INAUGURAL – DISSERTATION

zur

Erlangung der Doktorwürde

der

Gesamtfakultät für Mathematik, Ingenieur- und Naturwissenschaften

der

Ruprecht – Karls – Universität

Heidelberg

vorgelegt von

Mäder, Tim, M.Sc.

aus: Bremen

Tag der mündlichen Prüfung: .....

Stratified Topological Data Analysis: Theory, Methods and Applications in Image Patch Spaces

Betreuer: Prof. Dr. Markus Banagl

Abstract

Singular spaces have gained prominence in diverse applications, generating interest in investigating stratified structures using methods within the field of topological data analysis (TDA). While persistent homology is a fundamental concept in TDA, its limitations in distinguishing stratified spaces prompt the need for alternative invariants.

To address this need, we introduce the concept of persistent stratified homotopy types. We establish a series of properties analogous to those of the ordinary persistent homotopy type, such as stability, computability, and inference, which are crucial for TDA applications.

In tackling the challenge of detecting singularities in data, we present methods to approximate stratifications from non-stratified data, sampled in the vicinity of sufficiently well-behaved two-strata Whitney stratified spaces.

Our findings enable the formulation of a sampling theorem for approximately inferring (persistent) stratified homotopy types of suitably well-behaved two strata Whitney stratified spaces.

Bridging theory and practical implementation, we introduce two distinct approaches to measure singularity in data. Leveraging insights from local persistent homology studies and employing the Hausdorff distance serve this purpose.

Moreover, we delve into the topology of image patch spaces, providing novel insights and reassessing existing models, especially with regard to potential stratified structures.

Zusammenfassung

Singuläre Räume haben in verschiedenen Anwendungen an Bedeutung gewonnen, was das Interesse an der Untersuchung stratifizierter Strukturen unter Verwendung von Methoden im Bereich der topologischen Datenanalyse (TDA) geweckt hat. Obwohl die persistente Homologie ein grundlegendes Konzept in der TDA ist, erfordern ihre Einschränkungen bei der Unterscheidung stratifizierter Räume die Entwicklung alternativer Invarianten.

Um diesem Bedarf gerecht zu werden, führen wir das Konzept der persistenten stratifizierten Homotopietypen ein. Wir etablieren eine Reihe von Eigenschaften, die denen des gewöhnlichen persistenten Homotopietyps ähnlich sind, wie Stabilität, Berechenbarkeit und Inferenz, die für TDA-Anwendungen entscheidend sind.

Um die Herausforderung der Detektion von Singularitäten in Daten anzugehen, präsentieren wir Methoden zur Approximation von Stratifizierungen aus nicht stratifizierten Daten, die in der Nähe ausreichend gut verhaltener Whitney stratifizierter Räume mit zwei Strata abgetastet wurden. Unsere Erkenntnisse ermöglichen die Formulierung eines Sampling Theorems zur ungefähren Ableitung (persistenter) stratifizierter Homotopietypen in angemessen gut verhaltenen Whitney stratifizierten Räumen mit zwei Strata.

Die Verbindung von Theorie und praktischer Umsetzung zeigt sich in der Einführung von zwei unterschiedlichen Ansätzen zur Messung der Singularität in Daten. Die Nutzung von Erkenntnissen aus Studien zur lokalen persistenten Homologie sowie die Anwendung des Hausdorff-Abstands dienen diesem Zweck.

Darüber hinaus untersuchen wir die Topologie von Bildflickenräumen, indem wir neue Erkenntnisse liefern und bestehende Modelle neu bewerten, insbesondere im Hinblick auf mögliche stratifizierte Strukturen.

Acknowledgements

I extend my sincere gratitude to my fellow PhD candidates, Shahryar Ghaed Sharaf, Lukas Waas, and Martin Rabel, with whom I shared an office and transformed it into a welcoming space for camaraderie. My special thanks to Lukas, with whom I collaborated closely and to whom I owe a substantial debt. His profound impact on the success of my work is undeniable, and I cherish the friendship that blossomed during our shared journey as PhD candidates.

I am indebted to Markus Banagl, for being my supervisor, a constant well of helpful suggestions and being one of the reasons I became interest in the field of topology in the first place. Christoph Schnörr deserves my gratitude for generously sharing his knowledge and expertise in image analysis, contributing significantly to the progress of my project. Filip Saldo's assistance in the analysis of image data and scientific visualization is also acknowledged with appreciation. Special thanks to Roland Herzog for providing valuable advice on an optimization problem that emerged in my studies.

My sincere thanks to the German Research Foundation (DFG, Deutsche Forschungsgemeinschaft) for their financial support of my research through Germany's Excellence Strategy EXC-2181/1 - 390900948 (the Heidelberg STRUCTURES Excellence Cluster). I am grateful for the opportunity to be part of this scientific community, especially as a member of comprehensive Project 6.

Lastly, heartfelt thanks to my family, friends, and my girlfriend Lea for their unwavering support throughout this journey.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Topological data analysis (TDA) has demonstrated its reliability in extracting qualitative and quantitative data features that are not easily accessible through other methods. Singular spaces have appeared in several application contexts [LPM03, BMP+08, CIDSZ08, MW11, FW16, Xia16, STHN20, ESM+21]. The emergence of stratified structures in these real-world application has also sparked an increasing interest in the performance of TDA techniques for the exploration of stratified spaces in data. One widely used method and arguably the most important concept for the development of the field of TDA is persistent homology ([ELZ00, ZC05, CSEH07, Ghr08, NSW08, Car09, Oud15]). Some of the main properties of persistent homology that made it a powerful tool in applications are the following:

(1): The fact that persistent homology defined through thickenings is computable by the means of a filtered simplicial complex, e.g. the Čech complex.

(2): The stability of persistent homology with respect to Hausdorff and interleaving type distances (see [CSEH07, CCSG+09, BL15]).

(3): Persistent homology can retrieve topological information about a space through data samples from the space. This is usually based on the result that $(X)_\alpha \xrightarrow{\simeq} (X)_0 = X$ for α a sufficiently small thickening parameter and X a sufficiently well-behaved space (see [NSW08, CCSL09]).

Nevertheless, certain characteristics of persistent homology imply that it may not be the optimal invariant for distinguishing stratified spaces. To illustrate, consider the two subspaces within $\mathbb{R}^2$ portrayed in Figs. 1.1 and 1.2. These spaces can be proportionally resized to exhibit identical persistent homology.

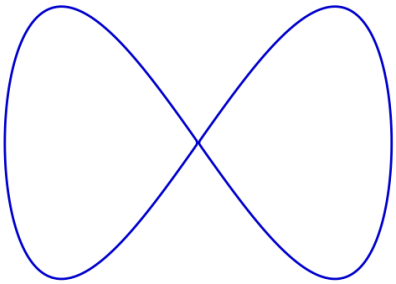

Figure 1.1. Lemniscate $x^4 - x^2 + y^2 = 0$

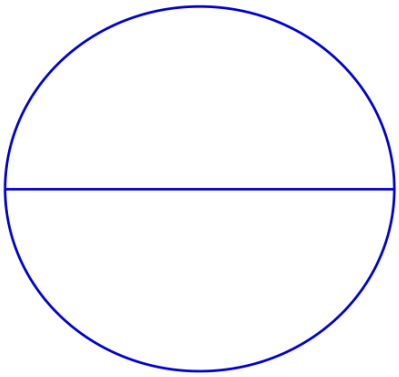

Figure 1.2. Circle with a diameter filament

However, the spaces themselves exhibit distinct topological characteristics, with the space depicted in Fig. 1.1 featuring one singularity and the space in Fig. 1.2 possessing two singularities. Depending on the specific application, there may be a need for an invariant capable of distinguishing between these scenarios. The field of stratified homotopy theory has recently experienced a resurgence, marked by notable breakthroughs [Woo09, Lur17, Mil13, Hai18, Dou21a, Dou21b, DW21]. These advancements have opened avenues for effectively applying these concepts in the realm of topological data analysis.

A declared objective of the collaborative project with Lukas Waas [MW22] was to establish a notion of persistent stratified homotopy type, demonstrating its fulfillment of properties akin to those outlined above.

Another challenge encountered when dealing with stratified invariants associated to point clouds is the absence of information within the data indicating which points should be classified as singular and which as regular. Consequently, it is imperative to develop methodologies for detecting singularities in the data. Substantial attention has been directed towards exploring various approaches to discerning the locations of singularities, as evident in numerous studies [Mil21, STHN20, Nan20, SW14, BWM12, FW16]. In the context of our collaborative project with Lukas Waas [MW22], our specific focus was on investigating the recovery of an approximation of stratifications from non-stratified data samples.

In the realm of application-focused disciplines such as Topological Data Analysis (TDA), we feel it is important to bridge the gap between theoretical foundations and practical implementation. Recognizing that theoretical insights form the backbone of any scientific endeavor, it is equally vital to delve into the realm of algorithms and real-world implementations. We therefore discuss algorithms tailored specifically for approximating the intricate concepts within our theoretical framework. We can report

on two different approaches. One is leveraging insights from prior studies that tackled stratification with the concept of local persistent homology [BWM12, SW14, Nan20, Mil21]. Our second method uses the Hausdorff distance of a local point neighborhood of every point in a point cloud to linear subspaces of Euclidean space in order to obtain a measure for singularity.

One of the primary motivators behind our advancement was also to explore the topology of image patch spaces that arise in various contexts such as natural image statistics or from single images. The topology of image patch spaces has been studied before [LPM03, SC04, CIDSZ08, Xia16, CG20] and uncovered a surprising yet incomplete picture of what topology the image patch space extracted from an image can have. We will reinvestigate a model proposed for the image patch space from natural images with our new methods. We also take a different look at the topology that can be found in image patch spaces by considering single images. Furthermore, we attempt to demonstrate the utility of stratification learning and persistent stratified homotopy type for real-world data analysis by studying a collection of retina artery images. In the following section we would like to present our results on the aforementioned topics in a short yet accessible way and also refer the reader to the various places where to find the respective results.

1.1 Persistent Stratified Homotopy Types

We would like to briefly explain our results regarding the concept of persistent stratified homotopy types by starting with ordinary persistent homotopy types. Recall that persistent homology can be split into a two-step process

$$\text{point clouds} \to \text{persistence modules} \to \text{persistence diagrams}.$$

For instance, consider a finite subset $\mathbb{X} \subset \mathbb{R}^N$ where a filtration of topological spaces, denoted as $(\mathbb{X}_\alpha)_{\alpha \geq 0}$, is assigned. This is typically achieved by progressively thickening the data set $\mathbb{X}$ into α-thickened spaces. Subsequently, persistence modules, which are filtered objects, are computed by determining homology at each filtration degree. Given the homotopy invariance of homology, the pertinent information in this computation is the *persistent homotopy type* of $\mathbb{X}$, often algorithmically captured using a Čech complex. Remarkably, the properties outlined in (1), (2), and (3) can be discerned at this level. However, the limitations in distinguishing stratified spaces, as illustrated in Figs. 1.1 and 1.2, also emanate from this approach. In order to advance toward a persistent invariant suitable for investigating stratified spaces from point clouds, our focus shifted to the initial step of the factorization, leading to the

identification of a persistent stratified homotopy type.

To begin, according to the findings in [Dou19, DW21], the stratified homotopy type of a suitably regular stratified space, like a Whitney stratified space W (Recollection 2.1.15) with two strata (or, more broadly, a conically stratified space), can be equivalently represented (Theorem 2.1.19 and Recollection 2.1.23) by the homotopy type of a *(stratification) diagram* of spaces

$$\{D(W)_p \leftarrow D(W)_{\{p,q\}} \rightarrow D(W)_q\}$$

where $D(W)_p$ and $D(W)_q$ correspond to the singular and the regular stratum of W respectively. The middle part $D(W)_{\{p,q\}}$ of the diagram corresponds to the homotopy type of the part that links the singular and the regular strata and is therefore called the *(homotopy) link* (Definition 2.1.8).

Depending on a parameter $v = (v_l, v_h)$, which determines the distance at which the link part of the diagram is located relative to the singular stratum, a stratification diagram for the Lemniscate with $v = (0.2, 0.3)$ is illustrated in the following:

Similarly, for the circle with diameter filament:

Besides serving as a discriminant invariant for stratified spaces, observe that comparing the amount of connected components in all the diagram parts distinguishes the spaces, a diagram also allows for a systematic way of thickening in order to make it a persistent invariant, i.e. to allow for the construction of a *persistent stratified homotopy type*. We construct such an object by thickening the parts of a diagram separately in an ambient Euclidean space. The following pictures give an illustration of this process for the previous case of the Lemniscate. For the scale $\varepsilon = 0.12$ we have

and for $\varepsilon = 0.24$ we have

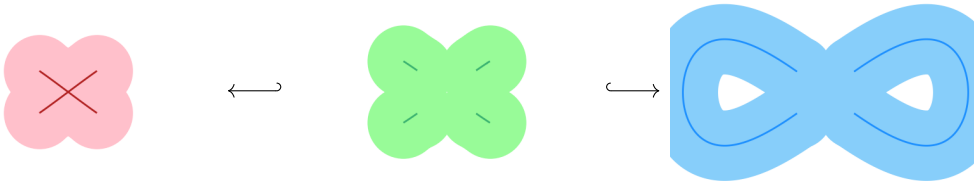

The first diagram indicates that thickening by a small amount does not change the type of the diagram while in the second thickened diagram we can see the point at which the type of the diagram changes.

In general, for a chosen parameter $v = (v_l, v_h)$, we can associate to a stratified subset $S \subset \mathbb{R}^n$ a persistent stratified homotopy type $\mathcal{P}_v(S)$. This means for any scale $\varepsilon \in \mathbb{R}_+$ we have a diagram indexed over $\{p\} \leftarrow \{p, q\} \rightarrow \{q\}$ which corresponds to a (weak) stratified homotopy type. The anticipated properties of this construction are captured in several results. For the stratified analogue of (1), the computability of thickenings of diagrams, see Remark 2.2.15. For the statement that the persistent stratified homotopy type does not change under small thickenings, i.e. the stratified version of (3), we included Proposition 2.2.16. Lastly, Theorems 2.2.31 and 2.2.32 provide the stability result, that is the substitution for (2). This implies that for a sequence of stratified spaces $\mathbb{S}_i \subset \mathbb{R}^n$ converging (in a stratified Hausdorff distance) to a two-strata Whitney stratified space W, the persistent stratified homotopy types $\mathcal{P}_v(\mathbb{S}_i)$ also converge to $\mathcal{P}_v(W)$. Subsequently, one can associate algebraic invariants such as persistent homology of the links and strata, which also converge in interleaving distance. In summary, we arrive at the following result:

Theorem 1.1.1. *The persistent stratified homotopy type $\mathcal{P}_v(S)$ associated to a stratified subset $S \subset \mathbb{R}^N$ with two strata (depending on a choice of parameter v) fulfills stratified analogues of (1), (2) and (3).*

1.2 Approximating Stratifications

We present a comprehensive and algorithmically executable pipeline designed to produce a stratification from a provided non-stratified data sample of a specific class of

Whitney stratified spaces with two strata. Our study notably investigates the essential conditions for the convergence of such an approximated stratification to the original, as depicted in Fig. 1.3.

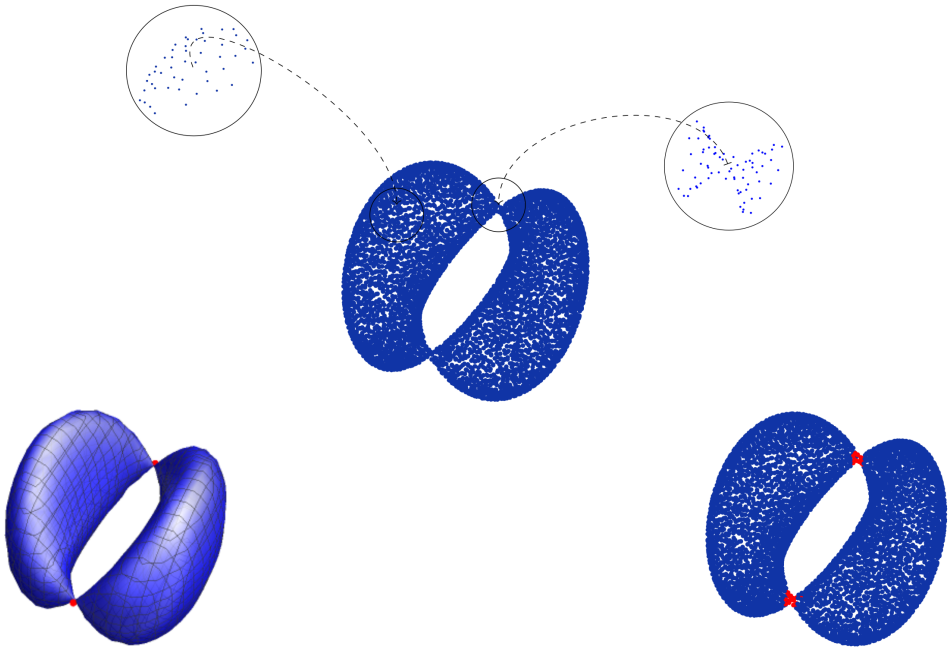

Figure 1.3. A stratified space, a sample with indicated magnifications and an approximated stratification

Our method demonstrates promising results in situations where singularities can be effectively identified by its local tangential geometry. It is clear that a stratification can only be retrieved from a data set if the actual stratification of the underlying space is recognisable on a (local) geometric level. The spaces for which we are able to apply our results are spaces whose singularities can be detected by the *extrinsic tangent cones*, denoted $\mathrm{T}_x^{\mathrm{ex}}(X)$ for a space X, a generalisation of the tangent space to the singular case (see Definition 3.1.5). Such spaces will be referred to as *tangentially stratified*. See Fig. 1.4 for an illustrations of a space that is tangentially stratified and see Fig. 1.5 for one that is not.

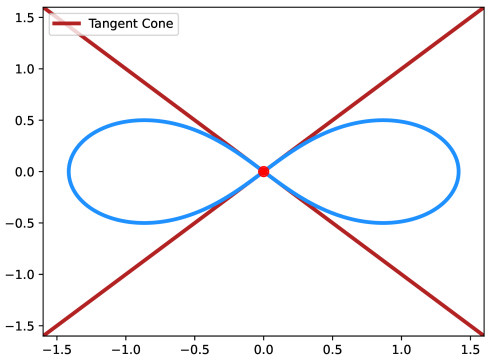

Figure 1.4. A curve and its tangential cone at the origin

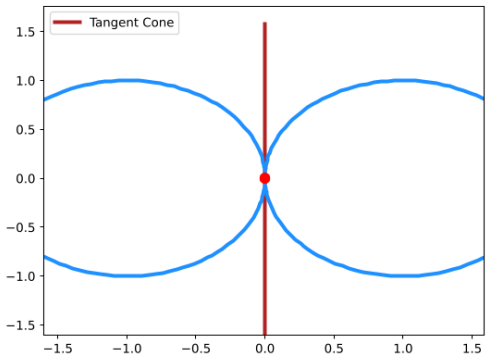

Figure 1.5. Two tangentially touching circles and the tangent cone at the origin

The tangent cones can not be determined directly from a sample but can be approximated by local information similar to what is used for the approach taken by [BWM12, SW14, FW16, Mil21] using local homology. For a space $W \subset \mathbb{R}^n$ one may consider the so-called *magnifications* (see Definition 3.1.3) given by

$$\mathcal{M}_x^r(W) := r(W - x) \cap \mathrm{B}_1(0).$$

at any point $x \in W$ (see also Fig. 1.3). For sufficiently well-behaved Whitney stratified spaces W such magnification converge to the extrinsic tangent cone for $r \to 0$ (when restricted to the unit ball). Such a result can, e.g., be found in [Hir69, BL07]. We prove a global version of this convergence result for the magnifications obtained from a close enough sample of W in Propositions 3.3.11 and 3.4.4.

Additionally what is needed is a measure to classify the local data contained in the magnifications at any given point into singular or regular. For this purpose we introduce continuous functions

$$\Phi : \{\text{local data}\} \to [0, 1]$$

giving a value close to 1 if the data is considered regular or a value close to 0 otherwise. The *local data* will usually be given by a magnification. Now, a Whitney stratified space W with two strata such that the singular stratum W_p agrees with exactly the set of points x at which $\Phi(\mathrm{T}_x^{\mathrm{ex}}(W)) < 1$ is called *(tangentially) Φ-stratified*. We discuss two specific examples of such functions Φ in detail. One is based on using local persistent homology (see Example 3.2.3 and Section 4.1.2) and the other is based on finding the minimal Hausdorff distance of a linear subspace of $\mathbb{R}^n$ to a given magnification (see Example 3.2.2 and Section 4.2.3). One of our main results (see

Theorem 3.5.8) is on the convergence behaviour in stratified Hausdorff distance of two strata samples $\mathcal{S}^r_{\Phi,u}(\mathbb{X})$ associated to a non-stratified samples $\mathbb{X}$ sampled from a sufficiently well-behaved Whitney stratified space. The singular stratum of $\mathcal{S}^\zeta_{\Phi,u}(\mathbb{X})$ is given by

$$\mathcal{S}^r_{\Phi,u}(\mathbb{X})_p = \{x \in \mathbb{X} \mid \Phi(\mathcal{M}^r_x(\mathbb{X})) \leq u\}$$

where u functions as a chosen threshold for the sensitivity of the function Φ. An illustration of such an object $\mathcal{S}^\zeta_{\Phi,u}(\mathbb{X})$ with Φ based on local persistent homology can be found in Fig. 1.3 on the right where $\mathcal{S}^r_{\Phi,u}(\mathbb{X})_p$ is marked red and $u = 0.87$, i.e. close to 1.

Theorem 1.2.1. *Let $W \subset \mathbb{R}^N$ be a compact and sufficiently well-behaved (see Definition 3.3.7) Whitney stratified space with two strata and underlying space denoted X, which is Φ-stratified with respect to a function Φ. Then there exists $u_0 \in (0,1)$ such that for all $u \in [u_0, 1)$ we have convergence*

$$\mathcal{S}^r_{\Phi,u}(\mathbb{X}) \to W$$

in stratified Hausdorff distance, for $r \to 0$ and $\frac{1}{r}d(\mathbb{X}, X) \to 0$.

In particular, this result can be applied to all compact, subanalytically Whitney stratified spaces with two strata. For the given space W it states that for any $\delta > 0$ there exists a radius R such that for all radii $r > R$ there exists $\varepsilon_r > 0$ such that

$$\max\{d_{Hd}(\mathcal{S}^r_{\Phi,u}(\mathbb{X}) = \mathbb{X}, X), d_{Hd}(\mathcal{S}^r_{\Phi,u}(\mathbb{X})_p, W_p)\} < \delta$$

for all $\mathbb{X}$ with $d_{Hd}(\mathbb{X}, X) < \varepsilon_r$. In other words, as $\mathbb{X}$ converges to X by assumption, the statement is really about the convergence of the singular parts $\mathcal{S}^r_{\Phi,u}(\mathbb{X})_p$ and W_p here.

Another thing to observe here is that we assume the samples $\mathbb{X}$ to converge to X faster than r converges to 0. This means, in order to increase the accuracy of the method by considering more and more local data we also have to increase the quality of the sample $\mathbb{X}$ in the sense that it gets closer to X in Hausdorff distance. To illustrate this coupled convergence consider the following case.

Example 1.2.2. We generated three sample from the Lemniscate, i.e. the space given by

$$V = \{x \in \mathbb{R}^2 \mid x_1^4 - x_1^2 + x_2^2 = 0\}.$$

First, for the sample $\mathbb{X}_1$ of the poorest quality, that is the Hausdorff distance of $\mathbb{X}_1$ to V is bounded by approximately 0.07, we can see the results for the approximated stratification first at $r = \frac{1}{3}$ in Fig. 1.6 and then at $r = \frac{1}{6}$ in Fig. 1.7.

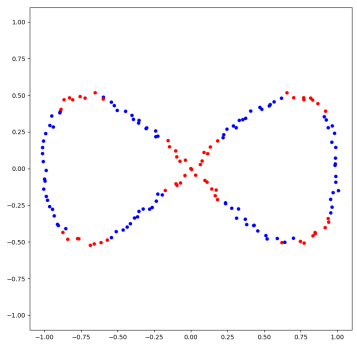

Figure 1.6. $\mathbb{X}_1$ with $\mathcal{S}^r_{\Phi,u}(()\mathbb{X}_1)_p$ marked red

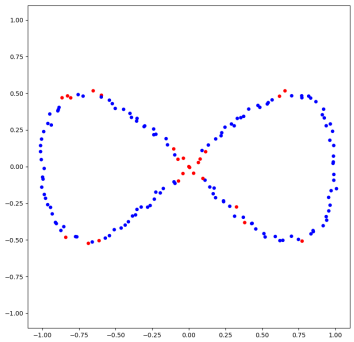

Figure 1.7. $\mathbb{X}_1$ with $\mathcal{S}^{\frac{1}{6}}_{\Phi,u}(\mathbb{X}_1)_p$ marked red

Here, we used a function Φ based on the Hausdorff distance to 1-dimensional linear subspaces of $\mathbb{R}^2$ (see Example 3.2.2). For a chosen $u = 0.4$ we observe that the radius $r = \frac{1}{3}$ is too large as we classify also points away from the actual singularity as singular, as this method is also susceptible to curvature. If we reduce the radius to $r = \frac{1}{6}$ we find no sensible classification anymore as the radius is too small for a coarse sample such as $\mathbb{X}_1$. Moving on, we improve the sample quality further, i.e. Hausdorff distance bounded by approximately 0.035. The results of our method for $\mathbb{X}_2$ with the radii $r = \frac{1}{6}$ and $r = \frac{1}{9}$ can be seen in Fig. 1.8 and Fig. 1.9 respectively.

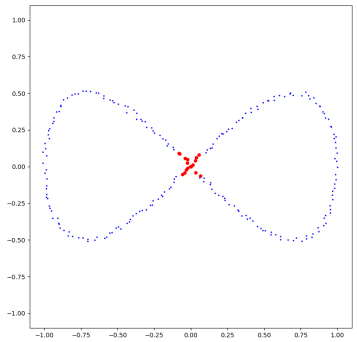

Figure 1.8. $\mathbb{X}_2$ with $\mathcal{S}^{\frac{1}{6}}_{\Phi,u}(\mathbb{X}_2)_p$ marked red

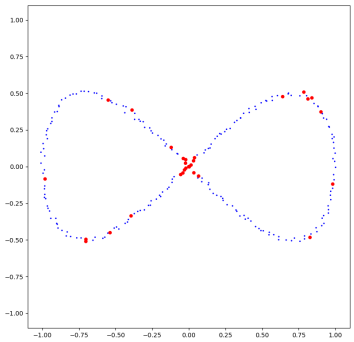

Figure 1.9. $\mathbb{X}_2$ with $\mathcal{S}^{\frac{1}{9}}_{\Phi,u}(\mathbb{X}_2)_p$ marked red

Although we only classifiy the region around the actual singularity of V, i.e. the origin, as singular we are not able to reduce the radius to $r = \frac{1}{9}$ to further improve the accuracy of the approximation for using $\mathbb{X}_2$. Lastly, our best sample $\mathbb{X}_3$ has Hausdorff distance below 0.023 from V. For this sample we are able to reduce the radius to $r = \frac{1}{9}$ and find a close approximation of the singular stratum of V, see Fig. 1.10.

We can also combine our results on stratification learning through a function Φ and our results on persistent stratified homotopy types to form a complete routine that

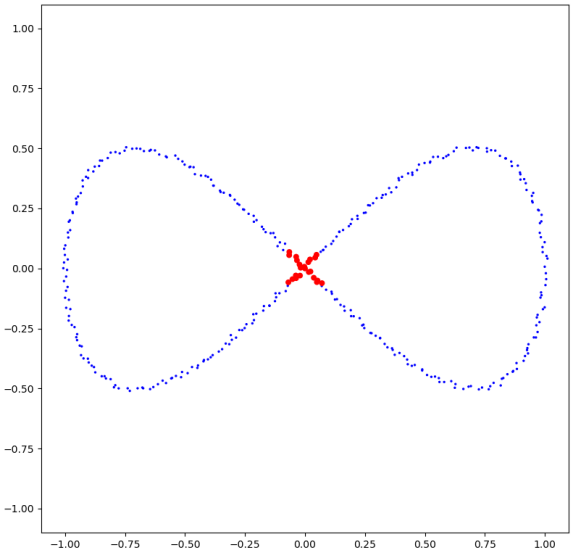

Figure 1.10. $\mathbb{X}_3$ with $\mathcal{S}^{\frac{1}{9}}_{\Phi,u}(\mathbb{X}_3)_p$ marked red

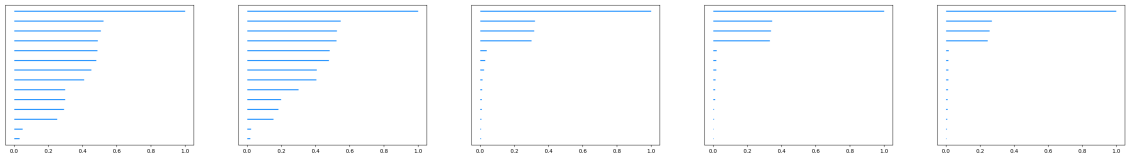

X_1 $r = \frac{1}{3}$ | X_1 $r = \frac{1}{6}$ | X_2 $r = \frac{1}{6}$ | X_2 $r = \frac{1}{9}$ | X_3 $r = \frac{1}{9}$

Figure 1.11. The 14 most persistent cycles of the 0-th homology of the link part of $\mathcal{P}_v(\mathcal{S}^r_{\Phi,u}(\mathbb{X}_i))$, for $i \in \{1,2,3\}$ and $r \in \{\frac{1}{3}, \frac{1}{6}, \frac{1}{9}\}$

takes a non-stratified sample as input and puts out a persistent stratified homotopy type (see Corollary 3.5.9). The convergence of persistent homotopy types associated to the non-stratified samples we just discussed in Example 1.2.2 can be illustrated by considering the persistent 0-homology of the the respective link parts, see Fig. 1.11. For X_2 and X_3 and radii $r = \frac{1}{6}$ and $r = \frac{1}{9}$ the persistent homology is close to the result expected for the link of the singular stratum of V (considered as stratified space).

1.3 Algorithmic Stratification Learning

While our initial foray into theoretical investigations was sparked by preliminary results on stratification learning from non-stratified point cloud data, employing a heuris-

tic approach grounded in local homology, the pivotal breakthrough occurred with the recognition of the capability to identify singularities through tangent cones. This discovery enabled us to establish rigorous conditions for various methods to theoretically guarantee the successful stratification of a point cloud. These conditions are articulated in Definition 3.2.1, defining Φ-stratifications, and in Examples 3.2.2 and 3.2.3, which present two principal candidates for Φ functions measuring singularity in practical applications.

Naturally, a purely theoretical exploration of this subject leaves the crucial question unanswered whether the proposed methods are genuinely useful for data analysis. Hence, we delve into a comprehensive discussion regarding the accuracy and implementation of two distinct algorithms designed to compute the two aforementioned Φ functions.

1.3.1 Local Persistent Homology

The general idea is based on the observation that on an n-manifold M^n local homology is constant, i.e. $H_i(M^n, M^n - p) \cong H_i(S^n)\ \forall p \in M^n$ which is due to the fact that a manifold is locally Euclidean. A stratified space such as a conically stratified space (see Recollection 2.1.17) does not possess this property and for the two strata case local homology can be used to identify the stratification (see Lemma 3.0.3).

As mentioned before, the idea of using local persistent homology for stratification learning is certainly not a new one (see e.g. [Mil21, STHN20, Nan20, SW14, BWM12, FW16]). However, what has not been done before is to describe a method that is provably able to detect the singular stratum of a two strata space. We intend to provide an algorithm that is readily implementable and applicable to real-world data to approximate such a method. For a space W, $x \in W$, radius r and some dimension d the function $\Phi^d_{\mathcal{P}\mathrm{L}}$ at the magnification $\mathcal{M}^r_x(W) = \frac{1}{r}(W - x) \cap \mathrm{B}_1(0)$ is given by

$$\Phi^d_{\mathcal{P}\mathrm{L}}(\mathcal{M}^r_x(W)) = 1 - 2 \max_{i \leq d} d(\mathcal{P}\mathrm{L}_i(\mathcal{M}^r_x(W)), \mathcal{P}\mathrm{L}_i(\mathbb{R}^d)).$$

where the distance on the right-hand side is given by the Bottleneck distance of persistence diagrams and $\mathcal{P}\mathrm{L}_i$ denotes local persistent homology at dimension i. There are two things that may not be immediately clear regarding the computability of $\Phi^d_{\mathcal{P}\mathrm{L}}$:

(*i*) How to compute the Bottleneck distance in this case?

(*ii*) How to compute local persistent homology based on a simplicial complex generated from point clouds?

The initial question finds a straightforward answer, as the Bottleneck distance has been explored to some extent, as evidenced by prior studies (cf., e.g., [EH10, KMN17]).

However, it is worth noting that this distance measure often involves significant computational complexity. In our specific scenario calculating the distance of a given barcode concerning the barcode provided by $\mathcal{PL}_i(\mathbb{R}^d)$ a notably simpler approach to compute the Bottleneck distance is available (refer to Remark 4.1.8).

While a response to Section 1.3.1 has been presented (compare with [SW14,BWM12, BB20]), these solutions predominantly hinge on Delaunay complexes or remain largely theoretical. In Definition 4.1.26, we introduce a singular filtered simplicial complex that, once established, facilitates the approximation of local persistent homology by applying the standard algorithm for computing persistent homology from a filtered simplicial complex. Notably, it shares a convenient feature with the Vietoris-Rips filtration, it can be constructed from the 1-skeleton. Our construction builds upon the groundwork laid by Skraba and Wang [SW14].

We provide pseudocode for a potential implementation of the filtered simplicial (see Algorithms 1, 3 and 4). Leveraging this, we present a comprehensive routine that performs all necessary computations for our version of a Φ function, based on the concept of using local persistent homology as in $\Phi^d_{\mathcal{PL}}$. This pseudocode is encapsulated in Algorithm 5. In Section 4.3.2 and Chapter 5, we furnish several practical examples to demonstrate the utility of our method for stratification learning.

One might choose to compute via a Čech complex, and we briefly discuss this option (see Remark 4.1.23), providing an avenue to directly compute the function $\Phi^d_{\mathcal{PL}}$ via simplicial complexes without approximation error. However, in practice, opting for a Vietoris-Rips type complex is prevalent for computational efficiency, as it can be constructed efficiently from its 1-skeleton (see [Zom10]) and we therefore also chose to focus on this. However, the drawbacks of Vietoris-Rips type complexes are also acknowledged, arising from not being homotopy equivalent to the Čech complex constructed from the same point cloud. We delve into this comparison, particularly in the worst-case scenario, summarized in the following statement (see Proposition 4.1.29).

Proposition 1.3.1. *Let $X, \mathbb{X} \in \mathbf{Sam}_\star$ with $\mathbb{X}$ a δ-sample of X. Then, for $\Phi_{\mathcal{PL}}(\mathbb{X}, d)$ denoting the output of Algorithm 5 and $\Phi^d_{\mathcal{PL}}$ the function described in Example 3.2.3 we have*

$$| \Phi_{\mathcal{PL}}(\mathbb{X}, d) - \Phi^d_{\mathcal{PL}}(X) | \leq 2\delta + \frac{1}{2}(\sqrt{\frac{2(d+1)}{d+2}} - 1).$$

1.3.2 Hausdorff Distance

Although the basic idea to use Hausdorff distances for stratification learning is similar to the observation that motivated the use of local homology for this purpose, the second function Φ we studied originated from the developed theory itself. Again, to a

manifold we can associate a tangent bundle whose fibers are given by Euclidean spaces of uniform dimension. For spaces that do not admit a well-defined tangent bundle one can still consider the extrinsic tangent cone at every point. Loosely speaking, points at which the tangent cone is far from being a linear subspace of some $\mathbb{R}^n$ will be considered singular. The distance is measured in Hausdorff distance. This also explains the notion of a tangentially stratified space to some extend. For $W \subset \mathbb{R}^n$ a q-dimensional Whitney stratified space, the function is given by

$$\Phi^d_{Hd}(\mathcal{M}^r_x(W)) = 1 - \min d(\mathcal{M}^r_x(W), V)$$

where $\mathcal{M}^r_x(W)$ is local data given by the magnification at $x \in W$ and we minimize over all d-dimensional linear subspaces V of $\mathbb{R}^n$. Here the Hausdorff distance is measured within the Euclidean unit ball in $\mathbb{R}^n$. Note that the magnification tends towards the tangential cone at x for W a sufficiently well-behaved Whitney stratified space (see Definition 3.3.7). Rephrasing Φ^d_{Hd} leads to the following optimization problem

$$\inf_{V \in \mathrm{Gr}(k,n)} \max\{\sup_{x \in X} d(x, V), \sup_{v \in V, ||v|| \leq 1} d(X, v)\}. \tag{1.1}$$

for a given topological space X. This prompts two apparent challenges:

(*i*) Computing the Hausdorff distance is computationally expansive or even intractable considering general and possibly infinite X and V.

(*ii*) With $\mathrm{Gr}(k, n)$ as search space we have to consider a matrix representation of the Grassmannians to reformulate the problem or use a randomized approach.

Primarily, our focus is directed towards scenarios where X represents a finite subset within a Euclidean space, specifically confined within the standard unit ball. To render the reformulation of 1.1 deterministic, we opt for a simplification by exclusively considering the one-sided Hausdorff distance. This choice results in a reduction of the class of Φ-stratified spaces, employing a simpler Φ-function. Additionally, we adopt a matrix representation of the Grassmannians to formulate a smooth constrained optimization problem, as detailed in Eq. (4.18).

In our randomized approach, we leverage a uniform distribution on the Grassmannians. The Hausdorff distance is approximated by generating controlled samples of linear subspaces $V \in \mathrm{Gr}(k, n)$ chosen randomly. Once again, we present a potential implementation of our approximation of Φ^d_{Hd} in the form of pseudocode in Algorithm 7. The outcome of our approximation for Φ^d_{Hd} (see Theorem 4.2.5) is stated as follows:

Theorem 1.3.2. *Let $X, \mathbb{X} \in \mathbf{Sam}_\star$ with $\mathbb{X}$ a δ-sample of X. Let γ be the sample accuracy for the linear subspaces. Then, for any $0 < \varepsilon \leq 1$, $s \in \mathbb{N}_+$ and $\Phi_{Hd}(\mathbb{X}, d)$*

denoting the output of Algorithm 7 we have that

$$| \Phi_{Hd}(\mathbb{X}, d) - \Phi^d_{Hd}(X) | \leq \varepsilon + \delta + \gamma$$

with probability $1 - (1 - \mathbb{P}(\theta < \arcsin \varepsilon))^s$.

In this context, $\mathbb{P}(\theta < \arcsin \varepsilon)$ signifies the probability of selecting a random $V \in \mathrm{Gr}(k, n)$ that is ε close, measured by the Hausdorff distance, to the optimal linear subspace. Moreover, the parameter s denotes the number of samples drawn, and its determination is contingent on the desired probability. We expound on the selection of the sample size to achieve a specified level of accuracy for the approximation of Φ^d_{Hd} at a specific point (refer to Remark 4.2.4) and for a comprehensive global approximation (refer to Remark 4.2.6).

1.4 Applications in Image Patch Spaces

The preceding literature primarily concentrated on the methodology employed to extract topological insights from spaces of image patches derived from extensive collections of natural images, exemplified by the van Hateren collection [vHvdS98]. An additional layer of statistical information or averaging on image patches was introduced by extracting the upper percentage of the densest subsets. These resulting spaces were then posited as representations of the most prevalent patches in natural images, leading to the proposition of various models for such spaces. Among the proposed models, the *three circle model*, initially introduced in [SC04], stood out as a stratified space with two strata. In our investigation, we revisited these image patch spaces, employing more refined invariants, available to us from our developed theory, to learn their stratification. Subsequently, we analyzed the persistent homology of the resulting regular, singular, and link components, providing compelling evidence that strongly supports the validity of the three circle model. See also Fig. 1.12 for a visualization of the image patch space generated for our analysis to support the three circle model.

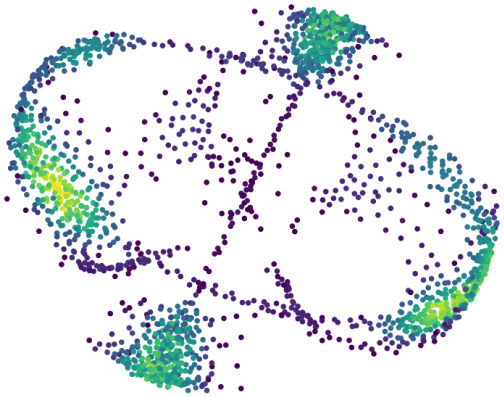

Figure 1.12. Visualization via PCA of the three circle image patch space $\mathbb{X}$ with coloring by results of $\Phi_{\mathcal{P}\mathrm{L}}(\mathbb{X}, 1)$

Furthermore, we deviated from examining image patch spaces generated from an array of images and shifted our focus towards the topology of individual image patch spaces, specifically those constituted by the most frequently occurring patches. Naturally, this undertaking is susceptible to failure if applied to a diverse array of images. To address this, we restricted our analysis to images with reduced structure, specifically a particular type of image patch known as *(ideal) step-edge patches* (refer to Definition 5.1.2). Images constructed exclusively from such step-edge patches, and characterized by a clear topological and stratified structure, are termed *cluster images* (refer to Definition 5.1.3). In broad terms, these images lack intricate textures and are described by dominant colors and the transitions between them. The resulting topology can be conceptualized as a union of 2-spheres, each capable of intersecting another 2-sphere at a single point—illustrated in Fig. 1.13.

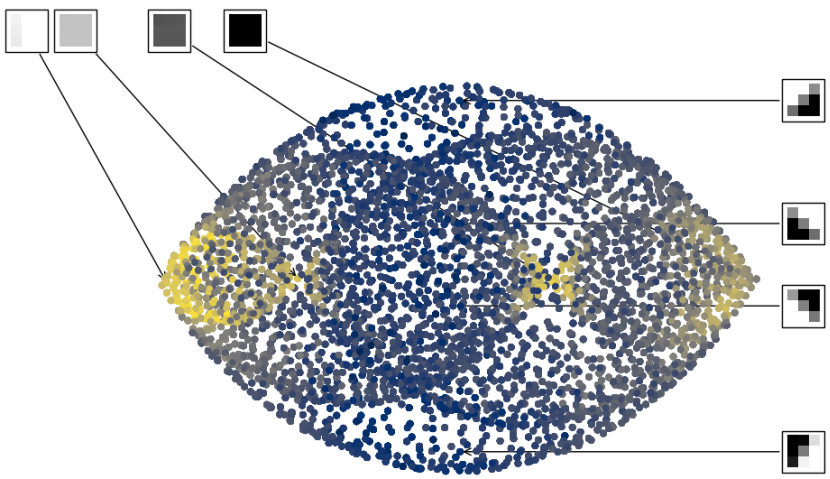

Figure 1.13. Image patch spaces $\mathbb{X}$ extracted from a cluster image with coloring by values of $\Phi_{\mathcal{PL}}(\mathbb{X}, 2)$

We used this observation to conduct an analysis that indicates a measurable topological difference between different types of images. The selected image groups were visually very different. One group showing landscapes and the other group of images showing people and most of the time the upper body. Both groups of images were taken from the CIFAR-100 dataset [KH09]. We prepared the data by extracting image patches from every image and projecting each patch to the closest step edge patch (see Preprocessing). For the resulting spaces we computed ordinary persistent homology and compared the most persistent cycles in dimension 1 and 2 (see Analysis). We can report a clear distinction of the groups of images based on these topological invariants.

Finally, we aim to illustrate another application scenario wherein our pipeline, which associates a persistent stratified homotopy type to a non-stratified point cloud, can be employed to unveil structures in real-world data. To achieve this, we considered a dataset comprising manually segmented retina photos. Specifically, the dataset encompasses 15 healthy retina images and 15 retinas from glaucoma patients [BBM$^+$13], as exemplified in Fig. 1.14. Our approach involved approximating stratifications for the 1×1 image patch spaces derived from each image, employing $\Phi_{\mathcal{PL}}(-, 1)$. In essence, we treated the pixels in the images as data points. An illustration of how the values of $\Phi_{\mathcal{PL}}(-, 1)$ quantify the singularity of areas in the artery images is provided in Fig. 1.15.

We can discern individuals with glaucoma from healthy test participants with a high degree of reliability using manually segmented retina images acquired through a fundus camera. The discriminative factor lies in the number of points in $\mathbb{X}_i^*$ for which

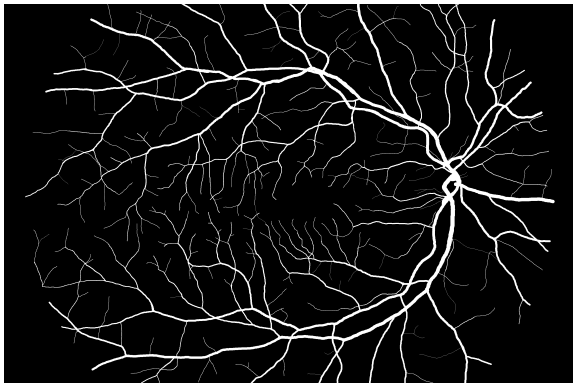

Figure 1.14. Example of a manually segmented retina photo

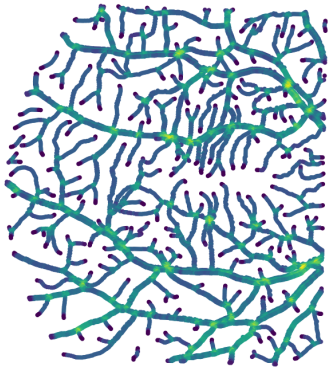

Figure 1.15. $\Phi_{\mathcal{P}\mathrm{L}}(-,1)$ of a pixel patches

$\Phi_{\mathcal{P}\mathrm{L}}(-,1)$ exceeds a specified threshold (refer to Table 5.1), as well as the quantity of persistent components within the determined singular parts of the patch space (refer to Table 5.2). A more detailed contextualization of these findings is provided in the corresponding section (refer to Section 5.2).

Although a complete answer to the topology of image patch spaces remains elusive, the project has provided insights into the topology of patch spaces. By devising appropriate methods to identify singular and regular points in the dataset, we have been able to analyze and understand the underlying stratified structure more effectively. In essence, we have hopes that our newly developed methods have the potential to unlock opportunities for advancing research in comprehending the topological and especially stratified structures of image patch spaces and other types of data.

Our contributions in this project, particularly in the collaborative effort with Lukas Waas, have evolved from the synergies between theoretical insights and practical exploration. These endeavors were facilitated and bolstered by the initiatives encapsulated within the comprehensive projects 6 and 7 of the excellence cluster STRUCTURES, emphasizing the integration of theoretical foundations and practical applications to advance our understanding and methodologies in this research domain.

Chapter 2

Persistent Stratified Homotopy Types

In this chapter, we provide an overview of stratified spaces and their homotopy theory as it pertains to our investigations. The presentation of stratified homotopy theory is aimed to be understandable for readers familiar with fundamental concepts in algebraic topology and category theory. For a more comprehensive understanding of both the details and the complete model categorical perspective, we direct the reader to [DW21, Dou21a, Dou21b]. Additionally, one may use sources like [Ban07] to find more in-depth information on stratified spaces and associated invariants. The results outlined in this chapter were featured in a collaborative article by the author and Lukas Waas [MW22]. This chapter approximately aligns with [MW22, Sections 2 and 3], and certain results and definitions are incorporated verbatim.

2.1 Stratified Homotopy Theory

We begin by reviewing some essential concepts that are relevant to the theory of stratified spaces.

Definition 2.1.1. A *stratified space* (over a poset P) is a pair $S = (X, s : X \to P)$ where X is a topological space and s is continuous with respect to the Alexandrov topology on P. The map s is called the *stratification of* S. The fiber of the stratification over $p \in P$

$$S_p := s^{-1}\{p\} = \{x \in X \mid s(x) = p\}$$

is called the *p-stratum of* S. The sub- and superlevel sets of s will be denoted by $S_{\leq p}$ and $S_{\geq p}$ respectively.

The above definition of a stratified space is rather general in comparison to e.g. more classical works on stratified spaces (as in e.g. [GM80, GM83]). According to the definition given here, simply any filtration $(X_{\leq 0} \subset ... \subset X_{\leq n} = X)$ by closed subsets of a topological space X induces a stratification over the poset $[n] = \{0 < ... < n\}$ with the k-stratum given by $X_k = X_{\leq k} \setminus X_{\leq k-1}$ for $k \in [n]$. This makes it a trivial task to equip a topological space with a stratification. Therefore, let us take a look at some of the places where stratified spaces arise in a more interesting, non-trivial way.

Example 2.1.2.

- Let $(M, \partial M)$ be a compact manifold with boundary and let X be the space obtained by coning off the boundary of M, i.e. $X = M \cup_{\partial M} C(\partial M)$, where $C(Y)$ denotes the cone on a space Y. One obtains a stratification of X by the map

$$s\colon X \to \{p < q\}; \qquad \begin{cases} x \mapsto q, & \text{for } x \in X \setminus \{\text{cone point}\}, \\ x \mapsto p, & \text{for } x = \text{cone point}. \end{cases}$$

  The resulting stratified space is locally euclidean away from one isolated singularity, at which arbitrarily small neighborhoods are homeomorphic to the open cone $\mathring{C}(\partial M)$.

- Let $S = (X, s\colon X \to P)$ be a stratified space and let $c(P) = \{*\} \cup P$ with $* < p$ for all $p \in P$. The (stratified) cone on X, $C(X)$, is naturally a stratified space over $c(P)$ by setting

$$X \to \{p < q\}; \qquad \begin{cases} [(x,t)] \mapsto s(x), & \text{for } t > 0, \\ [(x,t)] \mapsto *, & \text{for } t = 0\ . \end{cases}$$

  We always take cones to be equipped with the teardrop topology [Qui88, Definition 2.1]. Note, however, that for compact Hausdorff spaces, the latter agrees with the usual quotient space topology on the cone.

- Given a smooth manifold M with a compact Lie group G acting smoothly and properly on M. A non-trivial stratification for the orbit space M/G can then be obtained by stratification by orbit types (see, e.g., [Pfl01, Chapter 4] for more details).

- Any n-dimensional complex algebraic variety X can be equipped with a geometrically meaningful stratification by taking the the filtration by closed subsets given by iteratively taking singular loci of the variety.

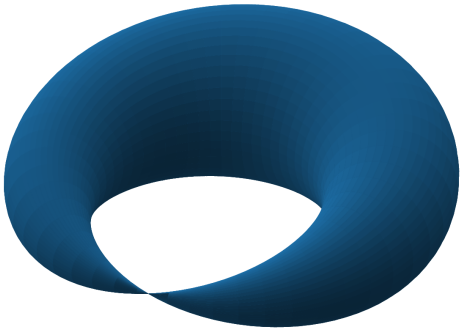

Figure 2.1. Pinched torus

- Another very important class of stratified spaces are so-called *pseudomanifolds*. The class of pseudomanifolds is much more restictive, e.g. the strata of a pseudomanifold have to be actual manifolds . Pseudomanifolds have been the subject of research to recover a form of generalized Poincaré duality for singular spaces ([GM80,GM83]). One specific example of a pseudomanifold is the pinched torus PT^2 which we illustrate in Fig. 2.1. Mathematically, PT^2 can be described as the quotient space of the torus $T^2 = S^2 \times S^1$ by collapsing one circle $* \times S^1$ to a point. The structure of a stratified space is then induced by taking the filtration $\{s\} \subset PT^2$ where s denotes the image of the collapsed circle in PT^2.

Definition 2.1.3. A *stratified map* between two stratified spaces $X \to P$ and $Y \to Q$ is a commutative square of continuous maps

$$\begin{array}{ccc} X & \xrightarrow{f} & Y \\ \downarrow & & \downarrow \\ P & \xrightarrow{\overline{f}} & Q\,. \end{array}$$

For ease of exposition we refer to such a diagram by simply writing f. In case $\overline{f}$ is the identity on posets, we call f *stratum preserving*.

