Geometry of Stratified Spaces for the Analysis of Complex Data

by

Shreya Arya

Department of Mathematics
Duke University

Defense Date: April 3, 2024

Approved:

Sayan Mukherjee, Supervisor

Ezra Miller, Supervisor

Kirsten Wickelgren

Katharine Turner

Dissertation submitted in partial fulfillment of the requirements for the degree of
Doctor of Philosophy in the Department of Mathematics
in the Graduate School of Duke University
2024

ABSTRACT

Geometry of Stratified Spaces for the Analysis of Complex Data

by

Shreya Arya

Department of Mathematics
Duke University

Defense Date: April 3, 2024

Approved:

Sayan Mukherjee, Supervisor

Ezra Miller, Supervisor

Kirsten Wickelgren

Katharine Turner

An abstract of a dissertation submitted in partial fulfillment of the requirements for the degree of Doctor of Philosophy in the Department of Mathematics in the Graduate School of Duke University
2024

Abstract

Traditionally, statistical data has been in the form of elements of Euclidean space. However, as data complexity increases, it is assumed to lie on lower dimensional non-linear space such as smooth manifolds, in what is known as the manifold hypothesis. Nonetheless, real-world data is not always in the form of smooth manifolds but, in general, can lie on stratified spaces. In this thesis, we explore the geometry of stratified spaces with the overall objective of enabling statistics on these spaces. More specifically, we provide answers to the following two problems.

A fundamental task in object recognition is to identify when two shapes are similar. One approach to rendering this as a precise mathematical problem is to look at the space of all shapes and define a metric on it. This approach has been taken by renowned statisticians and mathematicians like Kendall, Grenander, Mumford, Michor, and others. In this, we provide an algebraic construction of the moduli space of shapes and define metrics on it with the objective of developing a statistical theory on shapes. The construction is far more general than existing constructions, as it doesn't restrict 'shapes' to smooth manifolds and includes a broad category of spaces, including many stratified spaces. The foundation of this construction relies on the topological analogue of the Radon transform, building on the work of Schapira who showed that such transforms are injective.

This thesis also provides a starting point for developing a theory of diffusion processes on general stratified spaces. On Euclidean spaces, Brownian motion is constructed by taking scaled limits of random walks. This approach is challenging because stratified spaces are not only non-linear and lack addition but also the tangent spaces of stratified spaces are non-linear, unlike smooth manifolds. So, instead, we define Brownian motion on stratified spaces by taking appropriate limits of Dirichlet forms. Sturm took this approach for general metric measure spaces, where he came up with a measure-theoretic condition required for these Dirichlet forms to converge properly. We prove this is the case for certain compact subanalytic spaces.

Parts of the thesis are based on joint work with Justin Curry and Sayan Mukherjee.

Dedication

To mama, papa and didi.

Contents

List of Figures

Acknowledgements

First of all, I would like to thank my advisors, Sayan Mukherjee and Ezra Miller. I would also like to thank Justin Curry, whose guidance and support were akin to that of an advisor. I could not have asked for a more supportive team of advisors, and none of this would have been possible without the opportunities they provided. Sayan has always been inspiring, generous, and has consistently believed in me. His insight in identifying mathematical problems that have significant statistical consequences is remarkable. Ezra has deeply influenced the way I think about mathematics. I have learned so much by observing him function as a mathematician. I cherish all the mathematical conversations we've had in the past; his geometric way of thinking, combined with his exceptional exposition skills, continue to inspire me. Justin not only taught me a wealth of mathematics but also made research a lot of fun and exciting. I thank him for all the guidance, mathematical and otherwise.

I am deeply grateful to the Mathematics Department at Duke University for providing me with a nurturing and supportive environment, truly making it feel like home. I would like to thank several faculty members at Duke from whom I learned a lot: Jonathan Mattingly, Chad Schoen, Colleen Robles, and Kirsten Wickelgren. I thank Dorothy Buck for all the support and mentorship over the years and Sarah Schott for helping me improve as a teacher.

I also express gratitude to Facundo Mémoli, Paige North, Markus Pflaum, and Kate Turner for their guidance in various areas of mathematics and for the opportunities they provided to collaborate. I thank postdocs Ranthony Clark, Marzieh Eidi, Celia Hacker, and Ling Zhou for all their support and guidance.

I have thoroughly enjoyed my time as a graduate student at Duke, and one of the main reasons is my wonderful cohort mates: Aygul, Ching-Lung, Chun-Hsien, Michael, Yupeng, Yuqing, Yixin and Yixuan with whom I've shared countless memorable moments while navigating the ups and downs of graduate school together. Special thanks to Chun-Hsien for his invaluable support during the final years of the PhD.

I extend my heartfelt thanks to my friends in Durham, who have become like family to me: Achint, Akanksha, Andrea, Anurag, Anvita, Kahini, Lalit, Niven, Rohit, Shiv, Shreya Maini and especially my roommates Devang and Harshit, as well as Aygul, for our late-night crossword puzzles, movies and unwavering friendship, making my time in Durham truly memorable. I am also thankful to Devang and Shrigyan for always being ready for squash matches at a moment's notice.

I thank my best friends from undergrad Devashi for the weekly marathon conversations, Manu and Yash Pandya for all the laughter, trips and mathematical discussions over chai, and Ishani, Nandan, Urvashi, and Yash Mehta for always being there.

To my partner, Shlok, and all my friends in Waterloo, thank you for providing me with a second home. Finally and most importantly, I would like to thank my parents, my sister, and brother-in-law, and my nephew for all their love and support.

1. Introduction

Traditionally, statistical data has been in the form of elements of Euclidean space. The collection of k data points typically looks like,

$$\{x_1, x_2, \cdots, x_k\} \subset \mathbb{R}^n.$$

It is very reasonable to assume $\mathbb{R}^n$-valued data, as it greatly simplifies statistical analysis. Firstly, the linear structure of $\mathbb{R}^n$ makes it possible to add data points, allowing for the computation of sample or population means. Additionally, random walks can be defined by adding independent and identically distributed (i.i.d.) random variables, and by taking scaled limits of random walks, we obtain Brownian motion. Brownian motion or more generally diffusion can be used to model the time evolution of data, generate random samples from complex probability distributions, and enable statistical inference. Secondly, distances between two data points can be easily measured by, for instance, the Euclidean distance. This helps quantify the similarity between two data points and is very useful in data analysis, for example, in clustering and segmentation.

With increasing technological advances and the *datafication* of everything, data is becoming increasingly complex, not only in terms of size but also structure. Assuming that data points are $\mathbb{R}^n$-valued is overly simplistic and not feasible. For example, data objects could be 3D scans from medical imaging. This forms the basis for *object-oriented data analysis* [MA14; MD21], where the central question posed is "What are the atoms of statistical analysis?"

More recently, it is assumed that data is non-linear and can be represented as

$$\{x_1, x_2, \cdots, x_k\} \subset M \subset \mathbb{R}^n,$$

where M is a Riemannian manifold, in what is known as the *manifold hypothesis*. For example, data points could lie on the sphere. The sphere is a non-linear space: adding two points on the sphere may result in a point lying outside of the sphere. The manifold assumption takes this non-linearity into account. The manifold assumption also reduces

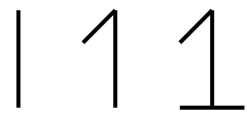

Figure 1.1: Variations of 1s, each constituting manifolds of different dimension.

the degrees of freedom of the system making computations more tractable.[1]

Since Riemannian manifolds locally look like Euclidean space, one would expect much of the theory applicable to Euclidean spaces to apply to the case of manifolds. The distances between manifold-valued data points can be measured by the geodesic distance induced by the Riemannian metric. So, there are well-developed tools for statistical analysis and inference on manifolds. For example, linear sample means can be replaced by Fréchet means on manifolds [Fré48], principal component analysis (PCA) becomes principal geodesic analysis [FLJ03; Fle+04], central limit theorems become central limit theorems for Fréchet means [BP03a; BP03b] and stochastic processes on $\mathbb{R}^n$ have an analogous counterpart on manifolds [Hsu02].

The assumption that data points lie on a smooth manifold offers significant generality. Even if there are some singular or non-smooth points like corners or edges, smooth manifolds are reasonable approximations of such spaces at generic (or smooth points). However, the problem occurs when the typical or average behavior of data points lies nearby singular points. This may occur in many spaces of interest, and so such an assumption is unreasonable *in a global sense*. For example, the space of handwritten digits is not a smooth manifold[2]. For simplicity, let's consider the space of "1s." In Figure 1.1, the first "1" can be described by its vertical line length, while the second one requires additional parameters for the angle and length of the top hat. The space of "1s" consists of different manifolds of

[1] This thesis does not pertain to manifold learning, where the objective is to find a low-dimensional data representation for high-dimensional Euclidean-valued data. Instead, we assume to work with manifold-valued data (or more generally data valued in stratified spaces) and aim to develop theory to enable statistical inference on these spaces.

[2] One of the most fundamental tasks in machine learning is to classify handwritten digits from the MNIST dataset.

varying dimensions glued together—a stratified space [Mil23]. So, if the average behavior of data points (for example, the mean of data points) congregates at the intersections of strata, then assuming that the space of "1s" lies on smooth manifolds may lead to incorrect statistical inferences. Therefore, understanding the geometry of data sampled nearby singular points is important.

So, we assume that data points are represented as

$$\{x_1, x_2, \cdots, x_k\} \subset X \subset \mathbb{R}^n,$$

where X is a stratified space. Roughly speaking, a stratified space $X \subset \mathbb{R}^n$ is a union of smooth submanifolds (called strata) of varying dimensions that are glued together in a consistent way.

Stratified spaces are spaces with *singularities* or non-smooth parts, with perhaps the most well-known example being black holes, which are singularities in space-time. However, one does not need to look so far to find examples of stratified spaces; simplicial complexes are stratified spaces. MacPherson, one of the inventors of intersection (co)homology of singular spaces, described singular spaces as spaces where the normal rules of algebra, geometry, and calculus break down. He gives an example of the intersection point of racetrack shaped like the numeral '8', where a fast racing car can drive full speed at all parts of the racetrack but at the intersection point will have to change the rules, slow down, and look in all directions for traffic.

1.1 Examples and Applications

Stratified spaces are ubiquitous, and below are some examples of stratified spaces.

Graphs and simplicial complexes

If X is an abstract simplicial complex, then the relatively open simplices define a stratification of its geometric realization.

Algebraic varieties

The set of zeroes of a finite set of polynomials is a stratified space. Statistical machines (mathematical objects or structures that can be described using statistical methods) like the parameter space of artificial neural networks, or more generally feasible solutions of optimization problems can be thought of as a semi-algebraic variety. The configuration space of a robot system can also be described by an semi-algebraic variety.

Moduli spaces

Moduli spaces are spaces where every point represents a distinct class of geometric objects of a fixed type. They often are singular spaces. Consider the moduli spaces of phylogenetic trees, as given in [BHV01], commonly known as BHV tree space. The space of such trees is used are used to track evolutionary changes in genetics, or represent biological pathways and are stratified spaces. The space of positive semi-definite matrices of dimension n can be stratified with each stratum of dimension r corresponding to positive semi-definite matrices of rank r for $0 \leqslant r \leqslant n$.

Orbit space of a smooth manifold

Let M be a smooth manifold and G be a compact Lie group. Suppose that G acts smoothly on M, then the orbit space M/G is a stratified space. More precisely the strata are

$$M_{(H)} = \{x \in M \mid G_x \text{ is conjuagte to } H\}$$

for every closed subgroup H of G where G_x is the is isotropy subgroup.

Orbit spaces of smooth manifolds can be used to represent data that exhibit symmetries. It is becoming increasingly common to do analysis on spaces that respect symmetries, for example in areas such as equivariant machine learning. Below are some examples of orbit spaces.

Example 1.1.1. (Sample space) The sample spaces of size k of a smooth manifold M

are orbit spaces under the action of the symmetric group S_k. Consider the set of k-point samples of M, denoted by M^k, i.e. each element x in M^k takes the form $x = (x_1, x_2, \cdots, x_k)$ where each $x_i \in M$. The orbit space M^k/S_k is a stratified space where the regular stratum is the set of points that are mutually distinct.

Example 1.1.2 (Shape spaces)**.** Shape spaces were pioneered in the works of Kendall [Ken77; Ken84] and Bookstein [Boo92] for the purpose of shape comparison and more generally statistical shape analysis. Each shape is defined by a set of k landmark points in $\mathbb{R}^n$, where typically $n = 2, 3$. The shape space is the quotient of the space of configurations of k landmark points by the groups of translations, re-scaling, and rotations of $\mathbb{R}^n$. The main point here is that shapes are invariant under rigid motions, for instance, a 90^o rotation of an image of the cat is still a cat and a metric defined on such a space should ensure that the distance between two such cat images is 0.

1.2 Overview

This thesis explores the geometry of data valued in stratified spaces with the overall objective of enabling statistics on these spaces. More specifically we, 1) provide an algebraic model of shape space (or the space of stratified spaces with the objective of shape analysis and comparison) and 2) provide a method to define Brownian motion on stratified spaces.

At a high level we are concerned with the following questions.

1. How do we represent stratified spaces? A computer scientist may rephrase this question as: What is a good data structure to represent such a space? What are the desirable properties of such a representation?

2. How do we compare such spaces to one another? A human eye is capable of telling when two shapes are similar, but how can we quantify this similarity?

3. How do we randomly sample from a stratified space? How can we describe distributions on stratified spaces? For this, we resort to diffusion processes, where the

transition kernels of such processes can serve as the probability density functions of the distributions.

Aside from the statistical interest, Item 3, i.e., diffusion on stratified spaces, can also be used to solve problems in geometry. Brownian motion and probabilistic methods inform the geometry of the space (see, for example, the probabilistic proof of the Atiyah-Singer index theorem [Hsu87]).

The next section discusses items 1 and 2 in more detail. Item 3 is discussed in section 1.4.

1.3 Shape space

Shape space has been a fundamental object of study in geometric morphometrics with countless applications to biology, biomedical imaging, and evolutionary anthropology (for review articles see [AO13], [Hal04], [MS22] and for a detailed list of application areas see [SK16, Section 1.2]). Shape spaces are not only intended to provide a single framework for comparing shapes but also allow for a mathematical and computational representation of shapes that can be used for statistical analysis. Different shapes are rendered as different points in shape space, and comparisons of shapes can be formalized in terms of distances between points.

There is no universally accepted mathematical definition of a shape. There have been several attempts to precisely define what shapes and their corresponding shape space are depending on context and applications. Historically, a shape has been defined as a collection of a fixed number of landmark points, in what is known as Kendall shape spaces and is denoted by

$$\Sigma_d^k := \{(\mathbb{R}^d)^{k-1} \backslash 0\}/\mathrm{Sim},$$

where Sim is the group of rotations and dilations and k denotes the number of landmarks. Each shape is defined by k points and the i-th point in one shape corresponds to the i-th point in every other shape in the space; this introduces the central notion of correspondences in shape space. Note that in Σ_d^k a shape is reduced to a $d \times k$ matrix, which is a very convenient representation. However, the downside of this approach is that a user will need to decide on landmarks before analysis can be carried out, and reducing modern databases of 3-dimensional micro-computed tomography (CT) scans [Gos15; Boy+16] to landmarks can result in a great deal of information loss.

There is another commonly accepted shape space, which is due to Grenander, and originates from his works on Pattern theory, although some aspects were anticipated by [CMM91]. In these works, a shape space is specified for each manifold M and dimension d. One then considers all possible immersions modulo the group of reparameterizations of M, i.e.

$$\mathrm{Shape}(M) := \mathrm{Imm}(M, \mathbb{R}^d)/\mathrm{Diff}(M).$$

Variation in shape is then modeled by the action of the Lie group of diffeomorphisms on $\mathbb{R}^d$. The advantage of Grenander's approach (or more generally known as the diffeomorphism-based approach) is that it avoids the need for landmarks, but the resulting spaces of interest are infinite-dimensional and shapes with different topology cannot be compared. However, many tools have been developed that efficiently compare the similarity between shapes in large databases via algorithms that continuously deform one shape into another [Boy+11; Ovs+12; Boy+15; GKD19].

1.3.1 The Shapes

Our definition of shape extends beyond smooth manifolds to include a large class of stratified spaces. Shapes found in nature often exhibit wondrous complexity, such as fractals. However, we make the assumption that these shapes adhere to the principles of Grothendieck's programme on tame topology [Gro97]. Essentially, this means excluding pathological examples like space-filling curves and Cantor-like sets. We work with stratified spaces that are triangulable and hence can be faithfully represented via a mesh on a computer. More precisely, a shape is represented by a compact definable set $M \subset \mathbb{R}^d$ (or a constructible set) with respect to a fixed o-minimal structure $\mathcal{O}$ on $\mathbb{R}^d$.

By considering shapes to be this general, we lose the one-one correspondence maps between any two shapes. For instance, in the landmark-based approach, the correspondence maps are mappings between landmarks, and in the diffeomorphism-based approach, the correspondence maps are diffeomorphisms. These maps were key in constructing metrics on these shape spaces, thereby making our task much more challenging.

1.3.2 The Representation

Once we have a precise definition of a shape, the next question to consider is how we should represent our shapes. More importantly, what properties must this representation satisfy? We must ensure that the representation is faithful or injective, as we do not want identical representations corresponding to different shapes. Secondly, we would want the representation to be equivariant with respect to rigid motions, i.e., the action of rigid motions such as translations and rotations on the input should produce corresponding rigid motions of the output. We build upon the fundamental work of Schapira [Sch91b; Sch95b], selecting our representatives as the two topological transforms—the Euler characteristic transform (ECT) and the persistent homology transform (PHT)—introduced in [TMB14] to facilitate comparison of non-diffeomorphic shapes. The ECT and PHT have two useful properties: standard statistical methods can be applied to the transformed shape and the transforms are injective [GLM18; CMT22b], so no information about the shape is lost via the transform. The utility of the transforms for applied problems in evolutionary anthropology, biomedical applications and plant biology were demonstrated in [Cra+20; Wan+21; Tan+22; Amé+22].

In 1917, Radon [Rad05] observed that differentiable functions in $\mathbb{R}^3$ could be explicitly determined by its integrals over all planes in $\mathbb{R}^3$. This was essentially the basis for the Radon transform, which has wide applications in computed tomography (CT), X-ray imaging, and related fields [Cor63; Cor64]. To illustrate the application of the Radon transform in tomography, consider the following scenario: Consider a convex body X with density

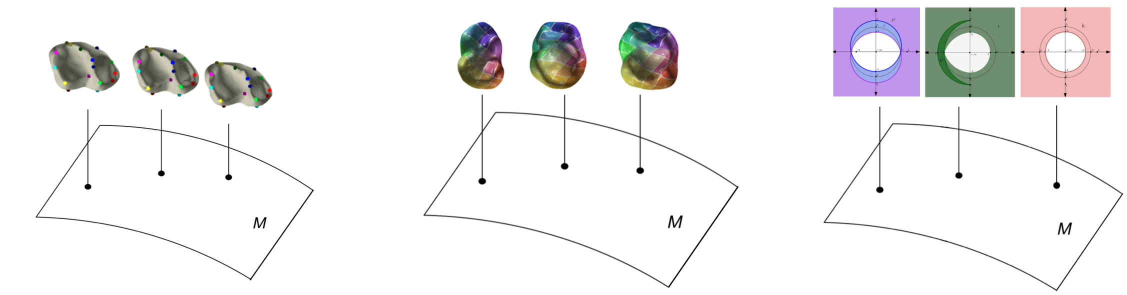

(a) Landmark-based Shape Space

(b) Diffeomorphism-based Shape Space

(c) PHT-based Shape Space

FIGURE 1.2: Previous constructions of shape space imply a fiber bundle perspective on shape space. The fibers encode similarity transformations or reparameterizations of a shape and the base space records the shape as an equivalence class. The PHT-based shape space introduced here uses a more algebraic construction where the base space is replaced by a base poset and the fibers are unique sheaf-theoretic representations of the shape.

function f through which a beam of X-rays is passed through a specific direction along line L. Cormack and Hounsfield observed that the line integral could be obtained through X-ray measurements, specifically as $\log(I_1/I_2)$, where I_1 and I_2 represent the intensities of the X-ray beam before and after interacting with the object X respectively. By applying Radon's observation, we can now determine the density completely. The idea of topological transforms stems from a similar concept: knowing the topological invariants such as the Euler characteristic or homology groups of shapes filtered in different directions can be used to reconstruct the shape. As one would expect, when the topological invariant is the Euler characteristic, we obtain the Euler characteristic transform (ECT), and when the invariant is the homology groups, we obtain the Persistent Homology Transform (PHT).

1.3.3 Our Construction

Now, we are ready to construct a truly general shape space using the topological transforms discussed above. In the case of previous models of shape space, the shape space can be thought of as a fiber bundle due to the presence of one-one correspondence maps between shapes. However, since we don't make the assumption that shapes are in one-one correspondence to each other, we transition from the land of fiber bundles to the world of *sheaves*, which replaces the local triviality condition of fiber bundles with the local continuity condition of sheaves (refer to Figure 1.2). This passage requires two preparatory steps of categorical generalization:

1. Instead of a "base manifold" of shapes we work with a "base poset" of constructible sets $\mathcal{CS}(\mathbb{R}^d)$ ordered by inclusion. This poset is equipped with a notion of continuity via a Grothendieck topology.

2. Each shape—that is, each point $M \in \mathcal{CS}(\mathbb{R}^d)$—is equivalently regarded via its persistent homology transform $\mathrm{PHT}(M)$, which is an object in the derived category of sheaves $\mathcal{D}^b(\mathbf{Shv}(\mathbb{S}^{d-1} \times \mathbb{R}))$.

1.3.4 Main results

With these observations in place, our main result are summarized below. This is joint work with Justin Curry and Sayan Mukherjee and is contained in [ACM23].

Theorem 1.3.1. *The following assignment is a homotopy sheaf:*

$$\mathcal{F} : \mathcal{CS}(\mathbb{R}^d)^{op} \to \mathcal{D}^b(\mathbf{Shv}(\mathbb{S}^{d-1} \times \mathbb{R})) \qquad M \mapsto \mathrm{PHT}(M).$$

This is precisely stated in Theorem 3.2.8. Intuitively, this result allows us to interpolate between shapes in a continuous way via their persistent homology transforms; continuity is mediated via the Grothendieck topology on $\mathcal{CS}(\mathbb{R}^d)$. Essentially, we can glue the PHT of smaller shapes to create the PHT of a larger shape. More precisely, our main result establishes *Čech descent* for the persistent homology transform, which is a generalization of the sheaf axiom that holds for higher degrees of homology.

In Theorem 3.1.1 we interpret the homotopy sheaf axiom for the PHT in terms of a generalized inclusion-exclusion principle for the Euler Characteristic Transform (ECT), which is the decategorification of the PHT.

Theorem 1.3.2. *For a finite cover $\mathcal{M} = \{M_i\}_{i \in \Lambda}$ of $M \subset \mathbb{R}^d$ by constructible subsets*

$$\mathrm{ECT}(M) = \sum_{I \subset \Lambda} (-1)^{|I|+1} \, \mathrm{ECT}(M_I)$$

where each M_I *denotes the intersection* $M_{i_1} \cap M_{i_2} \cap \cdots \cap M_{i_k}$ *for* $I = (i_1, ..., i_k)$.

The homotopy sheaf axiom also implies a nerve lemma for the PHTs, which says that degree-0 homology suffices to characterize polyhedral shapes (see Corollary 3.4.1 for a more precise statement).

Theorem 1.3.3 (Nerve Lemma for the PHT.). *If* $M \in \mathcal{CS}(\mathbb{R}^d)$ *is a polyhedron, i.e. it can be written as a finite union of closed linear simplices* $\mathcal{M} = \{\sigma_i\}_{i \in \Lambda}$*, then the* n*-th cohomology sheaf of the PHT (or the persistent homology transform in degree* n*), written* $\mathrm{PHT}^n(M)$*, is isomorphic to the* n*-th cohomology of the following complex of sheaves:*

$$0 \to \bigoplus_{I \subset \Lambda \, s.t. \, |I|=1} \mathrm{PHT}^0(\mathcal{M}_I) \to \bigoplus_{J \subset \Lambda \, s.t. \, |J|=2} \mathrm{PHT}^0(\mathcal{M}_J) \to \cdots$$

Here $\mathcal{M}_I$ *with* $|I| = k$ *denotes the disjoint union of depth* k *intersections of closed simplices appearing in the cover* $\mathcal{M}$.

As such it is desirable to have an approximation result that is provably stable under the persistent homology transform and allows us to work degree-0 homology alone. We do this by proving a general stability theorem for the PHT across shapes i.e. small perturbations of shapes result in at most small changes in the corresponding persistent homology transforms. To do this, we need to define metrics on our shape space. In Chapter 4, we define two classes of distances: the interleaving type distance and the Wasserstein type distances. We provide some comparisons between these distances and those used in the past. We also show that ϵ perturbations of shapes result in no more than ϵ perturbations of the persistent homology transforms with respect to the introduced distances. The result is stated as follows:

Theorem 1.3.4 (Stability of the PHT). *Any two constructible sets* $M, N \in \mathcal{CS}(\mathbb{R}^d)$ *that are homotopy equivalent via homotopies that move no point more than* ϵ *have* ϵ*-close PHTs of* M *and* N.

Using the sampling methods in [NSW08], construct a discrete approximation of the manifold such that with high confidence the homotopy type of manifold is the same as the nerve of the union of balls around sample points. So, any compact submanifold M can be approximated with arbitrary $\epsilon > 0$ precision by a polyhedral shape such that their peristent homology transforms are ϵ-close.

Corollary 1.3.5 (Approximation of the PHT)**.** *For any compact submanifold M and any $\epsilon > 0$ we can construct a polyhedron N so that with high probability* $\mathrm{PHT}(M)$ *and* $\mathrm{PHT}(N)$ *are ϵ-close.* $\mathrm{PHT}(N)$ *can then be computed using Theorem 1.3.3.*

To summarize, we present a mathematically rigorous construction of shape space, where shapes are not required to be smooth, unlike in previous constructions of shape spaces. Furthermore, shapes need not be diffeomorphic or even homeomorphic to one another. We study the structure of such a shape space and demonstrate that it possesses a desirable local-to-global property, witnessed by the homotopy sheaf axiom. Furthermore, we define metrics on this space and describe a method to approximate shapes using close polyhedral counterparts.

1.4 Diffusion on stratified spaces

We are interested in sampling randomly from stratified spaces. For example, this could allow us to put a distribution on shape spaces, see Example 1.1.2. Stratified spaces are not only non-linear, but also their tangent spaces are not always linear, unlike the situation with smooth manifolds. They are precisely non-linear at points that are singular (or not smooth). This makes statistical analysis much more challenging. Our approach involves defining Brownian motion (or diffusion in general) on these spaces. The transition kernels of these processes can then serve as probability distribution functions from which we can sample or construct statistical models.

Prior work on diffusion on stratified spaces include tree spaces [FD84], Euclidean 2-complexes [BK01] and cubical complexes [Nye20]. Nevertheless, the understanding of diffusion in general stratified spaces remains limited.

Diffusion on stratified spaces not only enables statistics on such spaces but can also be used to solve problems in geometry. Brownian motion and probabilistic methods inform the geometry of the space (see for example the probabilistic proof of the Atiyah-Singer index theorem [Hsu87]).

Sturm in [Stu98] introduced an approach for constructing diffusion processes on *metric measure spaces*. This construction relies on the theory of Dirichlet forms, due to the well known correspondence between regular Dirichlet forms and symmetric Markov processes [FOT11]. The key idea of this approach is to construct Brownian motion by taking limits of a sequence of non-local Dirichlet forms rather than taking a limit of an appropriated scaled random walks. The advantage of using Dirichlet forms is that it does not require the domain, and more importantly, the ambient space to be smooth.

Example 1.4.1. (Dirichlet forms on Riemannian manifolds) Let (M, g) be a n-dimensional smooth Riemannian manifold. Then the Dirichlet form is given by,

$$\mathcal{E}(u, u) = \frac{1}{2} \int_M |\nabla u(x)|^2 m(\mathrm{d}x)$$

where u is any (real-valued) continuous Lipschitz function on M with compact support and m is the Riemannian volume. The associated canonical diffusion process is the Brownian motion on M. It can be shown that the generator of this diffusion process is the Laplace-Beltrami operator Δ_M, and so this is consistent with Hsu's treatment of Brownian motion on manifolds [Hsu02].

Our goal is to generalize the above construction for stratified spaces and construct a canonical diffusion process on them. However, the methods used to prove these conditions for Riemannian manifolds cannot be directly applied to stratified spaces for the following reasons.

1. **Degenerate volume on strata.** To construct a well-defined volume on stratified spaces that restricts to the Riemannian volume on each of the strata calls for a definition that gives zero volume to singular strata. It is often the case, that due to

the presence of isolated singularities, the volume of certain subsets can blow up to infinity.

2. **Lack of exponential maps.** Non-linear spaces lack vector addition, and to some extent, exponential maps remedy this defect. However, on stratified spaces, tangent spaces are no longer vector spaces, rather they are cones. For these reasons, exponential maps are not well defined.

The first item can be mitigated by considering a nicer class of stratified spaces, for eg. subanalytic spaces. The second item is much more technical and so essentially we view stratified spaces as metric measure spaces. Sturm gives a general measure theoretic condition on the metric measure space (X, d, m) that guarantees the convergence of the sequence of non-local Dirichlet forms to a local one, thereby ensuring a canonical diffusion process on (X, d, m) which is point-wise well defined for any starting point $x \in X$. This measure theoretic condition is a weaker notion of the volume doubling condition, i.e. measure of balls of radius $2r$ is at most constant times the measure of the ball of radius r. In Theorem 5.3.6 we show that compact subanalytic spaces with open and dense stratum satisfy this condition. The implication of this result reads:

Theorem 1.4.2. *Let X be a compact subanalytic space with open and dense top-dimensional strata. Then there exists a strongly local regular Dirichlet form $(\mathcal{E}, \mathcal{F})$. This implies the existence of a diffusion process $(X_t, \mathbb{P}_t)$ that is uniquely determined for every starting point $x \in X$.*

This is joint work with Celia Hacker, Sayan Mukherjee and Markus Pflaum.

1.5 Outline

Chapters 2-4 discuss shape space and metrics on shape space, and can also be found in [ACM23]. Chapter 5 addresses diffusion on stratified spaces. Specifically, Chapter 2 provides background information on o-minimal sets, sheaves, cohomology, and topological transforms. In Chapter 3, we describe the sheaf-theoretic construction of shape space and prove the sheaf axiom for the persistent homology transform, along with the nerve lemma. Additionally, we present a short proof using infinity categories in Section 3.7. Chapter 4 discusses metrics on the constructed shape space, where we establish the stability and approximation theorem. In Chapter 5, we first give a brief introduction to subanalytic sets in Section 5.1. Then, in Section 5.2, we discuss diffusion on metric measure spaces, and finally, in Section 5.3, we describe diffusion on subanalytic sets.

2. Background

2.1 Stratified spaces

The spaces we consider in this thesis are in general not smooth and possess singularities. We still assume that these spaces can be stratified meaning that they can be decomposed into smooth parts such that the decomposition is locally finite and fulfills the condition of frontier. For the convenience of the reader we explain this in some detail below; for further information on stratified spaces see [Pfl01].

Definition 2.1.1 (Stratified subset of $\mathbb{R}^d$)**.** Let $X \subset \mathbb{R}^d$ be a locally closed subset. By a *stratification* one understands a locally finite partition $\mathscr{S} = \{S_i\}_{i \in I}$ of X into locally closed subspaces $S_i \subset X$ called *strata* such that the following properties hold true:

(D1) Each stratum is a finite-dimensional topological manifold in the induced topology and $X = \bigsqcup_{i \in I} S_i$.

(D2) The *condition of frontier* holds true that is if R, S are two strata such that $R \cap \overline{S} \neq \emptyset$, then $R \subset \overline{S}$.

The dimension of X is given by

$$\dim X = \sup\{\dim S \mid S \in \mathscr{S}\}.$$

Note that the dimension is independent of the stratification. The dimension of X is p if and only if X contains an open set homeomorphic to an open ball in $\mathbb{R}^d$, but not an open set homeomorphic to an open ball in $\mathbb{R}^e$, for $e > d$. The strata with maximal dimension are called *regular* strata, while those that are not regular are termed as *singular*. The union of singular strata is denoted as $\Sigma \subset X$.

Remark 2.1.2. Note that we can replace $\mathbb{R}^d$ with any smooth manifold M, and many of the results are true for arbitrary M. However, in this work, we are only interested in subsets of $\mathbb{R}^d$, so we work exclusively within the setting of $\mathbb{R}^d$.

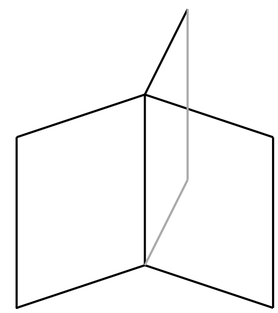

FIGURE 2.1: The 2-open book with 3 pages.

Remark 2.1.3. If R, S are two strata of a stratified space $(X, \mathscr{S})$ such that $R \subset \overline{S}$ one calls the stratum R *incident* to S. We sometimes write $R \leqslant S$ for the situation when R is incident to S. The incidence relation is an order relation on the set of strata.

Example 2.1.4 (Manifold with Boundaries)**.** The simplest example of stratified spaces are manifolds with boundaries. As you might expect a manifold with boundary X can be decomposed as a union of the interior and the boundary.

Example 2.1.5 (Glued Spaces)**.** A general recipe for creating stratified spaces is by gluing together manifolds with boundaries along their boundaries such that the axiom (D2) is satisfied. Some important examples include *k-spiders*, which consist of k copies of the real line glued together at the origin, and the *n-open book* with k pages, obtained by taking the Cartesian product of the k-spider with $\mathbb{R}^n$.

