatorname{diam}\sigma_b = r$.

By Lemma 3.18, we have $x, y \in S_r(a) \cup S_r(b)$ and therefore $d(a, y) = r$ (implying $a \neq y$) or $d(b, y) = r$ (implying $b \neq y$). As $\{a, y\} \subseteq \sigma_a$ and $\{b, y\} \subseteq \sigma_b$, this implies $\{a, y\} \leq e_{\sigma_a} = e = \{x, y\}$ or $\{b, y\} \leq e_{\sigma_b} = e = \{x, y\}$, respectively. In particular, we have $a \leq x$ or $a < b \leq x$, and if $a = x$, then $\tilde{e} = \{a, b\} \subseteq \sigma_b$. In any case, $\tilde{e} \leq e = e_{\sigma_b}$ as claimed. □

Consider a maximal simplex $\Delta \in C_r$, where $r > 0$. Note that all edges in E_Δ have length r and thus are ordered lexicographically. Enumerate them as $e_1 < \cdots < e_q$. We use the shorthand notation $\Sigma_i = \Sigma_{e_i}$.

Lemma 3.32 implies that $N_\Delta = \{[e_i, \Sigma_i]\}_{i=1}^q$ is a collection of disjoint intervals. It follows from Lemma 3.29 that for each $j \in \{1, \ldots, q\}$ the interval $[e_j, \Sigma_j]$ is the union

$$[e_j, \Sigma_j] = \bigcup\{[\cup S, (\cup S) \cup L_\Delta] \mid S \subseteq E_\Delta,\ e_j \text{ maximal element of } \mathrm{Cl}(\cup S) \cap E_\Delta\} \quad (3.14)$$

and that N_Δ partitions $C_r \cap \mathrm{Cl}(\Delta)$. Moreover, it is the discrete gradient of the function

$$f_\Delta\colon \mathrm{Cl}(\Delta) \to \mathbb{R},\ \sigma \mapsto \begin{cases} i & \sigma \in [e_i, \Sigma_i] \\ \dim\sigma - \dim\Delta & \sigma \in K_r \end{cases} \quad (3.15)$$

and the intervals are regular, because L_Δ is non-empty (Lemma 3.27). By Proposition 2.30, N_Δ induces a collapse $\mathrm{Cl}(\Delta) \searrow K_r \cap \mathrm{Cl}(\Delta)$. Therefore, the total order on X induces a symbolic perturbation scheme on the edges, establishing the situation of a generic tree metric as in Section 3.4.1.

Example 3.33. Recalling the tree metric from Example 3.30, we get

$$N_\Delta = \{[\{a, c\}, \{a, b, c\}], [\{a, d\}, \{a, b, d\}], [\{c, d\}, \{a, b, c, d\}]\}.$$

Note that this gradient is different from W_Δ.

Consider the union $N_r = \bigcup_\Delta N_\Delta$, where Δ runs over all maximal simplices in C_r. We call $N = \bigcup_r N_r$ the *perturbed gradient*. By (3.14), the perturbed gradient N coarsens the canonical gradient W. Analogously to Theorem 3.31, we obtain the following result.

Theorem 3.34. *The perturbed gradient is a discrete gradient on* $\mathrm{Cl}(X)$. *For every* $r > 0$, *it induces the collapses*

$$\mathrm{Rips}_r(X) \searrow \mathrm{Rips}_r^<(X) \cup T_r \searrow T_r.$$

Remark 3.35. As the lower bounds of the intervals in the perturbed gradient are edges, it follows from Theorem 3.34 that these collapses can be expressed as *edge collapses* [BP20], a notion that is similar to the elementary strong collapses described in Remark 3.15.

The Apparent Pairs Gradient Finally, we show that for a specific total order of X, which we describe next, the perturbed gradient is refined by the zero persistence apparent pairs of the diam-lexicographic order (see Section 2.3.1).

From now on, assume that the tree $T = (X, E)$ is rooted at an arbitrary vertex and orient every edge away from this point. Let $\leq_X$ be the partial order on X where u is smaller than w if there exists an oriented path $u \rightsquigarrow w$. In particular, we have the identity path $\mathrm{id}\colon u \rightsquigarrow u$. Note that for any two vertices $u, w \in X$ the unique shortest unoriented path $u \leftrightsquigarrow w$ can be written uniquely as a zig-zag $u \overset{\gamma}{\leftsquigarrow} z \overset{\eta}{\rightsquigarrow} w$, where z is the greatest point with $z \leq_X u$, $z \leq_X w$, and γ, η are oriented paths in T that intersect only in z. If $w \leftrightsquigarrow p$ is another unique shortest unoriented path with the zig-zag $w \overset{\varphi}{\leftsquigarrow} z' \overset{\mu}{\rightsquigarrow} p$, then we can form the following diagram

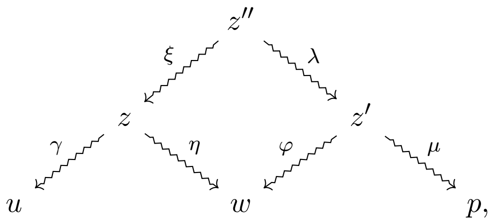

(3.16)

where z'' is the greatest point with $z'' \leq_X z, z'' \leq_X z'$. Moreover, as T has no cycles, it follows that either ξ or λ is the identity path and $\varphi \circ \lambda = \eta$ or $\eta \circ \xi = \varphi$, respectively.

Remark 3.36. Note that in general the oriented paths $z'' \rightsquigarrow u$ and $z'' \rightsquigarrow p$ can intersect in a point different from z''. In particular, the zig-zag $u \leftsquigarrow z'' \rightsquigarrow p$ is not necessarily a decomposition of the unique shortest unoriented path $u \leftrightsquigarrow p$.

Extend the partial order $\leq_X$ on X to a total order $<$ and consider the diam-lexicographic order on simplices. As this total order on the simplices extends $<$ under the identification $v \mapsto \{v\}$, we will also denote it by $<$.

Consider a maximal simplex $\Delta \in C_r$. Recall that N_Δ is the discrete gradient of the function $f_\Delta\colon \mathrm{Cl}(\Delta) \to \mathbb{R}$ defined in (3.15), using the same vertex order as above. By Proposition 2.34, the zero persistence apparent pairs with respect to the f_Δ-lexicographic order $<_{f_\Delta}$ are precisely the gradient pairs of the minimal vertex refinement of N_Δ.

Lemma 3.37. *Every apparent pair with respect to $<_{f_\Delta}$ of the form*

$$(\sigma, \tau = \sigma \cup \{v\}) \subseteq [e_i, \Sigma_i] \in N_\Delta,$$

where v is the minimal vertex in $\Sigma_i \setminus e_i$, is also an apparent pair with respect to $<$.

Proof. First, let $\sigma \cup \{p\} \in C_r$ be a cofacet of σ not equal to τ. We show that we must have $\tau < \sigma \cup \{p\}$, proving that τ is the minimal cofacet of σ with respect to $<$: If $p \in \Sigma_i$, then $p \in \Sigma_i \setminus e_i$, as $p \notin \sigma \supseteq e_i$, and the statement is true by minimality of v in the minimal vertex refinement. Now assume that $p \notin \Sigma_i$ and write $e_i = \{u, w\}$ with $u < w$. By (3.14), we have $L_\Delta \subseteq \Sigma_i$ and hence it follows that $p \notin L_\Delta$ and that the point p is contained in an edge in E_Δ, by definition of L_Δ. It follows from Lemma 3.18 that p together with at least one vertex of e_i forms an edge in E_Δ. Call this edge g; if there are two such edges, consider the larger one, and call it g. From $\{u, w, p\} \subseteq \Delta$ and $p \notin \Sigma_i$ we get $e_i < g$: The edge e_i is not the maximal edge of the two simplex $\{u, w, p\}$, since otherwise p would be contained in Σ_i. Hence, one of the two other edges is maximal, and that edge is g by definition. Considering the two possible cases $g = \{u, p\}$ and $g = \{w, p\}$, we must have $u < p$. We will argue that $v < p$ holds, which proves $\tau = \sigma \cup \{v\} < \sigma \cup \{p\}$.

Consider the diagram (3.16). If $\gamma \neq \mathrm{id}$, then it follows from the fact that $e_i = \{u, w\}$ is not a tree edge that along the unique shortest path $u \leftrightsquigarrow w$ there exists a vertex x distinct from u and w with $x < u < p$. Then $x \in L_\Delta \subseteq \Sigma_i \setminus e_i$ by Lemma 3.27, and as v is the minimal element in $\Sigma_i \setminus e_i$, we get $v \leq x < p$.

If $\gamma = \mathrm{id}$, then $u = z$, and it follows from $d(w, p) \leq r$ and $p \notin e_i = \{u, w\}$ that we must have $\lambda \neq \mathrm{id}$ and $\xi = \mathrm{id}$: Otherwise $\lambda = \mathrm{id}$ and $u = z$ lies on φ. Therefore, u lies on the unique shortest path from w to p and $d(w, p) = d(w, u) + d(u, p) = r + d(u, p) > r$, yielding a contradiction. Thus, the unique shortest path $(u = z) \leftrightsquigarrow p$ decomposes as $u \rightsquigarrow z' \rightsquigarrow p$, where $u \rightsquigarrow z'$ is contained in $u \rightsquigarrow z' \rightsquigarrow w$. Note that $u \neq z'$, because $\lambda \neq \mathrm{id}$. Hence, as e_i is not a tree edge, the immediate successor x of u on the path $u \rightsquigarrow w$ is distinct from u and w with $x \leq z'$. This point satisfies $x \leq z' \leq p$, and it follows from Lemma 3.27 that we have $x \in L_\Delta \subseteq \Sigma_i \setminus e_i$. Because $p \notin L_\Delta$ we even have $x < p$. Therefore, as v is the minimal vertex in $\Sigma_i \setminus e_i$, it follows that $v \leq x < p$.

It remains to prove that σ is the maximal facet of τ with respect to $<$. We write $e_i = \{u, w\}$ with $u < w$ and $\tau = \{b_0, \ldots, b_{\dim \tau}\}$ with $b_0 < \cdots < b_{\dim \tau}$. As $e_i \subseteq \tau$, there are indices $k_1 < k_2$ with $u = b_{k_1} < b_{k_2} = w$. If $k_1 > 0$, then $v = b_0$, so σ is of the form $\{b_1, \ldots, b_{\dim \tau}\}$ and is the maximal facet of τ with respect to $<$ as claimed. Now assume $k_1 = 0$. If τ contains no edges $e \in E_\Delta$ other than e_i, then the facets $\tau \setminus \{u\}$ and $\tau \setminus \{w\}$ are both contained in $\mathrm{Rips}_r^<(X)$, because they do not contain any edge of length r, and the maximal facet of τ is $\tau \setminus \{x\}$ with x the minimal vertex in $\tau \setminus e_i$. By assumption, we have $x = v$ and hence $\tau \setminus \{x\} = \tau \setminus \{v\} = \sigma$. If τ contains other edges $e \neq e_i$ with $e \in E_\Delta$, label them $s_1, \ldots, s_a$. As $e_i \subseteq \tau \subseteq \Sigma_i$, it follows from Lemma 3.32 that we have $s_b < e_i$ for all b. Because of this and our assumption $k_1 = 0$, i.e., u is the minimal vertex of τ, we have $s_b = \{u, x_b\} < \{u, w\} = e_i$ with $u < x_b < w$. Therefore, the facet $\{b_1, \ldots, b_{\dim \tau}\}$ contains no edges in E_Δ and hence it is contained in $\mathrm{Rips}_r^<(X)$. The facet $\{b_0, b_2, \ldots, b_{\dim \tau}\}$ of τ contains e_i, hence it is an element of C_r, and so it is maximal among the facets containing b_0, implying that it is the maximal facet of τ with respect to $<$. Because b_1 is the minimal vertex in $\tau \setminus e_i$ and $v \in \tau \setminus e_i$, it follows from the minimality of $v \in \Sigma_i \setminus e_i$ that we have $b_1 = v$, implying $\{b_0, b_2, \ldots, b_{\dim \tau}\} = \sigma$. Therefore, σ is the maximal facet of τ with respect to $<$. ☐

The following proposition directly implies Theorem B.

Proposition 3.38. *The intervals in the perturbed gradient N are refined by apparent pairs with respect to $<$. For every $r > 0$, the apparent pairs gradient for the lexicographically refined Vietoris–Rips filtration on the full simplicial complex* $\mathrm{Cl}(X)$ *induces a collapse*

$$\mathrm{Rips}_r(X) \searrow \mathrm{Rips}_r^<(X) \cup T_r.$$

Proof. The first statement follows readily from the facts that by Proposition 2.34 every interval $[e_i, \Sigma_i] \in N_\Delta$ of the perturbed gradient is refined by the zero persistence apparent pairs with respect to the f_Δ-lexicographic order $<_{f_\Delta}$, and that by Lemma 3.37 those are also apparent pairs with respect to $<$. As the zero persistence apparent pairs of $<_{f_\Delta}$, taken over all maximal simplices $\Delta \in C_r$, yield a partition of $C_r = \mathrm{Rips}_r(X) \setminus (\mathrm{Rips}_r^<(X) \cup T_r)$, the same is then true for the apparent pairs of $<$. Thus, by Proposition 2.30, the apparent pairs gradient of $<$ induces a collapse $\mathrm{Rips}_r(X) \searrow \mathrm{Rips}_r^<(X) \cup T_r$. $\square$

Remark 3.39. The preceding Proposition 3.38 also implies Theorem A in the special case of tree metrics: if $t > u \geq 2\,\mathrm{geod}_p(X) = 2\,\mathrm{geod}(X) = \max_{e \in E} l(e)$ are real numbers, then $T_u = T$ is the entire tree, and we obtain collapses $\mathrm{Rips}_t(X) \searrow \mathrm{Rips}_u(X) \searrow T \searrow \{*\}$. If all edges of T have the same length, it turns out that the collapse $T \searrow \{*\}$ is also induced by the apparent pairs gradient of the same order $<$.

Remark 3.40. For metrics other than tree metrics, the collapse $\mathrm{Rips}_t(X) \searrow \mathrm{Rips}_u(X)$ from Theorem A is not always achieved by the apparent pairs gradient. Consider the following weighted graph:

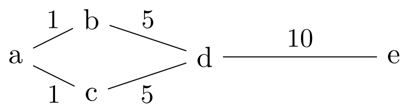

For this graph metric, the hyperbolicity is $\mathrm{hyp}(X) = 1$, and the geodesic defect is $\mathrm{geod}(X) = 5$; therefore, we have $4\,\mathrm{hyp}(X) + 2\,\mathrm{geod}(X) = 14$. The maximal Vietoris–Rips complex has 31 simplices in total. For the apparent pairs gradient only the simplices $\{b, e\}$ and $\{b, d, e\}$ are critical, and both have diameter 15. Thus, the collapse $\mathrm{Rips}_{15}(X) \searrow \mathrm{Rips}_{14}(X)$ is not induced by the apparent pairs gradient.

Example 3.41. Revisiting the tree metric from Example 3.30 once more, we see that the apparent pairs of the lexicographically refined Vietoris–Rips filtration with diameter two are

$$(\{a, c\}, \{a, b, c\}), (\{a, d\}, \{a, b, d\}), (\{c, d\}, \{a, c, d\}), (\{b, c, d\}, \{a, b, c, d\}).$$

Note that together with Example 3.33 this shows that the canonical gradient, the perturbed gradient, and the apparent pairs gradient can all be different in general.

4. Lexicographic Optimal Homologous Cycles in Delaunay Filtrations

In this section, we prove the main results summarized in Section 1.3, namely Theorem C and Corollary D. More concretely, in Section 4.1, we explain how persistence pairs form an algebraic gradient and how the associated gradient flow (see Section 2.3.4) relates to lexicographically minimal cycles (Definition 2.26). Together with results from Sections 2.3.2 and 2.3.4, which have some novel aspects, this is then used in Section 4.2 to relate lexicographically minimal cycles to the descending complexes of a generalized discrete Morse function.

4.1. Algebraic Morse Theory and Persistence

We saw that the apparent pairs are closely related to persistent homology (Lemma 2.24) and discrete Morse theory (Proposition 2.34). In this section, we show how all the persistence pairs are related to algebraic Morse theory (see Section 2.3.4). We also show how this approach connects to lexicographically minimal cycles, and to matrix reductions used in the context of persistent homology (see Section 2.2).

4.1.1. Gradient Pairs from Persistence Pairs

We now explain how all persistence pairs determine an algebraic gradient that relates to discrete Morse theory through apparent pairs (Lemma 2.24 and Proposition 2.34). This establishes the framework for a key step in our proof of Theorem C.

Let $(C_*, \Sigma_* = \sigma_1 < \cdots < \sigma_l)$ be a based chain complex with an elementwise filtration, and let $R = D \cdot S$ be a reduction of the filtration boundary matrix. For any chain $c \in C_n$, we denote by $\operatorname{Pivot}_{\Sigma_*} c = \max \operatorname{supp}_{\Sigma_*} c$ the maximal basis element in the basis representation of c with respect to Σ_*. If v is the coordinate vector in $\mathbb{F}^l$ representing c, we also write $\operatorname{Pivot}_{\Sigma_*} v$ for $\operatorname{Pivot}_{\Sigma_*} c = \sigma_{\operatorname{PivInd} v}$.

The direct sum decomposition of filtered chain complexes explained in Section 2.2 yields a straightforward interpretation of persistence pairs as an algebraic gradient, which, however, is not suitable for our purposes as it neither relates directly to apparent pairs nor lexicographically minimal cycles. For completeness of the exposition, we briefly explain this approach, nevertheless. To this end, we equip the chain complex C_* with the new ordered basis $E_* = \eta_1 < \cdots < \eta_l$ given by

$$\eta_i = \begin{cases} R_j & \text{if there exists an index persistence pair } (i,j), \\ S_i & \text{if } i \text{ is a death or essential index.} \end{cases}$$

We call E_* the *decomposition basis*, noting that it induces a decomposition of the filtered chain complex as described in Section 2.2. Note also that with respect to the original basis we have $\operatorname{Pivot}_{\Sigma_*} \eta_i = \sigma_i$ for all i. By pairing the death columns S_j with their boundaries $R_j = D \cdot S_j$, we obtain a set of disjoint pairs that we call the *decomposition gradient* of S:

$$\{(R_j, S_j) \mid j \text{ is a death index}\}.$$

Proposition 4.1. *The decomposition gradient is an algebraic gradient on the basis E_*.*

Proof. Consider the function $f\colon E_* \to \mathbb{N}$ with values $f(R_j) = f(S_j) = j$ for every death index j and $f(S_i) = i$ for every essential index i. Note first that for any death index j, the basis element R_j is the only facet of S_j, and R_j has no facets as it is a cycle. Similarly, for any essential index i, the basis element S_i has no facets, as well. Thus, it is immediate that f is a monotonic function whose algebraic gradient is the decomposition gradient. □

We now explain an alternative approach to how persistence pairs determine an algebraic gradient, which is used in the subsequent sections, and that utilizes the distinctness of non-zero pivot elements in the reduced matrix $R = D \cdot S$. To this end, we equip the chain complex C_* with the new ordered basis $\Omega_* = \tau_1 < \cdots < \tau_l$ given by

$$\tau_i = \begin{cases} \sigma_i & \text{if } i \text{ is a birth or essential index,} \\ S_i & \text{if } i \text{ is a death index.} \end{cases}$$

We call this basis Ω_* the *reduction basis*. Note that with respect to the original basis we have $\operatorname{Pivot}_{\Sigma_*} \tau_i = \sigma_i$ for all i. Moreover, note that for every death index j and $R_j = D \cdot S_j$ we have $\operatorname{Pivot}_{\Omega_*} R_j = \operatorname{Pivot}_{\Sigma_*} R_j$. By pairing the death columns S_j with the pivot elements $\operatorname{Pivot}_{\Omega_*} R_j$ of their boundaries $R_j = D \cdot S_j$, we obtain a set of disjoint pairs, which we call the *reduction gradient* of S:

$$\{(\operatorname{Pivot}_{\Omega_*} R_j, S_j) \mid j \text{ is a death index}\}.$$

Proposition 4.2. *The reduction gradient is an algebraic gradient on Ω_*.*

Proof. Consider the function $f\colon \Omega_* \to \mathbb{N}$ with values $f(\operatorname{Pivot}_{\Omega_*} R_j) = f(S_j) = \operatorname{PivInd} R_j$ for every death index j and $f(\tau_i) = i$ for every essential index i. Note that, as S is a full rank upper triangular matrix, the ordered basis $\Omega_* = \tau_1 < \cdots < \tau_l$ is compatible with the given elementwise filtration, in the sense that $\downarrow \tau_j$ induces the same subcomplex of C_* as $\downarrow \sigma_j$ for every j. In particular, whenever τ_i is a facet of τ_j, we have $i < j$. The function f assigns to any basis element of the form $\tau_j = S_j$, for j a death index, the index $i = \operatorname{PivInd} R_j$ of its maximal facet $\operatorname{Pivot}_{\Omega_*} R_j = \operatorname{Pivot}_{\Sigma_*} R_j = \tau_i \in \Omega_*$, and to all other basis elements τ_i their index i. This implies that f is monotonic with algebraic gradient the reduction gradient. □

The reduction gradient is closely related to apparent pairs (Definition 2.23) in the following sense. For any apparent pair (σ_i, σ_j) of the elementwise filtration, the column of σ_j in the filtration boundary matrix D is already reduced (Lemma 2.24). Therefore, we are led to make the following definition.

Definition 4.3. Let $(C_*, \Sigma_* = \sigma_1 < \cdots < \sigma_l)$ be a based chain complex with an elementwise filtration, and $R = D \cdot S$ a reduction of the filtration boundary matrix. We call the reduction matrix S *apparent pairs compatible* if the column S_j contains only one non-zero entry $S_{j,j} = 1$ for every apparent pair (σ_i, σ_j) of the elementwise filtration.

Note that, in analogy to Lemma 2.32, the collection of apparent pairs of the elementwise filtration forms an algebraic gradient on Σ_* [Lam20, Lemma 2.2; Bau21, Lemma 3.5] that we also call the *apparent pairs gradient*. The following is a direct consequence of the definitions.

Lemma 4.4. *Let $(C_*, \Sigma_* = \sigma_1 < \cdots < \sigma_l)$ be a based chain complex with an elementwise filtration, and $R = D \cdot S$ a reduction of the filtration boundary matrix. If the reduction matrix S is apparent pairs compatible, then the apparent pairs gradient of the elementwise filtration on Σ_* is a subset of the reduction gradient of S on the reduction basis Ω_*.*

Remark 4.5. Both reduction algorithms in Section 2.2 compute a reduction $R = D \cdot S$ of the filtration boundary matrix, noting explicitly that S is homogeneous. Furthermore, Algorithm 1 computes a reduction matrix S that is also apparent pairs compatible.

4.1.2. Gradient Flows on Elementwise Filtered Chain Complexes

We now study the flow determined by an algebraic gradient (Definition 2.51) in the context of based chain complexes with an elementwise filtration. Our main example is the flow determined by the reduction gradient associated to a reduction of the filtration boundary matrix, as explained in Section 4.1.1.

We first establish, in Proposition 4.8, a criterion for a cycle to be invariant under the flow determined by an algebraic gradient, which comes in handy when relating the gradient flow to lexicographically minimal cycles in Section 4.1.3. The following statement is a direct consequence of the definitions.

Lemma 4.6. *Let (C_*, Σ_*) be a based chain complex and let $\sigma_1 < \cdots < \sigma_l$ be any total order on the basis Σ_* corresponding to an elementwise filtration that refines the sublevel set filtration induced by an algebraic Morse function f. Then, with respect to this total order, for any gradient pair $(a, b) \in V$ of f the pivot of ∂b is equal to a, $\operatorname{Pivot}_{\Sigma_*} \partial b = a$.*

For a based chain complex (C_*, Σ_*) and algebraic gradient V on Σ_*, we denote by

$$\partial_n V = \{\partial b \mid \exists\, (a, b) \in V \text{ with } a \in \Sigma_n\}$$

the set of gradient cofacet boundaries in degree n. We say that two cycles $z, z' \in Z_n$ are *V-homologous* if there exists an element $\partial e \in \operatorname{span} \partial_n V$ such that $z - z' = \partial e$. Observe that for any cycle $z \in Z_n$, the cycle $\Phi(z) = z + \partial\,\mathrm{F}(z)$ is V-homologous to z.

Let L_n be the linear subspace of C_n spanned by the gradient facets of V and consider the canonical linear projection $\pi_n\colon \operatorname{span} \partial_n V \to L_n$ with respect to the basis Σ_n.

Lemma 4.7. *The subset $\partial_n V \subseteq C_n$ is linearly independent, and the canonical linear projection $\pi_n\colon \operatorname{span} \partial_n V \to L_n$ is an isomorphism.*

Proof. By Lemma 4.6, we can take a total order $\sigma_1 < \cdots < \sigma_l$ on the basis Σ_* such that for any pair $(a, b) \in V$ the pivot of ∂b is equal to a, $\operatorname{Pivot}_{\Sigma_*} \partial b = a$. As the pairs in V are disjoint, we know that the pivots $\operatorname{Pivot}_{\Sigma_*} \partial b = a$ for $(a, b) \in V$ are distinct. In particular, this implies that the subset $\partial_n V$ is linearly independent.