Definition 2.1.4. Let $f, f' \colon S = (X, s \colon X \to P) \to S'$ be stratified maps. We call f and f' *stratified homotopic*, if there exists a stratified map

$$\mathrm{H} \colon (X \times [0,1], X \times [0,1] \to X \xrightarrow{s} P) \to S'$$

such that $\mathrm{H}_{|X\times\{0\}} = f$ and $\mathrm{H}_{|X\times\{1\}} = f'$. Furthermore, f is called a *strict stratified homotopy equivalence*, if there exists another stratified map $g\colon S' \to S$ such that $f \circ g$ and $g \circ f$ are stratified homotopic to $\mathrm{id}_{S'}$ and id_S respectively.

Remark 2.1.5. The above definition of strict stratified homotopy equivalence agrees with the commonly used definition of stratified homotopy equivalence which is a rather rigid notion. The rigidity, further demonstrated in Example 2.1.6, is one of the issues one has to be aware of when working with stratified spaces in topological data analysis. It will be one of the main goals of this chapter to define another concept of stratified homotopy equivalence which is less rigid and therefore suitable for the use in topological data analysis, as we will see later. Because we will have to refer to different notions equivalences of stratified spaces, we use the convention of speaking of *strict stratified homotopy equivalences* instead of stratified homotopy equivalences. The class of all stratified spaces *strictly stratified homotopy equivalent* to a stratified space S is called the *strict stratified homotopy type of* S.

Example 2.1.6. Consider the space $X = S^1 \vee S^1$ embedded in $\mathbb{R}^2$ as a curve, shown in Fig. 2.2. It features a singular point at the self-crossing. Denote the resulting stratified space over $P = \{0 < 1\}$ with the singularity sent to 0 and the remainder to 1 by S. While there generally seems to be no canonical way to thicken such a space, one possibility is to thicken both the total space as well as the singularity as in Fig. 2.4. The resulting thickened space S'' is strictly stratified homotopy equivalent to the original curve with the singular stratum extended from a point to the crossing, denoted S', see Fig. 2.3. To sketch one way to see this, note that the original curve separates the space S'' into two inner parts and an outer part, decomposing S'' into three parts, all homeomorphic to a space $S^1 \times [0, 1]$. One then constructs respective homotopy equivalences of $S^1 \times [0, 1]$ to S^1 relative to the sections of the curve included in the three parts that we can compatibly glue to a homotopy equivalence between S'' and S'. However, S and S' (and hence S'') are not strictly stratified homotopy equivalent. To see this, note that a stratified homotopy equivalence between S and S' would also have to be a homotopy equivalence of the underlying spaces. Such a map has to send a circle S^1 with degree ± 1 onto another circle. But the image of any stratified map between S and S' is (non-stratifiedly) contractible.

Notation 2.1.7. Now that we have defined objects and maps let us fix some notation. Stratified spaces together with stratified maps define a category which we will denote by **Strat**. For the category of stratified spaces over a poset P and stratum preserving maps we write $\mathbf{Top}_P$. Isomorphisms in $\mathbf{Top}_P$ - i.e. stratum preserving homeomorphisms - will be denoted by $\cong_P$.

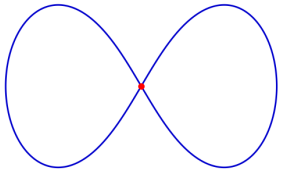

Figure 2.2. Singular curve, S

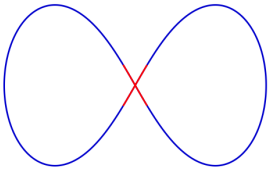

Figure 2.3. Thickening 1, S'

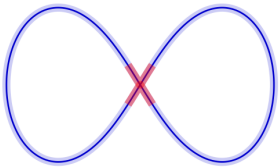

Figure 2.4. Thickening 2, S''

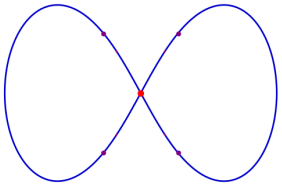

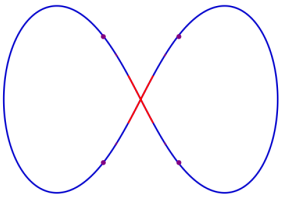

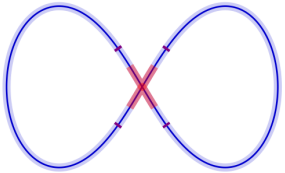

Figure 2.5. Geometric models of homotopy links marked in purple

Definition 2.1.8. Let S be a stratified space and $p, q \in P$ with $p < q$. The *homotopy link* of the p-stratum in the q-stratum is the space of so called *exit paths*

$$\mathrm{hoLink}_{\mathrm{p<q}}(S) = \{\gamma \colon [0,1] \to X \mid \gamma(0) \in S_p, \gamma(t) \in S_q, \forall t > 0\}$$

with its topology induced by $\mathrm{hoLink}_{\mathrm{p<q}}(S) \subset C^0([0,1], X)$, where the latter denotes the space of continuous functions, equipped with the the compact open topology. The induced functors

$$\mathbf{Top}_P \to \mathbf{Top}$$

come with natural transformations

$$S_p \leftarrow \mathrm{hoLink}_{\mathrm{p<q}}(S) \to S_q,$$

given by the endpoint and beginning point evaluation map.

Example 2.1.9. Let us return to Example 2.1.6 to give an illustration of the homotopy link. For the original singular curve and both thickenings, the homotopy links are all homotopy equivalent to four isolated points (see Fig. 2.5). This can be seen from Construction 2.1.36, which states that the homotopy links are homotopy equivalent to the boundary of a cylinder neighborhood of the singular stratum.

Definition 2.1.10. A stratum preserving map $f \colon S \to S'$ in $\mathbf{Top}_P$ is called a *weak equivalence of stratified spaces*, if it induces weak equivalences of topological spaces

$$\mathrm{hoLink}_{\mathcal{I}}(S) \to \mathrm{hoLink}_{\mathcal{I}}(S'),$$

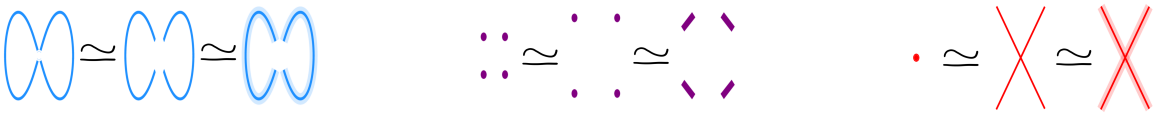

Figure 2.6. Regular strata, homotopy links and singular strata of the spaces in Example 2.1.6

for all regular flags $\mathcal{I} \subset P$. A stratified map $f\colon S = (X, s\colon X \to P) \to S'$ is called a *weak equivalence of stratified spaces*, if $\overline{f}$ is an isomorphism, and f is a stratum preserving weak equivalence in $\mathbf{Top}_P$, after identifying the target poset with P through $\overline{f}$.

Notation 2.1.11. We denote by ho$\mathbf{Top}_P$ and ho$\mathbf{Strat}$ the categories obtained by localizing $\mathbf{Top}_P$ and $\mathbf{Strat}$ respectively at the class of weak equivalences. The isomorphism class of $S \in$ ho$\mathbf{Strat}$ is called the *stratified homotopy type of* S. Isomorphisms in ho$\mathbf{Top}_P$ will be denoted by $\simeq_P$.

It is an immediate consequence of the fact that homotopy links map stratified homotopy equivalences to homotopy equivalences, that any strict stratified homotopy equivalence is also a weak equivalence of stratified spaces. The converse is generally false.

Example 2.1.12. Let us illustrate these concepts for the spaces from Example 2.1.6 where we already discussed that there is no strict stratified homotopy equivalence between the original curve and any of the described thickenings. However, all the spaces are weakly stratified homotopy equivalent. Indeed, this is already hinted at by the fact that we may find a homotopy equivalence between the respective regular and singular parts as well as the homotopy links as described in Example 2.1.9. Consider Fig. 2.6 for an illustration.

In the context of topological data analysis it would seem advantageous to work with a slightly weaker notion of equivalence of stratified spaces where the objects discussed in Example 2.1.12 are properly equivalent. In order to identify spaces for which the weaker equivalence agrees with the strict notion we employ the work of [DW21] on stratified homotopy theory. But first, we need to recall some further definitions from the field of stratified spaces.

Notation 2.1.13. Given a stratified space $S = (X, s : X \to P)$ and $p \in P$ we write

$$\begin{aligned} S_{\leq p} &:= s^{-1}(\{q \leq p\}), \\ S_{<p} &:= s^{-1}(\{q < p\}), \\ S^{\geq p} &:= s^{-1}(\{q \geq p\}), \\ S^{>p} &:= s^{-1}(\{q > p\}). \end{aligned}$$

For many theoretical as well as for our more applied investigations of stratified spaces, it is fruitful to impose additional regularity assumptions on the strata (such as manifold assumptions) and the way they interact. The notion central to this paper is the notion of a Whitney stratified space. These are characterized by the convergence behavior of secant lines around singularities.

Notation 2.1.14. Given two vectors $v, u \in \mathbb{R}^{\mathbb{N}}$ such that $u \neq v$. We denote by $l(v, u)$ the 1-dimensional subspace of $\mathbb{R}^N$ spanned by $v - u$.

Recollection 2.1.15. A stratified spaces $W = (X, s\colon X \to P)$ with $X \subset \mathbb{R}^N$ locally closed is called a *Whitney stratified space*, if it fulfills the following properties.

1. *Local finiteness*: Every point $x \in X$ has a neighborhood intersecting only finitely many of the strata of W.

2. *Frontier condition:* W_p is dense in $W_{\leq p}$, for all $p \in P$.

3. *Manifold condition:* W_p is a smooth submanifold of $\mathbb{R}^N$, for all $p \in P$.

4. *Whitney's condition (b):* Let $p, q \in P$ such that $p < q$ and x_n, y_n be sequences in W_q and W_p respectively, both converging to some $y \in W_p$. Furthermore, assume that the secant lines $l(x_n, y_n)$ converge to a 1-dimensional space $l \subset \mathbb{R}^N$ and that the tangent spaces $\mathrm{T}_{x_n}(W_q)$ converge to a linear subspace $\tau \subset \mathbb{R}^N$. Then $l \subset \tau$. (By convergence of vector spaces we mean convergence in the respective Grassmannians.)

Example 2.1.16. Whitney's work ([Whi65a], [Whi65b]) states that every algebraic and analytic variety admits a Whitney stratification. More general, Whitney stratifications can even be given to spaces such as semi-analytic sets (see e.g. [Ło65]) or o-minimally definable sets (see e.g. [Loi98]). Finally, if X is such that it has only isolated singularities and admits a Whitney stratification, then any stratification of X, fulfilling frontier and boundary condition, with smooth strata is automatically a Whitney stratification. In particular, any definable set with isolated singularities and a dense open submanifold is canonically Whitney stratified with two strata. Another

class of Whitney stratified spaces arises from G-manifolds, already noted in Example 2.1.2. For a proof, see [Pfl01, Theorem 4.3.7].

Whitney's condition (b) has a series of immanent topological consequences, which ultimately led to the more general notion of a conically stratified space. The latter are (with some additional assumptions) one of the main objects of interest in the algebro-topological study of stratified spaces [Sie72, GM80, GM83, Qui88, Lur17].

Recollection 2.1.17. [Lur17, Def. A.5.5] We call a stratified space $S = (X, s\colon X \to P)$ *conically stratified*, if for every point x that lies in some stratum S_p for $p \in P$ there exists

- a stratified space $(L, L \mapsto P_{>p} = \{q \in P \mid q > p\})$,
- an open neighborhood $U \subset X$ containing x,
- a space Z
- and a stratified homeomorphism $Z \times C(L) \cong U$ with the posets $c(P_{>p}) \cong P_{\geq p}$ identified.

In addition to the Whitney stratification assumption, we will frequently need additional control over how pathological the subsets of euclidean space we allow for can be. To obtain such additional control, we use the notion of a set $X \subset \mathbb{R}^N$, definable with respect to some o-minimal structure (see [vdD98] for a definition). For the reader entirely unfamiliar with these notions it suffices to know that all semialgebraic or compact subanalytic sets have this property. On the one hand, definability assumptions guarantee the existence of certain mapping cylinder neighborhoods (see Example 2.1.33) that allow thickenings that do not change the homotopy type (see Lemma 2.2.17). At the same time, asserting additional control over the functions defining a set (polynomially bounded), has several consequences for the convergence behavior of tangent cones, already noted in [Hir69, BL07]. We will use these to recover stratifications from samples in accordance to Section 3.5.

Definition 2.1.18. We say that a stratified space $S = (X, s\colon X \to P)$, with $X \subset \mathbb{R}^N$ and P finite, is *definable* (or *definably stratified*) if all of its strata are definable with respect to some fixed o-minimal structure.

Moreover, a stratified space is called triangulable, if it admits a triangulation compatible with the stratification (for details see [DW21]). What suffices for us to know is that Whitney stratified and (locally compact) definably stratified spaces even

admit a PL-structure compatible with the stratification and are thus triangulable, see [Gor78], [Shi05], [Cza12].

We have now all the notation at hand to state a theorem by Douteau and Waas:

Theorem 2.1.19. *[DW21, Theorem 1.2] Let* $\mathbf{Con}_P \subset \mathbf{Top}_P$ *be the full subcategory of triangulable conically stratified spaces over* P*, and* $\simeq$ *be the relation of stratified homotopy. Denote by* $\mathbf{Con}_P/_{\simeq}$ *the category obtained by identifying stratified homotopic morphisms in* $\mathbf{Con}_P$*. Then the induced functor*

$$\mathbf{Con}_P/_{\simeq} \to \mathrm{ho}\mathbf{Top}_P$$

is a fully faithful embedding.

An analogous statement holds for the case of ho**Strat**.

Loosely speaking, the above theorem identifies a class of stratified spaces for which no homotopic information is lost by considering the stratified homotopy type over the strict stratified homotopy type. That means we can study sufficiently regular spaces, e.g. Whitney stratified spaces, up to their strict stratified homotopy type with methods of topological data analysis that are well-behaved w.r.t. the stratified homotopy type. What we mean by well-behaved is captured in our results Proposition 2.2.16 and Theorems 2.2.32 and 3.5.8. What remains for now is to establish how to systematically thicken a stratified space without changing the (strict) stratified homotopy type. We want to illustrate the problem with the following example:

Example 2.1.20. In Fig. 2.7 we exhibit three different thickenings of the original space from Example 2.1.6. The first thickening is neither weakly nor strictly stratified homotopy equivalent to the original curve (as can be seen by comparing homotopy links). The second thickening, being only weakly equivalent to the unthickened space, was discussed in Example 2.1.12. However, note that the inclusion of the original curve into it is not a stratified map. Hence, this notion of thickening does not allow for a persistent approach. For the third thickening, the inclusion of the original curve is even a strict stratified homotopy equivalence. However, it seems unclear how to systematically achieve such a thickening, particularly when working with samples.

In analogy to the classical scenario, this should assign to a stratified space $S \subset \mathbb{R}^N$, a functor from the posetal category of (positive) reals with the usual order $\mathbb{R}_+$ into some category representing stratified homotopy types $\mathcal{C}$. In the classical scenario, $\mathcal{C}$ is often taken to be the homotopy category of simplicial complexes (or sets) using constructions such as the Čech or Vietoris-Rips complex. For now, let us refer to the image under such a functor $\mathcal{P}(S)$ as the *persistent stratified homotopy type* of S, and

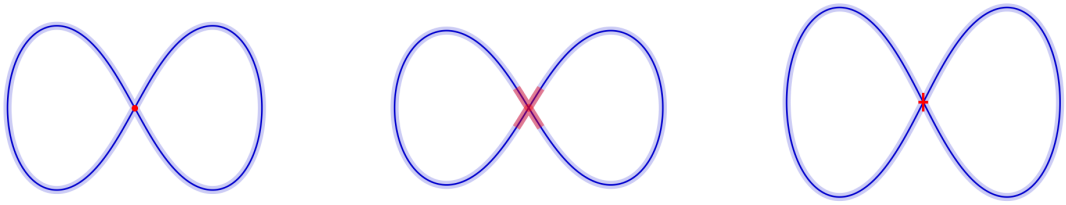

Figure 2.7. Three possible thickenings

similarly to the non-stratified construction using thickenings or Čech complexes as the *persistent homotopy type*. In order to make thickenings of stratified spaces in a systematic way we represent stratified homotopy types by what are called *stratification diagrams* as described in [DW21].

Definition 2.1.21. We denote by $\mathrm{R}(P)$ the category with objects given by regular (i.e. strictly increasing) flags $\mathcal{I} = \{p_0 < \cdots < p_k\}$ in P and morphisms given by inclusion relations of flags. We denote by

$$\mathbf{Diag}_P := \mathrm{Fun}(\mathrm{R}(P)^{\mathrm{op}}, \mathbf{Top})$$

the category of $\mathrm{R}(P)^{\mathrm{op}}$ indexed diagrams of topological spaces. We call elements of $\mathbf{Diag}_P$ *stratification diagrams*.

Definition 2.1.22. A morphism $f\colon D \to D'$ in $\mathbf{Diag}_P$, for which $f_{\mathcal{I}}$ is a weak equivalence, at all $\mathcal{I} \in \mathrm{R}(P)$, is called a *weak equivalence of stratification diagrams*.

We denote by $\mathrm{ho}\mathbf{Diag}_P$ the category obtained by localizing $\mathbf{Diag}_P$ at weak equivalences of diagrams. We arrived now at the main feature (at least to us here) of stratification diagrams and that is we can equivalently describe a stratified homotopy type by stratification diagrams. This is due to the following result, which is discussed in more detail [Dou19, DW21]:

Recollection 2.1.23. . (Generalized) homotopy links induce a functor

$$\begin{aligned} \mathrm{D}_P\colon \mathbf{Top}_P &\to \mathbf{Diag}_P \\ S &\mapsto \{\mathcal{I} \mapsto \mathrm{hoLink}_{\mathcal{I}}(S)\}. \end{aligned}$$

By definition, a stratum preserving map is a weak equivalence, if and only if its image under $\mathbf{Diag}_P$ is a weak equivalence. In particular, one obtains an induced functor

$$\mathrm{D}_P\colon \mathrm{ho}\mathbf{Top}_P \to \mathrm{ho}\mathbf{Diag}_P$$

which turns out to be an equivalence of categories. In this sense, the stratification diagram encodes the same homotopy theoretic information as the original space. We

will use this equivalence to identify these two homotopy categories and often not distinguish between a stratified space and its stratification diagram.

So far we collected several results from the field of stratified homotopy theory to describe a stratified space, up to weak equivalences, by another object, the stratification diagram, for which we claimed that it will allow us to define thickenings of stratified spaces that do not change the stratified homotopy type given sufficiently regularity conditions on the stratified space. But the reader may doubt that homotopy links (and thus also stratification diagrams) defined as subspaces of mapping spaces are suited to a computational or algorithmic approach. For this case we included a geometric description that will allow us to model a stratification diagram combinatorially. We will also use another equivalent description of stratified homotopy types, which occur naturally, particularly when trying to quantitatively recover stratifications from non-stratified data in Section 3.5. For this, observe that a poset can naturally be considered a simplicial complex by taking the elements of a poset P as vertices and its flags as simplices. Since we mostly consider the two strata case, we will only consider $P = \{p < q\}$ for the rest of this section and all the definitions therein will be given for this case. For this specific case the poset P regarded as a simplicial complex is canonically homeomorphic to the unit interval $[0, 1]$, where p corresponds to 0 and q corresponds to $(0, 1]$. We will use this observation to make the following definition.

Definition 2.1.24. A *strongly stratified space* (over $P = \{p < q\}$) is a pair

$$S = (X, s : X \to [0, 1])$$

where X is a topological space and s is continuous. A *strongly stratum preserving map* $f : S = (X, s) \to (X', s') = S'$ is a map of topological spaces $f : X \to X'$ making the diagram

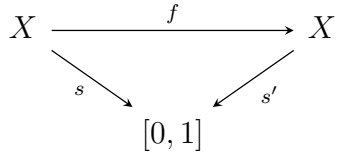

commute.

Notation 2.1.25. We denote by $\mathbf{Top}_{\mathrm{N}(P)}$ the category with objects given by strongly stratified spaces and morphisms given by strongly stratum preserving maps. Isomorphisms in this category - i.e. strongly stratum preserving homeomorphisms - will be denoted by $\cong_{\mathrm{N}(P)}$.

From a strongly stratified space, using the above notation, we can recover a stratification by:

$$\begin{aligned}[0,1] &\to \{p < q\} \\ t &\mapsto \begin{cases} p & t = 0; \\ q & t > 0 \end{cases}.\end{aligned}$$

This may also serve as explanation why a strong stratification is a stronger notion than a stratification. Strongly stratified spaces are a natural object to consider in the context of TDA because spaces usually appear equipped with a metric. Then, a stratification naturally induces a strong stratification by using the distance to the singular stratum as map s which, in some sense, gives a parametrization of the neighborhood of the singularity. To make things more precise, we have the following example:

Example 2.1.26. Let $S = (X, s)$ be a stratified space equipped with a metric $\mathrm{d}(-,-)$ on X. Then, S can be equipped with the structure of a strongly stratified space, compatible with the original stratification. The strong stratification map is given by the minimum of the distance to singularity function and 1, i.e. by

$$\begin{aligned}\mathrm{d}_{S_p} \colon X &\to [0,1] \\ x &\mapsto \min\{\mathrm{d}(x, S_p), 1\}.\end{aligned}$$

As a hint towards stratification learning, which we will detail in Section 3.5, note that if we were able to substitute or closely approximate the distance to the singular stratum by another functions that does not require knowledge about the stratification we would have a candidate for a method to stratify a given metric space.

In order to equivalently describe stratified homotopy types we need strongly stratified spaces that have the structure of a mapping cylinder close to the singularity (see Definition 2.1.30). The structure of such spaces near the singularity is specified by the following example.

Example 2.1.27. Given a map of topological spaces $r \colon L \to X$, we can consider the mapping cylinder of r

$$M_r := L \times [0,1] \cup_{L \times 0, r} X$$

equipped with the teardrop topology [Qui88, Definition 2.1] as a strongly stratified space via

$$\begin{aligned}\pi_{[0,1]} : M_r &\to [0,1] \\ [(x,t)] &\mapsto t.\end{aligned}$$

Note that if the above r is a proper map between locally compact Hausdorff spaces, then the usual quotient space topology agrees with the teardrop topology on the mapping cylinder (see e.g. [Hug99] for more details). When working with metric spaces, there is the following criterion for a map

$$f : M_r \to Z$$

into a metric space Z to be continuous. The map f is continuous, if and only if its restrictions to $L \times (0, 1]$ and X are continuous, and the family of maps $f(-, t) : L \to Z$ with $t > 0$, converges uniformly to $f\,|_X \circ r$, as $t \to 0$ (compare to [Qui88, Def. 2.1]). We will use this criterion later on in Example 2.1.33.

Let us now explain how strongly stratified spaces can serve us to describe stratified homotopy types.

Recollection 2.1.28. Without going into too much detail, note that the category of strongly stratified spaces can also be equipped with a notion of weak equivalence. Localizing at the weak equivalences then leads to a homotopy category denoted $\mathrm{ho}\mathbf{Top}_{\mathrm{N}(P)}$. The forgetful functor

$$\mathbf{Top}_{\mathrm{N}(P)} \to \mathbf{Top}_P,$$

obtained by post composing the strong stratification with the stratification of the interval

$$[0, 1] \to \{p < q\}$$

given by taking 0 as the p-stratum, then (by passing to derived functors with respect to the model structures explained in [Dou21a]) induces an equivalence of homotopy categories

$$\mathrm{ho}\mathbf{Top}_{\mathrm{N}(P)} \to \mathrm{ho}\mathbf{Top}_P.$$

We will regularly treat strongly stratified spaces as stratified spaces under this forgetful functor. The fact that we have an equivalence of categories here means that no homotopy theoretical information is lost by passing from one homotopy category to the other.

Although we will not be using this result in the following, it may be useful to understand the occurrences of strongly stratified spaces in our investigations of stratified homotopy types.

For ease of exposition we introduce some short-hand notation similar to what we have for the stratified scenario:

Notation 2.1.29. Let S be a strongly stratified space and $v' \leq v \in [0,1]$. We use the following notation:

$$\begin{aligned} S_v &:= s^{-1}\{v\}, \\ S_{\leq v} &:= s^{-1}[0,v], \\ S^{\geq v'} &:= s^{-1}[v',1], \\ S_v^{v'} &:= s^{-1}[v',v]. \end{aligned}$$

For values of v, v' outside of $[0,1]$ we define these as above, using the closest allowable value.

It turns out that for particularly nice strongly stratified spaces, these sub- and superlevel sets can be used to recover the stratification diagram and in particular the link. Such results can already be found in e.g. [Qui88, Mil94, DW21]. However, to have such an important result readily available in our notation here we include details of this behavior with Example 2.1.33. The particularly nice strongly stratified spaces we are looking for are those for which the strata have cylindrical neighborhoods.

Definition 2.1.30. We say a stratified space $S = (X, X \to P = \{p < q\})$ is *cylindrically stratified*, if there exists a neighborhood N of S_p in X, a space L and a map $r\colon L \to S_p$, such that

$$N \cong_P M_r,$$

where M_r denotes the stratified mapping cylinder of r from Example 2.1.27. We say a strongly stratified space $S = (X, s\colon X \to [0,1))$ is *cylindrically stratified*, if it is cylindrically stratified as a stratified space and the strongly stratified subspace

$$s^{-1}(0,1) = \bigcup_{v<1} S_{\leq v} \setminus S_0$$

is strongly stratum preserving homeomorphic to $S_{\frac{1}{2}} \times (0,1)$, i.e. making the the diagram

$$\begin{array}{ccc} s^{-1}(0,1) & \xrightarrow[\sim]{f} & S_{\frac{1}{2}} \times (0,1) \\ & {\scriptstyle s|_{s^{-1}(0,1)}} \searrow \quad \swarrow {\scriptstyle \pi_{(0,1)}} & \\ & (0,1) & \end{array} \tag{2.1}$$

commute.

Remark 2.1.31. Note, that the definition of a cylindrically stratified space in the strong case is slightly weaker then assuming a strongly stratified mapping cylinder neighborhood. We choose this definition for our purposes as it has precisely the same consequences for us and is much easier to verify. Nevertheless, it follows by

an application of the two-out-of-six property, as in Lemma 2.2.17, that the inclusion $S_0 \hookrightarrow S_{\leq v}$, for $v < 1$, are homotopy equivalences.

Definition 2.1.32. A *cylindrically stratified metric space* S over $P = \{p < q\}$ is a stratified space equipped with a metric $\mathrm{d}(-,-)$, which is cylindrically stratified when considered as a strongly stratified space, with respect to the strong stratification induced by the metric (compare to Example 2.1.26).

Example 2.1.33. At this point we owe the reader an argument that the property of being cylindrically stratified is something that many of the spaces that we are interested in actually posses. For example, Whitney stratified spaces, equipped with the metric induced by the inclusion into $\mathbb{R}^N$, are cylindrically stratified (up to a rescaling). They even admit neighborhoods that are strongly stratum preserving homeomorphic to a strongly stratified mapping cylinder of a fiber bundle (in particular, they are conically stratified). This is a classical result that can be found, for example, in [Tho69, Mat12]. We would like to sketch some of the ideas that go into a proof of this statement to get more familiar with this concept.
Let $W = (X, X \to \{p < q\})$ be a Whitney stratified space with W_p compact and $X \subset \mathbb{R}^N$. By passing to a sufficiently small neighborhood of W_p we may assume W_p to lie in a (standard) tubular neighborhood N of W_p in $\mathbb{R}^N$, such that the retraction map

$$\begin{aligned} r : N &\to W_p \\ x &\mapsto y_m, \end{aligned}$$

where y_m minimizes $\mathrm{d}(x, y)$, is well defined and smooth. Next, consider the distance to W_p map

$$\begin{aligned} \mathrm{d}_{W_p} \colon X &\to \mathbb{R} \\ x &\mapsto \mathrm{d}(x, W_p). \end{aligned}$$

It is then a consequence of Thom's first isotopy lemma (which in this two strata case amounts to Ehresmann's lemma [Ehr51]) see e.g. [Tho69] and [Ban07, Thm. 6.7] for a modern source), that the map

$$\begin{aligned} X \cap N &\to W_p \times \mathbb{R} \\ x &\mapsto (r(x), \mathrm{d}_{W_p}(x)) \end{aligned}$$

restricts to a fiber bundle over $W_p \times (0, \varepsilon]$, for ε small enough. If we denote by $\mathrm{N}_\varepsilon(W_p)$ a closed ε-neighborhood of W_p and set $L = \mathrm{d}_{W_p}^{-1}(\varepsilon)$ this means that there is a homeomorphism

$$f : \mathrm{N}_\varepsilon(W_p) \setminus W_p \xrightarrow{\sim} L \times (0, \varepsilon]$$

such that the diagram

$$\begin{array}{ccc} \mathrm{N}_\varepsilon(W_p) \setminus W_p & \xrightarrow{\quad f \quad} & L \times (0,\varepsilon] \\ & {\scriptstyle r\times \mathrm{d}_{W_p}} \searrow \quad \swarrow {\scriptstyle r\times \pi_{(0,\varepsilon]}} & \\ & W_p \times (0,\varepsilon] & \end{array} \tag{2.2}$$

commutes. By rescaling, we may assume without loss of generality that $\varepsilon = 1$ and let $N = \mathrm{N}_1(W_p)$, the closed neighborhood of points with distance ≤ 1 to W_p.
Now, consider the map

$$g\colon M_r \to N; \qquad \begin{cases} (x,t) \mapsto f(x,t), & \text{for } t > 0, \\ [(x,0)] = [y] \mapsto r(x) = y, & \text{for } t = 0. \end{cases}$$

The map g is clearly bijective and continuous on W_p and $L\times(0,1]$. Furthermore, by the commutativity of 2.2, for $t \to 0$, we find that $f(-,t) : L \to N$ converges uniformly to $r\mid_L$. By the alternative characterization of the mapping cylinder topology in Example 2.1.27 it follows that g is continuous. In total, g is a continuous bijection from a compactum to a Hausdorff space, and thus a homeomorphism.

Example 2.1.34. Another example of cylindrically stratified spaces are given by compact definably stratified spaces $S = (X, s : X \to \{p < q\})$. First, note that they are cylindrically stratified as topological spaces. This follows from the fact that they are triangulable in a way that is compatible with the strata (see [vdD98] for more details). In particular, S_p always admits a mapping cylinder neighborhood given by a regular neighborhood in the piecewise linear sense. Furthermore, note that the map

$$\mathrm{d}_{S_p} : X \to \mathbb{R}$$

again is definable. Thus, by Hardt's theorem for definable sets (see e.g. [vdD98] for a proof of such a statement), it restricts to a trivial fiber bundle over $(0,\varepsilon]$ for ε sufficiently small. After a possible rescaling, we indeed have a strongly stratum preserving homeomorphism

$$s^{-1}(0,1) \to S_{\frac{1}{2}} \times (0,1).$$

over $(0,1)$.

Remark 2.1.35. We will generally consider all compact definably or Whitney stratified spaces to be appropriately rescaled, such that they are cylindrically stratified. Similar assumptions will be made for definably stratified spaces, as we will see, when using Lemma 2.2.17.

We arrived at the construction of a geometric model for stratification diagrams that arises for a cylindrically stratified space by considering level-sets of the map that comes with a strongly stratified space. Such models, detailed in the following construction, together with Proposition 2.1.37 will make the concept of stratification diagrams, and thereby stratified homotopy types, useful for TDA.

Construction 2.1.36. *Given a stratified mapping cylinder M_r for $r : L \to Y$ a map of metrizable spaces. Consider the map*

$$\begin{aligned} \alpha \colon L &\to \mathrm{hoLink}_{\mathrm{p<q}} M_r \\ x &\mapsto \{t \mapsto [(x,t)]\}, \end{aligned}$$

that maps a point x to the corresponding line segment in M_r. A homotopy inverse to this map is given by

$$\begin{aligned} \beta \colon \mathrm{hoLink}_{\mathrm{p<q}} M_r &\to L \\ \gamma &\mapsto \pi_L(\gamma(1)). \end{aligned}$$

By construction, $\beta \circ \alpha = id_L$. A homotopy between $\alpha \circ \beta$ and $id_{\mathrm{hoLink}_{\mathrm{p<q}} M_r r}$ is given by

$$\begin{aligned} \mathrm{hoLink}_{\mathrm{p<q}} M_r \times [0,1] &\to \mathrm{hoLink}_{\mathrm{p<q}} M_r \\ (\gamma, s) &\mapsto \{t \mapsto (\pi_L(\gamma(s + (1-s)t), t). \end{aligned}$$

Compare to e.g. [DW21, Fri03, Qui88] for more details on such constructions and for rigorous arguments for the continuity of such maps. Now, if S is a metrizable, cylindrically stratified space over $P = \{p < q\}$ and $N \cong M_r$ is a stratified mapping cylinder neighborhood of S_p with boundary L, then the inclusion

$$\mathrm{hoLink}_{\mathrm{p<q}} N \hookrightarrow \mathrm{hoLink}_{\mathrm{p<q}} S$$

is a (weak) homotopy equivalence. Essentially, the idea of the proof is to continuously retract paths in the underlying space of S into N (see [Fri03, Appendix] for details under slightly stronger assumptions). In particular, we have a commutative diagram

$$\begin{array}{ccccc} S_q & = & & = & S_q \\ \uparrow & & & & \uparrow \\ L \times \{v\} & \xhookrightarrow{\simeq} & \mathrm{hoLink}_{\mathrm{p<q}}(N) & \xhookrightarrow{\simeq} & \mathrm{hoLink}_{\mathrm{p<q}}(S), \\ \downarrow r & & & & \downarrow \\ S_p & = & & = & S_p\,, \end{array}$$

for $v \in (0,1]$ making the homotopy link of S interpretable as the link of S_p.

Next, we show that the diagram of cylindrically stratified metric space given by the level-sets of the map that comes with the strong stratification induced by the metric, denoted $\{S_{\leq v_h} \hookleftarrow S^{v_l}_{v_h} \hookrightarrow S^{\geq v_l}\}$, is equivalent to the diagram given by $\mathrm{D}_P(S)$ (compare to Recollection 2.1.23).

Proposition 2.1.37. *Let S be a compact, cylindrically stratified metric space and (v_l, v_h) such that $0 < v_l \leq v_h < 1$. Then there is an isomorphism*

$$\{S_{\leq v_h} \hookleftarrow S^{v_l}_{v_h} \hookrightarrow S^{\geq v_l}\} \simeq \mathrm{D}_P(S)$$

in ho$\mathbf{Diag}_P$.

Proof. Let s be the strong stratification induced by the metric on S. By assumption, S admits a mapping cylinder neighborhood $N \cong M_r$, for some map $r\colon L \to S_p$. Let $\tilde{s}\colon N \to [0,1]$ denote the alternative strong stratification induced by this choice of mapping cylinder neighborhood. Since we assume that S_p is compact we may assume, without loss of generality, that $N \subset s^{-1}[0,1)$. By Construction 2.1.36 (using the same notation), it suffices to expose a (canonical) zigzag of weak equivalence to the diagram

$$\{S_p \leftarrow L \times \{v\} \hookrightarrow S_q\},$$

for some $v \in (0,1]$. Such a zigzag between diagrams is given as follows:

$$\begin{array}{ccccc}
\{S_{\leq v_h} & \longleftarrow & S^{v_l}_{v_h} & \longhookrightarrow & S^{\geq v_l}\} \\
 & & \downarrow \simeq & & \\
\{s^{-1}[0,1) & \longleftarrow & s^{-1}(0,1) & \longhookrightarrow & s^{-1}(0,1]\} \\
 & & \uparrow \simeq & & \\
\{s^{-1}[0,1) & \longleftarrow & L \times \{v\} & \longhookrightarrow & s^{-1}(0,1] = S_q\} \\
 & & \downarrow \simeq & & \\
\{S_p & \underset{r}{\longleftarrow} & L & \longhookrightarrow & S_q\}.
\end{array} \tag{2.3}$$

We describe the morphisms of diagrams from top to bottom, and show that they are weak equivalences. The first morphism is given by inclusions. Since we have assumed that $s^{-1}(0,1)$ has the shape of an open cylinder $L' \times (0,1)$ for some L' this morphism is clearly given by weak equivalences at the $\{p\}, \{q\}$ and $\{p < q\}$ parts respectively. The next morphism is again given by inclusions. To see that it is a weak equivalence, we need to show that

$$L \times \{v\} \hookrightarrow \tilde{s}^{-1}(0,1] \hookrightarrow s^{-1}(0,1) \cong L' \times (0,1)$$

is a weak equivalence. Since the first of these maps is a weak equivalence, it suffices to show that the second is one, too. Since S_p is compact we find $\varepsilon, \varepsilon' > 0$ such that

$$\tilde{s}^{-1}(0,\varepsilon') \subset s^{-1}(0,\varepsilon) \subset \tilde{s}^{-1}(0,1) \subset s^{-1}(0,1).$$

Now, note that since all sets involved are given by open cylinders (on L and L' respectively), these inclusions fulfill the requirements for the two-out-of-six property of homotopy equivalences. In particular, all maps involved are weak equivalences (even homotopy equivalences). Finally, the last morphism is constructed as follows. Both at q and $\{p < q\}$ it is given by the identity. Assume that $v \in (0,1]$ was taken such that

$$L \times \{v\} \cong \tilde{s}^{-1}\{v\} \subset s^{-1}(0,\varepsilon] \subset \tilde{s}^{-1}(0,1),$$

for some $\varepsilon > 0$ sufficiently small. Note that this is indeed possible by the compactness of S_p. Then, at $\{p\}$ the morphism is given by the composition

$$s^{-1}[0,1) \to s^{-1}[0,\varepsilon] \hookrightarrow N \cong M_r \to S_p$$

where the first of these maps is given by

$$(x,t) \mapsto (x, \min\{t,\varepsilon\})$$

under the identification $s^{-1}(0,1) \cong L' \times (0,1)$. By the assumption that $L \times \{v\}$ maps into $s^{-1}[0,\varepsilon]$, this map induces a morphism of diagrams. It remains to show that it is a weak equivalence. By the cylinder structure of $s^{-1}(0,1)$, the first map in this composition is a homotopy equivalence. The same holds true for the second map by a two-out-of-six argument, completely analogous to the one performed when comparing L and $L' \times (0,1)$. Finally, the last map is the retraction of a mapping cylinder and thus also a homotopy equivalence. Combining this, we have shown that the final morphism is also a weak equivalence of diagrams. □

2.2 Persistent Stratified Homotopy Types

In this section, we introduce the concept of persistent stratified homotopy type (Section 2.2.1) and delve into its stability properties (Section 2.2.3). The stability analysis is most pronounced in the case of Whitney stratified spaces with two strata, presenting a stronger result compared to (Theorem 2.2.32). Before we delve into the subject matter, we establish some notation for the remainder of this chapter to simplify the otherwise extensive discussion of persistent objects. When we initiate computational considerations at the beginning of Chapter 4, we will reassess some of the objects discussed in the upcoming sections from an alternative and potentially more recognizable perspective.

Notation 2.2.1. Let $\mathbf{T}$ denote some category of geometrical and or combinatorial objects, for example the categories **Top**, **sSet** (the category of simplicial sets), $\mathbf{Top}_P$, $\mathbf{Diag}_P$, equipped with a class of morphisms called weak equivalences. Let $\mathbf{A}$ denote some category of algebraic objects, for example, the category of vector spaces over some fixed field. Then, let

$$H\colon \mathbf{T} \to \mathbf{A}$$

be a functor that sends weak equivalences to isomorphisms, e.g. the homology functor in case $\mathbf{T} = \mathbf{Top}$.

- We denote by $(\mathbb{R}_+, \leq)$ (or briefly by $\mathbb{R}_+$ if the context is clear) the posetal category of non-negative reals.
- Given any category $\mathcal{C}$ and another category U (most prominently $\mathbb{R}_+$) denote by $\mathcal{C}^U$ the category of functors from U to $\mathcal{C}$ with morphisms given by natural transformations.
- Let $\mathbf{S}$ denote some categories of objects representing datasets which contain geometrical information. A persistent version of the homology functor $H = H_i(-; k)$ with k a field is constructed by taking a composition

  $$\mathbf{S} \to \mathbf{T}^{\mathbb{R}_+} \to \mathbf{A}^{\mathbb{R}_+},$$

  for some functor $\mathbf{S} \to \mathbf{T}^{\mathbb{R}_+}$ turning objects in $\mathbf{S}$ into persistent objects in $\mathbf{T}$, i.e. elements of $\mathbf{T}^{\mathbb{R}_+}$. Examples of such functors include sending a subspace of $\mathbb{R}^N$ to the family of its ε thickenings, filtered Čech- or Vietoris-Rips complexes. As an orientation, note that in case $\mathbf{A} = \mathbf{Vect}_{\mathbb{F}}$ with $\mathbb{F}$ a field, a functor $I \to \mathbf{A}$ with I a posetal category associated to some poset is commonly referred to as *persistence module indexed over* I.
- ho$\mathbf{T}^{\mathbb{R}_+}$ denotes the category obtained by localizing $\mathbf{T}^{\mathbb{R}_+}$ at natural transformations, given by weak equivalences at each $\alpha \in \mathbb{R}_+$. Such natural transformations will be called weak equivalences of functors. We will also refer to isomorphism in the homotopy category ho$\mathbf{T}^{\mathbb{R}_+}$ (which are always represented by zigzags of weak equivalences in $\mathbf{T}^{\mathbb{R}_+}$) as weak equivalences. The same notation and language is used when $\mathbb{R}_+$ is replaced by a more general indexing category.

The above functor $\mathbf{S} \to \mathbf{T}^{\mathbb{R}_+}$ can then be understood as assigning to an object in $\mathbf{S}$ a persistent homotopy type which we will be referring to as the *persistent homotopy type of* $\mathbb{X} \in \boldsymbol{S}$. We allow this inaccurate notion and let the type depend on the choice of a functor $\mathbf{S} \to \mathbf{T}^{\mathbb{R}_+}$ as it is also common to speak of *the* persistent homology

although it usually depends on the choices such as Čech or Vietoris-Rips complexes. Many properties of persistent homology can already be understood at the level of the persistent homotopy type. Another advantage of this approach is that it quickly allows to obtain results for different choices of H.

These considerations regarding the persistent homotopy type motivate our declared goal to uncover an analog functor $\mathbf{S} \to \mathrm{ho}\mathbf{Top}_P^{\mathbb{R}_+}$ with values in the homotopy category $\mathrm{ho}\mathbf{Top}_P^{\mathbb{R}_+}$ (or, by composing, in $\mathrm{ho}\mathbf{Strat}^{\mathbb{R}_+}$). Analogous to the non-stratified case a persistent stratified homotopy type of a space, possibly represented by thickenings or a Čech complex, we want the values of such a functor to fulfill certain properties (see (1), (2), (3) in the introduction) amongst which is a form of stability stability with respect to Hausdorff and interleaving type distances similar to the classical results for persistent homology (see [CSEH07, CCSG+09, BL15]). Our results in that direction are stated in Theorems 2.2.31 and 2.2.32. Furthermore, we want the persistent stratified homotopy type to be computable, i.e. to be equivalently represented by a filtered Čech complex (cf. Remark 2.2.15, a consequence of the Nerve Theorem in the non-stratified case.

2.2.1 Definition of persistent stratified homotopy types

Let us review the steps we have taken thus far to make progress toward our objective. Start with a Whitney stratified or definably stratified space $S \subset \mathbb{R}^N$ (or later on a sample of any of the latter) aiming to obtain a persistent version of its stratified homotopy type. By Recollection 2.1.23 the stratified homotopy type of S is equivalently described by its stratification diagram $\mathrm{D}_P(S)$. Then, one could employ [Dou21a, Thm. 2.12] (which is a stronger version of Recollection 2.1.23) to see that that the homotopy link functor from Recollection 2.1.23 sending a stratified space to a stratification diagram, induces an equivalence of categories

$$\mathrm{ho}\mathbf{Top}_P^{\mathbb{R}_+} \xrightarrow{\sim} \mathrm{ho}\mathbf{Diag}_P^{\mathbb{R}_+}. \tag{2.4}$$

Thus, we may equivalently work towards a functor valued in $\mathrm{ho}\mathbf{Diag}_P^{\mathbb{R}_+}$. To do so, we need to obtain a geometric description of the stratification diagram $\mathrm{D}_P(S)$, which admits thickenings. By Examples 2.1.33 and 2.1.34, S naturally admits the structure of a cylindrically stratified space (up to a rescaling). Then, by Proposition 2.1.37, up to a weak equivalence we may recover the stratification diagram $\mathrm{D}_P(S)$ of S by the diagram of sub- and superlevel sets

$$S_{\leq v_h} \hookleftarrow S_{v_h}^{v_l} \hookrightarrow S^{\geq v_l}.$$

The diagram obtained in this fashion has the advantage that it admits an obvious notion of thickening. One simply thickens the parts of the diagram contained in $\mathbb{R}^N$ separately.

Let us now set up the framework to analyze the stability properties of these constructions and their interactions with sampling. For the remainder of this subsection let $P = \{p < q\}$ be a poset with two elements. we define a series of space of subspaces of $\mathbb{R}^N$. One should mainly think of elements of these spaces as samples, sampled nearby a continuous space, whose convergence behavior we are investigating. Nevertheless, in the generality below all sorts of complicated, non-finite sets are permitted. We will follow the naming convention of using blackboard bold letters like $\mathbb{X}$, when suggesting an object conceptually takes the role of a sample. We use usual letters, like X, when an objects takes the role of a space nearby which samples are taken.

Notation 2.2.2. Throughout the remainder of this thesis, we will follow the convention of denoting all distances by the letter d. What metric is meant will be clear from the specific context. Furhtermore, the Euclidean norm on $\mathbb{R}^N$ will always be donted by $||-||$. For $\mathbb{X} \subset \mathbb{R}^N$, and $\alpha \geq 0$ the α thickening of $\mathbb{X}$, denoted $\mathbb{X}_\alpha$, is given by

$$\{y \in \mathbb{R}^N \mid \exists x \in \mathbb{X} : ||x - y|| \leq \alpha\}.$$

Definition 2.2.3. Denote by **Sam** the space of subspaces of $\mathbb{R}^N$, $\{\mathbb{X} \subset \mathbb{R}^N\}$, equipped with the (extended pseudo) metric given by the Hausdorff-distance. That is, for $\mathbb{X}, \mathbb{X}' \in \mathbf{Sam}$, we set

$$\mathrm{d}(\mathbb{X}, \mathbb{X}') = \inf\{\alpha > 0 \mid \mathbb{X} \subset \mathbb{X}'_\alpha, \mathbb{X}' \subset \mathbb{X}_\alpha\}.$$

We refer to **Sam** as the *space of samples* (of $\mathbb{R}^N$).

There is an obvious extension to the case where a sample comes equipped with a stratification, i.e. a sample version of the category $\mathbf{Top}_P$ with a metric.

Definition 2.2.4. Denote by $\mathbf{Sam}_P$ the space of pairs in $\mathbb{R}^N$

$$\{(\mathbb{X}, \mathbb{X}_p) \mid \mathbb{X}_p \subset \mathbb{X}\} \subset \mathbf{Sam}^2.$$

For $\mathbb{X}, \mathbb{X}' \in \mathbf{Sam}_P$, we set

$$\mathrm{d}((\mathbb{X}, \mathbb{X}_p), (\mathbb{X}', \mathbb{X}'_p)) := \max\{\mathrm{d}(\mathbb{X}, \mathbb{X}'), \mathrm{d}(\mathbb{X}_p, \mathbb{X}'_p)\}$$

which defines an (extended pseudo) metric on $\mathbf{Sam}_P$. We also refer to $\mathbf{Sam}_P$ as the *space of (P-)stratified samples* (of $\mathbb{R}^N$).

Next, we need an analogous construction for the category $\mathbf{Diag}_P$.

Definition 2.2.5. Denote by $\mathbf{D}_P\mathbf{Sam}$ the space of triples of subspaces of $\mathbb{R}^N$

$$\{(\mathbb{D}_p, \mathbb{D}_{\{p<q\}}, \mathbb{D}_q) \mid \mathbb{D}_q \supset \mathbb{D}_{\{p<q\}} \subset \mathbb{D}_p, \} \subset \mathbf{Sam}^3,$$

equipped with the subspace metric. That is, for $\mathbb{D}, \mathbb{D}' \in \mathbf{D}_P\mathbf{Sam}$ we set

$$\mathrm{d}(\mathbb{D}, \mathbb{D}') = \max_{\mathcal{I} \in \mathrm{R}(P)} \mathrm{d}(\mathbb{D}_{\mathcal{I}}, \mathbb{D}'_{\mathcal{I}}).$$

We also refer to $\mathbf{D}_P\mathbf{Sam}$ as the *space of stratification diagram samples* (of $\mathbb{R}^N$).

Finally, we repeat a similar process for $\mathbf{Top}_{\mathrm{N}(P)}$.

Definition 2.2.6. Denote by $\mathbf{Sam}_{\mathrm{N}(P)}$ the space

$$\{(\mathbb{X}, s\colon \mathbb{X} \to [0,1]) \mid \mathbb{X} \subset \mathbb{R}^N\},$$

equipped with the (extended pseudo) metric given as follows. Embed $\mathbf{Sam}_{\mathrm{N}(P)}$ into the space of subspaces of $\mathbb{R}^N \times [0,1]$, equipped with the Hausdorff distance on the product metric, by sending s to its graph. The metric on $\mathbf{Sam}_{\mathrm{N}(P)}$ is given by the subspace metric under this embedding.
Equivalently, this comes down to the following. For $(\mathbb{X}, s), (\mathbb{X}', s') \in \mathbf{Sam}_{\mathrm{N}(P)}$, we set

$$\mathrm{d}((\mathbb{X}, s), (\mathbb{X}', s')) := \max_{\mathbb{X}_0, \mathbb{X}_1 \in \{\mathbb{X}, \mathbb{X}'\}^2} \{\inf\{\varepsilon > 0 \mid \forall x \in \mathbb{X}_0 \exists y \in \mathbb{X}_1 : \\ ||x - y||, |s(x) - s'(y)| \leq \varepsilon\}\}.$$

We also refer to $\mathbf{Sam}_{\mathrm{N}(P)}$ as the *space of strongly (P-)stratified samples* (of $\mathbb{R}^N$).

Notation 2.2.7. For the remainder of this work, we denote $v = (v_l, v_h) \in (0,1)^2$ with $v_l < v_h$, and $u \in (0,1)$. Furthermore, we equip $(0,1)^2$ with the usual order in the second, and the opposite order in the first component, that is we write

$$v \leq v' \iff v'_l \leq v_l \quad \text{and} \quad v_h \leq v'_h.$$

We denote by $\Omega \subset (0,1)^2$ the subposet of $(0,1)^2$ with this order, consisting of (v_l, v_h) with $0 < v_l < v_h < 1$.

We close this sequence of definitions with constructions that connect stratified spaces, strongly stratified spaces and stratification diagrams.

Construction 2.2.8. *Consider the following three maps.*

$$\mathbf{Sam}_P \underset{\mathcal{F}_u}{\overset{\mathcal{N}}{\rightleftarrows}} \mathbf{Sam}_{\mathrm{N}(P)} \xrightarrow{\mathcal{D}_v} \mathbf{D}_P\mathbf{Sam}$$

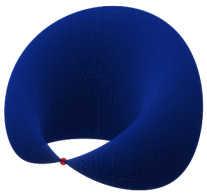

Figure 2.8. The pinched torus PT as an element of $\mathbf{Sam}_P$

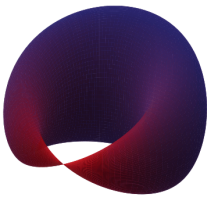

Figure 2.9. Illustration of $\mathcal{N}(PT)$. The colouring indicates the strong stratification.

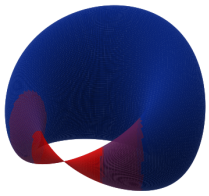

Figure 2.10. Illustration of $\mathcal{D}_v(\mathcal{N}(PT))$.