Example 2.1.6 (The Whitney cusp)**.** Consider the zero set of the real polynomial $y^2 + x^3 - z^2x^2$ i.e.

$$\mathcal{W} = \{(x, y, z) \in \mathbb{R}^3 \mid y^2 + x^3 - z^2x^2 = 0\}.$$

The gradient of the polynomial is $(3x^2 - 2xz^2, 2y, 2x^2z)$ and so the critical points lie on the z-axis. For every point p in $\mathcal{W}$ that does not lie on the z-axis, by the inverse function theorem there exists a neighbourhood of p that is diffeomorphic to a 2-dimensional Euclidean ball. Call this part the regular or smooth part of $\mathcal{W}$, denoted as $\mathcal{W}^{\mathrm{reg}}$. Hence, we have a

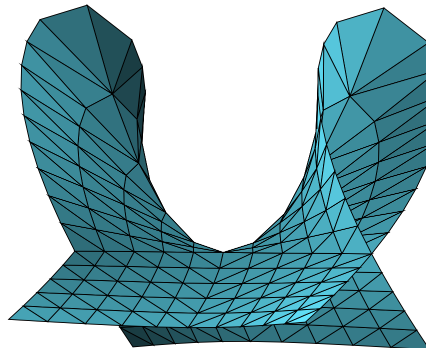

FIGURE 2.2: The Whitney cusp.

decomposition of $\mathcal{W}$ as follows

$$\mathcal{W} = \mathcal{W}^{\text{reg}} \cup \{z\text{−axis}\}.$$

There is another natural decomposition of $\mathcal{W}$ given by

$$\mathcal{W} = \mathcal{W}^{\text{reg}} \cup \{0\} \cup \{z\text{−axis}\backslash 0\}.$$

This turns out to be the 'right stratification' for the Whitney cusp, and one of the reasons is because this decomposition satisfies *local homogeneity*, which means that every neighbourhood around a point on each stratum looks the "same." These conditions are precisely formulated by Whitney and are commonly known as Whitney's conditions.

Stratified spaces, in general, constitute a very large class of spaces, and often, for applied problems, such a high degree of generality may not be necessary. For example, in the context of shape analysis, we are concerned with comparing shapes that can be represented by a computer via triangular mesh. Therefore, we are interested in working with a class of sets that are triangulable and have finite changes in topology.

In this thesis, we primarily consider two main stratified spaces: *sets definable with respect to o-minimal structures* on $\mathbb{R}^d$ (Section 2.2) and *subanalytic sets* (Section 5.1). Both of these sets of sets are tame in the sense of Grothendieck's *Esquisse d'un programme* [Gro97]. A relationship between these two classes of sets is discussed in Example 2.2.3. The explicit stratification for sets definable with respect to o-minimal structures is not important, and we do not utilize it in our work. In fact, a cell decomposition theorem can be found in [Van98b], and they admit Whitney stratifications [Lê 98].

2.2 O-minimal topology

We define a class of shapes that we are interested in working with– constructible sets. These sets precisely capture the notion of a "shape." Specifically, these shapes are stratified spaces but to ensure our shapes possess desirable properties and avoid pathological cases like Cantor-like sets, we impose certain tameness conditions. These conditions have been formalized into a broader class of sets known as o-minimal structures.

The main idea of an o-minimal structure is to provide a class of "nice" sets that do not degenerate (i.e. remain "nice") after performing elementary logical, geometric, and topological operations. Refer to [Van98b] for a full treatment of the subject.

Definition 2.2.1 (O-minimal structure)**.** An **o-minimal structure** $\mathcal{O} = \{\mathcal{O}_d\}$, is a specification of a boolean algebra of subsets $\mathcal{O}_d$ of $\mathbb{R}^d$ for each natural number $d \geqslant 0$ such that

1. $\mathcal{O}$ is closed under certain product operators, i.e. if $A \in \mathcal{O}_d$, then $A \times \mathbb{R}$ and $\mathbb{R} \times A$ are in $\mathcal{O}_{d+1}$;

2. $\mathcal{O}$ is closed under projection operators, i.e. if $A \in \mathcal{O}_{d+1}$, then $\pi(A) \in \mathcal{O}_d$ where $\pi : \mathbb{R}^{d+1} \to \mathbb{R}^d$ is the axis aligned projection;

3. $\{(x_1, ..., x_d) \in \mathbb{R}^d : x_i = x_j\} \in \mathcal{O}_d$ for $1 \leqslant i < j \leqslant d$;

4. The graphs of addition and multiplication are in $\mathcal{O}_3$;

5. $\mathcal{O}_1$ contains only finite unions of points and intervals (with respect to the natural order $<$ on $\mathbb{R}$).

Here a boolean algebra of sets on $\mathcal{O}_d$ means that $\mathcal{O}_d$ is closed under finite unions, intersections, and complements. Condition 1 ensures that $\mathcal{O}$ is closed under Cartesian products. To see this write $A \times B = (A \times \mathbb{R}^n) \cap (\mathbb{R}^m \times B)$ for $A \in \mathcal{O}_m$ and $B \in \mathcal{O}_n$, and $\mathcal{O}$ is closed under intersection since it is a Boolean algebra. Condition 3 ensures diagonals belong to $\mathcal{O}$. Conditions 1-3 are required to form a structure on $\mathbb{R}$ in the logical sense i.e. every set that is definable from a set in $\mathcal{O}$ is also in $\mathcal{O}$. Condition 4 is required to ensure that there is definable way of selecting the midpoint of a bounded interval, which in turn implies some deeper results, for instance, the triangulation theorem.

Condition 5 is the simplest assumption that makes the structure compatible with the order of $\mathbb{R}$. Since the order gives a topology to $\mathbb{R}$, these conditions express compatibility with the topology. Thus, the term "o-minimal": *minimal* assumptions required for the structure to behave well with the *order* of $\mathbb{R}$.

Example 2.2.2 (Semi-algebraic sets)**.** A semi-algebraic set in $\mathbb{R}^d$ is a finite union of sets of the form

$$\{x \in \mathbb{R}^d \mid f_1(x) = \cdots = f_k(x) = 0, g_1(x) > 0, \ldots, g_l(x) > 0\}.$$

Let $\mathcal{O}_d$ be the set of semi-algebraic sets in $\mathbb{R}^d$. Then $\mathbb{R}_{alg} = (\mathcal{O}_d)_d$ is a o-minimal structure on $\mathbb{R}$. One can readily check that semi-algebraic sets are closed under unions, intersections, complements, and products. Closure with respect to projections is due to Tarski-Seidenberg and involves model-theoretic notions to prove. The semi-algebraic sets of $\mathbb{R}$ are finite unions of points and intervals and so condition 5 is true.

In fact, due to this definition, every o-minimal structure must contain all semi-algebraic sets, so $\mathbb{R}_{alg}$ is the smallest o-minimal structure.

Example 2.2.3 (Globally subanalytic sets)**.** In general, subanalytic sets (see Section 5.1 for a definition) do not form an o-minimal structure on $\mathbb{R}$ due to the fact that they are not closed under projections. However, subsets of $\mathbb{R}^d$ that are subanalytic in the larger projective space $\mathbb{RP}^d$ given by the map

$$(x_1, \ldots, x_d) \mapsto [1, x_1, \ldots, x_d]$$

forms an o-minimal structure [Van98a].

A set $A \subseteq R^d$ is **definable** with respect to an o-minimal structure $\mathcal{O}$ if $A \in \mathcal{O}_d$. A map $f : A \to \mathbb{R}^n$ **definable** (with respect to $\mathcal{O}$) if its graph $\{(x, f(x))|x \in A\} \subset \mathbb{R}^{n+d}$ is definable. Furthermore, *compact* definable sets are called **constructible sets**. The sub-collection of constructible subsets in $\mathcal{O}_d$ is denoted $\mathcal{CS}(\mathbb{R}^d)$.

2.2.1 Topology of definable sets

Grothendieck envisaged a new foundation of topology and geometry– tame topology (or topologie modérée [Gro97])– grounded in axioms rather than analysis. Tame topology seeks to avoid "pathological" objects and counter-intuitive results found in traditional approaches. For example, one of the strong assumptions introduced in this programme is the *triangulability axiom* which says that all tame sets should be triangulable in a tame sense. In contrast, traditional topology readily allows the construction of non-triangulable spaces, such as the comb space,

$$S = \{(x, y) \mid 0 \leqslant y \leqslant 1; x = 0 \text{ or } x = 1/n \text{ for } n \in \mathbb{N}\} \cup ([0, 1] \times \{0\}).$$

This space, being locally disconnected, cannot be homeomorphic to a simplicial complex. Tame topology aims to eliminate such problematic sets.

O-minimal structures are a reasonable choice for an axiomatic approach that realizes a tame topology. For instance, each definable set (with respect to an o-minimal structure) admits cell decompositions and is triangulable. In other words, the definable topology of a set can be entirely expressed through finite combinatorial information.

From now on we fix an o-minimal structure $\mathcal{O}$ on $\mathbb{R}$. When we say definable set, we mean definable set with respect to the o-minimal structure $\mathcal{O}$.

A triangulation in $\mathbb{R}^n$ of a definable set $A \subseteq \mathbb{R}^d$ is a tuple (ϕ, K) where K is a simplicial complex in $\mathbb{R}^n$ and $\phi : A \to |K|$ is a definable homeomorphism. Here $|K|$ is the geometric realization of the simplicial complex K.

Theorem 2.2.4. *(Triangulation Theorem [Van98b]) Each definable set is definably homeomorphic to a simplicial complex.*

An important consequence of this theorem is that algebraic topological signatures such as Euler characteristic and homology, are well defined for any definable set.

Remark 2.2.5. In fact, a stronger version of this theorem exists. Suppose the definable set X has definable subsets $A_1, \ldots, A_k$, then the triangulation of X can be chosen such that it is compatible with each of the subsets [Van98b, Theorem 2.9].

Definition 2.2.6 (Definable Euler Characteristic)**.** If $A \subseteq \mathbb{R}^d$ is a definable set and if $\phi : A \to \sqcup_i \sigma_i$ is a definable bijection with a collection of (open) simplices, then the **definable Euler characteristic** of A is

$$\chi(A) := \sum_i (-1)^{\dim \sigma_i}$$

where $\dim \sigma_i$ denotes the dimension of the open simplex σ_i.

As you might predict, the definable Euler characteristic remains the same regardless of the specific triangulation chosen [Van98b, pp. 70-71]. This is because it is a definable homeomorphism invariant. However, as illustrated in the next example, the definable Euler characteristic is not a homotopy invariant.

Example 2.2.7. The open interval $(0, 1)$ has definable Euler Characteristic -1 but it is contractible to a point which has definable Euler characteristic 1. In general, the definable Euler characteristic of an open n-simplex is $(-1)^n$.

However, for compact definable sets (or constructible) sets the definable Euler characteristic is a homotopy invariant. This is because the definable Euler characteristic can also be written as an alternating sum of the Borel-Moore homology (See [BG10])

$$\chi(A) = \sum_{k=0}^{\infty} (-1)^k H_k^{\mathrm{BM}}(A; \mathbb{R}),$$

and for compact spaces the Borel-Moore homology coincides with the usual singular homology, a well-known homotopy invariant.

2.2.2 Euler Calculus in an o-minimal structure

In this section, a fixed o-minimal structure $\mathcal{O}$ is assumed. When referring to *constructible sets*, we mean constructible sets with respect to the o-minimal structure $\mathcal{O}$. For brevity, we drop the term 'definable' before Euler characteristic from now on.

Euler calculus is a topological integration theory for definable functions based on the Euler characteristic. In this theory, the (finitely additive) *measure* is the Euler characteristic and the *measurable sets* are the definable sets. This makes sense for the following reason:

Lemma 2.2.8. *For definable $A, B \subseteq \mathbb{R}^d$*

$$\chi(A \cup B) = \chi(A) + \chi(B) - \chi(A \cap B). \tag{2.1}$$

This can be seen by applying the triangulation theorem and counting cells.

Remark 2.2.9. Unlike a measure, this is only true for a *finite* unions of definable sets. However, there is still a well-defined notion of an integration operator for finitely additive measure but as you might expect, the integral theory does not behave well with limits. See [BG10] to see more details about this.

For a simple function f valued in $\mathbb{Z}$, $f = \sum_{\substack{A_i \in \mathcal{CS}(\mathbb{R}^d) \\ \text{finite}}} \lambda_{A_i} \mathbb{1}_{A_i}$, and $B \in \mathcal{CS}(\mathbb{R}^d)$ the Euler integration, reminiscent of usual measure-theoretic integration, is given by

$$\int_B f(x) d\chi(x) = \sum_{\substack{A_i \in \mathcal{CS}(\mathbb{R}^d) \\ \text{finite}}} \lambda_{A_i} \chi(A_i \cap B). \tag{2.2}$$

If $X \in \mathcal{CS}(\mathbb{R}^d)$, then a **constructible function** is an integer valued compactly supported function $f : X \to \mathbb{Z}$ with only finitely many non-empty level sets, each of which are definable, and hence triangulable. We denote the ring of constructible functions on X by $\mathcal{CF}(X)$. In this thesis, we specifically focus on constructible functions, so we limit ourselves to compactly supported functions. However, it's worth noting that a more general theory exists.

The stronger version of the triangulation theorem (Remark 2.2.5) allows us to view any function $f \in \mathcal{CF}(X)$ as a simple function i.e. $f = \sum_{\substack{\alpha \\ \text{finite}}} c_\alpha \mathbb{1}_{\sigma_\alpha}$ where $\{\sigma_\alpha\}$ are the simplices and $c_\alpha \in \mathbb{Z}$. So, this gives a well defined notion of Euler integration on $\mathcal{CF}(X)$ for any definable X.

Definition 2.2.10 (Euler Integration)**.** The **Euler integration** is a homomorphism $\int_A d\chi : \mathcal{CF}(X) \to \mathbb{Z}$ that satisfies Equation (2.2). More explicitly, after applying the definitions it can be readily checked that

$$\int_X f d\chi = \sum_{n=-\infty}^{n=\infty} n\chi(f = n) = \sum_{n=0}^{\infty} \chi(f > n) - \chi(f < -n)$$

where the latter equality follows from a telescoping expansion.

Another fundamental property satisfied by the Euler characteristic is the multiplicative property. For any two definable sets X and Y, it holds that $\chi(X \times Y) = \chi(X)\chi(Y)$. This gives two important theorems, namely a Fubini-type theorem and a Riemann-Hurwitz theorem for Euler integration [Vir06]. The Euler integration allows for various operations on constructible functions, and some of these operations are defined below.

Definition 2.2.11 (Pushforward and Pullback, cf. [Sch91a])**.** If $\varphi : X \to Y$ is a (definable) mapping between constructible sets, then we have **pushforward** $\varphi_* : \mathcal{CF}(X) \to \mathcal{CF}(Y)$ and **pullback** $\varphi^* : \mathcal{CF}(Y) \to \mathcal{CF}(X)$ operations via

$$\varphi_* f(y) := \int_{\varphi^{-1}(y)} f d\chi \quad \text{and} \quad \varphi^* g(x) := g(\varphi(x)).$$

Definition 2.2.12 (Radon Transform)**.** Let $S \subset X \times Y$ be a closed constructible subset of the product of two constructible sets. Let π_X and π_Y be the projections onto the indicated factors. The **Radon Transform** with respect to S is a group homomorphism $\mathcal{R}_S : \mathcal{CF}(X) \to \mathcal{CF}(Y)$ defined by

$$\mathcal{R}_S(\phi) := (\pi_Y)_*\big[((\pi_X)^*\phi)\mathbb{1}_S\big] \quad \text{where} \quad \mathcal{R}_S(\phi)(y) = \int_{\pi_Y^{-1}(y)} (\phi \circ \pi_X)\mathbb{1}_S d\chi.$$

A celebrated theorem of Schapira [Sch95a] gives a criterion for determining the invertibility of the Radon transform $\mathcal{R}_S$ in terms of the Euler characteristic of the fibers of $S \subset X \times Y$, when projected to each of these two factors:

Theorem 2.2.13 ([Sch95a] Theorem 3.1)**.** *If $S \subset X \times Y$ and $S' \subset Y \times X$ have fibers S_x and S'_x in Y satisfying*

1. $\chi(S_x \cap S'_x) = \chi_1$ *for all* $x \in X$, *and*

2. $\chi(S_x \cap S'_{x'}) = \chi_2$ *for all* $x' \neq x \in X$,

then for all $\phi \in \mathcal{CF}(X)$,

$$(\mathcal{R}_{S'} \circ \mathcal{R}_S)\phi = (\chi_1 - \chi_2)\phi + \chi_2 \left(\int_X \phi d\chi \right) 1_X.$$

2.2.3 Euler Characteristic Transform

Now, we are ready to define the Euler characteristic transform, a central object of the thesis. Intuitively, the Euler characteristic transform should be thought of as a representation of a shape- it encodes the Euler characteristic of 'enough' sub-levels of the shape, ensuring that the transform is injective.

Definition 2.2.14 (ECT: Map Version)**.** The **Euler Characteristic Transform (ECT)** of a constructible set $M \in \mathcal{CS}(\mathbb{R}^d)$ is the map that assigns to each direction $v \in \mathbb{S}^{d-1}$ the piece-wise constant integer-valued function on $\mathbb{R}$ that records the Euler characteristic of the

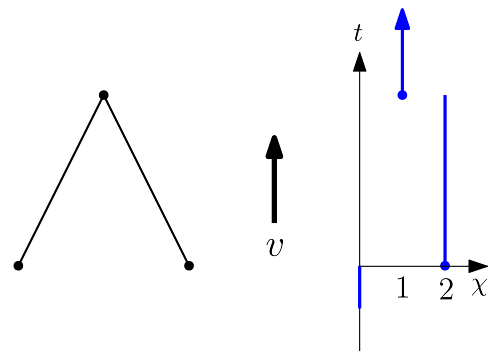

FIGURE 2.3: The inverted 'V' shape summarised by the *Euler characteristic transform.* The Euler characteristic transform in direction v records the Euler characteristic curve of the filtration of M in v.

sublevel set of M in direction v, i.e.

$$\mathrm{ECT}(M) : \mathbb{S}^{d-1} \to \mathrm{Fun}(\mathbb{R}, \mathbb{Z}) \qquad \mathrm{ECT}(M)(v,t) = \chi\big(M_{v,t}\big),$$

where $M_{v,t} := \{x \in M \mid x \cdot v \leqslant t\}$ is the intersection of M with the half-space $x \cdot v \leqslant t$.

We note that since M is constructible and the equation defining a sub-level set is semi-algebraic—and hence definable—the intersection $M_{v,t}$ is constructible as well, and thus has a well-defined Euler characteristic.

As our notation suggests, one can also view $\mathrm{ECT}(M)$ as a function from $\mathbb{S}^{d-1} \times \mathbb{R}$ to $\mathbb{Z}$ that assigns to each pair $(v,t) \in \mathbb{S}^{d-1} \times \mathbb{R}$ the Euler characteristic $\chi(M_{v,t})$. This perspective is important because it allows us to view the ECT as a type of integral transform, which takes a shape (viewed as an indicator function on $\mathbb{R}^d$) and produces a function on $\mathbb{S}^{d-1} \times \mathbb{R}$—a coordinate system built for tomographic comparison.

Definition 2.2.15 (ECT: Version 2)**.** Let $S = \{(x,v,t) \in \mathbb{R}^d \times \mathbb{S}^{d-1} \times \mathbb{R} \mid x \cdot v \leqslant t\}$. The Euler characteristic transform of M is equal to the Radon transform of its indicator function

$$\mathrm{ECT}(M) = \mathcal{R}_S(\mathbb{1}_M).$$

In [CMT22a] this was used to prove that the Euler Characteristic Transform is injective.

Theorem 2.2.16 ([CMT22a] Theorem 3.5)**.** *The Euler Characteristic Transform* ECT : $\mathcal{CS}(\mathbb{R}^d) \to \mathcal{CF}(\mathbb{S}^{d-1} \times \mathbb{R})$ *is injective, i.e. if* $\mathrm{ECT}(M) = \mathrm{ECT}(M')$*, then* $M = M'$*.*

The proof essentially follows from Schapira's inversion theorem (Theorem 2.2.13).

2.3 Persistent Homology

So far, we have discussed the Euler characteristic of constructible sets, which is a topological invariant, and developed an injective transform on constructible sets based on the Euler characteristic. Another useful topological invariant of a constructible set $M \in \mathcal{CS}(\mathbb{R}^d)$ is the homology in degree n with coefficients in a field $\Bbbk$[1], denoted as $H_n(M; \Bbbk)$. This algebraic summary is more powerful than the Euler characteristic because of *functoriality*, which means that for maps on spaces $g : X \to Y$, there is an induced map on homology $g_* : H_k(X) \to H_k(Y)$ and these maps compose correctly in the sense that if there is another map $h : Y \to Z$, then the induced map on homology $(h \circ g)_*$ is the same as $h_* \circ g_*$.

Persistent homology encodes how the homology of M changes under an $\mathbb{R}$-indexed filtration of M. More specifically, let $f : M \to \mathbb{R}$ be a real-valued function on constructible set M. We can define a family of $\mathbb{R}$-modules, typically called persistence modules in degree i by

$$X_t^f := H_i(f^{-1}(-\infty, t]; \Bbbk).$$

Inclusions $f^{-1}((-\infty, t]) \subseteq f^{-1}((-\infty, s])$ when $t < s$ induce homomorphisms of the homology groups[2] for each degree i

$$\iota_i^{t<s} : H_i(f^{-1}(-\infty, t]; \Bbbk) \to H_i(f^{-1}(-\infty, s]; \Bbbk).$$

Definition 2.3.1 (Height function)**.** The height function in the direction $v \in \mathbb{S}^{d-1}$ is denoted as $h_v : M \to \mathbb{R}$, defined by $h_v(x) = x \cdot v$. Denote the sub-level set $h_v^{-1}((-\infty, t]) = \{x \in M \mid x \cdot v \leqslant t\}$ as $M_{v,t}$.

[1] In general, the coefficients can be taken from any abelian group G, but we exclusively work with field coefficients.

[2] In the case where the coefficients are valued in a field $\Bbbk$, homology groups are actually vector spaces.

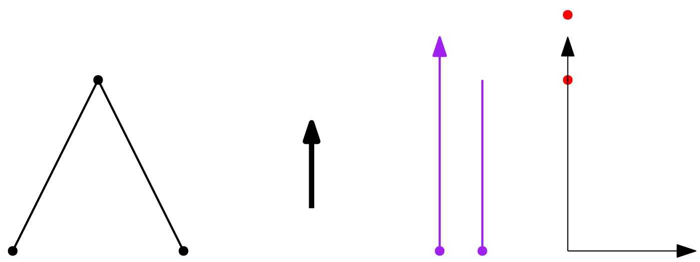

FIGURE 2.4: The barcode and persistence diagram of shape inverted 'V' filtered according to height function.

We call $r \in \mathbb{R}$ a homological **critical value** for the filtration function f if there exists a degree i such that the induced map $H_i(f^{-1}(-\infty, r-\delta]; \Bbbk) \to H_i(f^{-1}(-\infty, r]; \Bbbk)$ is not an isomorphism for any $\delta > 0$.

The filtration function f is **tame** if there are finitely many critical values, and the homology vector spaces $H_i(f^{-1}(-\infty, t]; \Bbbk)$ is finitely generated for all degrees i and parameters $t \in \mathbb{R}$.

For constructible sets $M \in \mathcal{CS}(\mathbb{R}^d)$, the height function h_v is tame for every direction $v \in \mathbb{S}^{d-1}$, as there are only finitely many critical values of $M_{v,t}$ as t varies [CMT22b, Lemma 3.4]. This ensures that the persistence diagrams, as defined in Definition 2.3.3, exist.

Definition 2.3.2 (Persistent homology group)**.** The ith sublevel set persistent homology group of the height function in direction $v \in \mathbb{S}^{d-1}$ is the image of the homomorphism $\iota_i^{t<s} : H_i(M_{v,t}; \Bbbk) \to H_i(M_{v,s}; \Bbbk)$.

These persistent homology groups contain important information about the birth and death of topological features of M when filtered in a direction v.

A homology class $\alpha \in H_i(M_{v,t})$ is born at time a if α lies in the cokernel of $\iota_i^{a'<a}$ for any $a' < a$. Often, the birth of a class α is denoted as $b(\alpha)$. The class α born at a dies at b (denoted as $d(\alpha)$) if for all $a' < a < b' < b$, $\iota_i^{a<b}(\alpha) \in \operatorname{im} \iota_i^{a'<b}$ but $\iota_i^{a<b'}(\alpha) \notin \operatorname{im} \iota_i^{a'<b'}$. The class that never dies is known as the essential class.

Definition 2.3.3 (Persistence Diagram)**.** Let $\mathbb{R}^{2+}$ be the extended plane that is above the diagonal, i.e. $\mathbb{R}^{2+} = \{(a,b) \in \{-\infty\} \cup \mathbb{R} \times \mathbb{R} \cup \{\infty\} \mid a < b\}$. Then the persistence diagram in degree i of a constructible set M filtered in direction $v \in \mathbb{S}^{d-1}$, is multi-set (set with repetition allowed) of points

$$\mathcal{B}_i := \mathrm{PH}_i(M, v) = \{(a, b); n)\} \subset \mathbb{R}^{2+} \times \mathbb{N}$$

such that $\dim \operatorname{im} \iota_i^{a<b}$ is equal to the number of points (counting multiplicity) in $(-\infty, a] \times [b, \infty)$. This is achieved by placing at each (a, b) number of points (or "dots"; see Figure 2.4) equal to dimension of the space of i-dimensional homology classes that are born at time a and die at time b.

We can consider any persistence diagram $\mathcal{B}$ as a set instead of a multi-set, referred to as a barcode, by utilizing another coordinate to enumerate copies of each interval I. i.e., $\mathcal{B} = \{(I; j) \mid (I; j) \in \mathbb{R}^{2+} \times \mathbb{N}\}$ where j indicates the j^{th} copy of I.

There is a persistent diagram for each direction $v \in \mathbb{S}^{d-1}$. We denote the set of all persistence diagrams as **Dgm**.

Definition 2.3.4 (Matchings and the p-Wasserstein Distances [CMT22b])**.** A **matching** of these is a partial bijection $\sigma : \mathcal{B} \to \mathcal{B}'$, i.e., a choice of subset $\mathcal{M} \subseteq \mathcal{B}$, called the domain $\operatorname{dom}(\sigma)$, and an injection $\sigma : \mathcal{M} \to \mathcal{B}'$. The complement of the domain of σ as $\operatorname{dom}^c(\sigma) := \mathcal{B} \backslash \operatorname{dom}(\sigma)$ and the complement of the image of σ as $\operatorname{im}^c(\sigma) := \mathcal{B}' \backslash \operatorname{im}(\sigma)$ together define the **unmatched points** of σ. We then promote a partial bijection σ to an actual bijection via the introduction of diagonal images; a point $I = (b, d) \in \mathbb{R}^{2+}$ where neither coordinate is ∞ has a **diagonal image** $\Delta(I) = (\frac{b+d}{2}, \frac{b+d}{2})$. With this convention, a partial bijection $\sigma : \mathcal{B} \to \mathcal{B}'$ becomes an actual bijection of augmented persistence diagrams $\tilde{\sigma} : \mathcal{B}(\sigma) \to \mathcal{B}'(\sigma)$ where

$$\mathcal{B}(\sigma) := \operatorname{dom}(\sigma) \cup \operatorname{dom}^c(\sigma) \cup \bigcup_{(I', j') \in \operatorname{im}^c(\sigma)} (\Delta(I'); j')$$

and

$$\mathcal{B}'(\sigma) := \operatorname{im}(\sigma) \cup \operatorname{im}^c(\sigma) \cup \bigcup_{(I, j) \in \operatorname{dom}^c(\sigma)} (\Delta(I); j).$$

The map $\tilde{\sigma}$ now matches points that were previously unmatched by σ with their corresponding diagonal images. By abuse of notation, we simply write σ for the extended map $\mathcal{B}(\sigma) \to \mathcal{B}'(\sigma)$. If $\sigma_j(I)$ denotes the $\mathbb{R}^{2+}$ coordinates of $\sigma(I;j)$ and $||I - \sigma_j(I)||_p$ is the ℓ^p metric on $\mathbb{R}^{2+}$, then the ***p*-cost** of this extended matching σ is

$$\operatorname{cost}_p(\sigma) = \left(\sum_{(I,j) \in \mathcal{B}(\sigma)} ||I - \sigma_j(I)||_p^p \right)^{1/p}.$$

For every $p \in [1, \infty]$ we define the **Wasserstein *p*-distance** between two diagrams $\mathcal{B}$ and $\mathcal{B}'$ is then the infimum of this matching cost over all matchings, i.e.

$$W_p(\mathcal{B}, \mathcal{B}') := \inf_{\sigma:\mathcal{B}\to\mathcal{B}'} \operatorname{cost}_p(\sigma).$$

We note that the Wasserstein ∞-distance is also called the **bottleneck distance**, for which we reserve the special notation

$$d_B(\mathcal{B}, \mathcal{B}') := W_\infty(\mathcal{B}, \mathcal{B}') = \inf_{\sigma:\mathcal{B}\to\mathcal{B}'} \max_{(I,j)\in\mathcal{B}(\sigma)} ||I - \sigma_j(I)||_\infty.$$

We now have enough preliminaries to provide the first, classical definition of the persistent homology transform. The persistent homology transform studies the persistent homology of a constructible subset $M \in \mathcal{CS}(\mathbb{R}^d)$ by considering the filtration $M_{v,t} = \{x \in M \mid x \cdot v \leqslant t\}$ for all directions $v \in \mathbb{S}^{d-1}$.

Definition 2.3.5 (PHT: Map Version)**.** Let $M \in \mathcal{CS}(\mathbb{R}^d)$ be a constructible set. The **persistent homology transform** of M is defined as the map

$$\operatorname{PHT}(M) : \mathbb{S}^{d-1} \to \mathbf{Dgm}^d \qquad v \mapsto (\operatorname{PH}^0(M,v), \operatorname{PH}^1(M,v), \ldots, \operatorname{PH}^d(M,v))$$

See Figure 2.5 for an example.

Remark 2.3.6. Note that we have defined the persistent homology transform using cohomology instead of homology, to ensure consistency with all the different versions of the definitions we provide later on. Typically, we assume that $\Bbbk$ is of characteristic 0 so that homology and cohomology are isomorphic, i.e. $H_n(M;\Bbbk) \cong H^n(M;\Bbbk)$.

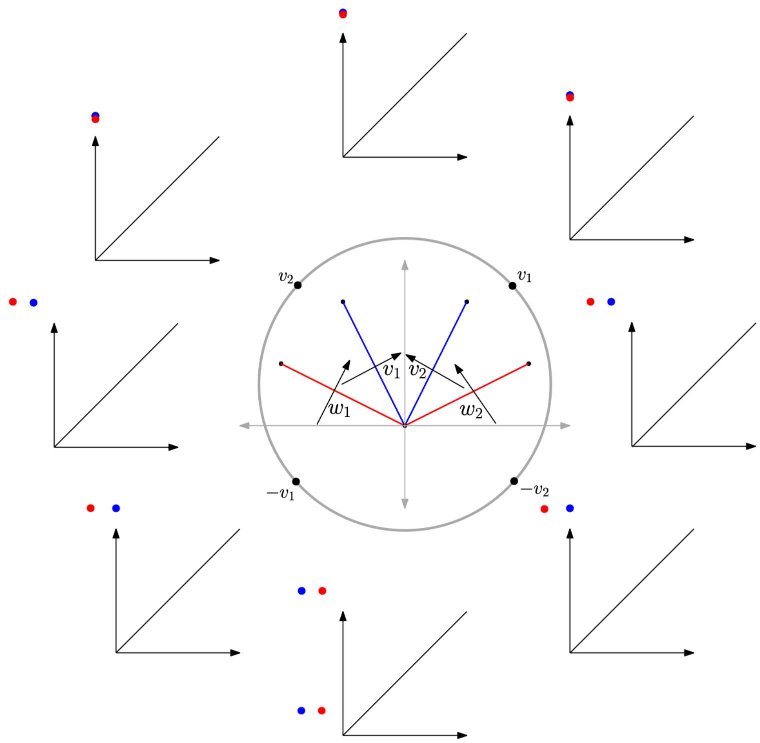

FIGURE 2.5: For an embedded shape in $\mathbb{R}^2$, the PHT (Definition 2.3.5) assigns to each direction $v \in \mathbb{S}^1$ a persistence diagram. In this example we super-impose the PHT of two different embeddings of the letter 'V', one in blue with normal directions v_1 and v_2 to the sides of the letter 'V' and one in red with flatter sides with steeper normal directions w_1 and w_2. Around the circle are eight persistence diagrams corresponding to the directions $\pm e_i$ and $\pm v_i$ for $i = 1, 2$; $e_1 = (1, 0)$ points in the positive x direction and $e_2 = (0, 1)$ points in the positive y direction. Notice that the blue persistence diagram associated to direction $-e_2$ has an earlier birth time than the red persistence diagram.

Although homology is usually a lossy summary of a shape, knowing the persistent homology of a shape in every direction completely determines the shape.

Theorem 2.3.7 (PHT: Injectivity, cf. [CMT22a])**.** *If* $\mathrm{PHT}(M) = \mathrm{PHT}(N)$, *then* $M = N$.

2.4 Sheaves and Cohomology

It is now understood, thanks to the foundational thesis by Curry [Cur14a], that traditional persistent homology can be viewed as a special case of sheaf theory. The main idea here is that the filtration parameter space $\mathbb{R}$ indexes the (co)chain complexes of spaces, such as the sub-level sets, rather than the (co)homology of these spaces. Moreover, as first outlined in [CMT22a], sheaves also allow us to view the persistent homology transform (PHT) as a collection of (co)chain complexes parameterized by the entire parameter space $\mathbb{S}^{d-1} \times \mathbb{R}$.

2.4.1 Sheaf Theory

In this section, we review basic sheaf theory and sheaf cohomology, necessary for understanding the concepts discussed above.

Presheaves and sheaves provide a way of associating "algebraic data" to topological spaces.

Definition 2.4.1 (Pre-Sheaves)**.** Let X be a topological space and let $\text{Open}(X)$ be the poset of open sets in X. A **pre-sheaf** valued in the category of abelian groups **Ab** on X is a functor $\mathcal{F} : \text{Open}(X)^{op} \to \mathbf{Ab}$. We sometimes write $\rho_{U,V} : \mathcal{F}(U) \to \mathcal{F}(V)$ for the restriction map associated to the inclusion $V \subseteq U$.