To prove the second claim, note that, similarly to before, the pivots $\operatorname{Pivot}_{\Sigma_*} \pi_n(\partial b) = a$ for $(a, b) \in V$ with $a \in \Sigma_n$ are distinct. This implies that the linearly independent vectors $\partial_n V$ are sent to linearly independent vectors in L_n. As $\operatorname{span} \partial_n V$ and L_n have the same dimension $\operatorname{card}(\{(a, b) \in V \mid a \in \Sigma_n\})$, this shows that π_n is an isomorphism. □

Proposition 4.8. *Let (C_*, Σ_*) be a based chain complex, and let V be an algebraic gradient on Σ_* with associated flow $\Phi\colon C_* \to C_*$. Then a cycle $z \in Z_n$ is Φ-invariant if and only if it contains no gradient facets of V. Moreover, if $z' \in Z_n$ is any Φ-invariant cycle that is V-homologous to z, then $z = z'$.*

Proof. We start by proving the first claim. On the one hand, we know from Proposition 2.54 and Lemma 2.53, that a Φ-invariant chain contains no gradient facets of V. On the other hand, if a cycle contains no gradient facets of V, then it is $\Phi^{(a,b)}$-invariant for every pair $(a, b) \in V$, and it follows from Corollary 2.56 that it is also Φ-invariant.

To prove the second claim, let $z' \in Z_n$ be any Φ-invariant cycle that is V-homologous to z. Then by definition, there exists a $\partial e \in \operatorname{span} \partial_n V$ with $z - z' = \partial e$. Recall from Lemma 4.7, that the projection $\pi_n\colon \operatorname{span} \partial_n V \to L_n$ onto the subspace of C_n spanned by the gradient facets of V is an isomorphism. By assumption, the difference $z - z' = \partial e$ is Φ-invariant, and therefore it contains no gradient facets of V, by our first claim. Thus, we have $\pi_n(\partial e) = 0$, implying $z - z' = \partial e = 0$, and proving the second claim $z = z'$. □

In order to relate the flow determined by the reduction gradient to the flow determined by the zero persistence apparent pairs gradient, and therefore to discrete Morse theory (Propositions 2.34 and 4.24), we make the following definition.

Definition 4.9. Let $(C_*, \Sigma_* = \sigma_1 < \cdots < \sigma_l)$ be a based chain complex with an elementwise filtration, and $R = D \cdot S$ a reduction of the filtration boundary matrix. We call the reduction matrix S *death-compatible* if for every death index j and non-zero entry $S_{i,j} \neq 0$ we also have that i is a death index.

Remark 4.10. Note that both of the algorithms in Section 2.2 compute a death-compatible reduction matrix S.

The following is a direct consequence of Lemma 4.4 and the definitions.

Lemma 4.11. *Let $(C_*, \Sigma_* = \sigma_1 < \cdots < \sigma_l)$ be a based chain complex with an elementwise filtration, and $R = D \cdot S$ a reduction of the filtration boundary matrix. If the reduction matrix S is apparent pairs and death-compatible, then the flows determined by the apparent pairs gradient of the elementwise filtration as an algebraic gradient on Σ_* and as an algebraic gradient on the reduction basis Ω_*, respectively, coincide.*

We now make two definitions that come in handy when relating the gradient flow to lexicographically minimal cycles in Section 4.1.3. Motivated by Lemma 4.6, we make the following definition.

Definition 4.12. Let $(C_*, \Sigma_* = \sigma_1 < \cdots < \sigma_l)$ be a based chain complex with an elementwise filtration. We call an algebraic gradient V on Σ_* *reduced in degree n*, or *n-reduced*, if for all pairs $(a, b) \in V$ with $a \in \Sigma_n$ we have $\operatorname{Pivot}_{\Sigma_*} \partial b = a$. Moreover, we call an algebraic gradient V on Σ_* *reduced* if it is n-reduced for all n.

Note that if an algebraic gradient V is reduced, then the matrix with columns ∂b for $(a, b) \in V$, written in the basis Σ_*, is reduced. Motivated by the notion of two cycles being V-homologous, and by Proposition 4.8, we make the following definition.

Definition 4.13. Let (C_*, Σ_*) be a based chain complex. We say that an algebraic gradient V on Σ_* *generates the n-boundaries* if the set of gradient cofacet boundaries $\partial_n V$ generates the entire subspace of boundaries $B_n \subseteq C_n$.

Remark 4.14. By construction, any reduction gradient is reduced in degree n and generates the n-boundaries for all n.

4.1.3. Relating the Gradient Flow and Lexicographically Minimal Cycles

We now relate the flow invariant cycles determined by a reduction gradient to lexicographically minimal homologous cycles. Let $(C_*, \Sigma_* = \sigma_1 < \cdots < \sigma_l)$ be a based chain complex with an elementwise filtration, and let $R = D \cdot S$ be a reduction of the filtration boundary matrix. The following result [CLV22, Lemmas 3.3 and 3.5] provides an equivalent condition for minimality.

Proposition 4.15. *A cycle is lexicographically minimal with respect to the elementwise filtration if and only if its support contains only death elements and essential elements.*

We now provide another characterization in terms of the gradient flow determined by the reduction gradient of S (defined on the reduction basis Ω_*), which we call the *reduction flow*.

Proposition 4.16. *Let $(C_*, \Sigma_* = \sigma_1 < \cdots < \sigma_l)$ be a based chain complex with an elementwise filtration, and let $R = D \cdot S$ be a reduction of the filtration boundary matrix by a death-compatible reduction matrix S. Then for a cycle $z \in Z_n$ the following are equivalent:*

- *z is lexicographically minimal with respect to the ordered basis Σ_*;*
- *z is invariant under the reduction flow.*

Remark 4.17. It is worth noting that the equivalence in Proposition 4.16 does not extend to chains. Consider the example in Fig. 13, with the simplex τ removed and such that σ has the function value 3. Then, the chain σ is not invariant under the reduction flow determined by the reduction gradient associated to the induced simplexwise filtration, but it is lexicographically minimal, since there are no 2-simplices.

Before proving Proposition 4.16, we state the following result that shows that the minimizers are unique:

Proposition 4.18. *Let $(C_*, \Sigma_* = \sigma_1 < \cdots < \sigma_l)$ be a based chain complex with an elementwise filtration, and let z be lexicographically minimal cycle with respect to the basis Σ_*. If $c \in Z_n$ is homologous to z and $c \sqsubseteq z$ with respect to Σ_*, then we necessarily have that $c = z$.*

We now prove Propositions 4.16 and 4.18 through a sequence of general statements. As we deal with two different bases in Proposition 4.16, we first introduce the *n-reduction gradient*, which provides an intermediate between lexicographically minimal cycles with respect to Σ_* and the reduction gradient defined on the reduction basis Ω_*.

Let $(C_*, \Sigma_* = \sigma_1 < \cdots < \sigma_l)$ be a based chain complex with an elementwise filtration, and let $R = D \cdot S$ be a reduction of the filtration boundary matrix. We equip the chain complex C_* with the new ordered basis $\Pi_* = \tau_1 < \cdots < \tau_l$ given by

$$\tau_i = \begin{cases} S_i & \text{if } i \text{ is a death index and } \sigma_i \in C_{n+1}, \\ \sigma_i & \text{otherwise.} \end{cases}$$

We call Π_* the *n-reduction basis.* Note that Σ_* and Π_* induce the same lexicographic preorder on C_n. In particular, note that for every death index j with $\sigma_j \in C_{n+1}$ we have $\operatorname{Pivot}_{\Pi_*} R_j = \operatorname{Pivot}_{\Sigma_*} R_j$, where $R_j = D \cdot S_j$. Similarly to Proposition 4.2, the collection of disjoint pairs

$$\{(\operatorname{Pivot}_{\Pi_*} R_j, S_j) \mid \ j \text{ is a death index and } \sigma_j \in C_{n+1}\}$$

is an algebraic gradient on Π_*, the *n-reduction gradient* of S. By construction, the n-reduction gradient is n-reduced and generates the n-boundaries.

We now prove some general statements about algebraic gradients that are n-reduced and generate the n-boundaries. Let $(C_*, \Pi_* = \tau_1 < \cdots < \tau_l)$ be any based chain complex with an elementwise filtration.

Lemma 4.19. *If V is an n-reduced algebraic gradient on Π_* that generates the n-boundaries, then a cycle $z \in Z_n$ is reducible if and only if its support in the basis Π_* contains a gradient facet of V.*

Proof. Assume that the cycle z is reducible. By definition, there exists a cycle $z' \sqsubset z$ and a chain $e \in C_{n+1}$ such that $z - z' = \partial e$. As V generates the n-boundaries, the boundary $\partial e \in B_n = \operatorname{span} \partial_n V$ can be written as a linear combination $\partial e = \sum \lambda_i \partial b_i$ for some $(a_i, b_i) \in V$ with $\lambda_i \neq 0$ and $a_i \in \Pi_n$. Since the chains ∂b_i have distinct pivots, as V is n-reduced and the pairs in V are disjoint, there exists a unique chain ∂b_j with $\operatorname{Pivot}_{\Pi_*} \partial e = \operatorname{Pivot}_{\Pi_*} \partial b_j = a_j$. As $\lambda_j \neq 0$, the gradient facet a_j must be contained in z, as otherwise $z \sqsubseteq z'$, contradicting the assumption.

For the converse, note that if the cycle z contains some gradient facet a with $(a, b) \in V$, then $z' = \Phi^{(a,b)}(z) = z + \langle z, a \rangle \cdot \partial \operatorname{F}(a)$, where $\operatorname{F}(a) = -\langle \partial b, a \rangle^{-1} \cdot b$, is a homologous cycle with $z' \sqsubset z$, since $\operatorname{Pivot}_{\Pi_*} \partial b = a$ and a is contained in $\operatorname{supp}_{\Pi_*} z$ but not in $\operatorname{supp}_{\Pi_*} z'$. ☐

Proposition 4.20. *If V is an n-reduced algebraic gradient on a basis Π_* that generates the n-boundaries, then for a cycle $z \in Z_n$ the following are equivalent:*

- *z is lexicographically minimal with respect to the basis Π_*;*
- *z is invariant under the gradient flow Φ determined by V, i.e., it satisfies $\Phi(z) = z$.*

Proof. Let $z \in Z_n$ be a Φ-invariant cycle. By Proposition 4.8, this is equivalent to z not containing a gradient facet, and that, in turn is equivalent by Lemma 4.19 to z being irreducible, which is the same as being lexicographically minimal. □

In particular, note that Proposition 4.20 implies that every n-reduced algebraic gradient on Π_* that generates the n-boundaries has the same set of invariant cycles.

Proof of Proposition 4.16. By construction of the gradient flow and under the assumption that S is death-compatible, the reduction flow agrees on cycles with the flow determined by the n-reduction gradient on the n-reduction basis Π_*. Thus, together with Proposition 4.20, this implies that the cycle z is invariant under the reduction flow if and only if z is lexicographically minimal with respect to the ordered basis Π_*. As mentioned before, the ordered bases Π_* and Σ_* induce the same lexicographic preorder on C_n, and hence the claim follows. □

Proof of Proposition 4.18. Let $R = D \cdot S$ be a reduction of the filtration boundary matrix, and let V be the n-reduction gradient on the n-reduction basis Π_*. Recall that V is n-reduced and generates the n-boundaries by construction. As mentioned before, the ordered bases Π_* and Σ_* induce the same lexicographic preorder on C_n, and thus, both cycles z and c are also lexicographically minimal with respect to the basis Π_*. By Proposition 4.20, this implies that both cycles are invariant under the flow determined by V. As V generates the n-boundaries, for z and c being homologous is the same as being V-homologous, and hence Proposition 4.8 implies $c = z$. □

Finally, we give a proof of Proposition 4.15 using these results.

Proof of Proposition 4.15. By Remark 4.10, we know that there exists a death-compatible reduction matrix S. Thus, by Proposition 4.16, a cycle z is lexicographically minimal with respect to the elementwise filtration if and only if z is invariant under the reduction flow. As in the proof of Proposition 4.16, this is equivalent to z being invariant under the flow of the n-reduction gradient. This in turn is, by Proposition 4.8, equivalent to z not containing any gradient facets of the n-reduction gradient. Finally, by the construction of the n-reduction gradient, this is equivalent to z not containing any (non-essential) birth elements, meaning that it only contains death elements and essential elements. □

4.1.4. Gradient flows as Matrix Reductions

We demonstrate how the gradient flow on a cycle can be interpreted as a variant of Gaussian elimination, tying it closely to the exhaustive reduction from Section 2.2.

Let $(C_*, \Sigma_* = \sigma_1 < \cdots < \sigma_l)$ be a based chain complex with an elementwise filtration, let V be an algebraic gradient on Σ_*, and denote the associated flow by $\Phi \colon C_* \to C_*$.

Proposition 4.21. *If V is n-reduced and $c \in Z_n$ is a cycle, then Algorithm 3 computes the image of the cycle under the flow Φ, i.e., it computes $\Phi(c)$.*

Input: $D = \partial$ the $l \times l$ filtration boundary matrix, c a cycle
for $i = 1, \dots, l$ **do**
  if $c_i \neq 0$ *and* $(\sigma_i, \sigma_j) \in V$ *a gradient pair* **then**
    $\mu = -c_i / D_{i,j}$;
    $c = c + \mu \cdot D_j$;
return c

Algorithm 3: Gradient flow reduction

Proof. By definition, for any cycle c the gradient flow is given by $\Phi(c) = c + \partial\,\mathrm{F}(c)$. We denote by L the set of pairs (i, j) with $c_i \neq 0$ and $(\sigma_i, \sigma_j) \in V$. Recall from Definition 2.51 that $\mathrm{F}\colon C_* \to C_{*+1}$ is the linear map with $\mathrm{F}(\sigma_i) = -\langle \partial\sigma_j, \sigma_i \rangle^{-1} \cdot \sigma_j$ if $(\sigma_i, \sigma_j) \in V$ and 0 on all other basis elements. Note that $\partial\sigma_j$ is represented by the column D_j and $\langle \partial\sigma_j, \sigma_i \rangle = D_{i,j}$. Thus, we have

$$\partial\,\mathrm{F}(c) = \sum_{(i,j)\in L} c_i \cdot \partial\,\mathrm{F}(\sigma_i) = \sum_{(i,j)\in L} (-c_i \langle \partial\sigma_j, \sigma_i \rangle^{-1}) \cdot \partial\sigma_j = \sum_{(i,j)\in L} (-c_i / D_{i,j}) \cdot D_j.$$

By assumption of V being n-reduced, these columns D_j have distinct pivots and thus, traversing and updating the cycle c from small to large index has the same effect as adding a summand from the sum above to the cycle c one after the other. Hence, Algorithm 3 computes $\Phi(c)$. □

To compute the image of a cycle under the stabilized flow, one can use Algorithm 4, which resembles the exhaustive reduction from Section 2.2, and that does not require any additional specific choices besides the algebraic gradient, in contrast to Algorithm 3.

Input: $D = \partial$ the filtration boundary matrix, c a cycle
while *there exists* i *with* $c_i \neq 0$ *and* $(\sigma_i, \sigma_j) \in V$ *a gradient pair* **do**
  $\mu = -c_i / D_{i,j}$;
  $c = c + \mu \cdot D_j$;
return c

Algorithm 4: Stabilized flow reduction

Proposition 4.22. *For any cycle $c \in Z_n$, Algorithm 4 computes the image of the cycle under the stabilized flow Φ^∞, i.e., it computes $\Phi^\infty(c)$.*

Proof. Note that we can choose a different elementwise filtration by the σ_i without affecting Algorithm 4. Therefore, by Lemma 4.6, we can assume that V is n-reduced. Hence, Algorithm 4 is essentially the same as Algorithm 2, implying that it terminates. Now note that Algorithm 4 computes a V-homologous cycle c' of c that contains no gradient facets, and thus it is Φ-invariant by Proposition 4.8. By definition, the stabilized cycle $\Phi^\infty(c)$ is also Φ-invariant and V-homologous to c. Therefore, both c' and $\Phi^\infty(c)$ are Φ-invariant cycles that are V-homologous. Proposition 4.8 now implies $c' = \Phi^\infty(c)$. □

Figure 12: Discrete gradient (blue) with corresponding descending complex (green). Left: Cycle c (red). Right: Stabilized cycle $\Phi(c) = \Phi^\infty(c)$ (red), supported on the descending complex (green).

4.2. Relating Algebraic Reduction Gradients and Discrete Gradients

We are now ready to relate our results for the gradient flow determined by the reduction gradient to discrete Morse theory. We show that for a generalized discrete Morse function $f\colon K \to \mathbb{R}$, defined on a finite simplicial complex, the lexicographically minimal cycles with respect to the f-lexicographic order on the simplices are supported on the descending complexes of f. In particular, this is true for the columns of the totally reduced filtration boundary matrix, considered as cycles in the sublevel set corresponding to the pivot index. We further show that the reduction chains, given as the columns of a suitable reduction matrix, are also supported on the descending complexes of f. This summarizes the relations between reductions of the filtration boundary matrix and the descending complexes of f. Recall from Section 2.3.4 that a discrete gradient V on a finite simplicial complex K gives rise to a flow $\Phi\colon C_*(K) \to C_*(K)$.

Lemma 4.23. *The flow Φ restricts to a chain map on the descending complex $D(V) \subseteq K$, i.e., for $c \in C_*(D(V))$ we also have $\Phi(c) \in C_*(D(V))$.*

Proof. Recall from Definition 2.51 that Φ is given by $\Phi(c) = c + \partial\,\mathrm{F}(c) + \mathrm{F}(\partial c)$, where $\mathrm{F}\colon C_*(K) \to C_{*+1}(K)$ is the linear map with $\mathrm{F}(\sigma) = -\langle \partial\tau, \sigma\rangle^{-1} \cdot \tau$ if $(\sigma, \tau) \in V$ and 0 on all other simplices. Thus, if $\eta \in D(V)$ is any simplex, then $\partial\eta$ and $\mathrm{F}(\eta)$ are both contained in $C_*(D(V))$, by definition of the descending complex. Therefore, if $c \in C_*(D(V))$ is any chain, then the chains $\partial c, \mathrm{F}(c)$, and $\mathrm{F}(\partial c)$ are also contained in $C_*(D(V))$. This shows that the chain $\Phi(c) = c+\partial\,\mathrm{F}(c)+\mathrm{F}(\partial c)$ is contained in $C_*(D(V))$, proving the claim. □

The Φ-invariant chains are supported on the descending complex, as illustrated in Fig. 12.

Proposition 4.24. *The Φ-invariant chains of $C_*(K)$ are supported on the descending complex $D(V)$, i.e., we have $C^\Phi_*(K) \subseteq C_*(D(V))$.*

Proof. Let $c \in C^\Phi_*(K)$ be any Φ-invariant chain; we show that c is contained in $C_*(D(V))$. By Proposition 2.54, we can assume without loss of generality that c is of the form $\Phi^\infty(\sigma)$ for a critical simplex σ. By definition, $\Phi^\infty = \Phi^r$ for a large enough r, and by Lemma 2.53, $\Phi^r(\sigma)$ is given by $\Phi^r(\sigma) = \sigma + w + \Phi(w) + \cdots + \Phi^{r-1}(w)$, where $w = \mathrm{F}(\partial\sigma)$. It follows directly from the definition of the descending complex (Definition 2.35) that $\sigma, \partial\sigma$, and

$w = \mathrm{F}(\partial\sigma)$ are contained in $C_*(D(V))$. By Lemma 4.23, we have $\Phi^k(w) \in C_*(D(V))$ for every k, and therefore

$$c = \Phi^r(\sigma) = \sigma + w + \Phi(w) + \cdots + \Phi^{r-1}(w) \in C_*(D(V)),$$

proving the claim. □

The following, together with Proposition 2.29, directly implies Theorem C.

Theorem 4.25. *Let f be a generalized discrete Morse function defined on a finite simplicial complex, let $r \in \mathbb{R}$ be a real number, and let $h \in H_*(S_r(f))$ be a homology class of the sublevel set $S_r(f)$. Then the lexicographically minimal cycle of h, with respect to the f-lexicographic order, is supported on the descending complex $D_r(f)$.*

Proof. Let V be the discrete gradient of f, and let W be the zero persistence apparent pairs gradient induced by the f-lexicographic order, which is a regular refinement of V by Proposition 2.34. Note that f is compatible with W.

Recall from Remarks 4.5 and 4.10 that Algorithm 1 computes a reduction $R = D \cdot S$ of the filtration boundary matrix, corresponding to the simplexwise filtration of $S_r(f)$ induced by the f-lexicographic order, such that the reduction matrix S is apparent pairs compatible and also death-compatible. Consider the corresponding reduction gradient on the corresponding reduction basis Ω_* with associated reduction flow $\Psi\colon C_*(S_r(f)) \to C_*(S_r(f))$.

The lexicographically minimal cycle $z \in Z_*(S_r(f))$ of h is, by Proposition 4.16, a Ψ-invariant cycle. We show that z is contained in $C_*(D_r(f))$, which proves the claim: As S is apparent pairs compatible, it follows from Lemma 4.4 and Proposition 2.55 that z is also invariant under the gradient flow determined by the zero persistence apparent pairs gradient W_r on Ω_*. Lemma 4.11 implies that this flow coincides with the gradient flow $\Phi\colon C_*(S_r(f)) \to C_*(S_r(f))$ determined by the zero persistence apparent pairs gradient W_r on the standard basis given by the simplices of $S_r(f)$. It now follows from Proposition 4.24, the definition of descending complex (Definition 2.35), and Corollary 2.39 that z is contained in

$$C_*^{\Phi}(S_r(f)) \subseteq C_*(D(W_r)) = C_*(D_r(W,f)) \subseteq C_*(D_r(V,f)) = C_*(D_r(f)).$$ □

Let D be the filtration boundary matrix of the simplexwise filtration of K induced by the f-lexicographic order, and let $R = D \cdot S$ be a totally reduced reduction of D. The following, together with Proposition 2.29, directly implies Corollary D.

Corollary 4.26. *Let f be a generalized discrete Morse function defined on a finite simplicial complex, and let (σ,τ) be a non-zero persistence pair of the simplexwise filtration induced by the f-lexicographic order. Let $r = f(\sigma)$ and $s = f(\tau)$ be the function values of σ and of τ, respectively. Then the lexicographically minimal cycle of $[\partial\tau]$ in the open sublevel set $S_s^<(f)$, given as the column R_τ of the totally reduced filtration boundary matrix, is supported on the descending complex $D_r(f)$.*

Proof. As (σ, τ) is a non-zero persistence pair, we know that $f(\sigma) < f(\tau) = s$ and that τ is a critical simplex. As R is a reduction of D, this implies that R_τ and $\partial\tau$ are homologous cycles in $S_s^<(f)$. Since R is totally reduced, the cycle R_τ does not contain a (non-essential) birth simplex of the (smaller) simplexwise filtration of $S_s^<(f)$ induced by the f-lexicographic order. Thus, Proposition 4.15 implies that R_τ is the lexicographically minimal cycle of $[\partial\tau]$ in $S_s^<(f)$. As $r = f(\sigma)$ and R is a reduction of D, the cycle R_τ is supported on the subcomplex $S_r(f)$ of $S_s^<(f)$, implying that it is also a lexicographically minimal cycle in $S_r(f)$. It now follows from Theorem 4.25 that the cycle R_τ is supported on the descending complex $D_r(f)$. □

Reduction Chains and Descending Complexes For a generalized discrete Morse function $f\colon K \to \mathbb{R}$, defined on a finite simplicial complex, we know from Corollary 4.26 that the columns of the totally reduced filtration boundary matrix corresponding to non-zero persistence pairs are supported on the descending complex. We now show that the reduction chains, i.e., the chains given as the columns of the reduction matrix, that correspond to essential and non-zero persistence pairs are also supported on the descending complexes.

Let D be the filtration boundary matrix of the simplexwise filtration of K induced by the f-lexicographic order, and let S be a death-compatible reduction matrix with $R = D \cdot S$ totally reduced.

Proposition 4.27. *Let $f\colon K \to \mathbb{R}$ be a generalized discrete Morse function defined on a finite simplicial complex, and let $\tau \in K$ be a critical simplex, i.e., a simplex that is either essential or contained in a non-zero persistence pair of the simplexwise filtration induced by the f-lexicographic order. Let $r = f(\tau)$ be the function value of τ. Then the chain given as the column S_τ of the reduction matrix is supported on the descending complex $D_r(f)$.*

Proof. Note that if τ is an essential simplex or a birth simplex contained in a non-zero persistence pair, then it is an essential simplex of the (smaller) simplexwise filtration of $S_r(f)$ induced by the f-lexicographic order. Let W_r be the zero persistence apparent pairs gradient of this simplexwise filtration and denote by $\Phi\colon C_*(S_r(f)) \to C_*(S_r(f))$, with $\Phi(c) = c + \mathrm{F}(\partial c) + \partial\,\mathrm{F}(c)$, the associated flow.

Assume first that τ is an essential simplex of the simplexwise filtration of $S_r(f)$. Then, the chain c given as the column S_τ of the reduction matrix is a cycle, i.e., $\partial c = 0$. Moreover, since S is death-compatible, we also have $\mathrm{F}(c) = 0$. Hence, we get

$$\Phi(c) = c + \mathrm{F}(\partial c) + \partial\,\mathrm{F}(c) = c + 0 + 0 = c$$

and the chain c is Φ-invariant. Analogously as in the proof of Theorem 4.25, we conclude that c is supported on the descending complex $D_r(f)$.