These are defined via:

$$\begin{aligned}(\mathbb{X}, \mathbb{X}_p) &\overset{\mathcal{N}}{\longmapsto} (\mathbb{X}, \min\{\mathrm{d}_{\mathbb{X}_p}, 1\}),\\ \mathbb{S} = (\mathbb{X}, s) &\overset{\mathcal{F}_u}{\longmapsto} (\mathbb{X}, \mathbb{S}_{\leq u}),\\ \mathbb{S} = (\mathbb{X}, s) &\overset{\mathcal{D}_v}{\longmapsto} (\mathbb{S}_{\leq v_h}, \mathbb{S}^{v_l}_{v_h}, \mathbb{S}^{\geq v_l})\,.\end{aligned}$$

The map $\mathcal{N}$ corresponds to the assignment of a strong stratification to a stratified metric space (see Example 2.1.26. $\mathcal{F}_u$ gives a family of models for the forgetful functor, $\mathbf{Top}_{\mathrm{N}(P)} \to \mathbf{Top}_P$, described in Recollection 2.1.28. Finally, by Proposition 2.1.37, $\mathcal{D}_v$ (composed with $\mathcal{N}$) provides a model for the functor assigning to a stratified space its stratification diagram, $\mathrm{D}_P \colon \mathbf{Top}_P \to \mathbf{Diag}_P$ (see Recollection 2.1.23).

Example 2.2.9. Consider the three pictures Figs. 2.8 to 2.10. Fig. 2.8 shows the pinched torus PT as a stratified subspace of $\mathbb{R}^3$, with the singularity marked in red. Fig. 2.9 shows $\mathcal{N}(PT)$, where the color coding indicated the strong stratification. Finally, Fig. 2.10 shows the image under $\mathcal{D}_v$ for $v = (0.2, 0.4)$. Specifically, the union of the red and purple part give the p-part of the diagram, the purple part the $\{p < q\}$-part, and the union of the purple and the blue one the q-part.

We will later make use of the following immediate elementary relation between $\mathcal{D}_v$ and $\mathcal{F}_u$.

Lemma 2.2.10. *Let $v = (v_l, u) \in \Omega$. Then,*

$$\mathcal{F}_u(\mathbb{S}) = (\mathcal{D}_v(\mathbb{S})_p \cup \mathcal{D}_v(\mathbb{S})_{\{p<q\}} \cup \mathcal{D}_v(\mathbb{S})_q, \mathcal{D}_v(\mathbb{S})_p)$$

for all $\mathbb{S} \in \mathbf{Sam}_{\mathrm{N}(P)}$.

Remark 2.2.11. Note, that all of the described sample spaces naturally admit the structure of a poset. In the case of $\mathbf{Sam}$, $\mathbf{Sam}_P$ and $\mathbf{D}_P\mathbf{Sam}$ it is simply given by

inclusions. In case of $\mathbf{Sam}_{\mathrm{N}(P)}$, it is obtained by treating elements of $\mathbf{Sam}_{\mathrm{N}(P)}$ as their graph, i.e. as a subset of $\mathbb{R}^N \times [0,1]$ and then using the inclusion relation. Equivalently, this means

$$(\mathbb{X}, s) \leq (\mathbb{X}', s') \iff (x \in \mathbb{X} \implies (x \in \mathbb{X}' \ \& \ s(x) = s'(x))).$$

In this fashion, the spaces of samples may also be treated as categories and the maps of Construction 2.2.8 are functors. Furthermore, from this perspective we can treat $\mathbf{Sam}_P$ as a subcategory of $\mathbf{Top}_P$ which allows for notation such as $\mathbf{Sam}^I$, where I is some indexing category to make sense, and we will use this freely. Furthermore, from this perspective the metrics on $\mathbf{Sam}_P$ and $\mathbf{D}_P\mathbf{Sam}$ are induced by the flow given by componentwise thickening (see Example 2.2.20 for details).

Remark 2.2.12. In the proof of Proposition 2.1.37 we have first constructed a weak equivalence of $\mathcal{D}_v \circ \mathcal{N}(S)$. Under the equivalence of homotopy categories $\mathrm{ho}\mathbf{Top}_P \cong \mathrm{ho}\mathbf{Diag}_P$, S and $\mathcal{D}_v \circ \mathcal{N}(S)$ represent the same stratified homotopy type. As a consequence, to define persistent stratified homotopy types, we can thicken stratification diagrams instead of stratified spaces.

Construction 2.2.13. *Define* the thickening of $\mathbb{D}$ by $\alpha \geq 0$ *via:*

$$\mathbb{D}_\alpha := ((\mathbb{D}_q)_\alpha, (\mathbb{D}_{\{p<q\}})_\alpha, (\mathbb{D}_q)_\alpha).$$

For $\alpha \leq \alpha'$ there are the obvious inclusions of diagrams

$$\mathbb{D}_\alpha \hookrightarrow \mathbb{D}_{\alpha'}$$

We thus obtain a map (functor from the categorical perspective)

$$\begin{aligned} \mathbf{D}_P\mathbf{Sam} &\to \mathbf{D}_P\mathbf{Sam}^{\mathbb{R}_+} \\ \mathbb{D} &\mapsto \{\alpha \mapsto \mathbb{D}_\alpha\} \end{aligned}$$

with the structure maps given by inclusions. Therefore, may treat the diagram samples as elements of $\mathbf{Diag}_P$*, ultimately obtaining the composition:*

$$\mathcal{DP} \colon \mathbf{D}_P\mathbf{Sam} \to \mathbf{D}_P\mathbf{Sam}^{\mathbb{R}_+} \to \mathrm{ho}\mathbf{Diag}_P^{\mathbb{R}_+} \simeq \mathrm{ho}\mathbf{Top}_P^{\mathbb{R}_+}.$$

With this everything is in position to define persistent stratified homotopy types.

Definition 2.2.14. The *persistent stratified homotopy type* of a stratified sample $\mathbb{S} \in \mathbf{Sam}_P$ (depending on the parameter v) is defined as the image of $\mathbb{S}$ under the composition

$$\mathcal{P}_v \colon \mathbf{Sam}_P \xrightarrow{\mathcal{N}} \mathbf{Sam}_{\mathrm{N}(P)} \xrightarrow{\mathcal{D}_v} \mathbf{D}_P\mathbf{Sam} \xrightarrow{\mathcal{DP}} \mathrm{ho}\mathbf{Top}_P^{\mathbb{R}_+}.$$

Remark 2.2.15. By construction $\mathcal{P}_v$ admits a combinatorial interpretation which can be stored on a computer given that its finite, i.e. for finite samples. In other words, we can immediately see that the persistent stratified homotopy type fulfills an analogue of property (1). Note that the persistent stratified homotopy type of $\mathbb{S} \in \mathbf{Sam}_P$ is equivalently represented by the image of $\mathbb{S}$ under

$$\mathcal{P}_v \colon \mathbf{Sam}_P \xrightarrow{\mathcal{N}} \mathbf{Sam}_{\mathrm{N}(P)} \xrightarrow{\mathcal{D}_v} \mathbf{D}_P\mathbf{Sam} \xrightarrow{\mathcal{DP}} \mathrm{ho}\mathbf{Diag}_P^{\mathbb{R}_+}.$$

We will not distinguish these two equivalent representation with different notation. Then, it is a consequence of the classical nerve theorem (see e.g. [Hat02, Prop. 4G.3] or [Bor48]), that for $\mathbb{S} \in \mathbf{Sam}_P$, $\mathcal{P}_v(\mathbb{S})$ is equivalently represented by the diagram of Čech complexes

$$\alpha \mapsto \{\mathcal{I} \mapsto \check{\mathrm{C}}_\alpha(\mathcal{D}_v(\mathcal{N}(\mathbb{S}))_\mathcal{I}\}$$

where $\check{\mathrm{C}}_\alpha$ denotes the Čech complex at scale α. In case $\mathbb{S}$ is finite, this is exactly the finite data that can be used used algorithmically, i.e. it can be stored as three filtered simplicial complexes and each part can for example be used for persistent homology calculation providing additional layers of information on the topology extracted from $\mathbb{S}$.

The following result guarantees that for sufficiently regular stratified spaces the homotopy type does not change under small thickenings. This is the stratified analogue to the aforementioned property that the homotopy type does not change under small thickenings, property (3). See also [NSW08] for the smooth case and [CCSL09] for the case of compact Euclidean subsets. This justifies the use of persistent stratified homotopy types as a means to investigate the geometry of stratified samples.

Proposition 2.2.16. *Let $S \in \mathbf{Sam}_P$ be a compact, definable stratified metric space. Then, for any $v \in \Omega$ there exists an $\varepsilon > 0$, such that the structure map*

$$\mathcal{P}_v(S)(\alpha) \to \mathcal{P}_v(S)(\alpha')$$

is a weak equivalence, for all $0 \leq \alpha \leq \alpha' < \varepsilon$. In particular,

$$\mathcal{P}_v(S) \mid_{[0,\varepsilon)} \simeq S.$$

In other words, the persistent stratified homotopy type of S at v restricted to $[0, \varepsilon)$, is weakly equivalent to the constant functor with value S.

Proof. Note that by definition of a weak equivalence in the category of stratification diagrams, this statement really just says there exists an $\varepsilon > 0$, such that for each flag $\mathcal{I}$ in P the inclusions

$$\mathcal{D}_v(S)_\mathcal{I} \hookrightarrow (\mathcal{D}_v(S)_\mathcal{I})_\alpha$$

are weak equivalences, for $\alpha \leq \varepsilon$. Note, however, that by the definability assumption $\mathcal{D}_v(S)_{\mathcal{I}}$ is again definable. Hence, this follows from the fact that the homotopy type of compact definable sets is invariant under slight thickenings (see Lemma 2.2.17 for the precise statement). For $\alpha = 0$, we have

$$\mathcal{P}_v(S)(0) \simeq S$$

as discussed in Remark 2.2.12. □

Lemma 2.2.17. *Let $X \subset Y \subset \mathbb{R}^N$ be definable with respect to some o-minimal structure and X compact. Then, there exists $\varepsilon > 0$ such for $0 < \alpha < \varepsilon$ the following holds:*

1. *$X \hookrightarrow X_\alpha \cap Y$ is a strong deformation retract.*
2. *There is a homeomorphism $(X_\alpha \cap Y) \setminus X \cong \mathrm{d}_X^{-1}(\frac{\alpha}{2}) \times (0, \alpha]$, such that the diagram*

$$\begin{array}{ccc} (X_\alpha \cap Y) \setminus X & \xrightarrow{\quad\sim\quad} & \mathrm{d}_X^{-1}(\tfrac{\alpha}{2}) \times (0, \alpha] \\ & {\scriptstyle \mathrm{d}_X} \searrow \quad \swarrow {\scriptstyle \pi_{(0,\alpha]}} & \\ & (0, \alpha] & \end{array}$$

commutes.

Proof. The claim 2 on the homememorphism type of the complements is an application of Hardt's theorem for definable sets together with the fact that d_X is definable (see e.g. [vdD98]). It remains to see that the inclusion is a strong deformation retraction. Note that by the triangulability of definable sets (see for example [vdD98, Theorem 2.9]), $\mathbb{R}^N$ may be equipped with a triangulation compatible with X and Y. In particular, by subdividing if necessary, X has arbitrarily small mapping cylinder neighborhoods in Y, given by piecewise linear regular neighborhoods. Furthermore, this means that $X \hookrightarrow X_\alpha \cap Y$ is a cofibration. Thus, it suffices to show that $X \hookrightarrow X_\alpha \cap Y$ is a homotopy equivalence. Let $\alpha < \alpha' < \varepsilon$ with ε such that 2 holds. Then, we have inclusions

$$X \hookrightarrow X_\alpha \cap Y \hookrightarrow N \hookrightarrow X_{\alpha'} \cap Y,$$

where N is a regular neighborhood with respect to the piecewise linear structure induced by the triangulation. By the open cylinder structure in 2 of the set $(X_{\alpha'} \cap Y) \setminus X$, the inclusion $X_\alpha \cap Y \hookrightarrow X_{\alpha'} \cap Y$ is a homotopy equivalence. The same holds for the inclusion $X \hookrightarrow N$. It follows by the two-out-of-six property of homotopy equivalences, that all maps are homotopy equivalences. □

2.2.2 Metrics on categories of persistent objects

One of the fundamental prerequisites for employing persistent homology in applications lies in its stability concerning Hausdorff and interleaving distance, initially demonstrated in [CCSG+09]. Metrics in persistent scenarios and the stability of functors with respect to these metrics have been the subject of ongoing study ([BW20, HKNU17, Les15, BL21, BSS20]). Remarkably, the stability of persistent homology with respect to interleaving distance can be expressed at the level of persistent homotopy types, which also extends to persistent spaces, as will be explained in the section that follows. Examining the stability properties of the stratified persistent homotopy type is the aim of Section 2.2.3. We do this by utilizing the idea of *flows*, which was first presented in [dSMS18]. Here, we provide a somewhat more condensed definition for the sake of clarity.

Recollection 2.2.18 ([dSMS18])**.** A *strict flow* on a category $\mathcal{C}$ is a strict monoidal functor $(-)_- \colon \mathbb{R}_+ \to \mathrm{End}(\mathcal{C})$. That is, to each $\varepsilon \in \mathbb{R}_+$ we assign an endofunctor $(-)_\varepsilon$ and whenever $\varepsilon \leq \varepsilon'$ we assign (functorially) a natural transformation $s_{\varepsilon\to\varepsilon'} \colon (-)_\varepsilon \to (-)_{\varepsilon'}$. Recall that being strict monoidal means that $(-)_0 = 1_{\mathcal{C}}$, $(-)_{\varepsilon'} \circ (-)_\varepsilon = (-)_{\varepsilon+\varepsilon'}$ and $(s_{\varepsilon\leq\varepsilon'})_\delta = s_{\varepsilon+\delta\leq\varepsilon'+\delta}$. Generally, one may think of flows as a notion of shift on $\mathcal{C}$. Then, just as in the scenario of the interleaving distance for persistence modules [CCSG+09], one says that $X, Y \in \mathcal{C}$ are *ε-interleaved* if there are morphisms $f \colon X \to Y_\varepsilon$ and $g \colon Y \to X_\varepsilon$ and such that the diagram

$$\begin{array}{ccc} X & & Y \\ {\scriptstyle f}\downarrow & \swarrow\!\!\searrow & \downarrow{\scriptstyle g} \\ Y_\varepsilon & & X_\varepsilon \\ {\scriptstyle g_\varepsilon}\downarrow & \swarrow\!\!\searrow & \downarrow{\scriptstyle f_\varepsilon} \\ X_{2\varepsilon} & & Y_{2\varepsilon} \end{array} \tag{2.5}$$

commutes (all unlabelled morphisms are given by the flow). One then obtains a (symmetric Lawvere) metric space by setting

$$d_I(X, Y) = \inf\{\varepsilon \geq 0 \mid X, Y \text{ are } \varepsilon\text{-interleaved}\}. \tag{2.6}$$

An immediate consequence of this definition is that any functor $F \colon \mathcal{C} \to \mathcal{D}$ between categories with a (strict) flow that fulfills $(-)_\varepsilon \circ F = F \circ (-)_\varepsilon$ for all $\varepsilon \in \mathbb{R}_+$, is necessary 1-Lippschitz with respect to the interleaving distance. For the specific case of the functor F being given by homology, we exhibit more detail in Section 4.1.

Example 2.2.19. Let $U \subset \mathbb{R}^k$ be a generalized interval (i.e. a subset of $\mathbb{R}^k$, such that $u, v \in U, u \leq v \leq w$ implies $w \in U$). Furthermore, let $\mathbf{T}$ be a category with a

terminal object $*$. Then, any functor category $\mathbf{T}^U$ is naturally equipped with a shift type flow, given by

$$X_\varepsilon(u) := \begin{cases} X(u+\varepsilon(1,\dots,1)) & \text{, for } u+\varepsilon(1,\dots,1) \in U; \\ * & \text{, for } u+\varepsilon(1,\dots,1) \notin U. \end{cases}$$

If $\mathbf{T}$ is equipped with a notion of weak equivalence (which includes all isomorphisms), then by construction the flow respects weak equivalences in $\mathbf{T}^U$. Thus, it descends to a flow on the homotopy category $\mathrm{ho}\mathbf{T}^U$ obtained by localizing weak equivalences of functors. In particular, this construction equips the persistent homotopy category $\mathrm{ho}\mathbf{Top}^{\mathbb{R}_+}$ with a symmetric Lawvere metric, such that the functor

$$\mathbf{Top}^{\mathbb{R}_+} \to \mathrm{ho}\mathbf{Top}^{\mathbb{R}_+}$$

is 1-Lippschitz. More generally, the same construction works for the cases $\mathbf{T} = \mathbf{Top}_P, \mathbf{Diag}_P, \mathbf{Strat}$. We call distances of this type *interleaving distances*. Furthermore, for any functor between two such categories $\mathbf{T}, \mathbf{T}'$, which descends to the homotopy category, the induced functors

$$\mathbf{T}^{\mathbb{R}_+} \to \mathbf{T}'^{\mathbb{R}_+},$$
$$\mathrm{ho}\mathbf{T}^{\mathbb{R}_+} \to \mathrm{ho}\mathbf{T}'^{\mathbb{R}_+}$$

commute with shifting and are thus 1-Lippschitz. Whenever we refer to a metric on such a functor category (or its homotopy category), we will mean the interleaving distance.

Example 2.2.20. Denote by **Sam** the category of subsets of $\mathbb{R}^N$, with morphisms given by inclusions. If we take

$$\mathbb{X}_\varepsilon := \mathrm{N}_\varepsilon(\mathbb{X}) := \{y \mid \exists x \in \mathbb{X} \colon |x-y| \leq \varepsilon\},$$

for $\varepsilon \in \mathbb{R}_+$ then this defines a strict flow on **Sam**. The distance induced by the flow is the Hausdorff distance (compare [dSMS18]). Clearly, the functor

$$\mathbf{Sam} \to \mathbf{Top}^{\mathbb{R}_+}$$
$$\mathbb{X} \mapsto \{\varepsilon \mapsto \mathbb{X}_\varepsilon\}$$

commutes with flows. In this notation, persistent homology is the composition

$$\mathrm{PH}_i \colon \mathbf{Sam} \to \mathrm{ho}\mathbf{Top}^{\mathbb{R}_+} \xrightarrow{\mathrm{H}_i^{\mathbb{R}_+}} \mathbf{Vec}_k^{\mathbb{R}_+} \tag{2.7}$$

where $\mathrm{H}_i^{\mathbb{R}_+}$ computes homology degree-wise which also commutes with flows. This is a way to prove the fact that persistent homotopy types are stable with respect

to Hausdorff and interleaving distance. This implies the stability of persistent homology (as in [CCSG+09]). Note that the interleaving distance on $\mathbf{Vec}_k^{\mathbb{R}_+}$ agrees with the Bottleneck distance (see Definition 4.1.6 and Eq. (4.2)). More generally, if we define thickening component-wise, then the distances on $\mathbf{Sam}_P$ and $\mathbf{D}_P\mathbf{Sam}$ (Definitions 2.2.4 and 2.2.5) are also induced by the thickening flow. However the interleaving distance on $\mathrm{ho}\mathbf{Top}^{\mathbb{R}_+}$, $\mathrm{ho}\mathbf{Sam}_P^{\mathbb{R}_+}$ or $\mathrm{ho}\mathbf{D}_P\mathbf{Sam}^{\mathbb{R}_+}$ may look like, i.e. the distance of their version of persistent homotopy types, we know that objects close in the interleaving distance will also be close in the interleaving distance in their associated persistence barcodes or similar invariants.

In the next subsection, the stability results as in Examples 2.2.19 and 2.2.20 are used repeatedly and often implicitly. For example, from the perspective of flows one immediately obtains:

Lemma 2.2.21. $\mathcal{DP}\colon \mathbf{D}_P\mathbf{Sam} \to \mathrm{ho}\mathbf{Top}_P^{\mathbb{R}_+}$ *is* 1*-Lipschitz.*

2.2.3 Stability of persistent stratified homotopy types

As illustrated in Section 2.2.2, particularly in Example 2.2.20, the stability of the persistent homotopy type concerning the Hausdorff distance is almost self-evident. When commencing with a stratification diagram of samples, the situation remains comparably straightforward for the stratified scenario, as indicated in Lemma 2.2.21.

However, the scenario becomes more nuanced for persistent stratified homotopy types in some sense when initiated from a stratified sample. This complexity arises from the fact that the operation of taking sublevel sets with respect to some function (or, equivalently, intersecting with a closed subset) is generally not a continuous operation in terms of the Hausdorff distance.

Example 2.2.22. Let $N = 1$, i.e. $\mathbb{R}^N = \mathbb{R}$ and $A = (-\infty, 0]$. Let $\mathbb{X} = \{0\}$ and $\mathbb{X}_n = \{\frac{1}{n}\}$. $\mathbb{X}_n$ converges to $\mathbb{X}$ in the Hausdorff distance. However, the intersections $A \cap \mathbb{X}_n$ are empty, while $A \cap \mathbb{X} = \mathbb{X}$. In particular, we have

$$\mathrm{d}(A \cap \mathbb{X}_n, A \cap \mathbb{X}) = \infty,$$

for all $n > 0$.

The issue highlighted in Example 2.2.22 arises from a lack of homogeneity in $\mathbb{X}$ while transitioning into the interior of A. By replacing $\mathbb{X}$ with $[-\varepsilon, 0]$, such phenomena are eliminated, and one can demonstrate the continuity of intersecting with A in $[-\varepsilon, 0]$. Expanding on this phenomenon in a more intricate context, we proceed to establish the stability of persistent stratified homotopy types when sampling around a compact

cylindrically stratified metric space as stability of persistent stratified homotopy types does generally not hold globally, as it does in the non-stratified case. We need a notion of local Lipschitz continuity for our upcoming exposition.

Definition 2.2.23. Let $K \in [0, \infty)$. We say that a function of (symmetric Lawvere) metric spaces $f : M \to M'$ is *K-Lipschitz (continuous) at $x \in M$*, if there is a $\delta > 0$, such that

$$\mathrm{d}(f(x), f(y)) \leq K\mathrm{d}(x, y),$$

for all $y \in M$ with $\mathrm{d}(x, y) \leq \delta$. We say *$f$ is K-Lipschitz (continuous)*, if it is K-Lipschitz for $\delta = \infty$, at every $x \in M$.

At this point it seems advisable to recall that we are still restricted to the two strata case. Let $P = p < q$ be a poset with two elements, and let $v = (v_l, v_h) \in \Omega$ and $u \in [0, 1]$ for the remainder of this section.

Proposition 2.2.24. *The map*

$$\mathcal{N} \colon \mathbf{Sam}_P \to \mathbf{Sam}_{\mathrm{N}(P)}$$

is 2-Lipschitz.

Proof. This is immediate from the triangle inequality. □

As an immediate consequence of the definition of the metric for strongly stratified samples, one obtains.

Lemma 2.2.25. *Let $\mathbb{S}, \mathbb{S}' \in \mathbf{Sam}_{\mathrm{N}(P)}$ and $v' \leq v \in [0, 1]$. Then*

$$(\mathbb{S}')_v^{v'} \subset (\mathbb{S}^{v'-\delta}_{v+\delta})_\delta,$$

for any $\delta > \mathrm{d}(\mathbb{S}, \mathbb{S}')$.

The dependency of the persistent stratified homotopy type on v introduces subtleties, primarily due to the absence of diagonal interleavings. Despite this challenge, the following technical lemma serves as the crucial argument in demonstrating stability.

Lemma 2.2.26. *Let $\mathbb{S}, \mathbb{S}' \in \mathbf{Sam}_{\mathrm{N}(P)}$. Let $\delta > \mathrm{d}(\mathbb{S}, \mathbb{S}')$ and suppose $v + \delta := (v_l - \delta, v_h + \delta) \in \Omega$ and $v - \delta := (v_l + \delta, v_h - \delta) \in \Omega$. Then*

$$\mathrm{d}(\mathcal{D}_v(\mathbb{S}), \mathcal{D}_v(\mathbb{S}')) \leq \delta + \max\{\mathrm{d}(\mathcal{D}_v(\mathbb{S}), \mathcal{D}_{v\pm\delta}(\mathbb{S}))\}.$$

Similarly, if $u \pm \delta \in (0, 1)$, then

$$\mathrm{d}(\mathcal{F}_u(\mathbb{S}), \mathcal{F}_u(\mathbb{S}')) \leq \delta + \max\{\mathrm{d}(\mathcal{F}_u(\mathbb{S}), \mathcal{F}_{u\pm\delta}(\mathbb{S}))\}.$$

Proof. We prove the diagram case, the other one can be shown completely analogously. Let $\alpha > \mathrm{d}(\mathcal{D}_v(\mathbb{S}), \mathcal{D}_{v\pm\delta}(\mathbb{S}))$. We then have inclusions

$$\mathcal{D}_v(\mathbb{S}) \hookrightarrow \mathcal{D}_{v-\delta}(\mathbb{S})_\alpha \hookrightarrow \mathcal{D}_v(\mathbb{S}')_{\delta+\alpha}\,,$$

$$\mathcal{D}_v(\mathbb{S}') \hookrightarrow \mathcal{D}_{v+\delta}(\mathbb{S})_\delta \hookrightarrow \mathcal{D}_v(\mathbb{S})_{\delta+\alpha}\,.$$

The upper left and lower right inclusion follow by the assumption on α. The lower left and upper right inclusions follow by Lemma 2.2.25. Hence, the result follows by considering the diagram distance as coming from a thickening flow as in Example 2.2.20. □

A way to interpret Lemma 2.2.26 is as the statement that the continuity of $\mathcal{D}_v$ in a strongly stratified sample $\mathbb{S}$ depends on the continuity of $\mathcal{D}_v(\mathbb{S})$ in the parameter v.

One consequence of the second part of Lemma 2.2.26 is captured in the next corollary for the purpose of referencing this result later on.

Corollary 2.2.27. *Let $\delta > 0$ such that $u \pm \delta \in (0,1)$. Let $\mathbb{S} = (\mathbb{X}, s) \in \mathbf{Sam}_{\mathrm{N}(P)}$ be such that $\mathbb{S}_{\leq u} = \mathbb{S}_{\leq u\pm\delta}$. Then,*

$$\mathcal{F}_u \colon \mathbf{Sam}_{\mathrm{N}(P)} \to \mathbf{Sam}_P$$

is 1-Lipschitz in $\mathbb{S}$ (on a ball with radius δ).

A handy fact is that the continuity of $\mathcal{D}_v(\mathbb{S})$ in v can be argued for by the continuity of the $\{p < q\}$ parts of diagrams as stated in the next lemma.

Lemma 2.2.28. *Let $\mathbb{S} \in \mathbf{Sam}_{\mathrm{N}(P)}$ and $v, v' \in \Omega$ and set $a = \min\{v_l, v'_l\}, b = \max\{v_h, v'_h\}$, then*

$$\mathrm{d}(\mathcal{D}_v(\mathbb{S}), \mathcal{D}_{v'}(\mathbb{S})) \leq \max\{\mathrm{d}(\mathbb{S}^{v_l}_{v_h}, \mathbb{S}^{v'_l}_{v'_h}), \mathrm{d}(\mathbb{S}^{a}_{v'_h}, \mathbb{S}^{a}_{v_h}), \mathrm{d}(\mathbb{S}^{v_l}_{b}, \mathbb{S}^{v'_l}_{b})\}.$$

Proof. This is an immediate consequence of the fact that

$$\mathrm{d}(\mathbb{X}, \mathbb{Y}) \leq \mathrm{d}(\mathbb{X} \setminus \mathbb{A}, \mathbb{Y} \setminus \mathbb{A}),$$

for $\mathbb{A} \subset \mathbb{X}, \mathbb{Y} \in \mathbf{Sam}$. □

For compact cylindrically stratified spaces we are now in position to show that $\mathcal{D}_v(\mathbb{S})$ and $\mathcal{F}_u(S)$ vary continuously in v and u respectively.

Proposition 2.2.29. *Let $S \in \mathbf{Sam}_{\mathrm{N}(P)}$ be compact and cylindrically stratified. Then*

$$(0,1) \to \mathbf{Sam}_P$$
$$u \mapsto \mathcal{F}_u(S)$$

and

$$\Omega \to \mathbf{D}_P\mathbf{Sam}$$
$$v \mapsto \mathcal{D}_v(S)$$

are continuous.

Proof. Note that it suffices to show the case of $\mathcal{D}_v$, since the nontrivial part of the continuity for $\mathcal{F}_u$ is given by the $(\mathcal{F}_u)_p$ component, and the latter is defined identically to the p-component of $\mathcal{D}_v(S)$. By Lemma 2.2.28 it suffices to show that for $v \to v^0$, we also have

$$\mathrm{d}(S^{v_l}_{v_h}, S^{v_l^0}_{v_h^0}) \to 0.$$

Next, note that the topology of the Hausdorff distance on the space of compact subspaces of a space only depends on the topology of the latter. Set $L := S^{\frac{1}{2}}_{\frac{1}{2}}$. Then by the cylinder assumption we may without loss of generality compute d in $L \times (0,1)$ equipped with the product metric. In particular, $S^{v_l}_{v_h} = L \times [v_l, v_h]$. We then have

$$\begin{aligned}\mathrm{d}(S^{v_l}_{v_h}, S^{v_l^0}_{v_h^0}) &= \mathrm{d}(L \times [v_l, v_h], L \times [v_l^0, v_h^0]) \\ &\leq \max\{|v_l - v_l^0|, |v_h - v_h^0|\} \xrightarrow{v \to v^0} 0.\end{aligned}$$

□

Invoking Proposition 2.2.29 and Lemma 2.2.26 we obtain the last result before stating one of our main results.

Corollary 2.2.30. *Let $S \in \mathbf{Sam}_{\mathrm{N}(P)}$ be compact and cylindrically stratified. Then*

$$\mathcal{F}_u \colon \mathbf{Sam}_{\mathrm{N}(P)} \to \mathbf{Sam}_P,$$
$$\mathcal{D}_v \colon \mathbf{Sam}_{\mathrm{N}(P)} \to \mathbf{D}_P\mathbf{Sam}$$

are continuous at S.
Furthermore, if $S_{\leq -} \colon (0,1) \to \mathbf{Sam}$ is K-Lipschitz in a neighborhood of u (and respectively S^-_- in a neighborhood of v), then $\mathcal{F}_u$ (and respectively $\mathcal{D}_v$) is $(K+1)$-Lipschitz at S.

In total, the next statement can be seen as a (slightly weaker) version of the classical, non-stratified Property (3) generalized to the persistent stratified homotopy type.

Theorem 2.2.31. *Let $S \in \mathbf{Sam}_P$ be compact and cylindrically stratified. Then*

$$\mathcal{P}_v \colon \mathbf{Sam}_P \to \mathrm{ho}\mathbf{Top}^{\mathbb{R}_+}$$

is continuous at S. Even more, if $S^-_- \colon \Omega \to \mathbf{Sam}$ is K-Lipschitz in a neighborhood of v, then $\mathcal{P}_v$ is $2(K+1)$-Lipschitz at S.

Proof. Recall that $\mathcal{P}_v = \mathcal{DP} \circ \mathcal{D}_v \circ \mathcal{N}$. Therefore, we simply collect that $\mathcal{N}$ is 2-Lipschitz by Proposition 2.2.24, $\mathcal{D}_v$ is continuous in $\mathcal{N}(S)$ by Corollary 2.2.30 and $\mathcal{DP}$ is 1-Lipschitz by Lemma 2.2.21. The second statement can be deduced in a similar fashion. □

Furthermore, the above statement can be significantly strengthened for the case of Whitney stratified spaces.

Theorem 2.2.32. *Let $P = \{p < q\}$ and $W \in \mathbf{Sam}_P$ be Whitney stratified with W_p compact. Then, for any $C > 1$, there exists some $R > 0$, such that the map*

$$\mathcal{P}_v : \mathbf{Sam}_P \to \mathrm{ho}\mathbf{Top}_P^{\mathbb{R}_+}$$

is $2(C+1)$-Lipschitz continuous at W, for all $v \in \Omega \cap (0, R)^2$.

Showing the above statement requires a detailed investigation of the methods found in [Hir69] in particular the concept of integral curves. We will discuss the concept of integral curves briefly at the time they will be necessary for our results in Section 3.3. The way to strengthen Theorem 2.2.31 for the case of Whitney stratified spaces is by showing that for any $C > 1$, there exists an $R > 0$, such that the function

$$\begin{aligned} \Omega \cap (0, R)^2 &\to \mathbf{D}_P\mathbf{Sam} \\ v &\mapsto \mathcal{D}_v(\mathcal{N}(W)) \end{aligned}$$

is C-Lipschitz continuous, where W is a Whitney stratified space with compact singular stratum, and then replace the K-Lipschitz continuity in the proof of Theorem 2.2.31 with the C-Lipschitz continuity. Because we feel we sufficiently demonstrated the main properties of the concept of persistent stratified homotopy types we would like to refer the reader to Appendix A.1 to see the technical details that go into proving Theorem 2.2.32.

Chapter 3

Approximate Stratifications

Clearly, there is no hope to be able to infer every stratification that one can define on a space. Even a manifold, like S^2, can be a stratified space by simply labelling a collection of closed subsets of S^2. Consider for example Fig. 3.1 where a little more than the upper hemisphere of S^2 and the whole of S^2 serve as strata. In this case the stratification is not based on a local geometric differences of the strata. In order to stratify a space without a priori knowledge on the actual strata and links, the stratification has to be measurable by local geometric means as in case of conically stratified spaces. Another classical example is given by homology stratifications, as used by Goresky and MacPherson in [GM83]. Because it will be a central concept, especially for the algorithmic and application chapters, we formally define local homology first. The results presented in this chapter appeared in an article written by the author together with Lukas Waas [MW22]. This chapter roughly corresponds to [MW22, Section 4] and certain results and definitions are incorporated verbatim.

Definition 3.0.1. The *local homology groups* at a point x of a space X are given by relative homology

$$\mathrm{H}_\bullet(X; x) := \mathrm{H}_\bullet(X, X \setminus \{x\}).$$

Local homology groups can be computed with direct limits

$$\varinjlim H(X, X \setminus U)$$

where the limit ranges over the open subsets of X containing x.

Example 3.0.2. For us it suffices to discuss the two strata case here but the concept of homology stratification is more general. Suppose $S = (X, s\colon X \to \{p < q\})$ is conically stratified, such that S_q is locally euclidean of dimension q, and S_p of dimension p. In particular, $x \in S_p$ admits a neighborhood

$$U \cong_P \mathbb{R}^p \times C(L)$$

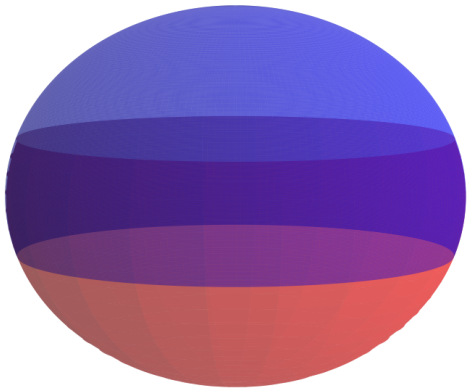

Figure 3.1. Artificial stratification of S^2

for some $q-(p+1)$ dimensional compact manifold L, called the link of x. Suppose further that L is not a homology sphere, and that the strata are connected. Then, the stratification of S can be recovered from the underlying space as follows.

By the assumption on the local geometry of X, for any $x \in X$, there exists a small open neighborhood U_x such that the natural map

$$\mathrm{H}_\bullet(X, X \setminus U_x) \to \mathrm{H}_\bullet(X; x)$$

is an isomorphism. In particular, for each $x \in X$ one obtains natural maps

$$\mathrm{H}_\bullet(X; x) \cong \mathrm{H}_\bullet(X, X \setminus U_x) \to \mathrm{H}_\bullet(X; y)$$

for $y \in U_x \cong_P \mathbb{R}^p \times C(L)$. If x, y are contained in the same stratum, then all of these maps are given by isomorphisms. By the path connectedness assumption any two points in the same strata are connected by such a sequence of isomorphisms. Conversely, since we assumed that L is not a homology sphere, we have

$$\mathrm{H}_\bullet(X; x) \cong \mathrm{H}_\bullet(U_x; x) \cong \tilde{\mathrm{H}}_{\bullet-(p+1)}(L) \neq \mathrm{H}_\bullet(X; y),$$

whenever $x \in S_p$ and $y \in S_q$. Thus, we can reobtain the stratification of S, by assigning to points the same stratum, if and only if they are connected through such a sequence of isomorphisms. Stratifications with the property that all the induced maps of local homologies on a stratum are isomorphisms are called *homology stratifications*.

Local homology as a means to obtain stratifications of point clouds (or combinatorial objects) have recently been investigated in several works ([BWM12,SW14,FW16, Nan20,STHN20,Mil21]). Both [BWM12] and [Nan20] make use of the structure maps $\mathrm{H}_\bullet(X;x) \to \mathrm{H}_\bullet(X;y)$ to determine the strata. Note however, that in the case of two strata, it really suffices to study the isomorphism type at each point, and there is no need to study the maps themselves, as stated by the following lemma.

Lemma 3.0.3. *Let $S = (X, s\colon X \to \{p < q\})$ be a conically stratified space with manifold strata of dimension q and p respectively. Then s is a homology stratification.*

Furthermore, if the local homology of X is different from $\mathrm{H}_\bullet(\mathbb{R}^q;0)$, at each $y \in S_p$, then s is the only homology stratification of X with two strata.

Conversely, one always obtains a homology stratification $\tilde{s} : X \to \{p < q\}$ defined by:

$$\tilde{s}(x) = q \iff \mathrm{H}_\bullet(X;x) \cong \mathrm{H}_\bullet(\mathbb{R}^q;0),$$

for $x \in X$.

Proof. The first result is immediate from the local conical structure of X. The second is immediate from the definition of a homology stratification, as clearly $X - S_p$ is a homology manifold. For the final result, note first that by the local conical structure, having local cohomology isomorphic to having $\tilde{\mathrm{H}}_\bullet(S^q)$ is an open condition on S_p. In particular, since this condition holds on all of $X - S_p$ it is an open condition on all of X. Thus, $s : X \to \{p < q\}$ as defined in the statement is actually a stratification of X. To see that this is indeed a homology stratification we need to see that the local isomorphism condition is fulfilled. By construction, we have $X - S_p \subset \tilde{s}^{-1}\{q\}$. Within $S_p \cap \tilde{s}^{-1}\{p\}$ the local isomorphism condition again holds by the local conical structure of X. Thus, it remains to consider the case where $x \in S_p$, and $\mathrm{H}_\bullet(X;x) \cong \tilde{\mathrm{H}}_\bullet(S^q)$. We need to show that, for $U_x \cong \mathbb{R}^{q-p-1} \times \mathring{C}(L_x)$, an open neighborhood of x, the natural map

$$\mathrm{H}_\bullet(X;x) \cong \mathrm{H}_\bullet(W, W - U_x) \to \mathrm{H}_\bullet(X;y)$$

is an isomorphism, for all $y \in U_x$. The only nontrivial degree in this case is $q = \dim W$. By an application of the Künneth formula L_x is again an orientable manifold. Hence, up to suspension, from this perspective, the claim reduces to the fact that if L_x is an orientable, closed manifold. Then, under the natural isomorphism

$$\mathrm{H}_\bullet(CL_x, L_x) \cong \tilde{\mathrm{H}}_{\bullet-1}(L_x)$$

the fundamental class of L_x induces a fundamental class of CL_x. □

This short excursion on local homology and homology stratifications serves as motivation to take a more in-depth look at conditions that have to be satisfied for methods such as local homology to distinguish strata. In the following we will also give a few options to choose from in practice that measure local geometric differences in a space.

3.1 Extrinsic tangent cones and magnifications

The distinction of singular and regular parts of a space will mostly be based on local information, such as local homology. As most of the spaces we consider are Whitney stratified and are thus embedded in Euclidean space we often want to compare spaces locally, based on intersections with Euclidean balls of some radius and around some center point. Making this information independent of the radius and location within ambient spaces would involve tedious spatial shifts, truncations and normalizations. For convenience and to keep notation neat, we will use the following:

Definition 3.1.1. Denote by $\mathbf{Sam}_\star$ the space of centered subspaces of in $\mathbb{R}^N$,

$$\mathbf{Sam}_\star := \{\mathbb{X} \mid \mathbb{X} \subset \mathbb{R}^N, x \in \mathbb{R}^N\}$$

equipped with the (extended pseudo) metric induced by pulling back the metric on **Sam** along

$$\begin{aligned} \mathbf{Sam}_\star &\to \mathbf{Sam} \\ \mathbb{X} &\mapsto \mathrm{B}_1(0) \cap \mathbb{X}. \end{aligned}$$

We call $\mathbf{Sam}_\star$ the *space of centered samples* (of $\mathbb{R}^N$).

Observe that the metric on $\mathbf{Sam}_\star$ is structured in such a way that it inherently identifies translations of a space, as well as a space with its intersection with the unit ball centered at the origin. In fact, as per its definition, $\mathbf{Sam}_\star$ is isometric to the space of subspaces within $\mathrm{B}_1(0) \subset \mathbb{R}^N$.

Remark 3.1.2. Let $\mathbb{X}$ represent a (potentially noisy) sample from a space X. Initially, let's refrain from assuming that either $\mathbb{X}$ or X is centered, but we will assume that both X and $\mathbb{X}$ belong to the set **Sam**. In this context, to derive non-trivial information, one cannot proceed all the way to the limit when computing (persistent) local homology. Specifically, for any thickening $\mathbb{X}_\varepsilon$, where $\varepsilon > 0$, $H_*(\mathbb{X}_\varepsilon, p) = H_*(\mathbb{R}^n, 0)$. Instead, one considers the persistent relative homology of the sample relative to a ball of radius

r around any chosen point (see [BCSE$^+$07, BWM12, SW14]). In other words, one computes persistent local homology using the spaces

$$(\mathbb{X}_\varepsilon, \mathbb{X}_\varepsilon \setminus B_r(p)).$$

For computational reasons, it is beneficial to use the intrinsically local notion of this structure. By the excision theorem, one may equivalently compute:

$$(\mathbb{X}_\varepsilon \cap \mathrm{B}_r(p), (\mathbb{X}_\varepsilon \cap \mathrm{B}_r(p)) \setminus \mathring{\mathrm{B}}_{\frac{r}{2}}(p)).$$

Now, at this point, the measurable outcome of the homology of the above space depends on the chosen radius r and will decrease as $r \to 0$. If we desire a measure of singularity that is comparable for different scales, then this needs to be normalized. One may use $\mathbb{M} = \frac{1}{r}(\mathbb{X} - p) \cap \mathrm{B}_1(0)$ and compute the homology of the stretched pair

$$(\mathbb{M}_\varepsilon, \mathbb{M}_\varepsilon \setminus \mathring{\mathrm{B}}_{\frac{1}{2}}(0)).$$

The above excursion is the inspiration for the following definition that will make the notation more compact. However, we do this so that the center point remains the same.

Definition 3.1.3. Let $\mathbb{X} \in \mathbf{Sam}$, $x \in \mathbb{R}^N$ and $r > 0$. We denote by

$$\mathcal{M}_x^r(\mathbb{X}) := \frac{1}{r}(\mathbb{X} - x) \cap \mathrm{B}_1(0) \in \mathbf{Sam}_\star$$

the *r-magnification of $\mathbb{X}$ at x*.

With the next example we want to illustrate the concept we will use to decide whether or not we are able to infer stratifications on a space.

Example 3.1.4. Consider the two real algebraic varieties

$$X = \{(x, y) \in \mathbb{R}^2 \mid (x^2 + y^2)^2 + 2x^2 + 2y^2 - 4x^2 = 0\}$$

and

$$Y = \{(x, y) \in \mathbb{R}^2 \mid (x^2 + y^2)^2 - 4x^2 = 0\}.$$

These varieties are Whitney stratified spaces with the singular set containing only the singular point located at the origin. In Fig. 3.2, we plotted three local magnifications of X at the origin $x = (0, 0)$, i.e. $\mathcal{M}_x^r(X)$ for three different r . We can observe that the homeomorphism type of the local magnifications stabilize as we decrease r. However, at the same time, the spaces $\mathcal{M}_x^r(X)$ converge for $r \to 0$, to a space homeomorphic to the previous (truncated) magnifications. In contrast, Y shows a different convergence behavior. Although the spaces $\mathcal{M}_x^r(Y)$ share the same homeomorphism type with the local magnifications of X at the origin for r small enough, Fig. 3.3 illustrates that the homeomorphism type changes when passing to the limit.

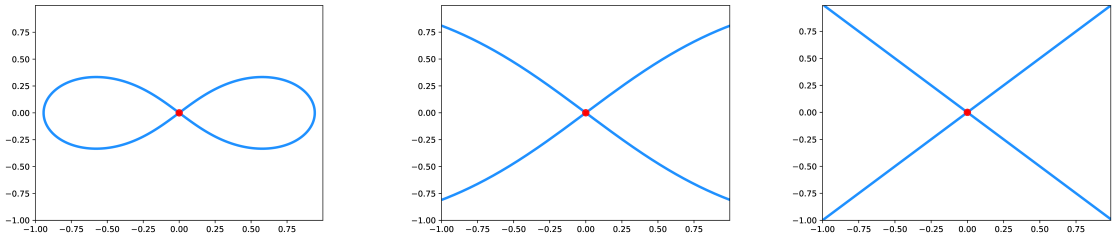

Figure 3.2. Three magnifications of $(x^2+y^2)^2+2x^2+2y^2-4x^2=0$ at the origin

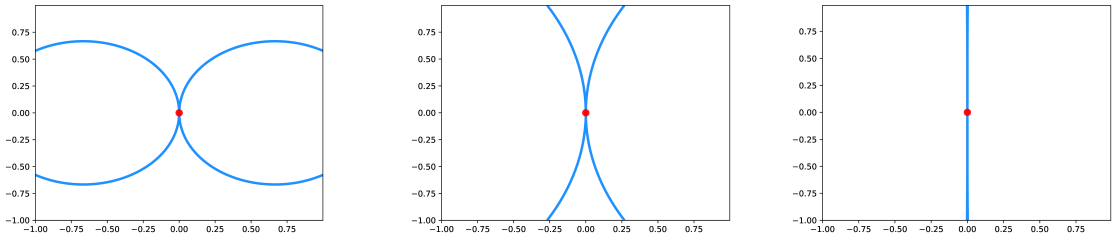

Figure 3.3. Three magnifications of $(x^2+y^2)^2-4x^2=0$ at the origin

One can immediately make out the limitations of this approach since even a singular space such as shown in Fig. 3.3 does not fall into the class of stratified spaces we can reliable investigate by local means. However, the limit spaces as described in Example 3.1.4 exist and are known as the (extrinsic) tangent cones of X at x, in the case where X admits a Whitney stratification and fulfills certain extra conditions as we will see in Proposition 3.4.4. For a more detailed investigation on metric tangent cones see [Lyt04, BL07]).

Definition 3.1.5. Let $X \subset \mathbb{R}^N$. The *(extrinsic) tangent cone of* X *at* $x \in X$ is defined as

$$\mathrm{T}^{\mathrm{ex}}_x(X) := \{v \in \mathbb{R}^N \mid \forall \varepsilon > 0 \exists y \in \mathrm{B}_\varepsilon(x) \cap X : v \in (\mathbb{R}_{\geq 0}(y-x))_\varepsilon\},$$

where $-_\varepsilon$ is to be understood as in Example 2.2.20. The extrinsic tangent cones define a map

$$\begin{aligned} \mathrm{T}^{\mathrm{ex}}(X) : X &\to \mathbf{Sam}_\star \\ x &\mapsto \mathrm{T}^{\mathrm{ex}}_x(X) \end{aligned}$$

Example 3.1.6. Invoking Taylor's expansion theorem yields

$$\mathrm{T}^{\mathrm{ex}}_x(X) = \mathrm{T}^{\mathrm{ex}}_x(U) = \mathrm{T}_x(U)$$

when $U \subset X \subset \mathbb{R}^N$ an open neighborhood of x such that U is a smooth submanifold of $\mathbb{R}^N$.

Example 3.1.7. For an (affine) complex algebraic variety X the tangent cone at the origin coincides with the algebraic tangent cone, i.e. the set of common zeroes of all polynomials in the ideal generated by the homogeneous elements of lowest degree of all polynomials that vanish identically on X.

It is a classical result, which can for example be found in [Hir69, BL07], that when X admits a subanalytic Whitney stratification, then

$$\mathcal{M}^r_x(X) \to \mathrm{T}^{\mathrm{ex}}_x(X)$$

in $\mathbf{Sam}_\star$ as $r \to 0$. This already gives an idea of what local geometric information one can obtain from methods such as local homology. Also, with the concept of tangent cones at hand we can define the central notion of spaces for which one can hope to infer stratifications by local geometric information.

Definition 3.1.8. Let $P = \{p < q\}$. Let $W = (X, s\colon X \to P) \in \mathbf{Sam}_P$ be q-dimensional Whitney stratified space. We say that W is *tangentially stratified* if

$$\mathrm{d}(\mathrm{T}^{\mathrm{ex}}_x(W), V) > 0,$$

for all $x \in W_p$ and for all q-dimensional linear subspaces $V \subset \mathbb{R}^N$.

The following example illustrates that it is not hard to find Whitney stratified spaces that are not tangentially stratified.

Example 3.1.9. Consider, again, $Y = \{(x, y) \in \mathbb{R}^2 \mid (x^2 + y^2)^2 - 4x^2 = 0\}$ from Example 3.1.4. In this case, the above condition specifies to $\mathrm{d}(\mathrm{T}^{\mathrm{ex}}_{(0,0)}(Y), V) > 0$, for all 1-dimensional subspaces $V \subset \mathbb{R}^N$. The tangent cone of Y at the origin is a 1-dimensional linear space given by

$$\mathrm{T}^{\mathrm{ex}}_{(0,0)}(Y) = \{(x, y) \in \mathbb{R}^2 \mid x = 0\},$$

see Fig. 3.4 on the right, which already serves as a subspace $V \subset \mathbb{R}^2$ such that $\mathrm{d}(\mathrm{T}^{\mathrm{ex}}_{(0,0)}(Y), V) = 0$. For the space

$$X = \{(x, y) \in \mathbb{R}^2 \mid (x^2 + y^2)^2 + 2x^2 + 2y^2 - 4x^2 = 0\}$$

on the other hand we find that the tangent cone at the origin is a 1-dimensional space given by

$$\mathrm{T}^{\mathrm{ex}}_{(0,0)}(X)\{(x, y) \in \mathbb{R}^2 \mid (x + y)(x - y) = 0\},$$

see Fig. 3.4 on the left. Clearly, there is no 1-dimensional subspace $V \subset \mathbb{R}^2$ such that $\mathrm{d}(\mathrm{T}^{\mathrm{ex}}_{(0,0)}(X), V) = 0$.

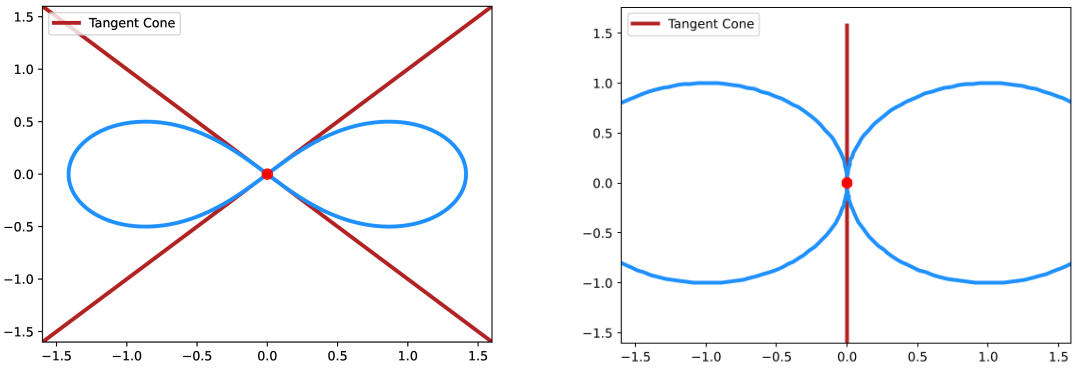

Figure 3.4. Two curves with their respective tangent cones at their singular stratum

In Section 3.5 we will see that for certain tangentially stratified spaces the stratification can be approximated arbitrarily accurate from a sufficiently good sample and by zooming-in far enough. Therefore, tangentially stratified spaces are the main object for our theoretical considerations but approximating a stratification from a sample is feasible also by other means.