We refer to elements of $\mathcal{F}(X)$ (commonly denoted as $\Gamma(X, \mathcal{F})$) as global sections and elements of $\mathcal{F}(U)$ (commonly denoted as $\Gamma(U, \mathcal{F})$) for $U \in \text{Open}(X)$ as local sections of $\mathcal{F}$ over U.

Remark 2.4.2. Categorically speaking, we can define a presheaf to be a contravariant functor from $\text{Open}(X)$ to **Ab**. More generally, we can replace the category of abelian groups to any fixed concrete category[3].

Definition 2.4.3 (Stalks)**.** The **stalk** of a pre-sheaf at a point $x \in \mathbb{X}$ is then defined as the direct limit of $\mathcal{F}$ over all open sets containing x, i.e.

$$\mathcal{F}_x := \varinjlim_{U \ni x} \mathcal{F}(U)$$

[3] A concrete category $\mathcal{C}$ can be thought of as the category of sets, with some additional structure. It is characterized by a faithful forgetful functor $\mathcal{C} \to \mathbf{Set}$. For example, the category of groups, rings, and vector spaces are concrete categories. See [nLa24] for a precise definition.

One can think of the stalk at x as the "local picture" of $\mathcal{F}$ at x. This can be constructed rigorously as follows

$$\mathcal{F}_x = \bigsqcup_{U \ni x} F(U) / \sim,$$

that is, the equivalence class of abelian groups given by the relation $(f, U) \sim (g, V)$ for $f \in \mathcal{F}(U)$ and $g \in \mathcal{F}(V)$ iff there is an open $W \subseteq U \cap V$ containing x such that $\rho_{U,V} f = \rho_{V,W} g$ i.e. two local sections around x are identified whenever they agree on some small enough neighborhood of x.

Remark 2.4.4. This definition remains applicable even if we replace abelian groups with a concrete category $\mathcal{C}$. However, for the rigorous definition of stalks for pre-sheaves valued in $\mathcal{C}$ to make sense, we must ensure that the forgetful functor from $\mathcal{C}$ to **Set** preserves filtered colimits.

Definition 2.4.5 (Čech Cohomology of a Cover)**.** Given a pre-sheaf $\mathcal{F}$ and an open cover $\mathcal{U} = \{U_i\}_{i \in \Lambda}$ of an open set U, one has the **Čech cochain complex** associated to $\mathcal{U}$ where the n-cochains are given by the product over intersections of $n+1$ cover elements, i.e.,

$$\check{C}^n(\mathcal{U}; \mathcal{F}) = \prod_{i_0, \ldots, i_n \in \Lambda} \mathcal{F}(U_{i_0, \ldots, i_n})$$

where we always assume $\mathcal{F}(\varnothing) = 0$. The n^{th}-coboundary operator is defined by specifying the contribution of a general element $s \in \check{C}^n(\mathcal{U}; \mathcal{F})$ to the factor $i_0, \ldots, i_{n+1}$ in $\check{C}^{n+1}(\mathcal{U}; \mathcal{F})$. This is given by the formula

$$(\delta^n s)_{i_0, \ldots, i_{n+1}} = \sum_{j=0}^{n+1} (-1)^j s_{i_0, \ldots, \hat{i_j}, \ldots, i_{n+1}} |_{U_{i_0, \ldots, i_{n+1}}},$$

where $\hat{i_j}$ denotes removal of that entry and the vertical line is the commonly accepted shorthand for the application of the restriction map from $U_{i_0, \ldots, \hat{i_j}, \ldots, i_{n+1}}$ to $U_{i_0, \ldots, i_{n+1}}$.

Definition 2.4.6 (Sheaves)**.** A pre-sheaf of abelian groups $\mathcal{F}$ is a **sheaf** if for every open set U and open cover $\mathcal{U} = \{U_i\}_{i \in \Lambda}$ we have that the value of $\mathcal{F}$ on U can be computed using the Čech cohomology of the cover:

$$\mathcal{F}(U) \cong H^0\Big[\check{C}^0(\mathcal{U};\mathcal{F}) \to \check{C}^1(\mathcal{U};\mathcal{F}) \to \check{C}^2(\mathcal{U};\mathcal{F}) \to \cdots\Big]$$

This is a *local-to-global principle*, because it guarantees that a sheaf is always determined locally by a cover. More generally, one can define sheaves valued in categories that are not necessarily abelian, such as **Set**, but where the categorical notion of an (inverse) limit makes sense. As a reminder, the limit[4] is a categorical operation that takes an entire diagram of objects—in this case, the objects $\mathcal{F}(U_i)$—and produces a single object that maps to each of the objects and commutes with the morphisms connecting the objects—in this case the restriction maps $\mathcal{F}(U_i) \to \mathcal{F}(U_{ij}) \leftarrow \mathcal{F}(U_j)$. One then modifies the sheaf axiom to instead require that

$$\mathcal{F}(U) \cong \varprojlim \Big[\prod_i \mathcal{F}(U_i) \rightrightarrows \prod_{i,j} \mathcal{F}(U_{ij}) \mathrel{\substack{\rightarrow\\\rightarrow\\\rightarrow}} \cdots\Big].$$

It is well known that the limit only depends on the objects $\{\mathcal{F}(U_i)\}$ associated to the cover elements along with the objects associated to the pairwise intersections $\{\mathcal{F}(U_{ij})\}$, in similar spirit to how H^0 only depends on the nodes and edges of a triangulated space. Higher-order information, such as intersections of three or more elements, is ignored by the usual sheaf axiom.

Remark 2.4.7. The sheaf condition can also be more familiarly expressed as follows: For any open covering $(U_i)_{i\in\Lambda}$ of $U \in \mathrm{Open}(X)$,

1. (Gluing) If sections $f_i \in \mathcal{F}(U_i)$ agree on overlaps i.e. $\rho_{U_{ij},U_i} f_i = \rho_{U_{ij},U_j} f_j$ for all $i,j \in \Lambda$, then there exists a unique lift $f \in \mathcal{F}(U)$ such that $\rho_{U_i,U} f = f_i$ for all i;
2. (Locality) Suppose that $f,g \in \mathcal{F}(U)$ and that $\rho_{U_i,U} f = \rho_{U_i,U} g$ for all $i \in \Lambda$, then $f = g$.

Definition 2.4.8 (Morphism of sheaves)**.** A morphism between two sheaves $\mathcal{F},\mathcal{G}$ on X valued in **Ab** is a natural transformation of $\eta : \mathcal{F} \to \mathcal{G}$ of the presheaves. A sheaf morphism

[4] Kernels, equalizers, and pullbacks are all examples of limits. See [Rie17, §3] for more background.

consists of a collection of morphisms $\eta_U : \mathcal{F}(U) \to \mathcal{G}(U)$ in **Ab** for every $U \in \mathrm{Open}(X)$ such that whenever there is an inclusion $U \hookrightarrow V$, there is a commutative diagram

$$\begin{array}{ccc} \mathcal{F}(V) & \xrightarrow[\eta_V]{} & \mathcal{G}(V) \\ \downarrow{\scriptstyle \rho_{V,U}} & & {\scriptstyle \rho_{V,U}}\downarrow \\ \mathcal{F}(U) & \xrightarrow{\eta_U} & \mathcal{G}(U) \end{array}$$

Example 2.4.9 (Constant and Locally Constant Sheaves)**.** The pre-sheaf that assigns the vector space $\Bbbk$ to every open set is not a sheaf because the value that should be associated to an open set $U = U_1 \cup U_2$ with two components should be $\Bbbk^2$. The sheafification of this constant pre-sheaf is called the **constant sheaf** and is usually written $\Bbbk$ as well. One can replace $\Bbbk$ with any vector space $\mathbb{V}$ and obtain an analogous constant sheaf. Moreover, if a sheaf $\mathcal{F} \in \mathbf{Shv}(X)$ has the property that for each point $x \in X$ there is a neighborhood U_x so that $\mathcal{F}$ restricts to a constant sheaf on U_x, then one calls $\mathcal{F}$ a **locally constant sheaf** or **local system**.

Definition 2.4.10. Let $\mathcal{C}$ be a complete "data" category, i.e., all limits in $\mathcal{C}$ exist. Denote the category of pre-sheaves and sheaves on X valued in **Dat** by $\mathbf{PShv}(X; \mathbf{Dat})$ and $\mathbf{Shv}(X; \mathbf{Dat})$, respectively. Since every sheaf is a pre-sheaf there is a natural inclusion of categories

$$\mathrm{ps} : \mathbf{Shv}(X; \mathbf{Dat}) \hookrightarrow \mathbf{PShv}(X; \mathbf{Dat}).$$

Under suitable hypotheses[5] on **Dat**, the ps has an "inverse"[6] called **sheafification**

$$\mathrm{sh} : \mathbf{PShv}(X; \mathbf{Dat}) \to \mathbf{Shv}(X; \mathbf{Dat}).$$

Typically we set $\mathbf{Dat} = \mathbf{Vect}$ and let $\mathbf{Shv}(X)$ denote the category of sheaves of vector spaces.

Remark 2.4.11 (Preservation of Stalks)**.** It is a fact that sheafification preserves stalks, so the "local picture" of a pre-sheaf $\mathcal{F}$ is unchanged by this process and only the failure of the local-to-global principle is repaired.

[5] See condition (17.4.1) in [KS94].

[6] i.e. a left adjoint, see [Rie17, §4].

Definition 2.4.12 (Pushforward and Pullback of Sheaves)**.** Suppose $f : X \to Y$ is a continuous map of spaces, $\mathcal{F}$ is a sheaf on X, and $\mathcal{G}$ is a sheaf on Y. The **pushforward** (or direct image) of $\mathcal{F}$ along f, written $f_*\mathcal{F}$ is the sheaf that assigns to each open set $V \in \text{Open}(Y)$ the value $\mathcal{F}(f^{-1}(V))$. Dually, the **pullback** (or inverse image) of $\mathcal{G}$ along f, written $f^*\mathcal{G}$, is the sheaf associated to the pre-sheaf that assigns to every open set $U \in \text{Open}(X)$ the "stalk at $f(U)$":

$$\text{sh}\left[U \mapsto \varinjlim_{V \supseteq f(U)} \mathcal{G}(V)\right]$$

These define functors

$$f_* : \mathbf{Shv}(X) \to \mathbf{Shv}(Y) \qquad \text{and} \qquad f^* : \mathbf{Shv}(Y) \to \mathbf{Shv}(X).$$

We now isolate a very important class of pushforward sheaves, which we call Leray sheaves.

Definition 2.4.13 (Leray Sheaves)**.** Suppose $f : Y \to X$ is a proper continuous map of spaces, then the i^{th} **Leray sheaf of** f, written $R^i f_* \Bbbk$, is the sheaf associated to the pre-sheaf

$$\text{sh}\left[U \in \text{Open}(X) \mapsto H^i(f^{-1}(U); \Bbbk)\right].$$

For $i = 0$ this is just the pushforward of the constant sheaf, but for $i \geqslant 1$ this definition arises naturally from the "derived" perspective considered in Section 2.4.2.

Remark 2.4.14 (Cohomology of the Fiber)**.** As Remark 2.4.11 indicates, the stalk of the Leray pre-sheaf is preserved under sheafification. Consequently, if

$$\varinjlim_{U \ni x} H^i(f^{-1}(U)) \cong H^i(f^{-1}(x)),$$

then this sheaf records the cohomology of the fiber. This last consequence follows from properness of f, which plays a key role in the proper base change theorem [Ive86].

2.4.2 Derived Sheaf Theory

In this section, instead of considering abelian groups as we did previously, we will work with abelian categories. Abelian categories provide a general framework for homological algebra; similar to abelian groups, they have kernels, cokernels, and various desirable properties such as the validity of the snake lemma.

Definition 2.4.15 (Chain complexes)**.** A **complex** in an abelian category $\mathbf{A}$ is a sequence $\mathbf{A}^\bullet$ of objects and maps, called differentials

$$\cdots \longrightarrow A^0 \xrightarrow{\delta^0} A^1 \xrightarrow{\delta^1} A^2 \longrightarrow \cdots$$

such that $\delta_{i+1} \circ \delta_i = 0$ for all i. A chain map $f : (\mathbf{A}^\bullet, \delta_\mathbf{A}) \to (\mathbf{B}^\bullet, \delta_\mathbf{B})$ of complexes consists of maps $f^i : A^i \to B^i$ that commute with the differentials. The complexes in $\mathbf{A}$ form an abelian category and it is denoted by $\mathcal{C}(\mathbf{A})$.

We often work with the category of complexes in $\mathbf{A}$ that are bounded, i.e. there exist integers n_1, n_2 such that $A^k = 0$ for all $k > n_1$ and $k < n_2$. Denote the category of bounded complexes of objects in $\mathbf{A}$ as $\mathcal{C}^b(\mathbf{A})$. The definitions below make sense if we replace $\mathcal{C}^b(\mathbf{A})$ with $\mathcal{C}(\mathbf{A})$.

The object $H^i(\mathbf{A}^\bullet) = \ker(\delta^i)/\operatorname{im}(\delta^{i-1})$ is the i-th cohomology of $\mathbf{A}^\bullet$.

Definition 2.4.16 (Quasi-isomorphism)**.** A chain map $f : \mathbf{A}^\bullet \to \mathbf{B}^\bullet$ is a **quasi-isomorphism** if f induces isomorphisms on $H^i(f) : H^i(\mathbf{A}^\bullet) \to H^i(\mathbf{B}^\bullet)$ for all i.

Define the shift of a complex, $\mathbf{A}[n]^\bullet$ as the complex with objects $A[n]^k = A^{k+n}$ with differentials $\delta^i_{\mathbf{A}[n]} = (-1)^n \delta^{i+n}_\mathbf{A}$.

Definition 2.4.17 (Mapping cone)**.** The **mapping cone** of a chain map $f : \mathbf{A}^\bullet \to \mathbf{B}^\bullet$ is the complex $C(f) = \mathbf{A}[1] \oplus \mathbf{B}$ with differentials $\delta^i(a, b) = (-\delta^{i+1}_\mathbf{A}(a), f^{i+1}(a) + \delta^i_\mathbf{B}(b))$.

Definition 2.4.18 (Homotopy Category)**.** Associated to $\mathcal{C}^b(\mathbf{A})$ is the **homotopy category** $\mathcal{K}^b(\mathbf{A})$ of complexes, which has the same objects as $\mathcal{C}^b(\mathbf{A})$, but where morphisms are homotopy classes of chain maps.

We now recall one definition of the derived category, as there are many.

Definition 2.4.19 (Derived category)**.** Let $\mathcal{Q}$ denote the class of quasi-isomorphisms. The **bounded derived category** of $\mathbf{A}$ is the localization of $\mathcal{K}^b(\mathbf{A})$ at the collection of morphisms $\mathcal{Q}$, i.e.

$$\mathcal{D}^b(\mathbf{A}) := \mathcal{K}^b(\mathbf{A})[\mathcal{Q}^{-1}].$$

More precisely, $\mathcal{D}^b(\mathbf{A})$ is endowed with a tautological functor $j : \mathcal{K}^b(\mathbf{A}) \to \mathcal{D}^b(\mathbf{A})$, and has the universal property that given any functor$F : \mathcal{K}^b(\mathbf{A}) \to \mathbf{B}$, that sends all quasi-isomorphisms in $\mathcal{K}^b(\mathbf{A})$ to isomorphisms in $\mathbf{B}$, then there exists a unique functor $F' : \mathcal{D}^b(\mathbf{A}) \to \mathbf{B}$ such that $F = F' \circ j$.

Remark 2.4.20. An alternative definition of the derived category makes use of the assumption that $\mathbf{A}$ has enough injectives, i.e. every object in $\mathbf{A}$ has an injective resolution or, said differently, every object in $\mathbf{A}$ is quasi-isomorphic to a complex of injective objects. Under this assumption, the derived category of $\mathbf{A}$ is equivalently defined as the homotopy category of injective objects in $\mathbf{A}$, i.e.

$$\mathcal{D}^b(\mathbf{A}) \simeq \mathcal{K}^b(\mathrm{Inj} - \mathbf{A}).$$

Remark 2.4.20 provides an easier-to-understand prescription for working with the derived category. One simply takes an object, e.g. sheaf, replaces it with its injective resolution, and works with the resolution instead.

The homotopy category $\mathcal{K}(\mathbf{A})$ and $\mathcal{D}(\mathbf{A})$ are not abelian categories. However, they are additive categories; i.e., the zero object exists, categorical products exist, and hom sets are abelian groups. However, not every morphism may exhibit kernels and cokernels, so an exact sequence does not make sense. Instead, the homotopy category and derived category have an extra structure described by the distinguished collection of *exact triangles* and are examples of *triangulated categories.* These distinguished exact triangles can be thought of as a replacement for exact sequences in abelian categories.

A triangle in any additive category $\mathcal{C}$ is a sequence of three composable maps:

$$\mathbf{A} \to \mathbf{B} \to \mathbf{C} \to \mathbf{A}[1].$$

A morphism of triangles f is a commutative diagram

$$\begin{array}{ccccccc} \mathbf{A} & \longrightarrow & \mathbf{B} & \longrightarrow & \mathbf{C} & \longrightarrow & \mathbf{A}[1] \\ {\scriptstyle f}\downarrow & & \downarrow{\scriptstyle f} & & \downarrow{\scriptstyle f} & & \downarrow{\scriptstyle f} \\ \mathbf{A}' & \longrightarrow & \mathbf{B}' & \longrightarrow & \mathbf{C}' & \longrightarrow & \mathbf{A}'[1] \end{array}$$

Remark 2.4.21. A triangle is also sometimes displayed as follows

$$\begin{array}{ccc} \mathbf{A} & \longrightarrow & \mathbf{B} \\ & {\scriptstyle +1}\nwarrow & \downarrow \\ & & \mathbf{C} \end{array}$$

Definition 2.4.22 (Distinguised triangles in $\mathcal{K}(\mathbf{A})$)**.** A **distinguished triangle** in $\mathcal{K}(\mathbf{A})$ is a triangle of the form,

$$\mathbf{A} \xrightarrow{f} \mathbf{B} \xrightarrow{i} C(f) \xrightarrow{p} \mathbf{A}[1],$$

where $i(b) = (0, b)$ and $p(a, b) = a$. An **exact triangle** in $\mathcal{K}(\mathbf{A})$ is a triangle isomorphic to a distinguished triangle.

Definition 2.4.23 (Exact triangles in $\mathcal{D}(\mathbf{A})$)**.** An **exact triangle** in $\mathcal{D}(\mathbf{A})$ is a triangle isomorphic in $\mathcal{D}(\mathbf{A})$ to a distinguished triangle in $\mathcal{K}(\mathbf{A})$.

A benefit of the derived category is that exact sequences in $\mathcal{C}(\mathbf{A})$ gives exact triangles in $\mathcal{D}(\mathbf{A})$, which is not true in $\mathcal{K}(\mathbf{A})$. Another feature of the derived category is that for every exact triangle in $\mathcal{D}(\mathbf{A})$

$$\mathbf{A} \to \mathbf{B} \to \mathbf{C} \to \mathbf{A}[1],$$

there is a long exact sequence of cohomology

$$\cdots \to H^0(\mathbf{A}) \to H^0(\mathbf{B}) \to H^0(\mathbf{C}) \to H^1(\mathbf{A}) \to \cdots.$$

Remark 2.4.24. By specializing to the setting where the abelian category is the category of abelian sheaves (i.e., sheaves valued in abelian categories), we obtain derived sheaf theory. The category of abelian sheaves has enough injectives, so Remark 2.4.20 applies.

An additive covariant functor $\mathcal{F}$ between two abelian categories is left exact if whenever $0 \to A \to B \to C \to 0$ is exact then $0 \to \mathcal{F}(A) \to \mathcal{F}(B) \to \mathcal{F}(C)$ is exact. The functor is exact if short exact sequences remains exact after the application of the functor.

Definition 2.4.25 (Derived functors)**.** Suppose $F : \mathbf{A} \to \mathbf{B}$ is an additive and left-exact functor, i.e. it commutes with direct sums and preserves kernels, then the **total right derived functor** of F, written $RF : \mathcal{D}^b(\mathbf{A}) \to \mathcal{D}^b(\mathbf{B})$ is defined by

$$RF(A^\bullet) := F(I^\bullet)$$

for $I^\bullet$ an injective resolution of $A^\bullet$. The cohomology objects $H^i(RF)$ denoted by R^iF are called the **classical right derived functors** of F .

Remark 2.4.26. An F-acylic object A is one that satisfies $R^iF(A) = 0$ for all $i > 0$. In general, one can substitute $I^\bullet$ in Definition 2.4.25 with any F-acyclic resolution of $A^\bullet$. Such resolutions are said to be **adapted** to F.

Recall that $\Gamma(X, -) : \mathbf{Shv}(X) \to \mathbf{Ab}$ is the global section functor that takes a an abelian sheaf $\mathcal{F}$ and outputs an abelian group $\mathcal{F}(X)$. This functor is left exact and so we can consider the right derived functor of Γ.

Definition 2.4.27 (Sheaf Cohomology)**.** The ith cohomology of the abelian sheaf $\mathcal{F}$ on X is defined as

$$H^i(X, \mathcal{F}) = R^i\Gamma(F).$$

2.5 Constructible Sheaves and their Functions

We finish our review of preliminary material by showing how the topological transform—the Euler characteristic—are recovered by the derived perspective. This is accomplished via the sheaf-to-function correspondence, which is best studied for constructible sheaves.

Definition 2.5.1 (Constructible and Cellular Sheaves)**.** A sheaf $\mathcal{F} \in \mathbf{Shv}(X)$ is said to be **constructible** if there is a decomposition of X into definable subsets $\{X_\alpha\}$ so that for each α the pullback of $\mathcal{F}$ along the inclusion $i_\alpha : X_\alpha \hookrightarrow X$ produces a locally constant sheaf $i_\alpha^* \mathcal{F} \in \mathbf{Shv}(X_\alpha)$, cf. Example 2.4.9. When the subsets $\{X_\alpha\}$ are cells in a triangulation of X, we can equivalently express $\mathcal{F}$ as a **cellular sheaf**, which simply assigns a vector space to each cell X_α and a linear map to each pair $X_\alpha \subseteq \overline{X_\beta}$; see [She85] for the first description of cellular sheaves and [Cur14b] for a modern treatment.

In similar spirit, if $\mathcal{F}^\bullet$ is a complex of sheaves, e.g., a derived sheaf, then it is said to be **cohomologically constructible** if the cohomology sheaves $\mathcal{H}^i\mathcal{F}^\bullet$ are constructible for every i. A derived cellular sheaf then associates to each cell in a cellulation a complex of vector spaces and a chain map for every face-relation pair $X_\alpha \subseteq \overline{X_\beta}$.

We now show how to associate to a constructible function to a constructible sheaf.

Definition 2.5.2 (Euler-Poincaré and Hilbert Functions)**.** Let $\mathcal{F}^\bullet$ be a complex of cohomologically constructible sheaves on $\mathbb{X}$. The **local Euler-Poincaré index** is the piecewise constant integer-valued function defined by

$$h(x) := \chi(\mathcal{F}^\bullet)(x) = \sum_i (-1)^i \dim \mathcal{H}^i(\mathcal{F}^\bullet)_x = \sum_i (-1)^i \dim(H^i \mathcal{F}_x^\bullet)$$

The first equality considers the stalks of the cohomology sheaves of $\mathcal{F}^\bullet$ and the second equality considers the cohomology of the stalk complex. By Remark 2.4.11 these are isomorphic and thus have the same dimension and Euler characteristic. In the simplest setting—constructible sheaves that are concentrated in a single cohomological degree—the local Euler-Poincaré index is just the **Hilbert function**—it records the dimension of the stalks of a sheaf.

2.6 The Persistent Homology Transform

We now have enough language to describe one of the central objects of thesis.

Definition 2.6.1 (PHT: Sheaf Version, cf. [CMT22a])**.** Let $M \in \mathcal{CS}(\mathbb{R}^d)$ be a constructible set. Associated to M is the **auxiliary total space**

$$Z_M := \{(x, v, t) \in M \times \mathbb{S}^{d-1} \times \mathbb{R} \mid x \cdot v \leqslant t\}.$$

The i^{th} **persistent homology transform sheaf** of M, written $\mathrm{PHT}^i(M)$, is the i^{th} Leray sheaf of the map $f_M : Z_M \to \mathbb{S}^{d-1} \times \mathbb{R}$ that projects onto the last two factors. Since M is compact and f_M is a projection, we see that f_M is proper. By Remark 2.4.14, the stalk of the i^{th} Leray sheaf at $(v, t) \in \mathbb{S}^{d-1} \times \mathbb{R}$ is isomorphic to the i^{th} cohomology of the fiber i.e., $H^i({f_M}^{-1}(v, t))$. See Figure 2.6 for a stalk-wise picture of PHT^0 of the shape 'V', which was considered from the map perspective in Figure 2.5.

Remark 2.6.2. As a consequence of the Trivialization Theorem [Van98b, Ch 9, 1.2] (or Lemma 3.4 in [CMT22a]) the PHT sheaf is constructible. This means that we can find a finite partition $\{A_i\}_{i=1}^N$ of $\mathbb{S}^{d-1} \times \mathbb{R}$ such that on each A_i, the PHT in degree k is a vector bundle with fibers $H^k(M_{v,t}; \Bbbk)$ where $(v, t) \in A_i$. Essentially, for every direction and height (v, t) in A_i we attach the cohomology of $M_{v,t}$ as the "fiber." The sheaf axiom provides a way to pass between two fibers from different partition elements. So one may view the PHT sheaf as stitching together vector bundles on different pieces of $S^{d-1} \times \mathbb{R}$.

Remark 2.6.3. If we restrict the sheaf $\mathrm{PHT}^i(M)$ to the subspace $\{v\} \times \mathbb{R}$, then one obtains a constructible sheaf that is equivalent to the persistent (co)homology of the filtration of M viewed in the direction of v. The persistence diagram in degree i is simply the expression of this restricted sheaf in terms of a direct sum of indecomposable sheaves.

Remark 2.6.4 (Iterated Pushforwards and the Equivariance Property)**.** An advantage of the sheaf-theoretic definition of the PHT is that there are further operations that can be performed on it. For example, one can iterate the Leray sheaf construction for any change of coordinates of $\mathbb{S}^{d-1} \times \mathbb{R}$. For example, a rigid motion $g \in \mathrm{SO}(d)$, induces an action on the corresponding PHT sheaves $\mathrm{PHT}^i(M)$ via pushforward along g. In particular, we have

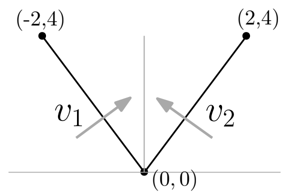

(a) M is an embedded 'V' shape.

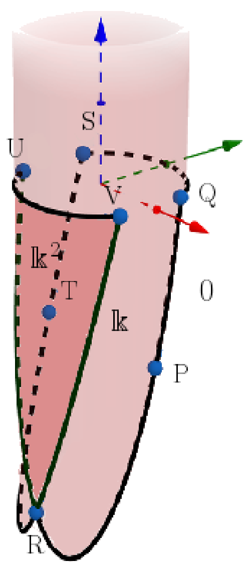

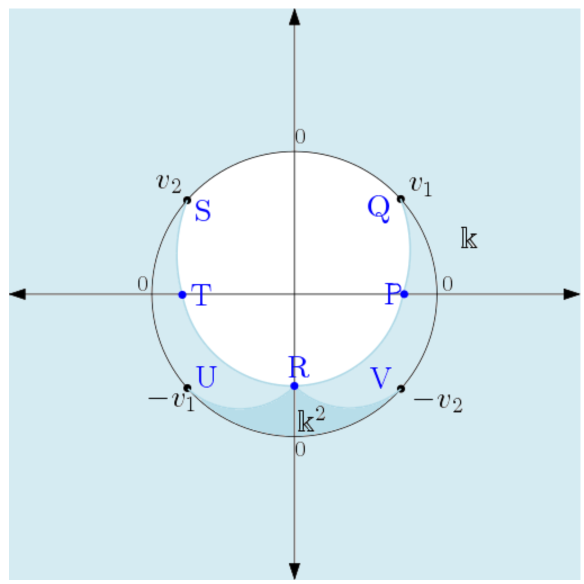

(b) $\mathrm{PHT}^0(M)$

(c) $\mathrm{PHT}^0(M)$ viewed on the plane.

FIGURE 2.6: Figure (B) is a visualization of the PHT sheaf in degree 0 of M = 'V' on the cylinder $\mathbb{S}^1 \times \mathbb{R}$. The dark pink region represents where there are two connected components and the light pink region represents one connected components. The green arrow represents the direction $e_2 = (0, 1)$ and the the red axis corresponds looking to the right $e_1 = (1, 0)$. To visualize the sheaf on the plane, as in Figure (C), we map the cylinder $\mathbb{S}^1 \times \mathbb{R}$ to the plane $\mathbb{R}^2$ by sending the circle at $t = -\infty$ to $(0, 0)$ and the circle at $t = +\infty$ goes to infinity in each direction.

the following equivariance formula:

$$\mathrm{PHT}^i(g \cdot M) = g_* \,\mathrm{PHT}^i(M)$$

To prove the above formula, it suffices to show for any test open $U \subseteq \mathbb{S}^{d-1} \times \mathbb{R}$,

$$f_{gM}^{-1}(U) = f_M^{-1}\left(g^{-1}(U)\right).$$

The equality of the sets follows from the definition of the projection map f_M (Definition 2.6.1) and the fact that $gx \cdot v = x \cdot g^{-1}v$ for any $x \in M$ and $v \in S^{d-1}$. In other words, filtering a shape rotated by g in direction v is equivalent to filtering the original shape in direction $g \cdot v$. A similar expression can be shown for translations. In particular, if g represents translation by $T \in \mathbb{R}^d$. Then it induces a map $g : \mathbb{S}^{d-1} \times \mathbb{R} \to \mathbb{S}^{d-1} \times \mathbb{R}$ by $(v,t) \mapsto (v, t + T \cdot v)$.

We can now define the persistent homology transform as a derived sheaf.

Definition 2.6.5 (PHT: Derived Version, cf. [CMT22a])**.** Let $M \in \mathcal{CS}(\mathbb{R}^d)$ be a constructible set. Let Z_M be the auxiliary space construction from Definition 2.6.1. Let $f_{M*} : \mathbf{Shv}(Z_M) \to \mathbf{Shv}(\mathbb{S}^{d-1} \times \mathbb{R})$ be the pushforward (or direct image) functor along the projection map $f_M : Z_M \to \mathbb{S}^{d-1} \times \mathbb{R}$. The **derived PHT sheaf** is

$$\mathrm{PHT}(M) := Rf_{M*}\mathbb{k}_{Z_M} \in \mathcal{D}^b(\mathbf{Shv}(\mathbb{S}^{d-1} \times \mathbb{R})).$$

More explicitly we can describe this right-derived pushforward as follows: For a topological space $\mathbb{X}$ we let $\mathcal{S}^p(U;\mathbb{k})$ denote the group of singular p-cochains of $U \subset \mathbb{X}$ with coefficients in $\mathbb{k}$. Define $\mathscr{S}^p(X,\mathbb{k}) = \mathrm{sh}(U \mapsto \mathcal{S}^p(U;\mathbb{k}))$ where sh stands for sheafification. The constant sheaf $\mathbb{k}_{Z_M}$ admits a flabby resolution by singular cochains:

$$0 \to \mathbb{k}_{Z_M} \to \mathscr{S}^0(Z_M;\mathbb{k}) \to \mathscr{S}^1(Z_M;\mathbb{k}) \to \mathscr{S}^2(Z_M;\mathbb{k}) \to \cdots$$

Because flabby resolutions form an adapted class for the pushforward functor [Bre12] we can describe $\mathrm{PHT}(M)$ as the pushforward of the complex of sheaves of singular cochains:

$$Rf_{M*}\mathbb{k}_{Z_M} := f_{M*}\mathscr{S}^0(Z_{M_i};\mathbb{k}) \to f_{M*}\mathscr{S}^1(Z_{M_i};\mathbb{k}) \to \cdots$$

Taking i^{th} cohomology of this complex—written $\mathcal{H}^i$ in the category of sheaves—produces the i^{th} PHT sheaf of Definition 2.6.1.

Remark 2.6.6. We remark that implicitly we have made use of the fact that singular cohomology is isomorphic to sheaf cohomology with constant coefficients. This is always true when X is paracompact and locally contractible, which is our case. More properly we would have written:

$$0 \to \Bbbk_{Z_M} \to \mathscr{S}^0(Z_M; \Bbbk_{Z_M}) \to \mathscr{S}^1(Z_M; \Bbbk_{Z_M}) \to \mathscr{S}^2(Z_M; \Bbbk_{Z_M}) \to \cdots$$

However, since we are in the derived category, we can replace it with a sequence that is quasi-isomorphic to it.

3. A Homotopy Sheaf on Shape Space

This chapter is derived from Chapter 3 of [ACM23], with a more detailed exposition of spectral sequences and infinity categories.

As mentioned in the introduction, we want to build a shape space using a sheaf-theoretic construction on the poset of constructible sets $\mathcal{CS}(\mathbb{R}^d)$. Naively one would like to prove that the association

$$\mathcal{F} : \mathcal{CS}(\mathbb{R}^d)^{op} \to \mathcal{D}^b(\mathbf{Shv}(\mathbb{S}^{d-1} \times \mathbb{R})) \qquad M \mapsto \mathrm{PHT}(M)$$

is a sheaf, but there are two main obstacles.

The first obstacle is that a topology on $\mathcal{CS}(\mathbb{R}^d)$ needs to be specified. Although sheaves on posets are well-defined via the Alexandrov topology—see [Cur14a] for a modern treatment—the poset under consideration is infinite and using the Alexandrov topology here would imply that a shape can be determined via a cover by its points; this is clearly impossible as there is not enough of an interface between points to determine homology. The second obstacle is fatal for a naive sheaf-theoretic approach: the pre-sheaf $U \mapsto H^i(U)$ is not a sheaf for $i \geqslant 1$. Indeed, the connecting homomorphism in the Mayer-Vietoris long-exact sequence quantifies precisely the failure of the sheaf axiom. Both of these obstacles are addressed via tools from "higher" sheaf theory: Grothendieck topologies, and homotopy sheaves. This is synonymous with the motto for derived categories "Homology Bad, Complexes Good."