Now assume that τ is a death simplex of the simplexwise filtration of $S_r(f)$ that is contained in a non-zero persistence pair (σ, τ), implying that τ is not contained in a zero-persistence apparent pair. By Definition 2.23, and as the matrix R is totally reduced, the column $R_\tau = D \cdot S_\tau$ does not contain any apparent facet of the simplexwise filtration. Therefore, the chain c given as the column S_τ of the reduction matrix satisfies $\mathrm{F}(\partial c) = 0$. As before, since S is death-compatible, we also have $\mathrm{F}(c) = 0$. Hence, we get

$$\Phi(c) = c + \mathrm{F}(\partial c) + \partial\,\mathrm{F}(c) = c + 0 + 0 = c$$

and the chain c is Φ-invariant. Analogously to before, we conclude that c is supported on the descending complex $D_r(f)$, proving the claim. □

Remark 4.28. Note that the assumption in Corollary 4.26 and Proposition 4.27 on the pair (σ, τ) to be a non-zero persistence pair can not be dropped. Consider, for example, the simplicial complex in Fig. 13 and the discrete Morse function $f\colon K \to \mathbb{R}$ that assigns to the simplices σ and τ the value 3 and to all other simplices the value as indicated. Then the pair (σ, τ) is a zero persistence apparent pair, but neither $S_\tau = \tau$ nor $S_\sigma = R_\tau = \partial\tau$ is supported on the descending complex $D_3(f)$, which only consists of the orange vertex with function value 0.

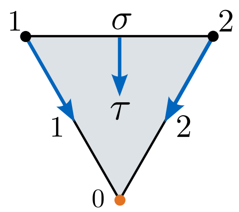

Figure 13: Discrete gradient (blue) together with the unique critical simplex (orange).

5. Functorial Nerve Theorems

In this section, we prove the functorial nerve theorems summarized in Section 1.4, namely Theorems E to H. More concretely, in Section 5.1, we argue that the blowup complex and the bar construction can be used interchangeably to prove functorial nerve theorems. In Section 5.2, we prove functorial nerve theorems for closed convex covers, and in Section 5.3, we prove (functorial) nerve theorems for simplicial and semi-algebraic covers. In Section 5.4, we prove the unified nerve theorem and discuss some counterexamples that illustrate the necessity of some of its assumptions. Finally, in Section 5.5, we give more direct proofs of results from work of Dugger and Isaksen [DI04] that we use.

5.1. Functorial Nerve Theorems via the Bar Construction

Theorems E, G, and H use the *blowup complex* of a cover $\mathcal{U}$ as an intermediate object to relate the nerve of $\mathcal{U}$ with the covered space. In this section, we define the blowup complex and its natural maps to the covered space and the nerve. The construction is not difficult, but there is an important point here: the blowup complex is closely related to the bar construction (Definition 2.57), and because of this, properties of the bar construction are used in many proofs of the nerve theorem, including Theorems E and H.

We first explain how the bar construction can be used to prove functorial nerve theorems. For any poset P, if we write $*^P$ for the diagram $P \to \mathsf{Top}$ with constant value the one-point space $*$, there is a canonical identification $\mathrm{Bar}(*^P) \cong |\,\mathrm{Flag}(P)|$. In particular, if $P_\mathcal{U}$ is the poset associated to a cover $\mathcal{U}$, we have $\mathrm{Bar}(*^{P_\mathcal{U}}) \cong |\,\mathrm{Flag}(P_\mathcal{U})| = |\,\mathrm{Sd}\,\mathrm{Nrv}(\mathcal{U})|$.

A morphism of covered spaces $(f, \varphi)\colon (X, \mathcal{U}) \to (Y, \mathcal{V})$ induces a poset map $g\colon P_\mathcal{U} \to P_\mathcal{V}$ between the posets associated to the covers (Definition 2.59), and a natural transformation $\lambda\colon \mathcal{D}_\mathcal{U} \Rightarrow \mathcal{D}_\mathcal{V} \circ g$. Thus, by definition of the category $\mathrm{Diag}_{\mathsf{Po}}(\mathsf{Top})$ in Section 2.4.1 and the discussion there, the operation $(X, \mathcal{U}) \mapsto \mathrm{Bar}(\mathcal{D}_\mathcal{U})$ defines a functor $\mathsf{Cov} \to \mathrm{Diag}_{\mathsf{Po}}(\mathsf{Top}) \to \mathsf{Top}$. Moreover, the unique natural transformation $\mathcal{D}_\mathcal{U} \Rightarrow *^{P_\mathcal{U}}$ induces a natural map

$$\pi_{\mathrm{Sd}\,N}\colon \mathrm{Bar}(\mathcal{D}_\mathcal{U}) \to \mathrm{Bar}(*^{P_\mathcal{U}}) \cong |\,\mathrm{Sd}\,\mathrm{Nrv}(\mathcal{U})|\,.$$

If every non-empty finite intersection of cover elements happens to be contractible, then this map is a homotopy equivalence by Proposition 2.62. Using also the natural map from the bar construction to the colimit as mentioned at the end of Section 2.4.1, we get a diagram that is natural in morphisms of covered spaces:

$$X \leftarrow \mathrm{colim}\,\mathcal{D}_\mathcal{U} \leftarrow \mathrm{Bar}(\mathcal{D}_\mathcal{U}) \to |\,\mathrm{Sd}\,\mathrm{Nrv}(\mathcal{U})|\,.$$

In Section 5.2 and Section 5.4, we will use this diagram to prove functorial nerve theorems by finding various sets of assumptions that make these natural maps equivalences of various kinds. This strategy – also employed in the well-known proof of the nerve theorem for open covers in Hatcher's textbook [Hat02, Section 4.G] – relies on the well-known good properties of the bar construction. We exploit this established theory repeatedly in Section 5.4, where we use the fact that the bar construction is homotopical in several contexts, including homological algebra. In Section 5.3, we will prove a functorial nerve theorem using a bar construction constructed in the category of posets, rather than the topological construction.

However, for purposes of computational topology, we want a nerve theorem to relate the space X directly with the nerve of $\mathcal{U}$, not the much larger subdivision of the nerve. In order to obtain a functorial nerve theorem that works for morphisms of covered spaces as we have defined them, in which the map of indexed covers need not be an inclusion, we cannot simply apply the usual homeomorphism $\alpha_K\colon |\operatorname{Sd}(K)| \to |K|$ defined for any simplicial complex K (see Section 2), as this map is natural only in inclusions of simplicial complexes. In the case of diagrams $\mathcal{D}_{\mathcal{U}}$ associated to a cover, the *blowup complex* is an efficient way to build a space homeomorphic to the bar construction, which comes with a natural map to $|\operatorname{Nrv}(\mathcal{U})|$ rather than $|\operatorname{Sd}\operatorname{Nrv}(\mathcal{U})|$. In the definition, for a non-empty finite set J, we write $|J|$ for the geometric realization of the full simplicial complex generated by J, which is homeomorphic to the standard topological simplex $\Delta^{|J|-1}$.

Definition 5.1. Let $\mathcal{U} = (U_i)_{i\in I}$ be a cover of a topological space X. The *blowup complex* $\operatorname{Blowup}(\mathcal{U})$ is the quotient space

$$\operatorname{Blowup}(\mathcal{U}) = \left(\coprod_{J\in\operatorname{Nrv}(\mathcal{U})} U_J \times |J| \right) \; / \; \sim,$$

where the disjoint union is taken over all simplices $J \in \operatorname{Nrv}(\mathcal{U})$, and the equivalence relation $\sim$ identifies, for all $J \subseteq J'$, the spaces $U_J \times |J|$ and $U_{J'} \times |J'|$ along their common subspace $U_{J'} \times |J|$.

Remark 5.2. For a finite cover $\mathcal{U} = (U_i)_{i\in[n]}$ the blowup complex can be defined as a subspace of the product $X \times \Delta^n$, as mentioned in Section 1.4. This is the approach used by Zomorodian and Carlsson [ZC08, Definition 3].

As before, the operation $(X, \mathcal{U}) \mapsto \operatorname{Blowup}(\mathcal{U})$ defines a functor $\mathsf{Cov} \to \mathsf{Top}$, and the projection maps $U_J \times |J| \to U_J$ define a natural map $\operatorname{Blowup}(\mathcal{U}) \to \operatorname{colim} \mathcal{D}_{\mathcal{U}}$, which gives us a natural map $\rho_S\colon \operatorname{Blowup}(\mathcal{U}) \to X$. But now the projection maps $U_J \times |J| \to |J|$ assemble to define a natural map $\rho_N\colon \operatorname{Blowup}(\mathcal{U}) \to |\operatorname{Nrv}(\mathcal{U})|$.

Example 5.3. As in Example 2.61, we consider the cover of the circle by three arcs.

Figure 14: The bar construction (left) and the blowup complex (right).

Note that the blowup complex is combinatorially simpler.

We can use the homeomorphism $\alpha_K : |\operatorname{Sd}(K)| \to |K|$ defined for any simplicial complex K to construct a homeomorphism $\operatorname{Bar}(\mathcal{D}_{\mathcal{U}}) \to \operatorname{Blowup}(\mathcal{U})$. For any simplex $J \in \operatorname{Nrv}(\mathcal{U})$, any flag $J \supset J_1 \supset \cdots \supset J_n$ in $P_{\mathcal{U}}$ indexes a piece $U_J \times \Delta^n$ in $\operatorname{Bar}(\mathcal{D}_{\mathcal{U}})$. The flags of this form glue together to give a copy of $U_J \times |\operatorname{Sd} J|$ inside $\operatorname{Bar}(\mathcal{D}_{\mathcal{U}})$, where $\operatorname{Sd} J$ is the subdivision of the simplicial complex generated by J. Now for all $J \in \operatorname{Nrv}(\mathcal{U})$ we have a map

$$U_J \times |\operatorname{Sd} J| \xrightarrow{\alpha_J \times \mathrm{id}} U_J \times |J| \subset \operatorname{Blowup}(\mathcal{U}) ,$$

and assembling these maps gives the homeomorphism $\operatorname{Bar}(\mathcal{D}_{\mathcal{U}}) \to \operatorname{Blowup}(\mathcal{U})$.

This homeomorphism is not natural in arbitrary morphisms of covered spaces, but it does fit into the following commutative diagram, where the solid arrows are natural in morphisms of covered spaces:

$$\begin{array}{ccccc} X & \xleftarrow{\pi_S} & \operatorname{Bar}(\mathcal{D}_{\mathcal{U}}) & \xrightarrow{\pi_{\operatorname{Sd} N}} & |\operatorname{Sd}\operatorname{Nrv}(\mathcal{U})| \\ \Vert & & \downarrow \cong & & \downarrow \cong \\ X & \xleftarrow[\rho_S]{} & \operatorname{Blowup}(\mathcal{U}) & \xrightarrow[\rho_N]{} & |\operatorname{Nrv}(\mathcal{U})| \end{array} \tag{5.1}$$

The somewhat subtle point here is that, even though the homeomorphism $\operatorname{Bar}(\mathcal{D}_{\mathcal{U}}) \to \operatorname{Blowup}(\mathcal{U})$ is not natural in arbitrary morphisms of covered spaces, we can use this homeomorphism and the good properties of the bar construction to prove functorial nerve theorems for the blowup complex: if some set of assumptions on the covered space $(X, \mathcal{U})$ imply that the top maps from $\operatorname{Bar}(\mathcal{D}_{\mathcal{U}})$ are equivalences of some kind, then the commutativity of the diagram 5.1 implies that the bottom maps from $\operatorname{Blowup}(\mathcal{U})$ are equivalences of the same kind.

5.2. Functorial Nerve Theorems for Closed Convex Covers

Now, we will discuss two ways of turning the result in Theorem 2.5 into a functorial nerve theorem. We start by giving a version that follows the strategy explained in Section 5.1. After this, we will give a version that is functorial on the nose but needs to use the concept of pointed covers.

The essential ideas of the following proof, which we include for completeness of the exposition, already appeared in the author's master's thesis [Rol20].

Theorem 5.4. *If $X \subset \mathbb{R}^d$, and $\mathcal{C} = (C_i)_{i \in [n]}$ is a cover of X by closed convex subsets, then the natural maps $\rho_S : \operatorname{Blowup}(\mathcal{C}) \to X$ and $\rho_N : \operatorname{Blowup}(\mathcal{C}) \to |\operatorname{Nrv}(\mathcal{C})|$, introduced in Section 5.1, are homotopy equivalences.*

Proof. As explained at the end of Section 5.1, it suffices to consider the (not necessarily commutative) diagram

$$\begin{array}{ccccc} & & \operatorname{Bar}(\mathcal{D}_{\mathcal{C}}) & & \\ & \swarrow^{\pi_S} & & \searrow^{\pi_{\operatorname{Sd} N}} & \\ X & & \xleftarrow[\Gamma]{} & & |\operatorname{Sd}\operatorname{Nrv}(\mathcal{C})|, \end{array}$$

where Γ is as in Theorem 2.5, and show that π_S and $\pi_{\operatorname{Sd} N}$ are homotopy equivalences.

By Proposition 2.62 and the fact that convex sets are contractible, we know that $\pi_{\operatorname{Sd} N}$ is a homotopy equivalence. Every point $p \in \operatorname{Bar}(\mathcal{D}_{\mathcal{C}})$ can be described as a pair $p = (x, \alpha)$, where α is a point in $|\sigma|$, for some $\sigma = (J_n \subset \cdots \subset J_0) \in \operatorname{Sd}\operatorname{Nrv}(\mathcal{C})$, and $x \in C_{J_0}$. By construction, we have $\Gamma(\pi_{\operatorname{Sd} N}(p)) = \Gamma(\alpha) \in C_{J_n}$ and $\pi_S(p) = x \in C_{J_0} \subseteq C_{J_n}$. Therefore, a straight line homotopy shows that the maps $\Gamma \circ \pi_{\operatorname{Sd} N} \simeq \pi_S$ are homotopic. As Γ and $\pi_{\operatorname{Sd} N}$ are homotopy equivalences the same is true for π_S. □

We will now describe the second way of obtaining a functorial nerve theorem. Recall from Definition 1.14 and the paragraphs afterwards the definition of Cov_* and its subcategory ClConv_*. Before stating the functorial nerve theorem let us consider one more important example of a pointed covered space.

Example 5.5. Let K be a simplicial complex. The cover $(\operatorname{bst}(v))_{v \in \operatorname{Vert} K}$ of $|K|$ by the closed barycentric stars is pointed by the vertices $(|w|)_{w \in \operatorname{Vert}\operatorname{Sd} K}$.

Let $(X, \mathcal{A}_*) \in \mathsf{ClConv}_*$ be a pointed covered space. Recall that the construction of Γ in Section 2.1.1 requires many choices. Those choices can be made such that Γ is a morphism of pointed covered spaces, where the nerve is a pointed covered space as described in Example 5.5, and such that it is affine linear on each simplex of the barycentric subdivision of the nerve.

Theorem 5.6. *The homotopy equivalence* $\Gamma\colon |\operatorname{Sd}\operatorname{Nrv}(\mathcal{A})| \to X$ *is natural with respect to the morphisms in* ClConv_**.*

Proof. To show naturality, let $(f, \varphi)\colon (X, \mathcal{A}_*) \to (Y, \mathcal{C}_*)$ be a morphism in ClConv_*. Then we need to prove that the diagram

$$\begin{array}{ccc} X & \xrightarrow{\quad f \quad} & Y \\ \uparrow & & \uparrow \\ |\operatorname{Sd}\operatorname{Nrv}(\mathcal{A})| & \xrightarrow{|\operatorname{Sd}\varphi_*|} & |\operatorname{Sd}\operatorname{Nrv}(\mathcal{C})| \end{array}$$

commutes. Both compositions are maps $|\operatorname{Sd}\operatorname{Nrv}(\mathcal{A})| \to Y$ that are affine linear on each simplex of the barycentric subdivision. Hence, they are completely determined by their values on the vertices, where both compositions coincide by construction. □

5.3. Nerve Theorems for Simplicial and Semi-Algebraic Covers

One can prove a nerve theorem for simplicial complexes as a corollary of Quillen's Theorem A for posets. In this section, we use combinatorial arguments to prove a functorial version of this result. Using a well-known triangulation theorem for semi-algebraic sets, this functorial nerve theorem for simplicial complexes implies such a theorem for finite, closed, semi-algebraic covers of compact semi-algebraic sets. Finally, we use the same combinatorial methods to prove a functorial version of a nerve theorem of Björner.

Posets and Homotopy Theory We begin by reviewing some basic facts about posets, following Quillen [Qui78].

Recall that the flag complex $\mathrm{Flag}(P)$ of a poset P is the simplicial complex whose vertices are the elements of P and whose n-simplices are the chains $x_0 < \cdots < x_n$ of elements of P. We will sometimes say that a poset has a certain topological property if its flag complex has that property. For example, we say that a poset P is contractible if $|\mathrm{Flag}(P)|$ is contractible, and we say a map $f\colon P \to Q$ of posets is a homotopy equivalence if the induced map $|\mathrm{Flag}(P)| \to |\mathrm{Flag}(Q)|$ is a homotopy equivalence. If P and Q are posets, then there is a canonical homeomorphism

$$|\mathrm{Flag}(P \times Q)| \xrightarrow{\cong} |\mathrm{Flag}(P)| \times |\mathrm{Flag}(Q)| \tag{5.2}$$

induced by the projection maps. As explained in [Qui78], the product must be taken in the category of compactly generated spaces, Definition 2.63. However, if one of P or Q is finite, then this agrees with the usual product. It follows that if $f, g\colon P \to Q$ are maps of posets such that $f(x) \leq g(x)$ for all $x \in P$, then $|\mathrm{Flag}(f)|, |\mathrm{Flag}(g)|\colon |\mathrm{Flag}(P)| \to |\mathrm{Flag}(Q)|$ are homotopic. To see this, observe that the relation $f \leq g$ determines a map of posets $H : P \times \{0 < 1\} \to Q$, and $|\mathrm{Flag}(\{0 < 1\})|$ is an interval.

The main result about posets that we need is Quillen's Theorem A [Qui73]. Given a map $f\colon P \to Q$ of posets and $y \in Q$, define the subposet of P: $f/y = \{x \in P \mid f(x) \leq y\}$.

Theorem 5.7 (Quillen's Theorem A)**.** *If $f\colon P \to Q$ is a map of posets, and f/y is contractible for all $y \in Q$, then f is a homotopy equivalence.*

It should be said that Quillen's theorem is more general than this result, but this is what we will use. For a nice proof at this level of generality, see [Wal81] or [Bar11], where it is shown that for finite posets the map f is even a simple homotopy equivalence. We now use Quillen's Theorem A to give a simple proof of the nerve theorem for covers of a simplicial complex by subcomplexes; this is similar to [Bjö81, Lemma 1.1], [Bjö03, Theorem 6], and [Bar11, Theorem 4.3].

Proposition 5.8. *Let K be a simplicial complex and let $\mathcal{A} = (K_i \subseteq K)_{i \in I}$ be a locally finite good cover of K by subcomplexes. Then $|K|$ is homotopy equivalent to $|\mathrm{Nrv}(\mathcal{A})|$.*

Proof. Define a map of posets $f\colon \mathrm{Pos}(K) \to \mathrm{Pos}(\mathrm{Nrv}(\mathcal{A}))^{\mathrm{op}}$ by

$$f(\sigma) = \{i \in I \mid \sigma \in K_i\} \ .$$

As $\mathcal{A}$ is locally finite, $f(\sigma)$ is finite for all $\sigma \in K$. We will show that f is a homotopy equivalence. As usual, for $J \subseteq I$, we write $K_J = \cap_{i \in J} K_i$. By Quillen's Theorem A, it suffices to show that, for all elements J of $\mathrm{Pos}(\mathrm{Nrv}(\mathcal{A}))^{\mathrm{op}}$, the poset f/J is contractible. Unwinding the definition, f/J is the subposet of $\mathrm{Pos}(K)$ with elements $\sigma \in \mathrm{Pos}(K)$ such that $J \subseteq f(\sigma)$. By definition, $J \subseteq f(\sigma)$ if and only if $\sigma \in K_J$. So, $f/J = \mathrm{Pos}(K_J)$. As the intersection K_J is non-empty, it is contractible by assumption, and so $|\mathrm{Flag}(\mathrm{Pos}(K_J))| \cong |K_J|$ is contractible. Thus, f is a homotopy equivalence. The homotopy equivalence of the proposition is the composition:

$$|K| \cong |\mathrm{Sd}(K)| \xrightarrow{f_*} |\mathrm{Sd}(\mathrm{Nrv}(\mathcal{A}))| \cong |\mathrm{Nrv}(\mathcal{A})| \ . \qquad \square$$

5.3.1. A Functorial Nerve Theorem for Simplicial Covers

In order to prove a functorial version of Proposition 5.8, we now introduce a poset PoBar that is intermediate between the posets $\mathrm{Pos}(K)$ and $\mathrm{Pos}(\mathrm{Nrv}(\mathcal{A}))^{\mathrm{op}}$ that appeared in the proof. We use the notation PoBar because this construction can be seen as a bar construction taken in the category of posets, as we explain in Example 5.41. An additional benefit of using this intermediate object is that it allows one to remove the assumption that the cover is locally finite. This is similar to the strategy of Björner [Bjö81, Lemma 1.1], which he attributes to Quillen.

Definition 5.9. If K is a simplicial complex and $\mathcal{A} = (K_i \subseteq K)_{i\in I}$ is a cover of K by subcomplexes, let $\mathrm{PoBar}(\mathcal{A})$ be the poset with the underlying set

$$\mathrm{PoBar}(\mathcal{A}) = \{(\sigma, J) \mid J \subseteq I \text{ finite}, \sigma \in K_J\}$$

where $(\sigma, J) \le (\sigma', J')$ if and only if $\sigma \subseteq \sigma'$ and $J \supseteq J'$.

Since $\mathrm{PoBar}(\mathcal{A})$ is a subposet of the product $\mathrm{Pos}(K) \times \mathrm{Pos}(\mathrm{Nrv}(\mathcal{A}))^{\mathrm{op}}$, it comes with projection maps $\lambda_S\colon \mathrm{PoBar}(\mathcal{A}) \to \mathrm{Pos}(K)$ and $\lambda_N\colon \mathrm{PoBar}(\mathcal{A}) \to \mathrm{Pos}(\mathrm{Nrv}(\mathcal{A}))^{\mathrm{op}}$. In the next lemma, f denotes the poset map defined in the proof of Proposition 5.8.

Proposition 5.10. *Let K be a simplicial complex and let $\mathcal{A} = (K_i \subseteq K)_{i\in I}$ be a cover of K by subcomplexes. Then, the map λ_S is a homotopy equivalence and if $\mathcal{A}$ is good, then λ_N is a homotopy equivalence, as well. Moreover, if $\mathcal{A}$ is locally finite, the map f is defined and the following diagram of posets commutes up to homotopy after taking flag complexes:*

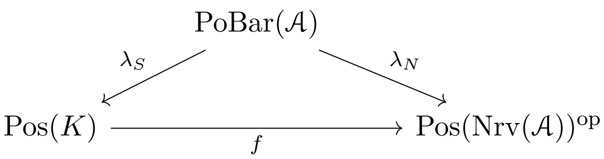

Proof. We begin by showing that λ_S is a homotopy equivalence. Since a map of posets $P \to Q$ is a homotopy equivalence if and only if the induced map on opposite posets $P^{\mathrm{op}} \to Q^{\mathrm{op}}$ is a homotopy equivalence, it suffices, by Quillen's Theorem A, to show that for any $\sigma \in \mathrm{Pos}(K)$ the subposet $\sigma\backslash\lambda_S = \{(\tau, J) \mid \sigma \subseteq \lambda_S(\tau, J) = \tau\} \subseteq \mathrm{PoBar}(\mathcal{A})$ is contractible: Consider the fiber $\lambda_S^{-1}(\sigma) = \{(\sigma, J) \mid \sigma \in K_J\} \subseteq \mathrm{PoBar}(\mathcal{A})$, and define the poset map $\mu\colon \sigma\backslash\lambda_S \to \lambda_S^{-1}(\sigma)$ by $\mu(\tau, J) = (\sigma, J)$. The map μ is a homotopy inverse to the inclusion of $\lambda_S^{-1}(\sigma)$ into $\sigma\backslash\lambda_S$, as, for any $(\tau, J) \in \sigma\backslash\lambda_S$ we have the relation $\mu(\tau, J) \le (\tau, J)$ in $\sigma\backslash\lambda_S$. The fiber $\lambda_S^{-1}(\sigma)$ is contractible, as it is isomorphic to the opposite face poset of a full simplicial complex. We conclude that $\sigma\backslash\lambda_S$ is contractible and thus λ_S is a homotopy equivalence.

Now, assume that $\mathcal{A}$ is good. We show that λ_N is a homotopy equivalence, using Quillen's Theorem A. So, we take $J \in \mathrm{Pos}(\mathrm{Nrv}(\mathcal{A}))^{\mathrm{op}}$, and we must check that λ_N/J is contractible. Consider the fiber $\lambda_N^{-1}(J) = \{(\sigma, J) \mid \sigma \in K_J\} \subseteq \mathrm{PoBar}(\mathcal{A})$, and define the poset map $\nu\colon \lambda_N/J \to \lambda_N^{-1}(J)$ by $\nu(\sigma, \tilde{J}) = (\sigma, J)$. The map ν is a homotopy inverse to the inclusion of $\lambda_N^{-1}(J)$ into λ_N/J, as, for any $(\sigma, \tilde{J}) \in \lambda_N/J$ we have the relation $\nu(\sigma, \tilde{J}) \geq (\sigma, \tilde{J})$ in λ_N/J. The fiber $\lambda_N^{-1}(J)$ is contractible, as it is isomorphic to $\operatorname{Pos}(K_J)$, which is contractible as $\mathcal{A}$ is good. We conclude that λ_N/J is contractible and thus Quillen's Theorem A implies that λ_N is a homotopy equivalence.