3.2 Φ-stratification

To obtain a practical access to stratifying a space we introduce the notion of tangentially Φ-stratified spaces. This extension on the concept of tangentially stratified space, reflects that in practice there are other local invariants usable to stratify a point cloud.

Definition 3.2.1. Let $\Phi : \mathbf{Sam}_\star \to [0,1]$ be a continuous function, such that $\Phi(V) = 1$, whenever V is a q-dimensional linear subspace of $\mathbb{R}^N$. Let $W \in \mathbf{Sam}_P$ be q-dimensional Whitney stratified space. We say that W is (tangentially) Φ-*stratified* if

$$\Phi(\mathrm{T}_x^{\mathrm{ex}}(W)) < 1,$$

for all $x \in W_p$.

Let us begin with a series of examples of functions $\Phi : \mathbf{Sam}_\star \to [0,1]$ that can be used (in theory) to detect local singularity. The prime example for the input to a function Φ are the magnifications of $\mathbb{X} \in \mathbf{Sam}_\star$ at every point $x \in \mathbb{X}$. We can think of $\mathcal{M}_x^r(\mathbb{X}) \in \mathbf{Sam}_\star$ as zooming into $\mathbb{X}$ at x by a magnification parameter r. To such local data we want to apply a function Φ to determine how far from a q-dimensional Euclidean unit disk $D^q \subset \mathbb{R}^q \hookrightarrow \mathbb{R}^N$ the space $\mathcal{M}_x^r(\mathbb{X})$ is.

Example 3.2.2. Consider the continuous map

$$\Phi^q_{Hd}\colon \mathbf{Sam}_\star \to [0,1] \tag{3.1}$$

$$\mathbb{B} \mapsto 1 - \inf\{\mathrm{d}(\mathbb{B}, V)\}, \tag{3.2}$$

where V ranges over the q-dimensional linear subspaces of $\mathbb{R}^N$. This is a continuous map as the distance function on domain and target is the Hausdorff distance. We promised the concept of Φ-stratification to be a practical concept and the reader is referred to Section 4.2 for the computational consideration of the map Φ^q_{Hd}. A q-dimensional Whitney stratified space $W \in \mathbf{Sam}_P$ is tangentially stratified if and only if it is Φ^q_{Hd}-stratified. Thus, Φ^q_{Hd} is universal in the sense that if W is Φ-stratified for some Φ as in Definition 3.2.1, then it is Φ^q_{Hd}-stratified.

An example of Φ-stratification that, in some sense, is already frequently considered in the field of TDA [BWM12, Mil21, Nan20, STHN20] is given by persistent local homology

Example 3.2.3. Define

$$\mathcal{P}\mathrm{L}_\bullet : M \mapsto \{\varepsilon \mapsto \mathrm{H}_\bullet(M_\varepsilon, M_\varepsilon \setminus \mathring{\mathrm{B}}_{\frac{1}{2}}(0))\}$$

to obtain a persistence module indexed over $[0, \frac{1}{2})$. Note that no bar in a persistence module $\mathcal{P}\mathrm{L}_i$ can be longer than $\frac{1}{2}$ as the diameter of a space in $\mathbf{Sam}_\star$ is at most 1. Therefore, the following map is well-defined:

$$\Phi^q_{\mathcal{P}\mathrm{L}}\colon \mathbf{Sam}_\star \to [0,1] \tag{3.3}$$

$$\mathbb{B} \mapsto 1 - 2\max_{i \leq q} d(\mathcal{P}\mathrm{L}_i(\mathbb{B}), \mathcal{P}\mathrm{L}_i(\mathbb{R}^q)). \tag{3.4}$$

where $\mathbb{R}^q \subset \mathbb{R}^N$ with a linear embedding in order to ensure $\mathbb{R}^q \in \mathbf{Sam}_\star$. Let $W \in \mathbf{Sam}_P$ be a (definable) q-dimensional Whitney stratified such that for all $x \in W_p$ we have $\mathcal{P}\mathrm{L}_\bullet(\mathrm{T}^{\mathrm{ex}}_x(W)) \neq \mathcal{P}\mathrm{L}_\bullet(\mathbb{R}^q)$, i.e. the two persistence modules are not isomorphic. Then, W is a Φ-stratified space.

As indicated before, one of the strengths of the general definition of Φ-stratified spaces is that in practice one may use a selection of possibly rougher or approximative invariants that allow algorithmic computation.

Example 3.2.4. Instead of using

$$\mathbb{B} \mapsto 1 - \inf\{\mathrm{d}(\mathbb{B}, V)\},$$

as in Example 3.2.2, one can only use the asymmetric Hausdorff distance, i.e. only consider

$$(\mathbb{X}, x) \mapsto \inf\{\varepsilon \mid \mathrm{B}_1(0) \cap (\mathbb{X} - x) \subset \mathrm{B}_1(0)) \cap V_\varepsilon\}$$

While this does narrow down the class of Φ-stratified spaces, the use of the asymmetric distance not only simplifies the computation itself but also enables the formulation of a deterministic optimization problem to identify a minimizing linear subspace (refer to Eq. (4.18)). In contrast to the general problem presented in Eq. (3.1), an approximation of the infimum can be obtained by randomly generating q-dimensional linear subspaces and recording the smallest distance (see Section 4.2.3).

Furthermore, rather than computing $\mathcal{PL}_\bullet(\mathbb{B})$ as illustrated in Example 3.2.3, one may opt for a Vietoris-Rips version of the latter, as discussed later in Section 4.1.2. Empirically, this Vietoris-Rips version of $\Phi^q_{\mathcal{PL}}$ has demonstrated its ability to distinguish a linear subspace from a singular region. We will also provide arguments supporting its efficacy in approximating the true function $\Phi^q_{\mathcal{PL}}$ (see Section 4.1.2). Additionally, focusing on specific dimensions of interest may yield further computational advantages.

Moreover, rather than utilizing the challenging-to-compute interleaving distance (compare to [BB18]), we will employ the Bottleneck distance. While the Bottleneck distance theoretically yields the same results as the interleaving distance, it is also computable and, notably, easily so in the special case of Eq. (3.3) (see Remark 4.1.8).

3.3 Lojasiewicz-Whitney stratified spaces

The preceding section emphasizes the necessity of gaining a more profound understanding of the convergence properties of magnification spaces to tangent cones for the purpose of reconstructing stratifications from sample data. These insights form the core content of Section 3.4. However, before delving into these aspects, we must establish a series of results concerning Whitney stratified spaces, which are defined within particularly well-behaved o-minimal structures.

Our approach heavily relies on the foundational work presented in [Hir69] and [BL07]. In particular, we will make use of the concept of *integral curves*. Prior to delving into this, we need to introduce some notation.

Definition 3.3.1. Let $V, U \subset \mathbb{R}^N$ be linear subspaces. The (asymmetric) *distance of* V *to* U is given by

$$\vec{d}(V,U) = \sup_{v\in V, ||v||=1} \inf\{||v-u|| \mid u \in U\} = \sup_{v\in V, ||v||=1} \{||\pi_{U^\perp}(v)||\},$$

where $\pi_{U^\perp}$ denotes the orthonormal projection to the orthogonal complement of U.

Whitney's condition (b) can be expressed in terms of a function, which measures the failure of secants being contained in the tangent space, as follows (compare [Hir69]).

Construction 3.3.2. *Let $S = (X, s\colon X \to P)$ be a stratified space with smooth strata, contained in $\mathbb{R}^N$. Consider the function*

$$\beta\colon X \times X \to \mathbb{R}; \qquad \begin{cases} (x,y) & \mapsto \vec{d}(l(x,y), \mathrm{T}_x(X_{s(x)})) \text{ if } x \neq y, \\ (x,x) & \mapsto 0 \text{ else} \end{cases}$$

where we consider all tangent spaces involved as linear subspaces of $\mathbb{R}^N$.

Note that for the time being we will not assume $P = \{p < q\}$. We will specifically say when we make the two strata restriction again. The results presented here and in Section 3.4 hold for general P. The following proposition basically describes the properties of what is called *integral curves* (compare to [Hir69]).

Proposition 3.3.3. *Let $S = (X, s\colon X \to P)$, be as in the assumption of Construction 3.3.2 and further so that the frontier and local finiteness condition are fulfilled. Then, S is a Whitney stratified space if and only if*

$$\beta\,|_{(X_q \times X_p) \cup \Delta_{X_p}}\colon (X_q \times X_p) \cup \Delta_{X_p} \to \mathbb{R},$$

is continuous, for all pairs $q \geq p \in P$. Here, Δ_{X_p} denotes the diagonal of $X_p \times X_p$,

Proof. Although this assertion is somewhat of a common understanding within the field we want to ensure a comprehensive coverage and provide a proof in Appendix A.2. □

Next, we need the notion of integral curves, as defined, for example, in [Hir69].

Proposition 3.3.4. *[Hir69, Lemma 4.1.1] Let W be a Whitney stratified space over the poset P and $y \in W_p$, for some $p \in P$. Let $B = \mathrm{B}_d(y) \subset \mathbb{R}^N$ be a ball of radius d around y such that $\beta(-, y)$ is bounded uniformly by some $\delta < 1$ on $W_{\geq p} \cap B$. Then, for any $x \in W_{\geq p} \cap B$, $x \neq y$, there exists a unique curve $\phi\colon [0, d] \to W \cap B$, fulfilling*

1. *$\phi(0) = y$ and $\phi(||y - x||) = x$,*
2. *ϕ is almost everywhere differentiable. At differentiable points, $t \neq 0$, the differential is given by*

$$\phi'(t) = \frac{||\phi(t) - y||}{||\pi_{\phi(t)}(\phi(t) - y)||^2} \pi_{\phi(t)}(\phi(t) - y),$$

where $\pi_{\phi(t)}$ denotes the projection to $\mathrm{T}_{\phi(t)}(W_{s(\phi(t))})$.

Definition 3.3.5. A curve ϕ as in Proposition 3.3.4 is called the *integral curve associated to the pair x, y.*

The following result due to Hironaka gives us additional amount of control over integral curves which we will need in order to control the convergence behavior of magnifications of samples.

Lemma 3.3.6. *Let W be a Whitney stratified space over P, $p \in P$ and $y \in W_p$. Suppose there exists $d_0 > 0$ such that there exists $\alpha > 0$, with*

$$\beta(x, y) \leq ||y - x||^{\alpha}$$

for all $x \in W_{\geq p} \cap \mathrm{B}_{d_0}(y)$. Then, for any $C > 0$ there exists $d > 0$ only depending on d_0, α (and the dimension of W), such that for any integral curve $\phi : [0, d] \to W$ starting in y and ending in $\mathrm{B}_d(y)$, the inequality

$$||\frac{1}{t}(\phi(t) - \phi(0)) - \frac{1}{s}(\phi(s) - \phi(0))|| \leq C|t - s|^{\alpha}$$

holds for all $t, s \in [0, d]$. In particular, all integral curves starting at y are differentiable in 0.

Proof. A complete proof of this statement can be found in [Hir69] □

Spaces that satisfy a local variant of the condition described above were examined in [Hir69]. In that work, it was referred to as the *strict Whitney condition*.

Definition 3.3.7. A Whitney stratified space fulfilling the requirements of Lemma 3.3.6 on any compactum K contained in some pure stratum W_p of W, is called a *Lojasiewicz-Whitney stratified space*. That is, W is called Lojasiewicz-Whitney stratified, if the following condition holds. Let $K \subset W_p$ be a compact, definable subset of some stratum W_p of W. Then there exist $\alpha > 0, d > 0$ such that

$$\beta(x, y) \leq |y - x|^{\alpha},$$

for all $y \in K$ and $x \in W \cap \mathrm{B}_d(y)$.

Put differently, Lojasiewicz-Whitney stratified spaces are Whitney stratified spaces where the rate at which secant lines deviate from the tangent spaces is constrained by a certain root.

Remarkably, the majority of the stratified spaces of interest (our interest), such as compact subanalytic or semialgebraic spaces, fall in the category of Lojasiewicz-Whitney stratified (Proposition 3.3.9).

Recollection 3.3.8. Recall that within an o-minimal structure, we refer to it as polynomially bounded if, for all $f : \mathbb{R} \to \mathbb{R}$ definable in the structure, there exists an $n \in \mathbb{N}$ such that for sufficiently large t,

$$|f(t)| \leq |t|^n.$$

Structures falling into the category of polynomially bounded include the structure of semialgebraic sets and finitely subanalytic sets (see [vdD86] and [Mil94]). Importantly, any compact subanalytically definable stratified space is definable within a polynomially bounded o-minimal structure.

To keep the pace towards our main results, we migrated the proof of the following statement to the appendix (see Appendix A.5).

Proposition 3.3.9. *Let W be a Whitney stratified space which is definable with respect to some polynomially bounded o-minimal structure. Then, W is Lojasiewicz-Whitney stratified.*

As an almost immediate consequence of Lemma 3.3.6 and Proposition 3.3.9 we obtain:

Proposition 3.3.10. *Let W be a Lojasiewicz-Whitney stratified space. Then, for any $x \in W$ every integral curve starting at x is differentiable in 0. Furthermore, we have*

$$\mathrm{T}_x^{\mathrm{ex}}(W) \cap \partial \mathrm{B}_1(x) = \overline{\{\phi'(0) \mid \phi \text{ is an integral curve starting at } x\}}.$$

Hence,

$$\mathrm{T}_x^{\mathrm{ex}}(W) = \overline{\{\alpha\phi'(0) \mid \phi \text{ is an integral curve starting at } x, \alpha \geq 0}.$$

Proof. Note first, that $\mathrm{T}_x^{\mathrm{ex}}(W)$ is closed by definition. The containment of the right hand side in the left hand side is immediate by definition of the derivative. For the converse inclusion, let $v \in \mathrm{T}_x^{\mathrm{ex}}(W) \cap \partial \mathrm{B}_1(x)$. For $\varepsilon > 0$ small enough we have $y \in W_{\geq p} \cap \mathrm{B}_\varepsilon(x)$ with $p = s(x)$ such that

$$||v - \lambda(y - x)|| < \varepsilon$$

for some $\lambda > 0$. We then also obtain

$$\big|1 - \lambda||y - x||\big| < \varepsilon.$$

Now, $t = ||y - x||$ and let $\phi : [0, d] \to W$ be the integral curve starting at x and passing through y. We then have

$$\begin{aligned} ||v - \phi'(0)|| &\leq ||v - \lambda(y - x)|| + ||\lambda(y - x) - \frac{(y - x)}{||y - x||}|| + ||\frac{(y - x)}{||y - x||} - \phi'(0)|| \\ &= ||v - \lambda(y - x)|| + \big|1 - \lambda||y - x||\big| + ||\frac{\phi(t) - x}{t} - \phi'(0)|| \\ &\leq \varepsilon + \varepsilon + C\varepsilon^\alpha, \end{aligned}$$

for some C, α independent of the choices above. In particular, we can choose ϕ such that $\phi'(0)$ is arbitrarily close to v and hence it lies in the closure. □

Next, we have a key technical result for to investigate the convergence behavior of magnifications of samples.

Proposition 3.3.11. *Let W be a Lojasiewicz-Whitney stratified space over P. Let $p \in P$ and $K \subset W_p$ be a compact subset. Then, there exist $d, C, \alpha > 0$, such that the following holds.*
For all r, such that $\frac{1}{r} \in [0, d]$ there exists $\varepsilon_0 > 0$, such that

$$\mathrm{d}(\mathrm{T}^{\mathrm{ex}}_x(W), \mathcal{M}^r_w(\mathbb{W})) \leq C(r^\alpha + \frac{\varepsilon}{r}),$$

for $\mathbb{W} \in \mathbf{Sam}$ with $\mathrm{d}(\mathbb{W}, W) = \varepsilon \leq \varepsilon_0$, $w \in \mathbb{R}^N$ and $x \in K$ with $|x - w| \leq \varepsilon$.

Proof. We work with the non-normalized spaces instead, that is instead of working in the unit ball of $\mathcal{M}^r_x(\mathbb{W})$, we work in the ball of radius r in $\mathbb{W}$. Furthermore, without loss of generality let $x = 0$. Again, choose d, C', α as in Lemma 3.3.6, possibly slightly decreasing d, such that the requirements on r still hold for $r + 2\varepsilon$. Let $c \in \mathrm{T}^{\mathrm{ex}}_x(W)$ with $|c| \leq r$. Let $\tilde{c} := \frac{r-2\varepsilon}{r} c$. We have

$$|c - \tilde{c}| \leq 2\varepsilon.$$

Next, using Proposition 3.3.10, consider the integral curve starting in 0 with initial direction $\frac{c}{|c|}$, $\phi : [0, d] \to W$ (or, by passing to the limit if necessary a curve with initial direction arbitrarily close to $\frac{c}{|c|}$). We then have

$$|\tilde{c} - \phi(|\tilde{c}|)| \leq C' r^{\alpha+1}$$

and

$$|\phi(|\tilde{c}|) - w| \leq |\tilde{c}| + \varepsilon \leq r - \varepsilon. \tag{3.5}$$

Choose $w' \in \mathbb{W}$ with $|w' - \phi(|\tilde{c}|| \leq \varepsilon$. Then, by, Eq. (3.5) $w' \in \mathrm{B}_r(w) \cap \mathbb{W}$. Summarizing, we have

$$|c + w - w'| \leq 2\varepsilon + \varepsilon + C' r^{\alpha+1} + \varepsilon \leq C(r^{\alpha+1} + \varepsilon),$$

for appropriate $C > 0$.

Conversely, let $w' \in \mathbb{W}$ with $|w - w'| \leq r$. By assumption, we find $y \in W$ with $|y - w'| \leq \varepsilon$ and have $|y| \leq r + 2\varepsilon$. Thus, for ϕ_y the integral curve starting in 0 through y we have

$$||y|\phi'_y(0) - y| \leq C'(r + 2\varepsilon)^{\alpha+1}.$$

Take $c = (r - \varepsilon)\frac{|y|}{r+2\varepsilon}\phi'_y(0) \in \mathrm{T}^{\mathrm{ex}}_x(W) \cap \mathrm{B}_{r-\varepsilon}(x)$. Note, that $|c + w| \leq |w| + |c| \leq r$ i.e. $c + w \in \mathrm{B}_r(w) \cap (\mathrm{T}^{\mathrm{ex}}_x(W) + w)$. We further have

$$|c - |y|\phi'_y(0)| \leq |y|(1 - \frac{r - \varepsilon}{r + 2\varepsilon}) \leq 3\varepsilon.$$

Summarizing, for $\varepsilon < r/2$, we have

$$\begin{aligned}|c+w-w'| \leq \varepsilon + |c-w'| &\leq \varepsilon + |c-|y|\phi_y'(0)| + ||y|\phi_y'(0)-y| + |y-w'| \\ &\leq \varepsilon + 4\varepsilon + C'(r+2\varepsilon)^{\alpha+1} \\ &\leq C(r^{\alpha+1}+\varepsilon)\end{aligned}$$

for appropriate $C > 0$. We obtain the result by multiplying with $\frac{1}{r}$ to pass to the magnification. □

As a first corollary of Proposition 3.3.11, we obtain that the tangent cones of a Lojasiewicz-Whitney stratified space vary continuously on each stratum.

Proposition 3.3.12. *Let W be a Lojasiewicz-Whitney stratified space over P and $p \in P$. Then, the map*

$$\begin{aligned}\mathrm{T}^{\mathrm{ex}}_{-}(W) : W_p &\to \mathbf{Sam}_\star \\ x &\mapsto \mathrm{T}^{\mathrm{ex}}_x(W)\end{aligned}$$

is continuous.

Proof. To see this, note that by Proposition 3.3.11, restricted to any compactum, $\mathrm{T}^{\mathrm{ex}}_{-}(X)$ is the uniform limit of the family of maps given by $f_r \colon x \mapsto \mathcal{M}^r_x(W)$. By exhausting W_p by compacta it suffices to see that the f_r are continuous for r small enough. Let $K \subset W_p$ be a compactum and let r be small enough, such that an r-neighborhood $\mathrm{N}_r(K) \cap W$ of K lies completely in $W_{\geq p}$. In other words, we may assume without loss of generality that W_p is the minimal stratum of W. Next, note that

$$\mathrm{d}(\mathcal{M}^r_x(W), \mathcal{M}^r_{x'}(W)) \leq r\mathrm{d}(\mathrm{B}_r(x) \cap W, \mathrm{B}_r(x') \cap W) + ||x - x'|| \tag{3.6}$$

for $x, x' \in W_p$. By an application of Thom's isotopy lemma, the map

$$\begin{aligned}\hat{g} : W \times W_p &\to (0,\infty) \times W_p \\ (x,y) &\mapsto (||x-y||, y)\end{aligned}$$

restricts to a fiber bundle over $(0,d]$ for d small enough. In particular, if we set $X = \hat{g}^{-1}(0,d] \cup \Delta_{W_p}$, we obtain an induced fiber bundle

$$\begin{aligned}g : X &\to W_p \\ (x,y) &\mapsto y\end{aligned}$$

with fiber $\mathrm{B}_r(y) \cap W$ over y. Again, locally using the independence of the Hausdorff-distance topology of the choice of metric, we obtain that $\mathrm{B}_r(y) \cap W$ varies continuously in y. Hence, by Eq. (3.6) so does $\mathcal{M}^r_y(W)$. □

Another consequence of Proposition 3.3.11 is that for a Lojasiewicz-Whitney stratified space W we have

$$\mathcal{M}_x^r(W) \xrightarrow{r\to 0} \mathrm{T}_x^{\mathrm{ex}}(W),$$

for all $x \in W$. This result is already present in a similar form in [Hir69]. However, our objective is to provide a description tailored to practical applications. Specifically, we aim to analyze the convergence behavior of magnifications for samples of X as $r \to 0$. Initially, this may seem like an illogical question. For a fixed sample $\mathbb{X}$, $\mathcal{M}_x^r(\mathbb{X})$ has a distance of 0 to a one-point (or empty) space when r is small enough. Instead, the correct notion of convergence is already suggested by the inequality in Proposition 3.3.11. What needs clarification is the convergence behavior where the quality of the sample is allowed to improve simultaneously as $r \to 0$. To be more precise, if $\mathbb{X}$ approaches X at a certain rate. We formalize this statement in the following corollary to Proposition 3.3.11:

Corollary 3.3.13. *Let $X \in$ **Sam** be a Lojasiewicz-Whitney stratifiable space. Let $x \in X$. Then,*

$$\mathcal{M}_x^r(\mathbb{X}) \to \mathrm{T}_x^{\mathrm{ex}}(X)$$

for $r \to 0$, $\mathbb{X} \to X$ and $\frac{1}{r}\mathrm{d}(\mathbb{X}, X) \to 0$. Furthermore, this convergence is uniform on any compactum K contained in a pure stratum.

3.4 Convergence of Tangentbundles

In order to establish a global result regarding the approximation of stratifications, it is imperative to derive a more global version of Corollary 3.3.13. This necessitates treating tangent cones not as isolated spaces but rather as a (stratified) bundle of cones. To articulate the ensuing convergence result, we require a space of samples corresponding to bundles.

Definition 3.4.1. Denote by **BSam** the set

$$\{(\mathbb{X}, F : \mathbb{X} \to \mathbf{Sam}_\star) \mid \mathbb{X} \in \mathbf{Sam}\},$$

equipped with the (extended pseudo) metric given by regarding F as a subset of $\mathbb{R}^N \times \mathbf{Sam}_\star$, equipping the latter with the product metric, and then using the resulting Hausdorff distance.
That is, for $(\mathbb{X}, F), (\mathbb{X}', F') \in \mathbf{BSam}$, we define

$$\mathrm{d}((\mathbb{X}, F), (\mathbb{X}', F')) := \max_{\mathbb{X}_0, \mathbb{X}_1 \in \{\mathbb{X}, \mathbb{X}'\}^2} \{\inf\{\varepsilon > 0 \mid \forall x \in \mathbb{X}_0 \exists y \in \mathbb{X}_1 : ||x - y||, \\ \mathrm{d}(F(x), F(y)) \leq \varepsilon\}.$$

We also refer to **BSam** as the *space of bundle samples* (of $\mathbb{R}^N$).

Definition 3.4.2. The *r-magnification bundle of* $\mathbb{X} \in \mathbf{Sam}$ is defined as the image of $\mathbb{X}$ under the map

$$\begin{aligned}\mathcal{M}^r : \mathbf{Sam} &\to \mathbf{BSam}\\ \mathbb{X} &\mapsto (\mathbb{X}, \{x \mapsto \mathcal{M}^r_x(\mathbb{X})\}).\end{aligned}$$

The *tangent cone bundle of* $X \in \mathbf{Sam}$ is defined as the image of X under the map

$$\begin{aligned}\mathrm{T}^{\mathrm{ex}} : \mathbf{Sam} &\to \mathbf{BSam}\\ X &\mapsto (X, \{x \mapsto \mathrm{T}^{\mathrm{ex}}_x(X)\}).\end{aligned}$$

We need to exercise caution regarding this nomenclature, as neither $\mathcal{M}^r$ nor $\mathrm{T}^{\mathrm{ex}}(\mathbb{X})$ closely resemble a fiber bundle for an arbitrary space $\mathbb{X}$.

Furthermore, it is important to recognize that the convergence asserted in Proposition 3.3.11 does not imply the convergence of magnification bundles in the metric on **BSam**. This convergence is only uniform on compact sets contained within pure strata. Nevertheless, we can endow the spaces **BSam** with alternative topologies, enabling the formulation of convergence notions on a compact set. Once again, for the remainder of this subsection, let $P = p < q$.

Construction 3.4.3. *Let* $K \in \mathbf{Sam}$ *and let* T *be any of the spaces* $\mathbf{Sam}_{\mathrm{N}(P)}$,**BSam**. *For* $\mathbb{X} \in \mathbf{Sam}$ *let* $\varepsilon : \mathbf{Sam} \to \mathbb{R}_+$ *be some continuous map. Define a map*

$$\begin{aligned}g^K_\varepsilon : T &\to T\\ (\mathbb{X}, f) &\mapsto (\mathbb{X} \cap K_{\varepsilon(\mathbb{X})}, f\mid_{K_{\varepsilon(\mathbb{X})}}).\end{aligned}$$

If $\mathcal{K} = (E, \varepsilon)$ *is a pair consisting of a set* $E \subset \mathbf{Sam}$*, together with a continuous map* $\varepsilon : \mathbf{Sam} \to \mathbb{R}_+$*, we denote by* $T^{\mathcal{K}}$*, the space with the same underlying set as* T*, but equipped with the initial topology with respect to the maps* g^K_ε *and* $T \to \mathbf{Sam} \xrightarrow{\varepsilon} \mathbb{R}_+$*. In particular, with respect to this topology, a sequence* $\mathbb{B}_n = (\mathbb{X}_n, F_n)$ *in* T *converges to* $\mathbb{B} = (\mathbb{X}, F) \in T$*, if and only if*

$$g^K_\varepsilon(\mathbb{B}_n) \xrightarrow{n\to\infty} g^K_\varepsilon(\mathbb{B}),$$

for all $K \in E$ *and*

$$\varepsilon(\mathbb{X}_n) \xrightarrow{n\to\infty} \varepsilon(\mathbb{X}).$$

We can now rephrase Proposition 3.3.11 as a global convergence result, which is essential for the our main result of this chapter, Theorem 3.5.8.

Proposition 3.4.4. *Let $X \in \mathbf{Sam}$ be equipped with a Lojasiewicz-Whitney stratification $W = (X, X \to P)$. Let $\mathcal{K} = (E, \varepsilon)$ with $\varepsilon := \mathrm{d}(X, -)$ and E a family of elements of* **Sam** *such that for all $K \in E$, there exist a decomposition into compacta $K = K_p \sqcup K_q$, such that $K_p \subset W_p$, $K_q \subset W_q$. Then,*

$$\mathcal{M}^r(\mathbb{X}) \to \mathrm{T}^{\mathrm{ex}}(X) \text{ in } \mathbf{BSam}^{\mathcal{K}}$$

as $r \to \infty$, $\mathbb{X} \to X$ and $\frac{1}{r}\mathrm{d}(\mathbb{X}, X) \to 0$.

Proof. Recall that for the convergence in $\mathbf{BSam}^{\mathcal{K}}$ we have to show

$$g_\varepsilon^K(\mathcal{M}^r(\mathbb{X})) \to g_\varepsilon^K(\mathrm{T}^{\mathrm{ex}}(X))$$

for any $K \in E$. Note that since $K \subset W$, $g_\varepsilon^K(\mathrm{T}^{\mathrm{ex}}(X)) = \mathrm{T}^{\mathrm{ex}}(X) \mid_K$ and we have $g_\varepsilon^K(\mathcal{M}^r(\mathbb{X})) = \mathcal{M}^r(\mathbb{X}) \mid_{\varepsilon(K)}$. By Proposition 3.3.11 we know that

$$\mathrm{d}(\mathrm{T}_x^{\mathrm{ex}}(W), \mathcal{M}_w^r(\mathbb{W})) \to 0$$

uniformly for all $x \in K$ and $||w - x|| \leq \varepsilon$ (i.e. $w \in \varepsilon(K)$) as $r \to \infty$, $\mathbb{X} \to X$ and $\frac{1}{r}\mathrm{d}(\mathbb{X}, X) \to 0$. □

3.5 The restratification theorem

We now have all the necessary tools to reconstruct stratifications from samples. As demonstrated in Theorem 2.2.32, the persistent stratified homotopy type is (Lipschitz) continuous in compact Whitney stratified spaces W over $P = p < q$. This implies that we can approximate the persistent stratified homotopy type of W from a stratified sample $\mathbb{W}$ that is close to W in the metric on $\mathbf{Sam}_P$. However, in practice, we typically receive non-stratified samples. In our exploration of magnifications and Φ-stratifications, we hinted at the use of local tangent cones to generate stratifications that approximate the original one. In this section, we will formalize and prove this assertion. But first, let us illustrate how the procedure works in the case where one is provided with a perfect sample, i.e., when working with the entirety of W. Once again, for the remainder of this section, let $P = p < q$.

Construction 3.5.1. *Let $W \in \mathbf{Sam}_P$ be a compact Lojasiewicz-Whitney Φ-stratified space, with respect to a function Φ as in Definition 3.2.1. Suppose we forget the stratification of $W = (X, s)$, and only have the data given by X. We can then associate to X its tangent cone bundle $\mathrm{T}^{\mathrm{ex}}X \in \mathbf{BSam}$. Next, we use the function Φ to decide which regions should be considered singular. We can do so, by applying Φ to $\mathrm{T}^{\mathrm{ex}}X$ fiberwise. As a result we obtain a strong stratification $\tilde{s}$ of X, given by*

$$x \mapsto \mathrm{T}_x^{\mathrm{ex}}(X) \mapsto \Phi(\mathrm{T}_x^{\mathrm{ex}}(X)).$$

By Proposition 3.3.12, this map is continuous on all strata. In particular, by assumption, it takes a maximum value $m < 1$ on W_p. Since W_q is a manifold, we have

$$\mathrm{T}^{\mathrm{ex}}_x(X) = \mathrm{T}^{\mathrm{ex}}_x(W_q) = \mathrm{T}_x(W_q) = \mathbb{R}^q$$

for $x \in W_q$, and thus the strong stratification has constant value 1 on W_q. Therefore, we may recover the stratification of s by choosing $u > m$ and applying $\mathcal{F}_u$:

$$\mathcal{F}_u(X, \tilde{s}) = W.$$

We now replicate the procedure described in Construction 3.5.1 in case of working with samples and investigate its convergence behavior.

Lemma 3.5.2. *Let $\Phi : \mathbf{Sam}_\star \to [0,1]$ be a continuous map. Then, the induced map*

$$\begin{aligned} \Phi_* : \mathbf{BSam} &\to \mathbf{Sam}_{\mathrm{N}(P)} \\ (\mathbb{X}, F) &\mapsto (\mathbb{X}, \Phi \circ F) \end{aligned}$$

is continuous. Even more, if Φ is C-Lipschitz, then so is Φ_*.

Proof. Since $\mathbf{Sam}_\star$ is isometric to the space of compact subspaces of $\mathrm{B}_1(0) \subset \mathbb{R}^N$ and thus compact, Φ is a uniformly continuous map. Hence, the result follows immediately by definition of the metrics on $\mathbf{BSam}$ and $\mathbf{Sam}_{\mathrm{N}(P)}$. □

Interestingly, Φ_* also extends to become a continuous map within the alternative topologies described in Construction 3.4.3.

Lemma 3.5.3. *Let $\Phi : \mathbf{Sam}_\star \to [0,1]$ be a continuous map. Let E be a family of subsets of $\mathbf{Sam}$, $\varepsilon : \mathbf{Sam} \to \mathbb{R}_+$ be some continuous function and $\mathcal{K} = (E, \varepsilon)$. Then, the map*

$$\begin{aligned} \Phi_* : \mathbf{BSam}^{\mathcal{K}} &\to \mathbf{Sam}^{\mathcal{K}}_{\mathrm{N}(P)} \\ (\mathbb{X}, F) &\mapsto (\mathbb{X}, \Phi \circ F) \end{aligned}$$

is continuous.

Proof. By definition of the topologies on $\mathbf{BSam}^{\mathcal{K}} \to \mathbf{Sam}^{\mathcal{K}}_{\mathrm{N}(P)}$ it suffices to show the claim on every element of E, i.e. we may without loss of generality assume $E = \{K\}$. Continuity of $\mathbf{BSam}^{\mathcal{K}} \to \mathbf{Sam}^{\mathcal{K}}_{\mathrm{N}(P)} \to \mathbf{Sam} \xrightarrow{\varepsilon} \mathbb{R}_+$ holds trivially. Furthermore, note that the diagram

$$\begin{array}{ccc} \mathbf{BSam}^{\mathcal{K}} & \xrightarrow{\Phi_*} & \mathbf{Sam}^{\mathcal{K}}_{\mathrm{N}(P)} \\ \downarrow{\scriptstyle g^K_\varepsilon} & & \downarrow{\scriptstyle g^K_\varepsilon} \\ \mathbf{BSam} & \xrightarrow{\Phi_*} & \mathbf{Sam}_{\mathrm{N}(P)} \end{array}$$

trivially commutes, since the g are given by restricting, i.e. precomposition and Φ_* by postcomposition. Then, for a sequence $\mathbb{B}_n \in \mathbf{BSam}^{\mathcal{K}}$ and $\mathbb{B} \in \mathbf{Sam}^{\mathcal{K}}_{\mathrm{N}(P)}$ we have:

$$\begin{aligned}
\mathbb{B}_n \xrightarrow{n\to\infty} \mathbb{B} \text{ in } \mathbf{BSam}^{\mathcal{K}} &\iff g^K_\varepsilon(\mathbb{B}_n) \xrightarrow{n\to\infty} g^K_\varepsilon(\mathbb{B}) \text{ in } \mathbf{BSam} \\
&\implies \Phi_*(g^K_\varepsilon(\mathbb{B}_n)) \xrightarrow{n\to\infty} \Phi_*(g^K_\varepsilon(\mathbb{B})) \text{ in } \mathbf{Sam}_{\mathrm{N}(P)} \\
&\iff g^K_\varepsilon(\Phi_*(\mathbb{B}_n)) \xrightarrow{n\to\infty} g^K_\varepsilon(\Phi_*(\mathbb{B}) \text{ in } \mathbf{Sam}_{\mathrm{N}(P)} \\
&\iff \Phi_*(\mathbb{B}_n) \xrightarrow{n\to\infty} \Phi_*(\mathbb{B}) \text{ in } \mathbf{Sam}^{\mathcal{K}}_{\mathrm{N}(P)},
\end{aligned}$$

where the implication in the second line follows by Lemma 3.5.2. □

As previously observed, concerning the alternative topologies, the magnification bundles uniformly converge to the tangent cone bundle. However, this is not the situation with the regular topologies. Therefore, in order to employ a magnification version of Construction 3.5.1 for stratification approximation, we must demonstrate the continuity of $\mathcal{F}_u$ within the corresponding tangent cone bundles with respect to the alternative topology.

Proposition 3.5.4. *Let $S = (X, s) \in \mathbf{Sam}_{\mathrm{N}(P)}$, X compact. Let $u \in [0,1]$ be such that $S_{\leq u}$ is closed and such that $S_{\leq u\pm\delta} = S_{\leq u}$ for δ sufficiently small. Let $\varepsilon = \mathrm{d}(X, -)$. Finally, let*

$$\mathcal{K} = (E = \{K \in \mathbf{Sam} \mid K = K_p \sqcup K_q, K_p, K_q \text{ compact}, K_p \subset S_u, K_q \subset s^{-1}(u,1]\}, \varepsilon).$$

Then,

$$\mathcal{F}_u : \mathbf{Sam}^{\mathcal{K}}_{\mathrm{N}(P)} \to \mathbf{Sam}_P$$

is continuous in S.

Proof. Let $\mathbb{S} = (\mathbb{X}, s') \in \mathbf{Sam}_{\mathrm{N}(P)}$. First, observe that $\varepsilon(\mathbb{X}) = \mathrm{d}(X, \mathbb{X})$ is continuous in $\mathbb{X}$ by assumption. Hence, it suffices to show convergence in the component $\mathbb{S}_{\leq u}$. We have

$$\mathrm{d}(S_{\leq u}, \mathbb{S}_{\leq u}) \leq \mathrm{d}(S_{\leq u}, \mathbb{S}_{\leq u} \cap K_{\varepsilon(\mathbb{X})}) + \mathrm{d}(S_{\leq u}, \mathbb{S}_{\leq u} \cap (S_{\leq u})_\gamma),$$

whenever $K = (X - (S_{\leq u})_{\frac{\gamma}{2}}) \sqcup S_{\leq u}$ and $\gamma > 0$ such that, $\mathbb{X} \subset K_{\varepsilon(\mathbb{X})} \cup (S_{\leq u})_\gamma$. Note, that for this to hold, it suffices that $\varepsilon(\mathbb{X}) \leq \frac{\gamma}{2}$. For the left summand we obtain,

$$\mathrm{d}(S_{\leq u}, \mathbb{S}_{\leq u} \cap K_{\varepsilon(\mathbb{X})}) = \mathrm{d}(\mathcal{F}_u(g^K_\varepsilon(S)), \mathcal{F}_u(g^K_\varepsilon(\mathbb{S}))) \leq \mathrm{d}(g^K_\varepsilon(S), g^K_\varepsilon(\mathbb{S})),$$

by Corollary 2.2.27, for $\varepsilon(\mathbb{X})$ sufficiently small and $g^K_\varepsilon(\mathbb{S})$ close to $g^K_\varepsilon(S)$. For the other summand we first split the Hausdorff distance into the directed distances

$$\mathrm{d}(S_{\leq u}, \mathbb{S}_{\leq u} \cap (S_{\leq u})_\gamma) \leq \vec{d}(S_{\leq u}, \mathbb{S}_{\leq u} \cap (S_{\leq u})_\gamma) + \vec{d}(\mathbb{S}_{\leq u} \cap (S_{\leq u})_\gamma, S_{\leq u})$$

where $\vec{d}(A,B) = \inf\{\delta \geq 0 \mid A \subset B_\delta\}$. Then, the second summand is bounded by γ and for the first summand we observe that

$$\vec{d}(S_{\leq u}, \mathbb{S}_{\leq u} \cap (S_{\leq u})_\gamma) \leq \vec{d}(S_{\leq u}, \mathbb{S}_{\leq u} \cap (S_{\leq u})_{\varepsilon(\mathbb{X})}).$$

This is due to the fact that $\varepsilon(\mathbb{X}) < \gamma$ and $\mathbb{S}_{\leq u} \cap (S_{\leq u})_{\varepsilon(\mathbb{X})} \subset \mathbb{S}_{\leq u} \cap (S_{\leq u})_\gamma$. If we set $K' = S_{\leq u}$ and invoke Corollary 2.2.27 again we obtain

$$\vec{d}(S_{\leq u}, \mathbb{S}_{\leq u} \cap (S_{\leq u})_{\varepsilon(\mathbb{X})}) = \vec{d}(\mathcal{F}_u(g_\varepsilon^{K'}(S)), \mathcal{F}_u(g_\varepsilon^{K'}(\mathbb{S}))) \leq \vec{d}(g_\varepsilon^{K'}, (S) g_\varepsilon^{K'}(\mathbb{S}))$$

for $g_\varepsilon^{K'}(\mathbb{S})$ close to $g_\varepsilon^{K'}(S)$. Summarizing, we have:

$$d(S_{\leq u}, \mathbb{S}_{\leq u}) \leq d(g_\varepsilon^K(S), g_\varepsilon^K(\mathbb{S})) + \vec{d}(g_\varepsilon^{K'}(S), g_\varepsilon^{K'}(\mathbb{S})) + \gamma.$$

In particular, we may first fix some γ while the other terms converge to 0 for $\mathbb{S} \to S$ in $\mathbf{Sam}^{\mathcal{K}}_{\mathrm{N}(P)}$ by assumption. Since γ can be taken arbitrarily small, the result follows. □

We arrived at the definition of a map which equips samples with stratifications, depending on their approximate tangential structure.

Definition 3.5.5. Let $\Phi : \mathbf{Sam}_\star \to [0,1]$ be a continuous map and $u \in [0,1)$, $r \in \mathbb{R}_+$. Let $\mathbb{X} \in \mathbf{Sam}_\star$. We call the image of $\mathbb{X}$ under the composition

$$\mathcal{S}^r_{\Phi,u} : \mathbf{Sam} \xrightarrow{\mathcal{M}^r} \mathbf{BSam} \xrightarrow{\Phi_*} \mathbf{Sam}_{\mathrm{N}(P)} \xrightarrow{\mathcal{F}_u} \mathbf{Sam}_P$$

the Φ-*stratification of* $\mathbb{X}$ *at* r *(with respect to* u*).* In the case where $r = 0$, replace $\mathcal{M}^r$ by T^{ex}.

Example 3.5.6. To illustrate the concepts in Definition 3.5.5 let us walk through every component of the composition defining $\mathcal{S}^r_{\Phi,u}$ for a specific sample. Let X denote the algebraic variety given by

$$\{(x,y,z) \in \mathbb{R}^3 \mid (x^2+y^2+z^2+1.44)^2 - 7.84x^2 + 1.44y^2 = 0\}. \tag{3.7}$$

In the bottom left of Fig. 3.5, a visual representation of X can be found. A finite sample from this variety, denoted $\mathbb{X}$ was obtained by randomly picking points from an enclosing rectangular cuboid and only keeping points that satisfy 3.7 up to a small error. Choosing a magnification parameter $r = 5$ we obtain the magnification bundle $\mathcal{M}^r(\mathbb{X})$ for $\mathbb{X}$, depicted in the top middle of Fig. 3.5. Φ was chosen as in in Example 3.2.3. Evaluating the fibers of $\mathcal{M}^r(\mathbb{X})$ we obtain a strongly stratified sample $\Phi_*(\mathcal{M}^r(\mathbb{X}))$, shown on the left of Fig. 3.5. Next, picking the threshold value $u \in [0,1)$ to be 0.83 induces a stratified sample via $\mathcal{F}_u$. A visual comparison indicates that the resulting stratified sample is close to the Whitney stratified space given by X with

two isolated singularity. This already points at the convergence behavior predicted by Theorem 3.5.8.

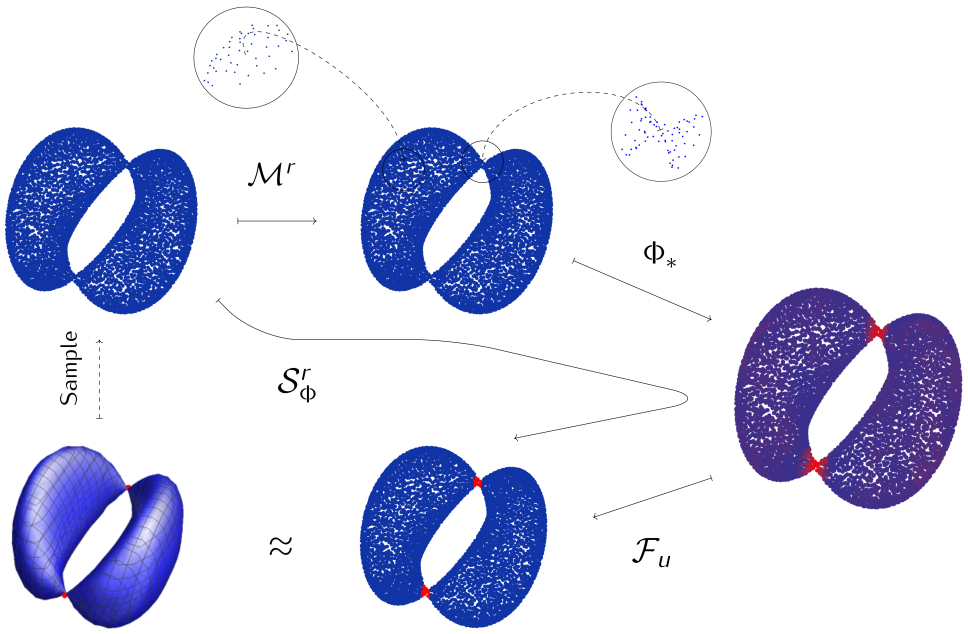

Figure 3.5. Illustration of $\mathcal{S}^r_{\Phi,u}$ for a sample from a 2-dimensional algebraic variety

We can then restate the content of Construction 3.5.1 as follows.

Proposition 3.5.7. *Let $W \in \mathbf{Sam}_P$ be a Lojasiewicz-Whitney stratified space, Φ-stratified with respect to $\Phi : \mathbf{Sam}_\star \to [0,1]$ as in Definition 3.2.1. Then,*

$$\sup\{\Phi(\mathrm{T}^{\mathrm{ex}}_x(X)) \mid x \in W_p\} < 1.$$

In particular,

$$\mathcal{S}^r_{\Phi,u}(X) = W,$$

for $\sup\{\Phi(\mathrm{T}^{\mathrm{ex}}_x(X)) \mid x \in W_p\} < u < 1$.

Proof. This was already covered in Construction 3.5.1. □

Now, we can formally present the primary theorem regarding the retrieval of stratification from samples of a Lojasiewicz-Whitney Φ-stratified space W. Essentially, in practical terms, it ensures that for sufficiently small r and with a suitably accurate sample, one can utilize the Φ-stratification at r to approximate the stratified space W. Notably, this result is applicable to all compact, subanalytically Whitney stratified spaces with two strata.

Theorem 3.5.8. *Let $P = \{p < q\}$ and let $W = (X, X \to P) \in \mathbf{Sam}_P$ be a compact Lojasiewicz-Whitney stratified space, Φ-stratified with respect to $\Phi : \mathbf{Sam}_\star \to [0,1]$. Then there exists $u_0 \in (0,1)$ such that for all $u \in [u_0, 1)$ we have convergence*

$$\mathcal{S}^r_{\Phi,u}(\mathbb{X}) \to W,$$

for $\frac{1}{r}d(\mathbb{X}, X) \to 0$ as $r \to 0$.

Proof. Let $\mathcal{K}$ be as in Proposition 3.5.4. It is the content of the latter, that

$$\mathcal{M}^r(\mathbb{X}) \to \mathrm{T}^{\mathrm{ex}}(X) \text{ in } \mathbf{BSam}^{\mathcal{K}}$$

for $r \to 0$ and $\frac{1}{r}d(\mathbb{X}, X) \to 0$. Applying, Φ_* to this and using Lemma 3.5.3, we obtain

$$\Phi_* \circ \mathcal{M}^r(\mathbb{X}) \to \Phi_* \circ \mathrm{T}^{\mathrm{ex}}(W) \text{ in } \mathbf{Sam}^{\mathcal{K}}_{\mathrm{N}(P)}.$$

Now, note that $\Phi_* \circ \mathrm{T}^{\mathrm{ex}}(X)$ fulfills the requirements of Proposition 3.5.4, if we take $1 > u > \max\{\Phi_* \circ \mathrm{T}^{\mathrm{ex}}(X)(x) \mid x \in X\}$. Hence,

$$\mathcal{S}^r_{\Phi,u}(\mathbb{X}) = \mathcal{F}_u \circ \Phi_* \circ \mathcal{M}^r(\mathbb{X})\mathcal{F}_u \circ \Phi_* \circ \mathrm{T}^{\mathrm{ex}}(X) = W,$$

for $r \to 0$ and $\frac{1}{r}d(\mathbb{X}, X) \to 0$, where the equality follows by Proposition 3.5.7. □

Using the notation as above, the Theorem makes a statement about the convergence of spaces in $\mathbf{Sam}_P$. For the given space W it states that for any $\delta > 0$ there exists a radius R such that for all radii $r > R$ there exists $\varepsilon_r > 0$ such that

$$\max\{d_{Hd}(\mathcal{S}^r_{\Phi,u}(\mathbb{X}) = \mathbb{X}, X), d_{Hd}(\mathcal{S}^r_{\Phi,u}(\mathbb{X})_p, W_p)\} < \delta$$

for all $\mathbb{X}$ with $d_{Hd}(\mathbb{X}, W_q) < \varepsilon_r$. In other words, as $\mathbb{X}$ converges to X by assumption, the statement is really about the convergence of the singular parts $\mathcal{S}^r_{\Phi,u}(\mathbb{X})_p$ and W_p here. In this sense, Theorem 3.5.8 can be used to approximate or to find singularities of a stratified space form a good enough sample and one may choose to use this method on its own for this purpose. However, we can now combine this result with Theorem 2.2.32 which guarantees us that $\mathcal{S}^r_{\Phi,u}$ can also be used to infer stratified homotopy types from non-stratified samples. Note that in the following we will not assume W to be rescaled in such a way that it is cylindrically stratified, hence one has to choose the parameter v sufficiently small to compensate for the rescaling.

Corollary 3.5.9. *Let $P = \{p < q\}$ and let $W = (X, X \to P = \{p < q\}) \in \mathbf{Sam}_P$ be a compact Lojasiewicz-Whitney stratified space, Φ-stratified with respect to $\Phi : \mathbf{Sam}_\star \to [0,1]$. Then*

$$\mathcal{P}_v \circ \mathcal{S}^r_{\Phi,u}(\mathbb{X}) \to \mathcal{P}_v(W),$$

for $r \to 0$ and $\frac{1}{r}d(\mathbb{X}, X) \to 0$, $u < 1$ close to 1 and $v \in \Omega$ sufficiently small.

This concludes our theoretical investigations on how to approximate the stratification of a compact Lojasiewicz-Whitney space W from a given sample space that lies sufficiently close the to W. At this point we owe the reader some examples that demonstrate the utility of our results and maybe also help understanding the meaning of certain statements. However, since we have a full chapter dedicated to the algorithmic implementation of the methods involved in computing Φ-functions and $\mathcal{S}^r_{\Phi,u}$ with Chapter 4 we decided to post-pone an exhaustive demonstration of our methods to Section 4.3.2 where we also discuss a longer example in Example 4.3.2 with several figures.

Chapter 4

Algorithmic Stratification Learning

This chapter is committed to detail the algorithmic aspect of this work. In essence, we want to take up upon the mathematical concepts we developed in the previous chapters and describe how to stratify a point cloud sampled from a stratified space algorithmically. In particular, we implement the computation of different functions $\Phi : \mathbf{Sam}_\star \to [0,1]$ based on local homology and tangential approximations as in Examples 3.2.2 and 3.2.3.

We provide an introduction to persistent homology in Section 4.1 with a focus on Vietoris-Rips complexes as we feel this to be necessary to explain the finer details of the implementation of our methods. The main result of Section 4.1 is the detailed description of an algorithm that approximates local persistent homology of a topological space X, that can be used to detect singularities in a given sample. We also provide theoretical arguments that our algorithm approximates $(X_\varepsilon, X_\varepsilon - B_r(x))$, from a good enough sample of X up to a certain error. These approximations results are based on the work of Skraba and Wang [SW14].

For the computation of Hausdorff distances of subspaces of $\mathbb{R}^n$, explained in Section 4.2, we will need very different mathematical concepts as to before and therefore included some results on Grassmannians. As a main result we derive an algorithm to approximate a solution to the problem of finding the k-dimensional linear subspace of $\mathbb{R}^n$ with the minimal Hausdorff distance to a given subspace of $\mathbb{R}^n$.