3.1 Inclusion-Exclusion for the ECT

As a reminder, Definition 2.2.15 presented the Euler characteristic transform as the Radon transform $\mathcal{R}_S$ associated to a particular relation $S \subseteq X \times Y$, where $X = \mathbb{R}^d$ and $Y = \mathbb{S}^{d-1} \times \mathbb{R}$. One of the stated properties of the Radon transform is that it defines a group homomorphism

$$\mathcal{R}_S : \mathcal{CF}(X) \to \mathcal{CF}(Y) \qquad \text{i.e.} \qquad \mathcal{R}_S(\phi + \psi) = \mathcal{R}_S(\phi) + \mathcal{R}_S(\psi).$$

This allows us to prove an immediate inclusion-exclusion principle for the ECT.

Theorem 3.1.1. *For a finite cover $\mathcal{M} = \{M_i\}_{i\in\Lambda}$ of $M \subset \mathbb{R}^d$ by constructible subsets*

$$\mathrm{ECT}(M) = \sum_{I\subset\Lambda} (-1)^{|I|+1}\,\mathrm{ECT}(M_I)$$

where each M_I denotes the intersection $M_{i_1} \cap M_{i_2} \cap \cdots \cap M_{i_k}$ for $I = (i_1, ..., i_k)$.

Proof. The inclusion-exclusion principle allows us to write the indicator function as

$$\mathbb{1}_M = \sum_{I\subset\Lambda} (-1)^{|I|+1}\mathbb{1}_{M_I}.$$

Linearity of the Radon transform then implies that

$$\mathcal{R}_S\mathbb{1}_M = \sum_{I\subset\Lambda} (-1)^{|I|+1}\mathcal{R}_S\mathbb{1}_{M_I},$$

which is the expression using ECTs written above. □

Remark 3.1.2. If we take the sheaf-to-function correspondence of Section 2.5 seriously, then Theorem 3.1.1 should be viewed as a decategorification of a deeper sheaf-theoretic result:

$$\mathrm{PHT}(M) \cong \text{“summary of”} \left[\prod \mathrm{PHT}(M_i) \to \prod \mathrm{PHT}(M_{ij}) \to \prod \mathrm{PHT}(M_{ijk}) \to \cdots\right]$$

However, unlike the sheaf axiom of Definition 2.4.6, this summary operation cannot only use the M_i and their pairwise intersections M_{ij}. Indeed the inclusion-exclusion result of Theorem 3.1.1 says that *all higher order terms must be considered.* This summary operation is accomplished by the homotopy limits of Definition **??** or via limits in the derived ∞-category (Section 3.7).

3.2 Sites and Homotopy Sheaves

Grothendieck topologies provide a way of generalizing sheaves to contravariant functors on a general category $\mathcal{C}$. Covers of an open set are replaced with collections of morphisms that have certain “cover-like” properties.

Definition 3.2.1 (Grothendieck Pre-topology, cf.[Art62])**.** Let $\mathcal{C}$ be a category with pullbacks. A **basis for a Grothendieck topology** (or a **pre-topology**) on $\mathcal{C}$ requires specifying for each object $U \in \mathcal{C}$ a collections of **admissible covers** of U. This collection of covers must be closed under the following operations:

1. (Isomorphism) If $f : U' \to U$ is an isomorphism then $\{f : U' \to U\}$ is a cover.

2. (Composition) If $\{f_i : U_i \to U\}$ is a cover of U and if for each i we have a cover $\{g_{i,j} : U_{i,j} \to U_i\}$ then the composition $\{f_i \circ g_{i,j} : U_{i,j} \to U\}$ is also a cover.

3. (Base Change) If $\{f_i : U_i \to U\}$ is a cover and $V \to U$ is any morphism then the pullback $\{\pi_i : V \times_U U_i \to V\}$ is a cover as well

As the name suggests, the above data specifies a genuine Grothendieck topology on $\mathcal{C}$. A category equipped with a Grothendieck topology is known as a **site**.

Remark 3.2.2 (Sheaves on Sites)**.** The classical definition of a pre-sheaf and sheaf can now be generalized to a site. A functor $\mathcal{F} : \mathcal{C}^{op} \to \mathbf{Dat}$ is a **pre-sheaf**. If **Dat** has all limits, we say a pre-sheaf is a **sheaf** if for every object $U \in \mathcal{C}$ and cover $\mathcal{U} = \{f_i : U_i \to U\}$

$$\mathcal{F}(U) = \varprojlim \left[\prod_i \mathcal{F}(U_i) \rightrightarrows \prod_{i,j} \mathcal{F}(U_{ij}) \right].$$

Here equality means isomorphic up to a unique isomorphism and U_{ij} is the pullback of $f_j : U_j \to U$ along $f_i : U_i \to U$ for any pair of morphisms f_i and f_j that participate in the cover $\mathcal{U}$.

Unfortunately the functor $\mathcal{F}$ specified in Theorem 1.3.1 is valued in the derived category of sheaves on $\mathbb{S}^{d-1} \times \mathbb{R}$. As discussed in Section 2.4.2, the derived category $\mathcal{D}^b(\mathbf{A})$ of an abelian category $\mathbf{A}$ is not abelian. Candidate kernels and co-kernels do not have canonical inclusion and projection maps, but one can work with *distinguished triangles* instead (refer to Definition 2.4.23). More generally, we can describe a sheaf axiom whenever the notion of a *homotopy limit* makes sense in the target category **Dat**. These limits allow for some homotopical leeway and is not as rigid as the ordinary limit. For example, the ordinary colimit of $* \leftarrow \mathbb{S}^2 \rightarrow *$ is a point and has no recollection of the homotopy type of $\mathbb{S}^2$ but the homotopy limit is the suspension of the sphere.

Definition 3.2.3 (Homotopy Limits)**.** Given an inverse system of objects (K_n, f_n) in $\mathcal{D}^b(\mathbf{A})$

$$\cdots \xrightarrow{f_{n-1}} K_{n-1} \xrightarrow{f_n} K_n \xrightarrow{f_{n+1}} \cdots$$

an object K is a **homotopy limit** if there is a distingushed triangle in the derived category

$$K \to \prod_n K_n \xrightarrow{\text{shift}} \prod_n K_n \to K[1].$$

The shift map being given by $(k_n) \mapsto (k_n - f_{n-1}(k_{n-1}))$. We note that the homotopy limit is not necessarily unique and so we say that K is *a* homotopy limit rather than it is *the* homotopy limit.

Remark 3.2.4. This definition applies more generally to any triangulated category. We specialize to the case of the derived category $\mathcal{D}^b(\mathbf{A})$.

We can now define sheaves valued in the derived category.

Definition 3.2.5 (Homotopy Sheaf)**.** A pre-sheaf $\mathcal{F} : \mathcal{C}^{op} \to \mathcal{D}^b(\mathbf{A})$ is a **homotopy sheaf** (or satisfies **Čech descent** [DHI04]) if for every object $U \in \mathcal{C}$ and cover $\mathcal{U} = \{U_i \to U\}$ the following map is a quasi-isomorphism:

$$\mathcal{F}(U) \xrightarrow{\simeq} \text{holim}\left[\prod_i \mathcal{F}(U_i) \rightrightarrows \prod_{i,j} \mathcal{F}(U_{ij}) \rightrightarrows \cdots\right]$$

We can now prove our main results.

Lemma 3.2.6. *The poset $\mathcal{CS}(\mathbb{R}^d)$ admits the structure of a site.*

Proof. For every object $M \in \mathcal{CS}(\mathbb{R}^d)$ we say that $\{M_i \hookrightarrow M\}$ is a covering if it is a finite closed cover of M in the usual sense, i.e. $\cup M_i = M$. Pullbacks exist by virtue of the fact that o-minimal sets are closed under intersection. □

Lemma 3.2.7. *The following assignment is a pre-sheaf*

$$\mathcal{F} : \mathcal{CS}(\mathbb{R}^d)^{op} \to \mathcal{D}^b(\mathbf{Shv}(\mathbb{S}^{d-1} \times \mathbb{R})) \qquad M \mapsto \mathrm{PHT}(M)$$

where $\mathrm{PHT}(M)$ *is the derived sheaf version of the PHT; see Definition 2.6.5.*

Proof. We want to show that $\mathcal{F}$ is a contravariant functor. Let $\iota : M_1 \hookrightarrow M_2$ be an inclusion of constructible sets of $\mathbb{R}^d$. Note that M_1 is a closed subspace of M_2. This induces an inclusion of the auxiliary total spaces $\iota : Z_{M_1} \hookrightarrow Z_{M_2}$ of Definition 2.6.1. This in turn determines a morphism of pre-sheaves $f_{M_2 *}\mathcal{S}^j(Z_{M_2}; \mathbb{k}) \to f_{M_1 *}\mathcal{S}^j(Z_{M_1}; \mathbb{k})$ for all j. To see this, take an open $U \subset \mathbb{S}^{d-1} \times \mathbb{R}$ and observe that $f_{M_1}^{-1}(U)$ is open in $f_{M_2}^{-1}(U)$. More generally, for $U \subset V$ in $\mathbb{S}^{d-1} \times \mathbb{R}$ we have a commutative diagram of cochain groups:

$$\begin{array}{ccc} \mathcal{S}^j(f_{M_1}^{-1}(U); \mathbb{k}) & \longleftarrow & \mathcal{S}^j(f_{M_2}^{-1}(U); \mathbb{k}) \\ \uparrow & & \uparrow \\ \mathcal{S}^j(f_{M_1}^{-1}(V); \mathbb{k}) & \longleftarrow & \mathcal{S}^j(f_{M_2}^{-1}(V); \mathbb{k}) \end{array}$$

Since sheafification is a functor, we get a morphism $f_{M_2 *}\mathscr{S}^j(Z_{M_2}; \mathbb{k}) \to f_{M_1 *}\mathscr{S}^j(Z_{M_1}; \mathbb{k})$ for all j. These fit together into a morphism between complexes of sheaves:

$$\begin{array}{ccccccc} f_{M_2 *}\mathscr{S}^0(Z_{M_2}; \mathbb{k}) & \longrightarrow & f_{M_2 *}\mathscr{S}^1(Z_{M_2}; \mathbb{k}) & \longrightarrow & f_{M_2 *}\mathscr{S}^2(Z_{M_2}; \mathbb{k}) & \longrightarrow & \cdots \\ \downarrow & & \downarrow & & \downarrow & & \\ f_{M_1 *}\mathscr{S}^0(Z_{M_1}; \mathbb{k}) & \longrightarrow & f_{M_1 *}\mathscr{S}^1(Z_{M_1}; \mathbb{k}) & \longrightarrow & f_{M_1 *}\mathscr{S}^2(Z_{M_1}; \mathbb{k}) & \longrightarrow & \cdots \end{array}$$

The canonical functor from $\mathcal{C}^b(\mathbf{Shv}(\mathbb{S}^{d-1} \times \mathbb{R})) \to \mathcal{D}^b(\mathbf{Shv}(\mathbb{S}^{d-1} \times \mathbb{R}))$ then induces the desired restriction morphism between derived PHT sheaves:

$$\mathcal{F}(M_2) := Rf_{M_2 *}\mathbb{k}_{Z_{M_2}} \to Rf_{M_1 *}\mathbb{k}_{Z_{M_1}} =: \mathcal{F}(M_1)$$

□

The following is the main result of the paper, which was stated as Theorem 1.3.1 in the introduction. We give a direct proof below, but Remark **??** gives a more intuitive and computationally flavored proof using spectral sequences.

Theorem 3.2.8. *The pre-sheaf $\mathcal{F}$ of Lemma 3.2.7 is a homotopy sheaf; see Definition 3.2.5.*

Proof. We have already specified a Grothendieck topology on $\mathcal{CS}(\mathbb{R}^d)$ in Lemma 3.2.6. Let $\mathcal{M} = \{M_i\}_{i\in\Lambda}$ be a finite closed cover of M. Since $\mathcal{F}$ is a pre-sheaf we have an inverse system of derived sheaves:

$$\bigoplus_{I\subset\Lambda\, s.t.\, |I|=1} Rf_{M_I*}\Bbbk_{Z_{M_I}} \rightrightarrows \bigoplus_{J\subset\Lambda\, s.t.\, |J|=2} Rf_{M_J*}\Bbbk_{Z_{M_J}} \mathrel{\substack{\rightarrow\\\rightarrow\\\rightarrow}} \bigoplus_{K\subset\Lambda\, s.t.\, |K|=3} Rf_{M_K*}\Bbbk_{Z_{M_K}} \to \cdots$$

We wish to show that $Rf_{M*}\Bbbk_{Z_M}$ is the homotopy limit of the above inverse system of derived sheaves, i.e. we want to show that

$$Rf_{M*}\Bbbk_{Z_M} \simeq \operatorname{holim}\left[\bigoplus_{|I|=1} Rf_{M_I*}\Bbbk_{Z_{M_I}} \rightrightarrows \bigoplus_{|J|=2} Rf_{M_J*}\Bbbk_{Z_{M_J}} \mathrel{\substack{\rightarrow\\\rightarrow\\\rightarrow}} \cdots\right].$$

By replacing each $Rf_{M_I*}\Bbbk_{Z_{M_I}}$ with its flabby resolution by singular cochains it suffices to prove that the following is a distinguished triangle:

$$\begin{array}{ccc} f_{M*}\mathscr{S}^{\boldsymbol{\cdot}}(Z_M;\Bbbk) & \longrightarrow & \prod_n \bigoplus_{|I|=n} f_{M_I*}\mathscr{S}^{\boldsymbol{\cdot}}(Z_{M_I};\Bbbk) \\ & \nwarrow^{+1} & \downarrow \text{shift} \\ & & \prod_n \bigoplus_{|I|=n} f_{M_I*}\mathscr{S}^{\boldsymbol{\cdot}}(Z_{M_I};\Bbbk) \end{array}$$

To show this we consider the following maps of complexes of sheaves:

$$f_{M*}\mathscr{S}^{\boldsymbol{\cdot}}(Z_M) \to \prod_n \bigoplus_{|I|=n} f_{M_I*}\mathscr{S}^{\boldsymbol{\cdot}}(Z_{M_I}) \xrightarrow{\text{shift}} \prod_n \bigoplus_{|I|=n} f_{M_I*}\mathscr{S}^{\boldsymbol{\cdot}}(Z_{M_I})$$

Where we drop the coefficient $\Bbbk$ for convenience. For every $(v,t) \in \mathbb{S}^{d-1}\times\mathbb{R}$, these morphisms induce a sequence on stalks, which give rise to a sequence of cochain complexes

$$S^{\boldsymbol{\cdot}}(M_{v,t}) \to \prod_n \bigoplus_{|I|=n} S^{\boldsymbol{\cdot}}((M_I)_{v,t}) \xrightarrow{\text{shift}} \prod_n \bigoplus_{|I|=n} S^{\boldsymbol{\cdot}}((M_I)_{v,t}) \tag{3.1}$$

where $(M_I)_{v,t}$ is the intersection of M_I with the half-space $\{x \mid x\cdot v \leqslant t\}$. The kernel of the shift map at each stalk is clearly the cochain complex of small co-chains $S^{\boldsymbol{\cdot}}_{\mathcal{M}_{v,t}}(M_{v,t})$;

these are cochains supported on singular simplices that are individually contained in some cover element $(M_i)_{v,t}$ of the fiber $M_{v,t}$. Consequently, we have the following distinguished triangle

$$\begin{array}{ccc} f_{M*}\mathscr{S}^{\cdot}_{\mathcal{M}}(Z_M) & \longrightarrow & \prod_n \bigoplus_{|I|=n} f_{M_I*}\mathscr{S}^{\cdot}(Z_{M_I}) \\ & {\scriptstyle +1}\nwarrow & \big\downarrow {\scriptstyle \text{shift}} \\ & & \prod_n \bigoplus_{|I|=n} f_{M_I*}\mathscr{S}^{\cdot}(Z_{M_I}) \end{array}$$

where $\mathscr{S}^{\cdot}_{\mathcal{M}}(Z_M)$ is the sheaf of cochains supported in the cover.

We now show we can replace $f_{M*}\mathscr{S}^{\cdot}_{\mathcal{M}}(Z_M)$ with $f_{M*}\mathscr{S}^{\cdot}(Z_M)$ above. For any $(v,t) \in \mathbb{S}^{d-1} \times \mathbb{R}$, we show the inclusion of cochain complexes

$$S^{\cdot}_{\mathcal{M}_{v,t}}(M_{v,t}) \hookrightarrow S^{\cdot}(M_{v,t}) \tag{3.2}$$

is a quasi-isomorphism. Since a map of (derived) sheaves is an isomorphism if and only if it induces an isomorphism on stalks [KS94, Prop. 2.2.2], we can conclude $f_{M*}\mathscr{S}^{\cdot}_{\mathcal{M}}(Z_M)$ and $f_{M*}\mathscr{S}^{\cdot}(Z_M)$ are isomorphic.

To prove inclusion (3.2) is a quasi-isomorphism, we appeal to simplicial cohomology. By the Triangulation Theorem (Theorem 2.2.4) we can triangulate $M_{v,t}$ in a way that is subordinate to the closed cover $\{(M_I)_{v,t}\}$ for arbitrary (yet finite) intersections M_I. Simplicial cochains for this triangulation form a sub-cochain complex of $S^{\cdot}_{\mathcal{M}_{v,t}}(M_{v,t})$, but the triangulation can be used to compute cohomology of $M_{v,t}$. This completes the proof. □

Remark 3.2.9 (Proof via infinity categories.)**.** There is yet another perspective to the homotopy sheaf axiom. This can be seen by promoting the derived category of sheaves on $\mathbb{S}^{d-1} \times \mathbb{R}$ to the derived *infinity* category. In Section 3.7, we show that our desired sheaf axiom for finite covers is automatically satisfied.

3.3 Spectral Sequences

Theorem 3.2.8 proves that the functor $\mathcal{F}$ is a sheaf but does not give an easy way to glue the sheaf data on cover elements. The sheaf axiom says that the cohomology of the fiber $M_{v,t}$ can be built out of the cohomology of the cover elements but computing the limit of singular chain complexes of the cover elements can be a difficult task. So we appeal to spectral sequences to build up the PHT of a shape from its cover elements.

3.3.1 Overview of Spectral Sequences

Definition 3.3.1 (Cohomological spectral sequence)**.** A **cohomological spectral sequence** in an abelian category is a sequence of objects $E_r^{p,q}$ with $r \geqslant 0$ (known as the sequence of pages) together with differentials,

$$\delta_r^{p,q} : E_r^{p,q} \to E^{p+r,q-r+1}$$

such that $(\delta_r^{p,q})^2 = 0$ and $E_{r+1} = H^{*,*}(E_r, \delta_r)$, that is,

$$E_{r+1}^{p,q} = \ker(\delta_r : E_r^{p,q} \to E_r^{p+r,q-r+1}) \big/ \operatorname{im}\left(\delta_r : E_r^{p-1,q+r-1} \to E_r^{p,q}\right)$$

We often suppress (p, q) notation for the differentials when it is clear. A spectral sequences **converges** to the sequence of objects H^*, written as $E_r^{p,q} \implies H^{p+q}$ if,

1. For each bidegree (p, q) there exists r_0 such that $\delta_r^{p,q} = 0$ for $r \geqslant r_0$.

2. There is a filtration of H^*,

$$0 \subseteq \cdots \subseteq F_{-1}H^* \subseteq F_0H^* \subseteq F_1H^* \subseteq \cdots \subseteq H^*,$$

   so that for each n, $E_\infty^{p,q} := \varinjlim_r E_r^{p,q}$ is isomorphic to $F_pH^{p+q}/F_{p-1}H^{p+q}$.

We say that the spectral sequences **collapses** at E_{r_0} if there exists r_0 such that for all p, q and $r \geqslant r_0$, $E_r^{p,q} = E_\infty^{p,q}$. If the the spectral sequence is in the first quadrant, i.e. when $p < 0, q < 0$, $E_r^{p,q} = 0$ for all r, it is automatic that the spectral sequence collapses at some r_0, because eventually the maps δ_r extend into the zero regions and hence become zero.

Next, we give a statement of the Grothendieck spectral sequence, which originated in Grothendieck's famous Tŏhuku paper. It is a computational tool to compute the composition of the derived functors $\mathcal{G} \circ \mathcal{F}$ from the information of the functors $\mathcal{G}$ and $\mathcal{F}$, and can be thought of as *chain rule* for derived functors. The power of the spectral sequence lies in the fact that many spectral sequences encountered in practice are just instances of the Grothendieck spectral sequence.

Theorem 3.3.2 (Grothendieck Spectral Sequence)**.** *Let $\mathcal{F} : \mathbf{A} \to \mathbf{B}$ and $\mathcal{G} : \mathbf{B} \to \mathbf{C}$ be two left exact exact functors between abelian categories. Assume that $\mathbf{A}$ and $\mathbf{B}$ have enough injective objects and that $\mathcal{F}$ takes injective objects of $\mathbf{A}$ to $\mathcal{G}$-acyclic objects, then for each X in $\mathcal{D}^+(\mathbf{A})$, there exists a spectral sequence,*

$$E_2^{p,q} = R^p G(H^q(RF(X))) \implies H^{p+q}(R(G \circ F)(X))$$

3.3.2 Proof via Spectral sequences

We now show how Theorem 3.2.8 can be inferred from a simple spectral sequence argument.

Proof. Consider the resolution of the constant sheaf supported on Z_M (Theorem 4.4.1 in [God58b]),

$$\Bbbk_{Z_M} \to \oplus_{|I|=1} \Bbbk_{Z_{M_I}} \to \oplus_{|I|=2} \Bbbk_{Z_{M_I}} \to \cdots \tag{3.3}$$

This already proves, in essence, Čech descent for the PHT. More specifically, the homotopy sheaf axiom is witnessed via a first quadrant spectral sequence. Taking flabby resolutions of the sheaves by singular sheaves we get the following double complex:

$$\begin{array}{ccccccc}
\vdots & & \vdots & & \vdots & & \\
\mathscr{S}^1(Z_M;\Bbbk) & \longrightarrow & \oplus_{|I|=1}\mathscr{S}^1(Z_{M_I};\Bbbk) & \longrightarrow & \oplus_{|I|=2}\mathscr{S}^1(Z_{M_I};\Bbbk) & \cdots & \\
\uparrow & & \uparrow & & \uparrow & & \\
\mathscr{S}^0(Z_M;\Bbbk) & \longrightarrow & \oplus_{|I|=1}\mathscr{S}^0(Z_{M_I};\Bbbk) & \longrightarrow & \oplus_{|I|=2}\mathscr{S}^0(Z_{M_I};\Bbbk) & \cdots & \\
\uparrow & & \uparrow & & \uparrow & & \\
\Bbbk_{Z_M} & \longrightarrow & \oplus_{|I|=1}\Bbbk_{Z_{M_I}} & \longrightarrow & \oplus_{|I|=2}\Bbbk_{Z_{M_I}} & \cdots &
\end{array}$$

Now we can apply the pushforward f_* to get

$$\begin{array}{ccccccc}
\vdots & & \vdots & & \vdots & & \\
f_*\mathscr{S}^1(Z_M;\Bbbk) & \longrightarrow & \oplus_{|I|=1} f_*\mathscr{S}^1(Z_{M_I};\Bbbk) & \longrightarrow & \oplus_{|I|=2} f_*\mathscr{S}^1(Z_{M_I};\Bbbk) & \cdots\cdots & \\
\uparrow & & \uparrow & & \uparrow & & \\
f_*\mathscr{S}^0(Z_M;\Bbbk) & \longrightarrow & \oplus_{|I|=1} f_*\mathscr{S}^0(Z_{M_I};\Bbbk) & \longrightarrow & \oplus_{|I|=2} f_*\mathscr{S}^0(Z_{M_I};\Bbbk) & \cdots\cdots & \\
\uparrow & & \uparrow & & \uparrow & & \\
f_*\Bbbk_{Z_M} & \longrightarrow & \oplus_{|I|=1} f_*\Bbbk_{Z_{M_I}} & \longrightarrow & \oplus_{|I|=2} f_*\Bbbk_{Z_{M_I}} & \cdots\cdots &
\end{array}$$

Since quasi-isomomorphisms can be detected at the level of stalks we take stalks of the above diagram at some point $(v,t) \in \mathbb{S}^{d-1} \times \mathbb{R}$,

$$\begin{array}{ccccccc}
\vdots & & \vdots & & \vdots & & \\
S^1(M_{v,t};\Bbbk) & \longrightarrow & \oplus_{|I|=1} S^1((M_I)_{v,t};\Bbbk) & \longrightarrow & \oplus_{|I|=2} S^1((M_I)_{v,t};\Bbbk) & \cdots\cdots & \\
\uparrow & & \uparrow & & \uparrow & & \\
S^0(M_{v,t};\Bbbk) & \longrightarrow & \oplus_{|I|=1} fS^0((M_I)_{v,t};\Bbbk) & \longrightarrow & \oplus_{|I|=2} S^0((M_I)_{v,t};\Bbbk) & \cdots\cdots & \\
\uparrow & & \uparrow & & \uparrow & & \\
(f_*\Bbbk_{Z_M})_{v,t} & \longrightarrow & \oplus_{|I|=1}(f_*\Bbbk_{Z_{M_I}})_{v,t} & \longrightarrow & \oplus_{|I|=2}(f_*\Bbbk_{Z_{M_I})_{v,t}} & \cdots\cdots &
\end{array}$$

where $M_{v,t}$ is the fiber of the map $f^{-1}(v,t) \cap Z_M$. Taking cohomology vertically gives us the E_1 page,

$$\oplus_{|I|=1} H^2(M_{I,v,t};\Bbbk) \longrightarrow \oplus_{|J|=2} H^2(M_{J,v,t};\Bbbk)$$

$$\oplus_{|I|=1} H^1(M_{I,v,t};\Bbbk) \longrightarrow \oplus_{|J|=2} H^1(M_{J,v,t};\Bbbk)$$

$$\oplus_{|I|=1} H^0(M_{I,v,t};\Bbbk) \longrightarrow \oplus_{|J|=2} H^0(M_{J,v,t};\Bbbk)$$

Taking cohomology of this complex horizontally gives us the E_2 page of the spectral sequence. The Grothendieck spectral sequence (Theorem 3.3.2) with $\mathcal{F} = f_* : \mathbf{Shv}(Z_M) \to \mathbf{Shv}(\mathbb{S}^{d-1} \times \mathbb{R})$, $\mathcal{G} = \Gamma : \mathbf{Shv}(\mathbb{S}^{d-1} \times \mathbb{R}) \to \mathbf{Ab}$ and X as the the sequence in 3.3 gives us that the E^2 page converges to the cohomology of $M_{v,t}$.

□

We now reconsider Theorem 3.1.1 from this spectral sequence perspective.

Remark 3.3.3 (Redux of the Inclusion-Exclusion Principle for the ECT)**.** The inclusion-exclusion principle expression for the indicator function

$$\mathbb{1}_M = \sum_{I \subset \Lambda} (-1)^{|I|+1} \mathbb{1}_{M_I}.$$

is exactly the local Euler-Poincaré index of Godemont's resolution from Equation 3.3. The Radon transform expression

$$\mathcal{R}_S \mathbb{1}_M = \sum_{I \subset \Lambda} (-1)^{|I|+1} \mathcal{R}_S \mathbb{1}_{M_I},$$

is exactly the local Euler-Poincaré index of the pushforward of the resolution in Equation 3.3. Checking on stalks reveals that for any $(v,t) \in \mathbb{S}^{d-1} \times \mathbb{R}$

$$\chi\big(f_M^{-1}(v,t)\big) = \sum_{I \subset \Lambda} (-1)^{|I|+1} \chi\big(f_{M_I}^{-1}(v,t)\big).$$

3.4 Nerve theorem

Corollary 3.4.1. *Suppose $M \in \mathcal{CS}(\mathbb{R}^d)$ is a polyhedron and suppose $\mathcal{M} = \{M_i\}_{i \in I}$ is a cover of M by PL subspaces, then $\mathrm{PHT}^n(M)$ is the n-th cohomology of the complex,*

$$0 \to \oplus_{|I|=1} \mathrm{PHT}^0(M_I) \to \oplus_{|I|=2} \mathrm{PHT}^0(M_I) \to \cdots \tag{3.4}$$

where the $\cdots$ represents PHT^0 of higher intersection terms.

Proof. By examining the E_1 page of the spectral sequence in Section 3.3.2 one can see that for a PL cover, the higher cohomologies, i.e. the higher PHTs, all vanish. The E_1 page looks like,

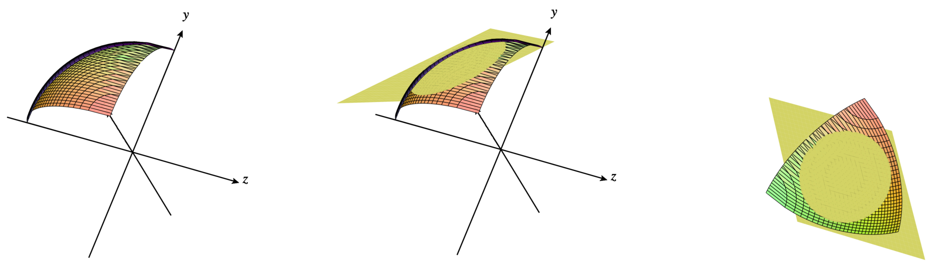

FIGURE 3.1: For a cover of $M = \mathbb{S}^2$ using the intersection with each orthant in $\mathbb{R}^3$, a cover element M_i can have non-trivial $\mathrm{PHT}^1(M_i)$ even though $\mathrm{PHT}^1(M)$ is trivial.

$$0 \longrightarrow 0$$

$$0 \longrightarrow 0$$

$$\oplus_{|I|=1} H^0(M_{I,v,t}; \Bbbk) \longrightarrow \oplus_{|J|=2} H^0(M_{J,v,t}; \Bbbk)$$

Consequently the spectral sequence collapses at the E_2 page.

□

Here M_I when $|I| = k$ represents the k- fold intersections. i.e $\{M_{i_1,\ldots,i_k} | i_1, \ldots, i_k \in I^k$ are distinct$\}$. Since taking stalks is an exact functor, the above theorem can be rephrased as the following: for any $(v,t) \in \mathbb{S}^{d-1} \times \mathbb{R}$, $\mathrm{PHT}^n(M)(v,t)$ is the n-th cohomology of

$$0 \to \oplus_{|I|=1} H^0(M_{I,v,t}, \Bbbk) \to \oplus_{|I|=2} H^0(M_{I,v,t}, \Bbbk) \to \cdots,$$

where $M_{I,v,t} = f_{M_I}^{-1}(v,t)$. This is the usual nerve lemma for persistent homology filtered in some direction v.

Remark 3.4.2. It should be noted that positive scalar curvature of a constructible set M (when defined) obstructs Theorem 3.4.1 from being directly applied. The cover elements may necessarily have higher homology when viewed in a direction along the normal vector to the point with positive scalar curvature. See Figure 3.1 for an example.

3.5 Example Calculation

A sheaf is not a collection of local data nor a collection of global data, but the glue that passes between both. In this section, we leverage the spectral sequence argument to illustrate an explicit calculation of the gluing process.

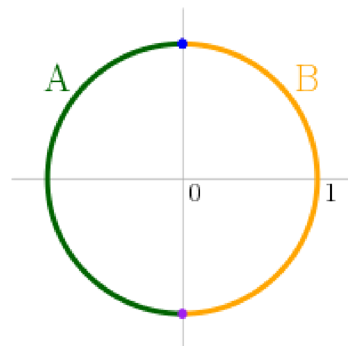

FIGURE 3.2: $M = S^1$ with cover elements A and B.

Let $M = \mathbb{S}^1$ be the unit circle in $\mathbb{R}^2$. Define a covering $\mathcal{M} = \{A, B\}$ by two closed half-circles, as indicated in Figure 3.2. First, we compute the PHT of each of the cover elements and their intersection. Because our PHT sheaves are on $\mathbb{S}^1 \times \mathbb{R}$ we can project this cylinder onto the plane $\mathbb{R}^2$ by following the instructions in the caption of Figure 3.3.

Now for every point $(v, t) \in \mathbb{S}^1 \times \mathbb{R}$ we write out the spectral sequence in Section 3.3.2. For example, let $(v, t) = (\uparrow, 0)$, then the E_1 page of the spectral sequence works out to be:

$$0 \longrightarrow 0 \longrightarrow 0$$

$$0 \longrightarrow 0 \longrightarrow 0$$

$$\Bbbk \oplus \Bbbk \longrightarrow \Bbbk \longrightarrow 0$$

This spectral sequence collapses after the E_2 page and converges to $H^*(M_{v,t}; \Bbbk)$. And so for this example taking cohomology horizontally gives us that $H^0(M_{v,t}; \Bbbk) = \Bbbk$ and $H^1(M_{v,t}; \Bbbk) = 0$. Since the PHT is a sheaf we can can do this at all (v, t) to find $PHT^*(M)$. Figure 3.4 shows the PHT of M.