Finally, assume $\mathcal{A}$ is locally finite, so that f is defined. If (σ, J) is in $\operatorname{PoBar}(\mathcal{A})$, then $\lambda_N(\sigma, J) = J \subseteq f(\sigma) = (f \circ \lambda_S)(\sigma, J)$. So, we have $(f \circ \lambda_S)(\sigma, J) \leq \lambda_N(\sigma, J)$, which implies that $|\operatorname{Flag}(f)| \circ |\operatorname{Flag}(\lambda_S)|$ and $|\operatorname{Flag}(\lambda_N)|$ are homotopic. ☐

The strategy now is to use what we have proved about the PoBar construction to show that the natural maps from the blowup complex to $|K|$ and $|\operatorname{Nrv}(\mathcal{A})|$ (defined in Section 5.1) are homotopy equivalences.

To do this, we will identify a subcomplex $\mathrm{T}(\mathcal{A}) \subseteq \operatorname{Flag}(\operatorname{PoBar}(\mathcal{A}))$ that is homeomorphic to the blowup complex after realization; we then show, using discrete Morse theory, that the inclusion $|\mathrm{T}(\mathcal{A})| \hookrightarrow |\operatorname{Flag}(\operatorname{PoBar}(\mathcal{A}))|$ is a homotopy equivalence. In particular, it follows that the blowup complex is homotopy equivalent to $|\operatorname{Flag}(\operatorname{PoBar}(\mathcal{A}))|$.

Definition 5.11. Let K be a simplicial complex, and let $\mathcal{A} = (K_i \subseteq K)_{i \in I}$ be a cover of K by subcomplexes. Let $\mathrm{T}(\mathcal{A})$ be the subcomplex of $\operatorname{Flag}(\operatorname{PoBar}(\mathcal{A}))$ consisting of the simplices $(\sigma_0, J_0) < \cdots < (\sigma_m, J_m)$ such that $\sigma_m \in K_{J_0}$.

The letter T stands for "triangulation", since $\mathrm{T}(\mathcal{A})$ turns out to be a triangulation of the blowup complex.

Lemma 5.12. *Let K be a simplicial complex, and let $\mathcal{A} = (K_i \subseteq K)_{i \in I}$ be a cover of K by subcomplexes; write $|\mathcal{A}| = (|K_i| \subseteq |K|)_{i \in I}$. Then there is a homeomorphism $\varphi\colon \operatorname{Blowup}(|\mathcal{A}|) \to |\mathrm{T}(\mathcal{A})|$ such that the following diagram commutes:*

$$\begin{array}{ccccc} & & |\operatorname{Flag}(\operatorname{PoBar}(\mathcal{A}))| & & \\ & \swarrow^{\lambda_S} & \uparrow & \searrow^{\lambda_N} & \\ |\operatorname{Sd}(K)| & \longleftarrow & |\mathrm{T}(\mathcal{A})| & \longrightarrow & |\operatorname{Sd}(\operatorname{Nrv}(\mathcal{A}))| \\ \cong\uparrow & & \varphi\uparrow & & \uparrow\cong \\ |K| & \xleftarrow[\rho_S]{} & \operatorname{Blowup}(|\mathcal{A}|) & \xrightarrow[\rho_N]{} & |\operatorname{Nrv}(\mathcal{A})| \end{array} \tag{5.3}$$

Here, the vertical maps on the left and right are the standard homeomorphisms.

Proof. The blowup complex $\operatorname{Blowup}(|\mathcal{A}|)$ is defined by glueing together pieces of the form $|K_J| \times |J|$ for $J \in \operatorname{Nrv}(\mathcal{A})$. We abuse notation and write J also for the full simplicial complex on J. For any such piece, define φ by the composition

$$\begin{aligned} |K_J| \times |J| &\cong |\operatorname{Sd} K_J| \times |\operatorname{Sd} J| \\ &= |\operatorname{Flag}(\operatorname{Pos}(K_J))| \times |\operatorname{Flag}(\operatorname{Pos}(J))| \\ &\cong |\operatorname{Flag}(\operatorname{Pos}(K_J))| \times |\operatorname{Flag}(\operatorname{Pos}(J)^{\mathrm{op}})| \\ &\cong |\operatorname{Flag}(\operatorname{Pos}(K_J) \times \operatorname{Pos}(J)^{\mathrm{op}})| \subseteq |\mathrm{T}(\mathcal{A})| \end{aligned}$$

where the last homeomorphism is an instance of 5.2. As these maps respect the equivalence relation from the definition of the blowup complex, together they define a continuous map $\varphi\colon \operatorname{Blowup}(|\mathcal{A}|) \to |\mathrm{T}(\mathcal{A})|$. By construction, the diagram 5.3 commutes.

To see that φ is a homeomorphism, we can construct its inverse. As J varies, the subcomplexes $\mathrm{Flag}(\mathrm{Pos}(K_J) \times \mathrm{Pos}(J)^{\mathrm{op}})$ cover $\mathrm{T}(\mathcal{A})$. For each J, we can reverse the homeomorphisms in the definition of φ to define φ^{-1} on $|\mathrm{Flag}(\mathrm{Pos}(K_J) \times \mathrm{Pos}(J)^{\mathrm{op}})|$. Since these maps agree on intersections, they glue together to define the inverse of φ, $\varphi^{-1}\colon |\mathrm{T}(\mathcal{A})| \to \mathrm{Blowup}(|\mathcal{A}|)$. $\square$

Lemma 5.13. *Let K be a simplicial complex, and let $\mathcal{A} = (K_i \subseteq K)_{i \in I}$ be a cover of K by subcomplexes. Then there exists a collapse $\mathrm{Flag}(\mathrm{PoBar}(\mathcal{A})) \searrow \mathrm{T}(\mathcal{A})$. In particular, the inclusion $|\mathrm{T}(\mathcal{A})| \hookrightarrow |\mathrm{Flag}(\mathrm{PoBar}(\mathcal{A}))|$ is a homotopy equivalence.*

Proof. We construct a discrete gradient V on $\mathrm{Flag}(\mathrm{PoBar}(\mathcal{A}))$ such that the set of critical simplices is $\mathrm{T}(\mathcal{A})$. Then we will argue that it follow from Proposition 2.43 that the inclusion $|\mathrm{T}(\mathcal{A})| \hookrightarrow |\mathrm{Flag}(\mathrm{PoBar}(\mathcal{A}))|$ is a homotopy equivalence.

To this end, let $L = \mathrm{Flag}(\mathrm{PoBar}(\mathcal{A})) \setminus \mathrm{T}(\mathcal{A})$ and consider the function $f\colon L \to \mathbb{N} \cup \{\infty\}$ that assigns to a simplex $\tau = ((\sigma_0, J_0) < \cdots < (\sigma_m, J_m))$ the value

$$f(\tau) = \min\{i \in \{0, \dots, m-1\} \mid \sigma_i < \sigma_{i+1},\ J_i < J_{i+1}\}$$

with $f(\tau) = \infty$ if no such i exists. Moreover, we consider the function $g\colon L \to \mathbb{N} \cup \{\infty\}$ that assigns to a simplex $\tau = ((\sigma_0, J_0) < \cdots < (\sigma_m, J_m))$ the value

$$g(\tau) = \min\{i \in \{0, \dots, m-2\} \mid \sigma_i = \sigma_{i+1} < \sigma_{i+2},\ J_i < J_{i+1} = J_{i+2}\}$$

with $g(\tau) = \infty$ if no such i exists. As $\tau \in L$ we have $\sigma_m \notin K_{J_0}$, by definition of $\mathrm{T}(\mathcal{A})$, and hence we get that if $f(\tau) = \infty$ then $g(\tau) < \infty$, as otherwise $\sigma_m \in K_{J_0}$. This and the definitions of the function f and g imply that either $f(\tau) < g(\tau)$ or $f(\tau) > g(\tau) + 1$.

We now define the discrete vector field $\hat{V}$ on $\mathrm{Flag}(\mathrm{PoBar}(\mathcal{A}))$ that partitions L into pairs of simplices and where we let every simplex in $\mathrm{T}(\mathcal{A})$ be critical. Take any simplex τ in $L = \mathrm{Flag}(\mathrm{PoBar}(\mathcal{A})) \setminus \mathrm{T}(\mathcal{A})$, i.e., a chain $(\sigma_0, J_0) < \cdots < (\sigma_m, J_m)$ in $\mathrm{PoBar}(\mathcal{A})$ such that $\sigma_m \notin K_{J_0}$. If $i = f(\tau) < g(\tau)$, consider the chain

$$(\sigma_0, J_0) < \cdots < (\sigma_i, J_i) < (\sigma_i, J_{i+1}) < (\sigma_{i+1}, J_{i+1}) < \cdots < (\sigma_m, J_m)$$

and pair the corresponding simplex μ in L with τ; note that $f(\mu) > f(\tau) + 1 = g(\mu) + 1$. We verify that V is a discrete vector field that partitions L: For any simplex $\mu = ((\tilde{\sigma}_0, \tilde{J}_0) < \cdots < (\tilde{\sigma}_m, \tilde{J}_m)) \in L$ with $f(\mu) > g(\mu) + 1 = j$ consider the facet τ of μ that skips the element $(\tilde{\sigma}_j, \tilde{J}_j)$; note that $f(\tau) = g(\mu) < g(\tau)$. It is straightforward to see that the sets $\{\tau \in L \mid f(\tau) < g(\tau)\}$ and $\{\tau \in L \mid f(\tau) > g(\tau) + 1\}$ partition L and that the above constructions yield mutually inverse bijections between those sets, implying that $\hat{V}$ is a discrete vector field that partitions L into pairs of simplices.

We prove that the collection of regular pairs V in $\hat{V}$ is a discrete gradient by showing that there are no non-trivial closed V-paths: Consider any V-path

$$\tau_0 \to \mu_0 \leftarrow \cdots \to \mu_r \leftarrow \tau_{r+1}$$

with $(\tau_i, \mu_i) \in V$ and $\tau_i \neq \tau_{i+1}$ for all i. To show that it is not closed, i.e., $\tau_{r+1} \neq \tau_0$, consider first any chain $((\sigma, J) < (\tilde{\sigma}, \tilde{J})) \in L$ of length 2 and the set

$$S = \{\tau \in L \mid \min \tau = (\sigma, J),\ \max \tau = (\tilde{\sigma}, \tilde{J})\}.$$

Note that $\hat{V}$ restricts to a partition of S. If $R \neq S$ is another such set of chains with $\tau_0 \in S$ and $\tau_i \in R$ for some i, then $\tau_j \notin S$ for all $j \geq i$ as at least one of the inequalities $\min \tau_0 \leq \min \tau_j$ and $\max \tau_j \leq \max \tau_0$ is strict. Therefore the V-path cannot be closed. Moreover, as S is finite, there are only finitely many possible other such R. Therefore, it is enough to show that for any such S there is no non-trivial closed V-path with $\tau_i \in S$ for all i.

To this end, we construct a lexicographic partial order on the set of chains S. First, we consider the product inclusion order on pairs of simplices of K and subsets of I, given by $(\sigma, J) \subseteq (\tilde{\sigma}, \tilde{J})$ if and only if $\sigma \subseteq \tilde{\sigma}$ and $J \subseteq \tilde{J}$. Now, we extend this partial order to a partial order on S: For any two chains $\tau = ((\sigma_0, J_0) < \cdots < (\sigma_m, J_m))$ and $\mu = ((\tilde{\sigma}_0, \tilde{J}_0) < \cdots < (\tilde{\sigma}_m, \tilde{J}_m))$ in S of equal length, we let $\tau \leq_{\text{lex}} \mu$ if $\tau = \mu$ or if for the smallest index j with $(\sigma_j, J_j) \neq (\tilde{\sigma}_j, \tilde{J}_j)$ we have $(\sigma_j, J_j) \subseteq (\tilde{\sigma}_j, \tilde{J}_j)$. We show that for any two gradient pairs $(\tau, \mu), (\tilde{\tau}, \tilde{\mu}) \in V$ with $\tilde{\tau}$ a facet of μ, we have $\mu >_{\text{lex}} \tilde{\mu}$, proving that the V-path above cannot be closed. We have

$$\mu = (\sigma_0, J_0) < \cdots < (\sigma_i, J_i) < (\sigma_i, J_{i+1}) < (\sigma_{i+1}, J_{i+1}) < \cdots < (\sigma_m, J_m)$$

with $i = f(\tau)$, and $\tilde{\tau}$ is a facet of μ that skips some element (σ_j, J_j) with $0 < j < m$. Note that j cannot be greater than $i+1$, as otherwise $g(\tilde{\tau}) = i < f(\tilde{\tau})$, contradicting the assumption that $\tilde{\tau}$ is a gradient facet of $\tilde{\mu}$ and therefore $f(\tilde{\tau}) < g(\tilde{\tau})$. Moreover, if $j \leq i$ then $f(\tilde{\tau}) = j - 1$, and if $j = i + 1$ then $f(\tilde{\tau}) = j$. In any case, $\tilde{\mu}$ is obtained by adding the element (σ_{j-1}, J_{j+1}) to $\tilde{\tau}$. Now $\sigma_{j-1} \subseteq \sigma_j$ and $J_j \supseteq J_{j+1}$, and thus $(\sigma_{j-1}, J_{j+1}) \subseteq (\sigma_j, J_j)$ and $\tilde{\mu} \leq_{\text{lex}} \mu$. As $\tilde{\tau} \neq \tau$, we must have $\tilde{\mu} \neq \mu$ and thus $\tilde{\mu} <_{\text{lex}} \mu$.

The reasoning above also implies that for every simplex in $\operatorname{Flag}(\operatorname{PoBar}(\mathcal{A}))$ its V-path height is finite and hence it follows from Lemma 2.40 and Proposition 2.43 that the inclusion $|\operatorname{T}(\mathcal{A})| \hookrightarrow |\operatorname{Flag}(\operatorname{PoBar}(\mathcal{A}))|$ is a homotopy equivalence. □

Theorem 5.14. *Let K be a simplicial complex and let $\mathcal{A} = (K_i \subseteq K)_{i \in I}$ be a good cover of K by subcomplexes. Then, the natural maps $\rho_S \colon \operatorname{Blowup}(|\mathcal{A}|) \to |K|$ and $\rho_N \colon \operatorname{Blowup}(|\mathcal{A}|) \to |\operatorname{Nrv}(\mathcal{A})|$ are homotopy equivalences.*

Proof. Consider Diagram 5.3. By Proposition 5.10 the maps λ_S and λ_N are homotopy equivalences. By Lemma 5.13, the inclusion $|\operatorname{T}(\mathcal{A})| \hookrightarrow |\operatorname{Flag}(\operatorname{PoBar}(\mathcal{A}))|$ is a homotopy equivalence, and by Lemma 5.12, φ is a homeomorphism. It follows that ρ_S and ρ_N are homotopy equivalences. □

5.3.2. A Functorial Nerve Theorem for Semi-Algebraic Covers

As a corollary of Theorem 5.14, we get a functorial nerve theorem for finite, closed, semi-algebraic covers of compact semi-algebraic sets. For this, we need a well known theorem on the existence of triangulations of semi-algebraic sets [BCR98, Theorem 9.2.1], which we now state. For K a simplicial complex and σ a simplex of K, we write $\operatorname{int}|\sigma| = |\sigma| \setminus |\partial\sigma| \subset |K|$ for the open simplex.

Lemma 5.15. *Let $S \subset \mathbb{R}^n$ be a compact semi-algebraic set, and let $(S_i)_{i=0}^q$ be a finite family of semi-algebraic subsets of S. There is a finite simplicial complex $K = \{\sigma_j\}_{j=0}^p$ and a homeomorphism $h\colon |K| \to S$, such that every S_i is the union of some images of simplices $h(\operatorname{int}|\sigma_j|)$.*

Theorem 5.16. *Let $S \subset \mathbb{R}^n$ be a compact semi-algebraic set, and let $\mathcal{A} = (S_i)_{i=0}^q$ be a finite good cover of S such that each S_i is semi-algebraic and closed in S. Then, the natural maps $\rho_S\colon \mathrm{Blowup}(\mathcal{A}) \to S$ and $\rho_N\colon \mathrm{Blowup}(\mathcal{A}) \to |\mathrm{Nrv}(\mathcal{A})|$ are homotopy equivalences.*

Proof. By Lemma 5.15, there is a simplicial complex K, a homeomorphism $h\colon |K| \to S$, and a cover of K by subcomplexes $\mathcal{B} = (K_i \subseteq K)_{i=0}^q$ such that $h|_{K_i}$ is a homeomorphism between K_i and S_i. Then, h induces a homeomorphism $\mathrm{Blowup}(|\mathcal{B}|) \to \mathrm{Blowup}(\mathcal{A})$ such that the following diagram commutes:

$$\begin{array}{ccccc} |K| & \xleftarrow{\rho_S} & \mathrm{Blowup}(|\mathcal{B}|) & \xrightarrow{\rho_N} & |\mathrm{Nrv}(\mathcal{B})| \\ {\scriptstyle h}\downarrow & & \downarrow & & \Vert \\ S & \xleftarrow[\rho_S]{} & \mathrm{Blowup}(\mathcal{A}) & \xrightarrow[\rho_N]{} & |\mathrm{Nrv}(\mathcal{A})| \end{array}$$

By Theorem 5.14 the top horizontal maps are homotopy equivalences, and the corollary follows. ☐

5.3.3. A Functorial Version of Björner's Nerve Theorem

If K is a simplicial complex, and $\mathcal{A}$ is a locally finite cover of K by subcomplexes, then we have a comparison map $|K| \to |\mathrm{Nrv}(\mathcal{A})|$ induced by the map of posets defined in the proof of Proposition 5.8, $f\colon \mathrm{Pos}(K) \to \mathrm{Pos}(\mathrm{Nrv}(\mathcal{A}))^{\mathrm{op}}$. In [Bjö03], Björner gives a detailed analysis of how the connectivity of this map is affected by the connectivity of the finite intersections of cover elements. For the final result of this section, we will use the PoBar construction and the blowup complex to prove a functorial version of Björner's theorem.

Definition 5.17. Let $k \geq 0$. A topological space X is *k-connected* if, for every $0 \leq r \leq k$, every map of the r-sphere into X is homotopic to a constant map.

Proposition 5.18 (Björner [Bjö03, Theorem 6])**.** *Let K be a simplicial complex, let $\mathcal{A}$ be a locally finite cover of K by subcomplexes, and let $k \geq 0$. Assume that every non-empty intersection $K_{i_1} \cap \cdots \cap K_{i_t}$ is $(k - t + 1)$-connected, for all $t \geq 1$. Then f induces a bijection*

$$\pi_0(K) \cong \pi_0(\mathrm{Nrv}(\mathcal{A})),$$

and for all $1 \leq j \leq k$, and for all $x \in |K|$, f induces an isomorphism

$$\pi_j(|K|, x) \cong \pi_j(|\mathrm{Nrv}(\mathcal{A})|, f_*(x)).$$

In fact, the theorem of Björner deals with regular CW complexes, rather than simplicial complexes. The assumption that $\mathcal{A}$ is locally finite is omitted from the original statement, but it is used in the proof. For convenience, Björner assumes that K is connected: Proposition 5.18 follows from Björner's theorem and the following lemma, which is easily proved.

Lemma 5.19. *Let K be a simplicial complex, and let $\mathcal{A} = (K_i)_{i\in I}$ be a locally finite cover of K by subcomplexes such that K_i is non-empty and connected for all $i \in I$. Then f induces a bijection $\pi_0(K) \cong \pi_0(\mathrm{Nrv}(\mathcal{A}))$.* ☐

Using the PoBar construction and the blowup complex as before, we obtain the following functorial version of Björner's theorem.

Theorem 5.20. *Let $k \geq 0$, let K be a simplicial complex, and let $\mathcal{A} = (K_i)_{i\in I}$ be a locally finite cover of K by subcomplexes. Assume that every non-empty intersection $K_{i_1} \cap \cdots \cap K_{i_t}$ is $(k - t + 1)$-connected, for all $t \geq 1$. The natural map $\rho_S\colon \mathrm{Blowup}(|\mathcal{A}|) \to |K|$ is a homotopy equivalence, and the natural map $\rho_N\colon \mathrm{Blowup}(|\mathcal{A}|) \to |\,\mathrm{Nrv}(\mathcal{A})|$ induces a bijection in path components, and for all $1 \leq j \leq k$, and for all $x \in \mathrm{Blowup}(|\mathcal{A}|)$, ρ_N induces an isomorphism $\pi_j(\mathrm{Blowup}(|\mathcal{A}|), x) \cong \pi_j(|\,\mathrm{Nrv}(\mathcal{A})|, \rho_N(x))$.*

Proof. Note that the proof of Proposition 5.10 shows that the poset map λ_S is a homotopy equivalence, and that the triangle commutes up to homotopy, without the assumption that $\mathcal{A}$ is good. So, by Björner's Proposition 5.18, $\lambda_N\colon |\,\mathrm{Flag}(\mathrm{PoBar}(\mathcal{A}))| \to |\,\mathrm{Sd}(\mathrm{Nrv}(\mathcal{A}))|$ induces a bijection in path components, and for all $1 \leq j \leq k$, and for all $x \in |\,\mathrm{Flag}(\mathrm{PoBar}(\mathcal{A}))|$, λ_N induces an isomorphism $\pi_j(|\,\mathrm{Flag}(\mathrm{PoBar}(\mathcal{A}))|, x) \cong \pi_j(|\,\mathrm{Nrv}(\mathcal{A})|, \lambda_N(x))$. The result follows from commutativity of Diagram 5.3. ☐

5.4. A Unified Nerve Theorem

We now prove the unified nerve theorem (Theorem H in Section 1.4), which subsumes Theorems 5.4 and 5.16 as special cases, and which implies Theorem 5.14 with the additional assumption that the cover by subcomplexes is locally finite and locally finite dimensional. Like in Theorem 5.4, we use the connection between the blowup complex and the bar construction (explained in Section 5.1) to deduce statements about the blowup complex from the corresponding statements about the bar construction. Since the bar construction is a standard tool in homotopy theory, we can use well-known results to prove the requisite properties. We begin by introducing the remaining notions from topology we need to state the unified nerve theorem; see also Section 2.4.2.

A simplicial complex K is sometimes said to be *locally finite dimensional* if every vertex v of K has a finite dimensional star, i.e., $\sup\{\dim\sigma \mid v \in \sigma\} < \infty$. Following this usage, we say that a cover of a topological space is locally finite dimensional if the nerve of the cover is so. More explicitly, we have the following:

Definition 5.21. If X is a topological space, and $\mathcal{A} = (A_i)_{i\in I}$ is a cover, then $\mathcal{A}$ is *locally finite dimensional* if for each cover element A_i there exists $k_i \in \mathbb{N}$ such that for any $J \subseteq I$ with $A_J \neq \emptyset$ and $i \in J$, we have $|J| \leq k_i$.

Definition 5.22. Let R be a commutative ring. We say that a continuous map f between topological spaces is an *R-homology isomorphism* if $H_n(f, R)$ is an isomorphisms for all $n \geq 0$. We say that a cover $\mathcal{A} = (A_i)_{i\in I}$ is *homologically good* with respect to R if, for all non-empty $J \subseteq I$ such that $A_J \neq \emptyset$, the map to the one point space $A_J \to *$ is an R-homology isomorphism.

Definition 5.23. We say that a cover $\mathcal{A} = (A_i)_{i\in I}$ is *weakly good* if, for all non-empty $J \subseteq I$ such that $A_J \neq \emptyset$, the map $A_J \to *$ is a weak homotopy equivalence, where $*$ is the one point space.

Definition 5.24. Let X be a topological space, and let $\mathcal{A} = (A_i)_{i\in I}$ be a cover. For $T \in \mathrm{Nrv}(\mathcal{A})$, the *latching space* is the subset

$$L(T) := \bigcup_{T \subsetneq J \subseteq I} A_J \subseteq A_T \ .$$

We can now state the unified version of the Nerve Theorem.

Theorem 5.25 (Unified Nerve Theorem)**.** *Let X be a topological space and let $\mathcal{A} = (A_i)_{i\in I}$ be a cover of X.*

1. *Consider the natural map $\rho_S\colon \mathrm{Blowup}(\mathcal{A}) \to X$.*
   a) *If $\mathcal{A}$ is an open cover, then ρ_S is a weak homotopy equivalence. If furthermore X is a paracompact Hausdorff space, or, more generally, if $\mathcal{A}$ is numerable, then ρ_S is a homotopy equivalence.*
   b) *Assume that X is compactly generated and that $\mathcal{A}$ is a closed cover that is locally finite and locally finite dimensional. If for any $T \in \mathrm{Nrv}(\mathcal{A})$ the latching space $L(T) \subseteq A_T$ is a closed subset and the pair $(A_T, L(T))$ satisfies the homotopy extension property, then ρ_S is a homotopy equivalence.*
2. *Consider the natural map $\rho_N\colon \mathrm{Blowup}(\mathcal{A}) \to |\mathrm{Nrv}(\mathcal{A})|$.*
   a) *If $\mathcal{A}$ is (weakly) good, then ρ_N is a (weak) homotopy equivalence.*
   b) *If for all $J \in \mathrm{Nrv}(\mathcal{A})$ the space A_J is compactly generated and $\mathcal{A}$ is homologically good with respect to a coefficient ring R, then ρ_N is an R-homology isomorphism.*

We prove Theorem 5.25 in Section 5.4.2.

Remark 5.26. The compactly generated assumption in 2(b) is satisfied for example if X is a locally compact Hausdorff space and $\mathcal{A}$ is an open cover. The assumption also holds if X is compactly generated and $\mathcal{A}$ is a closed cover; this also includes the case of a cover of a CW-complex by subcomplexes. See Example 2.64.