In Section 4.3, we put together our theoretical and our algorithmic results to a fully implementable stratification pipeline and apply it to a controlled example, that is to say having full control over the Hausdorff distance of the sample to the targeted space, in order to demonstrate the meaning and utility of what we described in the previous chapter and in particular in Theorem 3.5.8 and Corollary 3.5.9.

4.1 Persistent Homology

We start with a broad recap on the most fundamental concepts involved in persistent homology and recall some well-known results that we will use in this chapter. For a thorough introduction into this topic consider e.g. [EH10, CdSGO16]. We already discussed persistent objects in more general context in Section 2.2.2. However, as we want to work towards our algorithmic results we chose to include a more specific recap on the case of *persistence modules*.

Previously, we referred to general persistent objects as elements of a functor category $\mathcal{D}^U$ with U some poset category. If we take $U = (\mathbb{R}, \leq)$ with the usual order and $\mathcal{D} = \mathbf{Vec}_k$ with k some field, then we obtain the category of (1-dimensional) persistence modules. Thus, it contains the data of a family of vector spaces $V = (V_s)_{s\in\mathbb{R}}$ together with maps $v_t^s : V_s \to V_t$ such that $v_t^s \circ v_u^t = v_u^s$ for $s \leq t \leq u$ and $v_s^s = \mathrm{id}_{V_s}$. Accordingly, a *morphism* of persistence modules $\phi : V \to W$, both indexed over $\mathbb{R}$, is given by a natural transformations, i.e. by a collection of morphisms $\phi^s : V_s \to W_s$ such that the following diagrams commute

$$\begin{array}{ccc} V_s & \xrightarrow{v_t^s} & V_t \\ \downarrow{\scriptstyle \phi^s} & & \downarrow{\scriptstyle \phi_t} \\ W_s & \xrightarrow{w_t^s} & W_t \end{array}$$

for all $s \leq t$.

Definition 4.1.1. Let k be a field and let I denote an interval in $\mathbb{R}$. The module defined by

$$Q(I_s) = \begin{cases} k & \text{if } s \in I \\ 0 & \text{else,} \end{cases}$$

together with $v_s^t = \mathrm{id}$ for $s \leq t$ is called the *interval module* associated to I.

An interval in $\mathbb{R}$ could mean different things, i.e. $I = (x, y)$ or $I = [x, y)$ etc. We will stick to the bounded below intervals, i.e. $[x, y)$, from hereon out. Note that one could equivalently work with e.g. the open intervals (x, y) but it is important in this context that we do not allow an interval I in $\mathbb{R}$ to be either or. Furthermore, we call a persistence module V indexed over $\mathbb{R}$ *of finite type* if it contains only a finite number of unique finite-dimensional vector spaces.

The following statement is a consequence of [CB14, Theorem 1.1]:

Theorem 4.1.2. *Let V be a persistence module of finite type indexed over $\mathbb{R}$. Then,*

$$V \simeq \bigoplus_{I\in B(V)} Q(I)$$

where $B(V)$ denotes a finite multiset of intervals in $\mathbb{R}$.

We will refer to $B(V)$ as the *barcode of* V. Although the decomposition in Theorem 4.1.2 is not unique, the barcode $B(V)$ associated with V is, at least up to a reordering of the intervals. A barcode determines the isomorphism type of a persistence diagram up to isomorphism. Therefore, it makes sense to speak of the barcode of a persistence module under the conditions imposed in the theorem. Note that in case V is of finite type the barcode is only a finite multiset of intervals. We will mostly work in this scenario, however the above Theorem can be substantially generalized (compare to [CB14, BCB20]).

Let us take a look at some common cases where persistence modules and barcodes arise.

Example 4.1.3. Let X be space together with a function $f : X \to \mathbb{R}$. Then, one naturally obtains a filtration of X given by

$$\cdots \hookrightarrow X_\varepsilon = f^{-1}(-\infty, \varepsilon] \overset{i}{\hookrightarrow} X_\delta = f^{-1}(-\infty, \delta] \hookrightarrow \cdots$$

for $\varepsilon \leq \delta$. For such a filtration $(X_\varepsilon, f)_{\varepsilon \in \mathbb{R}}$ one obtains a persistence module indexed over $\mathbb{R}$ by applying a functor from **Top** to $\mathbf{Vec}_k$. More specifically, let $H(-) = H_d(-) = H_d(-; k)$ denote d-dimensional singular homology with coefficients in a field k. One obtains a persistence module

$$\cdots \to H(X_\varepsilon) \overset{i_\varepsilon^\delta}{\to} H(X_\delta) \to \cdots$$

where the maps are induced by inclusions. We denote such a persistence module by $H(X_\mathbb{R})$ or $H(f_\mathbb{R})$ in case we would like to stress the dependence on a certain filtration function f. In many cases the function will be given by the distance to X in Euclidean space. Furthermore, the function f is called *tame* if it has a finite number of homological critical values, i.e. points $\varepsilon \in \mathbb{R}$ where $H(X_{\varepsilon-\alpha}) \to H(X_{\varepsilon-\alpha})$ is not an isomorphism for all sufficiently small $\alpha > 0$, and all $H(X_\varepsilon)$ are finite-dimensional in all homological dimensions dimension which is to say that $H(f)_\mathbb{R}$ is of finite type. In this case we can give a very explicit description of the barcode associated to $H(f_\mathbb{R})$. We say a class $a \in H(X_\varepsilon)$ was *born* at ε if $a \notin im(i_{\varepsilon-\alpha}^\varepsilon)$ for any $\alpha > 0$. Such a class born at ε is said to *die* at δ if $i_\varepsilon^{\delta-\alpha}(a) \notin im(i_{\varepsilon-\alpha}^{\delta-\alpha})$ for all $\alpha > 0$ but $i_\varepsilon^\delta(a) \in im(i_{\varepsilon-\alpha}^\delta)$. The barcode of the persistent homology module, also referred to as *persistent homology*, is then given by the collections of the intervals ranging from the time of birth to the time of death of all persistent homology classes.

Example 4.1.4. In application of topological methods it is often the case that X is a finite simplicial complex and all sublevel sets X_ε given by a function on X are

subcomplexes. This has the consequence that all vector spaces $H(X_\varepsilon)$ are finite dimensional and also as ε increases there are only finitely many value at which the complexes change and thus there are only finitely many points at which the homology can change, i.e. at the critical values of the function on X. Assume the critical values are

$$\varepsilon_1 < \varepsilon_2 < \cdots < \varepsilon_n.$$

Then, the complete information of the persistence module is contained in the finite diagram

$$H(X_{\varepsilon_1}) \to H(X_{\varepsilon_2}) \to \cdots \to H(X_{\varepsilon_n})$$

where every vector space is finite-dimensional. Note that in this case we could also consider the persistence module induced by the sublevel sets of a function on X as a persistence module indexed over the poset $[n] = \{1, \ldots, n\}$. We will discuss this more precisely in the next two following section.

We have already introduced metrics on generalized persistence module in Section 2.2.2 and called them interleaving distances there. More specifically, for the case of persistence modules indexed of $\mathbb{R}_+$ we can define a flow (compare to Recollection 2.2.18) by functors $-_\varepsilon : (\mathbb{R}_+, \leq) \to (\mathbb{R}_+, \leq)$ for some $\varepsilon \in \mathbb{R}_+$, together with natural transformations $id_{(\mathbb{R}_+,\leq)} \to -_\varepsilon$ given by $s \mapsto s + \varepsilon$. We say two persistence modules V, W are ε-interleaved if there are morphisms $\phi : V \to W_\varepsilon$ and $\psi : W \to V_\varepsilon$ such that the diagram

$$\begin{array}{ccc} V & & W \\ {\scriptstyle\phi}\downarrow & \swarrow\searrow & \downarrow{\scriptstyle\psi} \\ W_\varepsilon & & V_\varepsilon \\ {\scriptstyle\psi_\varepsilon}\downarrow & \swarrow\searrow & \downarrow{\scriptstyle\phi_\varepsilon} \\ V_{2\varepsilon} & & W_{2\varepsilon} \end{array} \tag{4.1}$$

commutes. Informally speaking, the definition of an ε-interleaving can be understood as an ε-approximate isomorphism of persistence modules. Then, again, the interleaving distance of two persistence modules is defined as the infimum of all values for which an interleaving exists. From hereon out by interleaving distance we will be referring to the interleaving distance between persistence modules. We will not work with interleaving distances explicitly and therefore we will not go into more detail on how interleaving distances can be computed. An important property of the interleaving distance is that it agrees with the bottleneck distance of the associated barcodes which we will explain next. In practice interleaving distances are not useful because computing it is an NP-hard problem, cf. [BB18]. The structure of a barcodes

makes them much more amenable to measuring a distance between persistent objects algorithmically. First, we need to introduce the notion of matchings.

Notation 4.1.5. A *matching* of two multiset of intervals D_1 and D_2 is a collection of pairs $\mathcal{X} = \{(I, J) \in D_1 \times D_2\}$ with intervals I, J being only allowed to occur in one pair. We say I is *matched* to J if $(I, J) \in \mathcal{X}$ and we say I is *unmatched* if I does not belong to a pair in $\mathcal{X}$. For D_1 and D_2 multisets of intervals in $\mathbb{R}$ we define the cost of a matching $\mathcal{X}$ as

$$c(\mathcal{X}) = \max\{\sup_{(I,J)\in\mathcal{X}} c(I, J), \sup_{\text{unmatched } I\in D_1\cup D_2} c(I)\}$$

with

$$c(I, J) = \max\{|c - a|, |d - b|\}$$

and

$$c(I) = \frac{|b - a|}{2}$$

for $I = [a, b)$ and $J = [c, d)$.

Definition 4.1.6. For D_1 and D_2 multisets of intervals in $\mathbb{R}$ the *bottleneck distance* is given by

$$d_B(D_1, D_2) = \inf\{c(\mathcal{X}) \mid \mathcal{X} \text{ a matching of } D_1 \text{ and } D_2\}.$$

The Bottleneck distance is as informative as the interleaving distance, i.e. one can show

$$d_B(B(V), B(W)) = d(V, W) \tag{4.2}$$

for V, W persistence modules indexed over $\mathbb{R}$ such that V_ε and W_ε are finite dimensional for all $\varepsilon \in \mathbb{R}$ (compare to [CCSG$^+$09, Les15]).

Remark 4.1.7. The reader may have noticed that the Bottleneck is only perceptive to the maximum over the matchings and all other matchings and costs are discarded. Other metrics that are sensitive to a bigger part of the matchings include the so-called *degree p Wasserstein distance* defined by

$$W_p(D_1, D_2) = \inf\{c_p(\mathcal{X}) \mid \mathcal{X} \text{ a matching of } D_1, D_2\}.$$

with

$$c_p(\mathcal{X}) = (\sum_{(I,J)\in\mathcal{X}} c(I, J)^p + \sum_{\text{unmatched } I\in D_1\cup D_2} c(I)^p)^{\frac{1}{p}}$$

which is a metric commonly used also in other applications. However, the Wasserstein distance does not possess the same stability properties for persistence diagrams as the

Bottleneck distance. In other words, the two distances are generally not equal and therefore the Wasserstein distance does not equal the interleaving distance in the sense the Bottleneck distance does.

Remark 4.1.8. If we recall the definition of $\Phi^n_{\mathcal{PL}}$ in Example 3.2.3 we observe that the computation involves the computation of the interleaving distance of persistent homology modules which agrees with Bottleneck distance of the associated barcodes. We will not discuss the implementation of an algorithm to compute the Bottleneck distance and the matchings associated with it in general as this has been done in the literature to some extend [EH10, KMN17] and would lead to far afield at this point. Probably the most common approach is to approximate these matching distances which is also common for the computation of Wasserstein distances [Cut13, AWR17, CRL+20]. Apart from that, note that the Bottleneck distances in $\Phi^n_{\mathcal{PL}}$ only ever appeared with respect to $d_B(-, B(H(S^n_{[0,\alpha)})))$ where $\alpha > 0$ and S^n denotes the standard Euclidean sphere and thickened w.r.t. the standard metric. More specifically, we look at the barcode associated with the persistence module $H(S^n_{[0,\alpha)})$ induced by

$$\begin{aligned} d_{S^n} : \mathbb{R}^n &\to \mathbb{R} \\ x &\mapsto d_{\mathbb{R}^n}(x, S^n) \end{aligned}$$

is given by

$$B(H(S^n_{[0,\alpha)})) = \{[0, \frac{\alpha}{2})\}$$

for homology in dimensions 0 and n. For reduced homology there will only be one distinct interval. The Bottleneck distance to any other finite multiset of intervals $B(V)$ associated to a persistence module V of finite type as in Theorem 4.1.2 such that $c(I) \leq \frac{\alpha}{2}$ for all $I \in B(V)$ is then given by

$$d_B(B(V), B(H(S^n_{[0,\alpha)}))) = \max\{c([a,b), [0, \frac{\alpha}{2})), \max_{B(V)\ [a,b)} c(I)\} \tag{4.3}$$

where $[a,b) = \arg\min_{I \in B(V)} c(I, [0,\alpha))$ which exists because $B(V)$ is finite. Note that $[a,b) = \arg\min_{I \in B(V)} c(I, [0,\alpha))$ implies that $c(I) \leq c([a,b))$ for all $I \in B(V)$. The upshot of this remark is the observation that the right-hand side of Eq. (4.3) gives an easily implementable way to compute all Bottleneck distances involved in the computation of $\Phi^n_{\mathcal{PL}}$.

4.1.1 Filtered Simplicial Complexes

In the upcoming section, we provide a brief recap of the elements of filtered simplicial complexes that are essential for our investigations. We begin by setting the notation we will be utilizing in most of the parts that follow.

Definition 4.1.9. Let V be a finite set. A finite *(abstract) simplicial complex* K is a collection of finite non-empty sets of V, such that for every set $\sigma \in K$, and for every non-empty subset $\tau \subseteq \sigma$, τ also belongs to K.

The elements of a simplicial complex are called *simplices*. We often refer to a simplex $\{v_0, \ldots, v_n\}$ as an n-simplex. In other words, the *dimension* of a simplex is determined by its cardinality minus 1. An element of a simplex is called *vertex*. The *dimension* of a simplicial complex K is defined by $\dim(K) := \max_{\sigma \in K} \dim(\sigma)$. For $\sigma, \tau \in K$ we say τ is a *face* of α if $\tau \subseteq \sigma$. If τ is a proper subset of σ it is a *proper face*. Also, σ is then called a *(proper) coface* of β. Furthermore, there is also a notion of the *boundary of a simplex* $\{v_0, \ldots, v_n\}$ which is defined by the set $\partial\{v_0, \ldots, v_n\} := \{\{v_0, \ldots, \hat{v_i}, \ldots, v_n\} \mid i \in \{0, \ldots, n\}\}$, where $\hat{v_i}$ signifies omitting the vertex. A subcollection of simplices of K that satisfies the conditions of a simplicial complex itself is called a subcomplex of K.

Definition 4.1.10. A *filtration* of a finite simplicial complex K is a collection of subcomplexes $(K_i)_{i \in I}$ with I a totally ordered set and $i < j \Rightarrow K_i \subseteq K_j$.

Example 4.1.11.

- Let $f : K \to \mathbb{R}$ be a function on a simplicial complex K. For any $\varepsilon \in \mathbb{R}$, the sublevel set $f^{-1}(-\infty, \varepsilon]$ of a monotonic function f, i.e. if $\sigma \subset \tau \in K$ implies $f(\sigma) \leq f(\tau)$, is a subcomplex. The sublevel sets constitute a filtration of K indexed over $\mathbb{R}$. The following example is a special cases of this.

- Let $(\mathbb{X}, d)$ be a point cloud, i.e. a finite metric space. The *Vietoris-Rips complex* of $\mathbb{X}$ at scale $\varepsilon \in \mathbb{R}$ is the abstract simplicial complex given by

$$Rips_\varepsilon(\mathbb{X}) = \{\emptyset \neq \sigma \subset \mathbb{X} \mid \operatorname{diam}(\sigma) \leq \varepsilon\},$$

  with $\operatorname{diam}(\sigma) = \max_{\{u,v\} \subset \sigma} d(u, v)$. Then, $Rips_{\mathbb{R}}(\mathbb{X})$ is a filtration of the complex given by the full simplex $\mathbb{X}$. We will often refer to this as the *Vietoris-Rips filtration* of $\mathbb{X}$. In other words, the Vietoris-Rips filtration arises as the sublevel sets of the diameter function on simplices.

- The *Čech complex* of $\mathbb{X}$ at scale $\varepsilon \in \mathbb{R}$ is given by

$$\check{C}_\varepsilon(\mathbb{X}) = \{\emptyset \neq \sigma \subset \mathbb{X} \mid \bigcap_{v \in \sigma} B_\varepsilon(v) \neq \emptyset\}.$$

  The *Čech filtration of* $\mathbb{X}$, denoted $\check{C}_{\mathbb{R}}(\mathbb{X})$, is another filtration of the complex given by the full simplex $\mathbb{X}$.

Vietoris-Rips and Čech filtration are two of the most commonly used filtered simplicial complexes in the field of TDA. While the Vietoris-Rips filtration is considered

a more computationally efficient method, the Čech filtration guarantees a certain theoretical correctness. The specific purpose of a filtered simplicial complex in TDA is to resemble the topological structure inherent in a given point cloud $\mathbb{X}$. To motivate the Čech complex as an abstract simplicial complex for the use in TDA (although this may not have been the original purpose for it), assume $\mathbb{X}$ to be a set of randomly sampled points from an embedded manifold $M \subset \mathbb{R}^n$. It can be shown that for certain $\varepsilon > 0$ the union of all Euclidean ε-balls centered at x for all $x \in \mathbb{X}$, $\bigcup_{x\in\mathbb{X}} B_\varepsilon(x)$, has the same homotopy type as M if $\mathbb{X}$ is sufficiently close (in Hausdorff distance) to M. Even more is true, M is a deformation retract of $\bigcup_{x\in\mathbb{X}} B_\varepsilon(x)$. A proof of such a result can be found in [NSW08]. Therefore, the remaining question is whether or not the Čech complex resembles the topology of the union of ε-balls. To formulate the answer we need the concept of nerves.

Definition 4.1.12. Let F be a finite collection of sets. Define the nerve of F to consist of all non-empty subcollections whose sets have a non-empty common intersection,

$$\text{Nrv } F = \{S \subset F \mid S \neq \emptyset\}.$$

The nerve of a finite collection of sets F defines an abstract simplicial complex.

Immediately from the definition we see that the Čech complex of $\mathbb{X}$ at scale ε is the nerve of the collection of all closed ε-balls.

Theorem 4.1.13. *Let F be a finite collection of closed, convex sets in Euclidean space. Then the nerve of F and the union of the sets in F have the same homotopy type.*

Proof. For a proof consider e.g. [Hat02, Corollary 4G.3, p. 459]. □

Although we referenced [Hat02] as a modern source for the proof, the so-called Nerve theorem is commonly credited to [Bor48].

With this theorem the Čech complex of $\mathbb{X}$ at scale ε is homotopy equivalent to the union of the balls $\bigcup_{x\in\mathbb{X}} B_\varepsilon(x)$. Unfortunately, the Vietoris-Rips complex is in general not homotopy equivalent to the Čech complex at the same scales. However, the Vietoris-Rips complex is far more efficient to implement because it is a so-called flag (or clique) complex. That is, any $k+1$ vertices span a k-simplex if and only if any two of them span a 1-simplex. Consequently, the Vietoris-Rips complex is determined by its 1-skeleton. The verification of the condition imposed on simplices to be contained in the Vietoris-Rips also has a favorable complexity when compared to the Čech complex. It is for these reasons that the Vietoris-Rips complex lends itself to applications but it also seems appropriate to briefly evaluate the error of alignment

with the Čech complex:

Remark 4.1.14. For a simplex to be contained in $\check{C}_\varepsilon(\mathbb{X})$ we have to check if the metric ε-balls around every vertex have a common point of intersection. For a simplex to satisfy the Vietoris-Rips conditions we only require the same metric balls all to have pairwise intersection. As a consequence we find that $\check{C}_\varepsilon(\mathbb{X}) \subset Rips_\varepsilon(\mathbb{X})$. The inverse inclusion does not hold in general. But if a finite number of points of $\mathbb{X}$ are within distance $\frac{\varepsilon}{2}$ to each other, then all of them are contained in a metric closed ball of radius ε. Thus we have the containment $Rips_{\frac{\varepsilon}{2}}(X) \subset \check{C}_\varepsilon(\mathbb{X})$. In total we have

$$\check{C}_{\frac{\varepsilon}{2}}(\mathbb{X}) \subseteq Rips_{\frac{\varepsilon}{2}}(\mathbb{X}) \subseteq \check{C}_\varepsilon(\mathbb{X}).$$

With this observation we find that if the Čech complexes at scale $\frac{\varepsilon}{2}$ and at scale ε are good approximations of the underlying space then so is the Vietoris-Rips complex at scale $\frac{\varepsilon}{2}$. Also, if a topological feature persists in Čech complexes from scale ε to ε' such that $\frac{\varepsilon'}{\varepsilon} > 2$ then this feature is also recorded using the Vietoris-Rips complex at the respective scale. In other words, this brief argument shows that the Čech and the Vietoris-Rips filtration are ε-interleaved and this interleaving distance value can be further optimized. The result [SG07, Theorem 2.5] states that the above inclusions still hold true for $\frac{\varepsilon'}{\varepsilon} \geq \sqrt{\frac{2d}{d+1}}$ where d is the dimension of the ambient space of $\mathbb{X} \subset \mathbb{R}^d$. The interleaving distance is therefore reduced to $(\sqrt{\frac{2d}{d+1}} - 1)\varepsilon$.

We close this recap with some results on the computability of persistent homology of point clouds by using filtered simplicial complexes. What is needed for persistent homology computations from a filtered simplicial complex (possibly the Vietoris-Rips filtration) is an essential and simplex-wise refinement of the filtration.

Notation 4.1.15. A filtration of a filtered simplicial complex $(K_i)_{i\in I}$ is called *essential* if $i \neq j$ implies $K_i \neq K_j$ for $i, j \in I$. A simplex-wise filtration of K is a filtration such that for all $i \in I$ with $K_i \neq \emptyset$ there is some index $j < i \in I$ and some simplex $\sigma_i \in K$ such that $K_i \setminus K_j = \{\sigma_i\}$. A *reindexing* of $(K_i)_{i\in I}$ is given by a monotonic map $\rho : I \to J$ of totally ordered sets I, J such that $K_i = K_{\rho(i)}$. Such a reindexing is called a *refinement* if ρ is injective.

For the computation of persistent homology a simplex-wise filtration is needed as this gives rise to a *filtration boundary matrix*. The fact that we can refine a filtration to make it simplex-wise and still obtain the same persistence barcodes in due to the following fact:

Proposition 4.1.16. *[Bau21, Proposition 2.1.] Let $f : K \to \mathbb{R}$ be a monotonic function on a simplicial complex K, and let $K_I = (K_i)_{i\in I}$ with $K_i = \{\sigma_k \mid k \in I,\, k \geq i\}$ be an essential simplex-wise refinement of the sublevel set filtration $K_{\mathbb{R}} = (f^{-1}(\infty, \varepsilon])_{\varepsilon\in\mathbb{R}}$. The persistence barcode of $(K_i)_{i\in I}$ determines the persistence barcode of $(f^{-1}(\infty, \varepsilon])_{\varepsilon\in\mathbb{R}}$:*

$$B(H_k(K_{\mathbb{R}})) = \{r^{-1}[i, j) \neq \emptyset \mid [i, j) \in B(H_k(K_I))\}$$

for all k. For $j < \infty$, we have $r^{-1}[i, j) = [f(\sigma_i), f(\sigma_j))$, and $r^{-1}[i, \infty) = [f(\sigma_i), \infty)$.

Remark 4.1.17. Given a point cloud $\mathbb{X} = \{x_0, \ldots, x_{n-1}\}$ with a total order on its elements. Note that such an ordering can be obtained by simply numbering its elements. The simplices in the full Vietoris-Rips complex associated to $\mathbb{X}$ can then be ordered by the lexicographic order induced by the total order on the vertices (through the order on $\{x_0, \ldots, x_{n-1}\}$). An essential and simplex-wise filtration of the simplices in the Vietoris-Rips complex can then be obtained by

- first ordering by diameter of the simplices,
- then by simplex dimension
- and lastly by lexicographic order induced by the total order on the vertices.

4.1.2 Local Persistent Homology

Consider a topological space X, and let $\mathbb{X}$ be a sample extracted from X. Our goal is to quantify topological properties of X through $\mathbb{X}$, specifically focusing on computing local homology at some point $p \in X$.

In reality, attempting to calculate $H_d(\mathbb{X}, \mathbb{X} \setminus p)$ for a discrete space $\mathbb{X}$ is destined to fail. Instead, our approach involves computing $H_d(X, X \setminus B_r(p))$, equivalent to local homology of X at $p \in X$ for a sufficiently small radius r. A possible strategy is to approximate $H_d(X, X \setminus B_r(p))$ by constructing a simplicial approximation of X and $X \setminus B_r(p)$ from $\mathbb{X}$ and subsequently applying an algorithm to compute relative homology. Algorithms for relative persistent homology are available [BB20, RB21]. In [BCSE$^+$07, BWM12], the authors detailed how to compute local homology directly from a given point cloud, typically based on Delaunay complexes, considered suitable primarily for low-dimensional applications. In [SW14], the authors explored the computation of local homology from Vietoris-Rips complexes, providing theoretical ideas for an algorithmic approach. In the following, we will present the construction and implementation of a single filtered simplicial complex suitable for approximating local persistent homology.

Given our emphasis on computing persistent (local) homology from Vietoris-Rips complexes, we commence this section with a concise review of flag complexes and neighborhood graphs. A flag complex can be entirely characterized by its one-dimensional skeleton, represented as an undirected graph denoted as $G = (V, E)$, where V is a set of vertices and E is a set of edge tuples. Our chosen approach aligns with the methodology outlined in [Zom10], presenting a fast iterative algorithm based on representing the 1-skeleton of a simplicial complex through a neighborhood graph.

Definition 4.1.18. A *neigborhood graph* is a tuple (G, w) with $G = (V, E)$ an undirected graph and $w : E \to \mathbb{R}$ a function defined on its edges.

In general, from an undirected graph $G = (V, E)$ we can construct a complex as the collection of all cliques in the graph, that is all subsets $C \subset V$ such that $x, y \in C \Rightarrow \{x, y\} \in E$. Having a weight function w on the set of edges, i.e. having a neighborhood graph (G, w), naturally introduces a filtration. We do not need the full generality and will focus on more specific cases in the following.

Definition 4.1.19. Let $(\mathbb{X}, d)$ be a point cloud and $\varepsilon > 0$. The *Vietoris-Rips neighborhood graph at scale* ε, denoted $(G_\varepsilon^{Rips}(\mathbb{X}) = (\mathbb{X}, E_\varepsilon(\mathbb{X}), w)$, is defined by setting

$$E_\varepsilon(\mathbb{X}) = \{\{x, y\} \mid w(\{x, y\}) \leq \varepsilon \text{ and } x \neq y \in \mathbb{X}\}$$

and

$$w(\{x, y\}) = d(x, y) \text{ for } \{x, y\} \in E_\varepsilon.$$

In this case the Vietoris-Rips complex at scale ε is constructed with E_ε as 1-skeleton, i.e.

$$Rips_\varepsilon(\mathbb{X}) = V \cup E_\varepsilon(\mathbb{X}) \cup \{\sigma \mid \{x, y\} \subset \sigma \Rightarrow \{x, y\} \in E_\varepsilon(\mathbb{X})\}$$

and the function $f : Rips_\varepsilon(\mathbb{X}) \to \mathbb{R}$ defined by

$$f(\sigma) = \begin{cases} 0 & \text{if } \sigma = \{x\}, \\ d(x, y) & \text{if } \sigma = \{x, y\}, \\ \max_{\{x,y\} \subset \sigma} d(x, y) & \text{else.} \end{cases}$$

To keep the notation neat at this stage we will assume that $X, \mathbb{X} \in \mathbf{Sam}_\star$ and are centered at 0. We will omit the center point for that reason in the following. Note that we want to ultimately explain how to calculate $\Phi^d_{\mathcal{PL}*} \circ \mathcal{M}^r$ and we will therefore normalize and centralize any input for the function $\Phi^d_{\mathcal{PL}*}$. Thus, the aim is to explain how to calculate $H_*(X, X \setminus \operatorname{int} B_{\frac{1}{2}})$ with the B denoting the standard Euclidean ball, centered at 0. In other words, X will represent the magnification of a space at a certain point and $\mathbb{X}$ will be the magnification of a sample of X. Note that this also

assumes that we only look at spaces that live in some Euclidean space. This allows for the use certain geometric constructions, summarized in the following remark.

Remark 4.1.20. We make the work of Skraba and Wang in [SW14] our stating point, in the form of Definition 4.1.21 and Theorem 4.1.22. We omit a full discussion and a proof of their results as this is not our work. However, probably the essential point to prove the statement we give here in Theorem 4.1.22 is the choice of a weight function on all edges. Let $m = m(x,y) = \frac{x+y}{2}$ denote the *midpoint* for arbitrary $x, y \in \mathbb{X}$. Basic geometric considerations reveal that

$$\partial(x,y) = \sqrt{\frac{1}{4} - \big(\frac{m^T(x-y)}{||x-y||}\big)^2 - ||m||^2 - \big(\frac{m^T(x-y)}{||x-y||}\big)^2 + \frac{||x-y||^2}{4}}$$

is exactly the minimal value α such that $(B_\alpha(x) \cap B_\alpha(y)) \setminus B_{\frac{1}{2}} \neq \emptyset$ for $\alpha < \frac{1}{2}$. If we this formula to determine the filtration value of a 1-simplex (for example in a Vietoris-Rips complex), then we associated to it the length of an edge in the nerve complex of the union of balls of radius $\alpha < \frac{1}{2}$ centered at all points of $\mathbb{X} \setminus \text{int } B_{\frac{1}{2}}$. Therefore, this construction may serve as the simplicial approximation of $X \setminus B_{\frac{1}{2}}$ from a point cloud $\mathbb{X}$ close to X.

Definition 4.1.21. Let $\mathbb{X} \in \mathbf{Sam}_\star$ be a point cloud and $0 \leq \varepsilon < \frac{1}{2}$. Let $(G_\varepsilon^{\mathcal{F}}(\mathbb{X}), v)$ be the graph neighborhood given by

$$V = \{x \in \mathbb{X} \mid ||x|| \geq \frac{1}{2} - \varepsilon\},$$
$$E = \{\{x,y\} \mid v(\{x,y\}) \leq \varepsilon \text{ and } x \neq y \in V\}$$

and

$$v(\{x,y\}) = \begin{cases} ||x-y|| \text{ if } ||m(x,y)|| \geq \frac{1}{2}, \\ \partial(x,y) \text{ else.} \end{cases} \tag{4.4}$$

The flag complex induced by $(G_\varepsilon^{\mathcal{F}}(\mathbb{X}), v)$, denoted $\mathcal{F}_\varepsilon(\mathbb{X})$, is filtered by extending v to a function $f : \mathcal{F}_\varepsilon(\mathbb{X}) \to \mathbb{R}$ given by

$$f(\sigma) = \begin{cases} \frac{1}{2} - \varepsilon & \text{if } \sigma = \{x\}, \\ v(\{x,y\}) & \text{if } \sigma = \{x,y\}, \\ \max_{\{x,y\} \subset \sigma} v(\{x,y\}) & \text{else.} \end{cases}$$

The upshot of this above construction is the following immediate consequence of [SW14, Theorem 3.2] and [SG07, Theorem 2.5] (see Remark 4.1.14 for an explanation of this result):

Theorem 4.1.22. *Let $X \in \mathbf{Sam}_\star$ and $\mathbb{X} \in \mathbf{Sam}_\star$ a δ-sample of X. Then, the persistence modules $H_d(Rips_{[0,\frac{1}{2})}(\mathbb{X}), \mathcal{F}_{[0,\frac{1}{2})}(\mathbb{X}))$ and $\{H_d(X_{[0,\frac{1}{2})}, X_{[0,\frac{1}{2})} \setminus int\, B_{\frac{1}{2}})$ are $(2\delta + \frac{1}{2}(\sqrt{\frac{2(d+1)}{d+2}} - 1))$-interleaved.*

Remark 4.1.23. It's important to note that the statement about interleaving primarily stems from the interleaving of the Vietoris-Rips filtration with a Čech filtration, as discussed in Remark 4.1.14. Therefore, if one opts for a Vietoris-Rips type complex, one must acknowledge the interleaving distance described above. While it is mathematically feasible to construct a similar framework for a Čech type complex, the challenge lies in the fact that a Čech complex is not a flag complex. This limitation prevents the extension of the weight function on the edges of a graph neighborhood to all simplices. Consequently, computing the correct filtration value for simplices of dimension greater than 1 becomes significantly more intricate both geometrically and computationally.

Corollary 4.1.24. *Let $X \in \mathbf{Sam}_\star$ and $\mathbb{X} \in \mathbf{Sam}_\star$ a δ-sample of X. Additionally, assume that there exists $0 \leq \gamma < \frac{1}{2}$ such that $\tilde{H}_d(Rips_\varepsilon(\mathbb{X})) = 0$ for all $\gamma \leq \varepsilon \leq \frac{1}{2}$ and $d \geq 0$. Then, the persistence modules $\tilde{H}_{d-1}(\mathcal{F}_{[\gamma,\frac{1}{2})}(\mathbb{X}))$ and $\tilde{H}_d(X_{[\gamma,\frac{1}{2})}, X_{[\gamma,\frac{1}{2})} \setminus int\, B_{\frac{1}{2}})$ are $(2\delta + \frac{1}{2}(\sqrt{\frac{2(d+1)}{d+2}} - 1))$-interleaved for all $d > 0$.*

Proof. This follows directly from Theorem 4.1.22 and the long exact sequence for the pair $(Rips_\varepsilon(\mathbb{X}), \mathcal{F}_\varepsilon(\mathbb{X}))$. □

Remark 4.1.25. The above lemma indicates when it is possible to use the complex $\mathcal{F}_\varepsilon(\mathbb{X})$ based on the findings in [SW14] directly to approximate local persistent homology. But we would like to stress the point that this statement should be taken with caution. Having $\tilde{H}_d(Rips_\varepsilon(\mathbb{X})) = 0$ for all $\gamma \leq \varepsilon \leq \frac{1}{2}$ and $d \geq 0$ is in general not a trivial assumption. The assumption is certainly met, for example, if $\mathbb{X}_\varepsilon$ is a contractible space from some scale $\varepsilon \geq \gamma$ and on-wards. If X here actually plays the role of a magnification at a given point of a Lojasiewicz-Whitney space which converges to the tangent cone at the same point (see Section 3.3, then the magnification will, for sufficiently small radius, be close to a space that is contractible. Here, closeness is to be understood in Hausdorff distance which implies that the considered persistence modules are proportionally close in interleaving distance. Then, the persistence module $\tilde{H}_d(Rips_{[\gamma,\frac{1}{2})}(\mathbb{X}))$ will be close to zero if $\mathbb{X}$ is sufficiently close to X. However, this is not enough to deduct the conditions stated for Corollary 4.1.24 and therefore to legitimize the use of $\mathcal{F}_\varepsilon(\mathbb{X})$ to approximate local persistent homology.

For cases where the additional assumptions of Corollary 4.1.24 seem unreasonable or in case the additional step of approximation in form of an interleaving only from

certain scales on-wards is unacceptable, we provide another, more accurate method that assumes no vanishing persistence modules to function.

Definition 4.1.26. Let $\mathbb{X} \in \mathbf{Sam}_\star$, $0 \leq \varepsilon < \frac{1}{2}$ and let $\{c\}$ be a set disjoint from $\mathbb{X}$. Let $(G_\varepsilon^{LRips}(\mathbb{X}), e)$ denote the graph neighborhood given by

$$\begin{aligned} V &= \mathbb{X} \cup \{c\}, \\ E &= \{\{x, y\} \mid v(\{x, y\}) \leq \varepsilon \text{ and } x \neq y \in V\} \end{aligned}$$

and $e : \{\{x, y\} \subset \mathbb{X} \cup \{c\} \mid x \neq y\} \to \mathbb{R}$ by

$$e(\{x, y\}) = \begin{cases} \max\{0, \frac{1}{2} - ||x||\} \text{ if } y = c \\ \max\{0, \frac{1}{2} - ||y||\} \text{ if } x = c \\ ||x - y|| \text{ else,} \end{cases} \tag{4.5}$$

with v as in Eq. (4.4). The *local Vietoris-Rips* complex, $LRips_\varepsilon(\mathbb{X})$, is given by all simplices σ in the flag complex of the graph neighborhood $(G_\varepsilon^{LRips}(\mathbb{X}), e)$ such that $g(\sigma) < \frac{\varepsilon}{2}$ with g defined on the flag complex by

$$g(\sigma) = \begin{cases} 0 & \text{if } \sigma = \{x\}, \\ e(\{x, y\}) & \text{if } \sigma = \{x, y\}, \\ \max_{\{x,y\} \subset \sigma \setminus \{c\}} v(\{x, y\}) & \text{if } c \in \sigma, \\ \max_{\{x,y\} \subset \sigma} ||x, y|| & \text{else.} \end{cases} \tag{4.6}$$

Furthermore, to this single simplicial complex one could directly apply any available standard algorithm for persistent homology. It is possible to profit from all the optimizations done to compute persistent homology such as, e.g., clearing [CK11] in conjunction with using cohomology [CK11, SMVJ11, Bau21] with this approach and additionally using simplicial collapses [BMS20, BP20]. This is not clear for the case of using persistent relative homology.

Lemma 4.1.27. *Let $\mathbb{X} \in \mathbf{Sam}_\star$. Then, the persistence modules $\tilde{H}_d(LRips_{[\gamma, \frac{1}{2})}(\mathbb{X}))$ and $\tilde{H}_d(Rips_{[\gamma, \frac{1}{2})}(\mathbb{X}), \mathcal{F}_{[\gamma, \frac{1}{2})}(\mathbb{X}))$ are isomorphic for any $d \geq 0$.*

Proof. For ease of exposition let $K_\varepsilon = Rips_\varepsilon(\mathbb{X})$, $L_\varepsilon = \mathcal{F}_\varepsilon(\mathbb{X})$ and $M_\varepsilon = LRips_\varepsilon(\mathbb{X})$ and let $0 \leq \varepsilon < \frac{1}{2}$ in this proof. The main observation here is that our construction ensures that $M_\varepsilon = K_\varepsilon \cup_i C(L_\varepsilon)$ is the simplicial mapping cone with $L_\varepsilon \overset{i}{\hookrightarrow} K_\varepsilon$ the canonical inclusion. To show this, let $\sigma \subset \mathbb{X} \cup \{c\}$. First, if $c \notin \sigma$ we can immediately see that $\sigma \in K_\varepsilon \cup_i C(L_\varepsilon)$ if and only if $\sigma \in M_\varepsilon$ as the respective filtration functions f and g agree on σ. If, however, $c \in \sigma$ we consider two further cases. In case

$\sigma \setminus \{c\} = \{x\}$ we know that $\frac{1}{2} - ||x|| \leq \varepsilon$ if and only if $\sigma \in M_\varepsilon$ by definition of e and in $K_\varepsilon \cup_i C(L_\varepsilon)$ because $\{x\} \in L_\varepsilon$. Now, consider $\dim \sigma > 2$. The statement $\sigma \in K_\varepsilon \cup_i C(L_\varepsilon)$ is equivalent to $g(\sigma - \{c\}) \leq \varepsilon$, i.e. $\max_{\{x,y\} \subset \sigma \setminus \{c\}} w(x,y) \leq \varepsilon$. By definition of g, this is the case if and only if $g(\sigma) \leq \varepsilon$. We now only need the basic argument that links the homology of the cone and relative homology which we included at this point for the sake of completeness. We start with the long exact sequence of the pair $(K_\varepsilon \cup_i C(L_\varepsilon), C(L_\varepsilon))$ and because $C(L_\varepsilon) \cong \{*\}$ we have a commutative diagram

$$\begin{array}{ccccccccc}
H_d(C(L_\varepsilon)) & \longrightarrow & H_d(K_\varepsilon \cup_i C(L_\varepsilon)) & \longrightarrow & H_d(K_\varepsilon \cup_i C(L_\varepsilon), c(L_\varepsilon)) & \longrightarrow & H_{d-1}(C(L_\varepsilon)) \\
\downarrow \cong & & \downarrow = & & \downarrow = & & \downarrow \cong \\
0 & \longrightarrow & H_d(K_\varepsilon \cup_i C(L_\varepsilon)) & \xrightarrow{\cong} & H_d(K_\varepsilon \cup_i C(L_\varepsilon), C(L_\varepsilon)) & \longrightarrow & 0
\end{array}$$

asserting $\tilde{H}_d(K_\varepsilon \cup_i C(L_\varepsilon)) \simeq \tilde{H}_d(K_\varepsilon \cup_i C(L_\varepsilon), C(L_\varepsilon))$ for $0 < \varepsilon < \frac{1}{2}$. Furthermore, by excising the closed subset $\{c\}$ of $K_\varepsilon \cup_i C(L_\varepsilon)$ and $C(L_\varepsilon)$ we see that $\tilde{H}_d(K_\varepsilon \cup_i C(L_\varepsilon)) \cong \tilde{H}_d(K_\varepsilon \cup_i C(L_\varepsilon) - \{c\}, C(L_\varepsilon) - \{c\})$ which in turn is equivalent to $\tilde{H}_d(K_\varepsilon, L_\varepsilon)$ for all i because $C(L_\varepsilon) - \{c\}$ is a deformation retract of L_ε for all i and $0 < \varepsilon < \frac{1}{2}$. In total, we have two sequences of vector spaces

$$\begin{array}{ccccccc}
\cdots & \longrightarrow & \tilde{H}_d(K_\varepsilon, L_\varepsilon) & \longrightarrow & \tilde{H}_d(K_{\varepsilon+\alpha}, L_{\varepsilon+\alpha}) & \longrightarrow & \cdots \\
& & \downarrow \cong & & \downarrow \cong & & \\
\cdots & \longrightarrow & \tilde{H}_d(M_\varepsilon) & \longrightarrow & \tilde{H}_d(M_{\varepsilon+\alpha}) & \longrightarrow & \cdots
\end{array}$$

where each horizontal morphism is induced by inclusion and every vertical isomorphism is constructed as above from the long exact sequence and excision. □

Finally, we are ready to quantify the worst case error of approximation of local persistent homology of a space X at any point from a given sample $\mathbb{X}$ of X. with our construction in Definition 4.1.26.

Proposition 4.1.28. *Let $X \in \mathbf{Sam}_\star$ and let $\mathbb{X} \in \mathbf{Sam}_\star$ be δ-sample of X. Then the relative persistence module $\tilde{H}_d(X_{[\gamma,\frac{1}{2})}, (X_{[\gamma,\frac{1}{2})} \setminus B_{\frac{1}{2}}))$ and the persistence module $\tilde{H}_d(LRips_{[\gamma,\frac{1}{2})}(\mathbb{X}))$ are $(2\delta + \frac{1}{2}(\sqrt{\frac{2(d+1)}{d+2}} - 1))$-interleaved for all $d > 0$.*

Proof. By Theorem 4.1.22 we see that the persistence modules $\tilde{H}_d(X_{[\gamma,\frac{1}{2})}, (X_{[\gamma,\frac{1}{2})} \setminus \text{int } B_{\frac{1}{2}})$ and $\tilde{H}_d(Rips_{[\gamma,\frac{1}{2})}(\mathbb{X}), \mathcal{F}_{[\gamma,\frac{1}{2})}(\mathbb{X}))$ are $(2\delta + \frac{1}{2}(\sqrt{\frac{2(d+1)}{d+2}} - 1))$-interleaved. Lemma 4.1.27 tells us that $\tilde{H}_d(Rips_{[\gamma,\frac{1}{2})}(\mathbb{X}), \mathcal{F}_{[\gamma,\frac{1}{2})}(\mathbb{X}))$ has interleaving distance 0 from $\tilde{H}_d(LRips_{[\gamma,\frac{1}{2})}(\mathbb{X}))$. The triangle inequality for the interleaving distance then implies the desired statement. □

4.1.3 Implementation

We now turn towards the actual implementation of the method we just described. In the center of the computation of local homology with the Vietoris-Rips complex, a flag complex, lies the implementation of the graph neighborhood construction. Although the full complex is constructed from the flag complex of a neighborhood we also have to take into account that the filtration value of the simplices is not determined by the highest filtration value of all edge (see Eq. (4.5)) contained in a simplex (see Eq. (4.6)). Consequently, for a given filtration value it is possible to have constructed simplices that have higher filtration value than the highest weight of all edges contained in the simplex from the complete flag complex. More specifically, we will show how to construct the local Vietoris-Rips graph neighborhood form which one could then construct the local Vietoris-Rips complex.

For the more algorithmic point of view that we will take now we have to pay attention to the actual implementation of graph neighborhoods. A graph neighborhood is then given by a so-called *adjacency lists*, i.e. a hash table G with keys given by the points of a given space $\mathbb{X} = \{x_0, \ldots, x_{n-1}\} \in \mathbf{Sam}_\star$ (or equivalently by the index for ordered sets) such that $G[x_i]$ contains a sorted list of neighbors with higher index. To such a list $G[x_i]$ we will be referring to as neighborhood of x_i.

The algorithm in form of pseudocode is provided in Algorithm 1 in which we will follow the notation as adjacency list. The neighborhood is constructed to the maximal scale $\frac{1}{2}$. We iterate through all points of $\mathbb{X}$. At each iteration step we include the current point x_i to its own neighborhood A. For each x_i we iterate through all points x_j with $i < j$ and check at each step if we have to add an edge $\{x_i, x_j\}$ to the graph neighborhood. This is true if and only if $||x_i - x_j|| < \frac{\varepsilon}{2}$. If we include the edge we may safe the edge weight given by $||x_i - x_j||$, compare to Eq. (4.5). After we iterated through all x_j with $j > i$ we included all edges containing x_i of the Vietoris-Rips graph neighborhood at scale $\frac{1}{2}$ by adding the neighborhood A to G. At the end of each iteration step we include the cone point c in the adjacency list of x_i if x_i has at distance from the origin greater than $\frac{1}{2}$. We store A in G with x_i as key.

Moving on to Algorithm 2, we describe the alternative graph neighborhood to construct $\mathcal{F}_{\frac{1}{2}}$. As the routine to determine the graph neighborhood for $\mathcal{F}_{\frac{1}{2}}$ is very similar to what we previously described we only highlight the differences. First, we have no extra point, or cone point included in our point cloud and therefore no additional edges between points of $\mathbb{X}$ and c to be included. The other main difference is in the weight of an edge and thus the condition to be included in G. Because the edge weight for $\mathcal{F}_{\frac{1}{2}}$ depends on the location of the midpoint between the two points of a potential edge there is an extra condition in line 6-11 as well as a different weight assigned in

Algorithm 1 $G^{LRips}(\mathbb{X})$

Input: Vertices $V = \{c\} \cup \mathbb{X} = \{x_0, \ldots, x_{n-1}\} \in \mathbf{Sam}_\star$

Output: Local Vietoris-Rips graph neighborhood $(G^{LRips}_{\frac{1}{2}}(\mathbb{X} \cup \{c\}), e)$

1: $G \leftarrow \emptyset$; $e \leftarrow \emptyset$;
2: **for** $i \leftarrow 0$ to $n-1$ **do**
3: $\quad A \leftarrow \{x_i\}$;
4: $\quad$ **for** $j \leftarrow i+1$ to $n-1$ **do**
5: $\quad\quad$ **if** $||x_i - x_j|| < \frac{1}{2}$ **then**
6: $\quad\quad\quad A \leftarrow A \cup \{x_j\}$; $e(\{x_i, x_j\}) \leftarrow ||x_i - x_j||$;
7: $\quad$ **if** $\max\{0, \frac{1}{2} - ||x_i||\} < \frac{1}{2}$ **then**
8: $\quad\quad A \leftarrow A \cup \{c\}$; $e(\{x_i, c\}) \leftarrow \max\{0, \frac{1}{2} - ||x_i||\}$;
9: $\quad G[x_i] \leftarrow A$;

line 11 given by Eq. (4.4).

Algorithm 2 $G^{\mathcal{F}}(\mathbb{X})$

Input: Point cloud $\mathbb{X} = \{x_0, \ldots, x_n\} \in \mathbf{Sam}_\star$

Output: Graph neighborhood $(G^{\mathcal{F}}_{\frac{1}{2}}(\mathbb{X}), w)$

1: $G \leftarrow \emptyset$;
2: **for** $i \leftarrow 0$ to $n-1$ **do**
3: $\quad$ **if** $||x_i|| > 0$ **then**
4: $\quad\quad A \leftarrow \{x_i\}$;
5: $\quad\quad$ **for** $j \leftarrow i+1$ to n **do**
6: $\quad\quad\quad$ **if** $||m(x_i, x_j)|| \geq \frac{1}{2}$ **then**
7: $\quad\quad\quad\quad$ **if** $\frac{||x_i - x_j||}{2} < \frac{1}{2}$ **then**
8: $\quad\quad\quad\quad\quad A \leftarrow A \cup \{x_j\}$; $w(\{x_i, x_j\}) \leftarrow ||x_i - x_j||$;
9: $\quad\quad\quad$ **else**
10: $\quad\quad\quad\quad$ **if** $\partial(x_i, x_j) < \frac{1}{2}$ **then** ▷ Compare to Eq. (4.4)
11: $\quad\quad\quad\quad\quad A \leftarrow A \cup \{x_j\}$; $w(\{x_i, x_j\}) \leftarrow \partial(x_i, x_j)$;
12: $\quad\quad G[i] \leftarrow A$;

In [Zom10], Zomorodian described a fast algorithm to compute the Vietoris-Rips complex incrementally from a given neighborhood graph. We included a modified version of Zomorodian's algorithm suitable for the construction of the simplicial complexes used for local homology computations with Algorithm 3. We will go through the algorithm to highlight the important adaptations to work for *LRips* but for the statement that this incremental algorithm accurately computes the k-skeleton of the

Vietoris-Rips complex associated with a graph neighborhood consider [Zom10, Theorem2]. Algorithm 3 is essentially a loop over all vertices included in the input graph neighborhood G. At each iteration we call the function AddCofaces, see Algorithm 4, and pass a container of all simplicies K to be filled, all the information on neighboring vertices G and $G[v]$, a simplex $\{v\}$, the maximal diameter of a simplex ε allowed to be included in K and the maximal dimension of simplices k. The routine AddCofaces can be found in Algorithm 4. Before we turn to that routine we want to note that in the end of Algorithm 3 we order all simplices in K. This is done as described in Remark 4.1.17 first by dimension, then by diameter and finally by the order on the vertices so the output is an essential and simplex-wise filtration. This is necessary to compute persistent homology algorithmically, compare to Proposition 4.1.16.

Algorithm 3 IterativeRips($(G, w), d$)

Input: Neighborhood graph (G, w), maximal dimension d

Output: Filtered simplicial complex K

1: $K \leftarrow \emptyset$;
2: **for** $v \leftarrow 0$ to $G.size$ **do**
3: AddCofaces$(K, (G, w), G[v], \{v\}, d)$; ▷ See Algorithm 4
4: sort(K); ▷ Create essential simplex-wise filtration (Remark 4.1.17)

Moving on to Algorithm 4, the given neighborhood N represents the intersection of all vertex neighborhoods of all vertices in s. Therefore, the simplex $\tau = s \cup \{v\}$ for any $v \in N$ constitutes a simplex in the flag of G. We call AddCofaces$(-)$ for any τ and $v \cap N$ as neighborhood. Again, check [Zom10] for more details on the theoretical considerations. We want to focus on the addition of lines 4 and 5. Here we check whether or not the diameter of the given simplex s is smaller the prespecified scale ε. This would be an unnecessary step if we actually wanted to compute the Vietoris-Rips complex associated with G because the diameter of a simplex is then the maximal weight of all its edges and are therefore below the scale imposed on G. In the case of *LRips* the diameter of a simplex is not given by the maximal weight of all its edges. Instead *LRips* is given by the flag complex of G minus all simplices that lie above the scale imposed in G. Thus the additional condition in Algorithm 4. We will not detail an algorithm to compute the simplex diameter but it should be noted that in order to compute it we actually do not have to parse the geometrical information included in $\mathbb{X} \in \mathbf{Sam}_\star$ as the diameter can still be completely determined by the weight of all its edges which is included with w in case of a standard Vietoris-Rips complex or *LRips*. In case of $\mathcal{F}$ we have to be able to compute the distance to the central point, i.e. the origin for $\mathbb{X} \in \mathbf{Sam}_\star$.