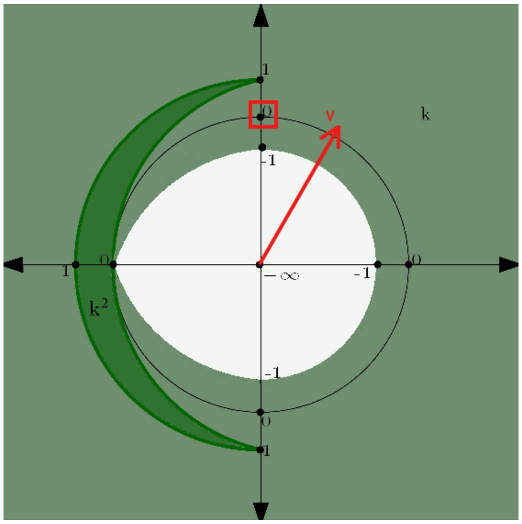

(a) $\mathrm{PHT}^0(A)$

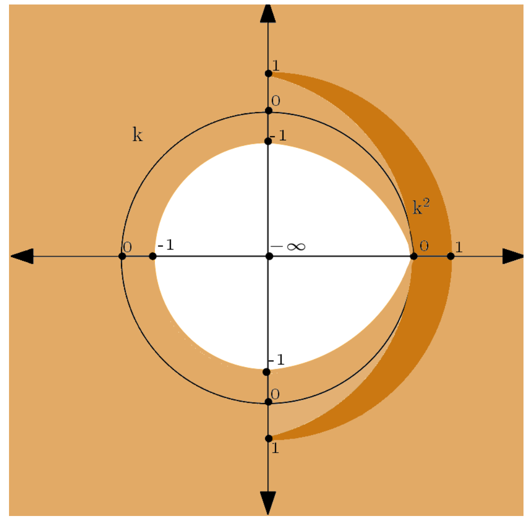

(b) $\mathrm{PHT}^0(B)$

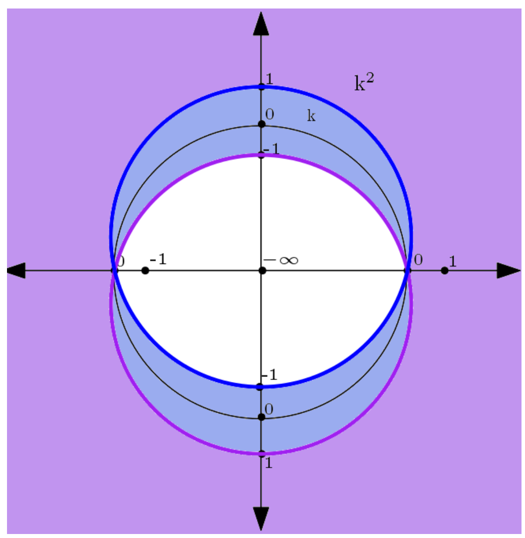

(c) $\mathrm{PHT}^0(A \cap B)$

FIGURE 3.3: To visualise the sheaf we have mapped the cylinder $\mathbb{S}^1 \times \mathbb{R}$ to the plane $\mathbb{R}^2$ by squashing down a tapering cylinder onto a plane in the following way: fix direction $v \in \mathbb{S}^1$ and then map $\{v\} \times \mathbb{R}$ onto $(0, \infty)$. So every direction v has a ray attached to it. For example, consider direction $w = \uparrow = (0,1)$ and $t = 0$ then in figure (A) we see $H^0(A_{w,t}, \mathbb{k}) = \mathbb{k}$ (see the red square) since $A_{w,t} = \{x \in A | x \cdot w \leqslant t\}$ has one connected component.

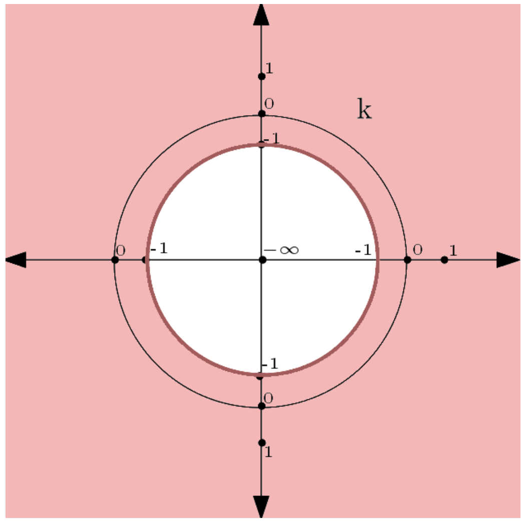

FIGURE 3.4: $\mathrm{PHT}^0(M)$

3.6 Relative PHT and ECT

We showed how to construct the PHT of a shape by gluing PHTs that from a cover. Intuitively this corresponds to "adding" several PHTs together in a precise way. A natural question to consider is if there is a process for "subtracting" one PHT from another. This is accomplished by using relative cohomology.

Definition 3.6.1 (Relative PHT). Let $M \in \mathcal{CS}(\mathbb{R}^d)$ be a constructible set and suppose $N \subset M$ is a closed constructible subset of M The **relative PHT** is defined to be the sheaf over $\mathbb{S}^{d-1} \times \mathbb{R}$ defined by

$$\mathrm{PHT}^i(M,N) = \mathrm{sh}\Big[U \to H^i\big(f_M^{-1}(U), f_N^{-1}(U); \Bbbk_{Z_M}\big)\Big].$$

The stalk at $(v,t) \in \mathbb{S}^{d-1} \times \mathbb{R}$ is the relative cohomology of the pair $(M_{v,t}, N_{v,t})$.

To prove that this definition suitably "subtracts" one PHT from another, consider the long exact sequence of pairs:

$$\cdots \to \mathrm{PHT}^i(M,N)_{v,t} \to \mathrm{PHT}^i(M)_{v,t} \to \mathrm{PHT}^i(N)_{v,t} \to \mathrm{PHT}^{i+1}(M,N)_{v,t} \to \cdots \quad (3.5)$$

Exactness at stalks implies exactness of sheaves and so we have the following long exact sequence of PHT sheaves:

$$\cdots \to \mathrm{PHT}^i(M,N) \to \mathrm{PHT}^i(M) \to \mathrm{PHT}^i(N) \to \mathrm{PHT}^{i+1}(M,N) \to \cdots$$

Similarly, we can also interpret the long exact sequence of pairs given in Equation 3.5 from the point of view of Euler characteristic.

Definition 3.6.2. (Relative ECT) The relative ECT of a pair (M,N) where N is a constructible subset of M is the function associated to the relative PHT sheaf under the function-to-sheaf correspondence. That is for all $(v,t) \in \mathbb{S}^{d-1} \times \mathbb{R}$,

$$\mathrm{ECT}(M,N)(v,t) = \chi(\mathrm{PHT}(M,N))(v,t) = \sum_i (-1)^i \dim H^i\big(f_M^{-1}(v,t), f_N^{-1}(v,t)\big)$$

Subtraction of ECTs characterizes the relative Euler Characteristic Transform.

Lemma 3.6.3. *For closed and constructible $N \subset M$*

$$\mathrm{ECT}(M,N) = \mathrm{ECT}(M) - \mathrm{ECT}(N).$$

Proof. Recall the LES of a pair from Equation 3.5,

$$\cdots \to \mathrm{PHT}^i(M,N)_{v,t} \to \mathrm{PHT}^i(M)_{v,t} \to \mathrm{PHT}^i(N)_{v,t} \to \mathrm{PHT}^{i+1}(M,N)_{v,t} \to \cdots$$

The long exact sequence implies that

$$\chi\big(\mathrm{PHT}(M)\big)(v,t) = \chi\big(\mathrm{PHT}(M,N)\big)(v,t) + \chi\big(\mathrm{PHT}(N)\big)(v,t),$$

which can be rewritten as

$$\mathrm{ECT}(M)(v,t) = \mathrm{ECT}(M,N)(v,t) + \mathrm{ECT}(N)(v,t).$$

□

3.7 Infinity Category Version

At a high-level, the ∞-category perspective works by maintaining a data structure that keeps track of maps between objects and separates out composition in a clear way. By doing so, it allows us to consider compositions that are identified with other morphisms only "up to homotopy" or via a similar identification. The usual refrain is that an ∞-category not only has 1-morphisms, like an ordinary category does, but also 2-morphisms that relate 1-morphisms, 3-morphisms that relate 2-morphisms, and all the way up *ad infinitum.* The data structure that organizes all these morphisms is called a simplicial set, and an ∞-category is a simplicial set that has certain "composition" properties that allow us to fill out a simplex.

3.7.1 Simplicial Sets and Infinity Categories

Let Δ denote the simplex category, whose objects are non-empty totally ordered finite sets and whose morphisms are order-preserving maps. More explicitly, the objects are of the form $[n]$, which denotes the totally ordered set $\{0 < 1 < \cdots < n\}$ for $n \geqslant 0$. The morphisms are maps $f : [n] \to [m]$ such that $i \leqslant j \implies f(i) \leqslant f(j)$.

For each $n \geqslant 0$ and $0 \leqslant i \leqslant n$, define the coface maps $d^i : [n-1] \to [n]$ and the codegeneracy maps $s^i : [n+1] \to [n]$ as follows:

$$d^i(k) = \begin{cases} k & \text{if } 0 \leqslant k < i \\ k+1 & \text{if } i \leqslant k \leqslant n-1 \end{cases} \qquad s^i(k) = \begin{cases} k & \text{if } 0 \leqslant k \leqslant i \\ k-1 & \text{if } i < k \leqslant n+1 \end{cases}$$

The coface map d^i skips i whereas the codegeneracy map s^i repeats i. It turns out that any order preserving morphism $f \in \hom_\Delta([n],[m])$ can be written as composites of s^i (insertions) and d^i (deletions). As a convention, if a morphism in $f \in \hom_\Delta([n],[m])$ contains a codegeneracy map in its factorization, then we say f is degenerate.

A **simplicial set** is a functor $S : \Delta^{\text{op}} \to \text{Set}$. It is typical to call $S([n])$ the set of n-simplices in S. In this setting, the **standard n-simplex** is the simplicial set $\Delta^n := \text{Hom}_\Delta(-,[n])$. Geometrically, once can think of a simplicial set as a bunch of non-degenerate simplices, with the face maps giving a way to glue simplices together. This observation is made precise in the following examples.

Example 3.7.1. (Standard 0-simplex) The standard 0-simplex (viewed as a simplicial set) consists of one element, namely the 0 map, for every $[n]$. Note that, there is only one non-degenerate morphism which is the identity morphism from $[0] \to [0]$. This can be thought of as a vertex, agreeing with standard 0-simplex from algebraic topology.

Example 3.7.2. (Standard 2-simplex) The 0-simplices of Δ^2 consists of order preserving maps from $[0] \to [2]$. There are 3 such maps each corresponding to $0 \mapsto i$ where $0 \leqslant i \leqslant 2$. These can be thought of as the three 0-simplices or vertices as indicated in Figure 3.5. There is only one non-degenerate map from the 2-simplex, namely the identity map from $[2] \to [2]$. This unique morphism can be thought of as filled-in triangle.

In this way, the standard n-simplex viewed as a simplicial set is a generalisation of the n-simplex from algebraic topology. Think of simplicial sets as the "singular" version of simplicial complexes from algebraic topology.

Definition 3.7.3. (Inner horn) The ith inner horn Λ^n_i is the subcomplex of Δ^n obtained by removing the n-simplex and its i^{th} face.

Simplicial sets form a category, with objects as simplicial sets and morphisms as natural transformations.

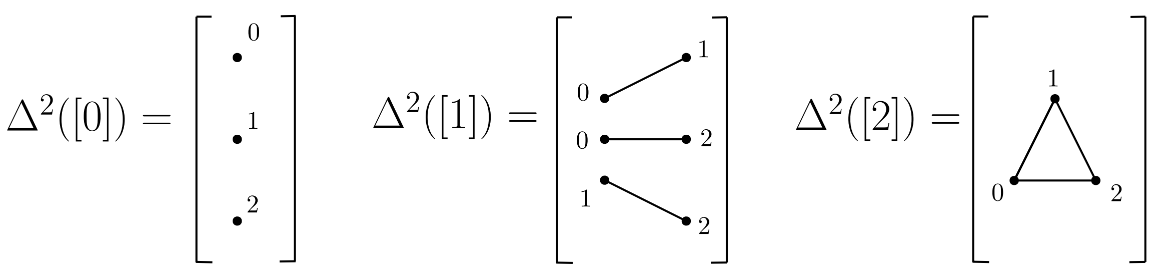

FIGURE 3.5: A representation of the standard 2-simplex: the non-degenerate maps of the standard 2-simplex are in one-one correspondence with this by the Yoneda Lemma.

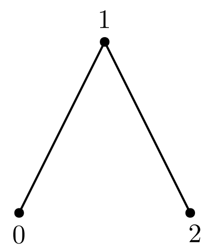

FIGURE 3.6: Inner horn Λ_1^2

Definition 3.7.4 (Infinity Category)**.** An ∞-category is a simplicial set $\mathcal{C}$ which has the following property: for any $0 < i < n$, any map $f : \Lambda_i^n \to \mathcal{C}$ admits an extension $f : \Delta^n \to \mathcal{C}$.

Example 3.7.5. (Standard n-simplex) It is easy to check from the definition that the standard n-simplex is an infinity category.

Example 3.7.6 (Nerve of a Category)**.** Every ordinary category $\mathbf{C}$ has an associated ∞-category via the nerve construction $\mathcal{N}(\mathbf{C})$. The 0-simplices of the nerve are objects of $\mathbf{C}$, the 1-simplices are morphisms of $\mathbf{C}$, the 2-simplices are strings of compositions of morphisms of length 2, e.g., $g \circ f : x \to y \to z$, the 3-simplices are strings of composition of length 3 and so on.
The nerve of a category is an infinity category [cite]. For an intuitive understanding, consider a map from $\Lambda_1^2 \to \mathcal{N}(\mathbf{C})$. From diagram 3.6, we can observe that this implies the existence of two 1-simplices: $g : C_0 \to C_1$ and $h : C_1 \to C_2$. The order of the maps is determined by the order-preserving structure of the simplex category, and the fact that the target of the first is the same as the origin of the second is due to the induced face maps. An extension from $\Delta^2 \to \mathcal{N}(\mathbf{C})$ always exists. This is true because there is a map $f : C_0 \to C_2$ by taking the composition $f = h \circ g$. The unique 2-simplex is the morphism $C_0 \to C_1 \to C_2$.

Morphisms $f, g : x \to y$ for any objects x, y in an infinity category are said to be

homotopic if there is a 2-simplex of the form:

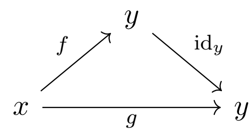

It can be seen that homotopy of morphisms is an equivalence relation on the set of morphisms between any two objects.

Definition 3.7.7. (Homotopy category) The homotopy category of an infinity category $\mathcal{C}$, denoted h$\mathcal{C}$, is the category whose objects are the objects of $\mathcal{C}$ and whose morphisms are homotopy classes of morphisms in $\mathcal{C}$.

One can verify that $[f \circ g] = [f] \circ [g]$ and $\mathrm{id}_{\mathcal{C}} = [\mathrm{id}_{\mathcal{C}}]$, confirming that the homotopy category is indeed a category.

3.7.2 Homological Algebra in Infinity categories

Let $\mathcal{A}$ be an abelian category with enough injective objects. We denote ch($\mathcal{A}$) to be the category of chain complexes valued in $\mathcal{A}$, i.e. a composable sequence of morphisms:

$$\cdots \to A^{-1} \xrightarrow{\partial^{-1}} A^0 \xrightarrow{\partial^0} A^1 \to \cdots$$

such that $\partial^n \circ \partial^{n-1} = 0$ for every integer n.

Remark 3.7.8. Note that we fix a cohomological grading, that is the differentials ∂ increase the index by 1.

This category can be promoted to a differentially graded category by enriching it in the category of chain complexes of abelian groups ch(**Ab**). In other words, the hom-sets in ch($\mathcal{A}$) can be replaced by mapping complexes as below: If $A^\bullet$ and $B^\bullet$ are chain complexes,

$$\underline{\mathrm{hom}}(A^\bullet, B^\bullet)[d] = \prod_{n \in \mathbb{Z}} \mathrm{hom}_{\mathcal{A}}(A^n, B^{n+d})$$

and the differentials $\partial : \underline{\mathrm{hom}}(A^\bullet, B^\bullet)[d-1] \to \underline{\mathrm{hom}}(A^\bullet, B^\bullet)[d]$ is given by $\partial f_n = \partial_B \circ f_{n-1} - (-1)^d f_n \circ \partial_A$.

We can build an infinity category from a category enriched in ch(**Ab**) by taking the differentially graded nerve, which is similar to the construction outlined above.

Example 3.7.9 (Differentially graded nerve of $\mathrm{ch}(\mathcal{A})$)**.**

- A 0-simplex of $N_{dg}(\mathrm{ch}(\mathcal{A}))$ is a chain complex $A^\bullet$.

- A 1-simplex of $N_{dg}(\mathrm{ch}(\mathcal{A}))$ is a pair of chain complexes $A^\bullet, B^\bullet$ and a morphism $f \in \underline{\mathrm{hom}}(A^\bullet, B^\bullet)_0$ such that $\partial f = 0$. In other words, 1- simplices between $A^\bullet$ and $B^\bullet$ are morphisms f such that each of the squares commute.

$$\begin{array}{ccccccccc}
\cdots & \longleftarrow & A^{-1} & \xrightarrow{\partial_A} & A^0 & \xrightarrow{\partial_A} & A^1 & \longrightarrow & \cdots \\
 & & \downarrow{\scriptstyle f} & & \downarrow{\scriptstyle f} & & \downarrow{\scriptstyle f} & & \\
\cdots & \longleftarrow & B^{-1} & \xrightarrow{\partial_B} & B^0 & \xrightarrow{\partial_B} & B^1 & \longrightarrow & \cdots
\end{array}$$

- A 2-simplex of $N_{dg}(\mathrm{ch}(\mathcal{A}))$ consists of a triple of chain complexes $A^\bullet, B^\bullet$ and $C^\bullet$ and morphisms $f \in \underline{\mathrm{hom}}(A^\bullet, B^\bullet)_0$, $g \in \underline{\mathrm{hom}}(B^\bullet, C^\bullet)_0$ and $h \in \underline{\mathrm{hom}}(A^\bullet, C^\bullet)_0$ such that $\partial f = \partial g = \partial h = 0$ and a morphism $h' \in \mathrm{hom}(A^\bullet, C^\bullet)$ such that $\partial h' = g \circ f - h$.

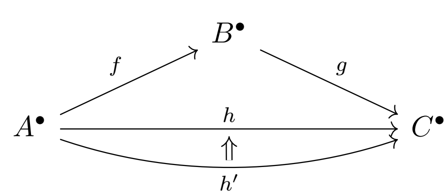

- and so on.

Remark 3.7.10. The homotopy category of the differentially graded nerve $hN_{dg}(\mathrm{ch}(\mathcal{A}))$ is the usual homotopy category of chain complexes in $\mathcal{A}$, denoted as $\mathcal{K}(\mathcal{A})$. So, the ∞-category $N_{dg}(\mathrm{ch}(\mathcal{A}))$ provides an ∞-categorical version of homological algebra in $\mathcal{A}$. In particular, there exists a derived infinity category which is the ∞-localization of the category of chain complexes in $\mathcal{A}$ at the class of quasi-isomorphisms. Equivalently, by [Lur17], the derived ∞-category can be defined to be the differentially graded nerve of the category of chain complexes of injectives in $\mathcal{A}$.

Let $\text{ch}^+(\mathcal{A})$ denote the full subcategory of $\text{ch}(\mathcal{A})$ spanned by those chain complexes that are bounded from below i.e. for small $n < 0$, $A^n \simeq 0$.

Definition 3.7.11. (Derived ∞-Category) The derived ∞-category of $\mathcal{A}$ is the differentially graded nerve of $\text{ch}(\mathcal{A}_{\text{inj}})$, that is $\mathcal{D}^+(\mathcal{A}) = \mathcal{N}_{dg}(\text{ch}^+(\mathcal{A}_{\text{inj}}))$ where $\mathcal{A}_{\text{inj}}$ is the full subcategory of $\mathcal{A}$ spanned by the injective objects.

This category is a stable ∞-category and carries a triangulated structure. In particular, the homotopy category of this category, $h\mathcal{D}^+(\mathcal{A})$ is the classical derived category in homological algebra. However, unlike the derived category in classical homological algebra, this category has finite limits and colimits by virtue of being a stable infinity category [Lur17, Prop. 1.1.3.4].

As in the case of classical homological algebra, one can also study the functoriality of the derived ∞-category. Suppose $\mathcal{F} : \mathcal{A} \to \mathcal{B}$ is an additive functor between abelian categories i.e. the functor preserves the 0 object and preserves direct sums. The object-wise application of the functor $\mathcal{F}$ gives a functor $\text{ch}(\mathcal{F}) : \text{ch}(\mathcal{A}) \to \text{ch}(\mathcal{B})$ which is an enriched functor over chain complexes of abelian groups. The functoriality of the differentially graded nerve construction gives a functor $\mathcal{N}_{dg}(\text{ch}(\mathcal{F})) : \mathcal{N}_{dg}(\text{ch}(\mathcal{A})) \to \mathcal{N}_{dg}(\text{ch}(\mathcal{B}))$. In particular, this functor commutes with finite limits and colimits because the functor preserves sums and the terminal object 0. Similar to the case of classical homological algebra this functor $\mathcal{N}_{dg}(\text{ch}(\mathcal{F}))$ induces a right derived functor $RF : D(\mathcal{A}) \to D(\mathcal{B})$ by taking its left Kan extension along the projection $\rho_A : \mathcal{N}_{dg}(\text{ch}(\mathcal{A})) \to D(A)$.

A surprising property of functors of stable infinity categories is that left exact (equivalently right exact) functors are exact, i.e., they commute with finite limits and colimits [Lur17, Prop. 1.1.4.1]. In particular, by the universal property of $\mathcal{D}^+(\mathcal{A})$ [Lur17, Thm. 1.3.3.2] we have that a left exact functor $\mathcal{F} : \mathcal{A} \to \mathcal{B}$, gives a left exact (and hence exact) functor, called the *right derived functor*, $RF : \mathcal{D}^+(\mathcal{A}) \to \mathcal{D}^+(\mathcal{B})$.

3.7.3 The PHT viewed as an object of the Derived ∞-Category.

We can now specialize to the case where the abelian categories are $\mathcal{A} = \mathbf{Shv}(M \times \mathbb{S}^{d-1} \times \mathbb{R})$ and $\mathcal{B} = \mathbf{Shv}(\mathbb{S}^{d-1} \times \mathbb{R})$ where $M \in \mathcal{CS}(\mathbb{R}^d)$. As defined before, let $f : M \times \mathbb{S}^{d-1} \times \mathbb{R} \to \mathbb{S}^{d-1} \times \mathbb{R}$ be the projection on the last two factors. Further, let $\iota_M : Z_M \to M \times \mathbb{S}^{d-1} \times \mathbb{R}$ be an embedding of spaces. The composition $f \circ \iota_M$ is written as f_M (see Definition 2.6.5).

Definition 3.7.12 (PHT: Infinity Category Version)**.** The ∞-categorical definition of the PHT of $M \in \mathcal{CS}(\mathbb{R}^d)$ is given by:

$$\mathrm{PHT}(M) = R(f_M)_* \Bbbk_{Z_M} \in \mathcal{D}^+(\mathbf{Shv}(\mathbb{S}^{d-1} \times \mathbb{R}))$$

where $\mathcal{D}^+(\mathbf{Shv}(\mathbb{S}^{d-1} \times \mathbb{R}))$ is the derived ∞-category of bounded below complexes.

3.7.4 Proof of the sheaf axiom via infinity categories

Since the right derived functor is exact (and hence commutes with finite limits and colimits), we have that the sheaf axiom for the assignment $M \mapsto \mathrm{PHT}(M) = R(f_M)_* \Bbbk_{Z_M}$ for $M \in \mathcal{CS}(\mathbb{R}^d)$ is automatic. Precisely, for a finite covering $\mathcal{M} = \{M_i \hookrightarrow M\}_{i \in I}$ we have:

$$\begin{aligned} R(f_M)_* \Bbbk_{Z_M} = Rf_* (\iota_M)_* \Bbbk_{Z_M} &= Rf_* \varprojlim_{I \in \mathcal{N}(\mathcal{M})} (\iota_{M_I})_* \Bbbk_{Z_{M_I}} \\ &= \varprojlim_{I \in \mathcal{N}(\mathcal{M})} Rf_* (\iota_{M_I})_* \Bbbk_{Z_{M_I}} \\ &= \varprojlim_{I \in \mathcal{N}(\mathcal{M})} R(f_{M_I})_* \Bbbk_{Z_{M_I}}. \end{aligned}$$

4. Metrics on Shape Spaces

Aside from the intrinsic theoretical interest in a gluing result for the PHT, a practical motivation is to parallelize PHT computations over a cover. This parallelization inevitably becomes more complex if our cover elements have higher homology when viewed in certain directions and at certain filtration values. See Figure 3.1 for an example. This motivates our need to replace our shapes with piecewise-linear (PL) approximations. The main result of this section proves that, up to some tolerance, we can always approximate a submanifold $M \in \mathcal{CS}(\mathbb{R}^d)$ via a PL shape K so that the PHT's of M and K are arbitrarily close. This requires several preparatory steps. Firstly, we introduce several novel distances[1] on the PHT and prove that the PHT is stable under small perturbations of the underlying shape. This stability property reaffirms our belief that the PHT is a good summary statistic for shapes. We also carry out some calculations for especially simple shapes (finite point clouds in $\mathbb{R}^d$) and compare these calculations with the Procrustes distance. Secondly, we use the sampling procedure from Niyogi-Smale-Weinberger [NSW08] to approximate a submanifold of $\mathbb{R}^d$ by a (PL) polyhedron. Finally, we conclude from the stability theorem that the PHT of the polyhedron is close to the PHT of the submanifold. This chapter is taken from Chapter 4 of [ACM23].

4.1 Distances on PHTs

In this section, we define distances on PHTs using both the sheaf perspective (Definition 2.6.5) and the map to persistence diagram space (Definition 2.3.5) perspective. We start with the sheaf perspective as the interleaving distance we introduce is more involved algebraically, but it is also simplest for proving our main stability result. Bounds on the interleaving distance then imply bounds on certain Wasserstein-type distances, but not others.

[1] By which we mean an extended—the value ∞ is allowed—metric.

4.1.1 Interleaving-type Distances on the PHT

As first introduced in [Cha+09], the interleaving distance provides a powerful generalization of the Hausdorff distance to functors from $(\mathbb{R}, \leqslant)$ to an arbitrary data category **Dat**. When **Dat** = **vect**, the celebrated isometry theorem of [Les15] proves that this distance is equivalent to the bottleneck distance (cf. Definition 2.3.4), which due to its combinatorial structure as a matching problem can be computed in polynomial time. However, the interleaving distance is far more general and permits the construction of distances in more general categorical settings, where no easy distance is to be found. Following a suggestion of Patel, the interleaving distance for sheaves first appeared in [Cur14b, §15], then for cosheaves of sets over $\mathbb{R}$ (equivalent to Reeb graphs) in [DMP16], and finally for derived sheaves over vector spaces in [KS18] as the *convolution distance*.

In parallel to these developments was the thesis work of Stefanou [Ste18], which generalized the interleaving distance to any category equipped with an action of $[0, \infty)$ on it—also called a category with a flow [SMS18]. This action is used to send any object F to its forward evolution F^ϵ. With this in hand, one defines an ϵ-interleaving between two objects F and G to be any commutative diagram of the form

$$\begin{array}{ccccc} F & \longrightarrow & F^{\epsilon} & \longrightarrow & F^{2\epsilon} \\ & \searrow\!\!\!\!\nearrow & & \searrow\!\!\!\!\nearrow & \\ G & \longrightarrow & G^{\epsilon} & \longrightarrow & G^{2\epsilon} \end{array}$$

although [SMS18; Ste18] consider more general diagrams then this one. Regardless of the particulars, one then defines the **interleaving distance** $d_I(F, G)$ to be the infimum over all $\epsilon \in [0, \infty)$ where such diagrams exist. If there is no such diagram relating two objects, then $d_I(F, G) = \infty$.

Returning to interleavings of pre-sheaves, the suggestion of Patel was to define the ϵ-thickening/smoothing of a presheaf on a metric space X to be $F^\epsilon(U) = F(U^\epsilon)$, where U^ϵ is the thickening of an open set by $\epsilon \geqslant 0$. However, for our applications we leverage the more general perspective of [SMS18] to work with a specialized shift operation in order to define our interleavings.

Definition 4.1.1 (ϵ-Shift of the Derived PHT)**.** The **ϵ-shift** of the derived PHT sheaf of a shape $M \in \mathcal{CS}(\mathbb{R}^d)$ is

$$\mathrm{PHT}(M)^\epsilon := R(f_{M^\epsilon})_* \Bbbk_{Z_{M^\epsilon}}$$

where Z_{M^ϵ} is an ϵ-shift of the set Z_M in the filtration parameter, i.e.,

$$Z_{M^\epsilon} = \{(x, v, t) \in M \times \mathbb{S}^{d-1} \times \mathbb{R} \mid x \cdot v \leqslant t + \epsilon\}.$$

The map f_{M^ϵ} is the usual projection of Z_{M^ϵ} onto its last two factors. Notice that if $(x, v, t) \in Z_M$, then it certainly is contained in Z_{M^ϵ}, which implies that $Z_M \subseteq Z_{M^\epsilon}$. By functoriality of cohomology, there is a restriction map of constant sheaves $\Bbbk_{Z_{M^\epsilon}} \to \Bbbk_{Z_M}$ and thus a map of sheaves $\mathrm{PHT}(M)^\epsilon \to \mathrm{PHT}(M)$. Further details that this defines an ϵ-shift functor, starting on the image of $\mathcal{CS}(\mathbb{R}^d)^{op} \to \mathcal{D}^b(\mathbf{Shv}(\mathbb{S}^{d-1} \times \mathbb{R}))$, which further satisfies the axioms of [SMS18] is left to the reader. Note that an ϵ-shift of the derived PHT sheaf is still a derived sheaf although it might not correspond to the PHT of any particular shape.

Since our sheaves are defined used cohomology, our interleaving diagram goes in the opposite direction of the one stated above, thus closer in spirit to the interleaving diagrams of [Cur14b, §15] and [KS18].

Definition 4.1.2 (Interleaving Distance between PHTs)**.** Let $M, N \in \mathcal{CS}(\mathbb{R}^d)$. An ϵ-interleaving of $\mathrm{PHT}(M)$ and $\mathrm{PHT}(N)$ is a pair of morphisms $\varphi : \mathrm{PHT}(M)^\epsilon \to \mathrm{PHT}(N)$ and $\psi : \mathrm{PHT}(N)^\epsilon \to \mathrm{PHT}(M)$ such that the following diagram is commutative:

$$\begin{array}{ccccc} \mathrm{PHT}(M)^{2\epsilon} & \longrightarrow & \mathrm{PHT}(M)^{\epsilon} & \longrightarrow & \mathrm{PHT}(M) \\ & \searrow\!\!\!\!\!\!\nearrow & & \searrow\!\!\!\!\!\!\nearrow & \\ \mathrm{PHT}(N)^{2\epsilon} & \longrightarrow & \mathrm{PHT}(N)^{\epsilon} & \longrightarrow & \mathrm{PHT}(N) \end{array} \tag{4.1}$$

The arrows $\mathrm{PHT}(M)^{2\epsilon} \to \mathrm{PHT}(N)^\epsilon$ and $\mathrm{PHT}(N)^{2\epsilon} \to \mathrm{PHT}(M)^\epsilon$ being given by the image of φ and ψ under the ϵ-shift functor. The interleaving distance between PHT sheaves is then

$$d_I(\mathrm{PHT}(M), \mathrm{PHT}(N)) := \inf\{\epsilon \geqslant 0 \mid \exists\, \epsilon\text{-interleaving}\}.$$

If no such interleaving exists, then $d_I(\mathrm{PHT}(M), \mathrm{PHT}(N)) = \infty$.

4.1.2 Wasserstein-type Distances on the PHT

We can also define a metric on the PHTs viewed as map (Definition 2.3.5). This is a generalisation of the p-PHT distance in [ST20, Definition 5.4]. Let $\mathrm{PH}^i(M, v)$ be the persistence diagram in degree i associated to shape M with sub-level set filtration given by the height function in direction $v \in \mathbb{S}^{d-1}$, i.e. $h_v(x) = x \cdot v$.

Definition 4.1.3. ((p, q)-PHT distance) The (p, q)-PHT distance between $M, N \in \mathcal{CS}(\mathbb{R}^d)$ for $p, q \geqslant 1$ in degree i for $i > 0$ and is defined as:

$$d_{p,q}^{\mathrm{PHT}^i}(M, N) = \left(\int_{v \in \mathbb{S}^{d-1}} \mathcal{W}_p^q \left(\mathrm{PH}^i(M, v), \mathrm{PH}^i(N, v) \right) d\mu \right)^{\frac{1}{q}}$$

where μ is the Lebesgue measure on $\mathbb{S}^{d-1}$. When $q = \infty$,

$$d_{p,\infty}^{\mathrm{PHT}^i}(M, N) = \max_{v \in \mathbb{S}^{d-1}} \mathcal{W}_p \left(\mathrm{PH}^i(M, v), \mathrm{PH}^i(N, v) \right).$$

Note that when $p = \infty$, we have the bottleneck distance between persistence diagrams.

Below, we consider the cases where $p = 2$ or $p = \infty$, i.e., the Wasserstein 2-distance or the bottleneck distance on diagrams. Additionally, we will restrict ourselves to $q = 2$, which computes the squared average diagram distance over all directions, or $q = \infty$, which takes the biggest diagram distance over all directions. We refer to $d_{\infty,\infty}^{\mathrm{PHT}^i}$ as the **PHT bottleneck distance** in degree i. The next lemma explains the relationship between the PHT sheaf interleaving distance and the PHT bottleneck distance.

Lemma 4.1.4.

$$d_I(\mathrm{PHT}(M), \mathrm{PHT}(N)) \geqslant \max_{i \geqslant 0} d_{\infty,\infty}^{\mathrm{PHT}^i}(M, N).$$

Proof. Evaluating diagram 4.1 on stalks $(v, t) \in \mathbb{S}^{d-1} \times \mathbb{R}$ gives

$$\begin{array}{ccccc} H^i\left(f_{M^{2\epsilon}}^{-1}(v,t); \Bbbk\right) & \longrightarrow & H^i\left(f_{M^{\epsilon}}^{-1}(v,t); \Bbbk\right) & \longrightarrow & H^i\left(f_{M}^{-1}(v,t); \Bbbk\right) \\ & \searrow\!\!\!\!\nearrow & & \searrow\!\!\!\!\nearrow & \\ H^i\left(f_{N^{2\epsilon}}^{-1}(v,t); \Bbbk\right) & \longrightarrow & H^i\left(f_{N^{\epsilon}}^{-1}(v,t); \Bbbk\right) & \longrightarrow & H^i\left(f_{N}^{-1}(v,t); \Bbbk\right) \end{array}$$

This can be interpreted as an ϵ-interleaving of persistence modules in degree i obtained by filtering M, N in direction v. The isometry theorem [Cha+16; Les15] guarantees that the interleaving distance between persistence modules is equal to the the bottleneck distance between their corresponding persistence diagrams. In other words,

$$d_I(\text{PHT}(M), \text{PHT}(N)) = \epsilon \implies \forall i\ \forall v \in \mathbb{S}^{d-1} \quad \mathcal{W}^i_\infty\,(\text{PH}(M,v), \text{PH}(N,v)) \leqslant \epsilon,$$

where $\mathcal{W}^i_\infty$ is the bottleneck distance in degree i. □

4.2 Comparison with other Distances on PHTs and Shape Spaces

We now proceeed to compute the above mentioned distances on some simple examples and attempt a comparison with other shape space metrics, primarily the Procrustes distance. A more detailed computation of the PHT of the two embeddings of the letter 'V' from Figure 2.5 is carried out in Example 4.3.4 in Section 4.3 after our main stability result is proved.