Remark 5.27. If X is a regular CW-complex and $\mathcal{A}$ is a cover of subcomplexes, then 2(b) can also be proven using spectral sequence techniques [Bro94, Chapter VII, Section 4]. Note that in this reference, the total complex of the double complex associated to the cover is isomorphic to the cellular chain complex of $\mathrm{Blowup}(\mathcal{A})$. Moreover, these techniques can also be used to prove an analogous statement to Proposition 5.18 for homology groups [Mes01, Theorem 2.1].

Remark 5.28. To illustrate the role of the technical assumptions in the unified nerve theorem, we now discuss some counterexamples when these assumptions are violated.

- The classical nerve theorem for a good open cover of a paracompact Hausdorff space is proven using 1(a) and 2(a). We will now give an example that shows that this paracompactness Hausdorff assumption, which ensures that the open cover is numerable, cannot be omitted in order to establish a homotopy equivalence between space and nerve. Consider the *long ray* L, which is constructed as follows: Take the first uncountable ordinal ω_1, which is a well-ordered set and its elements are all countable ordinals, and insert a unit interval $(0,1)$ between each countable ordinal α and its successor $\alpha+1$, yielding a totally ordered set. The topology on L is given by the order topology. The long ray is a standard example for a non-paracompact Hausdorff space that is also not contractible [SS95; Jos83]. However, L is weakly contractible and for any point $p \in L$ the open set $L_{<p} = \{t \in L \mid t < p\} \subset L$ is homeomorphic to the interval $[0,1)$. Thus, the open cover $\mathcal{A} = (L_{<p})_{p\in\omega_1}$ is a good cover of L and it follows from 1(a) and 2(a) that the nerve $|\operatorname{Nrv}\mathcal{A}|$ is weakly contractible and hence contractible by Whitehead's theorem [Hat02, Theorem 4.5]. This implies that the space L and the nerve $|\operatorname{Nrv}\mathcal{A}|$ are not homotopy equivalent.

  Note that the Hausdorff assumption cannot be dropped either; there exist non-Hausdorff paracompact spaces with good open covers that are not homotopy equivalent to the nerve of the cover. Specifically, any finite simplicial complex K is weakly homotopy equivalent to a finite topological space X whose points correspond to the simplices and whose open sets are upsets in the face poset [McC66]. The open sets corresponding to vertex stars form a good open cover of X whose nerve is isomorphic to K. However, it is straightforward to verify that every map $X \to |K|$ is locally constant, and therefore $|K|$ and X are not homotopy equivalent in general.

- The finiteness conditions in 1(b) control the size of the cover. If $\mathcal{A}$ is the cover of the circle S^1 by its points, then all conditions in 1(b) and 2(a) are satisfied except the locally finiteness assumption. As the nerve $|\operatorname{Nrv}\mathcal{A}|$ is a disjoint union of points, it is not homotopy equivalent to S^1.

- Even if we are only interested in finite good and closed covers, the covered space does not need to be homotopy equivalent to the nerve of the cover. Consider the *double comb space* C and denote the two combs by A_1 and A_2 (see Fig. 15). Then, the nerve $|\operatorname{Nrv}\mathcal{A}|$ of the finite good and closed cover $\mathcal{A} = \{A_1, A_2\}$ of C is contractible. Hence, it can not be homotopy equivalent to C, because the latter is not contractible. In this example, the pairs $(A_1, A_1 \cap A_2)$ and $(A_2, A_1 \cap A_2)$ do not satisfy the homotopy extension property. This shows that the conditions on the latching spaces are crucial, as all other assumptions in 1(b) and 2(a) are satisfied.

- If $\mathcal{A}$ is any homologically good open cover of a locally compact Hausdorff space X, then it follows from 1(a) and 2(b) that the space X and the nerve $\operatorname{Nrv}\mathcal{A}$ have isomorphic homology groups. This conclusion does not hold if one replaces the open cover by a closed cover. Consider the *Warsaw circle* $W \subseteq S^2$ that separates the sphere into two connected components U_1 and U_2 (see Fig. 15). The closed sets $A_1 = U_1 \cup W$ and $A_2 = U_2 \cup W$ cover the sphere and are contractible. Moreover, the intersection $A_1 \cap A_2 = W$ is acyclic and hence $\mathcal{A} = \{A_1, A_2\}$ is a homologically good closed cover of S^2. Nevertheless, the space S^2 and the nerve $\operatorname{Nrv} \mathcal{A}$ do not have isomorphic homology groups, as $H_2(S^2) \cong \mathbb{Z}$ and $H_2(\operatorname{Nrv} \mathcal{A}) \cong 0$. Hence, the conditions on the latching spaces are crucial, as all other assumptions in 1(b) and 2(b) are satisfied. This counterexample also shows that the nerve of a weakly good closed cover is not necessarily weakly equivalent to the space it covers.

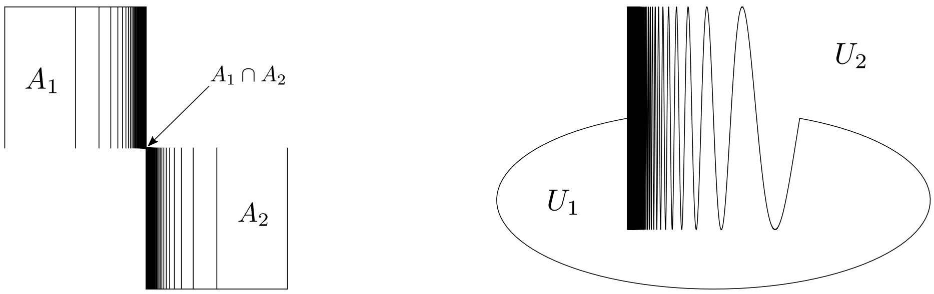

Figure 15: The double comb space C (left) and the Warsaw circle W (right).

5.4.1. Applications of the Unified Nerve Theorem

The assumptions on the latching spaces in 1(b) of Theorem 5.25 might not be easy to check in all situations. We now give a reformulation, and a union theorem for pairs that satisfy the homotopy extension property, which help to verify these assumptions. We also show in this subsection that Theorem 5.25 implies the functorial nerve theorems Theorem 5.4 and Theorem 5.14.

To study pairs that satisfy the homotopy extension property it suffices to consider neighborhood deformation retracts.

Definition 5.29. A pair of topological spaces (X, A) is called an *NDR-pair* if there exist continuous maps $u\colon X \to [0,1]$ and $h\colon X \times [0,1] \to X$ such that

(i) $A = u^{-1}(0)$

(ii) $h(-,0) = \mathrm{id}_X$

(iii) $h(a,-) = a$ for all $a \in A$

(iv) $h(x,1) \in A$ for all $x \in X$ with $u(x) < 1$.

Proposition 5.30 ([Koz08, Proposition 7.7])**.** *Let A be a closed subspace of X. Then (X, A) is an NDR-pair if and only if (X, A) satisfies the homotopy extension property.*

The following union theorem is due to Lillig [Lil73].

Proposition 5.31. *Let $A_0, \dots, A_n \subseteq X$ be closed subsets and assume that for all $J \subseteq [n]$ the pair (X, A_J) satisfies the homotopy extension property. Then the pair $(X, \bigcup_{i=0}^n A_i)$ also satisfies the homotopy extension property.*

This proposition, together with 1(b) in Theorem 5.25, implies the following corollary, which does not involve the latching spaces.

Corollary 5.32. *Let X be a compactly generated topological space and $\mathcal{A} = (A_i)_{i\in[n]}$ a finite closed cover. Assume that for all $I \subseteq J \subseteq [n]$ the pair (A_I, A_J) satisfies the homotopy extension property. Then $\rho_S\colon \mathrm{Blowup}(\mathcal{A}) \to X$ is a homotopy equivalence.*

We will now illustrate how these statements can be used to deduce the functorial nerve theorem for closed convex sets in $\mathbb{R}^d$ (Theorem 5.4) from the unified nerve theorem (Theorem 5.25). The proof of the following lemma is elementary and left to the reader.

Lemma 5.33. *Let $K \subseteq \mathbb{R}^d$ be compact and convex. Assume that $\operatorname{aff} K = \mathbb{R}^d$, where $\operatorname{aff} K$ is the affine hull of K. Then $\operatorname{int} K$ is convex and $\overline{\operatorname{int} K} = K$.*

Proposition 5.34. *Let K and K' be compact and convex sets in $\mathbb{R}^d$ with $K \subseteq K'$. Then the pair (K', K) satisfies the homotopy extension property.*

Proof. Without loss of generality, we can assume that $\operatorname{aff} K' = \mathbb{R}^d$ holds.

First of all, let us assume that K is the intersection of K' with an affine subspace. Now, choose a point x in K. By Lemma 5.33 and the proof of [Mun84, Lemma 1.1], we see that there exists a homeomorphism $\varphi\colon \mathbb{R}^d \to \mathbb{R}^d$ such that $\varphi(K') = \overline{\mathbb{B}}_1(0)$, $\varphi(x) = 0$ and $\varphi(K) = \overline{\mathbb{B}}_1(0) \cap \mathbb{R}^l \times \{0\}^{d-l}$, with $l = \dim \operatorname{aff} K$. The pair

$$(\varphi(K'), \varphi(K)) = (\overline{\mathbb{B}}_1(0), \overline{\mathbb{B}}_1(0) \cap \mathbb{R}^l \times \{0\}^{d-l})$$

is a CW-pair and hence satisfies the homotopy extension property (Remark 2.66).

Now let $K \subseteq K'$ be such that $\operatorname{aff} K = \operatorname{aff} K' = \mathbb{R}^d$. As before, choose a point $x \in K$ and let $\varphi\colon \mathbb{R}^d \to \mathbb{R}^d$ be a homeomorphism with $\varphi(K) = \overline{\mathbb{B}}_1(0)$, $\varphi(x) = 0$ and $\varphi(K')$ star-shaped with respect to 0. It is easy to see that $(\varphi(K'), \varphi(K))$ is an NDR-pair. Hence, it follows from Proposition 5.30 that the pair (K', K) satisfies the homotopy extension property.

Finally, for arbitrary compact convex sets $K \subseteq K'$ we factor the inclusion as $K \hookrightarrow \operatorname{aff} K \cap K' \hookrightarrow K'$. The claim now follows from the previous two cases together with transitivity of the homotopy extension property: if two pairs (X, Y) and (Y, Z) satisfy the homotopy extension property, then so does the pair (X, Z). ☐

Using a truncation argument, we extend this result to any pair of closed convex sets.

Proposition 5.35. *Let K and K' be closed and convex sets in $\mathbb{R}^d$ with $K \subseteq K'$. Then the pair (K', K) satisfies the homotopy extension property.*

Proof. We verify Definition 2.65: Let $f\colon K' \to Y$ be any continuous map and $H\colon K \times [0,1] \to Y$ be any homotopy with $H(-,0) = f|_K$. We inductively construct an extension $\tilde{H}\colon K' \times [0,1] \to Y$ of H with $\tilde{H}(-,0) = f$. To this end, consider for every $n \in \mathbb{N}$ the compact convex sets $K_n = D_n(0) \cap K \subseteq K$ and $K'_n = D_n(0) \cap K' \subseteq K'$, where $D_n(0)$ is the closed ball of radius n centered at the origin. Denote by $f_n\colon K'_n \to Y$ and $H_n\colon K_n \times [0,1] \to Y$ the restrictions of f and H, respectively.

By Proposition 5.34, the pair (K'_1, K_1) satisfies the homotopy extension property. Hence, there exists an extension $\tilde{H}_1\colon K'_1 \times [0,1] \to Y$ of H_1 that satisfies $\tilde{H}_1(-,0) = f_1$.

Let $n \in \mathbb{N}$ be arbitrary and consider the following diagram of inclusions:

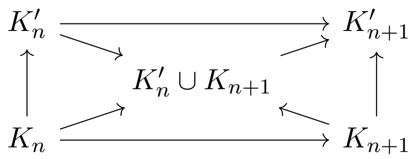

By Proposition 5.31 and Proposition 5.34, the pair $(K'_{n+1}, K'_n \cup K_{n+1})$ satisfies the homotopy extension property. Hence, we can extend the homotopy on $K'_n \cup K_{n+1}$ induced by $(\tilde{H}_n, H_{n+1})$ to a homotopy $\tilde{H}_{n+1}\colon K'_{n+1} \times [0,1] \to Y$ that satisfies $\tilde{H}_{n+1}(-,0) = f_{n+1}$. Taking to the colimit over all n yields the desired extension $\tilde{H}$. □

Let $X \subset \mathbb{R}^d$ be a subset and let $\mathcal{B} = (C_i)_{i\in[n]}$ be a finite cover of X by closed convex sets. The previous corollary shows, together with the fact that any closed subset of $\mathbb{R}^d$ is compactly generated (Example 2.64), that the assumptions in Corollary 5.32 are satisfied and hence, the map $\rho_S\colon \operatorname{Blowup}(\mathcal{B}) \to X$ is a homotopy equivalence. As any cover by convex sets is good, it follows from 2(a) in Theorem 5.25 that the map $\rho_N\colon \operatorname{Blowup}(\mathcal{B}) \to |\operatorname{Nrv}(\mathcal{B})|$ is a homotopy equivalence as well. This proves Theorem 5.4.

If in the functorial nerve theorem for covers by subcomplexes (Theorem 5.14) we additionally assume that the cover is locally finite dimensional, then this theorem also follows readily from the unified nerve theorem (Theorem 5.25): the realization of a simplicial complex is compactly generated (Example 2.64); moreover, the latching spaces are subcomplexes and thus satisfy the homotopy extension properties (Remark 2.66).

5.4.2. Proof of the Unified Nerve Theorem

We now define the bar construction in the setting of a simplicial model category (see Section 2.4.2), generalizing Definition 2.57.

Definition 5.36. Let P be a poset and let $\mathcal{M}$ be a simplicial model category. The *simplicial bar construction* of a functor $F\colon P \to \mathcal{M}$ is the simplicial object

$$\operatorname{Bar}_\bullet(F)\colon \Delta^{\mathrm{op}} \to \mathcal{M}$$

whose n-simplices $\operatorname{Bar}_n(F)$ are defined by the coproduct

$$\operatorname{Bar}_n(F) = \coprod_{v_0 \le v_1 \le \cdots \le v_n} F(v_0).$$

Equivalently, the coproduct is indexed by functors of the form $\gamma\colon [n] \to P$. For any map $\theta\colon [m] \to [n]$ in Δ, $\theta^*\colon \operatorname{Bar}_n(F) \to \operatorname{Bar}_m(F)$ takes the summand indexed by γ to the summand indexed by $\gamma \circ \theta$, via the map $F(\gamma(0)) \to F(\gamma(\theta(0)))$.

The identifications that were used to define the bar construction for topological spaces are achieved in this setting by the categorical notion of a *coend*:

Definition 5.37. Let $\mathcal{C}$ be a small category, $\mathcal{E}$ any category, and $H\colon \mathcal{C}^{\mathrm{op}} \times \mathcal{C} \to \mathcal{E}$ a functor. The *coend* $\int^{\mathcal{C}} H$, sometimes denoted $\int^{c\in\mathcal{C}} H(c,c)$, is an object of $\mathcal{E}$ equipped with arrows $H(c,c) \to \int^{\mathcal{C}} H$ for each $c \in \mathcal{C}$ that are collectively universal with the property that the diagram

$$\begin{array}{ccc} H(c',c) & \xrightarrow{f_*} & H(c',c') \\ {\scriptstyle f^*}\downarrow & & \downarrow \\ H(c,c) & \longrightarrow & \int^{\mathcal{C}} H \end{array}$$

commutes for each $f\colon c \to c'$ in $\mathcal{C}$.

We can use this notion to define the geometric realization functor $|-|\colon \mathsf{sSet} \to \mathsf{CGSpc}$ mentioned in Section 2.4.2. Writing Δ for the simplex category as before, there is a functor $\Delta \to \mathsf{CGSpc}$ that takes $[n]$ to the standard n-simplex. If $X\colon \Delta^{\mathrm{op}} \to \mathsf{Set}$ is a simplicial set, $|X|$ is the coend of the functor $\Delta^{\mathrm{op}} \times \Delta \to \mathsf{CGSpc}$ that takes $([n],[m])$ to $X_n \times |\Delta^m|$, where X_n has the discrete topology. This is written

$$|X| = \int^{[n]\in\Delta} X_n \times |\Delta^n|\,.$$

We now apply the same idea to the simplicial bar construction:

Definition 5.38. Let P be a poset and let $\mathcal{M}$ be a simplicial model category. The *bar construction* of a functor $F\colon P \to \mathcal{M}$ is the coend

$$\mathrm{Bar}(F) = \int^{[n]\in\Delta} \mathrm{Bar}_n(F) \otimes \Delta^n.$$

Note that there are canonical maps $\mathrm{Bar}_n(F) \otimes \Delta^n \to \mathrm{Bar}_n(F) \to \operatorname{colim} F$ for all $n \geq 0$, which induce a map $\mathrm{Bar}(F) \to \operatorname{colim} F$ by the universal property of the coend.

Remark 5.39. Let $F\colon P \to \mathsf{CGSpc}$ be a functor valued in the category CGSpc of compactly generated spaces. Since CGSpc is a simplicial model category (with either the Hurewicz or Quillen model structures), we can consider the bar construction $\mathrm{Bar}(F)$. However, if one thinks of F as a functor valued in Top, forgetting that $F(p)$ is compactly generated for each $p \in P$, then one could instead consider the bar construction of Definition 2.57; one could also mimic the construction of Definition 5.38, but taking the various limits and colimits in Top rather than in CGSpc. Fortunately, all of these constructions coincide, as we will now explain.

Let $F\colon P \to \mathsf{Top}$ be a functor valued in the category of all topological spaces. Just in this remark, let $\mathrm{Bar}^*(F)$ be the coend

$$\mathrm{Bar}^*(F) = \int^{[n]\in\Delta} \mathrm{Bar}_n(F) \times |\Delta^n|$$

in Top, where $\mathrm{Bar}_\bullet(F)\colon \Delta^{\mathrm{op}} \to \mathsf{Top}$ is defined as in Definition 5.36.

Following Dugger–Isaksen [DI04, Appendix A], we can compute $\mathrm{Bar}^*(F)$ as a sequential colimit of pushouts in Top, as follows. For $k = 0$, we define $\mathrm{Bar}(F)(0) = \coprod_{v \in P} F(p)$, and for $k > 0$ we inductively define $\mathrm{Bar}(F)(k)$ as the pushout

$$\begin{array}{ccc} \coprod_{v_0 < \cdots < v_k} F(v_0) \times |\partial \Delta^k| & \longrightarrow & \mathrm{Bar}(F)(k-1) \\ \big\downarrow & \ulcorner & \Big\downarrow f_{k-1} \\ \coprod_{v_0 < \cdots < v_k} F(v_0) \times |\Delta^k| & \longrightarrow & \mathrm{Bar}(F)(k) \end{array} \tag{5.4}$$

where the top horizontal map is defined using the face maps $\mathrm{Bar}_k(F) \to \mathrm{Bar}_{k-1}(F)$. Then we have an isomorphism $\mathrm{Bar}^*(P) \cong \mathrm{colim}_k \mathrm{Bar}(F)(k)$.

Now, say F is valued in compactly generated spaces. We will use this characterization of $\mathrm{Bar}^*(F)$ to show that it coincides with the bar construction computed in CGSpc. The key fact we need is that, for any diagram in CGSpc, if its (co)limit computed in Top happens to be compactly generated, then this is also the (co)limit in CGSpc. This follows from the existence of the pair of adjunctions relating CGSpc and Top, mentioned right below Definition 2.63: because of these adjunctions, CGSpc is a reflective subcategory of k-spaces, and k-spaces is a coreflective subcategory of Top [Rie16, Definition 4.5.12]. As the disjoint union of compactly generated spaces is compactly generated, and because $|\partial \Delta^n|$ as well as $|\Delta^n|$ are locally compact Hausdorff, it follows from [Str09, Proposition 2.6. and Corollary 2.16] that the spaces on the left hand side of diagram 5.4 and $\mathrm{Bar}(F)(0)$ are compactly generated. Moreover, because the pushout of compactly generated spaces along a closed inclusion is again compactly generated [May99, p.40], it follows that the spaces $\mathrm{Bar}(F)(k)$ are compactly generated. By [Str09, Proposition 2.35], the maps f_k are closed inclusions. Finally, it follows again from [May99, p.40] that $\mathrm{Bar}^*(F) \cong \mathrm{colim}_k \mathrm{Bar}(F)(k)$ is compactly generated. Thus, $\mathrm{Bar}^*(F)$ agrees with the bar construction $\mathrm{Bar}(F)$ computed in CGSpc.

Furthermore, one can use this method for building the bar construction as a sequential colimit of pushouts to check that, given $F\colon P \to \mathsf{Top}$, the bar construction we just discussed is naturally homeomorphic to the bar construction of Definition 2.57, which justifies using the same notation in both places.

Remark 5.40. The coend construction used to define $\mathrm{Bar}(F)$ in Definition 5.38 is an example of the *geometric realization* of a simplicial space. We note that the blowup complex of Definition 5.1 can also be seen as the geometric realization of a simplicial space, the *ordered Čech complex* of [DI04, Section 2.5].

The map $\pi_{\mathrm{Sd}\,N}\colon \mathrm{Bar}(\mathcal{D}_{\mathcal{U}}) \to |\,\mathrm{Sd}\,\mathrm{Nrv}(\mathcal{U})|$ is the geometric realization of a map of simplicial spaces $\mathrm{Bar}_\bullet(\mathcal{D}_{\mathcal{U}}) \to \mathrm{Bar}_\bullet(*^{P_{\mathcal{U}}})$, and the map $\pi_S\colon \mathrm{Bar}(\mathcal{D}_{\mathcal{U}}) \to X$ is also the geometric realization of a map of simplicial spaces $\mathrm{Bar}_\bullet(\mathcal{D}_{\mathcal{U}}) \to X_\bullet$, where $X_n = X$ for all n. Analogously, the maps $\rho_N\colon \mathrm{Blowup}(\mathcal{U}) \to |\,\mathrm{Nrv}(\mathcal{U})|$ and $\rho_S\colon \mathrm{Blowup}(\mathcal{U}) \to X$ can be seen as geometric realizations of maps of simplicial spaces.

Example 5.41. If we leave aside the requirement that we work in a simplicial model category, the PoBar construction from Section 5.3 is a bar construction. In more detail, let K be a simplicial complex, and let $\mathcal{A} = (K_i)_{i \in I}$ be a cover by subcomplexes.

Let $\mathcal{D}_{\mathcal{A}}\colon P_{\mathcal{A}} \to \mathsf{Po}$ be the functor with $\mathcal{D}_{\mathcal{A}}(J) = \mathrm{Pos}(\cap_{i\in J} K_i)$. There is a simplicial object $\mathrm{Bar}_\bullet(\mathcal{D}_{\mathcal{A}})$ in Po defined as in Definition 5.36, and the inclusion $\Delta \subset \mathsf{Po}$ defines a cosimplicial object, i.e., a functor $\Delta \to \mathsf{Po}$. Then

$$\mathrm{PoBar}(\mathcal{A}) = \int^{[n]} \mathrm{Bar}_n(\mathcal{D}_{\mathcal{A}}) \times [n].$$

Given a sufficiently well-behaved diagram $F\colon P \to \mathcal{M}$ in a simplicial model category $\mathcal{M}$, the bar construction of F computes the homotopy colimit of F. This appears as [Rie14, Corollary 5.1.3]. In the proof of Theorem 5.25, we will use two statements closely related to this result. The first one says that the bar construction is homotopical for pointwise cofibrant diagrams (see [Rie14, Corollary 5.2.5]):

Proposition 5.42. *Let $\mathcal{M}$ be a simplicial model category, let P be a poset, and let $F, G\colon P \to \mathcal{M}$ be pointwise cofibrant diagrams. For a natural transformation $F \Rightarrow G$ that is a pointwise weak equivalence, the induced map $\mathrm{Bar}(F) \to \mathrm{Bar}(G)$ is a weak equivalence.*

Recall that in Proposition 2.62 we already saw a similar statement for topological spaces, saying that the bar construction respects pointwise homotopy equivalences, without any pointwise cofibrancy or compactly-generated assumptions. There is an analogous result for weak homotopy equivalences, which follows from work of Dugger–Isaksen [DI04]:

Proposition 5.43. *Let P be a poset, and let $F, G\colon P \to \mathsf{Top}$ be diagrams of topological spaces. For a natural transformation $F \Rightarrow G$ that is a pointwise weak homotopy equivalence, the induced map $\mathrm{Bar}(F) \to \mathrm{Bar}(G)$ is a weak homotopy equivalence.*

Proof. The natural transformation $F \Rightarrow G$ induces a map $\mathrm{Bar}_\bullet(F) \to \mathrm{Bar}_\bullet(G)$ such that $\mathrm{Bar}_n(F) \to \mathrm{Bar}_n(G)$ is a weak homotopy equivalence for all $n \geq 0$. It is straightforward to check that the simplicial bar construction has free degeneracies in the sense of [DI04, Definition A.4], and therefore the induced map $\mathrm{Bar}(F) \to \mathrm{Bar}(G)$ is a weak homotopy equivalence by Remark 5.40 together with [DI04, Corollary A.6]. ☐

We will soon use this last result to prove the "weak" version of Theorem 5.25 2(a). If we did not know this result, and applied only Proposition 5.42, we would need the additional assumption that the intersection A_J is cofibrant in the Quillen model structure on CGSpc for all $J \in \mathrm{Nrv}(\mathcal{A})$.