Algorithm 4 AddCofaces($K, (G, w)$,N,s, d);

Input: Simplicial complex K, neighborhood graph (G, w), neighborhood N, simplex s, dimension d

```
if dim(s) ≥ d then
    Return;
if diam(s) < 1/2 then                          ▷ Not always necessary
    K ← K ∪ {σ}
for v ∈ N do
    if v ∉ s then
        τ ← s ∪ {v};
        AddCofaces(K, (G, w), G[v] ∩ N, τ, d);
```

With Algorithm 1 and Algorithm 2 at hand, we are in position to apply Algorithm 3 in order to construct $LRips$ and $\mathcal{F}$. The only point that is missing from the full picture are the specific filtration functions, compare to Eq. (4.6) and Eq. (4.4) respectively. However, we feel that computing these values is sufficiently clear so we omit the discussion of further pseudocode sequences.

Putting all the previously presented algorithms together we have Algorithm 5 which described how to approximate the $\Phi^d_{\mathcal{PL}}$. We record the closeness of approximation in Proposition 4.1.29.

Algorithm 5 $\Phi_{\mathcal{PL}}(\mathbb{X}, d)$

Input: Point cloud $\mathbb{X} \in \mathbf{Sam}_\star$, dimension d

Output: $\Phi_{\mathcal{PL}}(\mathbb{X})$

```
(G, w) ← G^{LRips}(X);                          ▷ compare to Algorithm 1
K ← IterativeRips((G, w), d);                   ▷ compare to Algorithm 3
i ← 1
for i ≤ d do
    D ← PH_i(K);          ▷ Persistent homology of filtered simplicial complex
    if d_B(D, PH_i(S^d)) is maximal then
        Φ_PL(X, d) ← 1 − 2 d_B(D, PH_i(S^d))    ▷ compare to Remark 4.1.8
```

Proposition 4.1.29. *Let $X, \mathbb{X} \in \mathbf{Sam}_\star$ and $\mathbb{X}$ a δ-sample of X. Then, for $\Phi_{\tilde{\mathcal{PL}}}(\mathbb{X}, d)$ denoting the output of Algorithm 5 for reduced homology and $\Phi^d_{\tilde{\mathcal{PL}}}$ the function described in Example 3.2.3 with reduced homology we have*

$$| \Phi_{\tilde{\mathcal{PL}}}(\mathbb{X}, d) - \Phi^d_{\tilde{\mathcal{PL}}}(X) | \leq 2\delta + (\sqrt{\frac{2(d+1)}{d+2}} - 1).$$

Proof. Let $D^i_{S^d} = B(\mathcal{P}\tilde{H}_i(S^d))$, $LD^i_X = B(\mathcal{P}\tilde{H}_i(X, X_\varepsilon \backslash B_{\frac{1}{2}}))$ and $LD^i_{\mathbb{X}} = B(\mathcal{P}\tilde{H}_i(LRips(\mathbb{X})))$ the barcodes associated to the persistent homology of S^d, X and $LRips(\mathbb{X})$ respectively (compare to Theorem 4.1.2). Then, we have

$$\Phi^d_{\tilde{\mathcal{P}\mathrm{L}}}(X) = 1 - 2\max_{i\leq d} d_B(LD^i_X, D^i_{S^d})$$

and

$$\Phi_{\tilde{\mathcal{P}\mathrm{L}}}(\mathbb{X}, d) = 1 - 2\max_{i\leq d} d_B(LD^i_{\mathbb{X}}, D^i_{S^d})$$

Using this we see that

$$\begin{aligned} \mid \Phi^d_{\tilde{\mathcal{P}\mathrm{L}}}(X) - \Phi_{\tilde{\mathcal{P}\mathrm{L}}}(\mathbb{X}, d) \mid &= \mid \max_{i\leq d} d_B(LD^i_X, D^i_{S^d}) - \max_{i\leq d} d_B(LD^i_{\mathbb{X}}, D^i_{S^d}) \mid \\ &\leq \max_{i\leq d} d_B(LD^i_X, LD^i_{\mathbb{X}}). \end{aligned}$$

Invoking Lemma 4.1.27 and Eq. (4.2) we have

$$\mid \Phi_{\tilde{\mathcal{P}\mathrm{L}}}(\mathbb{X}, d) - \Phi^d_{\tilde{\mathcal{P}\mathrm{L}}}(X) \mid \leq 2\delta + (\sqrt{\frac{2(d+1)}{d+2}} - 1)$$

which is what we claimed. □

Remark 4.1.30. The above result aims to quantify the worst-case error associated with our proposed method for approximating the function $\Phi_{\mathcal{P}\mathrm{L}}$, as outlined in Example 3.2.3. It is crucial to evaluate this result independently of the utility of Algorithm 5. The interleaving error, as detailed in Proposition 4.1.29, does not provide a satisfactory approximation of $\Phi_{\mathcal{P}\mathrm{L}}$ when using Vietoris-Rips complexes. Further exploration in this direction, possibly employing a different approach, may yield a more nuanced assessment. However, the true effectiveness of Algorithm 5 can also be assessed through practical results. In practical terms, the algorithmic Φ demonstrates its capacity to discern regions resembling the topology of a linear subspace from spaces with persistent homology barcodes that are distant from these. We will illustrate these practical aspects in Section 4.3.2 and the subsequent Chapter 5.

4.2 Tangential Approximation

This section is primarily concerned with computing the Hausdorff distance between arbitrary subspaces of Euclidean space $\mathbb{R}^n$ and linear subspaces, with the goal of identifying linear subspaces in $\mathbb{R}^n$ that minimize this distance. To achieve this objective, we will explore the Hausdorff distance to some extend and subsequently focus on optimization techniques within Grassmannians. The overarching aim of this section is to pave the way for the implementation of an algorithm aimed at approximating the function Φ^d_{Hd} introduced earlier in Example 3.2.2.

4.2.1 Hausdorff distance between subsets of $\mathbb{R}^n$

We have used Hausdorff distances before as the metric on the non-empty compact subsets of $\mathbb{R}^n$ for the definition of $\mathbf{Sam}, \mathbf{Sam}_\star, \mathbf{Sam}_P, \mathbf{D}_P\mathbf{Sam}$. However, we feel that it is advisory to recall this concept at this point before we go any further. Let $A, B \subset \mathbf{Sam}_\star$. The Hausdorff distance is then given by

$$d_{Hd}(A,B) = \inf\{\varepsilon > 0 \mid A_\varepsilon \subset B \text{ and } B_\varepsilon \subset A\}. \tag{4.7}$$

which is equivalently given by

$$d_{Hd}(A,B) = \max\{\sup_{a\in A} d(a,B), \sup_{b\in B} d(A,b)\} \tag{4.8}$$

which is finite in case A and B are finite. Eventually, we want to compute $\Phi^d_{Hd*} \circ \mathcal{M}^r$ and thus the Hausdorff distance only needs to be considered for elements of $\mathbf{Sam}_\star$, i.e. spaces isometric to subsets of the standard Euclidean n-ball B^n. To abbreviate the notation at this point we write $\mathbb{X} \in \mathbf{Sam}_\star$ if we want to indicate that $\mathbb{X}$ is centered at 0 and scaled to lie in B^n.

If the first argument of the distance function is a linear subspace of $\mathbb{R}^n$ then for a given basis $[v_1, \ldots, v_k]$ of $V \in \mathrm{Gr}(k,n)$ we can simplify 4.8 to

$$\max\{\max_{x\in\mathbb{X}} ||(x - \pi_V(x))||, \sup_{v\in V, ||v||=1} d(\mathbb{X}, v)\} \tag{4.9}$$

where $\pi_V(-)$ is the projection onto V which can be done by multiplying x with the matrix $(v_1, \ldots, v_k)^T$. Observe that if we allow to deviate from the Hausdorff distance by only computing one side of the distance, that is

$$\vec{d}(V,\mathbb{X}) = \max_{x\in\mathbb{X}} ||(x - \pi_V(x))||, \tag{4.10}$$

we are left with a much simpler objective function. However, note that $\vec{d}$ is not a metric anymore as the symmetry axiom is not fulfilled. We also find $\vec{d}(V,\mathbb{X}) \leq d_{Hd}(V,\mathbb{X})$ for all V which shows that $\tilde{d}$ is continuous w.r.t. to the Hausdorff distance, i.e. convergence in d_{Hd} implies convergence in $\vec{d}$. However, consider the alternative characterization of d_{Hd} by Eq. (4.7). The distance $\vec{d}(V,\mathbb{X})$ being less than some $\alpha > 0$ implies $V \subset \mathbb{X}_\alpha$ but not necessarily $V_\alpha \subset \mathbb{X}$ as well. Thus, it is not hard to generate examples that converge in $\vec{d}$ but not in d_{Hd}. As a consequence, the class of Φ-stratified spaces using $\vec{d}$ instead of d_{Hd} is strictly smaller than the class of Φ^d_{Hd}-stratified spaces (compare to Examples 3.2.2 and 3.2.4). This brief discussion prepares the upcoming discussion on how to formulate an optimization problem with a deterministic approach.

4.2.2 Optimization on Grassmannians

Algorithmically computing Φ^k_{Hd} for a given $\mathbb{X} \in \mathbf{Sam}_\star$ involves addressing the optimization problem stated as follows:

$$\inf_{V \in \mathrm{Gr}(k,n)} \max\{\sup_{x \in \mathbb{X}} d(x, V), \sup_{v \in V, ||v|| \leq 1} d(\mathbb{X}, v)\}. \tag{4.11}$$

It is essential to note that we must confine ourselves to $||v|| \leq 1$ for all $v \in V$ to ensure a well-posed problem, utilizing the metric of $\mathbf{Sam}_\star$ instead of the general Hausdorff distance. Additionally, to streamline the discussion and eliminate the need to address the case of $\mathbb{X} = \emptyset$ each time (which is not crucial for our purposes), we assume $\mathbb{X} \neq \emptyset$.

Regarding the classification of Eq. (4.11) as an optimization problem, it is worth observing that the objective function is neither smooth nor convex. Furthermore, the optimization variable reappears as an infinite search space within the objective function, specifically in $\sup_{v \in V} d(\mathbb{X}, v)$. Consequently, the above optimization problem seems challenging to approach using many deterministic algorithmic optimization strategies. In this section, we will briefly discuss how to simplify the problem to obtain a smooth quadratic constrained optimization problem based on the one-sid. We will also introduce a random optimization approach to tackle the general problem. To this end, we will review pertinent results on Grassmannian optimization and random distributions on Grassmannians.

Deterministic approach

First, with the previous considerations we can simplify the optimization problem Eq. (4.11) with Eq. (4.10) to

$$\inf_{V \in \mathrm{Gr}(k,n)} \vec{d}(V, X) = \max_{x \in \mathbb{X}} ||(x - \pi_V(x))||. \tag{4.12}$$

if we are willing to reduce the class of stratified spaces whose stratification we will be able to approximate. To formulate an optimization problem within the space of k-dimensional linear subspaces of $\mathbb{R}^n$, denoted as $\mathrm{Gr}(k, n)$, let's take a brief detour to explore alternative descriptions of Grassmannians. Our primary references for this exploration are [EAS98, AEK06, LLY20]. The initial and arguably the most natural characterization of the space of linear subspaces is through the so-called *Stiefel manifold*. It is noteworthy that every linear subspace has an orthonormal basis and can be effectively identified with it. We represent the space of all k orthonormal bases in an n-dimensional space by

$$V(k, n) = \{M \in \mathbb{R}^{n \times k} \mid M^T M = I_{n \times n}\} \cong O(n)/O(n - k). \tag{4.13}$$

Of course, orthonormal bases are not unique but related to each other by orthogonal transformation and thus we have

$$\mathrm{Gr}(k,n) \cong V(k,n)/O(k). \tag{4.14}$$

First considered for algorithmic consideration in [EAS98], this approach is commonly used in optimization (see e.g. [JD15, WY13]).

Furthermore, as every linear subspace admits an orthonormal basis, there is a unique projection matrix associated with such a basis. Hence, one can show

$$\mathrm{Gr}(k,n) \cong P_{k,n} = \{P \in \mathbb{R}^{n\times n} \mid P^T = P = P^2 \text{ and } \mathrm{rank}(P) = k\}. \tag{4.15}$$

We will concentrate on this representation as we develop our randomized algorithm later, drawing inspiration from the work of [AEK06]. In addition to this, projection matrices find widespread use in various applications [Zha10, Chi03, Nic07]. For a more in-depth comparison of these methods and the advantages they offer, both in the context of optimization and alternative approaches for representing $\mathrm{Gr}(k,n)$ using so-called *symmetric involution matrices*, refer to [LLY20].

The above matrix representations of $\mathrm{Gr}(k,n)$ allow us to work with an element of the optimization search space as elements of a vector space either by row or by column major order. To summarize this discussion, we can reformulate the simplified distance in Eq. (4.12) by representing the Grassmannians as matrices e.g. by using the Stiefel manifold. With Eqs. (4.13) and (4.14) we obtain

$$\inf_{M\in\mathbb{R}^{n\times k}} \tilde{d}(v,X) = \inf_{M\in\mathbb{R}^{n\times k}} \max_{x\in\mathbb{X}} ||(x - MM^T x)|| \tag{4.16}$$

$$\text{subject to: } M^T M = I_{n\times n} \tag{4.17}$$

Here, the matrix $M \in \mathbb{R}^{n\times k}$ has rows representing the basis vectors of a k-linear subspace of $\mathbb{R}^n$. The described problem gives rise to a constrained, quadratic optimization problem with n^2 variables. The function within the maximum is convex, rendering the objective function convex, as the maximum of a convex function is itself convex. However, because Grassmannians are not convex, this problem is non-convex, and finding a global minimum may pose challenges. It's important to note that the presence of the maximum makes the problem non-smooth. To avoid dealing with the maximum in the objective function, we can alternatively consider the equivalent epigraph formulation

$$\inf_{M\in\mathbb{R}^{n\times k}, t\in\mathbb{R}} t \tag{4.18}$$

$$\text{subject to: } ||(x - MM^T x)|| \leq t \text{ for each } x \in \mathbb{X}, \tag{4.19}$$

$$M^T M = I_{n\times n}, \tag{4.20}$$

which is a smooth problem with inequality constraints equal to the cardinality of $\mathbb{X}$. Note that we could also get rid of the norm in the constraints by simply squaring it. While there are various existing implementations for solving such problems, it could be advantageous to explore specialized algorithms designed for optimization on matrix manifolds, such as the Stiefel manifold. As mentioned earlier, optimization within Grassmannians has been explored previously, commonly relying on the differential structures of Grassmannians represented by matrices. This approach involves defining tangent spaces and gradients and formulating an optimization algorithm based on these terms. For a more in-depth understanding of the implementation of optimization on matrix manifolds, refer to sources like [LLY20, KWT23]. The remaining challenges in solving the problem 4.18 are predominantly associated with devising a computationally efficient and stable implementation. We feel that at this point we have delineated the mathematical essence of the problem when attempting to solve the simplified optimization problem and leave the exploration of possible implementation and evaluations for future work.

Randomized approach

Let us now turn towards our randomized approach to approximate a solution for the original problem in 4.11. The fundamental concept behind a randomized approach is straightforward: Generate a set of elements $V_1, \dots, V_n \subset \mathrm{Gr}(k, n)$ randomly and record $\min f(V_i) \mid i = 1, \dots, n$. Consequently, we need a method for randomly generating elements V_i, and it would be advantageous to quantify the probability of being close to the optimal solution with respect to the measure used to generate V_i. This quantification enables us to approximate the required sample size n. To begin, we will review some findings on the uniform distribution on $P_{k,n}$.

Proposition 4.2.1. *Let $Z \in \mathbb{R}^{n\times k}$ be a matrix such that the elements are independent and identically $\mathcal{N}(0,1)$ distributed. Then, $P = Z(Z^TZ)^{-1}Z$ is uniformly distributed on $P_{k,n}$.*

Proof. See e.g. [Chi03, Theorem 2.2.2 (iii)]. □

Remark 4.2.2. This statement already suggests the algorithmic feasibility of obtaining random samples from Grassmannians. As mentioned in Eq. (4.15), the subset of $\mathbb{R}^{n\times n}$ containing rank k projection matrices serves as an equivalent representation of Grassmannians. Alternatively, if one opts to represent linear subspaces using the column space of rank k matrices from $\mathbb{R}^{n\times k}$, the matrix Z from above can be directly employed. This is because Z^TZ almost surely has full rank, and its inverse amounts

to a change of bases, thus representing the same point in the Grassmannians. This approach may lead to a more efficient method of generating random samples.

At this juncture, one might inquire about the likelihood that a randomly selected V from the uniform distribution on $\mathrm{Gr}(k,n)$ is in close proximity to a given $W \in \mathrm{Gr}(k,n)$ which could potentially be the optimal point for the objective function in 4.12. While various distance metrics exist for linear subspaces of $\mathbb{R}^n$, describing $\mathbb{P}(d(V,W)) < \delta$ for uniformly chosen V and W from $\mathrm{Gr}(k,n)$ with $\delta > 0$ requires employing the canonical metric on $\mathrm{Gr}(k,n)$ known as the *projection 2-norm*

$$d_{\mathrm{Gr}}(V,W) = ||P(V) - P(W)||_2.$$

Here, $P(V), P(W)$ are the respective projection matrices associated to V, W and $||-||_2$ denotes the matrix 2-norm. Furthermore, one can show the following equalities

$$||P(V)-P(W)||_2 = \sin(\theta_k) = \max\{\sup_{\substack{v\in V\\ ||v||\leq 1}} d(v,W), \sup_{\substack{w\in W\\ ||w||\leq 1}} d(w,V)\} = d_{Hd}(V,W) \quad (4.21)$$

where θ_k denotes the largest principal (or canonical) angle. The principal angles $\theta_1, \ldots, \theta_k \in [0, \pi/2)$ between V and W can be defined recursively by

$$\cos(\theta_j) = \max_{\substack{v\in V\\ ||v||=1}} \max_{\substack{w\in W\\ ||w||=1}} v^T w = v_j^T w_j$$

with $v^T v_i = 0$ and $w^T w_i = 0$ for $i = 1, \ldots, j-1$. See e.g. [BG73, Drm00, AEK06] for more details on these concepts. Put differently, the final equality in 4.21 implies that the objective function in 4.12 exhibits continuity with respect to the canonical distance on the Grassmannians. Before delving into the details of our random optimization algorithm, we investigate the probability of selecting a linear subspace in proximity to a specified subspace, drawing on findings from Edelman [AEK06].

Proposition 4.2.3. *Let $W \in Gr(k,n)$, $k < \frac{n+1}{2}$ and let V be randomly sampled from the uniform distribution on $Gr(k,n)$ endowed with the canonical metric. Then, for any $\varepsilon \in [0,1)$ we have*

$$\mathbb{P}(d_{Hd}(W,V) < \varepsilon) = \frac{\Gamma(\frac{p+1}{2})\Gamma(\frac{n-p+1}{2})}{\Gamma(\frac{1}{2})\Gamma(\frac{n+1}{2})} \varepsilon_2^{p(n-p)} {}_2F_1(\frac{n-p}{2}, \frac{1}{2}; \frac{n+1}{2}; \varepsilon^2 I_{p\times p})$$

where ${}_2F_1$ denotes the hypergeometric function with matrix argument.

Proof. This immediately follows from Eq. (4.21) and the probability distribution of the largest principal angle [AEK06, Theorem 1]. □

(n,k)	$\varepsilon = 0.1$	$\varepsilon = 0.2$
(2,1)	$s = 25$	$s = 12$
(3,1)	$s = 321$	$s = 80$
(4,1)	$s = 3786$	$s = 471$
(4,2)	$s = 48175$	$s = 2990$
(5,1)	$s = 42840$	$s = 2662$
(5,2)	$s = 6420174$	$s = 99487$

Table 4.1. Required approximate sample sizes for certainty of 80 percent

Remark 4.2.4. The rationale behind the constraint $p < \frac{n+1}{2}$ may raise questions for the reader. It can be demonstrated that all principal angles within the range of $(0, \pi/2)$ between two subspaces are congruent to the angles between their orthogonal complements, which also fall within the interval $(0, \pi/2)$. Additionally, the probability function presented above theoretically enables the determination of the required sample size to find a linear subspace that is ε close to a given subspace with a specified probability.

In practical calculations, however, we faced challenges in precisely computing the hypergeometric function with matrix arguments. Our efforts led to an approximation using the algorithm outlined in [KE06]. Consequently, the reported sample sizes are also approximations. We present these results primarily as a reference to the order of magnitude.

Given a specific ambient dimension n, subspace dimension k, and distance ε from a linear subspace $W \in \mathrm{Gr}(k,n)$, one can determine the sample size s. This ensures a desired level of confidence, for instance, $p = 0.8$, in selecting another linear subspace V from the random sample such that $d(V,W) < \varepsilon$. Detailed results are provided in Table 4.1.

4.2.3 Implementation

Now, we describe an implementable random algorithm that carries out a random optimization program. A possible pseudocode sequencing can be seen in Algorithm 6. The prespecified sample size s determines how many random samples we draw. This is done inside the loop starting at line 3. In lines 4–6 we generate a random element from the uniform distribution on $\mathrm{Gr}(k,n)$, according to Proposition 4.2.1. The element is represented as projection matrix, compare to Eq. (4.15). The column span of the matrix P, denoted $V =< P >$, is a k-dimensional linear subspace of $\mathbb{R}^n$ and elements in V are given by linear combinations of the column vectors of P. By

generating a δ sample of the unit interval k times and taking all possible combinations of length k we can generate a finite set of points $\mathbb{V}$ such that $d_{Hd}(\mathbb{V}, V) \leq \delta$. This is done in line 7. At line 8 we have two finite subsets of $\mathbb{R}^n$, i.e. V and $\mathbb{X}$, for which the Hausdorff distance can be computed directly (compare to 4.10). Note that $d_{Hd}(V, \mathbb{X}) \leq d_{Hd}(\mathbb{V}, \mathbb{X}) + d_{Hd}(V, \mathbb{V}) = d_{Hd}(\mathbb{V}, \mathbb{X}) + \delta$. We record $d_{Hd}(\mathbb{V}, \mathbb{X})$ if we found a new minimum amongst the set of all previous linear spaces. Finally, in line 14 for the output value F we can say the following: Let V_{opt} denote the arg min for our optimization problem and V be the element such that $d_{Hd}(\mathbb{X}, \mathbb{V}) = F$. For any $\varepsilon > 0$ we have that $\mid F - d_{Hd}(\mathbb{X}, V_{opt}) \mid \leq \varepsilon + \delta$ with a probability of $1 - (1 - \mathbb{P}(d_{Hd}(V_{opt}, V) < \varepsilon))^s$ (compare to Proposition 4.2.3).

Algorithm 6 RandHd$(\mathbb{X}, n, k, s, \delta)$

Input: Sample $\mathbb{X} \in \mathbf{Sam}_\star$, dimensions n, k, sample size s, density δ

Output: Minimized Hausdorff distance to $\mathbb{X}$ of random linear subspace

1: $i \leftarrow 0$
2: $F \leftarrow \infty$
3: **while** $i < s$ **do**:
4: Generate $\{z_j\}_{j=1,\ldots,nk}$ i.i.d. $\mathcal{N}(0,1)$ distributed
5: $M \leftarrow ((z_1, \ldots, z_n)^T, \ldots, (z_{n(k-1)}, \ldots, z_{nk})^T)$
6: $P \leftarrow M(MM^T)^{-1}M$
7: Generate finite δ-sample $\mathbb{V}$ from $V =< P >$
8: $f = d_{Hd}(\mathbb{V}, X)$
9: **if** $f < F$ **then**:
10: $F \leftarrow f$
11: i++
12: Return F

Putting together the results for our random optimization approach we obtain Algorithm 7 which described how to approximate the Φ_{Hd}^d. We record the level of approximation in Theorem 4.2.5.

Algorithm 7 $\Phi_{Hd}(\mathbb{X}, d)$

Input: Point cloud $\mathbb{X} \in \mathbf{Sam}_\star$, dimension d

optional parameters: sample size s, sample density δ

Output: $\Phi_{Hd}(\mathbb{X})$

$n \leftarrow \dim(\mathbb{X})$ ▷ Dimension of ambient space

$\Phi_{Hd}(\mathbb{X}) \leftarrow \text{RandHD}(\mathbb{X}, n, d, s, \delta)$

Theorem 4.2.5. *Let $X, \mathbb{X} \in \mathbf{Sam}_\star$ with $\mathbb{X}$ a δ-sample of X. Let γ be the sample accuracy for the linear subspaces. Then, for any $0 < \varepsilon \leq 1$, $s \in \mathbb{N}_+$ and $\Phi_{Hd}(\mathbb{X}, d)$ denoting the output of Algorithm 7 we have that*

$$\mid \Phi_{Hd}(\mathbb{X}, d) - \Phi_{Hd}^d(X) \mid \leq \varepsilon + \delta + \gamma$$

with probability $1 - (1 - \mathbb{P}(\theta < \arcsin \varepsilon))^s$.

Proof. Let $V_{opt} \in \mathrm{Gr}(k, n)$ be such that $d_{Hd}(X, V_{opt}) = \inf_{V \in \mathrm{Gr}(k,n)} d_{Hd}(X, V)$. The existence of the optimum is due to the fact that the Grassmannians are compact and the Hausdorff distance is continuous, see Eq. (4.21). For the given sample size s we generate random elements $\{W_1, \ldots, W_s\}$ of $\mathrm{Gr}(k, n)$ w.r.t. the uniform distribution by the procedure described in Algorithm 6 lines 4–7. The probability that there exists $j \in \{1, \ldots, s\}$ such that $d_{Gr}(W_j, V_{opt}) \leq \arcsin \varepsilon$ is given by

$$\mathbb{P}(\min_{i=1,\ldots,s} \{d_{Hd}(W_i, V_{opt})\} \leq \tilde{\theta})) = 1 - (1 - \mathbb{P}(\theta \leq \arcsin \varepsilon)^s.$$

In Algorithm 6 we record the minimum of $d_{Hd}(\mathbb{W}, \mathbb{X})$ where $\mathbb{W}$ denotes an γ-close sample of W. Therefore, with the stated probability, Algorithm 6 produces a sample $\mathbb{W}$ such that

$$d_{Hd}(\mathbb{W}, V_{opt}) \leq \gamma + \varepsilon,$$

by the triangle inequality and Eq. (4.21). Furthermore, by assumption $d_{Hd}(\mathbb{X}, X) \leq \delta$ and therefore, again by triangle inequality, we have

$$\mid \Phi_{Hd}(X) - \Phi_{Hd}(X) \mid \leq \varepsilon + \delta + \gamma$$

□

Remark 4.2.6. This outcome provides the degree of approximation for the value of $\Phi_{Hd}(\mathbb{X}, d)$, where $\mathbb{X}$ is derived from the magnification of a sample space $\mathbb{Y}$ at a specific point. We have demonstrated how to derive an estimator for the necessary sample size based on the chosen value of $1 - (1 - \mathbb{P}(\theta \leq \arcsin \varepsilon))^s$ in Remark 4.2.4. It is important to note that the probability of achieving the same level of accuracy at every point in $\mathbb{Y}$ is expressed as $(1-(1-\mathbb{P}(\theta \leq \arcsin \varepsilon))^s)^{|Y|}$. This elucidates the process of determining a required sample size for the function $\Phi_{Hd}(\mathbb{X}, d)$ to approximate $\Phi_{Hd}^d(\mathbb{X})$ to a chosen level of accuracy across all points in $\mathbb{Y}$.

4.3 Point cloud stratification

In this section, we delve into the practical intricacies of stratifying a given point cloud. We outline a straightforward routine for implementing a stratification of a

given sample space, resulting in a diagram comprised of links and strata that conveys the persistent stratified homotopy type. Although we have already introduced all the requisite tools, providing a comprehensive explanation of how to integrate them and establish a pipeline, as previously hinted at in Example 3.5.6, is still valuable.

Subsequently, we showcase the practical utility of the implementation. We focus on a specific algebraic surface to generate samples with controlled sample density, allowing us to assess the accuracy of the obtained stratifications relative to the quality of the sample. Additionally, we leverage this opportunity to compare the effectiveness of the various methods presented.

The modular nature of Φ-stratifications prompts an exploration of whether other well-established methods, such as PCA, could be seamlessly integrated into this context. This aspect is briefly touched upon towards the end of this section.

4.3.1 Implementation

In Corollary 3.5.9 we combined the convergence of $\mathcal{S}^r_{\Phi,u}$ with the continuity of $\mathcal{P}_v$ to see that $\mathcal{P}_v \circ \mathcal{S}^r_{\Phi,u}(\mathbb{X})$ converges to $\mathcal{P}_v(W)$ for $\frac{1}{r}d(X,\mathbb{X}) \to 0$ with X the underlying topological space of W. In total, the process always has four parameters to specify and these are dimension d, zoom parameter (or radius) r, threshold u and a tuple $v = (v_l, v_h)$ for the sub-level diagram. If we use our algorithm based on Hausdorff distances (see Algorithm 7) we would have another two parameters that specify the level of accuracy. We will first explain $\mathcal{S}^r_{\Phi,u}$ and then how to generate a model for $\mathcal{P}_v$ with their respective parameters.

Recall that $\mathcal{S}^r_{\Phi,u}$ is given by

$$\mathcal{S}^r_{\Phi,u}(\mathbb{X}) = \mathcal{F}_u \circ \Phi_* \circ \mathcal{M}^r(\mathbb{X})$$

which prompts a logical sequence for computing a stratification for a given sample $\mathbb{X} \in \mathbf{Sam}_\star$. Consider a sample space $\mathbb{X}$, potentially sampled from a stratified space of interest denoted as X. To proceed with determining the parameters for the routine, it is essential to formulate a hypothesis about X. In cases where no hypothesis is available, the process of stratifying a sample space becomes more explorative, and all parameters should be considered within their respective ranges.

The selection of a function Φ may hinge, in part, on computational preferences, influenced by the specific implementation of the method. For instance, while the computation of high-dimensional persistent homology becomes impractical rapidly, the size of the Grassmannians and thus the complexity of Algorithm 6 depends on the codimension of the spaces optimized over. However, the primary determinant

guiding the choice for Φ should be for which Φ the space X is Φ-stratified (refer to Definition 3.2.1). At present, with only $\Phi_{\mathcal{PL}}$ or Φ_{Hd} available, this decision simplifies to whether or not X is $\Phi_{\mathcal{PL}}$-stratified. If not, one must resort to using Φ_{Hd} (refer to Example 3.2.2).

The parameter d denotes the expected dimension of the top stratum of X. The zoom parameter or radius r is ideally set as low as possible. However, as previously discussed, practical constraints prevent setting it too low, in order to avoid $\mathcal{M}_x^r(\mathbb{X}) = x$ for all $x \in \mathbb{X}$.

Recall from Theorem 3.5.8 that for convergence $\mathcal{S}_{\Phi,u}^r(\mathbb{X}) \to X$, we require, among other conditions for X to be fulfilled, that $\frac{1}{r}d(\mathbb{X}, X) \to 0$. In other words, for smaller radii, we necessitate sufficiently good samples in that sense. In practice, it is prudent to conduct a sweep of r over a specified range.

The parameter u serves as the lower threshold for $\Phi(\mathcal{M}_x^r(\mathbb{X}), d)$, determining whether x is classified as singular. Theoretically, a good choice for u is given by $\sup\{\Phi(\mathrm{T}_x^{\mathrm{ex}}(X)) \mid x \in X_p\} < u < 1$. Notably, if information regarding the stratification defined by (X, X_p) were available, obviating the need for this routine, it would provide a clear interpretation of u. It is pertinent to note that u can be explored across its full range, i.e., $u \in [0, 1)$, as the thresholding occurs at the conclusion of the routine, and all computationally intensive tasks are completed at this point.

To formalize the algorithmic procedure, we provide pseudocode in Algorithm 8 outlining the steps to determine $\mathcal{S}_{\Phi,u}^r$. Commencing with $\mathcal{M}^r(\mathbb{X})$, which denotes the collection of local magnifications $\mathcal{M}_x^r(\mathbb{X})$ for all $x \in \mathbb{X}$, the algorithm proceeds as follows:

The construction of $\mathcal{S}_{\Phi,u}^r$ involves evaluating the function Φ on $\mathcal{M}_x^r(\mathbb{X})$ at each point $x \in \mathbb{X}$. Points where $\Phi(\mathcal{M}_x^r(\mathbb{X}))$ falls below the threshold u are classified as singular, corresponding to the singular stratum of the given point cloud (refer to lines 4–5). The magnification $\mathcal{M}_x^r(X)$ is determined by normalizing the neighborhood of radius r around the current point x, expressed as $\mathcal{M}_x^r(\mathbb{X}) = \frac{1}{r}(\mathbb{X} - x) \cap B_r$ (see line 3).

Upon iterating over all points in $\mathbb{X}$, the algorithm collects all singular points in the set $\mathbb{S} \subset \mathbb{X}$ at line 6, representing $(\mathcal{S}_{\Phi,u}^r(\mathbb{X}))_p$.
We can construct a diagram from $\mathcal{S}_{\Phi,u}^r\mathbb{X}$ which we will denote $\mathbb{D} = \{\mathbb{D}^{\leq v_l} \hookleftarrow \mathbb{D}_{v_h}^{v_l} \hookrightarrow \mathbb{D}_{\leq v_h}\}$ for brevity. This diagram is given by the level sets of the distance function to

Algorithm 8 $\mathcal{S}^r_{\Phi,u}(\mathbb{X}, 2)$

Input: Point cloud $\mathbb{X} \in \mathbf{Sam}$, $\Phi \in \{\Phi_{\mathcal{PL}}, \Phi_{Hd}\}$, dimension d, zoom r, threshold u
Output: A tuple $(\mathbb{X}, \mathbb{S})$ comprising $\mathcal{S}^r_{\Phi,u}(\mathbb{X})$

```
S ← ∅;
for x ∈ X do
    M^r_x(X) ← (1/r)((X ∩ B_r(x)) − x);
    if Φ^d(M^r_x(X)) ≥ u then                ▷ Compare to Algorithms 5 and 7
        S ← S ∪ {x}
```

the previously determined set $\mathbb{S}$ and that is

$$\mathbb{D}^{\geq v_l}(\mathcal{S}^r_{\Phi,u}\mathbb{X}) = \{x \in \mathbb{X} \mid \min(1, d(x, \mathbb{S})) \geq v_h\}$$
$$\mathbb{D}^{v_l}_{v_h}(\mathcal{S}^r_{\Phi,u}\mathbb{X}) = \{x \in \mathbb{X} \mid v_l \leq \min(1, d(x, \mathbb{S})) \leq v_h\}$$
$$\mathbb{D}_{\leq v_h}(\mathcal{S}^r_{\Phi,u}\mathbb{X}) = \{x \in \mathbb{X} \mid \min(1, d(x, \mathbb{S})) \leq v_h\}$$

whose component-wise thickenings can be modelled by a filtered simplicial complex and may serve as a geometric model for the associated persistent stratified homotopy type $\mathcal{P}_v \circ \mathcal{S}^r_{\Phi,u}\mathbb{X}$. As this process is straightforward, we omit the provision of pseudocode. It is essential to note that the distance to $\mathbb{S}$ must be rescaled to ensure a meaningful cutoff at parameter values $v = (v_l, v_h) \in (0, 1)^2$. Recall from Corollary 3.5.9 the assertion of the convergence of $\mathcal{P}_v \circ \mathcal{S}^r_{\Phi,u}\mathbb{X}$ to $\mathcal{P}_v(X)$ for X being associated to a compact Lojasiewicz-Whitney space, given that v is sufficiently small.

Additionally, the ordering on tuples (v_l, v_h) follows $(v^1_l, v^1_h) \leq (v^2_l, v^2_h) \iff v^1_l \leq v^2_l$ and $v^1_h \geq v^2_h$, while still satisfying $0 < v^i_l < v^i_h < 1$. Therefore, the aim is to choose this parameter as small as possible, contingent upon the quality of the given sample $\mathbb{X}$. This implies setting the values of v_l and v_h close to each other, with a small value for v_h. In the scenario where $\mathbb{X}$ represents a continuous space, one would ideally set $v_h = v_l$. However, this often results in $\mathbb{D}^{v_l}_{v_h}$ being empty for discrete spaces, such as a point cloud.

Moreover, considering our convergence results, which assume X to be associated to a compact Lojasiewicz-Whitney space, we seek a more persistently geometric link of the singular stratum at a certain non-zero distance from the singular stratum (refer to Example 2.1.33).

4.3.2 Method Evaluation

Throughout this research, our underlying motivation has been the development of a method capable of approximating the stratification of a space from a sufficiently close

sample and subsequently representing the outcome in a stable manner that captures pertinent homotopy theoretical information. In order to assess the practical utility of our methods in the context of Topological Data Analysis (TDA), we believe it is prudent to conduct demonstrations with artificial data where the correct results are known. This approach allows for a thorough examination of whether our methods function as intended before their direct application to real-world scenarios. Nevertheless, the efficacy of our methods when applied to real-world data is discussed in detail in Chapter 5.

Moreover, we revisit the previously highlighted flexibility in the selection of a function $\Phi : \mathbf{Sam}_\star \to [0, 1]$ for stratifying a point cloud. While there may be numerous suitable functions to choose from, we focused on two distinct functions (and mentioned some variants). The evaluation of both methods is based on local homology and Hausdorff distances, respectively, and we briefly comment on alternative methods that have been employed in practice for stratification learning.

Remark 4.3.1. To establish the convergence of $\mathcal{S}^r_{\Phi,u}$ for $\mathbb{X} \in \mathbf{Sam}_\star$, with Φ and X satisfying the conditions outlined in Theorem 3.5.8, it is essential to furnish a sequence $\mathbb{X}_i$ comprising δ_i-samples of X, where the density δ_i increases, signifying $\delta_i < \delta_j$ for $i < j$. Precision in controlling the increments is desirable to establish a discernible relationship between enhancements in sample density and improvements in the distance of the stratifications generated from $\mathbb{X}_i$. While generating samples of geometric objects, such as real algebraic varieties, is feasible, achieving a sample with specific density requirements requires careful consideration. Notably, the task has been accomplished for real algebraic varieties, as evidenced by the work in [DEHH18].

The task becomes substantially more straightforward if the space possesses a parametrization over the real numbers. An illustrative example is the sphere S^2, for which a dense sample of real intervals $[0, \pi)$ and $[0, 2\pi)$ suffices to generate an equally dense sample of S^2. Given that the sphere is a prototypical example of a smooth manifold, there is no necessity to undertake efforts to stratify a point cloud sampled from it. Instead, our focus shifts to spaces arising from intersections of regions parametrized over $\mathbb{R}^2$ and embedded in $\mathbb{R}^3$.

Example 4.3.2. We generated four samples from the standard Euclidean sphere $S = \{p = (x, y, z) \in \mathbb{R}^3 \mid \|p\| = 1\}$ and defined $X = S \cup (S + (0.5, 0.5, 0.5))$. In other words, the space under examination is a composite of non-disjoint Euclidean spheres with a common radius of $r = 1$. The spheres have centers at the origin $(0, 0, 0)$ and the point $(0.5, 0.5, 0.5)$. Consequently, the singular stratum in the natural stratification of X is characterized by $\Sigma = \{(x, y, z) \in S^2 \mid x + y + z = 0.75\}$, which is

topologically equivalent to S^1. The link of Σ is a space that can be represented as a disjoint union $S^1 \sqcup S^1$. The sample densities δ_i for $i \in \{1, 2, 3\}$ exhibit an ascending pattern, specifically $\delta_{i+1} = 2\delta_i$. Achieving this progression involves managing the sampling of intervals $[0, \pi)$ and $[0, 2\pi)$ that parameterize the standard 2-sphere in spherical coordinates. It is noteworthy that the scaling factor for the spherical coordinates always remains less than or equal to the radius, which is 1 in this context. Consequently, although points may come closer to each other after transformation, there is no stretching-out of points.

Let $\mathbb{X}_i$ denote the sample of density δ_i of X. Theorem 3.5.8 suggests that for increasing magnification parameter r and simultaneously improving sample quality, in the sense of increasing δ_i, we can expect a better approximation of the stratified space X, that is the distance of $\mathcal{S}^r_{\Phi,u}$ and (X, Σ) as elements of $\mathbf{Sam}_P$ decreases. Note that we abuse notation here and refer to X as a stratified space and the underlying space itself in hope to abbreviate the notation at this point. We want to demonstrate the convergence of the approximate stratifications for two choices of $\Phi \in \{\Phi_{PLH}, \Phi_{Hd}\}$. For the other parameters we chose $r \in \{0.1, 0.3, 0.4\}$ and $u = 0.7$ in all cases. A visual representation of the resulting function values of $\Phi_{\mathcal{PL}}(\mathbb{X}_i, 2)$ at every point of $\mathbb{X}_i$, i.e. $\Phi_* \circ \mathcal{M}(\mathbb{X}_i)$, can be seen in Figs. 4.1 to 4.8 for $r \in \{0.3, 0.4\}$. In Fig. 4.9 we can see $\Phi_{\mathcal{PL}*} \circ \mathcal{M}^{0.1}(\mathbb{X}_1, 2)$. Note that $\Phi_* \circ \mathcal{M}(\mathbb{X}_3)$ can be computed with Algorithm 8 by returning all the calculated function values instead of the classification. Furthermore, we chose $v = (v_l = 0.48, v_h = 0.52)$ to generate a diagram from $\mathcal{S}^r_{\Phi,u}(\mathbb{X}_i) = \mathcal{F}_u \circ \Phi_* \circ \mathcal{M}^r(\mathbb{X}_i)$ for $\mathcal{M}^{0.4}(\mathbb{X}_2)$, $\mathcal{M}^{0.4}(\mathbb{X}_1)$, $\mathcal{M}^{0.3}(\mathbb{X}_1)$ and $\mathcal{M}^{0.1}(\mathbb{X}_1)$ with $\Phi_{\mathcal{PL}}$. Here the distance to Σ has to be rescaled for $(v_l, v_h) \in (0, 1)^2$ to make sense. As this example is intended to demonstrate the convergence results the diagram parameters were chosen for the persistent homology of the the parts of the corresponding diagram $\mathcal{P}_v(X)$ to clearly reflect the topology of the respective strata and the link. The results have been visualized by marking the respective geometric link, i.e. $\mathbb{D}^{v_l}_{v_h}$, in yellow color in Figs. 4.11 to 4.14. With the case of $\mathcal{M}^{0.4}\mathbb{X}_3$ we actually want to demonstrate the case where our methods were unable to retrieve a stratification of the sample because the radius was set to low for the rather coarse sample. We only exhibit the diagrams and further investigation for $\Phi_{\mathcal{PL}}$ as the results are almost identical to the results with Φ_{Hd}. However, the strongly stratified spaces shown in Figs. 4.3 to 4.8 actually show a very distinct behaviour of the Φ functions. The function $\Phi_{\mathcal{PL}}$ attains the highest values very close to the actual singular stratum of X but Φ_{Hd} exhibits highest values near the boundary of an r-neighborhood of Σ. Finally, we measured the distance of the singular part of the stratified samples, i.e. $\mathcal{F}_{0.7} \circ \Phi_{\mathcal{PL}*} \circ \mathcal{M}^{0.4}(\mathbb{X}_3)$, to the space Σ, consider Table 4.2. The main observation here is that the distance to Σ, which

Figure 4.1. $\Phi_{\mathcal{P}\mathrm{L}*} \circ \mathcal{M}^{0.4}(\mathbb{X}_3, 2)$

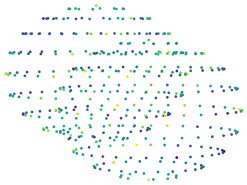

Figure 4.2. $\Phi_{Hd*} \circ \mathcal{M}^{0.4}(\mathbb{X}_3, 2)$

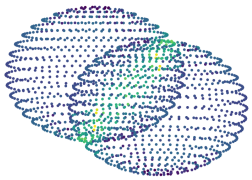

Figure 4.3. $\Phi_{\mathcal{P}\mathrm{L}*} \circ \mathcal{M}^{0.4}(\mathbb{X}_2, 2)$

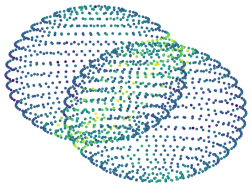

Figure 4.4. $\Phi_{Hd*} \circ \mathcal{M}^{0.4}(\mathbb{X}_2, 2)$

is the only interesting part as the distance of X and $\mathbb{X}_i$ equals δ_i, decreases almost linearly with the radius but only if the quality of the sample allows. For $r = 0.1$ and $\mathbb{X}_2$ we find the distance to be rather big because in that case the Algorithm 8 malfunctioned for the radius was to small for the provided sample density δ_2. We also calculated the persistent homology of the approximated link, i.e. $\mathbb{D}_{0.52}^{0.48}(\mathcal{S}^r_{\Phi_{\mathcal{P}\mathrm{L}\bullet,0.7}}(\mathbb{X}_i, 2))$ marked yellow, respectively, consider Table 4.3. The table contains the length of the four most persistent homology cycles, as we would expect four persistent 1-cycles. For reference the expected persistence of the four most persistent 1-cycles for the geometric link determined by v in X is approximately $(0.65, 0.65, 0.32, 0.32)$. These results are intended to illustrate how to interpret Theorem 3.5.8 and Theorem 3.5.8 in practice. In this example one can see that stratifications can be reconstructed for the specified parameters given the sample is close enough. We have also seen that the results for the persistent homology of the approximated links approach the expectation for the persistent stratified homotopy type of X represented by the specified diagrams.

Remark 4.3.3. The motivation behind the introduction of a novel tangential approximation method and the discussion on computing local persistent homology might prompt readers to question the necessity, considering the existence of widely accepted methods for such tasks in practical applications. Notably, alternatives such as localized versions of principal component analysis (PCA) have been employed for dimen-

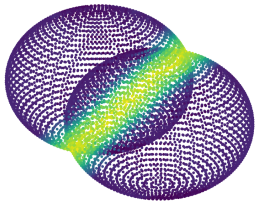

Figure 4.5. $\Phi_{\mathcal{PL}*} \circ \mathcal{M}^{0.4}(\mathbb{X}_1, 2)$

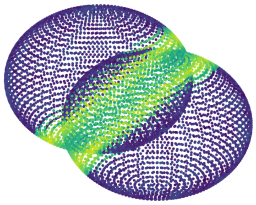

Figure 4.6. $\Phi_{Hd*} \circ \mathcal{M}^{0.4}(\mathbb{X}_1, 2)$

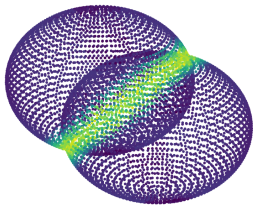

Figure 4.7. $\Phi_{\mathcal{PL}*} \circ \mathcal{M}^{0.3}(\mathbb{X}_1, 2)$

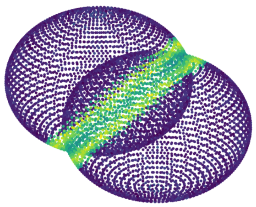

Figure 4.8. $\Phi_{Hd*} \circ \mathcal{M}^{0.3}(\mathbb{X}_1, 2)$

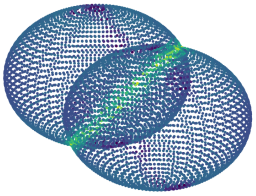

Figure 4.9. $\Phi_{\mathcal{PL}*} \circ \mathcal{M}^{0.1}(\mathbb{X}_1, 2)$

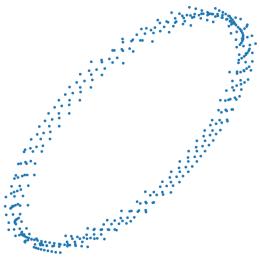

Figure 4.10. Approximated singular stratum from $\Phi^2_{\mathcal{PL}*} \circ \mathcal{M}^{0.1}(\mathbb{X}_1, 2)$

(i, r)	$d_{Hd}(\Sigma, \mathcal{F}_{0.7} \circ \Phi^2_{\mathcal{PL}*} \circ \mathcal{M}^r(\mathbb{X}_i))$
$(2, 0.4)$	0.273
$(2, 0.3)$	0.198
$(2, 0.1)$	1.049
$(1, 0.4)$	0.272
$(1, 0.3)$	0.196
$(1, 0.1)$	0.065

Table 4.2. Hausdorff distances of Σ to the approximated singular strata

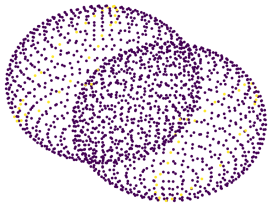

Figure 4.11. $\mathbb{D}_{0.52}^{0.48}(\mathcal{S}^{0.4}_{\Phi^2_{\mathcal{PL},0.7}}(\mathbb{X}_2,2))$ marked yellow

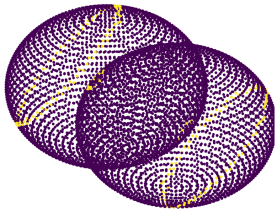

Figure 4.12. $\mathbb{D}_{0.52}^{0.48}(\mathcal{S}^{0.4}_{\Phi^2_{\mathcal{PL},0.7}}(\mathbb{X}_1,2))$ marked yellow

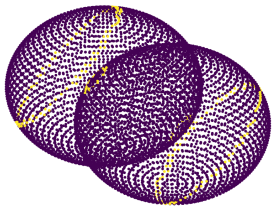

Figure 4.13. $\mathbb{D}_{0.52}^{0.48}(\mathcal{S}^{0.3}_{\Phi^2_{\mathcal{PL},0.7}}(\mathbb{X}_1,2))$ marked yellow

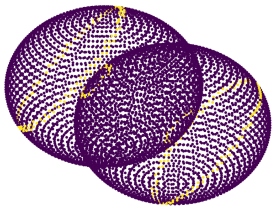

Figure 4.14. $\mathbb{D}_{0.52}^{0.48}(\mathcal{S}^{0.1}_{\Phi^2_{\mathcal{PL},0.7}}(\mathbb{X}_1,2))$ marked yellow

(i,r)	$\mathcal{P}H_1(\mathbb{D}_{0.52}^{0.48}(\mathcal{S}^{r}_{\Phi^2_{\mathcal{PL},0.7}}(\mathbb{X}_i)))$
$(2,0.4)$	$(0.559,0.559,0.071,0.071)$
$(1,0.4)$	$(0.707,0.707,0.077,0.077)$
$(1,0.3)$	$(0.662,0.662,0.179,0.179)$
$(1,0.1)$	$(0.646,0.646,0.299,0.299)$

Table 4.3. Persistence of the four most persistent homology cycles of the approximated links

sion estimation and stratification learning, as illustrated in [LON22]. To conclude this chapter, we offer insights into the distinctions between our introduced methods and other established approaches.

A fundamental point to underscore is that the methods discussed here, both in this context and in our previous work [MW22], come with theoretical assurances for their efficacy in stratification learning (Theorem 3.5.8 and Corollary 3.5.9). Additionally, we observed that the method for approximating stratification from a point cloud is adaptable based on the selection of a function Φ, with a key property being its continuity concerning the Hausdorff distance.

A crucial differentiation between the aforementioned methods and localized versions of principal component analysis, as employed in stratification learning [LON22], lies in the optimization problem inherent in PCA. PCA yields a linear subspace that minimizes the standard Euclidean norm distance, a metric that is generally not continuous with respect to the Hausdorff metric. Moreover, while PCA is a well-established method for approximating tangential spaces in cases where the space of interest is sufficiently smooth, it is important to emphasize that our theoretical framework does not support the claim that principal component analysis can effectively approximate stratifications.