4.2.1 Distances between Point Clouds

We begin with the simplest possible shapes: a finite collection of points in $\mathbb{R}^d$. It should be noted that one quirk of the persistent homology transform is that it is very sensitive to the global homology of a shape. Consequently, if two point clouds have differing numbers of points and no further construction is performed on them, then their PHTs are infinitely far apart. To this end, let $A = \{a_1, \cdots, a_n\}$ and $B = \{b_1, \cdots, b_n\}$ be two point clouds in $\mathbb{R}^d$.

Proposition 4.2.1. *If $A = [a_1, \cdots a_n]$ and $B = [b_1, \cdots, b_n]$ are point clouds regarded as matrices where the vectors that coordinatize each point are stored as columns, then*

$$d_I(\text{PHT}(A), \text{PHT}(B)) = d^{\text{PHT}^0}_{\infty,\infty}(A, B) = \inf_{\substack{\phi: A \to B \\ \textit{matching}}} \max_{a_i \in A} \|a_i - \phi(a_i)\|_{\mathbb{R}^d} \tag{4.2}$$

$$
\begin{aligned}
d_{2,2}^{\mathrm{PHT}^0}(A,B) &= \left(\int_{\|v\|=1} \inf_{\substack{\phi:A\to B\\ matching}} \sum_{i=1}^{n} |a_i \cdot v - \phi(a_i)\cdot v|^2 d\mu \right)^{1/2} \\
&\leqslant \inf_{\substack{\phi:A\to B\\ matching}} \frac{\sqrt{Area(\mathbb{S}^{d-1})} \times \|A-\phi(A)\|_F}{\sqrt{n}}.
\end{aligned} \tag{4.3}
$$

$$
d_{2,\infty}^{\mathrm{PHT}^0}(A,B) = \inf_{\substack{\phi:A\to B\\ matching}} \|A-\phi(A)\|_\infty. \tag{4.4}
$$

where $\|\cdot\|_F$ *and* $\|\cdot\|_\infty$ *stands for the Frobenius norm and Schatten* ∞*-norm of a matrix respectively.*

Proof. To prove Equation 4.2, we prove that the sheaf interleaving distance is equal to the max bottleneck distance over all directions. By Lemma 4.1.4, we have that an ϵ-interleaving of PHT sheaves implies that the (∞,∞)-PHT distance is less than ϵ. Suppose the latter distance is ϵ in degree 0, since the sublevel sets are finite point sets, for any $v \in \mathbb{S}^{d-1}$ and $t \in \mathbb{R}$ the following diagram commutes,

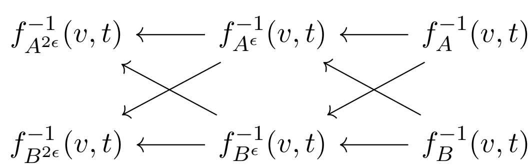

Since the sets in the above diagram are finite, it extends to a commutative diagram for every open $U \in \mathbb{S}^{d-1} \times \mathbb{R}$ and so, by the functoriality of the right derived functor there is an ϵ-interleaving of PHT(A) and PHT(B). So, $d_I(\mathrm{PHT}(A),\mathrm{PHT}(B)) = d_{\infty,\infty}^{\mathrm{PHT}^0}(A,B)$. For the case of points sets, the PHT Bottleneck distance turns out to be:

$$
d_{\infty,\infty}^{\mathrm{PHT}^0}(A,B) = \max_{v\in\mathbb{S}^{d-1}} \inf_{\substack{\phi:A\to B\\ \text{matching}}} \max_{a_i\in A} |a_i\cdot v - \phi(a_i)\cdot v| = \inf_{\substack{\phi:A\to B\\ \text{matching}}} \max_{a_i\in A} \|a_i - \phi(a_i)\|_{\mathbb{R}^d}
$$

To prove Equation 4.3, we need to calculate,

$$d_{2,2}^{\mathrm{PHT}^0}(A,B) = \left(\int_{v \in \mathbb{S}^{d-1}} \inf_{\substack{\phi: A \to B \\ \text{matching}}} \sum_{i=1}^{n} |a_i \cdot v - \phi(a_i) \cdot v|^2 d\mu \right)^{1/2}$$

where μ is the Lebesgue measure on $\mathbb{S}^{d-1}$. Let $X_i = a_i - \phi(a_i)$, and so rewriting the above expression in terms of matrices gives,

$$\begin{aligned}
\left(d_{2,2}^{\mathrm{PHT}^0}(A,B)\right)^2 &\leqslant \inf_{\substack{\phi: A \to B \\ \text{matching}}} \int_{v \in \mathbb{S}^{d-1}} v^T \left(\sum_{i=1}^{n} X_i X_i^T \right) v \; d\mu \\
&= \inf_{\substack{\phi: A \to B \\ \text{matching}}} \int_{v \in \mathbb{S}^{d-1}} v^T Y v \; d\mu \\
&= \inf_{\substack{\phi: A \to B \\ \text{matching}}} \int_{v \in \mathbb{S}^{d-1}} \operatorname{tr}(v^T Y v) \; d\mu \\
&= \inf_{\substack{\phi: A \to B \\ \text{matching}}} \int_{v \in \mathbb{S}^{d-1}} \operatorname{tr}(Y v v^T) \; d\mu \\
&= \inf_{\substack{\phi: A \to B \\ \text{matching}}} \operatorname{tr}\left(Y \int_{v \in \mathbb{S}^{d-1}} v v^T \; d\mu \right)
\end{aligned}$$

where Y is $\sum_{i=1}^{n} X_i X_i^T$. The integral $A := \int_{v \in \mathbb{S}^{d-1}} v v^T \; d\mu$ is invariant under conjugation by orthogonal matrices, since the sphere is rotationally invariant. So, in particular for orthogonal matrix U we have,

$$A = \int_{v \in \mathbb{S}^{d-1}} v v^T \; d\mu = \int_{v \in \mathbb{S}^{d-1}} U v v^T U^T \; d\mu = U A U^T.$$

The matrix A commutes with orthogonal matrices and so by Schur's Lemma, A must be a scalar times the identity matrix.

$$A = \int_{v \in \mathbb{S}^{d-1}} v v^T = \lambda I$$

where λ can be found my taking trace on both sides i.e.,

$$\int_{v \in \mathbb{S}^{d-1}} d\mu = \lambda n.$$

So, we have,

$$\inf_{\substack{\phi:A\to B\\ \text{matching}}} \operatorname{tr}\left(Y\int_{v\in\mathbb{S}^{d-1}} vv^T\, d\mu\right) = \inf_{\substack{\phi:A\to B\\ \text{matching}}} \frac{1}{n}\operatorname{tr}(Y)\operatorname{Area}(\mathbb{S}^{d-1}) = \inf_{\substack{\phi:A\to B\\ \text{matching}}} \frac{1}{n}\|A-\phi(A)\|_F^2\operatorname{Area}(\mathbb{S}^{d-1}).$$

Finally, to prove Equation 4.4, we note that

$$\begin{aligned}
\left(d_{2,\infty}^{\mathrm{PHT}^0}(A,B)\right)^2 &= \max_{v\in\mathbb{S}^{d-1}} \inf_{\substack{\phi:A\to B\\ \text{matching}}} \sum_{i=1}^n |a_i\cdot v-\phi(a_i)\cdot v|^2\\
&= \inf_{\substack{\phi:A\to B\\ \text{matching}}} \max_{v\in\mathbb{S}^{d-1}} v^TYv\\
&= \inf_{\substack{\phi:A\to B\\ \text{matching}}} \lambda_{\max}(Y)\\
&= \inf_{\substack{\phi:A\to B\\ \text{matching}}} \lambda_{\max}\left((A-\phi(A)(A-\phi(A))^T\right)\\
&= \inf_{\substack{\phi:A\to B\\ \text{matching}}} \|A-\phi(A)\|_\infty,
\end{aligned}$$

where $\|\cdot\|_\infty$ stands for the ∞-Schatten norm of the matrix. □

Proof. To prove Equation 4.2, we prove that the sheaf interleaving distance is equal to the max bottleneck distance over all directions. By Lemma 4.1.4, we have that an ϵ-interleaving of PHT sheaves implies that the (∞,∞)-PHT distance is less than ϵ. Suppose the latter distance is ϵ in degree 0, since the sublevel sets are finite point sets, for any $v\in\mathbb{S}^{d-1}$ and $t\in\mathbb{R}$ the following diagram commutes,

$$\begin{array}{ccccc}
f^{-1}_{A^{2\epsilon}}(v,t) & \longleftarrow & f^{-1}_{A^{\epsilon}}(v,t) & \longleftarrow & f^{-1}_{A}(v,t)\\
 & \nwarrow\!\!\!\!\swarrow & & \nwarrow\!\!\!\!\swarrow & \\
f^{-1}_{B^{2\epsilon}}(v,t) & \longleftarrow & f^{-1}_{B^{\epsilon}}(v,t) & \longleftarrow & f^{-1}_{B}(v,t)
\end{array}$$

Since the sets in the above diagram are finite, it extends to a commutative diagram for every open $U\in\mathbb{S}^{d-1}\times\mathbb{R}$ and so, by the functoriality of the right derived functor there is an ϵ-interleaving of $\mathrm{PHT}(A)$ and $\mathrm{PHT}(B)$. So, $d_I(\mathrm{PHT}(A),\mathrm{PHT}(B)) = d_{\infty,\infty}^{\mathrm{PHT}^0}(A,B)$. For the case of points sets, the PHT Bottleneck distance turns out to be:

$$d_{\infty,\infty}^{\mathrm{PHT}^0}(A,B) = \max_{v\in\mathbb{S}^{d-1}} \inf_{\substack{\phi:A\to B\\ \text{matching}}} \max_{a_i\in A} |a_i \cdot v - \phi(a_i)\cdot v| = \inf_{\substack{\phi:A\to B\\ \text{matching}}} \max_{a_i\in A} \|a_i - \phi(a_i)\|_{\mathbb{R}^d}$$

To prove Equation 4.3, we need to calculate,

$$d_{2,2}^{\mathrm{PHT}^0}(A,B) = \left(\int_{v\in\mathbb{S}^{d-1}} \inf_{\substack{\phi:A\to B\\ \text{matching}}} \sum_{i=1}^{n} |a_i \cdot v - \phi(a_i)\cdot v|^2 d\mu\right)^{1/2}$$

where μ is the Lebesgue measure on $\mathbb{S}^{d-1}$. Let $X_i = a_i - \phi(a_i)$, and so rewriting the above expression in terms of matrices gives,

$$\begin{aligned}
\left(d_{2,2}^{\mathrm{PHT}^0}(A,B)\right)^2 &\leqslant \inf_{\substack{\phi:A\to B\\ \text{matching}}} \int_{v\in\mathbb{S}^{d-1}} v^T \left(\sum_{i=1}^{n} X_i X_i^T\right) v\ d\mu \\
&= \inf_{\substack{\phi:A\to B\\ \text{matching}}} \int_{v\in\mathbb{S}^{d-1}} v^T Y v\ d\mu \\
&= \inf_{\substack{\phi:A\to B\\ \text{matching}}} \int_{v\in\mathbb{S}^{d-1}} \mathrm{tr}(v^T Y v)\ d\mu \\
&= \inf_{\substack{\phi:A\to B\\ \text{matching}}} \int_{v\in\mathbb{S}^{d-1}} \mathrm{tr}(Y v v^T)\ d\mu \\
&= \inf_{\substack{\phi:A\to B\\ \text{matching}}} \mathrm{tr}\left(Y \int_{v\in\mathbb{S}^{d-1}} v v^T\ d\mu\right)
\end{aligned}$$

where Y is $\sum_{i=1}^n X_i X_i^T$. The integral $A := \int_{v\in\mathbb{S}^{d-1}} vv^T\ d\mu$ is invariant under conjugation by orthogonal matrices, since the sphere is rotationally invariant. So, in particular for orthogonal matrix U we have,

$$A = \int_{v\in\mathbb{S}^{d-1}} vv^T\ d\mu = \int_{v\in\mathbb{S}^{d-1}} Uvv^T U^T\ d\mu = UAU^T.$$

The matrix A commutes with orthogonal matrices and so by Schur's Lemma, A must be a scalar times the identity matrix.

$$A = \int_{v\in\mathbb{S}^{d-1}} vv^T = \lambda I$$

where λ can be found my taking trace on both sides i.e.,

$$\int_{v\in\mathbb{S}^{d-1}} d\mu = \lambda n.$$

So, we have,

$$\begin{aligned}\inf_{\substack{\phi:A\to B\\ \text{matching}}} \operatorname{tr}\left(Y\int_{v\in\mathbb{S}^{d-1}} vv^T\, d\mu\right) &= \inf_{\substack{\phi:A\to B\\ \text{matching}}} \frac{1}{n}\operatorname{tr}(Y)\text{Area}(\mathbb{S}^{d-1})\\ &= \inf_{\substack{\phi:A\to B\\ \text{matching}}} \frac{1}{n}\|A-\phi(A)\|_F^2\text{Area}(\mathbb{S}^{d-1}).\end{aligned}$$

Finally, to prove Equation 4.4, we note that

$$\begin{aligned}\left(d_{2,\infty}^{\text{PHT}^0}(A,B)\right)^2 &= \max_{v\in\mathbb{S}^{d-1}} \inf_{\substack{\phi:A\to B\\ \text{matching}}} \sum_{i=1}^n |a_i\cdot v-\phi(a_i)\cdot v|^2\\ &= \inf_{\substack{\phi:A\to B\\ \text{matching}}} \max_{v\in\mathbb{S}^{d-1}} v^TYv\\ &= \inf_{\substack{\phi:A\to B\\ \text{matching}}} \lambda_{\max}(Y)\\ &= \inf_{\substack{\phi:A\to B\\ \text{matching}}} \lambda_{\max}\left((A-\phi(A)(A-\phi(A))^T\right)\\ &= \inf_{\substack{\phi:A\to B\\ \text{matching}}} \|A-\phi(A)\|_\infty,\end{aligned}$$

where $\|\cdot\|_\infty$ stands for the ∞-Schatten norm of the matrix. □

The appearance of matrix norm expressions for our PHT distances is a welcome development, as it permits a qualitative comparison with a certain class of Procrustes distances.

Remark 4.2.2 (Comparison with Procrustes Distances)**.** The general **Procrustes distance** [DM98] between two *ordered* point clouds in $\mathbb{R}^d$, scaled by their centroid sizes[2], $A=[a_1,\cdots,a_n]$ and $B=[b_1,\cdots,b_n]$ is

$$d_P(A,B)=\inf_{R\in\mathcal{R}}\left(\sum_{i=1}^n\|Ra_i-b_i\|^2\right)^{1/2}=\inf_{R\in\mathcal{R}}\|RA-B\|_F$$

[2] The centroid size is given by $\left(n^{-1}\sum_{i=1}^n\|a_i-\bar{a}\|^2\right)^{\frac{1}{2}}$ where $\bar{a}$ is the mean of $\{a_i\}_{i=1}^n$.

where $\mathcal{R}$ is the group of rigid motions on $\mathbb{R}^d$. One advantage of the PHT distances considered here is that apriori no ordering needs to be put on the points in the point clouds. On the other hand, the PHT distances are sensitive to the embedding and consequently point clouds are not compared via their optimal alignments.

However, there is a closer connection between the ECT/PHT and the *orthogonal* Procrustes distance, which is given by

$$d_{OP}(A, B) = \min_{R \text{ s.t. } R^T R = I} \|RA - B\|_F,$$

and whose solution $R = UV^T$ is determined by the singular value decomposition (SVD) of $B^T A = U\Sigma V^T$. By comparison, [CMT22a, Thm. 6.7] states that if two generic simplicial complexes M and N in $\mathbb{R}^d$ have identical pushforward measures on the space of Euler curves (or persistence diagrams), i.e., if $\mathrm{ECT}(M)_*\mu = \mathrm{ECT}(N)_*\mu$, then M and N are related by an $O(d)$ action. This, along with Remark 2.6.4, suggests that one could modify the PHT distances here to also consider a minimization procedure along orthogonal transformations or rigid motions.

4.3 The Stability Theorem

In general the PHT distances are hard to compute, so often times we need to use other notions to bound the PHT distance. In this section we prove that if two shapes are homotopic through an ϵ-controlled homotopy, then their derived PHTs (Definition 2.6.5) are also ϵ-close in the interleaving distance.

Theorem 4.3.1 (Stability of the PHT under Controlled Homotopies)**.** *Let $M, N \subseteq \mathbb{R}^d$ be constructible sets and let $\varphi : M \to N$ and $\psi : N \to M$ be a homotopy equivalence of M and N; that is, there are homotopies $H_M : M \times I \to M$ and $H_N : N \times I \to N$ connecting Id_M to $\psi \circ \varphi$ and Id_N to $\varphi \circ \psi$. If there is some $\epsilon > 0$ such that $\|x - \varphi(x)\|^2_{\mathbb{R}^d} \leqslant \epsilon$ for all $x \in M$ and $\|y - \psi(y)\|^2_{\mathbb{R}^d} \leqslant \epsilon$ for all $y \in N$, where further $\|x - H_M(x, s)\|^2_{\mathbb{R}^d} \leqslant 2\epsilon$ for all $x \in M$ and $\|y - H_N(y, s)\|^2_{\mathbb{R}^d} \leqslant 2\epsilon$ for all $y \in N$ and $s \in I$, then the PHT of M and N are ϵ-interleaved.*

Proof. We show the following diagram is commutative.

$$\begin{array}{ccc}
R(f_{M^{2\epsilon}})_*\Bbbk_{Z_{M^{2\epsilon}}} & & R(f_{N^{2\epsilon}})_*\Bbbk_{Z_{N^{2\epsilon}}} \\
\downarrow & \swarrow\!\!\!\searrow & \downarrow \\
R(f_{M^{\epsilon}})_*\Bbbk_{Z_{M^{\epsilon}}} & & R(f_{N^{\epsilon}})_*\Bbbk_{Z_{N^{\epsilon}}} \\
\downarrow & \swarrow\!\!\!\searrow & \downarrow \\
R(f_{M})_*\Bbbk_{Z_{M}} & & R(f_{N})_*\Bbbk_{Z_{N}}
\end{array} \tag{*}$$

We prove this for the left triangle as the commutativity of the right triangle follows from a similar argument. Let $U \subset \mathbb{S}^{d-1} \times \mathbb{R}$ be a test open set and so $Rf_*\Bbbk_{Z_{M^\epsilon}}(U) = \left[\mathscr{S}^\bullet(f_{M^\epsilon}^{-1}(U);\Bbbk)\right]$ where $\left[\mathscr{S}^\bullet(f_{M^\epsilon}^{-1}(U);\Bbbk)\right]$ represents the class of complexes quasi-isomorphic to the singular cochain complex on $f_{M^\epsilon}^{-1}(U)$. We first prove that we have the following *non-commutative* diagram of topological spaces

$$\begin{array}{ccc}
f_{M^{2\epsilon}}^{-1}(U) & & \\
\uparrow & \overset{h}{\nwarrow} & \\
f_{M^{\epsilon}}^{-1}(U) & & f_{N^{\epsilon}}^{-1}(U) \\
\uparrow & \underset{g}{\nearrow} & \\
f_{M}^{-1}(U) & &
\end{array}$$

that commutes up to homotopy; that is $h \circ g$ is homotopic to $\iota : f_M^{-1}(U) \hookrightarrow f_{M^{2\epsilon}}^{-1}(U)$. If we apply the singular cochain functor to this diagram and then take quotients by quasi-isomorphisms, we will get the desired commutative trianglea in Equation 4.1.

We now explicitly describe the maps g and h. For $(x, v, t) \in f_M^{-1}(U)$ define g such that $g(x, v, t) = (\varphi(x), v, t + \epsilon)$. It remains to verify that $(\varphi(x), v, t)$ is in $f_{N^\epsilon}^{-1}(U)$. Note that

$$\varphi(x) \cdot v - x \cdot v = (\varphi(x) - x) \cdot v \leqslant \|\varphi(x) - x\|^2 \leqslant \epsilon.$$

Since $x \cdot v \leqslant t$, we have that $\varphi(x) \cdot v - t \leqslant \varphi(x) \cdot v - x \cdot v \leqslant \epsilon$ and so $\varphi(x) \cdot v \leqslant t + \epsilon$. Similarly for $(y, v, t) \in f_{N^\epsilon}^{-1}(U)$, let $h(y, v, t) = (\psi(y), v, t + \epsilon)$. After composing we get that $h \circ g(x, v, t) = (\psi \circ \varphi(x), v, t + 2\epsilon)$. To see that $h \circ g$ is homotopic to the inclusion

$\iota : f_M^{-1}(U) \hookrightarrow f_{M^{2\epsilon}}^{-1}(U)$, define a map $K : f_M^{-1}(U) \times I \to f_{M^{2\epsilon}}^{-1}(U)$ by

$$K((x, v, t), s) = (H_M(x, s), v, t + 2\epsilon).$$

By Cauchy-Schwarz,

$$H_M(x, s) \cdot v - x \cdot v = (H_M(x, s) - x) \cdot v \leqslant \|H_M(x, s) - x\|^2 \leqslant 2\epsilon$$

for all $s \in I$. Since $x \cdot v \leqslant t$, we have $H_M(x, s) \cdot v \leqslant t + 2\epsilon$. This means that $K((x, v, t), s) = (H_M(x, s), v, t + 2\epsilon) \in f_{M^{2\epsilon}}^{-1}(U)$ for $(x, v, t) \in f_M^{-1}(U)$. Further, $K((x, v, t), 0) = (x, v, t + 2\epsilon)$ and $K((x, v, t), 1) = (\psi \circ \varphi(x), v, t + 2\epsilon) = h \circ g(x, v, t)$. Continuity follows from continuity of H_M and so K is a homotopy between the inclusion map and $h \circ g$. □

Application of Lemma 4.1.4 implies the following,

Corollary 4.3.2. *(Stability of the PHT Bottleneck distance) Under the assumptions of Theorem 4.3.1 for all $i \geqslant 0$,*

$$d_{\infty,\infty}^{\mathrm{PHT}^i}(M, N) \leqslant \epsilon.$$

Remark 4.3.3. One can also bound the (∞, q)-PHT distance when $q \neq \infty$. To see this, the (∞, q)-PHT distance is the q-th integral norm of the bottleneck distance, integrated over all directions. Consequently, the (∞, q)-PHT distance is bounded by $\epsilon \cdot \mathrm{Area}(\mathbb{S}^{d-1})^{\frac{1}{q}}$, thereby establishing stability.

Example 4.3.4. We now calculate and bound some PHT distances between the shapes A (in blue) and B (in red) in Figure 4.1. The normals v_1, v_2 are $(\frac{2}{\sqrt{5}}, \frac{1}{\sqrt{5}})$ and $(-\frac{2}{\sqrt{5}}, \frac{1}{\sqrt{5}})$. Since all the sublevel sets of the shapes are contractible, it suffices to only consider PHT^0 of the two shapes. In particular, the PHT bottleneck distance in degree 0 is

$$\begin{aligned} d_{\infty,\infty}^{\mathrm{PHT}^0}(A, B) &= \max_{\theta \in \mathbb{S}^1} \mathcal{W}_\infty \left(\mathrm{PH}^0(A, \theta), \mathrm{PH}^0(B, \theta)\right) \\ &= \left|(-4, 2) \cdot \left(\frac{2}{\sqrt{5}}, \frac{1}{\sqrt{5}}\right)\right| \\ &= \frac{6}{\sqrt{5}}. \end{aligned}$$

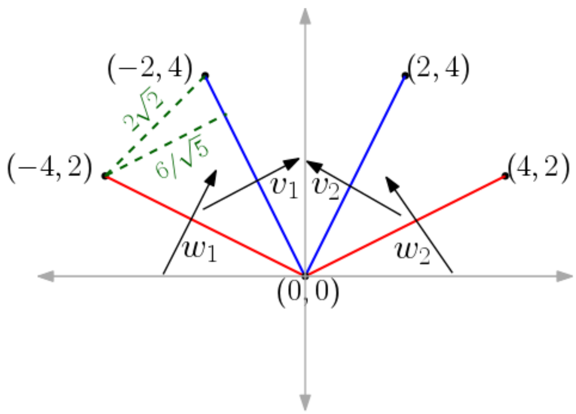

FIGURE 4.1: Distances between the PHTs of A (in blue) and B (in red).

The direction at which the maximum occurs is v_1 and v_2.

The PHT sheaf interleaving distance is hard to compute in practice. However, the application of Theorem 4.3.1 gives an upper bound on the PHT sheaf interleaving distance. In this example, we have that A and B are homeomorphic to each other. Explicitly, the linear map sending $(-4,2)$ to $(-2,4)$ can be extended to a homeomorphism $\varphi : B \to A$ where $(x, \frac{1}{2}|x|) \mapsto (\frac{1}{2}x, |x|)$. Further, the maximum movement of points under map φ is, $\|(x, \frac{1}{2}|x|) - \varphi(x, \frac{1}{2}|x|)\|_{\mathbb{R}^2} \leqslant 2\sqrt{2}$. So, by Theorem 4.3.1 and Lemma 4.1.4

$$2\sqrt{2} \geqslant d_I\left(\mathrm{PHT}(A)), \mathrm{PHT}(B)\right) \geqslant \frac{6}{\sqrt{5}}.$$

4.4 Point Samples for Approximating the PHT

After having established various distances on the PHT, we are now in a position to describe how to approximate a shape with point samples so that the resulting PHTs are close. We note that this is the only section where we require a manifold hypothesis on our shapes M. This is because the problem of approximating a general constructible set (or stratified space) is not well understood. Instead we rely on the following sampling and inference result, which makes implicit use of the injectivity radius of an embedded submanifold. This is encoded via the condition number τ, but we refer to [NSW08, §2] for a more detailed description of this.

Theorem 4.4.1 (Theorem 3.1 in [NSW08])**.** *Let M be a compact submanifold of $\mathbb{R}^d$ with condition number τ. Let $\bar{x} = \{x_1, .., x_n\}$ be a set of n points drawn independently and identically from a uniform probability measure on M. Let $0 < \epsilon < \tau/2$. Let $U = \cup_{x \in \bar{x}} B_\epsilon(x)$ be the union of the open balls of radius ϵ around the sample points. Then for all*

$$n > \beta_1 \left(\log \beta_2 + \log \frac{1}{\delta} \right)$$

the homology of U equals the homology of M with probability $> 1 - \delta$. Here β_1 and β_2 are constants that depend on the condition number τ, ϵ and the volume of M. The bound on n ensures that with high probability the sample is $\frac{\epsilon}{2}$-dense in M.

We let $\mathcal{U} := \{B_\epsilon(x)\}$ be the balls produced by Theorem 4.4.1. As the set of sampled points $X = \{x\}$ is embedded in $\mathbb{R}^d$, we can consider the Voronoi cell decomposition $\mathcal{V} = \{V_x\}$ of $\mathbb{R}^d$ generated by X—this is the decomposition of $\mathbb{R}^d$ into convex regions where every point $y \in \text{int}V_x$ is closest to x and no other point $x' \in X$. The **alpha complex** [Ede93] is defined to be the (geometric) nerve of the cover $\{B_\epsilon(x) \cap V_x\}$ of U; see Figure 4.2. By the nerve lemma, the alpha complex is homotopy equivalent to union of the balls U and so with high probability Theorem 4.4.1 says that the homology of K is equal to homology of M. We now promote this to an observation about the PHT.

Corollary 4.4.2. *(Approximation) Let M be a compact submanifold of $\mathbb{R}^d$ with condition number τ. Let $\bar{x} = \{x_1, .., x_n\}$ be a set of n points drawn independently and identically from a uniform probability measure on M. Let $0 < \epsilon < \frac{\tau}{2}$. Let $U = \cup_{x \in \bar{x}} B_\epsilon(x)$ be the union of the open balls of radius ϵ around the sample points. Let K be the alpha complex of U. Then for all*

$$n > \beta_1 \left(\log \beta_2 + \log \frac{1}{\delta} \right)$$

we have that, $d_I(\text{PHT}(M), \text{PHT}(K)) \leqslant \epsilon^2$ with high confidence i.e. probability > 1- δ.

Proof. We show that the assumptions of Theorem 4.3.1 are satisfied with $0 < \epsilon < \tau/2$ and then apply Theorem 4.3.1 to conclude the result. We need to find a homotopy equivalence

$\varphi : K \to M$ and $\psi : M \to K$ that satisfies the assumptions of Theorem 4.3.1. We do this by passing through the union of balls U as an intermediary.

Homotopy equivalence of M and U: Since the sample is $\epsilon/2$-dense in M with high probability, there is an inclusion of M into U with high probability. Let ι be this inclusion and let f be the projection that sends $x \mapsto \arg\min_{p \in M} \|p - x\|$. By the definition of condition number, the distance between any $q \in M$ to the medial axis is greater than $\tau > 2\epsilon$. The well-definedness of f is equivalent to $x \in U$ not contained in the medial axis. Suppose $x \in U$ is in the medial axis, so $\|p - x\| > 2\epsilon$ for every $p \in M$. Since $\bar{x}$ is $\epsilon/2$-dense in M w.h.p and $U = \cup_{y \in \bar{x}} B_\epsilon(y)$, it must be that for any $p \in M$, $\|p - x\| \leqslant \|p - y\| + \|y - x\| < \frac{\epsilon}{2} + \epsilon < \frac{3}{2}\epsilon$. This is a contradiction to x in the medial axis. This proves that f is well-defined. The map f is a deformation retraction and can be seen by taking the homotopy $H_U(x,t) = tx + (1-t)f(x)$ for all $x \in U$ and $t \in [0,1]$. Further, $\|H_U(t,x) - x\| = \|tx + (1-t)f(x) - x\| = (1-t)\|x - f(x)\| < (1-t)\epsilon/2 \leqslant \epsilon$.

Homotopy equivalence of U and K: We have the inclusion map $j : K \to U$. There is a retract $g : U \to K$ which follows the lines in U connected to the nearest point in K; see Figure 4.2. We call the homotopy that follows these lines G_U. Since U is a union of balls of radius ϵ, the homotopy G_U does not move the points of U more than ϵ.

Let $\varphi := f \circ j$ and $\psi := g \circ \iota$. On composing, $\varphi \circ \psi = f \circ j \circ g \circ \iota \sim f \circ \mathrm{Id}_U \circ \iota = f \circ \iota \sim \mathrm{Id}_M$ and similarly $\psi \circ \varphi \sim \mathrm{Id}_K$. The radius of balls are less than ϵ and so for $x \in K$, $\|x - f \circ j(x)\| \leqslant \epsilon$ and so $\|x - \varphi(x)\|^2 < \epsilon^2$. Since the sample points are $\epsilon/2$-dense, for $y \in M$, $\|y - g \circ \iota(y)\| \leqslant \epsilon/2 < \epsilon$ and so $\|y - \psi(y)\|^2 \leqslant \epsilon^2$. Since the homotopy maps H_U and G_U satisfy the conditions $\|H_U(x,t) - x\| \leqslant \epsilon$ and $\|G_U(x,t) - x\| \leqslant \epsilon$ for all $x \in U$, the homotopy map K_1 connecting $\varphi \circ \psi$ and Id_M satisfy $\|K(x,t) - x\| \leqslant \epsilon$ for all $x \in M$. The case for K_2 connecting $\psi \circ \varphi$ and Id_K is similar. Applying Theorem 4.3.1 gives ϵ^2-interleaving of the PHTs of M and K. □

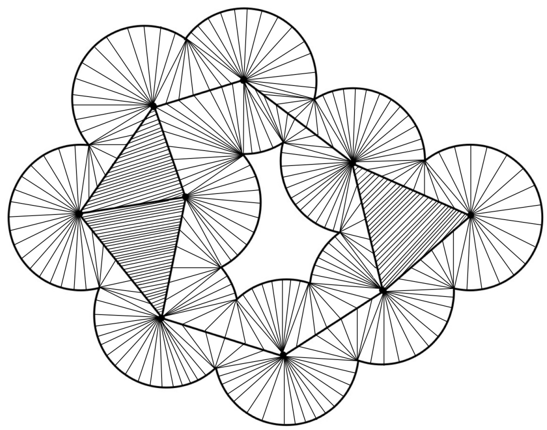

FIGURE 4.2: A collection of disks in the plane has a simpler summary given by its alpha complex, which is the geometric simplicial complex shown with shaded simplices above. The radial lines in each disk shows how points in the union of the disks deformation retract onto its alpha complex. Figure inspired by Fig. 3.1 of [Ede93].

5. Diffusion on stratified spaces

In this chapter, we describe a method of constructing diffusion on stratified spaces. More precisely, we consider subanalytic spaces because the volume measure on such spaces is finite. We briefly give an introduction to subanalytic spaces.

5.1 Subanalytic sets

Stratified spaces constitute a large category of spaces, and dealing with such generality can be very challenging. Therefore, we restrict to a subclass of singular spaces—subanalytic sets—since they satisfy many regularity properties, which make them amenable to analysis. Subanalytic sets align with Grothendieck's program on tame topology [Gro97]; for instance, subanalytic sets are triangulable. In this section, we define subanalytic sets and discuss some of the regularity conditions they satisfy.

Definition 5.1.1. (Semi-analytic sets) A subset $X \subset \mathbb{R}^n$ is called **semi-analytic** if it is locally defined by finitely many equalities and inequalities of real analytic functions. Precisely, for each $p \in \mathbb{R}^n$ there is a neighbourhood U of p and real analytic functions f_{ij} and g_{ij} on U, where $i = 1, \ldots, r$ and $j = 1, \ldots s_i \in \mathbb{N}$ such that

$$X \cap U = \cup_{i=1}^{r} \cap_{j=1}^{s_i} \{x \in U : g_{ij}(x) > 0, f_{ij}(x) = 0\}.$$

It is important to note that p is any point in $\mathbb{R}^n$ and are not just points in X.