For the proof of Theorem 5.25, we will also need a second result related to the general fact that the bar construction computes the homotopy colimit. This is similar to [WZŽ99, Lemma 4.5], for example. Recall that $\mathcal{D}_{\mathcal{A}}\colon P_{\mathcal{A}} \to \mathsf{Top}$ denotes the nerve diagram of the cover $\mathcal{A}$ of a topological space X.

Proposition 5.44. *Let X be a compactly generated space, and let $\mathcal{A}$ be a closed cover that is locally finite and locally finite dimensional. If for all $T \in \mathrm{Nrv}(\mathcal{A})$ the latching space $L(T) \subseteq A_T$ is closed and the pair $(A_T, L(T))$ satisfies the homotopy extension property, then the natural map $\mathrm{Bar}(\mathcal{D}_{\mathcal{A}}) \to \mathrm{colim}\, \mathcal{D}_{\mathcal{A}} \cong X$ is a homotopy equivalence.*

Only in this proof, we will make use of model structures on functor categories, in particular, the projective and Reedy model structures, which we do not introduce in this thesis. For the interested reader, we refer to [Dug08, Section 13] and [Hir03, Chapter 15].

Proof. A closed subspace of a compactly generated space is also compactly generated, so $\mathcal{D}_{\mathcal{A}} \colon P_{\mathcal{A}} \to \mathsf{Top}$ takes values in the subcategory CGSpc of compactly generated spaces. Hence, we can take the bar construction of $\mathcal{D}_{\mathcal{A}}$, as discussed in Remark 5.39.

As $\mathcal{A}$ is locally finite dimensional, the poset $P_{\mathcal{A}}$ is an upwards-directed Reedy category, with $\deg(J) = \sup\{(|J'| - |J|) \mid J' \in \operatorname{Nrv}(\mathcal{A}) \text{ with } J \subseteq J'\}$. Working with the Hurewicz model structure on CGSpc, the Reedy model structure on the functor category $\mathsf{Fun}(P_{\mathcal{A}}, \mathsf{CGSpc})$ coincides with the projective model structure as $P_{\mathcal{A}}$ is upwards-directed. This is immediate from the definition of the Reedy model structure; see [Dug08, Proposition 13.12] for a clear discussion of the relationship with the projective model structure.

The condition on the latching spaces implies that all inclusions $L(T) \subseteq A_T$ are Hurewicz cofibrations, so that $\mathcal{D}_{\mathcal{A}}$ is Reedy cofibrant and thus projective cofibrant. As the bar construction $\operatorname{Bar}(\mathcal{D}_{\mathcal{A}})$ computes the homotopy colimit of $\mathcal{D}_{\mathcal{A}}$ [Rie14, Corollary 5.1.3], the natural map $\operatorname{Bar}(\mathcal{D}_{\mathcal{A}}) \to \operatorname{colim} \mathcal{D}_{\mathcal{A}}$ is a homotopy equivalence. As the cover is locally finite and X is compactly generated, it follows from [Str09, Corollary 2.23] that the colimit calculated in CGSpc coincides with the one in Top, and by Remark 2.60 this is naturally homeomorphic to X. □

We are now ready to prove the unified nerve theorem.

Proof of Theorem 5.25. In Section 5.1, we explained how one can prove that the natural maps ρ_S and ρ_N from the blowup complex are equivalences by proving that the natural maps π_S and $\pi_{\operatorname{Sd} N}$ from the bar construction are equivalences (see Diagram 5.1). So, we work with the bar construction in this proof.

The first part of 1(a) follows from work of Dugger and Isaksen [DI04, Theorem 2.1 and Proposition 2.7]; in Section 5.5 we give a short proof using their ideas. The second part of 1(a) is essentially the content of [Hat02, Proposition 4G.2]; note that the author uses the convention that paracompact spaces are assumed to be Hausdorff. Statement 1(b) is the content of Proposition 5.44.

We now prove 2(a). By assumption, the unique natural transformation $\mathcal{D}_{\mathcal{A}} \Rightarrow *^{P_{\mathcal{A}}}$ from the nerve diagram of the cover $\mathcal{A}$ to the constant diagram on the one-point space is a pointwise (weak) homotopy equivalence. The results now follow from Proposition 2.62 and Proposition 5.43, respectively. Alternatively, see Section 5.5 for a short proof of the fact that ρ_N is a weak homotopy equivalence whenever $\mathcal{A}$ is a weakly good cover.

We now prove 2(b). For every compactly generated space Z, there is a natural weak homotopy equivalence $|\operatorname{Sing}(Z)| \to Z$, given by the counit of the adjunction $(|-|, \operatorname{Sing})$ as explained in Section 2.4.2. So, there is a pointwise weak homotopy equivalence

$$|-| \circ \operatorname{Sing} \circ \mathcal{D}_{\mathcal{A}} \Rightarrow \mathcal{D}_{\mathcal{A}},$$

that induces, by Proposition 5.43, a weak homotopy equivalence

$$\operatorname{Bar}(|-| \circ \operatorname{Sing} \circ \mathcal{D}_{\mathcal{A}}) \to \operatorname{Bar}(\mathcal{D}_{\mathcal{A}}). \tag{5.5}$$

We work with the model structure on the category $\mathsf{s}(R\text{-}\mathsf{Mod})$ of simplicial R-modules described in Example 2.78; recall that a map $X \to Y$ of compactly-generated spaces is an R-homology isomorphism if and only if the induced map $\mathsf{R}(\mathrm{Sing}(X)) \to \mathsf{R}(\mathrm{Sing}(Y))$ is a weak equivalence of simplicial R-modules. By our assumption that the natural transformation $\mathcal{D}_{\mathcal{A}} \Rightarrow *^{P_{\mathcal{A}}}$ is a pointwise R-homology isomorphism, the natural transformation $\mathsf{R} \circ \mathrm{Sing} \circ \mathcal{D}_{\mathcal{A}} \Rightarrow \mathsf{R} \circ \mathrm{Sing} \circ *^{P_{\mathcal{A}}}$ is a pointwise weak equivalence of simplicial R-modules. As R preserves cofibrant objects (since it is the left adjoint of a Quillen adjunction), and every simplicial set is cofibrant, both diagrams are pointwise cofibrant. So, by Proposition 5.42, the induced map $\mathrm{Bar}(\mathsf{R} \circ \mathrm{Sing} \circ \mathcal{D}_{\mathcal{A}}) \to \mathrm{Bar}(\mathsf{R} \circ \mathrm{Sing} \circ *^{P_{\mathcal{A}}})$ is a weak equivalence. Furthermore, for any poset P and any functor $F \colon P \to \mathsf{sSet}$, there is a natural isomorphism $\mathrm{Bar}(\mathsf{R} \circ F) \cong \mathsf{R}(\mathrm{Bar}(F))$, using the definition of the tensor structure on $\mathsf{s}(R\text{-}\mathsf{Mod})$ and the fact that R preserves colimits (as it is a left adjoint). So we have a commutative diagram:

$$\begin{array}{ccc}
\mathrm{Bar}(\mathsf{R} \circ \mathrm{Sing} \circ \mathcal{D}_{\mathcal{A}}) & \xrightarrow{\cong} & \mathsf{R}(\mathrm{Bar}(\mathrm{Sing} \circ \mathcal{D}_{\mathcal{A}})) \\
{\scriptstyle\simeq}\big\downarrow & & \big\downarrow \\
\mathrm{Bar}(\mathsf{R} \circ \mathrm{Sing} \circ *^{P_{\mathcal{A}}}) & \xrightarrow{\cong} & \mathsf{R}(\mathrm{Bar}(\mathrm{Sing} \circ *^{P_{\mathcal{A}}}))
\end{array}$$

and hence, the morphism on the right is a weak equivalence by 2-of-3.

For any simplicial set K, the unit map $K \to \mathrm{Sing}(|K|)$ is a natural weak equivalence, as explained in Section 2.4.2. The functor R preserves all weak equivalences, as every simplicial set is cofibrant [Hov99, Lemma 1.1.12 (Ken Brown's lemma)]. So we have a commutative square, in which the indicated maps are weak equivalences:

$$\begin{array}{ccc}
\mathsf{R}(\mathrm{Bar}(\mathrm{Sing} \circ \mathcal{D}_{\mathcal{A}})) & \xrightarrow{\simeq} & \mathsf{R}(\mathrm{Sing}(|\mathrm{Bar}(\mathrm{Sing} \circ \mathcal{D}_{\mathcal{A}})|)) \\
{\scriptstyle\simeq}\big\downarrow & & \big\downarrow \\
\mathsf{R}(\mathrm{Bar}(\mathrm{Sing} \circ *^{P_{\mathcal{A}}})) & \xrightarrow[\simeq]{} & \mathsf{R}(\mathrm{Sing}(|\mathrm{Bar}(\mathrm{Sing} \circ *^{P_{\mathcal{A}}})|))
\end{array}$$

It follows, by 2-of-3, that the fourth map in the square is a weak equivalence as well, and so $|\mathrm{Bar}(\mathrm{Sing} \circ \mathcal{D}_{\mathcal{A}})| \to |\mathrm{Bar}(\mathrm{Sing} \circ *^{P_{\mathcal{A}}})|$ is an R-homology isomorphism.

For any poset P and any functor $F \colon P \to \mathsf{sSet}$, there is a natural isomorphism $\mathrm{Bar}(|-| \circ F) \cong |\mathrm{Bar}(F)|$, again using the definitions and the fact that geometric realization preserves colimits, being a left adjoint. So we have the following commutative diagram:

$$\begin{array}{ccc}
|\mathrm{Bar}(\mathrm{Sing} \circ \mathcal{D}_{\mathcal{A}})| & \xrightarrow{\sim_R} & |\mathrm{Bar}(\mathrm{Sing} \circ *^{P_{\mathcal{A}}})| \\
\big\downarrow{\scriptstyle\cong} & & \big\downarrow{\scriptstyle\cong} \\
\mathrm{Bar}(|-| \circ \mathrm{Sing} \circ \mathcal{D}_{\mathcal{A}}) & \longrightarrow & \mathrm{Bar}(|-| \circ \mathrm{Sing} \circ *^{P_{\mathcal{A}}}) \\
& \searrow & \big\uparrow \\
& & \mathrm{Bar}(\mathcal{D}_{\mathcal{A}})
\end{array}$$

Recall that the map $\mathrm{Bar}(|-| \circ \mathrm{Sing} \circ \mathcal{D}_{\mathcal{A}}) \to \mathrm{Bar}(\mathcal{D}_{\mathcal{A}})$ from line 5.5 is a weak equivalence. Together with the fact that weak homotopy equivalences are also R-homology isomorphisms [Hat02, Proposition 4.21], we get that the canonical map

$$\pi_{\mathrm{Sd}\,N}\colon \mathrm{Bar}(\mathcal{D}_{\mathcal{A}}) \to |\,\mathrm{Sd}\,\mathrm{Nrv}(\mathcal{A})| \cong \mathrm{Bar}(*^{P_{\mathcal{A}}}) \cong \mathrm{Bar}(|-|\circ \mathrm{Sing}\circ *^{P_{\mathcal{A}}})$$

is, once more by 2-of-3, an R-homology isomorphism as well. □

5.5. The Blowup Complex for Open Covers

The parts of 1(a) and 2(a) in Theorem 5.25 that establish weak homotopy equivalences follow from work of Dugger and Isaksen [DI04]. In this section, we adapt their proof strategy to give a more direct proof of the fact that the natural map $\rho_S\colon \mathrm{Blowup}(\mathcal{A}) \to X$ is a weak homotopy equivalence whenever $\mathcal{A}$ is an open cover of X. Moreover, we use the same ideas to give a short proof of the fact that $\rho_N\colon \mathrm{Blowup}(\mathcal{A}) \to |\,\mathrm{Nrv}(\mathcal{A})|$ is a weak homotopy equivalence whenever $\mathcal{A}$ is a weakly good cover.

By the following lemma, a map is a weak homotopy equivalence if it is so locally.

Lemma 5.45 ([Gra75, Lemma 16.24]; [Die08, Theorem 6.7.11])**.** *Let $f\colon Y \to X$ be a continuous map and let $\mathcal{A} = (A_i)_{i\in I}$ be an open cover of X. If for every $\sigma \in \mathrm{Nrv}(\mathcal{A})$ the restricted map $f^{-1}(A_\sigma) \to A_\sigma$ is a weak homotopy equivalence, then so is f.*

In order to apply this lemma to ρ_S, we need to determine the preimages of the finite intersections of cover elements in $\mathcal{A} = (A_i)_{i\in I}$. Recall from Definition 5.1 that the blowup complex is defined as

$$\mathrm{Blowup}(\mathcal{A}) = \left(\bigsqcup_{J\in \mathrm{Nrv}(\mathcal{A})} A_J \times |J| \right) \Big/ \sim$$

and ρ_S is induced by the projections of the products $A_J \times |J|$ onto the first coordinate. For any $\sigma \in \mathrm{Nrv}(\mathcal{A})$ the preimage is the subspace

$$\rho_S^{-1}(A_\sigma) = \left(\bigsqcup_{J\in \mathrm{St}_{\mathrm{Nrv}(\mathcal{A})}(\sigma)} A_J \times |J| \right) \Big/ \sim \,,$$

where $\mathrm{St}_{\mathrm{Nrv}(\mathcal{A})}(\sigma) = \{J \in \mathrm{Nrv}(\mathcal{A}) \mid \sigma \subseteq J\}$ is the star of σ in $\mathrm{Nrv}(\mathcal{A})$. This can be seen as follows: Whenever $A_J \cap A_\sigma = A_{J\cup\sigma}$ is non-empty for some J, the union $J \cup \sigma$ is a simplex in $\mathrm{St}_{\mathrm{Nrv}(\mathcal{A})}(\sigma)$ and so

$$(A_J \cap A_\sigma) \times |J| \subseteq A_{J\cup\sigma} \times |J \cup \sigma|$$

is contained in the right hand side of the equality above. Conversely, $A_J \subseteq A_\sigma$ for every $J \in \mathrm{St}_{\mathrm{Nrv}(\mathcal{A})}(\sigma)$, and so the above equality holds.

Proposition 5.46. *Let $\mathcal{A} = (A_i)_{i\in I}$ be an open cover of the topological space X. Then the natural map $\rho_S\colon \mathrm{Blowup}(\mathcal{A}) \to X$ is a weak homotopy equivalence.*

Proof. By Lemma 5.45 it suffices to prove that for every $\sigma \in \mathrm{Nrv}(\mathcal{A})$ the restricted map $\rho_S^{-1}(A_\sigma) \to A_\sigma$ is a weak homotopy equivalence. We show that this map is in fact a homotopy equivalence.

Choose any point $z \in |\sigma|$ and consider the following subspace

$$A_\sigma \times \{z\} \hookrightarrow A_\sigma \times |\sigma| \hookrightarrow \rho_S^{-1}(A_\sigma).$$

Note that the space $|\mathrm{St}_{\mathrm{Nrv}(\mathcal{A})}(\sigma)|$ is star-shaped with respect to z and hence it deformation retracts onto this point. As $A_\tau \subseteq A_\sigma$ for every $\tau \in \mathrm{St}_{\mathrm{Nrv}(\mathcal{A})}(\sigma)$, this lifts to a deformation retraction of $\rho_S^{-1}(A_\sigma)$ onto $A_\sigma \times \{z\}$. Therefore, we get the following commutative diagram, where the horizontal maps are homotopy equivalences:

$$\begin{array}{ccc} A_\sigma \times \{z\} & \xrightarrow{\simeq} & \rho_S^{-1}(A_\sigma) \\ \downarrow{\scriptstyle \pi_1} & & \downarrow{\scriptstyle \rho_S} \\ A_\sigma & \xrightarrow{=} & A_\sigma \end{array}$$

Obviously, π_1 is a homotopy equivalence and hence so is the map on the right. This proves that $\rho_S\colon \mathrm{Blowup}(\mathcal{A}) \to X$ is weak homotopy equivalence. □

Recall now that ρ_N is induced by the projections of the products $A_J \times |J|$ onto the second coordinate. To apply Lemma 5.45 to ρ_N, we cover $|\mathrm{Nrv}(\mathcal{A})|$ by the open simplex stars $(S_\sigma)_{\sigma \in \mathrm{Nrv}(\mathcal{A})}$, where

$$S_\sigma = \bigcup \{\mathrm{int}\,|J| \mid J \in \mathrm{St}_{\mathrm{Nrv}(\mathcal{A})}(\sigma)\}.$$

Note that this cover is closed under finite intersections. Hence, it suffices to consider for any $\sigma \in \mathrm{Nrv}(\mathcal{A})$ the preimage

$$\rho_N^{-1}(S_\sigma) = \left(\coprod_{J \in \mathrm{St}_{\mathrm{Nrv}(\mathcal{A})}(\sigma)} A_J \times \mathrm{int}\,|J| \right) \;/\; \sim .$$

Proposition 5.47. *Let $\mathcal{A} = (A_i)_{i \in I}$ be a weakly good cover of the topological space X. Then the natural map $\rho_N\colon \mathrm{Blowup}(\mathcal{A}) \to |\mathrm{Nrv}(\mathcal{A})|$ is a weak homotopy equivalence.*

Proof. By Lemma 5.45 it suffices to prove that for every $\sigma \in \mathrm{Nrv}(\mathcal{A})$ the restricted map $\rho_N^{-1}(S_\sigma) \to S_\sigma$ is a weak homotopy equivalence.

Similarly to the the proof of Proposition 5.46, we get the following commutative diagram, where $z \in \mathrm{int}\,|\sigma|$ is any point and the horizontal maps are homotopy equivalences:

$$\begin{array}{ccc} A_\sigma \times \{z\} & \xrightarrow{\simeq} & \rho_N^{-1}(S_\sigma) \\ \downarrow{\scriptstyle \pi_2} & & \downarrow{\scriptstyle \rho_N} \\ \{z\} & \xrightarrow{\simeq} & S_\sigma \end{array}$$

As $\mathcal{A}$ is a weakly good cover, the map π_2 is a weak homotopy equivalence and hence so is the map on the right. This proves that $\rho_N\colon \mathrm{Blowup}(\mathcal{A}) \to |\mathrm{Nrv}(\mathcal{A})|$ is a weak homotopy equivalence. □

A. Nerve Theorem for Closed Convex Covers

We prove Theorem 2.5 using some key ideas that already appeared in the author's master's thesis [Rol20], where the theorem was proven under the additional assumption on the cover elements to be compact. While the key ideas are essentially the same, the exposition below is more streamlined. We include the material here for the sake of readability and completeness.

Let $\mathcal{C} = (C_i)_{i\in[n]}$ be a collection of closed convex subsets of $\mathbb{R}^d$, and let X be their union. We briefly recall the construction of the continuous map $\Gamma\colon |\operatorname{Sd}\operatorname{Nrv}(\mathcal{C})| \to X$ that is claimed to establish the homotopy equivalence in Theorem 2.5: Each vertex $J \in \operatorname{Sd}\operatorname{Nrv}(\mathcal{C})$ represents a simplex in the nerve $\operatorname{Nrv}(\mathcal{C})$, and hence we can choose a point p_J from the non-empty intersection $C_J = \bigcap_{j\in J} C_j$. By convexity of the cover elements in $\mathcal{C}$, this choice extends uniquely to a map $\Gamma\colon |\operatorname{Sd}\operatorname{Nrv}(\mathcal{C})| \to X$ that is affine linear on each simplex of the barycentric subdivision; see Fig. 6 for an illustration.

We construct a homotopy inverse Ψ to Γ. For this construction, we work with an open cover and a subordinate partition of unity, as in the familiar proof of the nerve theorem for open covers [Hat02, Proposition 4G.2]. To this end, we thicken the subsets C_i slightly so that the nerve remains unchanged. If the C_i are compact, it is possible to choose an ε such that the open ε-neighborhoods of the C_i have this desired property. More generally, we can choose such neighborhoods according to the following lemma.

Lemma A.1. *Let $\mathcal{C} = (C_i)_{i\in[n]}$ be a collection of closed and convex subsets of $\mathbb{R}^d$. Then there exists a collection of open sets $\mathcal{G} = (U_i \supseteq C_i)_{i\in[n]}$ satisfying* $\operatorname{Nrv}(\mathcal{C}) = \operatorname{Nrv}(\mathcal{G}_\varepsilon)$.

Proof. Since $\mathbb{R}^d$ is a normal space, every disjoint pair of closed sets C_i, C_j admits disjoint open neighborhoods $V_{i,j} \supseteq C_i$, $V_{j,i} \supseteq C_j$. Taking the finite intersection $U_i = \bigcap_{j : C_i \cap C_j = \emptyset} V_{i,j}$ for every i yields the desired open cover $(U_i)_{i\in[n]}$. □

We choose $\mathcal{G}$ according to Lemma A.1. For each $i \in [n]$, there exists a Urysohn function $\varphi_i\colon \mathbb{R}^d \to [0,1]$ that takes on the value 0 outside of U_i and the value 1 on C_i. For example, we may take

$$x \mapsto \frac{d(x, \mathbb{R}^d \setminus U_i)}{d(x, C_i) + d(x, \mathbb{R}^d \setminus U_i)}.$$

Normalizing these functions φ_i yields a partition of unity on X subordinate to the cover $(U_i \cap X)_{i\in[n]}$ of X: $\psi_i = \varphi_i / \sum_{j=0}^n \varphi_j \colon X \to [0,1]$. We define the map $\Phi\colon X \to |\operatorname{Nrv}(\mathcal{C})|$ in barycentric coordinates for $|\operatorname{Nrv}(\mathcal{C})|$ as

$$\Phi\colon x \mapsto \sum_{i=0}^{n} \psi_i(x) \cdot |v_i|, \tag{A.1}$$

where $v_i = \{i\}$ is the vertex in $\operatorname{Nrv}(\mathcal{C})$ corresponding to i and $|v_i|$ is the corresponding point in the geometric realization. The map $\Psi\colon X \to |\operatorname{Sd}\operatorname{Nrv}(\mathcal{C})|$ is then given as the composite $\alpha_{\operatorname{Nrv}(\mathcal{C})} \circ \Phi$, where $\alpha_{\operatorname{Nrv}(\mathcal{C})}\colon |\operatorname{Nrv}(\mathcal{C})| \to |\operatorname{Sd}\operatorname{Nrv}(\mathcal{C})|$ is the standard homeomorphism from the nerve to its barycentric subdivision.

In order to show that Ψ is a homotopy inverse to Γ, we analyze more closely how these maps are related combinatorially. To this end, we use the following construction.

Definition A.2. For every vertex v of a simplicial complex K, define the *closed barycentric star* as the subspace

$$\operatorname{bst} v = |\operatorname{Cl}\operatorname{St}_{\operatorname{Sd} K} v| \subseteq |\operatorname{Sd} K|,$$

where $\operatorname{Cl}\operatorname{St}_{\operatorname{Sd} K} v = \{\sigma \in \operatorname{Sd} K \mid \sigma \cup \{v\} \in \operatorname{Sd} K\}$ is the closure of the star of v in the barycentric subdivision of K.

We now state two lemmas about the closed barycentric stars, deferring the proofs to the end of this section.

Lemma A.3. *Let K be a simplicial complex and let $\sigma \in K$ be a simplex. Then the intersection $\bigcap_{v\in\sigma} \operatorname{bst} v$ is contractible. In particular, the collection of closed barycentric stars forms a good cover of $|K|$.*

It is not hard to see that the nerve of this cover is isomorphic to K. The following statement describes the closed barycentric stars in terms of barycentric coordinates.

Lemma A.4. *Let K be a simplicial complex and let v be a vertex of K. The closed barycentric star $\operatorname{bst} v$ consists of all points $x \in |K|$ that satisfy*

$$b_v(x) \geq b_w(x) \quad \text{for all } w \in \operatorname{Vert} K, \tag{A.2}$$

where b_v denotes the barycentric coordinate with respect to the vertex v.

The following propositions use the language established in Definition 1.11.

Proposition A.5. *The pair of maps $(\Psi, \operatorname{id}_{[n]})$ constitutes a morphism of covered spaces*

$$(X, \mathcal{C} = (C_i)_{i\in[n]}) \to (|\operatorname{Sd}\operatorname{Nrv}(\mathcal{C})|, (\operatorname{bst} v_i)_{i\in[n]}).$$

Proof. Recall that $\Psi = \alpha_{\operatorname{Nrv}(\mathcal{C})} \circ \Phi$, where $\alpha_{\operatorname{Nrv}(\mathcal{C})}$ is the isomorphism $|\operatorname{Nrv}(\mathcal{C})| \cong |\operatorname{Sd}\operatorname{Nrv}(\mathcal{C})|$ and $\Phi\colon X \to |\operatorname{Nrv}(\mathcal{C})|$, $x \mapsto \sum_{i=0}^{n} \psi_i(x) \cdot |v_i|$. Note that if $x \in C_i$, then $\phi_i(x) = 1$ and thus $\psi_i(x)$ is maximal among the $\psi_j(x)$. Hence, by Lemma A.4 we know that $\Psi(x) \in \operatorname{bst}(v_i)$ and the claim follows. □

Proposition A.6. *The pair of maps $(\Gamma, \operatorname{id}_{[n]})$ constitutes a morphism of covered spaces*

$$(|\operatorname{Sd}\operatorname{Nrv}(\mathcal{C})|, (\operatorname{bst} v_i)_{i\in[n]}) \to (X, \mathcal{C} = (C_i)_{i\in[n]}).$$

Proof. By definition, the map Γ sends the vertices of a geometric simplex σ in $\operatorname{bst} v_i$ to C_i. As the cover element C_i is convex and Γ is affine linear on σ, it follows that $\Gamma(\sigma)$ is also contained in C_i. This shows $\Gamma(\operatorname{bst} v_i) \subseteq C_i$, proving the claim. □

We will now show that Ψ is a homotopy inverse to Γ, which implies $|\operatorname{Nrv}(\mathcal{C})| \cong |\operatorname{Sd}\operatorname{Nrv}(\mathcal{C})| \simeq X$.