Regarding the use of local persistent homology for stratification learning, as demonstrated in works such as [STHN20, Nan20, BWM12, BCSE$^+$07], we focused on an implementation of a function that is based on, and approximates, a method capable of (provably) approximating the stratification of a two-strata space from a sufficiently accurate sample.

Chapter 5

Applications

This chapter will explore the application of topological methods on real-world data. Although each section exhibits a different dataset, our main field of applications lay on images. After a thorough assessment of previous work on the different types of data we continue by investigating possible stratified structures. In every case we detail the preprocessing, the methods used for the analysis and give an interpretation of our results.

5.1 Image patch data

The study of image patch spaces involves representing an image as collection of smaller, sub-region patches that can be thought of as vectors in a high-dimensional space, where each dimension corresponds to a pixel intensity value. The space of all possible image patches is commonly referred to as *image patch space*.

The idea of distance between patches is crucial to understanding picture patch space mathematics. The similarity or dissimilarity between patches can be measured using a variety of distance metrics. The standard Euclidean norm, for instance, can be used to quantify the pixel-level variations between patches. Other metrics such as, e.g., the correlation coefficient or the structural similarity index measure (SSIM) may also be considered if other relationships are to be inferred.

We will concentrate on the topology of image patch spaces here. To determine the topological characteristics of image patches and their connections to one another, topological tools like persistent homology can be utilized. For instance, persistent homology can be used to detect components, tunnels and voids in an image patch space, which can be helpful for processes like image segmentation or classification [DMV17, AGV+18, VNG20, ESM+21].

Other mathematical techniques, such as stratification learning (or manifold learn-

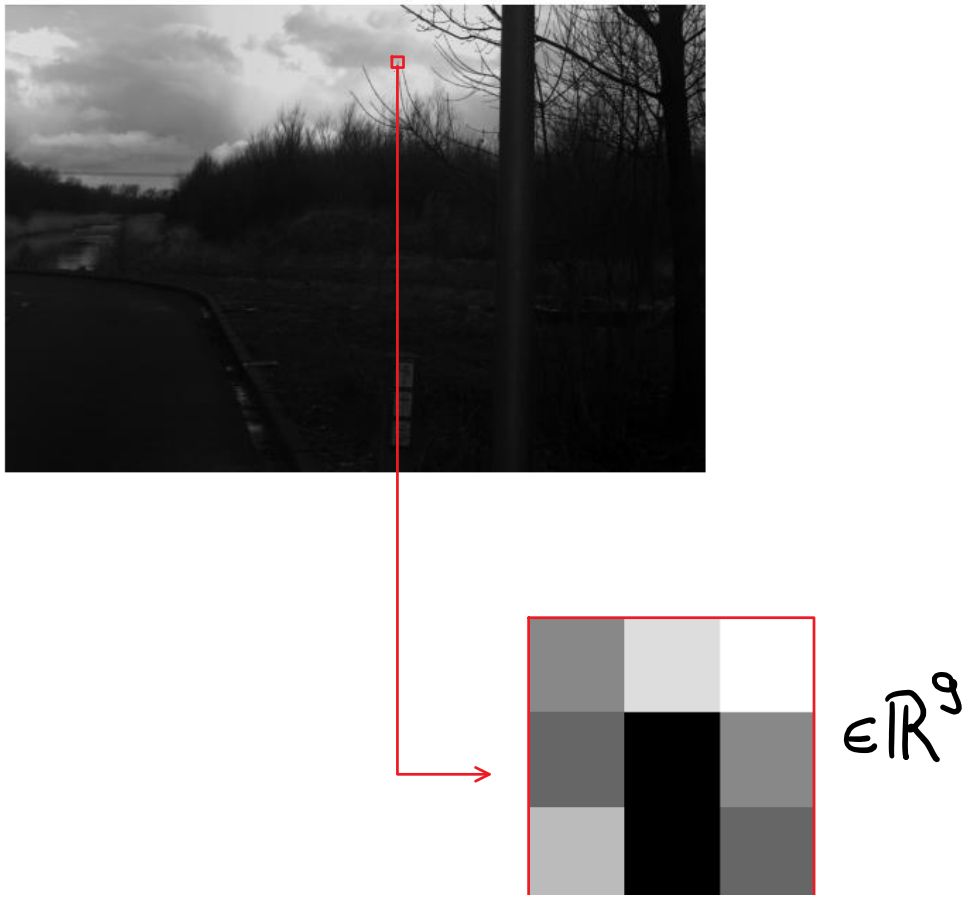

Figure 5.1. Enlarged image patch in a natural grayscale image

ing for that matter) and clustering, can also be applied to image patch spaces. While clustering techniques can be used to group related patches together, stratification learning algorithms seek to reveal the underlying low-dimensional structures or a stratification of the patch space [LPM03, CIDSZ08, Xia16, CG20]. We will base our stratification learning techniques on the methods described in the previous chapters. This will also serve as a demonstration of the utility of $\Phi_{\mathcal{PL}}$ in this application context.

To generate some intuition about image patches we included Fig. 5.1. To make things precise will state what we will refer to as an image. Here, an *image* is represented by a map

$$f : \Omega = \{0, \dots, n\} \times \{0, \dots, m\} \subset \mathbb{N}^2 \to \mathcal{T} \subset \mathbb{R}^c$$

where n, m are the height and width of the image respectively, $c \in \mathbb{N}_{>0}$ is the number of color channels and $\mathcal{T} \subset \mathbb{R}^c$ is a non-emtpy set of color vectors. For an image as such, which we will often simply refer to as f, the associated *image patch space of*

$p \times q$ *patches*, denoted $P_{(p,q)}$, can be represented by the map

$$F_{(p,q)} : \{0, \dots, n-p\} \times \{0, \dots, m-q\} \to P_{(p,q)} \subset \mathbb{R}^{pqc}$$
$$x \mapsto (f(x+t))_{t \in [0,p] \times [0,q]},$$

with $0 < p < (n-p)$ and $0 < q < (m-q)$ natural numbers. The topology of P critically depends on the metric on the image patches and we will discuss different choices in the subsequent sections to a small extend. We start with one of the first approaches towards the topological analysis of image patch spaces and that is the analysis of high contrast image patch spaces.

5.1.1 High contrast image patch spaces

Background and other work

Earlier studies examining the topological structures within image patch spaces, particularly those we will build upon, are found in [LPM03, SC04, CIDSZ08]. While these sources primarily focused on natural grayscale images, similar investigations into colored images are documented elsewhere [Xia16]. Specifically, the datasets utilized in [LPM03, SC04, CIDSZ08] were derived from the van Hateren image collection [vHvdS98], comprising over 4000 large calibrated grayscale still images capturing outdoor scenes around the city of Groningen, taken with a Kodak DCS420 camera.

In these studies, patches extracted from the image collection were of size 3×3 and interpreted as elements in 9-dimensional Euclidean space $\mathbb{R}^9$. A sequence of preprocessing steps was implemented to sparsify the dataset through random subsampling and position it on a geometrically manageable object in Euclidean space, namely the sphere. This transformation facilitated statistical analyses based on volume calculations, as described in [LPM03], and computations of distances to another object situated on a sphere, as in [CIDSZ08].

The primary topological finding in [LPM03] is the concentration of most patches around blurred step-edge patches, resembling the topology of an annulus, often referred to as the "main circle." Filtering the dataset by density allowed for the targeting and further investigation of dense subsets, potentially revealing more complex spatial structures. While [LPM03] primarily employed statistical methods, [CIDSZ08] utilized persistent homology. They confirmed that blurred step-edge patches constituted the densest subset in the extracted image patch space and proposed a 3-circle model for a slightly higher percentage of densest patches.

To refine and build upon the analyses conducted in [LPM03, CIDSZ08], we aim to provide a more precise account and transition into our own investigations of image

patch spaces.

Dataset

We will briefly describe the approach by [LPM03] and [CIDSZ08] which focused on high contrast patches extracted from van Hateren's natural image dataset [vHvdS98].

The dataset was systematically generated through the random extraction of a predetermined number of patches, denoted as N, from each patch space characterized by a width w associated with an image $f : \Omega \to \mathcal{T}$ selected from the van Hateren database. Subsequently, the collection of all patch spaces into a unified, comprehensive patch space denoted as $\mathcal{A}$ was performed as per the methodology proposed in [LPM03, CIDSZ08]. However, it is pertinent to mention that we also opted to incorporate instances where the patch space originates from a single image in the dataset collection. The initial step in dataset size reduction involves the selection of N patches from each chosen image.

For the purpose of establishing a metric space, the patches are construed as real w^2-vectors arranged column-wise, representing elements of $\mathbb{R}^{w^2}$ endowed with the standard Euclidean metric. To mitigate variations in luminance among diverse images within the database, a logarithmic transformation is applied to the intensity values of the patches. It is imperative to acknowledge that this adjustment is not applicable to individual images, notwithstanding the logarithmic transformation applied. Even in such cases, where a logarithmic transformation is employed, darker intensity values are elevated, rendering features in the darker regions more perceptible to the human eye. Simultaneously, brighter intensity levels are attenuated to lower values due to the logarithmic transformation.

In summary, we have:

$$\mathcal{A} \to P \subset \mathbb{R}^{w^2}$$
$$a \mapsto x = (\log 1 + a_{11}, \log 1 + a_{21}, \ldots \log 1 + a_{w1}, \log 1 + a_{w2}, ..., \log 1 + a_{ww}).$$

Preprocessing

The so-called *D-norm* given by $||x||_D = \sqrt{\sum_{i\sim j}(x_i - x_j)^2}$ with $i \sim j$ if the pixels are adjacent in the patch associated to x, induces an ordering on $\mathcal{P}'$. The D-norm can be

computed by $||x||_D = \sqrt{x^T D x}$ with a matrix D given by

$$D = \begin{pmatrix} 2 & -1 & 0 & -1 & 0 & 0 & 0 & 0 & 0 \\ -1 & 3 & -1 & 0 & -1 & 0 & 0 & 0 & 0 \\ 0 & -1 & 2 & 0 & 0 & -1 & 0 & 0 & 0 \\ -1 & 0 & 0 & 3 & -1 & 0 & -1 & 0 & 0 \\ 0 & -1 & 0 & -1 & 4 & -1 & 0 & -1 & 0 \\ 0 & 0 & -1 & 0 & -1 & 3 & 0 & 0 & -1 \\ 0 & 0 & 0 & -1 & 0 & 0 & 2 & -1 & 0 \\ 0 & 0 & 0 & 0 & -1 & 0 & -1 & 3 & -1 \\ 0 & 0 & 0 & 0 & 0 & -1 & 0 & -1 & 2 \end{pmatrix}.$$

The D-norm is a measure for contrast and can therefore be used to filter out low contrast patches, e.g. by only keeping the T elements with highest D-norm. In order to place P' on the surface of an $w^2 - 2$-dimensional ellipsoid the dataset is centralized and normalized with the D-norm

$$P' \to \tilde{P} \tag{5.1}$$

$$x \mapsto \frac{x - \frac{1}{w^2}\sum_{i=1} x_i}{||x - \frac{1}{w^2}\sum_{i=1} x_i||_D}. \tag{5.2}$$

It is crucial to note that the D-norm necessitates a non-zero value, a requirement assured through the meticulous selection of high-contrast patches, ensuring that T does not fall below a certain threshold. As of the current stage, the transformative impact on the topology of patch spaces remains inconsequential, as alterations in inter-point distances are reversible, a consideration integral to the computation of persistent homology.

In order to land on an actual $S^{w^2-2} \subset \mathbb{R}^{w^2-1}$, the authors in [LPM03, CIDSZ08] undertake a coordinate transformation with respect to the 2-dimensional Discrete Cosine Transform (DCT) basis inherent to a 3×3 image patch. This calculated transformation effectively diagonalizes the matrix D". Within $\mathbb{R}^9$, this basis characterized by

the following eight vectors:

$$
\begin{aligned}
e_1 &= \frac{1}{\sqrt{6}}\begin{pmatrix}1 & 0 & 1 & 1 & 0 & 1 & 1 & 0 & 1\end{pmatrix}^T,\\
e_2 &= \frac{1}{\sqrt{6}}\begin{pmatrix}1 & 1 & 1 & 0 & 0 & 0 & -1 & -1 & 1\end{pmatrix}^T,\\
e_3 &= \frac{1}{\sqrt{54}}\begin{pmatrix}1 & -2 & 1 & 1 & -2 & 1 & 1 & -2 & 1\end{pmatrix}^T,\\
e_4 &= \frac{1}{\sqrt{54}}\begin{pmatrix}1 & 1 & 1 & -2 & -2 & -2 & 1 & 1 & 1\end{pmatrix}^T,\\
e_5 &= \frac{1}{\sqrt{8}}\begin{pmatrix}1 & 0 & 1 & 0 & 0 & 0 & -1 & 0 & 1\end{pmatrix}^T,\\
e_6 &= \frac{1}{\sqrt{48}}\begin{pmatrix}1 & 0 & 1 & -2 & 0 & 2 & 1 & 0 & 1\end{pmatrix}^T,\\
e_7 &= \frac{1}{\sqrt{48}}\begin{pmatrix}1 & -2 & 1 & 0 & 0 & 0 & -1 & 2 & 1\end{pmatrix}^T,\\
e_8 &= \frac{1}{\sqrt{216}}\begin{pmatrix}1 & -2 & 1 & -2 & 4 & 2 & 1 & -2 & 1\end{pmatrix}^T.
\end{aligned}
$$

Density Filtration

An essential step in the investigation of patch spaces from natural images is to not look at the whole dataset as previously described in the above but at percentages of the densest points in $\tilde{P}$. This is achieved by computing the k-nearest neighbors density at every point $x \in \tilde{P}$, denoted $\rho_k(x)$, which is given by the maximum distance between x and the k closest points to x in $\tilde{P}$. These density values induce an ordering on $\tilde{P}$ by $x \leq y$ iff $\rho_k(x) \leq \rho_k(y)$. For a chosen value k and once $\tilde{P}$ is ordered in descending order by the k-nearest density, we can select a percentage $p \in [0, 1]$ of highest density points for further investigation. We will refer to such a subset as $X(k, p) \subseteq \tilde{P}$.

Analysis

To illustrate the importance of this filtration step let $k = 15$ and $p = 0.3$, i.e. the top 30 percent of densest points. See Fig. 5.2 for a visualization of a point cloud, denoted $\mathbb{X}$, embedded into $\mathbb{R}^2$ via PCA that was generated from a single image (image number 1111 in the van Hateren collection). Note that this considers only a single image not a whole collection of images but the preprocession was essentially the same. In Fig. 5.3, we showcase the effect of not selecting the denser subsets by setting $p = 1$, which results in a larger dataset that also has less resemblance of the S^1 then the smaller percentage.

Another point cloud, denoted $\mathbb{X}'$, was taken from a single image from the van Hateren collection (number 120) with preprocession parameters $p = 0.3$, $k = 15$. An

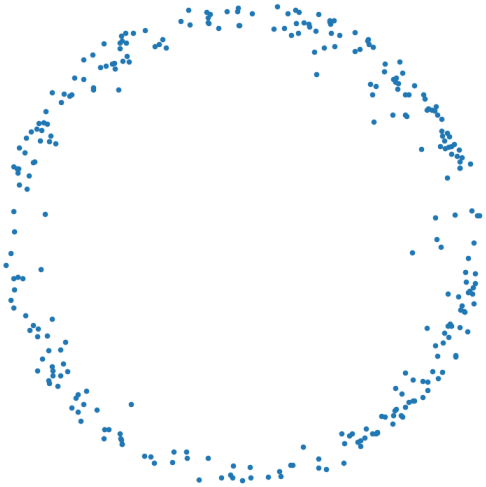

Figure 5.2. Point cloud from van Hateren image number 1111 $p = 0.3$

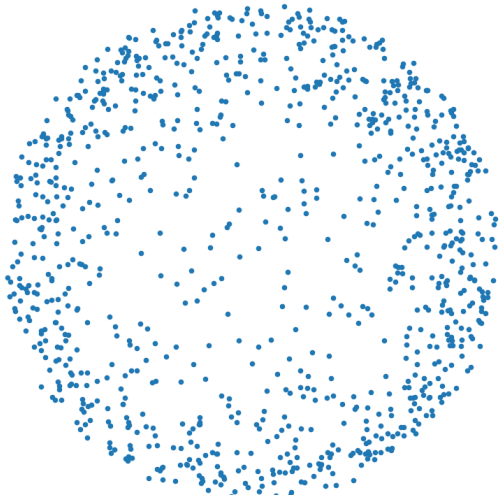

Figure 5.3. Point cloud from van Hateren image number 1111 $p = 1.0$

interesting point here is the difference in the topology. While in the case of $\mathbb{X}$ the patch space very much resembled the topology of S^1, the main circle, in case of $\mathbb{X}'$ we can record three persistent homology 1-cycles (see Fig. 5.4.

Apart from that, we generated a dataset from all images in the van Hateren collection exactly as described previously. To further investigate the stratified structure in a dense subset of the whole patch space we chose the parameters $k = 15$ and $p = 0.1$, for reference denote this patch space as $\mathbb{X}$. We applied Algorithm 5 in order to compute $\Phi_{\mathcal{PL}}(\mathbb{X}, 1)$ of the point cloud, depicted in Fig. 5.5. Furthermore, we selected radius $r = 0.5$, $u = 0.3$ as threshold in order to obtain a stratified sample by $\mathcal{S}^r_{\Phi,u}$. With this procedure we were able to localize four connected regions of significant lower $\Phi_{\mathcal{PL}}(\mathbb{X}, 1)$ value, thus resulting in four connected singular regions. This is in line with the model of having three nested circles where two of the three intersect the third only in two distinct points. The identified singular regions may correspond to theses intersection points. A picture of $\mathcal{S}^r_{\Phi,u}(\mathbb{X}, 1)$ with its singular stratum indicated by yellow color is in included with Fig. 5.6. To further test the hypothesis whether or not the patch space resembles the topology of a circle to which two distinct circles are attached at two distinct points each, we proceeded by investigating the persistent homology of the parts of the corresponding stratification diagram. That is, we selected $(v_l, v_h) = (0.3, 0.4)$ to generate a stratification diagram from the stratified sample $\mathcal{S}^r_{\Phi,u}(\mathbb{X}, 1)$. The results for persistent homology up to dimension 1 of the three parts of the stratification diagram can be seen in Fig. 5.7 for $\mathbb{D}^{v_l}\mathcal{S}^r_{\Phi,u}(\mathbb{X}, 1)$ (the regular part), in Fig. 5.8 for $\mathbb{D}^{v_l}_{v_h}\mathcal{S}^r_{\Phi,u}(\mathbb{X}, 1)$ (the link part) and in Fig. 5.9 for $\mathbb{D}_{v_h}\mathcal{S}^r_{\Phi,u}(\mathbb{X}, 1)$ (the singular part). We only calculated up to dimension 1 as we want to see if the hypothesis can be further

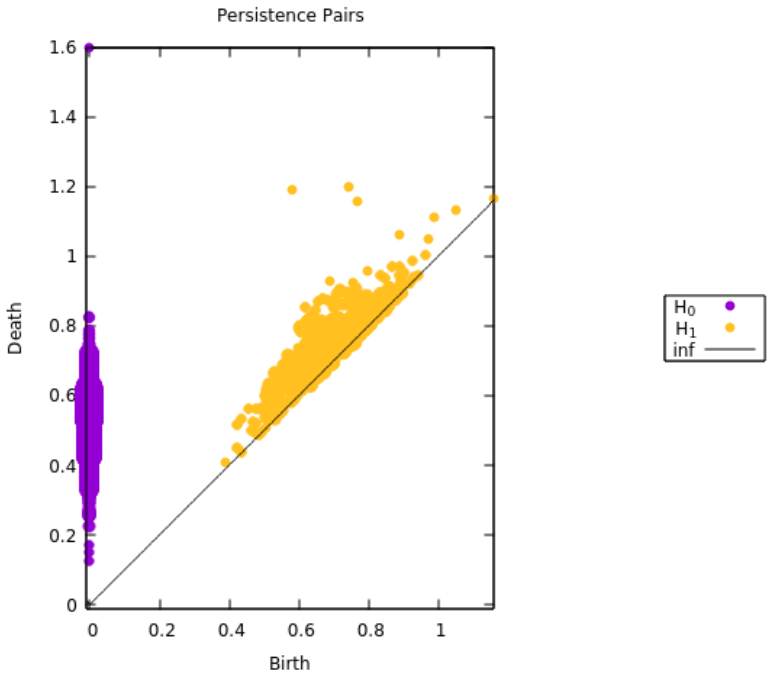

Figure 5.4. Vietoris-Rips persistence diagram of $\mathbb{X}'$

consolidated which would mean that we find only non-trivial 0 homology cycles. If we only look at the most persistent features, we record eight persistent 0-cycles for the regular part, 16 persistent 0 cycles in the link and four persistent 0-cycles in the singular part.

Interpretation

Our analysis had mainly two objectives. First, we wanted to reconsider the analysis originally done in [SC04, CIDSZ08] and take a closer look at the data preprocessing as well as the effect on the topology of patch spaces. The transformations up to the density filtration have no significant effect on the topology as it is more of a rescaling of the distances between the patches and therefore the impact on persistent homology can be adjusted for. The data was filtered by density estimations as a means to extract subsets of an image patch space for further investigation. This followed the idea that the space of most frequent patches in a collection of natural images possess non-trivial topological features which can then be measured by persistent homology. Previous publications reported that most of the patches concentrate around blurred step edges [LPM03] and suggested a so-called three circle model [SC04]. We also considered the topology of individual images with a simple comparison of the images with noticeably different homological features as demonstrated by comparing comparing image number 1111 and 120.

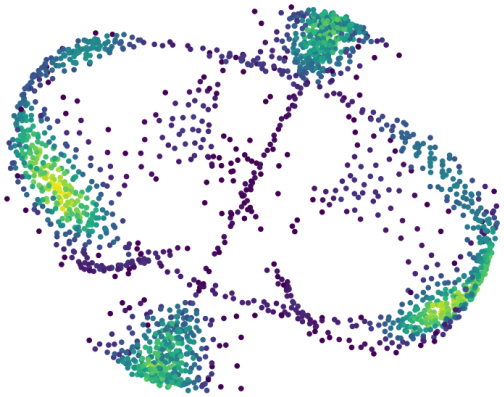

Figure 5.5. Visualization via PCA of $\mathbb{X}$ with coloring by results of $\Phi_{\mathcal{PL}}(\mathbb{X}, 1)$

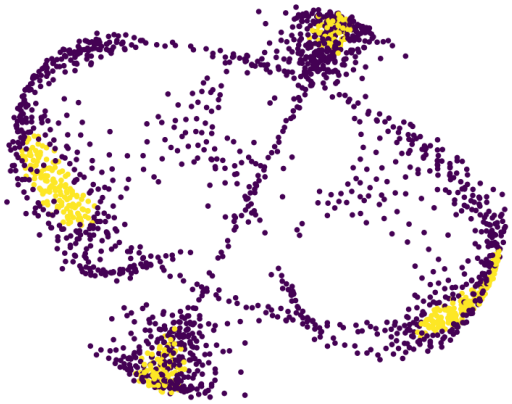

Figure 5.6. Stratification indicated by color of the PCA embedding of $\mathbb{X}$

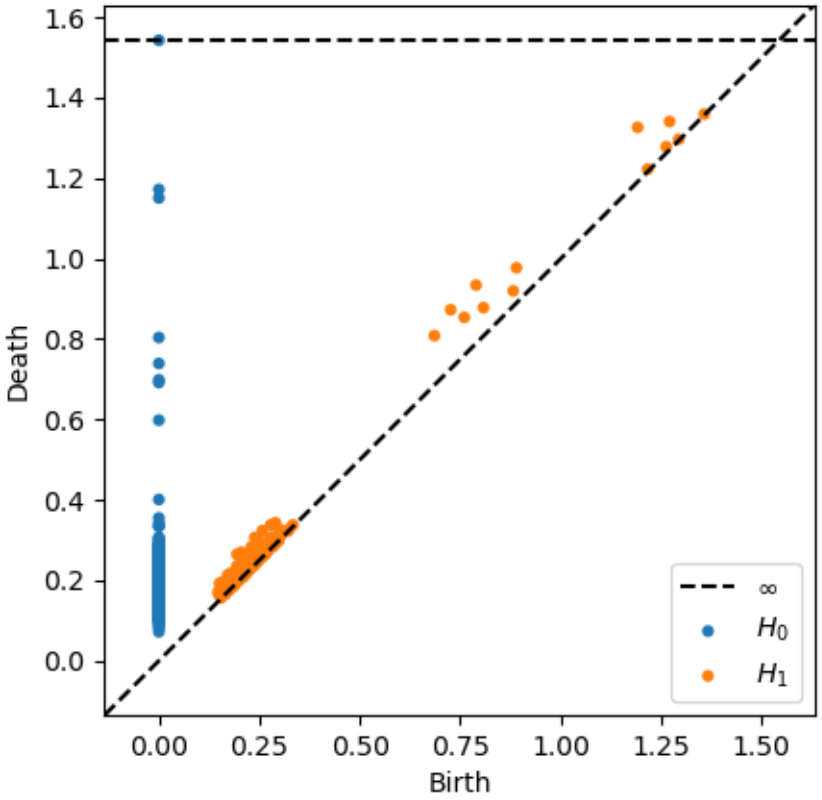

Figure 5.7. Persistent homology of $\mathbb{D}^{v_l}\mathcal{S}^r_{\Phi,u}(\mathbb{X}, 1)$

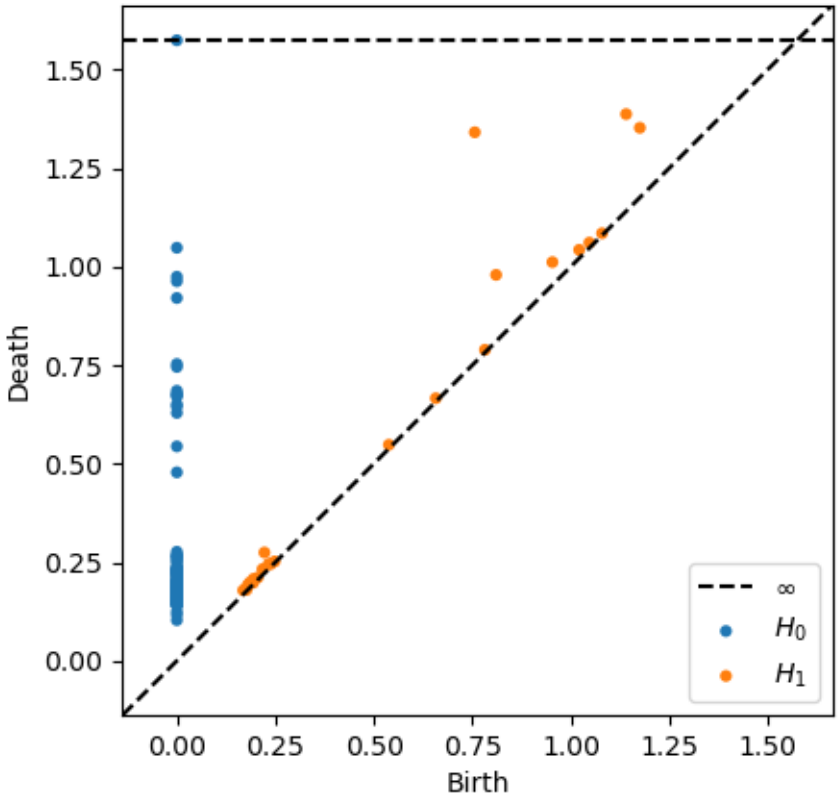

Figure 5.8. Persistent homology of $\mathbb{D}_{v_h}^{v_l}\mathcal{S}_{\Phi,u}^{r}(\mathbb{X},1)$

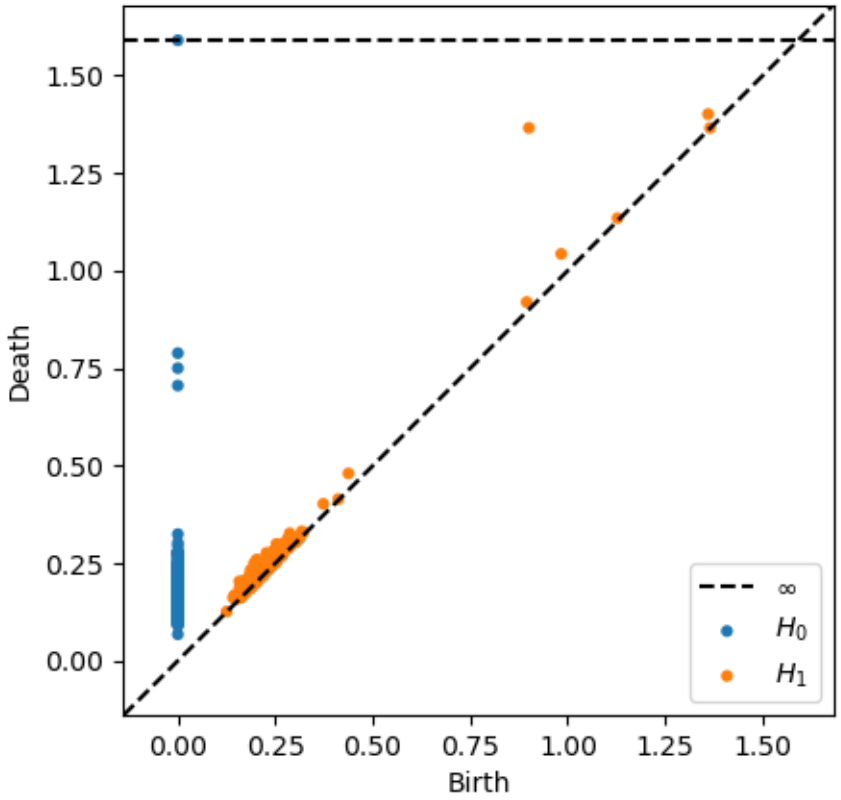

Figure 5.9. Persistent homology of $\mathbb{D}_{v_h}\mathcal{S}^r_{\Phi,u}(\mathbb{X}, 1)$

We then investigated the dataset constructed as in [CIDSZ08] with our own methods to test the circle model which would be a stratified space. We were able to uncover a stratification for the sample space given by image patches and then generate a stratification diagram. We then investigated the persistent homology of the parts of the stratification diagram. All of our results strongly support the hypothesis of the threes circle model by [SC04, CIDSZ08].

5.1.2 Cluster Images

In contrast to the previously examined image patch spaces, we now turn our attention to individual image patch spaces. Specifically, we aim to explore the topology associated with a single image, contemplating the question of what form the topology of an individual image might assume. The task of establishing a comprehensive and meaningful topological model for every conceivable type of image appears daunting without the incorporation of statistical considerations. Consequently, we direct our focus towards specific features within images, as elucidated at the outset of this section.

In a broad sense, our approach involves regarding an image as an amalgamation of colors and the transitions between them, adopting a reductionist perspective. A category of images that readily lends itself to this descriptive framework is what we will term as "cluster images." To foster initial understanding and to foreshadow the eventual mathematical consolidation, we shall embark upon an example to provide intuition regarding the model toward which this discussion converges.

Example 5.1.1. As a simple academic example consider the grayscale image in Fig. 5.10. It contains four distinct grayscale colors and clearly defined and uniformly colored regions. The transition between different colors is not immediate but smoothed-out, i.e. a patch that lies between a black and a white region also contains colors that lie in between both colors on the grayscale. A few of the patches present in the image are shown in Fig. 5.11. An embedding of the corresponding image patch space of width 3 equipped with the Euclidean distance into $\mathbb{R}^3$ via PCA can be seen in Fig. 5.12. We colored the points according to the value of the function

$$\Phi_{\mathcal{PL}}(\mathcal{M}_x^r(P), 2)$$

at any point x in the patch space P to indicate the stratified topology of the patch space. Although resembling the shape of nested spheres there is no intersection between the smaller spheres inside the largest sphere as indicated by the local homology not being irregular at the visual intersection in the 3-dimensional embedding. Furthermore, the global persistent homology also supports the hypothesis of five nested

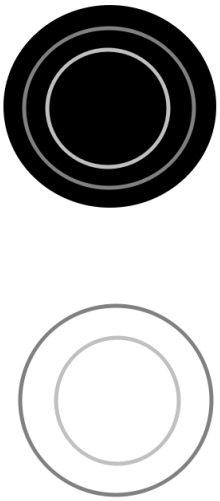

Figure 5.10. Simple grayscale cluster image

spheres representing the transition between four different colors. Consider Fig. 5.13 for the persistence diagram of the patch space up to a finite radius. In dimension two we see exactly five persistent cycles that represent the fact that there are five different complete transitions between two colors. Furthermore, the two persistent one-dimensional cycles reflect that there are two different possibilities for cyclic transition form one color to another.

The illustrated image in Example 5.1.1 depicts a particularly unique case for an image, and it is not within our expectations for a natural image patch space to comprise such a limited array of patches. This expectation is substantiated by existing scientific findings, as evidenced in, for instance, [LPM03]. Additionally, the prevalence of (ideal) step edges predominantly occurs in regions of an image devoid of textures.

Moreover, in the context of Example 5.1.1, we have already observed that each color transition within a given image may correspond to an S^2 within the image patch space. This preliminary observation hints at the computational intricacies inherent in the examination of the image patch space encompassing all color transitions.

Background

To refine this concept, we aim to introduce a model for the image patch space comprising exclusively what we will term as *step edge patches*. In developing a mathematical model for step edge patches, it is crucial to recognize that an ideal step edge patch can be defined by two parameters—the distance from the patch's center and the angle of the edge. Consequently, the patch space exclusively composed of step edges situated between two colors is isomorphic to S^2. It is worth noting that the exclusion of

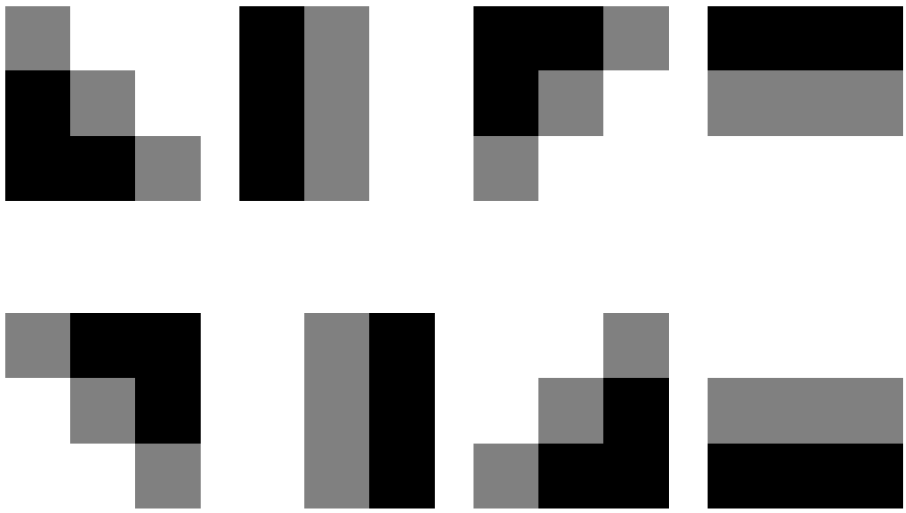

Figure 5.11. Step edge patches from simple image

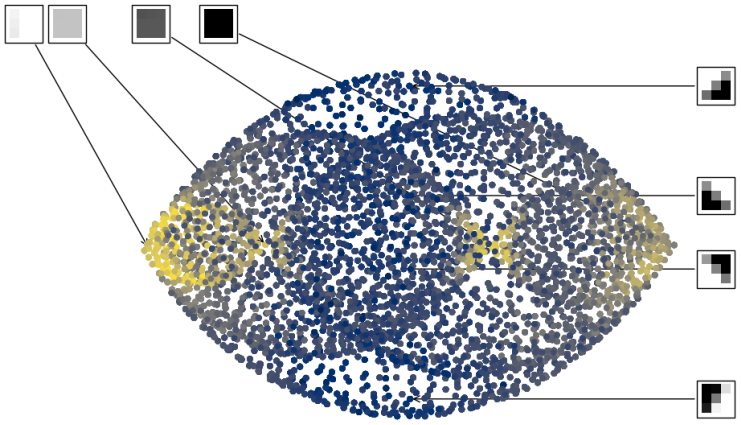

Figure 5.12. PCA embedding of step edge patch space with coloring by $\Phi_{\mathcal{PL}}(\mathcal{M}_x^r(P), 2)$

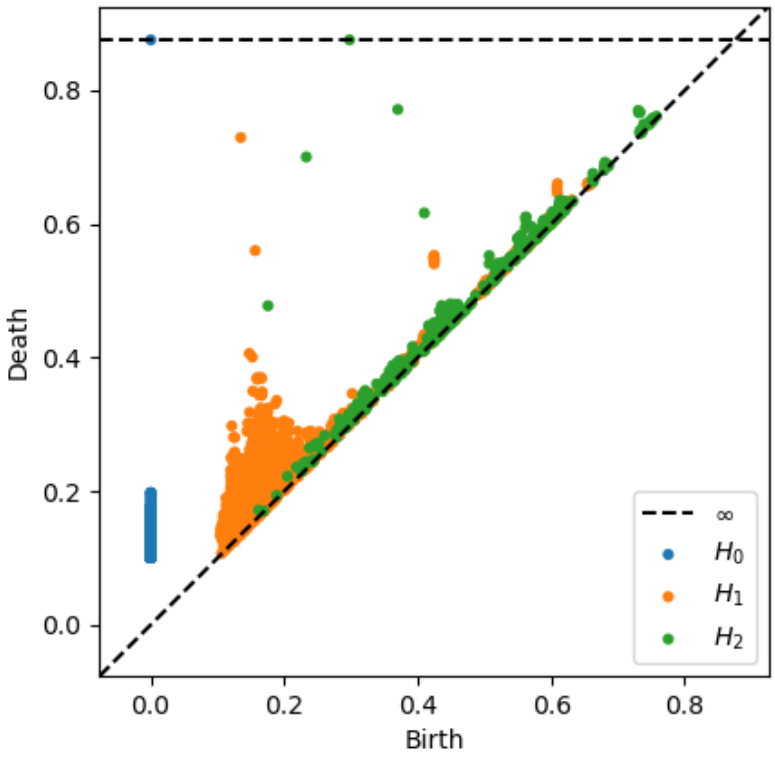

Figure 5.13. Persistence diagram up to a finite scale and dimension two of the simple image patch space

low-contrast patches would alter the topology to that of an annulus, which is, at the very least, homotopy equivalent to S^1.

For a given number of colors, we can systematically generate all step edge patches facilitating transitions between two colors. For the sake of simplicity, we consider colors to be represented on the grayscale, designated by a real number within the range of 0 to 1. Importantly, different approaches to defining a step edge patch may exist, potentially yielding distinct outcomes concerning patch spaces. In our context, we opt to characterize step edge patches through the following polynomial description:

Definition 5.1.2. Let $p^{\phi,\theta} : \{-1, 0, 1\}^2 \to [-1, 1]$ be the polynomial given by

$$q^{\phi,\theta}(x, y) = \sin(\phi)(1 + \cos(\theta)x + \sin(\theta)y)$$

for $\theta \in [0, 2\pi]$ and $\phi \in [0, \frac{\pi}{2})$. An *(ideal) step edge patch* between colors $t_1, t_2 \in \mathbb{R}^c$ is given by

$$p_{ij}^{\phi,\theta}(t_1, t_2) = \frac{\mid t_1 - t_2 \mid}{2} q^{\phi,\theta}(i, j) + t_1$$

for $(i, j) \in \{-1, 0, 1\}^2$.

For a given set of colors $\mathcal{T} = \{t_1, \ldots, t_k\}$ we denote the image patch space of all (ideal) step edge patches by $\mathcal{G}(\mathcal{T}) = \{p^{\phi,\theta}(t_i, t_j) \mid \phi \in [0, \frac{\pi}{2}), \theta \in [0, 2\pi](t_i, t_j) \in \binom{\mathcal{T}}{2}\}$.

Definition 5.1.3. An image $f : \Omega \to \mathcal{T} \subset \mathbb{R}^c$ is called *cluster image* if the associated patch space P contains only step edge patches, i.e. $P \subset \mathcal{G}(\mathcal{T})$.

Remark 5.1.4. The definition of the ideal step edge patches as well as the definition of cluster images in this context is intended to give a formal object to refer to that models a certain class of patches that actually occur naturally in images. It is highly unlikely that a patch space from a natural image, e.g. one from the van Hateren collection, can be classified as a cluster image. However, we will use our model of cluster images to describe the topology of a certain subset of patch spaces from arbitrary images and to explore which visual information is contained in image patches that lie close to a space $\mathcal{G}(\mathcal{T})$. The topology of the space of all step edge patches between colors $\mathcal{T}$ is can be described as a union of $\binom{|\mathcal{T}|}{2}$ 2-spheres whose pairwise intersection is either given by a single point or the empty set.

We have seen an indication of what type of image can be faithfully classified as cluster image in Example 5.1.1 and what effect the replacement of natural image patches with step edge patches has in Example 5.1.5. More general, cartoonish images fit the model of a cluster image very well. However, we would like to briefly demonstrate that even a rather natural image can be represented as a cluster image without changing the visuals detrimentally.

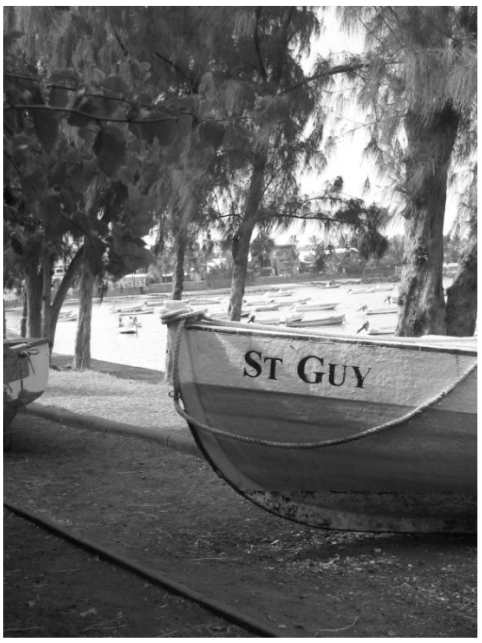

Figure 5.14. Grayscale example image from the INRIA Holidays dataset

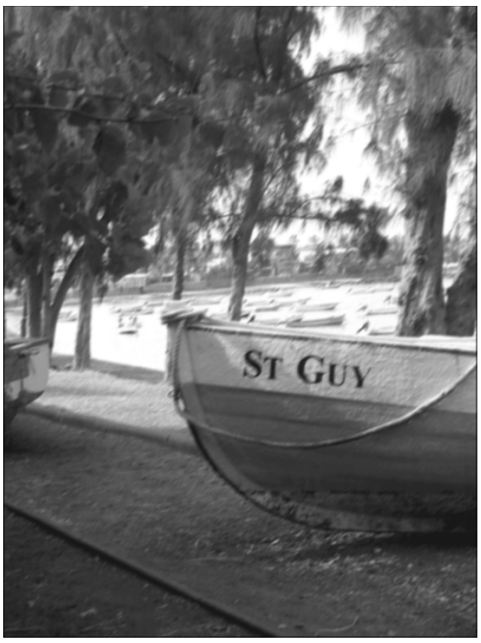

Figure 5.15. Grayscale INRIA Holidays image after transformation to $\mathcal{G}(\mathcal{T})$

Example 5.1.5. To illustrate what portion of a natural image can be represented solely with step edge patches, we have incorporated an example image extracted from the INRIA Holidays image dataset [JDS08]. In Fig. 5.14, a grayscale image from the dataset is depicted. Initially, we constructed the patch space of width 3, designated as P, from the grayscale image. Subsequently, each patch in P was substituted with its nearest point in the space $\mathcal{G}(\mathcal{T})$, where $|\mathcal{T}| = 15$, representing the 15 most dominant colors in the given image. Following this replacement, an image can be reconstructed from the collection of step edge patches, allowing an examination of the visual impact of this transformation. The outcome is portrayed in Fig. 5.15. The image exhibits an overall smoother or blurred appearance and has lost many of its textures. This example underscores the challenge of capturing the topology of an image patch space, particularly with the metric imposed on the image patch space. Although not identical, the images presented in Figs. 5.14 and 5.15 are visually akin while harboring distinct topological features.

The upshot of the transformation of an image as in Example 5.1.5 is to make certain topological characteristics of an image accessible that were not available by other means. How informative this characterization of an image can be will depend on

the application and remains open at this point. Nevertheless, we made a first attempt to demonstrate the utility of this topological feature of image data in the next part.

Dataset

We considered the CIFAR-100 dataset [KH09] as a source of rather natural images. The dataset is a collection of small photographs of various natural scene types, e.g. landscape, man-made outdoor scenes, animals, flower, people. The dataset is organized by groups of images. Each group contains images of a certain type, e.g. one group contains only images of people and another contains images showing landscapes. This is ideal for our purposes as we are interested in seeing if we are able to associate distinct topological features with different types of images. Additionally, this dataset contains images that are all of size 32×32 which makes the dataset uniform.

Preprocessing

We want to compare two classes of images contained in the CIFAR-100 dataset. For every image in a group we determined the five most dominant colors $t_1, \ldots, t_5$ and used these colors to generate spaces $\mathcal{G}(\mathcal{T})$ with $\mathcal{T} = \{t_1, \ldots, t_5\}$. Furthermore, we generated the associated patch space P of width 3 for every image. Then, for every patch $p \in P$ we determined the patch in $\mathcal{G}(\mathcal{T})$ with minimal distance to p. In that way we generated a patch space $P_c \subset \mathcal{G}(\mathcal{T})$ whose patches are as closely matched to the patches in P.

Analysis

We computed the persistent homology of the spaces P_c for every image in the respective two groups from the dataset. We extracted the persistence bars from the persistence diagrams that were longer than a threshold determined from the distances between colors in $\mathcal{T}$. This was done separately in every dimension to preserve most of the topological information. The persistence was summed up to a single value in dimension one and two. For every image we generated a 4-dimensional vector containing the summed persistence and the number of homology 1- and 2-cycles that persisted for longer than a fixed and uniform threshold. Let α_i^k denote a persistent homology cycle in dimension $k = 1, 2$ that persists for $\mathrm{Pers}(\alpha_i^k)$. The format of the vectors are $(\sum_i^m \mathrm{Pers}(\alpha_i^2), \sum_i^l \mathrm{Pers}(\alpha_i^1), m, l)$. In the CIFAR-100 dataset the groups are of size 2500. We selected only the first 500 images of each group as data sample to work with. We then merged the two sets of 4-dimensional vectors containing the homological information, i.e. we simply forgot the grouping, and then applied a standard k-means

Figure 5.16. Excerpt image from the CIFAR-100 landscape image group

Figure 5.17. Excerpt image from the CIFAR-100 people image group

clustering with $k = 2$. In Figs. 5.16 and 5.17, we show two examples images from the respective groups along-side their associated step edge patch space Figs. 5.18 and 5.19. In that exemplary case the difference in topology is quite drastic which is clearly not always the case. To be more precise, we can report that the two clusters formed by the k-means algorithm separated the two groups again with high accuracy. That is, 84.4 percent of the images from the people image group were assigned to the first cluster whose center was calculated as $(91, 80, 5, 77)$, i.e. a total persistence value of 91 from 5 persistent homology 2-cycles. In the second cluster with center vector given by $(32, 128, 2, 31)$ we find 87.8 percent of the cluster are images from the landscape group. We record an overall larger amount of short to medium persistent 1-cycles in this dataset which we attribute to the fact that 1-cycles need less points to be formed and are therefore much more susceptible to noise. Still, these results indicate a clear trend in the topology of these image patch spaces that we will interpret in view of the visual differences of the image groups in the next paragraph.

Interpretation

The results of our analysis indicate a measurable topological difference between different types of images. Our short investigation was by no means intended to develop a new method for image classification but rather to demonstrate the presence of topological structure in image data. The selected image groups were visually very different. One group showing landscapes which are organized on few (mostly two) large areas of similar color and texture separated horizontally or vertically. This corresponds to the strong tendency of these images, after transformation to the cluster image patch space, to have few but very persistent homology cycles, especially in dimension 1. This stems from the fact that the direction of step edge patches present in the associated

Figure 5.18. Step edge patch space generated from landscape image in Fig. 5.18

Figure 5.19. Step edge patch space generated from people image in Fig. 5.17

patch spaces are restricted in their directions. The persistent 1-cycles may come from incomplete 2-spheres in the space of all step edge patches.

The group of images showing people (and most of the time the upper body) usually feature a strong background to foreground contrast. The foreground is dominated by the outline of a person and the color transitions occur in several directions. Furthermore, we can often record more color transition occurring in these types of images. This results in an increased number of 2-cycles in the image patch space generated from the images. Note, however, that although the number of 2-cycles was usually larger than in landscape images, the cycles were also less persistent indicating a smaller distance in color space between which the transitions happen.

5.2 Artery images and pixel patches

Background and other work

Image patch spaces of width 1 are another special case as the patches are the pixels of an image itself. The Euclidean distance between pixels, i.e. the color values, is topologically not significant. However, another obvious choice of metric would be the one induced by the metric on the lattice Ω associated with an image. Combined with a filtration on the color values of the pixels or with a distance between them, we can generate point clouds from the 2-dimensional geometric shapes visible in an image.

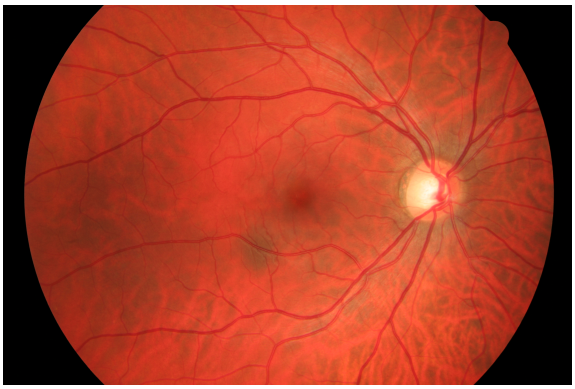

Figure 5.20. Fundus camera image of healthy retina

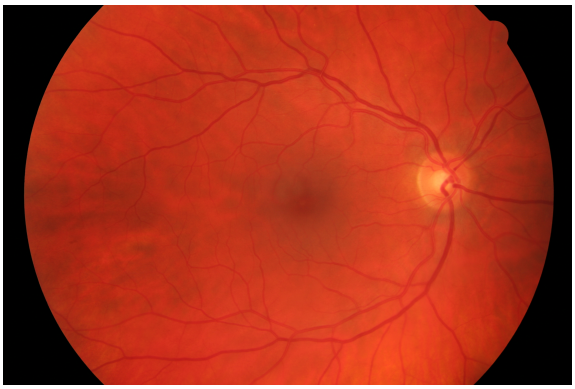

Figure 5.21. Fundus camera image of glaucoma patient's retina

The analysis of images based on varying local geometry in an image is common in the field of image segmentation and clustering. In [ESM+21], for example, the authors used persistent homology of local image regions for the segmentation of microscopy images of cells into functional regions, e.g. cell border or nuclei. Because microscopy images of cells are organized into different functional areas that result in different geometrical features present in a cell image, persistent homology can be used to distinguish these regions and is therefore used in [ESM+21] to determine which regions contribute to the classification in a machine learning task.

Dataset

The dataset we want to investigate here is a set of manually segmented retina photos. After the photos were taken with a fundus camera, a group of retinal image analysis experts and clinicians from ophthalmology clinics manually segmented the photos into blood vessels and background. The dataset is composed of 15 healthy retina images, 15 retinas from glaucoma patients and 15 images of diabetes patients, which we did not consider in our analysis, and was created by [BBM+13]. Examples of the original fundus camera pictures can be seen in Figs. 5.20 and 5.21 and for examples of manually segmented images, see Figs. 5.22 and 5.23.

Preprocessing

For a given segmented retina photo $f : \Omega \to \mathcal{T}$ we consider the space of patches P of width 1 equipped with the induced metric from the lattice Ω. Because the segmented images are actually binary images and the patches are literally the pixels of the image we can select only the patches that make the blood vessels in the image by thresholding the pixel values, i .e. the point cloud we are working with is $P' := \{p \in P \mid p[0] > 0\}$.