Example 5.1.2. Since polynomials are semi-analytic functions, semi-algebraic sets are clearly semi-analytic.

We work with a bigger class of sets, namely with the class of sets being locally proper projections of semi-analytic sets.

Definition 5.1.3. (Subanalytic sets) A subset $X \subset \mathbb{R}^n$ is subanalytic if there exists a semi-analytic set $Z \subset \mathbb{R}^{n+p}, p \in \mathbb{N}$ such that $X = \pi(Z)$ where $\pi : \mathbb{R}^{n+p} \to \mathbb{R}^n$ is the projection onto the first n coordinates.

Example 5.1.4. Clearly, semi-analytic sets are subanalytic, and so semi-algebraic sets are examples of subanalytic sets.

Every subanalytic set admits a Whitney stratification [Hir73]. This property extends to definable sets in arbitrary o-minimal structures [Lê 98]. We remind the reader that subanalytic sets, in general, are not definable sets, however, globally subanalytic sets as defined in 2.2.3 are definable. Generally, we observe the following inclusions of classes of stratified spaces.

$$\text{Semi-algebraic sets} \hookrightarrow \text{Semi-analytic sets} \hookrightarrow \text{Subanalytic sets} \hookrightarrow \text{Stratified spaces}$$

Remark 5.1.5. Below are some important properties of subanalytic sets.

1. The intersection and union of a finite collection of subanalytic sets are subanalytic.

2. A subanalytic set is locally connected.

3. The closure, complement, and interior of a subanalytic set is subanalytic.

4. Let $X \subset \mathbb{R}^n$ be a closed subanalytic set. Then for every point in X, there exists a neighbourhood U such that $X \cap U = \pi(A)$ where A is a closed analytic subset of $U \times R^m$ for some m, and $\dim A = \dim X \cap U$ and $\pi : U \times \mathbb{R}^m \to U$ is a proper map when restricted to A [BM88].

5.1.1 Lengths on stratified spaces

Recall that in any path connected metric space (X, d), the geodesic distance between any two points can be defined as

$$\delta(x, y) = \inf\{l(\gamma) \mid \gamma \in \mathcal{C}([0,1], X), \gamma(0) = x \text{ and } \gamma(1) = y\},$$

where $\mathcal{C}([0,1], X)$ is the set of all curves that are defined on $[0,1]$ and the *length* of a curve $\gamma : [0,1] \to X$ is given by

$$l(\gamma) = \sup\left\{ \sum_{0\leqslant i\leqslant k} d\left(\gamma(t_j), \gamma(t_{j-1})\right) \,\middle|\, k \in \mathbb{N},\ 0 = t_0 < t_1 \cdots < t_{k-1} < t_k = 1 \right\}.$$

Let $X \subseteq \mathbb{R}^n$ and let d be the Euclidean distance restricted to X. Then it is clear that the geodesic distance δ between any two points $x, y \in X$ is at least the Euclidean distance, that is,

$$\|x - y\| = d(x, y) \leqslant \delta(x, y).$$

However, the reverse is not necessarily true, even up to some constant factor. In fact, one can come up with examples where the geodesic distance is infinite. So, the metrics need not be equivalent. Tougeron came up with a notion of l-regularity which gives an upper bound on the geodesic distance in terms of the Euclidean distance.

A compact set $K \subseteq \mathbb{R}^n$ is said to be ***l*-regular** with $l \geqslant 1$, if K is finitely path connected and there exists a constant $C > 0$ such that

$$\delta(x, y) \leqslant C d(x, y)^{1/l},$$

for all $x, y \in K$.

Definition 5.1.6 (Whitney-Tougeron regular sets)**.** Let X be a subset of $\mathbb{R}^n$. We say X is **Whitney-Tougeron regular** if for every $x \in X$ there is an $l(x) \geqslant 1$ depending on x and a $l(x)$-regular compact neighbourhood $K \subset \mathbb{R}^n$.

Remark 5.1.7. We remark that all these definitions carry forward to general stratified sets using singular charts as defined in [Pfl01].

Example 5.1.8. Every subanalytic set in $\mathbb{R}^n$ is regular [KO97].

Remark 5.1.9. In fact, there is a more intrinsic way to define distances on stratified spaces. There exists a smooth Riemannian metric μ on subanalytic sets such that (X, δ_μ) becomes a length space. Here δ_μ is the geodesic distance induced by the metric μ. Moreover, the metric reflects the underlying topology of X. See Section 2.4 in [Pfl01].

5.1.2 Volume on stratified spaces

In this section, we define volumes of subsets of stratified spaces based on Ferrarotti's work in [Fer86].

Let $(X, \mathscr{S})$ be a closed stratified subset of $\mathbb{R}^d$. Recall a σ-compact subset of a topological space is a set that is expressible as a countable union of compact sets.

Definition 5.1.10 (Volume)**.** The volume of a σ-compact set Y of X is defined as

$$\mathbf{Vol}\,(Y, \mathscr{S}) := \varinjlim_{N \in \mathcal{N}} \mathbf{Vol}\,(Y \backslash N)\,, \tag{5.1}$$

where $\mathcal{N}$ is the directed system of family of open neighbourhoods of Σ in X, ordered by $N \leqslant N'$ if and only if $N \supseteq N'$.

The volume $\mathbf{Vol}\,(Y \backslash N)$ makes sense because $Y \backslash N$ belongs to a submanifold which has a well-defined notion of volume.

Remark 5.1.11. Below are some properties of volumes on stratified spaces.

1. The limit in Equation 5.1 exists (possibly being infinite) because the volume of $Y \backslash N$ is increasing in $\mathcal{N}$.
2. The volume of σ-compact subset of Σ is 0.
3. If X is a Riemannian manifold then the volume coincides with the usual definition of volume on Riemannian manifolds.

Surprisingly, this definition of volume does not depend on the stratification $\mathscr{S}$ [Fer86, Proposition I.5]. Since the volume of Riemannian manifolds are measures, we expect that the volume on stratified spaces are also measures.

Lemma 5.1.12 ([Fer86])**.** *Let $\mathcal{Y}$ be a family of σ-compact subsets of X. Then*

1. *For every $Y \in \mathcal{Y}$, $\mathbf{Vol}\,(Y) \geqslant 0$.*
2. *If $Y \subseteq Y'$ for $Y, Y' \in \mathcal{Y}$, $\mathbf{Vol}\,(Y) \leqslant \mathbf{Vol}\,(Y')$.*

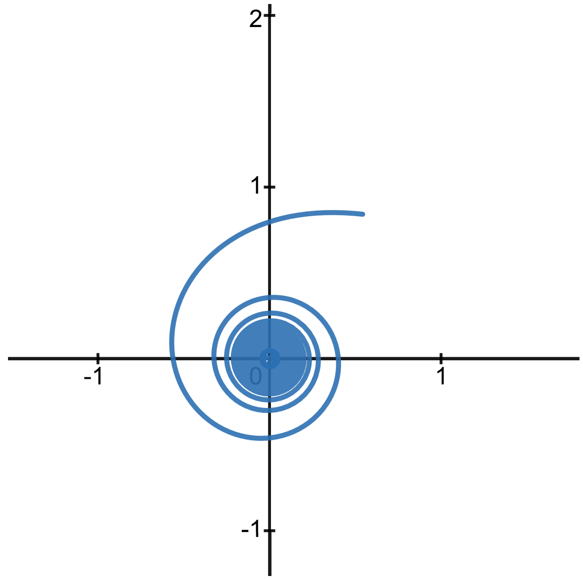

FIGURE 5.1: The spiral X defined as image of the curve given by Equation 5.2 and the origin. The volume of a ball around the origin has infinite volume.

3. *Countable additivity: If* $\{Y_i\}_{i\in\mathbb{N}}$ *is a sequence of disjoint sets of* $\mathcal{Y}$*, then* $\mathbf{Vol}\,(\cup_i Y_i) = \sum_{i=1}^{\infty} \mathbf{Vol}\,(Y_i)$.

4. *For every* $Y \in \mathcal{Y}$, $\mathbf{Vol}\,(Y) = \mathbf{Vol}\left(Y \cap \overline{X\backslash\Sigma}\right)$.

However, the volume can blow up to infinity. Below is an example describing this situation.

Example 5.1.13 (Spiral)**.** Let $X \subseteq \mathbb{R}^2$ be defined as the union of the image of the curve

$$\gamma(t) = \left(\sqrt{t}\cos\left(\frac{1}{t}\right), \sqrt{t}\sin\left(\frac{1}{t}\right)\right), \tag{5.2}$$

when $0 < t < (2\pi)^{-1}$ and the origin $(0,0)$. The origin is the singular stratum and the one-dimensional curve is the regular stratum (See Figure 5.1). In this case

$$\mathbf{Vol}\,(X) = \lim_{s\to 0}\int_s^{2\pi^{-1}} \frac{1}{\sqrt{t}}\sqrt{\frac{1}{4}+\frac{1}{t^2}}dt \geqslant \lim_{s\to 0}\int_s^{2\pi^{-1}} \frac{1}{t\sqrt{t}} \geqslant \lim_{s\to 0} \left.\frac{-2}{\sqrt{t}}\right|_s^{2\pi^{-1}} = \infty.$$

This suggests even for simple one-dimensional stratified spaces, the volume may be infinite. We are interested in spaces where the volume of compact subsets does not blow up to infinity. One of the main results of Ferrarotti's paper [Fer86] is that the subanalytic spaces have locally finite volume. In fact, subanalytic spaces satisfy a stronger volume regularity property known as v-regularity.

Definition 5.1.14 (v-regularity)**.** Let $(X, \mathscr{S})$ be a closed stratified subset of $\mathbb{R}^d$ of dimension n. If for every relatively compact open set K of $\mathbb{R}^d$ there is a positive constant $c(K)$ depending only on K such that

$$\mathbf{Vol}\left(X \cap B_r^{\mathrm{Euc}}\right) \leqslant c(K) r^n,$$

for each ball B_r^{Euc} in $\mathbb{R}^d$ with $B_r^{\bar{\mathrm{Euc}}} \subset K$, then X is said to be v-regular.

Theorem 5.1.15. *([Fer86, Theorem IV.2]) Let X be a subanalytic set in $\mathbb{R}^d$. Then X is v-regular.*

In summary, the geodesic distance between any two points in a subanalytic set is finite and the volume of a subset subanalytic set is finite. This suggests that the subanalytic sets can be regarded as metric measure spaces.

For us, a metric measure space means the following.

Definition 5.1.16 (Metric measure spaces)**.** A *metric measure space* is a triple (X, d, m) where (X, d) is a locally compact separable metric space and m is a Radon measure on X with $m(U) \geqslant 0$ for each non-empty open set $U \subset X$.

Lemma 5.1.17. *Any compact subanalytic set of $\mathbb{R}^d$ is a metric measure space.*

Proof. Let X be a subanalytic space with metric given by the geodesic distance δ and measure of any σ-compact subset A of X given by $\mathbf{Vol}(A)$. Since X is compact and the metric δ induces the topology of X ([Pfl01, Theorem 2.4.7]), every open subset of X is σ-compact. Theorem 5.1.15 guarantees that the volume of every open set $U \subset X$ is finite. The tightness of the measure follows from the fact that X is compact and so the measure is Radon. □

5.2 Analysis on metric measure spaces

So far, we have established that compact subanalytic sets in $\mathbb{R}^d$ are metric measure spaces. There is a well developed theory of analysis on metric measure spaces, to the extent that there is a notion of diffusion and heat kernels on such spaces. Much of this is due to Sturm [Stu98], where they use the theory of Dirichlet forms to construct diffusion on such spaces. In this section, we give a brief account of the theory of Dirichlet forms and how they relate to diffusion and further show that Dirichlet forms on subanalytic sets are well defined and of *diffusion type*, meaning they give rise to diffusion. The argument essentially boils down to ensuring that the volume measure of balls in subanalytic spaces grows and shrinks in a controlled manner.

5.2.1 Dirichlet Forms

Dirichlet forms are analytic objects that are used to construct certain Markov processes, especially in cases where typical tools of analysis such as the Laplacians are difficult to define. In this sense, they can be regarded as a generalized version of the Laplacian.

Let (X, m) be a measure space. For the sake of convenience assume that X is compact with Radon measure m and $\text{supp}(m) = X$. The assumption that X is compact is not necessary and with minor modifications, we can replace it with locally compact X. The space of all square-integrable functions with respect to the measure m is a Hilbert space

$$L^2(X, m) = \left\{ m\text{−measurable } f : X \to \overline{\mathbb{R}} \;\middle|\; \int_X f^2(x) m(dx) < \infty \right\},$$

with inner product

$$\langle f, g \rangle_{L^2(X,m)} := \int_X f(x) g(x) m(dx).$$

A **symmetric form** $\mathcal{E}$ on Hilbert space $L^2(X, m)$ with domain $\mathcal{D}(\mathcal{E})$ is a map $\mathcal{E}$: $\mathcal{D}(\mathcal{E}) \times \mathcal{D}(\mathcal{E}) \to \mathbb{R}$ such that:

1. $\mathcal{D}(\mathcal{E})$ is a dense subspace of $L^2(X, m)$,

2. $\mathcal{E}(au + bv, w) = a\mathcal{E}(u, w) + b\mathcal{E}(v, w)$ for $u, v, w \in \mathcal{D}(\mathcal{E})$ and $a, b \in \mathbb{R}$,

3. $\mathcal{E}(u, u) \geqslant 0$ for $u \in \mathcal{D}(\mathcal{E})$,

4. $\mathcal{E}(u, v) = \mathcal{E}(v, u)$ for any $u, v \in \mathcal{D}(\mathcal{E})$.

Some additional properties need to be satisfied if one wants the domain $\mathcal{D}(\mathcal{E})$ to also be a Hilbert space. Define for every $\alpha > 0$, a symmetric form $\mathcal{E}_\alpha$ by

$$\mathcal{E}_\alpha(u, v) = \mathcal{E}(u, v) + \alpha\langle u, v\rangle_{L^2(X,m)} \qquad u, v \in L^2(X, m).$$

Note that $\mathcal{D}(\mathcal{E})$ with inner product $\mathcal{E}_1$ forms an inner product space. We say the form $\mathcal{E}$ is *closed* if $\mathcal{D}(\mathcal{E})$ is complete with respect to $\mathcal{E}_1$. Since, the forms $\mathcal{E}_\alpha$ and $\mathcal{E}_\beta$ are comparable for different $\alpha, \beta > 0$, closure of $\mathcal{E}$ implies that $\mathcal{D}(\mathcal{E})$ is a Hilbert space with inner product $\mathcal{E}_\alpha$ for every $\alpha > 0$.

Definition 5.2.1 (Dirichlet Form)**.** A **Dirichlet form** is a symmetric form on Hilbert space $L^2(X, m)$ with domain $\mathcal{D}(\mathcal{E})$ such that:

1. $\mathcal{E}$ is closed,

2. For every $u \in \mathcal{D}(\mathcal{E})$, $v = (0 \vee u) \wedge 1$ also belongs to $\mathcal{D}(\mathcal{E})$ and $\mathcal{E}(v, v) \leqslant \mathcal{E}(u, u)$. We then say that unit contraction acts on $\mathcal{E}$.

Example 5.2.2 (Dirichlet form on $\mathbb{R}^n$)**.** The Dirichlet form on $\mathbb{R}^n$ with Lebesgue measure that corresponds to the standard Brownian motion is given by

$$\mathcal{E}(u, v) = \frac{1}{2}\int \nabla u \cdot \nabla v dx \qquad u, v \in W^{1,2}(\mathbb{R}^n),$$

where $W^{1,2}(\mathbb{R}^n) = \{u \in L^2(\mathbb{R}^n) \mid \partial_i u \in L^2(\mathbb{R}^n) \text{ for all } i\}$. One can readily check that $\mathcal{E}$ defined above is a Dirichlet form.

As suggested before, Dirichlet forms are in one-one correspondence with symmetric Markov processes. Refer to Appendix A for a short account of the theory of Markov semigroups and generators, and their correspondence with Markov processes. The following is a basic result in the theory of Dirichlet forms (See [FOT11]):

Theorem 5.2.3. *The following classes of objects are in one-one correspondence with each other:*

1. *closed symmetric forms* $(\mathcal{E}, \mathcal{D})$ *on* $L^2(X, m)$,
2. *strongly continuous semigroups on* $L^2(X, m)$,
3. *non-positive definite self-adjoint operators (or generators)* $\mathcal{L}$ *on* $L^2(X, \mu)$*).*

This correspondence is essentially given by the Riesz representation theorem.

$$\mathcal{D} = \mathcal{D}(\sqrt{-\mathcal{L}}); \qquad \mathcal{E}(u, v) = \langle \sqrt{-\mathcal{L}}u, \sqrt{-\mathcal{L}}v \rangle \text{ for } u, v \in \mathcal{D}.$$

When $u \in \mathcal{D}(\mathcal{L})$, then the correspondence can be characterized by $\mathcal{E}(u, v) = \langle -\mathcal{L}u, v \rangle$.

When the closed symmetric form satisfies the unit contraction property, i.e. it is a Dirichlet form, the corresponding stochastic process is a Markov process.

Example 5.2.4. (Connection between Dirichlet forms and generators on $\mathbb{R}^n$) It is well known that the generator of standard Brownian motion on $\mathbb{R}^n$ is $\frac{1}{2}\Delta$. Using the Dirichlet form on $\mathbb{R}^n$ defined as in Example 5.2.2, a simple integration by parts reveals that

$$\mathcal{E}(u, v) = \frac{1}{2} \int \nabla u \cdot \nabla v dx = \frac{1}{2} \int -\Delta u v dx = \left\langle -\frac{1}{2}\Delta u, v \right\rangle.$$

A Dirichlet form $\mathcal{E}$ is **regular** if there exists a **core** $\mathcal{C} \subset \mathcal{D}(\mathcal{E}) \cap \mathcal{C}(X)$ such that $\mathcal{C}$ is dense in $\mathcal{C}(X)$ with respect to the $\| \cdot \|_\infty$- norm and dense in $\mathcal{D}(\mathcal{E})$ with respect to the norm induced by $\mathcal{E}_1$. Here $\mathcal{C}(X)$ denotes the space of continuous functions on X.

A Dirichlet form $\mathcal{E}$ is **strongly local** if whenever u is constant on some neighbourhood of the support of v or vice-versa, then $\mathcal{E}(u, v) = 0$.

A Dirichlet form is called of **diffusion type** if it is both regular and strongly local. This is because when Dirichlet forms are regular, they correspond to certain types of Markov processes called Hunt processes. Hunt processes are special Markov processes that have right continuity of sample paths and satisfy the strong Markov property amongst others. When the Dirichlet forms are also strongly local, they correspond specifically to diffusion processes, meaning the sample paths are continuous.

Theorem 5.2.5 ([FOT11])**.** *Let $(\mathcal{E}, \mathcal{D})$ be a Dirichlet form of diffusion type on $L^2(X, m)$. Then, there exists a m-symmetric diffusion process $(\{X_t\}_{t>0}, (\mathbb{P}_x)_{x\in X})$ on X with Dirichlet form $\mathcal{E}$.*

Example 5.2.6 (Dirichlet forms on metric measure spaces [Stu98])**.** Let (X, d, m) be a metric measure space as in Definition 5.1.16. For each scale $r > 0$ and $u, v \in L^2(X, m)$

$$\mathcal{E}^r(u, v) = \frac{1}{2} \int_X N(x) \int_{B_r^*(x)} \frac{(u(x) - u(y))((v(x) - v(y)))}{d^2(x, y)} \frac{m(dy)m(dx)}{\sqrt{m(B_r(y))}\sqrt{m(B_r(x))}}, \quad (5.3)$$

where $B_r^*(x)$ is the ball of radius r around x excluding x itself and N can be any normalization function and can be seen to play the role of the local dimension.

Observe that when x is in a region that is locally like $\mathbb{R}^d$, then we can write by essentially the generalized Lebesgue differentiation theorem

$$|\nabla u(x)|^2 = d \lim_{r\to 0} \frac{1}{m(B_r(x))} \int_{B_r^*(x)} \left[\frac{u(y) - u(x)}{d(x, y)}\right]^2 m(dy),$$

and so Equation 5.3 is a generalization of $\int_{\mathbb{R}^d} |\nabla u|^2 dx$.

5.2.2 Gamma Limits

The study of Gamma limits or variational limits was introduced by E. de Giorgi in [De 75]. To grasp the main idea, consider a sequence of real-valued functions $\{F_h\}$ on a space X that converge pointwise to a function F. This does not necessarily imply that the minimum value of $\{F_h\}$ converges to the minimum value of F. The *gamma limit* or *variational limit*

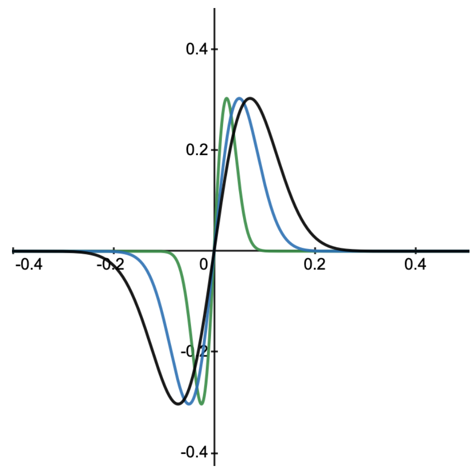

FIGURE 5.2: The sequence of functions $F_h(x) = hxe^{-2h^2x^2}$ for different values of h.

is a special kind of limit that remedies this; namely, the minimum value of $\{F_h\}$ converges to the minimum value of the gamma limit (if it exists). In this sense, these limits are stable under 'variational' or optimization problems. In this section, we give a brief introduction to gamma limits and list some of its properties.

Definition 5.2.7. (Γ-limit) The Γ- lower limit and the Γ- upper limit of the sequence of functions $\{F_h\}$ from a topological space X to $\overline{\mathbb{R}}$ are the continuous functions from X into $\overline{\mathbb{R}}$ defined as follows:

$$(\Gamma - \liminf_{h\to\infty} F_h)(x) := \sup_{U\in\mathcal{N}(x)} \liminf_{h\to\infty} \inf_{y\in U} F_h(y), \tag{5.4}$$

$$(\Gamma - \limsup_{h\to\infty} F_h)(x) := \sup_{U\in\mathcal{N}(x)} \limsup_{h\to\infty} \inf_{y\in U} F_h(y), \tag{5.5}$$

where $\mathcal{N}(x)$ is the set of all open neighbourhoods of x in X. Further, if the functions are equal, i.e. $\Gamma - \liminf_{h\to\infty} F_h = \Gamma - \limsup_{h\to\infty} F_h = F$, then write F as $\Gamma - \lim_{h\to\infty} F_h$. We say that $\{F_h\}$ converges (in the sense of Γ-convergence) to F or F is the Γ-limit of $\{F_h\}$.

Remark 5.2.8. The Γ-limit of a sequence of functions does not depend on the pointwise limit.

Examples 5.2.9. 1. Consider the sequence of functions $F_h(x) = hxe^{-2h^2x^2}$ defined on $X = \mathbb{R}$ (See Figure 5.2). The pointwise limit of this sequence of functions as $h \to \infty$

is the zero function. However, the Γ-limit when $x = 0$ is the infimum value of the function $-\frac{1}{2}e^{-1/2}$ at value $-\frac{1}{2h}$. When $x \neq 0$, the Γ-limit is just 0.

2. The Γ-limit of the sequence of functions $F_h(x) := \cos(hx)$ is the constant function -1.

Theorem 5.2.10. *(Fundamental theorem of* Γ*-convergence) If* $F_h \to F$ *in the sense of* Γ*-convergence, and* $x_h \in X$ *minimizes* F_h*, then every limit point of the sequence* $\{x_h\}$ *is a minimizer of* F*.*

The Γ-limit and the pointwise limit are related in the following way.

Lemma 5.2.11. *If* $F_h \to F$ *in the sense of* Γ*-convergence and* $F_h \to G$ *pointwise, then* $F \leqslant G$*. In particular, if the* Γ*-limit of* F_h *is* F *then the pointwise limit is* F *if and only if* $\limsup_{h\to\infty} F_h(x) \leqslant \Gamma - \lim_{h\to\infty} F_h(x)$*.*

Theorem 5.2.12 (Sequential Characterization of Gamma limits; Proposition 8.1 [Dal12])**.** *The* Γ*-limit of a sequence of functions* $\{F_h\}$ *is* F *if and only if the following conditions are satisfied:*

1. *For every* $x \in X$ *and every sequence* $x_h \to x$ *in* X

$$F(x) \leqslant \liminf_{h\to\infty} F_h(x_h);$$

2. *For every* $x \in X$*, there is a sequence* $x_h \to x$ *in* X *such that*

$$F(x) = \lim_{h\to\infty} F_h(x_h).$$

Theorem 5.2.13 (Existence of Gamma limits; Theorem 8.5 [Dal12])**.** *Assume* X *is second countable. Every sequence of functions* $\{F_h\}$ *has a* Γ*-convergent subsequence.*

5.3 Existence of strongly local forms on subanalytic sets

Let (X, d, m) be a metric measure space as in Definition 5.1.16. Let $\mathcal{E}^r$ be a Dirichlet form on X at scale r as given in Example 5.2.6. The pointwise limit of this family $\mathcal{E}^r$ as

$r \to 0$ may not be a reasonable object, for instance, it might not be a regular Dirichlet form, and so we instead work with gamma limits. The next theorem gives a criteria for when the gamma limit of Dirichlet forms $\mathcal{E}^r$ is strongly local and hence corresponds to a diffusion process.

Theorem 5.3.1 (Theorem 3.3 [Stu98]). *Let $(r_n)_n$ be a sequence with $r_n \to 0$ as $n \to \infty$ for which limit $\mathcal{E}^0 = \Gamma - \lim_{n\to\infty} \mathcal{E}^{r_n}$ exists. Then the closure of $(\mathcal{E}^0, \mathcal{C}_0^{\mathrm{Lip}})$ is a regular Dirichlet form. Further, it is strongly local if for every compact $Y \subset X$*

$$\limsup_{r\to 0} \sup_{\substack{x,y\in Y,\\ d(x,y)<r}} \frac{m(B_r(x))}{m(B_r(y))} < \infty \tag{5.6}$$

Remark 5.3.2. The above equation 5.6 can be thought of as a weaker notion of the volume doubling condition. The volume doubling condition given by

$$\limsup_{r\to 0} \sup_{y\in Y} \frac{m(B_{2r}(y))}{m(B_r(y))} < \infty \tag{5.7}$$

for each compact $Y \subset X$, implies the volume condition 5.6 because $m(B_{2r}(y)) > m(B_r(y))$. Riemannian manifolds satisfy the asymptotically volume doubling condition, as given in (5.7). The main idea is that in a manifold M of dimension n, the volume of geodesic balls behave as

$$vol(B_r) \sim r^n,$$

for r small enough, i.e., in the radius where the exponential map is defined.

More generally, they also satisfy the global volume doubling condition (i.e., for each r and not only for small r) when the Ricci curvature is non-negative. This follows from the Bishop-Gromov comparison theorem.

Let $X \subset \mathbb{R}^d$ be a compact subanalytic set with measure defined as in section 5.1.2 equipped with the geodesic distance δ. Let p be the dimension of the top dimensional stratum. Let $B_r(x) = \{y \in X \mid \delta(x, y) < r\}$ be an open ball of radius r in X and let $B^{\mathrm{Euc}}(x)$ be an open ball in $\mathbb{R}^d$ with the Euclidean distance.

Lemma 5.3.3. *For X defined as above, there exists a positive constant $C > 0$ such that*

$$\mathbf{Vol}\left(B_r(x)\right) \leqslant Cr^p, \tag{5.8}$$

for every $x \in X$ and $r > 0$.

Proof. Let $x \in X$, and $r > 0$. Since the volume is v-regular by Theorem 5.1.15, we have that there exists a positive constant C, such that $\mathbf{Vol}\left(B_r^{\mathrm{Euc}}(x) \cap X\right) \leqslant Cr^p$. Since the $d_{\mathrm{Euc}}(x,y) \leqslant \delta(x,y)$ for every pair of points $x, y \in X$, $\mathbf{Vol}\left(B_r(x)\right) \leqslant \mathbf{Vol}\left(B_r^{\mathrm{Euc}}(x) \cap X\right) \leqslant Cr^p$ for all $x \in X$.

□

This provides an upper bound on the volume growth rate of balls in X. However, the lower bound given by $\mathbf{Vol}\left(B_r(x)\right) > Cr^p$ may not hold for every r. Nevertheless, we can demonstrate the existence of a constant C_1 such that $\mathbf{Vol}\left(B_r(x)\right) \geqslant C_1 r^p$ for *small* r.

Lemma 5.3.4. *If the top-dimensional strata denoted as $X - \Sigma$ is open and dense in X, then*

$$\liminf_{r \to 0} \inf_{x \in X} \frac{\mathbf{Vol}\left(B_r(x)\right)}{r^p} > 0.$$

Proof. Suppose that

$$\liminf_{r \to 0} \inf_{x \in X} \frac{\mathbf{Vol}\left(B_r(x)\right)}{r^p} = 0.$$

Then there exists a sequence (r_i) of positive radii converging to 0 such that sequence $\inf_{x \in X} \frac{\mathbf{Vol}\left(B_{r_i}(x)\right)}{{r_i}^p}$ converges to 0. Without loss of generality, since there are finitely many strata, we can assume that there is only one connected p-dimensional stratum, denoted as S^p. Since S^p is dense in X, $B_{r_i}(x) \cap S^p$ is non-empty, and

$$\inf_{x \in X} \frac{\mathbf{Vol}\left(B_{r_i}(x)\right)}{{r_i}^p} \geqslant \inf_{x \in X} \frac{\mathbf{Vol}\left(B_{r_i}(x) \cap S^p\right)}{{r_i}^p}.$$

Since S^p is dense in X, for every $x \in X$ there exists a sequence $\{x_n\}_{n \in \mathbb{N}}$ in S^p such that $x_n \to x$. In particular, for $\epsilon = \frac{r_i}{2}$, there is a number $N(r_i)$ such that $\delta(x_n, x) \leqslant \frac{r_i}{2}$ for

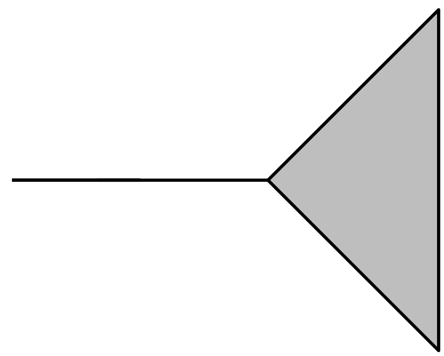

FIGURE 5.3: The Balalaika shape, represented by a solid triangle glued to a closed interval $[-1, 0]$ at its endpoint 0. The 2-dimensional stratum is not dense in the space.

$n \geqslant N = N(r_i)$. This implies that $B_{r_i/2}(x_N) \cap S^p \subseteq B_{r_i}(x) \cap S^p$ and so

$$\inf_{x \in X} \frac{\mathbf{Vol}\,(B_{r_i}(x) \cap S^p)}{r_i{}^p} \geqslant \inf_{x \in X} \frac{\mathbf{Vol}\left(B_{r_i/2}(x_N) \cap S^p\right)}{r_i{}^p},$$

where x_N denotes the point on S^p such that $d(x, x_N) \leqslant r_i/2$. By the continuity of Radon measures, the function $x \mapsto \frac{\mathbf{Vol}\left(B_{r_i/2}(x_N) \cap S^p\right)}{r_i{}^p}$ is continuous. Since X is compact there is a minimizer $\tilde{x}$, that depends only on r_i which realizes the infimum, call it $C(r_i)$. Since $B_{r_i/2}(\tilde{x}_N) \cap S^p$ is an open ball in the smooth manifold S^p, we know that as $r_i \to 0$, $\frac{\mathbf{Vol}\left(B_{r_i/2}(\tilde{x}_N) \cap S^p\right)}{(r_i/2)^p}$ is a constant greater than 0 which gives a contradiction.

□

Remark 5.3.5. The top-dimensional stratum being dense in X is a necessary condition. For example, in Figure 5.3, for a point x in the interval, the volume of a small ball is 0 since it completely lies within the singular stratum.

Combining lemmas 5.3.3 and 5.3.4 we can conclude the following.

Theorem 5.3.6. *For compact subanalytic sets X, with an open and dense top dimensional stratum, the volume condition 5.6 holds.*

Proof. Observe by Lemma 5.3.3, $\frac{m(B_r(x))}{m(B_r(y))} < \frac{Cr^p}{m(B_r(y))}$ and then apply the previous lemma 5.3.4 to prove the statement of the theorem. □

Corollary 5.3.7. *For each limit point $\mathcal{E}^0$ in the sense of Γ-convergence of a sequence of Dirichlet forms $\{\mathcal{E}^r\}_{r>0}$, there is a symmetric diffusion process $(X_t, \mathbb{P}_x)$ on X with*

$$\mathcal{E}^0(u,u) = \lim_{t\to 0} \frac{1}{2t} \int \mathbb{E}_x \left[u(X_t) - u(x)\right]^2 m(dx)$$

for all lipschitz u in $L^2(X, m)$.

Remark 5.3.8. In the case of Riemannian manifolds, where the distance metric is induced by the Riemannian metric and the measure is the usual notion of Riemannian volume, it satisfies a stronger volume regularity property known as the *measure contraction property* [Stu98, Definition 4.1]. It turns out that this volume property implies that the pointwise limit is a reasonable object, in particular, the pointwise limit equals the gamma limit. Unfortunately, as we will see later on in Example 5.3.13, for the case of compact subanalytic sets with open and dense top dimensional stratum, this property is not true.