Proof of Theorem 2.5. It follows from Proposition A.6 and Proposition A.5 that the pair of maps $(\Gamma \circ \Psi, \mathrm{id}_{[n]})$ constitutes a morphism of covered spaces. Hence, $\Gamma \circ \Psi$ is carried by the identity on $\mathcal{C}$ and thus it is homotopic to the identity id_X by a straight line homotopy: for every $x \in C_i$, we have $\Gamma \circ \Psi(x) \in C_i$, and since the C_i are convex, the line segment joining x and $\Gamma \circ \Psi(x)$ lies in C_i. Similarly, the pair of maps $(\Psi \circ \Gamma, \mathrm{id}_{[n]})$ constitutes a morphism of covered spaces. That the composition $\Psi \circ \Gamma$ is homotopic to $\mathrm{id}_{|\operatorname{Sd}\operatorname{Nrv}(\mathcal{C})|}$ now follows from Lemma A.3 and the following Proposition A.7. ☐

Recall that any two maps into a contractible space are homotopic (to a constant map). The following statement generalizes this fact to good covers, where contractibility is only guaranteed locally.

Proposition A.7. *Let K be a finite simplicial complex and let Y be a topological space. Assume we have two morphisms of covered spaces*

$$(f, \varphi), (g, \varphi)\colon (|K|, (|L_i|)_{i \in [n]}) \to (Y, (V_j)_{j \in J}),$$

with the same map of index sets $\varphi\colon [n] \to J$, where $(|L_i|)_{i \in [n]}$ is a cover by subcomplexes and $(V_j)_{j \in J}$ is a good cover. Then f is homotopic to g.

Proof. Let $I = [0, 1]$ denote the unit interval. We inductively construct homotopies $H^m\colon |\operatorname{sk}_m K| \times I \to Y$ between $f|_{|\operatorname{sk}_m K|}$ and $g|_{|\operatorname{sk}_m K|}$ such that H^m is carried by the map of indexed covers $\varphi\colon (|\operatorname{sk}_m L_i| \times I)_{i \in [n]} \to (V_j)_{j \in J}$ induced by the given map of index sets $\varphi\colon [n] \to J$. If $m = \dim K$ is the dimension of the simplicial complex, the map $H = H^m$ is the desired homotopy between f and g.

To establish the base case $m = 0$, let p be any vertex of K and let $i_0, \dots, i_k \in [n]$ be those indices i with $|p| \in |L_i|$. By the assumption that f and g are carried by φ, we know that both $f(|p|)$ and $g(|p|)$ are contained in $S := \bigcap_{l=0}^{k} V_{\varphi(i_l)}$, which is contractible by assumption, and thus we can choose a path in S that connects these two points. This defines the desired homotopy H^0. To see that the map H^0 is carried by P^0, let $(|p|, t) \in |p| \times I$ be a point. If $(|p|, t) \in |L_i| \times I$, then $i = i_l$ is one of the indices above. Thus, by construction, $H^0((|p|, t)) \in S \subseteq V_{\varphi(i_l)} = V_{\varphi(i)}$, and the claim is proven.

For the induction step from $(m-1)$ to m, let H^{m-1} satisfy the induction hypothesis. Let σ be an m-simplex in $\operatorname{sk}_m K$. Furthermore, let $i_0, \dots, i_k \in [n]$ be those indices i with $\sigma \in L_i$. By the induction hypothesis, we have

$$H^{m-1}(|\partial\sigma| \times I) \subseteq W := \bigcap_{l=0}^{k} V_{\varphi(i_l)}.$$

By the assumption that (V_j) is good, the space W is contractible, and so we can extend the homotopy $H^{m-1}|_{|\partial\sigma| \times I}$ to a homotopy $H^m|_{|\sigma| \times I}$ from $f|_{|\sigma|}$ to $g|_{|\sigma|}$:

$$\begin{array}{ccc} (|\partial\sigma| \times I) \cup (|\sigma| \times \{0,1\}) \cong S^m & \xrightarrow{(H^{m-1}, (f,g))} & W \subseteq Y \\ \big\downarrow & \nearrow_{H^m|_{|\sigma| \times I}} & \\ |\sigma| \times I \cong B^{m+1} & & \end{array}$$

Because the m-simplex σ was arbitrary, we can extend $H^{m-1}\colon |\operatorname{sk}_{m-1} K| \times I \to Y$ to $H^m\colon |\operatorname{sk}_m K| \times I \to Y$.

By construction, this map is carried by $\varphi\colon (|\operatorname{sk}_m L_i| \times I)_{i\in[n]} \to (V_j)_{j\in J}$. To see this, we verify that for any i, every point $(x,t) \in |\operatorname{sk}_m L_i| \times I$ is mapped to $H^m(x,t) \in V_{\varphi(i)}$. By induction, this is true whenever $x \in |\operatorname{sk}_{m-1} L_i|$, so it remains to show the claim for x in the interior of some m-simplex $\sigma \in L_i$. Now $i = i_l$ is one of the indices above, and by construction of H^m, we have

$$H^m(x,t) \in H^m(|\sigma| \times I) \subseteq W \subseteq V_{\varphi(i_l)} = V_{\varphi(i)},$$

proving the claim. □

We now prove Lemma A.3 and Lemma A.4. To this end, we use some auxiliary lemmas about geometric simplicial complexes.

Lemma A.8. *Let $\sigma = \{v_0, \dots, v_k\} \in K$ be a simplex and consider the subcomplex $L = \{\tau_0 \subset \cdots \subset \tau_m \mid \sigma \subseteq \tau_0\} \subseteq \operatorname{Sd} K$. Then $\bigcap_{i=0}^k \operatorname{bst} v_i = |L|$.*

Proof of Lemma A.8. First, let $\phi = (\tau_0 \subset \cdots \subset \tau_m) \in L$ be a simplex. By definition, ϕ is contained in the simplex $\sigma \subseteq \tau_0 \subset \cdots \subset \tau_m$ of $\operatorname{Sd} K$. Thus, the realization of ϕ is contained in $\operatorname{bst} v_i$ for all i, and so we have $|L| \subseteq \bigcap_{i=0}^k \operatorname{bst} v_i$.

Now, let $|\phi = (\tau_0 \subset \cdots \subset \tau_m)| \subseteq \bigcap_{i=0}^k \operatorname{bst} v_i$. Since for all i we have $|\tau_0 \subset \cdots \subset \tau_m| \subseteq \operatorname{bst} v_i$, we know that $v_i \in \tau_0$. Thus, the simplex σ is also contained in τ_0. Therefore, $\phi \in L$ and so we have $\bigcap_{i=0}^k \operatorname{bst} v_i \subseteq |L|$. □

Proof of Lemma A.3. By Lemma A.8, every (geometric) simplex in $\bigcap_{v\in\sigma} \operatorname{bst} v \subseteq |\operatorname{Sd} K|$ has a coface in this intersection with $z(|\sigma|)$ as a vertex, where $z(|\sigma|)$ is the barycenter of $|\sigma|$. Thus, $\bigcap_{v\in\sigma} \operatorname{bst} v$ is star-shaped with respect to $z(|\sigma|)$ and hence contractible. □

The following two lemmas are straightforward calculations (compare [ES52, p.62]).

Lemma A.9. *Let K be a simplicial complex and let $x \in |K|$. Write x in barycentric coordinates of K as*

$$x = \sum_{j=0}^{m} \nu_j \cdot |w_j|$$

with $w_i \in \operatorname{Vert} K$, $\nu_i > 0$ and $\sum\limits_{j=0}^{m} \nu_j = 1$ as well as $\nu_0 \geq \nu_1 \geq \cdots \geq \nu_m$. Then, using the (geometric) simplices

$$|\tau_i| = \operatorname{conv}\{|w_0|, \dots, |w_i|\} \quad \text{for all } i \in \{0, \dots, m\} \tag{A.3}$$

in the realization $|K|$ and by writing $z(|\tau_i|)$ for the barycenter of $|\tau_i|$, we have $x \in \operatorname{conv}\{z(|\tau_0|), \dots, z(|\tau_m|)\}$. Specifically, writing x in barycentric coordinates of $\operatorname{Sd} K$ as $x = \sum_{j=0}^m \mu_j z(|\tau_j|)$, we have

$$\begin{aligned} \mu_i &= (i+1)\Big(\nu_i(x) - \nu_{i+1}(x)\Big) \text{ for } i = 0, \dots, m-1 \\ \mu_m &= (m+1)\nu_m(x). \end{aligned}$$

Lemma A.10. *Let $x \in |\operatorname{Sd} K|$, written in barycentric coordinates as $x = \sum_{j=0}^{m} \mu_j z(|\tau_j|)$ for some flag of simplices $\tau_0 \subset \cdots \subset \tau_m$ in K, where*

$$|\tau_i| = \operatorname{conv}\{|w_0|, \ldots, |w_i|\} \quad \text{for all } i \in \{1, \ldots, m\}$$

and $w_i \in \operatorname{Vert} K$. Then we have $x \in |\tau_m| = \operatorname{conv}\{|w_0|, \ldots, |w_m|\}$. Specifically, the barycentric coordinates ν_i of x in K with respect to $|w_0|, \ldots, |w_m|$ take the form

$$\nu_i = \sum_{j=i}^{m} \frac{1}{j+1} \mu_j. \tag{A.4}$$

Proof of Lemma A.4. Let $x \in |K|$ be a point satisfying Eq. (A.2). It suffices to show that x is contained in a simplex of $|\operatorname{Sd} K|$ having $|v|$ as a vertex. Let $v_0, \ldots, v_m$ be the vertices in K with $b_{v_i}(x) > 0$ in descending order of barycentric coordinates. By Eq. (A.2) we may choose $v_0 = v$. Now, by Lemma A.9, we know that the point x is contained in $\operatorname{conv}\{|v| = z(|\tau_0|), \ldots, z(|\tau_m|)\}$ for the simplices $|\tau_i| \subseteq |K|$ specified as in Eq. (A.3). Hence, by definition the point x is contained in $\operatorname{bst} v$.

Conversely, let $x \in \operatorname{bst} v$ for some vertex $v \in \operatorname{Vert} K$. Then there exists a simplex $\tau \in \operatorname{Sd} K$ with v as a vertex such that $x \in |\tau|$ and that τ corresponds to a flag $v = \tau_0 \subset \cdots \subset \tau_m$ of simplices in K. From Lemma A.10, or more specifically Eq. (A.4), we deduce that the barycentric coordinate $\nu_0 = b_v(x)$ of x in K with respect to v is maximal. □

References

[Ada13] Michał Adamaszek. “Clique complexes and graph powers”. In: *Israel J. Math.* 196.1 (2013). DOI: 10.1007/s11856-012-0166-1.

[AA17] Michał Adamaszek and Henry Adams. “The Vietoris–Rips complexes of a circle”. In: *Pacific Journal of Mathematics* 290.1 (July 2017). DOI: 10.2140/pjm.2017.290.1.

[Ada+20] Michał Adamaszek, Henry Adams, Ellen Gasparovic, Maria Gommel, Emilie Purvine, Radmila Sazdanovic, Bei Wang, Yusu Wang, and Lori Ziegelmeier. “On homotopy types of Vietoris–Rips complexes of metric gluings”. In: *J. Appl. Comput. Topol.* 4.3 (2020). DOI: 10.1007/s41468-020-00054-y.

[Ale28] Paul Alexandroff. “Über den allgemeinen Dimensionsbegriff und seine Beziehungen zur elementaren geometrischen Anschauung”. In: *Math. Ann.* 98.1 (1928). DOI: 10.1007/BF01451612.

[Ame+02] Nina Amenta, Sunghee Choi, Tamal K. Dey, and Naveen Leekha. “A simple algorithm for homeomorphic surface reconstruction”. In: *Internat. J. Comput. Geom. Appl.* 12.1-2 (2002). 16th Annual Symposium on Computational Geometry (Kowloon, 2000). DOI: 10.1142/S0218195902000773.

[ACK01] Nina Amenta, Sunghee Choi, and Ravi Krishna Kolluri. “The power crust, unions of balls, and the medial axis transform”. In: *Comput. Geom.* 19.2-3 (2001). Combinatorial curves and surfaces. DOI: 10.1016/S0925-7721(01)00017-7.

[Att+23] Dominique Attali, Hana Dal Poz Kouřimská, Christopher Fillmore, Ishika Ghosh, André Lieutier, Elizabeth Stephenson, and Mathijs Wintraecken. *Tight bounds for the learning of homotopy à la Niyogi, Smale, and Weinberger for subsets of Euclidean spaces and of Riemannian manifolds.* 2023. arXiv: 2206.10485.

[AL22] Dominique Attali and André Lieutier. “Delaunay-like triangulation of smooth orientable submanifolds by ℓ_1-norm minimization”. In: *38th International Symposium on Computational Geometry.* Vol. 224. LIPIcs. Leibniz Int. Proc. Inform. Schloss Dagstuhl. Leibniz-Zent. Inf., Wadern, 2022. DOI: 10.4230/lipics.socg.2022.8.

[ALS13] Dominique Attali, André Lieutier, and David Salinas. “Vietoris-Rips complexes also provide topologically correct reconstructions of sampled shapes”. In: *Comput. Geom.* 46.4 (2013). DOI: 10.1016/j.comgeo.2012.02.009.

[ALS19] Dominique Attali, André Lieutier, and David Salinas. “When convexity helps collapsing complexes”. In: *35th International Symposium on Computational Geometry.* Vol. 129. LIPIcs. Leibniz Int. Proc. Inform. Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern, 2019. DOI: 10.4230/LIPIcs.SoCG.2019.11.

[Bar11] Jonathan Ariel Barmak. “On Quillen’s Theorem A for posets”. In: *J. Combin. Theory Ser. A* 118.8 (2011). DOI: 10.1016/j.jcta.2011.06.008.

[BM12] Jonathan Ariel Barmak and Elias Gabriel Minian. “Strong homotopy types, nerves and collapses”. In: *Discrete Comput. Geom.* 47.2 (2012). DOI: 10.1007/s00454-011-9357-5.

[Bar02] Michael Barr. *Acyclic models.* Vol. 17. CRM Monograph Series. American Mathematical Society, Providence, RI, 2002. DOI: 10.1090/crmm/017.

[Bat02] Ekkehard Batzies. *Discrete Morse Theory for Cellular Resolutions.* Philipps-Universität Marburg. 2002. DOI: 10.17192/z2002.0115.

[Bau+20] U. Bauer, H. Edelsbrunner, G. Jabłoński, and M. Mrozek. "Čech–Delaunay gradient flow and homology inference for self-maps". In: *Journal of Applied and Computational Topology* 4 (Dec. 2020). DOI: 10.1007/s41468-020-00058-8.

[Bau21] Ulrich Bauer. "Ripser: efficient computation of Vietoris-Rips persistence barcodes". In: *J. Appl. Comput. Topol.* 5.3 (2021). Software available at https://github.com/Ripser/ripser. DOI: 10.1007/s41468-021-00071-5.

[BE17] Ulrich Bauer and Herbert Edelsbrunner. "The Morse theory of Čech and Delaunay complexes". In: *Trans. Amer. Math. Soc.* 369.5 (2017). DOI: 10.1090/tran/6991.

[Bau+23a] Ulrich Bauer, David Hien, Oliver Junge, and Konstantin Mischaikow. *Cycling Signatures: Identifying Oscillations from Time Series using Algebraic Topology*. 2023. arXiv: 2312.04734.

[Bau+23b] Ulrich Bauer, Michael Kerber, Fabian Roll, and Alexander Rolle. "A unified view on the functorial nerve theorem and its variations". In: *Expo. Math.* 41.4 (2023). DOI: 10.1016/j.exmath.2023.04.005.

[BLL23] Ulrich Bauer, Fabian Lenzen, and Michael Lesnick. "Efficient two-parameter persistence computation via cohomology". In: *39th International Symposium on Computational Geometry*. Vol. 258. LIPIcs. Leibniz Int. Proc. Inf. Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern, 2023. DOI: 10.4230/lipics.socg.2023.15.

[BL20] Ulrich Bauer and Michael Lesnick. "Persistence diagrams as diagrams: a categorification of the stability theorem". In: *Topological data analysis—the Abel Symposium 2018*. Vol. 15. Abel Symp. Springer, Cham, 2020. DOI: 10.1007/978-3-030-43408-3_3.

[BR22] Ulrich Bauer and Fabian Roll. "Gromov Hyperbolicity, Geodesic Defect, and Apparent Pairs in Vietoris-Rips Filtrations". In: *38th International Symposium on Computational Geometry (SoCG 2022)*. Vol. 224. Leibniz International Proceedings in Informatics (LIPIcs). 2022. DOI: 10.4230/LIPIcs.SoCG.2022.15.

[BR24] Ulrich Bauer and Fabian Roll. "Wrapping Cycles in Delaunay Complexes: Bridging Persistent Homology and Discrete Morse Theory". In: *40th International Symposium on Computational Geometry (SoCG 2024)*. Vol. 293. Leibniz International Proceedings in Informatics (LIPIcs). 2024. DOI: 10.4230/LIPIcs.SoCG.2024.15.

[BS23] Ulrich Bauer and Luis Scoccola. *Generic Multi-Parameter Persistence Modules are Nearly Indecomposable*. 2023. arXiv: 2211.15306.

[Ben+07] Paul Bendich, David Cohen-Steiner, Herbert Edelsbrunner, John Harer, and Dmitriy Morozov. "Inferring Local Homology from Sampled Stratified Spaces". In: *48th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2007), Providence, Proceedings*. IEEE Computer Society, 2007. DOI: 10.1109/FOCS.2007.45.

[Bjö81] Anders Björner. "Homotopy type of posets and lattice complementation". In: *J. Combin. Theory Ser. A* 30.1 (1981). DOI: 10.1016/0097-3165(81)90042-X.

[Bjö03] Anders Björner. "Nerves, fibers and homotopy groups". In: *J. Combin. Theory Ser. A* 102.1 (2003). DOI: 10.1016/S0097-3165(03)00015-3.

[Ble+21] Michael Bleher, Lukas Hahn, Juan Angel Patino-Galindo, Mathieu Carriere, Ulrich Bauer, Raul Rabadan, and Andreas Ott. *Topology Identifies Emerging Adaptive Mutations in SARS-CoV-2*. Preprint. 2021. arXiv: 2106.07292.

[BL22] Andrew J. Blumberg and Michael Lesnick. "Stability of 2-Parameter Persistent Homology". In: *Found. Comp. Math.* (2022). DOI: 10.1007/s10208-022-09576-6.

[BL17] Andrew J. Blumberg and Michael Lesnick. *Universality of the Homotopy Interleaving Distance*. 2017. arXiv: 1705.01690.

[BCR98] Jacek Bochnak, Michel Coste, and Marie-Françoise Roy. *Real algebraic geometry*. Vol. 36. Results in Mathematics and Related Areas (3). Translated from the 1987 French original, Revised by the authors. Springer, Berlin, 1998. DOI: 10.1007/978-3-662-03718-8.

[BCY18] Jean-Daniel Boissonnat, Frédéric Chazal, and Mariette Yvinec. *Geometric and topological inference*. Cambridge Texts in Applied Mathematics. Cambridge University Press, Cambridge, 2018. DOI: 10.1017/9781108297806.

[BG14] Jean-Daniel Boissonnat and Arijit Ghosh. "Manifold reconstruction using tangential Delaunay complexes". In: *Discrete Comput. Geom.* 51.1 (2014). DOI: 10.1007/s00454-013-9557-2.

[BP20] Jean-Daniel Boissonnat and Siddharth Pritam. "Edge collapse and persistence of flag complexes". In: *36th International Symposium on Computational Geometry*. Vol. 164. LIPIcs. Leibniz Int. Proc. Inform. Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern, 2020. DOI: 10.4230/LIPIcs.SoCG.2020.19.

[BS11] M. Bonk and O. Schramm. "Embeddings of Gromov hyperbolic spaces". In: *Selected works of Oded Schramm. Volume 1, 2*. Sel. Works Probab. Stat. With a correction by Bonk. Springer, NY, 2011. DOI: 10.1007/978-1-4419-9675-6_10.

[BS73] A. Borel and J.-P. Serre. "Corners and arithmetic groups". In: *Comment. Math. Helv.* 48 (1973). DOI: 10.1007/BF02566134.

[Bor48] Karol Borsuk. "On the imbedding of systems of compacta in simplicial complexes". eng. In: *Fund. Mathematicae* 35.1 (1948). URL: http://eudml.org/doc/213158.

[BC20] Magnus Bakke Botnan and William Crawley-Boevey. "Decomposition of persistence modules". In: *Proc. Amer. Math. Soc.* 148.11 (2020). DOI: 10.1090/proc/14790.

[BOO22] Magnus Bakke Botnan, Steffen Oppermann, and Steve Oudot. *Signed Barcodes for Multi-Parameter Persistence via Rank Decompositions and Rank-Exact Resolutions*. 2022. arXiv: 2107.06800.

[BS15] Magnus Bakke Botnan and Gard Spreemann. "Approximating persistent homology in Euclidean space through collapses". In: *Applied Algebra in Engineering, Communication and Computing* 26.1-2 (2015). DOI: 10.1007/s00200-014-0247-y.

[BT82] Raoul Bott and Loring W. Tu. *Differential forms in algebraic topology*. Vol. 82. Grad. Texts in Mathematics. Springer, NY-Berlin, 1982. ISBN: 0-387-90613-4.

[BH99] Martin R. Bridson and André Haefliger. *Metric spaces of non-positive curvature*. Vol. 319. Fundamental Principles of Mathematical Sciences. Springer, Berlin, 1999. DOI: 10.1007/978-3-662-12494-9.

[Bro94] Kenneth S. Brown. *Cohomology of groups*. Vol. 87. Grad. Texts in Mathematics. Corr. reprint of the 1982 original. Springer, NY, 1994. ISBN: 0-387-90688-6.

[Bro92] Kenneth S. Brown. "The geometry of rewriting systems: a proof of the Anick-Groves-Squier theorem". In: *Algorithms and classification in combinatorial group theory (Berkeley, CA, 1989)*. Vol. 23. Math. Sci. Res. Inst. Publ. Springer, NY, 1992. DOI: 10.1007/978-1-4613-9730-4_6.

[BDK23] Mickaël Buchet, Bianca B. Dornelas, and Michael Kerber. "Sparse higher order Čech filtrations". In: *39th International Symposium on Computational Geometry*. Vol. 258. LIPIcs. Leibniz Int. Proc. Inform. Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern, 2023. DOI: 10.4230/lipics.socg.2023.20.

[BBI01] Dmitri Burago, Yuri Burago, and Sergei Ivanov. *A course in metric geometry*. Vol. 33. Graduate Studies in Mathematics. American Mathematical Society, Providence, RI, 2001. DOI: 10.1090/gsm/033.

[BP87] Herbert Busemann and B. B. Phadke. *Spaces with distinguished geodesics*. Vol. 108. Monographs and Textbooks in Pure and Applied Mathematics. Marcel Dekker, Inc., NY, 1987. ISBN: 0-8247-7545-7.

[Car09] Gunnar Carlsson. "Topology and data". In: *Bull. Amer. Math. Soc. (N.S.)* 46.2 (2009). DOI: 10.1090/S0273-0979-09-01249-X.

[Car+08] Gunnar Carlsson, Tigran Ishkhanov, Vin de Silva, and Afra Zomorodian. "On the local behavior of spaces of natural images". In: *Int. J. Comput. Vis.* 76.1 (2008). DOI: 10.1007/s11263-007-0056-x.

[CS18] Nicholas J. Cavanna and Donald R. Sheehy. *The Generalized Persistent Nerve Theorem*. 2018. arXiv: 1807.07920.

[Cha+23] Wojciech Chacholski, Andrea Guidolin, Isaac Ren, Martina Scolamiero, and Francesca Tombari. *Koszul complexes and relative homological algebra of functors over posets*. 2023. arXiv: 2209.05923.

[CCR13] Joseph Minhow Chan, Gunnar Carlsson, and Raul Rabadan. "Topology of viral evolution". In: *Proceedings of the National Academy of Sciences* 110.46 (2013). DOI: 10.1073/pnas.1313480110.

[Cha+09] Frédéric Chazal, David Cohen-Steiner, Marc Glisse, Leonidas J. Guibas, and Steve Y. Oudot. "Proximity of Persistence Modules and Their Diagrams". In: *Proceedings of the Twenty-Fifth Annual Symposium on Computational Geometry*. SCG '09. NY: Ass. for Computing Machinery, 2009. DOI: 10.1145/1542362.1542407.

[CO08] Frédéric Chazal and Steve Oudot. "Towards Persistence-Based Reconstruction in Euclidean Spaces". In: *Proceedings of the Annual Symposium on Computational Geometry* (Jan. 2008). DOI: 10.1145/1377676.1377719.

[Cha+16] Frédéric Chazal, Vin de Silva, Marc Glisse, and Steve Oudot. *The structure and stability of persistence modules*. SpringerBriefs in Mathematics. Springer, Cham, 2016. DOI: 10.1007/978-3-319-42545-0.

[CSO14] Frédéric Chazal, Vin de Silva, and Steve Oudot. "Persistence stability for geometric complexes". In: *Geom. Dedicata* 173 (2014). DOI: 10.1007/s10711-013-9937-z.

[Che+00] Siu-Wing Cheng, Tamal K. Dey, Herbert Edelsbrunner, Michael A. Facello, and Shang-Hua Teng. "Sliver exudation". In: *J. ACM* 47.5 (2000). DOI: 10.1145/355483.355487.