For two-fold reasons we thinned-out the point cloud by iterating over all patches and including the current patch if the distance to all previously included patches is less or equal a chosen minimum δ. We obtain a space of patches P'' such that $p, q \in P''$ implies $d(p, q) >= \delta$. This is a common procedure in data analysis to reduce noise or as in our case to improve computational performance. Besides this effect we wanted to make sure that the points in the datasets extracted from both classes of images are equally dense and our persistent homology results are not biased by such an effect. The resulting space is a collection of points in the plane that sample the shape of blood vessels in a participants retina.

Analysis

Topologically speaking, this space presents itself with many singularities of different kind and we investigated its type with the methods we described in the Chapter 4.

- Sticking to usual notion for sample spaces, let $\mathbb{X}_i^g$ and $\mathbb{X}_i^h$ denote the spaces P'' as described in the preprocessing for the i-th image from the set of glaucoma or healthy images respectively.
- The metric on $\mathbb{X}_i^*$ is induced by the natural image lattice Ω.
- As measure for singularity, i.e. as Φ-function we chose $\Phi_{\mathcal{PL}}$.
- The magnification parameter was set to $r = 30$, thus, by evaluating the fibers of $\mathcal{M}^r(\mathbb{X}_i^*)$, we obtain a strongly stratified sample $\Phi_{\mathcal{PL}*}(\mathcal{M}^r(\mathbb{X}_i^*), 1)$ for $i = 1, \ldots, 15$ and for both image classes.
- The threshold value u was set to $u = 0.58$ for all spaces considered.
- We record two values obtained by this method. The first is the total amount of points in $\mathbb{X}_i^*$ for which the threshold was surpassed. Second, we calculated persistent homology in dimension 0 of the sample given by all points above the threshold to retrieve the amount of singular regions.

Results

In Table 5.1, we see the amount of points in $\mathbb{X}_i^*$ for which the threshold was surpassed by $\Phi_{\mathcal{PL}}(\mathcal{M}^r(\mathbb{X}_i^*), 1)$. For every healthy test person the calculated value lies above the value of any glaucoma patient. We proceeded by calculating the persistent connected components of the approximated singular strata. See Table 5.2 for the recorded numbers.While the presentation of our results may not be as straightforward as in the case

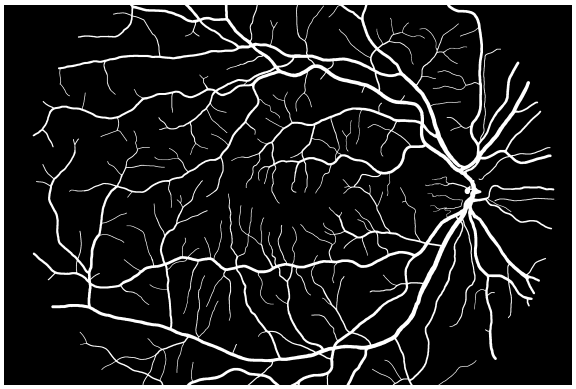

Figure 5.22. Healthy

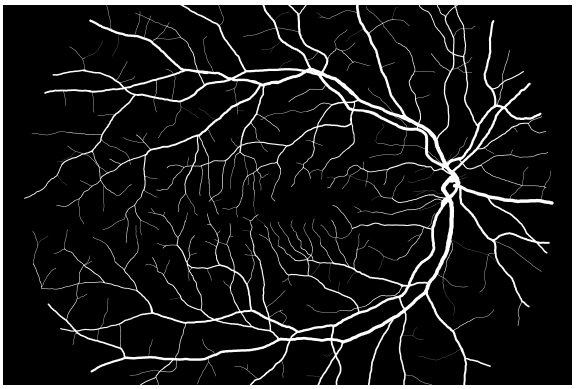

Figure 5.23. Glaucoma

of the point count surpassing the threshold, it is evident that, in general, the healthy group displays a greater prevalence of singular regions.

Interpretation

Our findings indicate that local persistent homology proves effective in distinguishing between glaucoma patients and healthy test participants using manually segmented retina images obtained through a fundus camera. The key discriminative factors lie in both the quantity of points in $\mathbb{X}_i^*$ reaching the threshold (Table 5.1) and the persistence of singular regions within the sample space (Table 5.2).

To contextualize our results, a study by Rudnicka et al. [ROW+20] reported smaller average area and width of arterioles and venules in glaucoma retina photography compared to healthy arteries. However, the study did not address a reduction in the number of bifurcations, branches, and crossings in arteries caused by glaucoma. Consequently, we hypothesize that components within the singular stratum of $\mathcal{S}^r_{\Phi,u}$, corresponding to regions around crossings, branchings, and bifurcations in a retina image, exhibit greater persistence in pixel data extracted from healthy groups. This may be attributed to our method's capability to detect singular aspects in regions where arterioles and venules are narrow, a feature that may be less discernible in the glaucoma group where these vessels are too thin.

Exploring the robustness of our findings, it would be intriguing to apply our analysis to algorithmically segmented retina photos to assess the reliability of our results under less-than-perfect image segmentation conditions.

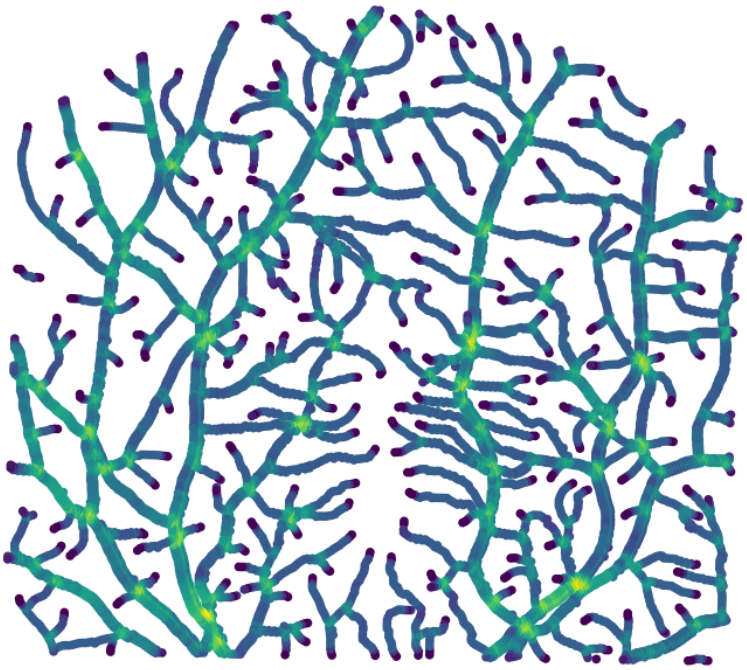

Figure 5.24. $\Phi_{\mathcal{P}\mathrm{L}*}(\mathcal{M}^r(\mathbb{X}_i^*), 1)$ of retina image

Image number	Healthy	Glaucoma
1	9147	3728
2	9646	4526
3	9251	2692
4	7957	2694
5	8485	3618
6	7367	3922
7	7180	3888
8	9557	3891
9	5571	2884
10	5980	3066
11	8016	5069
12	11575	5010
13	6833	3533
14	5726	4114
15	5154	4559

Table 5.1. Number of points where $\Phi_{\mathcal{P}\mathrm{L}*}(\mathcal{M}^r(\mathbb{X}_i^*), 1)$ is above threshold

Image number	Healthy	Glaucoma
1	199	132
2	194	142
3	200	116
4	182	103
5	193	107
6	191	124
7	159	128
8	173	129
9	174	114
10	177	105
11	174	139
12	180	**185**
13	184	148
14	191	128
15	**133**	125

Table 5.2. Number of persistent connected components in the singular strata of $\mathcal{S}^r_{\Phi,u}(\mathbb{X}^*_i, 1)$

Chapter 6

Conclusion

In this final chapter we would like to discuss open questions and some promising directions based on the results we presented. We address the results of every previous chapter separately and in the same order they appeared in this thesis.

6.1 Persistent Stratified Homotopy Types

In Section 2.2, we have validated that the concept of persistent stratified homotopy type indeed satisfies the expected stratified analogs of the properties delineated for the persistent homotopy type, specifically encompassing properties (1), (2), and (3) mentioned in the introduction. In concise terms, we have demonstrated its computability (as discussed in Remark 2.2.15), stability (as elucidated in Theorems 2.2.31 and 2.2.32), and its ability to infer information through invariance under small thickenings (detailed in Proposition 2.2.16).

An avenue for future research may involve exploring extensions beyond the two strata case. The deliberate choice to concentrate on the two-strata case at the outset serves as a strategic starting point, facilitating the introduction of fundamental concepts such as persistent stratified homotopy types and stratification learning. This deliberate focus helps manage technical complexity in the initial stages.

Expanding the scope to more than two strata represents a logical progression but comes with an anticipated increase in complexity. Although the specific techniques for investigating persistent stratified homotopy types in the context of more intricate poset structures remain uncertain, it is noteworthy that certain aspects of the associated abstract homotopy theory have already been investigated ([Dou19, DW21, Hai18, AFR19]). The theoretical groundwork in abstract homotopy theory provides a foundation and potential guidance for future explorations. Nevertheless, extending the analysis to the multi strata case holds substantial promise for advancing the field

of topological data analysis. This extension allows for the examination of even more complex stratified structures that may emerge in real-world applications.

A crucial future direction involves exploring suitable algebraic invariants for stratified spaces within the framework of Topological Data Analysis (TDA). Although we have already used algebraic invariants capable of discerning singular spaces, readily computable from point cloud data, namely, the persistent homology of links and strata, there exists a persistent version of intersection homology (as discussed in, for instance, [BH11]) that demands further investigation, particularly in terms of its computability from stratified samples. Based on our investigations of persistent stratified homotopy types and their representation by thickenings of stratification diagrams, this work provides results that further indicate possible approaches to address this aspect.

Consequently, there is considerable room for future refinements and generalizations in subsequent research pursuits.

6.2 Approximate Stratification

We have worked towards a theoretical framework that allows to identify methods that can be used in practice to infer stratified structures in point cloud data. With Theorem 3.5.8 we reported the result that proves that it is possible to approximate the stratification for a large class of two strata Whitney stratified spaces from sufficiently close non-stratified samples.

Parallel to the conclusions drawn from our investigations into persistent stratified homotopy types, extending our results to the multi-strata case presents a promising avenue for further exploration. Similarly, challenges emerge in this extension, particularly concerning the increased difficulty in learning stratifications from non-stratified data when confronted with a more complex underlying stratification poset. This heightened complexity is exemplified by instances like the Whitney Umbrella, as documented [Ban07, p. 128-129].

It is noteworthy that certain outcomes in this direction transcend the confines of the two-strata case, as demonstrated by the convergence of magnifications to tangent cones (Proposition 3.3.11). This observation hints at the potential to generalize our approximation result (Theorem 3.5.8) for suitable functions Φ and sufficiently well-behaved spaces. Such generalizations may even offer avenues to address the challenges posed by examples like the Whitney Umbrella.

6.3 Algorithmic Stratification Learning

Our overarching aim was to provide a comprehensive exposition of the algorithmic aspects inherent in our work. The primary objective was to establish a connection between the advanced mathematical concepts elucidated in preceding chapters and the pragmatic implementation of stratifying a point cloud derived from a stratified space. In particular, we meticulously expound upon the computation of two specific choices for function Φ that can be used as a measure of singularity in point clouds. One of the functions leveraged local persistent homology and the other was based on the local distance to linear subspace of some Euclidean space.

Local persistent homology:

The function $\Phi_{\mathcal{PL}}$, introduced in *Example* 3.2.3, presents a viable method for stratifying a point cloud through the use of local persistent homology. Our goal was to devise an algorithm capable of executing the computations associated with the application of $\Phi_{\mathcal{PL}}$.

- We expounded upon and implemented a methodology for quantifying singularity within point cloud data, grounded in the application of local persistent homology (compare to Sections 4.1.2 and 4.3).
- The requisite sensitivity and robustness to noise of this methodology were demonstrated through assessments on both artificial and real-world datasets. This validation involved verifying anticipated outcomes (refer to Example 4.3.2 and Analysis of the three circle model) and providing additional insights into unexplored data domains (refer to Section 5.2).
- We also delineated the capacity to approximate genuine local persistent homology (see Proposition 4.1.29) through the use of a Vietoris-Rips type construction. Although the Vietoris-Rips complex is commonly used as it allows for efficient construction and storage, the identified worst-case scenario error (based on the interleaving of Vietoris-Rips and Čech complexes) suggests opportunities for enhancement, possibly through the exploration of alternative methods to evaluate the accuracy of our approach.
- Another avenue of exploration involves approximating $\Phi_{\mathcal{PL}}$ by utilizing a Čech type complex, briefly mentioned in Remark 4.1.23. Achieving this would necessitate further investigation into the efficient computation of Čech complexes.

Tangential Approximation

Once again, we introduced a function Φ_{Hd} in Example 3.2.2, which serves the purpose of identifying potential singularities in non-stratified samples. We outlined an initial algorithm for approximating the values of Φ_{Hd} on a given sample space and evaluated the accuracy of these approximations. The computations involve solving an optimization problem, formulated in both a deterministic and randomized manner.

While we explained how to formulate a smooth constrained optimization problem for calculating approximations of Φ_{Hd}, we did not provide an actual implementation. Although existing implementations for similar problems are available, exploring dedicated algorithms tailored for optimization on matrix representations of Grassmannians could prove advantageous. Previous investigations into optimization within Grassmannians have often relied on the differential structures of Grassmannians. Moreover, existing implementations of these methods may offer a promising starting point [LLY20, KWT23].

Regarding the implementation we provided for our random optimization approach, optimizing computational efficiency was not our primary goal, and there are certainly several areas for improvement. We highlight a few of them here:

- The generation of random elements in the Grassmannians $\mathrm{Gr}(n, k)$ can be achieved using algorithms of varying complexities. For instance, [Ste80] provides a description of an algorithm with a complexity of $\mathcal{O}(n^2)$ for generating random orthogonal matrices, which could be applicable in this context.

- Various sources discuss efficient methods for computing the Hausdorff distance, as outlined in [GBK05, TH15]. Approximative approaches may be particularly relevant in practice, given that the described algorithm inherently is of an approximative nature.

- Unlike our fundamental random optimization approach, more sophisticated designs such as genetic algorithms or stochastic gradient descent are worth considering [Bot04, Mit98, LJ73] to further improve on the applicability of our methods in future.

In total, the modular form of our stratification pipeline in the choice of functions Φ to measure singularity leaves room for future investigations. Different methods that may fit into this context and may be realized as a function Φ include methods based on curvature and density (compare to [RBSL20]) as well as tangent spaces and local dimension estimation (compare to [LON22]).

Our efforts reached fruition with the integration of theoretical insights and algorithmic advancements into a complete implementable stratification pipeline. Applying this pipeline to a controlled example, where we meticulously controlled the Hausdorff distance of the sample to the intended space (see Example 4.3.2), not only underscores its practical utility but also facilitates the demonstration of previously developed concepts.

6.4 Applications

The presence of stratified structures or singularities in data examples, providing distinctive features, is evident in various applications ([LPM03, BMP+08, CIDSZ08, MW11, FW16, Xia16, STHN20, ESM+21]). It is essential to acknowledge that point clouds in general cannot be expected to be homogeneous or sampled solely from a manifold. The exploration of stratifications in non-structured datasets is expanding, and this work aims to contribute to that growing body of knowledge. The availability of more appropriate methods empowers scientists to discover and recognize stratified structures within data.

Our analysis of data examples illustrated how established models, such as the three-circle model, could be reevaluated and supported by our methods, offering additional insights into the underlying topology of the data space (see Analysis of the three circle model).

Furthermore, we demonstrated the identification of stratified structures in simple applications, such as manually segmented fundus camera images (discussed in Section 5.2). The presence of distinctive features, characterized by the number of crossings and branchings, allowed us to organize the dataset into distinct groups based on this structure.

Our investigation into the topology of image patch spaces, addressing the question of whether the underlying structure resembles that of a manifold or a stratified space, is not exhaustive. We presented a model for characterizing image patch spaces extracted from a specific class of images, namely cluster images (see Definition 5.1.3), and found indications of a singular structure. Future research could explore additional elements to augment the proposed model for cluster image patch spaces. For instance, incorporating not only step edge patches (see Definition 5.1.2) but also other frequently observed image patches depicting lines and cross formations could provide further insights into how the topology changes. This may enable us to represent a broader class of images or potentially identify conditions under which this is feasible for certain images.

Bibliography

[AEK06] P.-A. Absil, A. Edelman, and P. Koev. On the largest principal angle between random subspaces. *Linear Algebra Appl.*, 414(1):288–294, 2006.

[AFR19] David Ayala, John Francis, and Nick Rozenblyum. A stratified homotopy hypothesis. *J. Eur. Math. Soc. (JEMS)*, 21(4):1071–1178, 2019.

[AGV+18] Rabih Assaf, Alban Goupil, Valeriu Vrabie, Thomas Boudier, and Mohammad Kacim. Persistent homology for object segmentation in multidimensional grayscale images. *Pattern Recognition Letters*, 112:277–284, 2018.

[AWR17] J. Altschuler, J. Weed, and P. Rigollet. Near-Linear Time Approximation Algorithms for Optimal Transport via Sinkhorn Iteration. In *NIPS*, 2017.

[Ban07] Markus Banagl. *Topological invariants of stratified spaces*. Springer, Berlin, Germany, 2007.

[Bau21] Ulrich Bauer. Ripser: efficient computation of Vietoris-Rips persistence barcodes. *J. Appl. Comput. Topol.*, 5(3):391–423, 2021.

[BB18] Håvard Bakke Bjerkevik and Magnus Bakke Botnan. Computational Complexity of the Interleaving Distance. In *34th International Symposium on Computational Geometry (SoCG 2018)*, volume 99 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 13:1–13:15, Dagstuhl, Germany, 2018. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.

[BB20] Nello Blaser and Morten Brun. Relative persistent homology. In *36th International Symposium on Computational Geometry*, volume 164 of *LIPIcs. Leibniz Int. Proc. Inform.*, pages Art. No. 18, 10. Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern, 2020.

[BBM+13] Attila Budai, Rüdiger Bock, Andreas Maier, Joachim Hornegger, and Georg Michelson. Robust vessel segmentation in fundus images. *International journal of biomedical imaging*, 2013, 2013.

[BCB20] Magnus Bakke Botnan and William Crawley-Boevey. Decomposition of persistence modules. *Proc. Amer. Math. Soc.*, 148(11):4581–4596, 2020.

[BCSE+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'07)*, pages 536–546, 2007.

[BG73] Ake Björck and Gene Golub. Numerical methods for computing angles between linear subspaces. *Mathematics of Computation*, 27:123, 07 1973.

[BH11] Paul Bendich and John Harer. Persistent intersection homology. *Found. Comput. Math.*, 11(3):305–336, 2011.

[BL07] Andreas Bernig and Alexander Lytchak. Tangent spaces and Gromov-Hausdorff limits of subanalytic spaces. *J. Reine Angew. Math.*, 608:1–15, 2007.

[BL15] Ulrich Bauer and Michael Lesnick. Induced matchings and the algebraic stability of persistence barcodes. *J. Comput. Geom.*, 6(2):162–191, 2015.

[BL21] Håvard Bakke Bjerkevik and Michael Lesnick. l^p-distances on multiparameter persistence modules, 2021. arXiv preprint available at arXiv:2106.13589.

[BMP+08] W Brown, Shawn Martin, Sara Pollock, Evangelos Coutsias, and Jean-Paul Watson. Algorithmic dimensionality reduction for molecular structure analysis. *The Journal of chemical physics*, 129:064118, 09 2008.

[BMS20] Markus Banagl, Tim Mäder, and Filip Sadlo. Stratified formal deformations and intersection homology of data point clouds, 2020. arXiv preprint available at arXiv:2005.11985.

[Bor48] Karol Borsuk. On the imbedding of systems of compacta in simplicial complexes. *Fund. Math.*, 35:217–234, 1948.

[Bot04] Léon Bottou. Stochastic learning. In Olivier Bousquet and Ulrike von Luxburg, editors, *Advanced Lectures on Machine Learning*, Lecture Notes

in Artificial Intelligence, LNAI 3176, pages 146–168. Springer Verlag, Berlin, 2004.

[BP20] Jean-Daniel Boissonnat and Siddharth Pritam. Edge collapse and persistence of flag complexes. In *International Symposium on Computational Geometry*, 2020.

[BSS20] Peter Bubenik, John A. Scott, and Don Stanley. An algebraic wasserstein distance for generalized persistence modules, 2020. arXiv preprint available at arXiv:1809.09654v2.

[BW20] David Bramer and Guo-Wei Wei. Atom-specific persistent homology and its application to protein flexibility analysis. *Comput. Math. Biophys.*, 8:1–35, 2020.

[BWM12] Paul Bendich, Bei Wang, and Sayan Mukherjee. Local homology transfer and stratification learning. In *Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1355–1370, New York, NY, USA, 2012. Association for Computing Machinery.

[Car09] Gunnar Carlsson. Topology and data. *Bull. Amer. Math. Soc. (N.S.)*, 46(2):255–308, 2009.

[CB14] William Crawley-Boevey. Decomposition of pointwise finite-dimensional persistence modules, 2014. arXiv preprint available at arXiv:1210:0819.

[CCSG+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, pages 237–246, New York, NY, USA, 2009. Association for Computing Machinery.

[CCSL09] Frédéric Chazal, David Cohen-Steiner, and André Lieutier. A sampling theorem for compact sets in Euclidean space. *Discrete Comput. Geom.*, 41(3):461–479, 2009.

[CdSGO16] Frédéric Chazal, Vin de Silva, Marc Glisse, and Steve Oudot. *The structure and stability of persistence modules*. SpringerBriefs in Mathematics. Springer, [Cham], 2016.

[CG20] Gunnar Carlsson and Rickard Brüel Gabrielsson. Topological approaches to deep learning. In Nils A. Baas, Gunnar E. Carlsson, Gereon Quick,

Markus Szymik, and Marius Thaule, editors, *Topological Data Analysis*, pages 119–146. Springer International Publishing, 2020.

[Chi03] Yasuko Chikuse. *Statistics on special manifolds*, volume 174 of *Lecture Notes in Statistics*. Springer-Verlag, New York, 2003.

[CIDSZ08] Gunnar Carlsson, Tigran Ishkhanov, Vin De Silva, and Afra Zomorodian. On the local behavior of spaces of natural images. *International journal of computer vision*, 76(1):1–12, 2008.

[CK11] Chao Chen and Michael Kerber. Persistent homology computation with a twist. 2011.

[Cos00] Michel Coste. *An introduction to o-minimal geometry*. Istituti editoriali e poligrafici internazionali, Pisa, Italy, 2000.

[CRL+20] L. Chizat, P. Roussillon, F. Léger, F.-X. Vialard, and G. Peyré. Faster Wasserstein Distance Estimation with the Sinkhorn Distance. *CoRR abs/2006.08172*, 2020.

[CSEH07] David Cohen-Steiner, Herbert Edelsbrunner, and John Harer. Stability of persistence diagrams. *Discrete Computational Geometry*, 37:103–120, 2007.

[Cut13] M. Cuturi. Sinkhorn Distances: Lightspeed Computation of Optimal Transport. In *Proc. NIPS*, Wasserstein distance; entropy regularization; discrete optimal transport; 2013.

[Cza12] Mał gorzata Czapla. Definable triangulations with regularity conditions. *Geom. Topol.*, 16(4):2067–2095, 2012.

[DEHH18] Emilie Dufresne, Parker B. Edwards, Heather A. Harrington, and Jonathan D. Hauenstein. Sampling real algebraic varieties for topological data analysis, 2018. arXiv preprint available at arXiv:1802.07716.

[DMV17] Tamal Dey, Sayan Mandal, and William Varcho. Improved image classification using topological persistence. In *Proceedings of the Conference on Vision, Modeling and Visualization*, page 161–168, Goslar, DEU, 2017. Eurographics Association.

[Dou19] Sylvain Douteau. Stratified homotopy theory, 2019. arXiv preprint available at arXiv:1908.01366.

[Dou21a] Sylvain Douteau. Homotopy theory of stratified spaces. *Algebr. Geom. Topol.*, 21(1):507–541, 2021.

[Dou21b] Sylvain Douteau. A stratified kan-quillen equivalence, 2021. arXiv preprint available at arXiv:2102.04876.

[Drm00] Zlatko Drmac. On principal angles between subspaces of euclidean space. *Siam Journal on Matrix Analysis and Applications - SIAM J MATRIX ANAL APPLICAT*, 22, 06 2000.

[dSMS18] Vin de Silva, Elizabeth Munch, and Anastasios Stefanou. Theory of interleavings on categories with a flow. *Theory Appl. Categ.*, 33:Paper No. 21, 583–607, 2018.

[DW21] Sylvain Douteau and Lukas Waas. From homotopy links to stratified homotopy theories, 2021. arXiv preprint available at arXiv:2112.02394.

[EAS98] Alan Edelman, Tomás A. Arias, and S.T. Smith. The geometry of algorithms with orthogonality constraints. *SIAM J. Matrix Anal. Appl.*, 20:303–353, 1998.

[EH10] Herbert Edelsbrunner and John L. Harer. *Computational topology*. American Mathematical Society, Providence, RI, 2010. An introduction.

[Ehr51] Charles Ehresmann. Les connexions infinitésimales dans un espace fibré différentiable. In *Colloque de topologie (espaces fibrés), Bruxelles, 1950*, pages 29–55. Georges Thone and Masson & Cie, Liège and Paris, 1951.

[ELZ00] Herbert Edelsbrunner, David Letscher, and Afra Zomorodian. Topological persistence and simplification. In *41st Annual Symposium on Foundations of Computer Science (Redondo Beach, CA, 2000)*, pages 454–463. IEEE Comput. Soc. Press, Los Alamitos, CA, USA, 2000.

[ESM+21] Parker Edwards, Kristen Skruber, Nikola Milićević, James B. Heidings, Tracy-Ann Read, Peter Bubenik, and Eric A. Vitriol. TDAExplore: Quantitative analysis of fluorescence microscopy images through topology-based machine learning. *Patterns*, page 100367, 2021.

[Fri03] Greg Friedman. Stratified fibrations and the intersection homology of the regular neighborhoods of bottom strata. *Topology Appl.*, 134(2):69–109, 2003.

[FW16] Brittany Terese Fasy and Bei Wang. Exploring persistent local homology in topological data analysis. In *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6430–6434, New York, NY, USA, 2016. Institute of Electrical and Electronics Engineers (IEEE).

[GBK05] Michael Guthe, Pavel Borodin, and Reinhard Klein. Fast and accurate hausdorff distance calculation between meshes. volume 13, pages 41–48, 01 2005.

[Ghr08] Robert Ghrist. Barcodes: the persistent topology of data. *Bull. Amer. Math. Soc. (N.S.)*, 45(1):61–75, 2008.

[GM80] Mark Goresky and Robert MacPherson. Intersection homology theory. *Topology*, 19(2):135–162, 1980.

[GM83] Mark Goresky and Robert MacPherson. Intersection homology. II. *Invent. Math.*, 72(1):77–129, 1983.

[Gor78] R. Mark Goresky. Triangulation of stratified objects. *Proc. Amer. Math. Soc.*, 72(1):193–200, 1978.

[Hai18] Peter J. Haine. On the homotopy theory of stratified spaces, 2018. arXiv preprint available at arXiv:1811.01119.

[Hat02] Allen Hatcher. *Algebraic topology*. Cambridge University Press, Cambridge, UK, 2002.

[Hir69] Heisuke Hironaka. Normal cones in analytic Whitney stratifications. *Inst. Hautes Études Sci. Publ. Math.*, 36:127–138, 1969.

[HKNU17] Christoph Hofer, Roland Kwitt, Marc Niethammer, and Andreas Uhl. Deep learning with topological signatures. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30, page 1633–1643, Red Hook, NY, USA, 2017. Curran Associates, Inc.

[Hug99] Bruce Hughes. Stratifications of mapping cylinders. *Topology and its Applications*, 94(1-3):127–145, 1999.

[JD15] Bo Jiang and Yu-Hong Dai. A framework of constraint preserving update schemes for optimization on Stiefel manifold. *Math. Program.*, 153(2):535–575, 2015.

[JDS08] Herve Jegou, Matthijs Douze, and Cordelia Schmid. Hamming embedding and weak geometry consistency for large scale image search. In *Proceedings of the 10th European conference on Computer vision*, 2008.

[KE06] Plamen Koev and Alan Edelman. The efficient evaluation of the hypergeometric function of a matrix argument. *Math. Comp.*, 75(254):833–846, 2006.

[KH09] Alex Krizhevsky and Geoffrey Hinton. Learning multiple layers of features from tiny images, 2009.

[KMN17] Michael Kerber, Dmitriy Morozov, and Arnur Nigmetov. Geometry helps to compare persistence diagrams. *ACM J. Exp. Algorithmics*, 22:Art. 1.4, 20, 2017.

[KWT23] Lingkai Kong, Yuqing Wang, and Molei Tao. Momentum stiefel optimizer, with applications to suitably-orthogonal attention, and optimal transport. In *International Conference on Learning Representations*, 2023.

[Les15] Michael Lesnick. The theory of the interleaving distance on multidimensional persistence modules. *Found. Comput. Math.*, 15(3):613–650, 2015.

[LJ73] Rein Luus and T. H. I. Jaakola. Optimization by direct search and systematic reduction of the size of search region. *AIChE Journal*, 19(4):760–766, 1973.

[LLY20] Zehua Lai, Lek-Heng Lim, and Ke Ye. Simpler grassmannian optimization, 2020. arXiv preprint available at arXiv:2009.13502.

[Łoj65] Stanislaw Łojasiewicz. Ensembles semi-analytiques, 1965. IHES notes, available at `http://perso.univ-rennes1.fr/michel.coste/Lojasiewicz.pdf`.

[Loi98] Ta Lê Loi. Verdier and strict Thom stratifications in o-minimal structures. *Illinois J. Math.*, 42(2):347–356, 1998.

[Loi16] Ta Lê Loi. Łojasiewicz inequalities in o-minimal structures. *Manuscripta Math.*, 150(1-2):59–72, 2016.

[LON22] Uzu Lim, Harald Oberhauser, and Vidit Nanda. Tangent space and dimension estimation with the wasserstein distance, 2022. arXiv preprint available at arxiv:2110.06357.

[LPM03] Ann B Lee, Kim S Pedersen, and David Mumford. The nonlinear statistics of high-contrast patches in natural images. *International Journal of Computer Vision*, 54(1):83–103, 2003.

[Lur17] Jacob Lurie. Higher algebra, 2017. available at: `https://www.math.ias.edu/~lurie/papers/HA.pdf`.

[Lyt04] Alexander Lytchak. Differentiation in metric spaces. *Algebra i Analiz*, 16(6):128–161, 2004.

[Mat12] John Mather. Notes on topological stability. *Bull. Amer. Math. Soc. (N.S.)*, 49(4):475–506, 2012.

[Mil94] Chris Miller. Expansions of the real field with power functions. *Annals of Pure and Applied Logic*, 68(1):79–94, 1994.

[Mil13] David A. Miller. Strongly stratified homotopy theory. *Trans. Amer. Math. Soc.*, 365(9):4933–4962, 2013.

[Mil21] Yuriy Mileyko. Another look at recovering local homology from samples of stratified sets. *J. Appl. Comput. Topol.*, 5(1):55–97, 2021.

[Mit98] Melanie Mitchell. *An introduction to genetic algorithms*. MIT press, 1998.

[MW11] Shawn Martin and Jean-Paul Watson. Non-manifold surface reconstruction from high-dimensional point cloud data. *Computational Geometry*, 44(8):427–441, 2011.

[MW22] Tim Mäder and Lukas Waas. From samples to persistent stratified homotopy types, 2022. arXiv preprint available at arxiv:2206.08926.

[Nan20] Vidit Nanda. Local cohomology and stratification. *Found. Comput. Math.*, 20(2):195–222, 2020.

[Nic07] Liviu Nicolaescu. *Lectures on the Geometry of Manifolds*. 09 2007.

[NSW08] Partha Niyogi, Stephen Smale, and Shmuel Weinberger. Finding the homology of submanifolds with high confidence from random samples. *Discrete Comput. Geom.*, 39(1-3):419–441, 2008.

[NV21] Guglielmo Nocera and Marco Volpe. Whitney stratifications are conically smooth, 2021. arXiv preprint available at arXiv:2105.09243.

[Oud15] Steve Y. Oudot. *Persistence theory: from quiver representations to data analysis*, volume 209. American Mathematical Society, Providence, RI, USA, 2015.

[Pfl01] Markus J. Pflaum. *Analytic and geometric study of stratified spaces*, volume 1768 of *Lecture Notes in Mathematics*. Springer, Berlin, Germany, 2001.

[Qui88] Frank Quinn. Homotopically stratified sets. *J. Amer. Math. Soc.*, 1(2):441–499, 1988.

[RB21] Yohai Reani and Omer Bobrowski. A coupled alpha complex, 2021. arXiv preprint available at arxiv:2105.08113.

[RBSL20] Bastian Rieck, Markus Banagl, Filip Sadlo, and Heike Leitte. Persistent intersection homology for the analysis of discrete data. In *Topological Methods in Data Analysis and Visualization V: Theory, Algorithms, and Applications 7*, pages 37–51. Springer, 2020.

[ROW+20] Alicja R. Rudnicka, Christopher G. Owen, Roshan A. Welikala, Sarah A. Barman, Peter H. Whincup, David P. Strachan, Michelle P.Y. Chan, Anthony P. Khawaja, David C. Broadway, Robert Luben, Shabina A. Hayat, Kay-Tee Khaw, and Paul J. Foster. Retinal Vasculometry Associations With Glaucoma: Findings From the European Prospective Investigation of Cancer–Norfolk Eye Study. *American Journal of Ophthalmology*, 220:140–151, 2020.

[SC04] Vin de Silva and Gunnar Carlsson. Topological estimation using witness complexes. In Markus Gross, Hanspeter Pfister, Marc Alexa, and Szymon Rusinkiewicz, editors, *SPBG'04 Symposium on Point - Based Graphics 2004*. The Eurographics Association, 2004.

[SG07] Vin Silva and Robert Ghrist. Coverage in sensor networks via persistent homology. *Algebraic & Geometric Topology*, 7, 04 2007.

[Shi05] Masahiro Shiota. Whitney triangulations of semialgebraic sets. *Ann. Polon. Math.*, 87:237–246, 2005.

[Sie72] Laurence Carl Siebenmann. Deformation of homeomorphisms on stratified sets. I, II. *Comment. Math. Helv.*, 47:123–136; ibid. 47 (1972), 137–163, 1972.

[SMVJ11] Vin Silva, Dmitriy Morozov, and Mikael Vejdemo-Johansson. Dualities in persistent (co)homology. *Inverse Problems - INVERSE PROBL*, 27, 07 2011.

[Ste80] G. W. Stewart. The efficient generation of random orthogonal matrices with an application to condition estimators. *SIAM J. Numer. Anal.*, 17(3):403–409 (loose microfiche suppl.), 1980.

[STHN20] Bernadette J. Stolz, Jared Tanner, Heather A. Harrington, and Vidit Nanda. Geometric anomaly detection in data. *Proc. Natl. Acad. Sci. USA*, 117(33):19664–19669, 2020.

[SW14] Primoz Skraba and Bei Wang. Approximating local homology from samples. In *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 174–192, New York, NY, USA, 2014. Association for Computing Machinery.

[TH15] Abdel Aziz Taha and Allan Hanbury. An efficient algorithm for calculating the exact hausdorff distance. In *IEEE Trans Pattern Anal Mach Intell*, 2015.

[Tho69] René Thom. Ensembles et morphismes stratifiés. *Bull. Amer. Math. Soc.*, 75:240–284, 1969.

[vdD86] Lou van den Dries. A generalization of the Tarski-Seidenberg theorem, and some nondefinability results. *Bull. Amer. Math. Soc. (N.S.)*, 15(2):189–193, 1986.

[vdD98] Lou van den Dries. *Tame topology and o-minimal structures*, volume 248 of *London Mathematical Society Lecture Note Series*. Cambridge University Press, Cambridge, UK, 1998.

[vHvdS98] Johannes Hendrik van Hateren and Arjen van der Schaaf. Independent component filters of natural images compared with simple cells in primary visual cortex. *Proceedings: Biological Sciences*, 265(1394):359–366, 1998.

[VNG20] Robin Vandaele, Guillaume Adrien Nervo, and Olivier Gevaert. Topological image modification for object detection and topological image processing of skin lesions. *Scientific Reports*, 10(21061), 2020.

[Whi65a] Hassler Whitney. Local properties of analytic varieties. In *Differential and Combinatorial Topology (A Symposium in Honor of Marston Morse)*, pages 205–244. Princeton Univ. Press, Princeton, NJ, USA, 1965.

[Whi65b] Hassler Whitney. Tangents to an analytic variety. *Ann. of Math. (2)*, 81:496–549, 1965.

[Woo09] Jon Woolf. The fundamental category of a stratified space. *J. Homotopy Relat. Struct.*, 4(1):359–387, 2009.

[WY13] Zaiwen Wen and Wotao Yin. A feasible method for optimization with orthogonality constraints. *Math. Program.*, 142(1-2):397–434, 2013.

[Xia16] Shengxiang Xia. A topological analysis of high-contrast patches in natural images. *The Journal of Nonlinear Sciences and Applications*, 09:126–138, 2016.

[ZC05] Afra Zomorodian and Gunnar Carlsson. Computing persistent homology. *Discrete Comput. Geom.*, 33(2):249–274, 2005.

[Zha10] Gongyun Zhao. Representing the space of linear programs as the Grassmann manifold. *Math. Program.*, 121(2):353–386, 2010.

[Zom10] Afra Zomorodian. Fast construction of the vietoris-rips complex. *Computers & Graphics*, 34(3):263–271, 2010. Shape Modelling International (SMI) Conference 2010.

Appendix A

Further results on Whitney stratified spaces and definability

This appendix presents supplementary technical proofs that were intentionally omitted from the main body of this thesis. The content herein consists of statements and proofs previously featured in a collaborative paper authored by the present writer and Lukas Waas [MW22], with the content reproduced in a largely verbatim manner.

A.1 Stability result for Whitney stratified spaces

Proposition A.1.1. *[Hir69, Proof of 4.1.1] Let W be a Whitney stratified space over P and let $\phi\colon [0,d] \to W$ be the integral curve associated to $x \in W_q$, $y \in W_p$, $q \geq p \in P$, with notation as in Proposition 3.3.4. Then ϕ has the following properties.*

1. *$||\phi(t) - y|| = t$, for $t \in [0,d]$.*
2. *$||\phi(t) - \phi(t')|| \leq \frac{1}{\sqrt{1-\delta^2}}|t - t'|$, for $t, t' \in [0,d]$.*

As a consequence of this result, the continuity result of Theorem 2.2.31 can be improved to Lipschitz continuity.

Proposition A.1.2. *Let $P = \{p < q\}$ and let $W \in \mathbf{Sam}_P$ be a Whitney stratified space with compact singular stratum W_p. Then, for any $C > 1$, there exists an $R > 0$, such that the function*

$$\begin{aligned} \Omega \cap (0,R)^2 &\to \mathbf{D}_P\mathbf{Sam} \\ v &\mapsto \mathcal{D}_v(\mathcal{N}(W)) \end{aligned}$$

is C-Lipschitz continuous.

Proof. We omit the $\mathcal{N}$, to keep notation concise. By Lemma 2.2.28, it again suffices to consider the link part of the diagrams given by $W^{v_l}_{v_h}$. Choose $\delta < 1$ such that $\frac{1}{\sqrt{1-\delta^2}} < C$. Next, take R small enough such that $\mathrm{N}_R(W_p)$, with retraction $r\colon \mathrm{N}_R(W_p) \to W_p$ is a standard tubular neighborhood of W_p. By [NV21, Lemma 2.1], for R small enough the spaces $W^y = r^{-1}(y) \cap W$ of y are given by Whitney stratified spaces with singular stratum given by a point. Then, using Construction A.3.1, we may also choose R so small, that

$$\beta(x, y) \leq \delta,$$

for the respective β on the fiber W^y. Now, let $v, v' \in \Omega \cap [0, R]$. Let $x \in W^{v_l}_{v_h}$ and assume that $v_h > v'_h$ (the other cases work similarly). Now, consider the integral curve ϕ from $y := r(x) \in W_p$ to x in $r^{-1}(y) \cap W$. By Proposition A.1.1 we have,

$$|x - \phi(v'_h)| = |\phi(|x|) - \phi(v'_h)| \leq C||x| - v'_h| \leq C|v_h - v'_h| \leq C|v - v'|.$$

Since $\phi(v'_h) \in W^{v'_l}_{v'_h}$, going through all the cases, we obtain

$$W^{v_l}_{v_h} \subset (W^{v'_l}_{v'_h})_{C|v-v'|}.$$

Thus, the result follows by symmetry. □

We thus obtain, as a corollary of Theorem 2.2.31, that for v sufficiently small the persistent stratified homotopy type $\mathcal{P}_v$ is even Lipschitz continuous at a Whitney stratified space.

Theorem A.1.3. *Let $P = \{p < q\}$ and $W \in \mathbf{Sam}_P$ be Whitney stratified with W_p compact. Then, for any $C > 1$, there exists some $R > 0$, such that the map*

$$\mathcal{P}_v : \mathbf{Sam}_P \to \mathrm{ho}\mathbf{Strat}^{\mathbb{R}_+}$$

is $2(C+1)$-Lipschitz continuous at W, for all $v \in \Omega \cap (0, R)^2$.

A.2 Proof of Proposition 3.3.3

Proof of Proposition 3.3.3. The mapping β clearly exhibits continuity on $S_q \times S_p$. The condition on β is thus tantamount to the extension by 0 to Δ_{S_p} being continuous. Indeed, due to the continuity of $\vec{d}(-,-)$, this extension condition immediately implies condition (b). For the converse, given that $\beta \geq 0$, it suffices to demonstrate upper semi-continuity. This is the essence of Proposition A.2.1. □

Proposition A.2.1. *Let $W = (X, s : X \to P)$ be a Whitney stratified space. Then, the restriction of β to $W_{\geq p} \times W_p \to \mathbb{R}$ is upper semi-continuous.*

Proof. The function β is evidently continuous on the strata of $W \times W$. Now, suppose $(x_n, y_n) \in W_{\geq p} \times W_p$ is a sequence converging to a point $(x, y) \in W_{p'} \times W_p$ for some $p' \geq p$. Then, for sufficiently large $n \in \mathbb{N}$, we have $s(x_n) \geq p'$. To demonstrate upper semi-continuity, we may, without loss of generality, assume that x_n lies in the same stratum W_q. We show that any subsequence of (x_n, y_n) has a further subsequence (all named the same by abuse of notation), for which $\beta(x_n, y_n)$ converges to a value less than or equal to $\beta(x, y)$.

By compactness of Grassmannians, we may first restrict to a subsequence such that $\mathrm{T}_{x_n}(W_q)$ and $l(x_n, y_n)$ converge to τ and l respectively. By Whitney's condition (a) [Whi65a,Whi65b] - which, by [Mat12], follows from condition (b) - we have $\mathrm{T}_x(W_{p'}) \subset \tau$. Summarizing, this gives:

$$\lim \beta(x_n, y_n) = \vec{d}(l, \tau) \leq \vec{d}(l, \mathrm{T}_x(W_{p'})).$$

Now, in the case when $x \neq y$, the last expression equals $\beta(x, y)$ by definition. In the case when $x = y$, then, by condition (b), $l \subset \tau$. Thus, again, we have:

$$\lim \beta(x_n, y_n) = \vec{d}(l, \tau) = 0 = \beta(y, y)$$

finishing the proof. □

A.3 A normal bundle version of β

Furthermore, we are going to make use of the following fiberwise version of β.

Construction A.3.1. *In the context of Construction 3.3.2, assume that $W = (X, s\colon X \to P)$ is a Whitney stratified space, with W_p compact. Take N to be a standard tubular neighborhood of W_p in $\mathbb{R}^N$ with retraction $r : N \to W_p$. Note that by Whitney's condition (a), for N sufficiently small, $r_{|W_q}$ is a submersion for $q \geq p$. In particular, by [NV21, Lemma 2.1], the fiber of*

$$W^y := (r)^{-1}_{|N \cap W^{\geq p}}(y)$$

is a Whitney stratified space over $\{q \in P \mid q \geq p\}$ with the p-stratum given by $\{y\}$. Furthermore, we have

$$\mathrm{T}_x(W_q) \cap \nu_{r(x)}(W_p) = \mathrm{T}_x(W_q^{r(x)}),$$

where $\nu_{r(x)}(W_p)$ denotes the normal space of W_p at $r(x)$. In particular, the dimension of these spaces is constant, and they vary continuously in x. Then, consider the

following function:

$$\tilde{\beta}_p(-) : N \cap W^{\geq p} \to \mathbb{R} \qquad \begin{cases} x & \mapsto \vec{d}(l(x, r(x)), \mathrm{T}_x(W^{r(x)}_{s(x)})), \text{ for } s(x) > p \\ x & \mapsto 0, \text{ for } s(x) = p. \end{cases}$$

Noting that $l(x, r(x)) \in \nu_{r(x)}(W_p)$, by an analogous argument to the proof of Proposition 3.3.3, one obtains that $\tilde{\beta}_p(-)$ is continuous on $W_q \cup W_p$. Note that if we restrict $\tilde{\beta}_p(-)$ to W^y, then we obtain the function $\beta(-, y)$ associated with W^y. Let us denote this β_y. In particular, by compactness of W_p, we obtain that the functions β_y can be globally bounded by any $\delta > 0$, for N sufficiently small.

A.4 Definability of β

Proposition A.4.1. *Let $S = (X, s\colon X \to P)$ be as in Construction 3.3.2. Then, if $X \subset \mathbb{R}^N$ is definable, then so is β.*

Proof. As all the strata of $X \times X$ are again definable, it suffices to show that β is definable on the strata of $X \times X$. Furthermore, as β is 0 along Δ_X, it suffices to show definability away from the diagonal. Here β is equivalently given by

$$\beta(x, y) = \inf_{v \in \mathrm{T}_x(X_{s(x)})} ||\frac{x - y}{||x - y||} - v||.$$

It follows from the fact that for $q \in P$, $\mathrm{T}(X_q) \subset \mathbb{R}^N \times \mathbb{R}^N$ is definable (see [Cos00] and Lemma A.5.1) that this defines a definable function $X_q \times X_p \to \mathbb{R}$. □

A.5 Proof of Proposition 3.3.9

We begin by proving a series of technical lemmas.

Lemma A.5.1. *Consider two definable maps $f : X \to \mathbb{R}$, $\pi : X \to Y$ such that f is bounded from above on every fiber of π. Then the map*

$$\begin{aligned} g : Y &\to \mathbb{R} \\ y &\mapsto \sup_{x \in \pi^{-1}(y)} f(x) \end{aligned}$$

is again definable.

Proof. This is immediate if one interprets the graph of g in terms of a formula being expressible with respect to the o-minimal structure. □

Lemma A.5.2. *Let $X \to \{p < q\}$ be a stratified metric space and Y a first countable, locally compact Hausdorff space. Let $\pi : X \to Y$ be a proper map, such that both the fibers of π, as well as the fibers of $\pi_{|X_p}$ vary continuously in the Hausdorff distance. Let $f : X \to \mathbb{R}$ be upper semi-continuous and continuous on the strata. Then,*

$$\begin{aligned} g : Y &\to \mathbb{R} \\ y &\mapsto \sup_{x \in \pi^{-1}(y)} f(x) \end{aligned}$$

is continuous.

Proof. Note first that as the fibers of π are compact and f is upper semi-continuous, it takes its maximum on every fiber. Now, let $y_n \to y$ be a convergent sequence in Y. We show that any of its subsequences y'_n, has a further subsequence $\tilde{y}_n \to y$, with

$$\sup_{x \in \pi^{-1}(\tilde{y}_n)} f(x) \to \sup_{x \in \pi^{-1}(y)} f(x).$$

Let $x'_n \in \pi^{-1}(y_n)$ for all n such that $f(x'_n) = \sup_{x \in \pi^{-1}(y'_n)} f(x)$. As Y is locally compact and π is proper, x'_n has a convergent subsequence $\tilde{x}_n \to \tilde{x}$. Define $\tilde{y}_n := \pi(\tilde{x}_n)$. Since the fibers of π vary continuously and $\tilde{y}_n \to y$, we also have $\tilde{x} \in \pi^{-1}(y)$. Thus, we have

$$\limsup \sup_{x \in \pi^{-1}(\tilde{y}_n)} f(x) = \limsup f(\tilde{x}_n) \leq f(\tilde{x}) \leq g(y).$$

It remains to see the converse inequality for a subsequence of $\tilde{y}_n$. Let $\hat{x} \in \pi^{-1}(y)$ be such that $f(\hat{x}) = \sup_{x \in \pi^{-1}(y)} f(x)$. By assumption we can find a sequence x''_n with $x''_n \in \pi^{-1}(\tilde{y}_n)$ converging to $\hat{x}$. If $\hat{x} \in X_p$, then x''_n can be taken to be in X_p, as $\pi^{-1}(\tilde{y}_n) \cap X_p$ converges to $\pi^{-1}(y) \cap X_p$. If $\hat{x} \in X_q$, then, as the latter is open, x''_n ultimately lies in X_q. Hence, by continuity of f on the strata, we have

$$g(y) = f(\hat{x}) = \lim f(x''_n) = \liminf f(x''_n) \leq \liminf \sup_{x \in \pi^{-1}(\tilde{y}_n)} f(x).$$

□

As a consequence of the prior two lemmas, we obtain:

Lemma A.5.3. *If W is a definably Whitney stratified over $P = \{p < q\}$. Then the map*

$$\begin{aligned} \hat{\beta} : W_p \times \mathbb{R}_{\geq 0} &\to \mathbb{R} \\ (y, d) &\mapsto \sup_{||x-y||=d, x \in W} \beta(x, y) \end{aligned}$$

is continuous in a neighborhood of $W_p \times \{0\}$, definable and vanishes on $W_p \times \{0\}$.

Proof. The proof of definability immediately follows from Lemma A.5.1. Consider the map

$$B : W \times W_p \to W_p \times \mathbb{R}_{\geq 0} \ (x, y) \mapsto (y, ||x - y||).$$

Over $W_p \times \mathbb{R}_{>0}$, it is given by submersion on each stratum of $W \times W_p$. Specifically, by Thom's first isotopy lemma [Mat12, Proposition 11.1], it is a fiber bundle with fibers $\partial \mathrm{B}_d(y)$ at (y, d) over $\mathbb{R}_{>0}$. Moreover, the fibers of B vary continuously over $W_p \times \mathbb{R}_{>0}$. Additionally, for $(y_n, d_n) \to (y, 0)$, the fiber converges to the point y. Thus, B fulfills the requirements of Lemma A.5.2. Furthermore, $\beta : W \times W_p \to \mathbb{R}$ also satisfies the conditions of Lemma A.5.2, demonstrating the continuity of $\hat{\beta}$. Lastly, $\hat{\beta}$ vanishes on $W_p \times \mathbb{R}_{\leq 0}$ by the definition of β. □

We now possess all the tools required to establish a proof of Proposition 3.3.9,

Proof of Proposition 3.3.9. We conduct this proof for the case of $P = p < q$ and $K = W_p$ (with notation as in Definition 3.3.7). The general case follows analogously by working strata-wise and then passing to maxima. By Lemma A.5.3, for d small enough, the function $\hat{\beta} : W_p \times \mathbb{R}_{\geq 0} \to \mathbb{R}$ fulfills the requirements of Lojasiewicz' theorem for (polynomially bounded) o-minimal structures [Loi16]. Hence, we find $\hat{\phi} : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ to be a definable and monotonous bijection such that on $W_p \times [0, d]$ we have

$$\hat{\phi}(\hat{\beta}(y, t)) \leq t.$$

If the relevant o-minimal structure is polynomially bounded, then there exist $n > 0$, such that

$$t^n \leq \hat{\phi}(t)$$

for $t \in [0, d']$. Hence, we obtain

$$\begin{aligned} \hat{\beta}(y, t)^n \leq \hat{\phi}(\hat{\beta}(y, t)) \leq t. \\ \implies \hat{\beta}(y, t) \leq t^\alpha \end{aligned}$$

for $t \in [0, d]$, $\alpha = \frac{1}{n}$ and $d := \phi^{-1}(d')$. □