[CDR05] Siu-Wing Cheng, Tamal K. Dey, and Edgar A. Ramos. "Manifold reconstruction from point samples". In: *Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms*. ACM, NY, 2005. DOI: 10.5555/1070432.1070579.

[Coh73] Marshall M. Cohen. *A course in simple-homotopy theory*. Vol. 10. Graduate Texts in Mathematics. Springer, NY-Berlin, 1973.

[CEH07] David Cohen-Steiner, Herbert Edelsbrunner, and John Harer. "Stability of persistence diagrams". In: *Discrete Comput. Geom.* 37.1 (2007). DOI: 10.1007/s00454-006-1276-5.

[CEM06] David Cohen-Steiner, Herbert Edelsbrunner, and Dmitriy Morozov. "Vines and vineyards by updating persistence in linear time". In: *Computational geometry (SCG'06)*. ACM, NY, 2006. DOI: 10.1145/1137856.1137877.

[CLV23] David Cohen-Steiner, André Lieutier, and Julien Vuillamy. "Delaunay and regular triangulations as lexicographic optimal chains". In: *Discrete Comput. Geom.* 70.1 (2023). DOI: 10.1007/s00454-023-00485-1.

[CLV19] David Cohen-Steiner, André Lieutier, and Julien Vuillamy. *Lexicographic optimal chains and manifold triangulations*. Research Report. INRIA Sophia-Antipolis Méditerranée, Dec. 2019. URL: https://hal.science/hal-02391190.

[CLV22] David Cohen-Steiner, André Lieutier, and Julien Vuillamy. "Lexicographic optimal homologous chains and applications to point cloud triangulations". In: *Discrete Comput. Geom.* 68.4 (2022). Special Issue: 36th Annual Symposium on Computational Geometry (SoCG 2020). DOI: 10.1007/s00454-022-00432-6.

[Dey07] Tamal K. Dey. *Curve and surface reconstruction: algorithms with mathematical analysis*. Vol. 23. Cambridge Monographs on Applied and Computational Mathematics. Cambridge University Press, Cambridge, 2007. DOI: 10.1017/CBO9780511546860.

[Dey+05] Tamal K. Dey, Joachim Giesen, Edgar A. Ramos, and Bardia Sadri. "Critical points of the distance to an epsilon-sampling of a surface and flow-complex-based surface reconstruction". In: *Computational geometry (SCG'05)*. ACM, NY, 2005. DOI: 10.1145/1064092.1064126.

[Die08] Tammo tom Dieck. *Algebraic topology*. EMS Textbooks in Mathematics. European Mathematical Society (EMS), Zürich, 2008. DOI: 10.4171/048.

[GLR22] Barbara Giunti, Jānis Lazovskis, and Bastian Rieck. *DONUT: Database of Original & Non-Theoretical Uses of Topology*. https://donut.topology.rocks. 2022.

[Dua14] Ran Duan. "Approximation algorithms for the Gromov hyperbolicity of discrete metric spaces". In: *LATIN 2014: theoretical informatics*. Vol. 8392. Springer, Heidelberg, 2014. DOI: 10.1007/978-3-642-54423-1_25.

[DP09] Ran Duan and Seth Pettie. "Fast algorithms for (max,min)-matrix multiplication and bottleneck shortest paths". In: *Proceedings of the Twentieth Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM, Philadelphia, PA, 2009.

[Dug08] Daniel Dugger. *A primer on homotopy colimits*. Preprint (updated in 2017). 2008. URL: https://_pages.uoregon.edu/ddugger/hocolim.pdf.

[DI04] Daniel Dugger and Daniel Isaksen. "Topological hypercovers and $\mathbb{A}^1$-realizations". In: *Mathematische Zeitschrift* 246 (Apr. 2004). DOI: 10.1007/s00209-003-0607-y.

[DS95] W. Dwyer and J. Spalinski. "Homotopy theories and model categories". In: *Handbook of Algebraic Topology*. Ed. by I.M. James. 1995.

[EKS83] H. Edelsbrunner, D. Kirkpatrick, and R. Seidel. "On the shape of a set of points in the plane". In: *IEEE Transactions on Information Theory* 29.4 (1983). DOI: 10.1109/TIT.1983.1056714.

[Ede03] Herbert Edelsbrunner. "Surface reconstruction by wrapping finite sets in space". In: *Discrete and computational geometry*. Vol. 25. Algorithms Combin. Springer, Berlin, 2003. DOI: 10.1007/978-3-642-55566-4_17.

[Ede95] Herbert Edelsbrunner. "The union of balls and its dual shape". In: *Discrete Comput. Geom.* 13.3-4 (1995). DOI: 10.1007/BF02574053.

[EH10] Herbert Edelsbrunner and John L. Harer. *Computational topology*. An introduction. American Mathematical Society, Providence, RI, 2010. DOI: 10.1090/mbk/069.

[ELZ02] Herbert Edelsbrunner, David Letscher, and Afra Zomorodian. "Topological persistence and simplification". In: vol. 28. 4. Discrete and computational geometry and graph drawing (Columbia, SC, 2001). 2002. DOI: 10.1007/s00454-002-2885-2.

[EM13] Herbert Edelsbrunner and Dmitriy Morozov. "Persistent homology: theory and practice". In: *European Congress of Mathematics*. Eur. Math. Soc., Zürich, 2013.

[EÖ20] Herbert Edelsbrunner and Katharina Ölsböck. "Tri-partitions and bases of an ordered complex". In: *Discrete Comput. Geom.* 64.3 (2020). DOI: 10.1007/s00454-020-00188-x.

[EZ03] Herbert Edelsbrunner and Afra Zomorodian. "Computing linking numbers of a filtration". In: *Homology Homotopy Appl.* 5.2 (2003). Algebraic topological methods in computer science (Stanford, CA, 2001). DOI: 10.4310/HHA.2003.v5.n2.a2.

[EH80] David A. Edwards and Harold M. Hastings. "Čech theory: its past, present, and future". In: *Rocky Mount. J. Math.* 10.3 (1980). DOI: 10.1216/RMJ-1980-10-3-429.

[ES52] Samuel Eilenberg and Norman Steenrod. *Foundations of algebraic topology*. Princeton University Press, Princeton, New Jersey, 1952.

[EK01] R. Espínola and M. A. Khamsi. "Introduction to hyperconvex spaces". In: *Handbook of metric fixed point theory*. Kluwer Acad. Publ., Dordrecht, 2001. DOI: 10.1007/978-94-017-1748-9_13.

[Eva08] Steven N. Evans. *Probability and real trees*. Vol. 1920. Lecture Notes in Mathematics. Lectures from the 35th Summer School on Probability Theory held in Saint-Flour, July 6–23, 2005. Springer, Berlin, 2008. DOI: 10.1007/978-3-540-74798-7.

[Far03] Michael Farber. "Topological Complexity of Motion Planning". In: *Discrete Comput. Geom.* 29.2 (2003). DOI: 10.1007/s00454-002-0760-9.

[FMN14] Steve Ferry, Konstantin Mischaikow, and Vidit Nanda. "Reconstructing functions from random samples". In: *Journal of Computational Dynamics* 1.2 (2014). DOI: 10.3934/jcd.2014.1.233.

[For02] Robin Forman. "A user's guide to discrete Morse theory". In: *Sém. Lothar. Combin.* 48 (2002).

[For98] Robin Forman. "Morse theory for cell complexes". In: *Adv. Math.* 134.1 (1998). DOI: 10.1006/aima.1997.1650.

[FIV15] Hervé Fournier, Anas Ismail, and Antoine Vigneron. "Computing the Gromov hyperbolicity of a discrete metric space". In: *Inform. Process. Lett.* 115.6-8 (2015). DOI: 10.1016/j.ipl.2015.02.002.

[Fre09] Ragnar Freij. "Equivariant discrete Morse theory". In: *Discrete Math.* 309.12 (2009). DOI: 10.1016/j.disc.2008.10.029.

[Fri12] Greg Friedman. "Survey article: An elementary illustrated introduction to simplicial sets". In: *Rocky Mount. J. Math.* 42.2 (2012). DOI: 10.1216/RMJ-2012-42-2-353.

[GZ67] P. Gabriel and M. Zisman. *Calculus of fractions and homotopy theory*. Ergebnisse der Mathematik und ihrer Grenzgebiete, Band 35. Springer NY, 1967.

[Ghr14] Robert Ghrist. *Elementary Applied Topology*. 2014. ISBN: 978-1502880857. URL: https://www2.math.upenn.edu/~ghrist/notes.html.

[GS07] Paul Goerss and Kristen Schemmerhorn. "Model categories and simplicial methods". In: *Interactions between homotopy theory and algebra*. Vol. 436. Contemp. Math. Amer. Math. Soc., 2007.

[GJ09] Paul G. Goerss and John F. Jardine. *Simplicial homotopy theory*. Modern Birkhäuser Classics. Reprint of the 1999 edition. Birkhäuser Verlag, Basel, 2009. DOI: 10.1007/978-3-0346-0189-4.

[GM23] Mario Gómez and Facundo Mémoli. *Curvature Sets Over Persistence Diagrams*. 2023. arXiv: 2103.04470.

[GS18] Dejan Govc and Primoz Skraba. "An approximate nerve theorem". In: *Found. Comput. Math.* 18.5 (2018). DOI: 10.1007/s10208-017-9368-6.

[Gra75] Brayton Gray. *Homotopy theory*. Pure and Applied Mathematics, Vol. 64. An introduction to algebraic topology. Academic Press, NY-London, 1975.

[Gro87] M. Gromov. "Hyperbolic groups". In: *Essays in group theory*. Vol. 8. Math. Sci. Res. Inst. Publ. Springer, NY, 1987. DOI: 10.1007/978-1-4613-9586-7_3.

[Hat02] Allen Hatcher. *Algebraic topology*. Cambridge University Press, Cambridge, 2002. ISBN: 9780521795401.

[Hau95] Jean-Claude Hausmann. "On the Vietoris-Rips complexes and a cohomology theory for metric spaces". In: *Prospects in topology (Princeton, NJ, 1994)*. Vol. 138. Ann. of Math. Stud. Princeton Univ. Press, Princeton, NJ, 1995.

[HMR21] Felix Hensel, Michael Moor, and Bastian Rieck. "A Survey of Topological Machine Learning Methods". In: *Frontiers in Artificial Intelligence* 4 (2021). DOI: 10.3389/frai.2021.681108.

[HG16] G. Henselman and R. Ghrist. *Matroid Filtrations and Computational Persistent Homology*. Software available at https://github.com/Eetion/Eirene.jl. 2016. arXiv: 1606.00199.

[Her05] Patricia Hersh. "On optimizing discrete Morse functions". In: *Adv. in Appl. Math.* 35.3 (2005). DOI: 10.1016/j.aam.2005.04.001.

[Hir03] Philip S. Hirschhorn. *Model categories and their localizations*. Vol. 99. Mathematical Surveys and Monographs. American Mathematical Society, 2003.

[Hov99] Mark Hovey. *Model categories*. Vol. 63. Mathematical Surveys and Monographs. American Mathematical Society, 1999.

[Hul] Michael Hull. *Hyperbolic Groups Lecture Notes*. URL: http://homepages.math.uic.edu/~mbhull/hyperbolic%20lecture%20notes.pdf (visited on 02/13/2024).

[Isb64] J. R. Isbell. "Six theorems about injective metric spaces". In: *Comment. Math. Helv.* 39 (1964). DOI: 10.1007/BF02566944.

[JW09] Michael Jöllenbeck and Volkmar Welker. "Minimal resolutions via algebraic discrete Morse theory". In: *Mem. Am. Math. Soc.* 197.923 (2009). DOI: 10.1090/memo/0923.

[Jon08] Jakob Jonsson. *Simplicial complexes of graphs*. Vol. 1928. Lecture Notes in Mathematics. Springer, Berlin, 2008. DOI: 10.1007/978-3-540-75859-4.

[Jos83] K. D. Joshi. *Introduction to general topology*. A Halsted Press Book. John Wiley & Sons, Inc., NY, 1983. ISBN: 0-470-27556-1.

[Kah11] Matthew Kahle. "Random geometric complexes". In: *Discrete Comput. Geom.* 45.3 (2011). DOI: 10.1007/s00454-010-9319-3.

[KM21] Woojin Kim and Facundo Mémoli. "Generalized persistence diagrams for persistence modules over posets". In: *J. Appl. Comput. Topol.* 5.4 (2021). DOI: 10.1007/s41468-021-00075-1.

[Koz08] Dmitry Kozlov. *Combinatorial algebraic topology.* Vol. 21. Algorithms and Computation in Mathematics. Springer, Berlin, 2008. DOI: 10.1007/978-3-540-71962-5.

[Koz20] Dmitry N. Kozlov. *Organized collapse: an introduction to discrete Morse theory.* Vol. 207. Graduate Studies in Mathematics. American Mathematical Society, Providence, RI, 2020. ISBN: 978-1-4704-5701-3.

[Lam20] Leon Lampret. *Chain complex reduction via fast digraph traversal.* 2020. arXiv: 1903.00783.

[Lan13] Urs Lang. "Injective hulls of certain discrete metric spaces and groups". In: *J. Topol. Anal.* 5.3 (2013). DOI: 10.1142/S1793525313500118.

[Lat01] Janko Latschev. "Vietoris-Rips complexes of metric spaces near a closed Riemannian manifold". In: *Arch. Math. (Basel)* 77.6 (2001). DOI: 10.1007/PL00000526.

[Lef42] Solomon Lefschetz. *Algebraic Topology.* AMS books online. American Mathematical Society, 1942. ISBN: 9780821810279.

[LV23] Boštjan Lemež and Žiga Virk. *Finite reconstruction with selective Rips complexes.* 2023. arXiv: 2205.05525.

[Ler45] Jean Leray. "Sur la forme des espaces topologiques et sur les points fixes des représentations". In: *J. Math. Pures Appl. (9)* 24 (1945). ISSN: 0021-7824.

[LRR20] Michael Lesnick, Raúl Rabadán, and Daniel I. S. Rosenbloom. "Quantifying genetic innovation: mathematical foundations for the topological study of reticulate evolution". In: *SIAM J. Appl. Algebra Geom.* 4.1 (2020). DOI: 10.1137/18M118150X.

[Lil73] Joachim Lillig. "A union theorem for cofibrations". In: *Arch. Math. (Basel)* 24 (1973). DOI: 10.1007/BF01228231.

[LMO22] Sunhyuk Lim, Facundo Memoli, and Osman Berat Okutan. *Vietoris-Rips Persistent Homology, Injective Metric Spaces, and The Filling Radius.* Preprint. 2022. arXiv: 2001.07588.

[LMS23] Sunhyuk Lim, Facundo Mémoli, and Zane Smith. "The Gromov–Hausdorff distance between spheres". In: *Geom. Topol.* 27.9 (2023). DOI: 10.2140/gt.2023.27.3733.

[May99] J.P. May. *A Concise Course in Algebraic Topology.* Univ. of Chicago Press, 1999.

[MP12] J.P. May and K. Ponto. *More Concise Algebraic Topology: Localization, Completion, and Model Categories.* The University of Chicago Press, 2012.

[McC67] Michael C. McCord. "Homotopy type comparison of a space with complexes associated with its open covers". In: *Proc. Amer. Math. Soc.* 18 (1967). DOI: 10.2307/2035443.

[McC66] Michael C. McCord. "Singular homology groups and homotopy groups of finite topological spaces". In: *Duke Math. J.* 33 (1966). DOI: 10.1215/S0012-7094-66-03352-7.

[MOW21] Facundo Memoli, Osman Berat Okutan, and Qingsong Wang. *Metric Graph Approximations of Geodesic Spaces.* Preprint. 2021. arXiv: 1809.05566.

[Mes01] Roy Meshulam. “The clique complex and hypergraph matching”. In: *Combinatorica* 21.1 (2001). DOI: 10.1007/s004930170006.

[Mod+24] Tejasvinee Atul Mody, Alexander Rolle, Nico Stucki, Fabian Roll, Ulrich Bauer, and Kay Schneitz. “Topological analysis of 3D digital ovules identifies cellular patterns associated with ovule shape diversity”. In: *Development* 151.20 (2024). DOI: 10.1242/dev.202590.

[Mor08] Dmitriy Morozov. “Homological Illusions of Persistence and Stability”. Theses. Duke University, 2008. URL: https://dukespace.lib.duke.edu/items/0714688b-fcf7-4977-b936-6aa5b9ab78e9.

[Mun84] James R. Munkres. *Elements of algebraic topology*. Addison-Wesley Publishing Company, Menlo Park, CA, 1984. ISBN: 0-201-04586-9.

[Nag07] Andrzej Nagórko. “Carrier and nerve theorems in the extension theory”. In: *Proc. Amer. Math. Soc.* 135.2 (2007). DOI: 10.1090/S0002-9939-06-08477-2.

[NSW08] Partha Niyogi, Stephen Smale, and Shmuel Weinberger. “Finding the homology of submanifolds with high confidence from random samples”. In: *Discrete Comput. Geom.* 39.1-3 (2008). DOI: 10.1007/s00454-008-9053-2.

[Oud15] Steve Y. Oudot. *Persistence theory: from quiver representations to data analysis*. Vol. 209. Mathematical Surveys and Monographs. American Mathematical Society, Providence, RI, 2015. DOI: 10.1090/surv/209.

[PH15] Jose A. Perea and John Harer. “Sliding windows and persistence: an application of topological methods to signal analysis”. In: *Found. Comput. Math.* 15.3 (2015). DOI: 10.1007/s10208-014-9206-z.

[Por+22] Cédric Portaneri, Mael Rouxel-Labbé, Michael Hemmer, David Cohen-Steiner, and Pierre Alliez. “Alpha wrapping with an offset”. In: *ACM Trans. Graph.* 41.4 (2022). DOI: 10.1145/3528223.3530152.

[Qui73] Daniel Quillen. “Higher algebraic K-theory. I”. In: *Algebraic K-theory, I: Higher K-theories (Proc. Conf., Battelle Memorial Inst., Seattle, Wash., 1972)*. 1973.

[Qui67] Daniel Quillen. *Homotopical Algebra*. Vol. 43. Lect. Notes in Math. Springer, 1967.

[Qui78] Daniel Quillen. “Homotopy properties of the poset of nontrivial p-subgroups of a group”. In: *Adv. in Math.* 28.2 (1978). DOI: 10.1016/0001-8708(78)90058-0.

[RS07] Edgar A. Ramos and Bardia Sadri. “Geometric and topological guarantees for the WRAP reconstruction algorithm”. In: *Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms*. ACM, NY, 2007.

[Rie14] Emily Riehl. *Categorical homotopy theory*. Vol. 24. New Mathematical Monographs. Cambridge University Press, Cambridge, 2014. DOI: 10.1017/CBO9781107261457.

[Rie16] Emily Riehl. *Category Theory in Context*. Aurora: Dover Modern Math Originals. Dover Publications, Inc., 2016.

[Rob99] Vanessa Robins. “Towards computing homology from finite approximations”. In: *Proceedings of the 14th Summer Conference on General Topology and its Applications (Brookville, NY, 1999)*. Vol. 24. Summer. 1999.

[Rol20] Fabian Roll. “Variations of the Nerve Theorem”. MA thesis. Technische Universität München, 2020. URL: https://mediatum.ub.tum.de/node?id=1730967.

[Rol23] Fabian Roll. *Wrapping Cycles*. Executable directly by running `docker build -o output github.com/fabian-roll/wrappingcycles` on any machine with Docker installed and configured with sufficient memory (16GB recommended). 2023. URL: https://github.com/fabian-roll/wrappingcycles.

[Rol22] Alexander Rolle. "The degree-Rips complexes of an annulus with outliers". In: *38th International Symposium on Computational Geometry*. Vol. 224. LIPIcs. Leibniz Int. Proc. Inform. Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern, 2022. DOI: 10.4230/lipics.socg.2022.58.

[RS23] Alexander Rolle and Luis Scoccola. *Stable and consistent density-based clustering via multiparameter persistence*. 2023. arXiv: 2005.09048.

[Sad09] Bardia Sadri. "Manifold Homotopy via the Flow Complex". In: *Comput. Graph. Forum* 28.5 (2009). DOI: 10.1111/J.1467-8659.2009.01513.X.

[Sch22] Maximilian Schmahl. *Comparison of persistent singular and Čech homology for locally connected filtrations*. 2022. arXiv: 2207.00848.

[Seg68] Graeme Segal. "Classifying spaces and spectral sequences". In: *Inst. Hautes Études Sci. Publ. Math.* 34 (1968). ISSN: 0073-8301.

[SMV11] Vin de Silva, Dmitriy Morozov, and Mikael Vejdemo-Johansson. "Dualities in persistent (co)homology". In: *Inverse Problems* 27.12 (2011). DOI: 10.1088/0266-5611/27/12/124003.

[SC04] Vin de Silva and Gunnar Carlsson. "Topological estimation using witness complexes". In: *SPBG'04 Symposium on Point-Based Graphics 2004*. The Eurographics Association, 2004. DOI: 10.2312/SPBG/SPBG04/157-166.

[Skö06] Emil Sköldberg. "Morse theory from an algebraic viewpoint". In: *Trans. Amer. Math. Soc.* 358.1 (2006). DOI: 10.1090/S0002-9947-05-04079-1.

[Sta] The Stanford 3D Scanning Repository. URL: http://graphics.stanford.edu/data/3Dscanrep/ (visited on 10/18/2022).

[SS95] Lynn Arthur Steen and J. Arthur Seebach Jr. *Counterexamples in topology*. Reprint of the second (1978) edition. Dover Publications, Inc., 1995. ISBN: 0-486-68735-X.

[Str09] N. P. Strickland. *The Category of CGWH Spaces*. 2009. URL: https://ncatlab.org/nlab/files/StricklandCGHWSpaces.pdf.

[Str72] Arne Strøm. "The homotopy category is a homotopy category". In: *Archiv der Mathematik* 23 (1972).

[Stu+23] Nico Stucki, Johannes C. Paetzold, Suprosanna Shit, Bjoern Menze, and Ulrich Bauer. "Topologically Faithful Image Segmentation via Induced Matching of Persistence Barcodes". In: *Proceedings of the 40th International Conference on Machine Learning*. Vol. 202. Proceedings of Machine Learning Research. PMLR, 2023.

[Tay+15] Dane Taylor, Florian Klimm, Heather A. Harrington, Miroslav Kramár, Konstantin Mischaikow, Mason A. Porter, and Peter J. Mucha. "Topological Data Analysis of Contagion Maps for Examining Spreading Processes on Networks". In: *Nature Communications* 6.1 (2015), p. 7723. DOI: 10.1038/ncomms8723.

[The15] The GUDHI Project. *GUDHI User and Reference Manual*. GUDHI Editorial Board, 2015. URL: http://gudhi.gforge.inria.fr/doc/latest/.

[UZ16] Michael Usher and Jun Zhang. "Persistent homology and Floer-Novikov theory". In: *Geom. Topol.* 20.6 (2016). DOI: 10.2140/gt.2016.20.3333.

[Vie27] Leopold Vietoris. “Über den höheren Zusammenhang kompakter Räume und eine Klasse von zusammenhangstreuen Abbildungen”. In: *Math. Ann.* 97.1 (1927). DOI: 10.1007/BF01447877.

[Vir21] Žiga Virk. “Rips Complexes as Nerves and a Functorial Dowker-Nerve Diagram”. In: *Medit. J. of Math.* 18 (2021). DOI: 10.1007/s00009-021-01699-4.

[Vui21] Julien Vuillamy. “Planimetric simplification and lexicographic optimal chains for 3D urban scene reconstruction”. Theses. Université Côte d’Azur, Sept. 2021. URL: https://hal.science/tel-03339931.

[Wal81] James W. Walker. “Homotopy type and Euler characteristic of partially ordered sets”. In: *Europ. J. Combin.* 2.4 (1981). DOI: 10.1016/S0195-6698(81)80045-5.

[Wei94] Charles A. Weibel. *An introduction to homological algebra*. Vol. 38. Cambridge Studies in Advanced Mathematics. Cambridge University Press, Cambridge, 1994. DOI: 10.1017/CBO9781139644136.

[Wei52] André Weil. “Sur les théorèmes de de Rham”. In: *Comment. Math. Helv.* 26 (1952). DOI: 10.1007/BF02564296.

[WZŽ99] Volkmar Welker, Günter M. Ziegler, and Rade T. Živaljević. “Homotopy colimits – comparison lemmas for combinatorial applications”. In: *Journal für die reine und angewandte Mathematik* 509 (1999). DOI: 10.1515/crll.1999.509.117.

[Whi50] J. H. C. Whitehead. “Simple homotopy types”. In: *Amer. J. Math.* 72 (1950). DOI: 10.2307/2372133.

[Wu62] W. Wu. “On a theorem of Leray”. In: *Chinese Math.* 2 (1962). ISSN: 0577-909X.

[Zar22] Matthew C. B. Zaremsky. “Bestvina-Brady discrete Morse theory and Vietoris-Rips complexes”. In: *Amer. J. Math.* 144.5 (2022). ISSN: 0002-9327.

[Zav24] Nicolò Zava. *Coarse and bi-Lipschitz embeddability of subspaces of the Gromov-Hausdorff space into Hilbert spaces*. 2024. arXiv: 2303.04730.

[ZC05] Afra Zomorodian and Gunnar Carlsson. “Computing persistent homology”. In: *Discrete Comput. Geom.* 33.2 (2005). DOI: 10.1007/s00454-004-1146-y.

[ZC08] Afra Zomorodian and Gunnar Carlsson. “Localized homology”. In: *Comput. Geom.* 41.3 (2008). DOI: 10.1016/j.comgeo.2008.02.003.