Definition 5.3.9 (Weak measure contraction property [Stu98])**.** A metric measure space (X, d, m) satisfies the *weak measure contraction property* (MCP for short) with exceptional set $Z \subset X$ with $m(Z) = 0$ if for every compact set $Y \subset X \backslash Z$ there exists

- $R > 0$,
- $\Theta < \infty$,
- $\theta < \infty$,
- m^2- measurable maps $\Phi_t : X \times X \to X$ for all $t \in [0, 1]$,

such that:

1. For $x, y \in Y$ with $d(x, y) < R$ and all s, t, the maps Φ_t have the following properties:

$$\Phi_0(x, y) = x, \tag{5.9}$$

$$\Phi_t(x, y) = \Phi_{1-t}(y, x), \tag{5.10}$$

$$\Phi_s(x, \Phi_t(x, y)) = \Phi_{st}(x, y), \tag{5.11}$$

$$d(\Phi_s(x, y), \Phi_t(x, y)) \leqslant \theta|s - t|d(x, y), \tag{5.12}$$

2. For all $r < R$, $x \in Y$, all measurable sets $A \subset B_r(x) \cap Y$ and all $t \in [0, 1]$

$$\frac{m_r(A)}{\sqrt{m(B_r(x))}} \leqslant \Theta \frac{m_{rt}(\Phi_t(x, A))}{\sqrt{m(B_{rt}(x))}}.$$

Equivalently this can be replaced by the condition

$$\frac{m(A))}{m(B_r(x))} \leqslant \Theta \frac{m(\Phi_t(x, A))}{m(B_{rt}(x))}. \tag{5.13}$$

Remark 5.3.10. If (X, d) is a geodesic space then conditions (5.9) - (5.12) are always satisfied with $\theta = 1$.

Remark 5.3.11. Sturm proved in Proposition 4.7 of [Stu98] that the measure contraction property without any exceptional set is true for Riemannian manifolds equipped with the Riemannian distance and the Riemannian volume.

Lemma 5.3.12. *Compact subanalytic sets equipped with the geodesic distance and the volume form defined in Defintion 5.1.10 satisfy the weak measure contraction property with an exceptional set equal to the singular stratum* Σ.

Proof. The volume form on subanalytic spaces is defined in such a way that the volume of the singular stratum is 0 and that the restriction of the volume form to each stratum coincides with the Riemannian volume. Then, the statement follows by Remark 5.3.11. □

However, we are interested in the case where the exceptional set is the empty set. This is because a non-trivial exceptional set results in a diffusion process whose sample paths do not cross the exceptional set. So, Lemma 5.3.12 implies that there is diffusion in each of the regular strata but it does not cross the singularities. We are interested in the situation in which heat diffuses through the singularities.

Example 5.3.13 (Failure of MCP for subanalytic sets)**.** Consider the stratified space consisting of two solid triangles glued together at the origin, as shown in Figure 5.4. If we consider the set A as depicted in the figure, then the set of geodesics at some time t

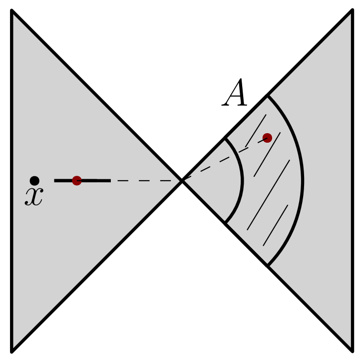

FIGURE 5.4: A demonstration of the failure of the measure contraction property without exceptional set.

with endpoints in A and starting at x is denoted by $\phi_t(x, A)$. There exists some value of $t > 0$ for which the volume of $\phi_t(x, A)$ is 0, whereas the volume of A is non-zero, thereby violating equation (5.13).

We conjecture that the MCP condition is true without exceptional set if the singular stratum of compact subanalytic sets has codimension at most 1. In particular, this would mean for spaces such as spiders and open books in Example 2.1.5, the pointwise limit of the Dirichlet forms coincides with the variational limit and gives rise to diffusion. In this case, one can explicitly write out the generator for the diffusion process (refer to [KS05]).

6. Conclusion

To summarise, we develop some preliminary methods to enable statistical analysis on stratified spaces. As mentioned before, stratified spaces possess singularities where fundamental geometry breaks down. This makes analysis on such spaces challenging and necessitates the use of novel methods to understand the structure of such spaces. We have outlined initial approaches to two key questions.

1. How to compare such spaces with one another?

2. How do we develop statistical models on stratified spaces?

Specifically, we introduced a general construction of shape space. The shapes are tame stratified spaces, more accurately described as constructible sets with respect to an o-minimal structure. The shape descriptors are topological analogues of the Radon transform. Much of this theory can be understood as a theory of signal processing on shapes, where the signals are shapes. We have proven that the assignment of each shape to its descriptor is a sheaf on the poset of shapes, ordered by inclusions. Additionally, we proposed metrics on shape spaces and demonstrated that these metrics are stable with respect to ϵ-perturbations.

We also provide a starting point for understanding diffusion processes on stratified spaces. Specifically, we provide a method to construct diffusion on compact subanalytic spaces with open and dense top-dimensional stratum via a sequence of Dirichlet forms. However, there is still much to understand. For example, what is the associated jump process when the top-dimensional stratum is not dense in the space? In what cases is the pointwise limit of the sequence of Dirichlet forms the same as the gamma limit? How does the diffusion process interact with the singularities? How can we make this theory more amenable for applications? Recently, a Central Limit Theorem was proved on stratified spaces (more specifically CAT(κ)-spaces) as evidenced in [MMT24; MMT23a; MMT23b; MMT23c]. Currently, I am working (in collaboration with Ezra Miller) on defining Brownian motion on such spaces, using approaches similar to those in the referenced papers.

Appendix A. Markov processes, Semigroups and Generators

For proofs see [EK09].

Let (X, m) be a compact measure space. Let H be a Hilbert space. For us the Hilbert space H is $L^2(X, m)$. The assumption that X is compact is not required and with some modification, the theory works for locally compact spaces, however for simplicity, we just state it for the compact case. Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a filtered probability space and let $\{X_t\}$ be a stochastic process on Ω adapted to $\mathcal{F}$ taking values in X.

Definition A.0.1 (Markov processes)**.** A stochastic process $\{X_t\}$ is a Markov process if for $s < t$ in $[0, \infty)$ and bounded measurable $f : X \to \mathbb{R}$, $\mathbb{E}\left[f(X_t) \mid \mathcal{F}_s\right] = \mathbb{E}\left[f(X_t) \mid X_s\right]$.

For a Markov process X_t, $p_t(x, B) = \mathbb{P}_x(X_t \in B)$, where B is a measurable set, is called the transition function or transition density of X.

Intuitively, a Markov process is symmetric with respect to m if the distribution of the process started in m is the same when moving forward or backward in time.

Definition A.0.2 (m-symmetric Markov process)**.** A Markov process $\{X_t\}$ on X is m-symmetric if $(X_s)_{0 \leqslant s \leqslant t} \sim (X_{t-s})_{0 \leqslant s \leqslant t}$ under $\mathbb{P}_m$, where $\mathbb{P}_m$ denotes the law of the process started from initial distribution m.

Markov processes give rise to semigroups, known as Markov semigroups.

Definition A.0.3 (Strongly continuous semigroups)**.** A collection of linear operators $\{T_t; t > 0\}$ on H is a strongly continuous semigroup if:

1. Each T_t is a symmetric operator with domain H,
2. (Semigroup) $T_s(T_t f) = T_{s+t} f$ for all $s, t > 0$ and $f \in H$,
3. (Strongly continuous, i.e. right continuous at 0) $\langle T_t f - f, T_t f - f\rangle \to 0$ as $t \to 0$ from above for all $f \in H$,
4. (Contraction) $\langle T_t f, T_t f\rangle \leqslant \|f\|$ for all $f \in H, t > 0$.

Note that this definition requires that the semigroup operators are right continuous at 0, but using the semigroup property one can show that these operators are continuous everywhere.

There is a one-one correspondence between abstract semigroups and symmetric Markov processes and one can confirm that Markov processes are related to semigroups in the following way

$$T_t f(x) = \mathbb{E}_x[f(X_t)].$$

One can also say that a Markov process with semigroup $\{T_t\}$ is m-symmetric if

$$\int_X f(x)T_t(g(x))m(dx) = \int_X T_t(f(x))g(x)m(dx),$$

for f, g in $L^2(X, m)$.

For a construction of a Markov process from a strongly continuous semigroup refer to [EK09].

Next, we define a generator of a semigroup. The generator of the semigroup is the derivative of the semigroup at $t = 0$, and so it describes the behavior of the semigroup in an infinitesimal time interval.

Definition A.0.4 (Generator of a semigroup)**.** The generator $\mathcal{L}$ of semigroup $\{T_t\}$ is a linear operator on H with domain $\mathcal{D}(\mathcal{L})$ defined as follows

$$\mathcal{D}(\mathcal{L}) = \left\{ f \in H \;\middle|\; \lim_{t\to 0} \frac{T_t f - f}{t} \text{ exists in } H \right\}; \qquad \mathcal{L}(f) = \lim_{t\to 0} \frac{T_t f - f}{t}.$$

One can check that the generator of strongly continuous semigroup is a non-positive definite self-adjoint operator. The converse of this statement is also true, for every such operator $\mathcal{L}$, we have a semigroup associated to it given by

$$T_t f = \lim_{n\to\infty} \left(I - \frac{t}{n}\mathcal{L} \right)^{-n} f.$$

So the generator defines the semigroup and the semigroup determines the process. This gives a three-way correspondence between Markov processes, semigroups and generators.

Bibliography

[ACM23] Shreya Arya, Justin Curry, and Sayan Mukherjee. “A Sheaf-Theoretic Construction of Shape Space”. In: *Foundations of Computational Mathematics* (Accepted on September 4, 2023).

[Ale18] Semyon Alesker. *Introduction to the Theory of Valuations*. Vol. 126. American Mathematical Soc., 2018.

[Ale35] JW Alexander. “On the chains of a complex and their duals”. In: *Proceedings of the National Academy of Sciences* 21.8 (1935), pp. 509–511.

[Amé+22] Erik J Amézquita et al. “Measuring hidden phenotype: Quantifying the shape of barley seeds using the Euler Characteristic Transform”. In: *in silico Plants* 4.1 (2022), diab033.

[AO13] Dean C Adams and Erik Otárola-Castillo. “geomorph: an R package for the collection and analysis of geometric morphometric shape data”. In: *Methods in ecology and evolution* 4.4 (2013), pp. 393–399.

[Arg+15] Sylvain Arguillere et al. “Shape deformation analysis from the optimal control viewpoint”. In: *Journal de mathématiques pures et appliquées* 104.1 (2015), pp. 139–178.

[Art62] Michael Artin. *Grothendieck topologies: notes on a seminar*. Harvard University, Department of Mathematics, 1962.

[BBM14] Martin Bauer, Martins Bruveris, and Peter W Michor. “Overview of the geometries of shape spaces and diffeomorphism groups”. In: *Journal of Mathematical Imaging and Vision* 50 (2014), pp. 60–97.

[Bec+03] J. Becker et al. “Complex dewetting scenarios captured by thin-film models”. In: *Nat. Mater.* 2.1 (Jan. 2003), pp. 59–63.

[Beg+05] M Faisal Beg et al. “Computing large deformation metric mappings via geodesic flows of diffeomorphisms”. In: *International journal of computer vision* 61 (2005), pp. 139–157.

[Ben+16] Paul Bendich et al. “Persistent homology analysis of brain artery trees”. In: *The annals of applied statistics* 10.1 (2016), p. 198.

[BG10] Yuliy Baryshnikov and Robert Ghrist. “Euler integration over definable functions”. In: *Proceedings of the National Academy of Sciences* 107.21 (May 2010), pp. 9525–9530. ISSN: 1091-6490. DOI: 10.1073/pnas.0910927107. URL: http://dx.doi.org/10.1073/pnas.0910927107.

[BHV01] Louis J Billera, Susan P Holmes, and Karen Vogtmann. “Geometry of the space of phylogenetic trees”. In: *Advances in Applied Mathematics* 27.4 (2001), pp. 733–767.

[BK01] Michael Brin and Yuri Kifer. “Brownian motion, harmonic functions and hyperbolicity for Euclidean complexes”. In: *Mathematische Zeitschrift* 237 (2001), pp. 421–468.

[BL17] Rabi Bhattacharya and Lizhen Lin. “Omnibus CLTs for Fréchet means and nonparametric inference on non-Euclidean spaces”. In: *Proceedings of the American Mathematical Society* 145.1 (2017), pp. 413–428.

[BM88] Edward Bierstone and Pierre D. Milman. “Semianalytic and subanalytic sets”. en. In: *Publications Mathématiques de l'IHÉS* 67 (1988), pp. 5–42. URL: http://www.numdam.org/item/PMIHES_1988__67__5_0/.

[Bon+07] D. Bonn et al. “Wetting and spreading”. Submitted to *Rev. Mod. Phys.* 2007. URL: http://itf.fys.kuleuven.be/~joi/papers/Wetting%20and%20spreading.pdf.

[Boo92] Fred L. Bookstein. *Morphometric Tools for Landmark Data: Geometry and Biology*. Cambridge University Press, 1992. DOI: 10.1017/CBO9780511573064.

[Boy+11] Doug M Boyer et al. “Algorithms to automatically quantify the geometric similarity of anatomical surfaces”. In: *Proceedings of the National Academy of Sciences* 108.45 (2011), pp. 18221–18226.

[Boy+15] Doug M Boyer et al. “A new fully automated approach for aligning and comparing shapes”. In: *The Anatomical Record* 298.1 (2015), pp. 249–276.

[Boy+16] Doug M Boyer et al. “Morphosource: archiving and sharing 3-D digital specimen data”. In: *The Paleontological Society Papers* 22 (2016), pp. 157–181.

[BP03a] Rabi Bhattacharya and Vic Patrangenaru. “Large sample theory of intrinsic and extrinsic sample means on manifolds”. In: *The Annals of Statistics* 31.1 (2003), pp. 1–29. DOI: 10.1214/aos/1046294456. URL: https://doi.org/10.1214/aos/1046294456.

[BP03b] Rabi Bhattacharya and Vic Patrangenaru. “Large sample theory of intrinsic and extrinsic sample means on manifolds. I”. In: *Ann. Statist.* 31.1 (2003), pp. 1–29. ISSN: 0090-5364. DOI: 10.1214/aos/1046294456. URL: https://doi.org/10.1214/aos/1046294456.

[Bre12] Glen E Bredon. *Sheaf theory*. Vol. 170. Springer Science & Business Media, 2012.

[Bro+10] Alexander M Bronstein et al. "A Gromov-Hausdorff framework with diffusion geometry for topologically-robust non-rigid shape matching". In: *International Journal of Computer Vision* 89.2-3 (2010), pp. 266–286.

[CD02] H.-C. Chang and E.A. Demekhin. "Complex Wave Dynamics on Thin Films". In: Elsevier, 2002. Chap. 9, pp. 271–285.

[CGR12] Justin Curry, Robert Ghrist, and Michael Robinson. "Euler calculus with applications to signals and sensing". In: *Proceedings of Symposia in Applied Mathematics*. Vol. 70. 2012, pp. 75–146.

[Cha+09] Frédéric Chazal et al. "Proximity of persistence modules and their diagrams". In: *Proceedings of the twenty-fifth annual symposium on Computational geometry*. 2009, pp. 237–246.

[Cha+16] Frédéric Chazal et al. *The structure and stability of persistence modules*. Vol. 10. Springer, 2016.

[Cha06] Isaac Chavel. *Riemannian geometry: a modern introduction*. Vol. 98. Cambridge university press, 2006.

[CMM91] Vicente Cervera, Francisca Mascaro, and Peter W Michor. "The action of the diffeomorphism group on the space of immersions". In: *Differential Geometry and its Applications* 1.4 (1991), pp. 391–401.

[CMT22a] Justin Curry, Sayan Mukherjee, and Katharine Turner. "How many directions determine a shape and other sufficiency results for two topological transforms". In: *Transactions of the American Mathematical Society, Series B* 9.32 (2022), pp. 1006–1043.

[CMT22b] Justin Curry, Sayan Mukherjee, and Katharine Turner. "How many directions determine a shape and other sufficiency results for two topological transforms". In: *Transactions of the American Mathematical Society, Series B* 9.32 (2022), pp. 1006–1043.

[Col99] Stëphen Colbét. "Name with proper names Hello Kitty". In: *Imag. Res.* 1 (1999), pp. 123–127.

[Cor63] Allan Macleod Cormack. "Representation of a function by its line integrals, with some radiological applications". In: *Journal of applied physics* 34.9 (1963), pp. 2722–2727.

[Cor64] Allan Macleod Cormack. "Representation of a function by its line integrals, with some radiological applications. II". In: *Journal of Applied Physics* 35.10 (1964), pp. 2908–2913.

[Cra+20] Lorin Crawford et al. "Predicting clinical outcomes in glioblastoma: an application of topological and functional data analysis". In: *Journal of the American Statistical Association* 115.531 (2020), pp. 1139–1150.

[Cra15] William Crawley-Boevey. "Decomposition of pointwise finite-dimensional persistence modules". In: *Journal of Algebra and its Applications* 14.05 (2015), p. 1550066.

[CT13] Nicolas Charon and Alain Trouvé. "The varifold representation of nonoriented shapes for diffeomorphic registration". In: *SIAM journal on Imaging Sciences* 6.4 (2013), pp. 2547–2580.

[Cur+22] Justin Curry et al. "Decorated merge trees for persistent topology". In: *Journal of Applied and Computational Topology* 6.3 (2022), pp. 371–428.

[Cur14a] Justin Michael Curry. *Sheaves, cosheaves and applications*. University of Pennsylvania, 2014.

[Cur14b] Justin Michael Curry. *Sheaves, cosheaves and applications*. University of Pennsylvania, 2014.

[Dal12] Gianni Dal Maso. *An introduction to* Γ*-convergence*. Vol. 8. Springer Science & Business Media, 2012.

[De 75] Ennio De Giorgi. "Sulla convergenza di alcune successioni d'integrali del tipo dell'area". In: *Rend. Mat* 8.6 (1975), pp. 277–294.

[DF75] Ennio De Giorgi and Tullio Franzoni. "Su un tipo di convergenza variazionale". In: *Atti della Accademia Nazionale dei Lincei. Classe di Scienze Fisiche, Matematiche e Naturali. Rendiconti* 58.6 (1975), pp. 842–850.

[DGM98] Paul Dupuis, Ulf Grenander, and Michael I Miller. "Variational problems on flows of diffeomorphisms for image matching". In: *Quarterly of applied mathematics* (1998), pp. 587–600.

[DHI04] Daniel Dugger, Sharon Hollander, and Daniel C Isaksen. "Hypercovers and simplicial presheaves". In: *Mathematical Proceedings of the Cambridge Philosophical Society*. Vol. 136. 1. Cambridge University Press. 2004, pp. 9–51.

[DM98] Ian L Dryden and Kanti V Mardia. "Statistical shape analysis: Wiley series in probability and statistics". In: *New York, NY: John Wiley & Sons, Ltd* (1998).

[DMP16] Vin De Silva, Elizabeth Munch, and Amit Patel. "Categorified reeb graphs". In: *Discrete & Computational Geometry* 55.4 (2016), pp. 854–906.

[EC56] Samuel Eilenberg and Henri Paul Cartan. *Homological algebra*. Princeton University Press, 1956.

[Ede93] Herbert Edelsbrunner. "The union of balls and its dual shape". In: *Proceedings of the ninth annual symposium on Computational geometry*. 1993, pp. 218–231.

[EK09] Stewart N Ethier and Thomas G Kurtz. *Markov processes: characterization and convergence*. John Wiley & Sons, 2009.

[ELZ02] Herbert Edelsbrunner, David Letscher, and Afra Zomorodian. "Topological persistence and simplification". In: *Discrete & computational geometry* 28 (2002), pp. 511–533.

[FD84] Douglas H Frank and Stephen Durham. "Random motion on binary trees". In: *Journal of applied probability* 21.1 (1984), pp. 58–69.

[Fer86] Massimo Ferrarotti. "Volume on stratified sets". In: *Annali di Matematica Pura ed Applicata* 144 (1986), pp. 183–201.

[Fle+04] P Thomas Fletcher et al. "Principal geodesic analysis for the study of nonlinear statistics of shape". In: *IEEE transactions on medical imaging* 23.8 (2004), pp. 995–1005.

[FLJ03] P Thomas Fletcher, Conglin Lu, and Sarang Joshi. "Statistics of shape via principal geodesic analysis on Lie groups". In: *2003 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2003. Proceedings.* Vol. 1. IEEE. 2003, pp. I–I.

[FLM17] Patrizio Frosini, Claudia Landi, and Facundo Mémoli. "The persistent homotopy type distance". In: *arXiv preprint arXiv:1702.07893* (2017).

[FOT11] Masatoshi Fukushima, Yoichi Oshima, and Masayoshi Takeda. *Dirichlet forms and symmetric Markov processes*. extended. Vol. 19. De Gruyter Studies in Mathematics. Walter de Gruyter & Co., Berlin, 2011, pp. x+489. ISBN: 978-3-11-021808-4.

[Fow] Charless C Fowlkes. "Surveying Shape Spaces". In: *survey article for Alan Weinstein's course on Reimannian Geometry. https://www. ics. uci. edu/˜ fowlkes/papers/f-ma240. pdf* ().

[Fré48] Maurice Fréchet. "Les éléments aléatoires de nature quelconque dans un espace distancié". In: *Annales de l'institut Henri Poincaré*. Vol. 10. 4. 1948, pp. 215–310.

[GBQ03] P.-G. de Gennes, F. Brochard-Wyart, and D. Quere. *Capillarity and wetting phenomena: drops, bubbles, pearls, waves.* New York: Springer Verlag, 2003.

[GKD19] Tingran Gao, Shahar Z Kovalsky, and Ingrid Daubechies. "Gaussian process landmarking on manifolds". In: *SIAM Journal on Mathematics of Data Science* 1.1 (2019), pp. 208–236.

[GLM18] Robert Ghrist, Rachel Levanger, and Huy Mai. "Persistent homology and Euler integral transforms". In: *Journal of Applied and Computational Topology* 2 (2018), pp. 55–60.

[GM98] Ulf Grenander and Michael I Miller. "Computational anatomy: An emerging discipline". In: *Quarterly of applied mathematics* 56.4 (1998), pp. 617–694.

[God58a] G Godement. "Topologies algébrique et théorie des faisceaux". In: *Actualites Scientifiques et Industrielles 1252* (1958).

[God58b] Roger Godement. "Topologie algébrique et théorie des faisceaux". In: *Actualites Scientifiques et Industrielles 1252* (1958).

[Gos15] A Goswami. "Phenome10K: a free online repository for 3-D scans of biological and palaeontological specimens". In: *www.phenome10k.org* (2015).

[Gra06] John W Gray. "Fragments of the history of sheaf theory". In: *Applications of Sheaves: Proceedings of the Research Symposium on Applications of Sheaf Theory to Logic, Algebra, and Analysis, Durham, July 9–21, 1977.* Springer. 2006, pp. 1–79.

[Gre96] Ulf Grenander. *Elements of pattern theory.* JHU Press, 1996.

[Gro57] Alexandre Grothendieck. "Sur quelques points d'algèbre homologique". In: *Tohoku Mathematical Journal, Second Series* 9.2 (1957), pp. 119–183.

[Gro97] Alexandre Grothendieck. "Esquisse d'un programme". In: *London Mathematical Society Lecture Note Series* (1997), pp. 5–48.

[Hal04] Demetrios J Halazonetis. "Morphometrics for cephalometric diagnosis". In: *American Journal of Orthodontics and Dentofacial Orthopedics* 125.5 (2004), pp. 571–581.

[Hat01] Allen Hatcher. *Algebraic Topology.* Cambridge University Press, 2001.

[Hir73] Heisuke Hironaka. "Subanalytic sets, Number theory". In: *Algebraic Geometry and Commutative Algebra* (1973), pp. 453–493.

[Hör17] Fritz Hörmann. “Fibered multiderivatories and (co)homological descent”. In: *Theory and Applications of Categories* 32.38 (2017), pp. 1258–1362. URL: http://www.tac.mta.ca/tac/volumes/32/38/32-38.pdf.

[Hör19] Fritz Hörmann. *The ∞-categorical interpretation of Abelian and non-Abelian derived functors*. Preliminary notes. 2019. URL: https://fhoermann.org/infty.pdf.

[Hsu02] Elton P Hsu. *Stochastic analysis on manifolds*. 38. American Mathematical Soc., 2002.

[Hsu87] Pei Hsu. “Brownian motion and the Atiyah-Singer index theorem”. In: *Preprint (present address: University of Illinois at Chicago)* (1987).

[Ive12] Birger Iversen. *Cohomology of sheaves*. Springer Science & Business Media, 2012.

[Ive86] Birger Iversen. *Cohomology of sheaves*. Universitext. Springer-Verlag, Berlin, 1986, pp. xii+464. ISBN: 3-540-16389-1. DOI: 10.1007/978-3-642-82783-9. URL: https://doi.org/10.1007/978-3-642-82783-9.

[JM00] Sarang C Joshi and Michael I Miller. “Landmark matching via large deformation diffeomorphisms”. In: *IEEE transactions on image processing* 9.8 (2000), pp. 1357–1370.

[Jos16] Jürgen Jost, ed. *Bernhard Riemann: On the Hypotheses which Lie at the Bases of Geometry*. Basel: Birkhäuser, 2016.

[Jos17] Jürgen Jost. “Relations and dependencies between morphological characters”. In: *Theory Biosci.* 136.1 (2017), pp. 69–83. URL: https://doi.org/10.1007/s12064-017-0248-z.

[Ken77] David G Kendall. “The diffusion of shape”. In: *Advances in applied probability* 9.3 (1977), pp. 428–430.

[Ken84] David G Kendall. “Shape manifolds, procrustean metrics, and complex projective spaces”. In: *Bulletin of the London mathematical society* 16.2 (1984), pp. 81–121.

[Kla04] Daniel A. Klain. “The Minkowski problem for polytopes”. In: *Advances in Mathematics* 185.2 (2004), pp. 270–288.

[KO97] Krzysztof Kurdyka and Patrice Orro. “Distance géodésique sur un sous-analytique”. In: *Rev. Mat. Univ. Complut. Madrid* 10.Special Issue, suppl. (1997), pp. 173–182.

[KS05] Takashi Kumagai and Karl-Theodor Sturm. “Construction of diffusion processes on fractals, d-sets, and general metric measure spaces”. In: *Journal of Mathematics of Kyoto University* 45.2 (2005), pp. 307–327.

[KS18] Masaki Kashiwara and Pierre Schapira. “Persistent homology and microlocal sheaf theory”. In: *Journal of Applied and Computational Topology* 2.1-2 (2018), pp. 83–113.

[KS94] Masaki Kashiwara and Pierre Schapira. *Sheaves on manifolds*. Vol. 292. Grundlehren der mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]. With a chapter in French by Christian Houzel, Corrected reprint of the 1990 original. Springer-Verlag, Berlin, 1994, pp. x+512. ISBN: 3-540-51861-4.

[Lê 98] Ta Lê Loi. “Verdier and strict Thom stratifications in o-minimal structures”. In: *Illinois Journal of Mathematics* 42.2 (1998), pp. 347–356.

[Ler45a] Jean Leray. “Sur la forme des espaces topologiques et sur les points fixes des représentations: Première partie d’un cours de topologie algébrique professé en captivité”. In: *Journal de Mathématiques Pures et Appliquées* 9 (1945), pp. 95–167.

[Ler45b] Jean Leray. “Sur la position d’un ensemble fermé de points d’un espace topologique: Deuxième partie d’un cours de topologie algébrique professé en captivité”. In: *Journal de Mathématiques Pures et Appliquées* 24 (1945), pp. 169–199.

[Ler45c] Jean Leray. “Sur les équations et les transformations: Troisième partie d’un cours de topologie algébrique professé en captivité”. In: *Journal de Mathématiques Pures et Appliquées* 24 (1945), pp. 200–248.

[Les15] Michael Lesnick. “The theory of the interleaving distance on multidimensional persistence modules”. In: *Foundations of Computational Mathematics* 15.3 (2015), pp. 613–650.

[Lur17] Jacob Lurie. *Higher algebra*. 2017. URL: `https://www.math.ias.edu/~lurie/papers/HA.pdf`.

[LY86] Peter Li and Shing-Tung Yau. “On the parabolic kernel of the Schrödinger operator”. In: *Acta Math.* 156.3-4 (1986), pp. 153–201. ISSN: 0001-5962,1871-2509. DOI: `10.1007/BF02399203`. URL: `https://doi.org/10.1007/BF02399203`.

[MA14] J. Steve Marron and Andrés M. Alonso. “Overview of object oriented data analysis”. In: *Biom. J.* 56.5 (2014), pp. 732–753. ISSN: 0323-3847,1521-4036. DOI: `10.1002/bimj.201300072`. URL: `https://doi.org/10.1002/bimj.201300072`.

[MD21] James Stephen Marron and Ian L Dryden. *Object oriented data analysis*. CRC Press, 2021.

[Mem07] Facundo Memoli. “On the use of Gromov-Hausdorff Distances for Shape Comparison”. In: *Eurographics Symposium on Point-Based Graphics*. Ed. by M. Botsch et al. The Eurographics Association, 2007. ISBN: 978-3-905673-51-7. DOI: 10.2312/SPBG/SPBG07/081-090.

[Mic15] Peter Michor. *Manifolds of mappings and shapes*. 2015. eprint: arXiv:1505.02359.

[Mil00] Haynes Miller. “Leray in Oflag XVIIA: the origins of sheaf theory, sheaf cohomology, and spectral sequences”. In: *Kantor 2000*. 2000, pp. 17–34. URL: https://math.mit.edu/~hrm/papers/ss.pdf.

[Mil15] Ezra Miller. “Fruit flies and moduli: interactions between biology and mathematics”. In: *Notices of the American Mathematical Society* 62.10 (2015), pp. 1178–1184.

[Mil23] Ezra Miller. personal communication. 2023.

[MM05] Peter Michor and David Mumford. “Vanishing geodesic distance on spaces of submanifolds and diffeomorphisms”. In: *DOCUMENTA MATHEMATICA* 10 (2005), pp. 217–245.

[MM06] David Mumford and Peter Michor. “Riemannian geometries on spaces of plane curves”. In: *Journal of the European Mathematical Society* 8.1 (2006), pp. 1–48.

[MMT23a] Jonathan C. Mattingly, Ezra Miller, and Do Tran. *Central limit theorems for Fréchet means on stratified spaces*. 2023. arXiv: 2311.09455 [math.PR].

[MMT23b] Jonathan C. Mattingly, Ezra Miller, and Do Tran. *Geometry of measures on smoothly stratified metric spaces*. 2023. arXiv: 2311.09453 [math.MG].

[MMT23c] Jonathan C. Mattingly, Ezra Miller, and Do Tran. *Shadow geometry at singular points of CAT(k) spaces*. 2023. arXiv: 2311.09451 [math.MG].

[MMT24] Jonathan C. Mattingly, Ezra Miller, and Do Tran. *A central limit theorem for random tangent fields on stratified spaces*. 2024. arXiv: 2311.09454 [math.PR].

[Mor08] Dmitriy Morozov. *Homological illusions of persistence and stability*. Duke University, 2008.

[MS22] Philipp Mitteroecker and Katrin Schaefer. "Thirty years of geometric morphometrics: Achievements, challenges, and the ongoing quest for biological meaningfulness". In: *American journal of biological anthropology* 178 (2022), pp. 181–210.

[MSJ07] Washington Mio, Anuj Srivastava, and Shantanu Joshi. "On shape of plane elastic curves". In: *International Journal of Computer Vision* 73 (2007), pp. 307–324.

[MTY15] Michael I Miller, Alain Trouvé, and Laurent Younes. "Hamiltonian systems and optimal control in computational anatomy: 100 years since D'Arcy Thompson". In: *Annual review of biomedical engineering* 17 (2015), pp. 447–509.

[Mum94] David Mumford. "Pattern theory: a unifying perspective". In: *First European Congress of Mathematics: Paris, July 6-10, 1992 Volume I Invited Lectures (Part 1)*. Springer. 1994, pp. 187–224.

[nLa24] nLab authors. *concrete category*. `https://ncatlab.org/nlab/show/concrete+category`. Revision 40. Feb. 2024.

[NSW08] Partha Niyogi, Stephen Smale, and Shmuel Weinberger. "Finding the homology of submanifolds with high confidence from random samples". In: *Discrete & Computational Geometry* 39 (2008), pp. 419–441.

[Nye20] Tom MW Nye. "Random walks and Brownian motion on cubical complexes". In: *Stochastic Processes and their Applications* 130.4 (2020), pp. 2185–2199.

[Ovs+12] Maks Ovsjanikov et al. "Functional maps: a flexible representation of maps between shapes". In: *ACM Transactions on Graphics (ToG)* 31.4 (2012), pp. 1–11.

[Pat18] Amit Patel. "Generalized persistence diagrams". In: *Journal of Applied and Computational Topology* 1.3-4 (2018), pp. 397–419.

[Pfl01] Markus Pflaum. *Analytic and geometric study of stratified spaces: contributions to analytic and geometric aspects*. 1768. Springer Science & Business Media, 2001.

[Rad05] Johann Radon. "1.1 über die bestimmung von funktionen durch ihre integralwerte längs gewisser mannigfaltigkeiten". In: *Classic papers in modern diagnostic radiology* 5.21 (2005), p. 124.

[Rie17] E. Riehl. *Category Theory in Context*. Aurora: Dover Modern Math Originals. Dover Publications, 2017. ISBN: 9780486820804. URL: `https://books.google.com/books?id=6B9MDgAAQBAJ`.

[Rie67] Bernhard Riemann. *Ueber die Hypothesen, welche der Geometrie zu Grunde liegen*. Vol. 13. Dietrich, 1867.

[Sch91a] Pierre Schapira. "Operations on constructible functions". In: *Journal of pure and applied algebra* 72.1 (1991), pp. 83–93.

[Sch91b] Pierre Schapira. "Operations on constructible functions". In: *Journal of pure and applied algebra* 72.1 (1991), pp. 83–93.

[Sch95a] Pierre Schapira. "Tomography of constructible functions". In: *International Symposium on Applied Algebra, Algebraic Algorithms, and Error-Correcting Codes*. Springer. 1995, pp. 427–435.

[Sch95b] Pierre Schapira. "Tomography of constructible functions". In: *International Symposium on Applied Algebra, Algebraic Algorithms, and Error-Correcting Codes*. Springer. 1995, pp. 427–435.

[She85] Allen Dudley Shepard. *A cellular description of the derived category of a stratified space*. Brown University, 1985.

[SK16] Anuj Srivastava and Eric P Klassen. *Functional and shape data analysis*. Vol. 1. Springer, 2016.

[SM06] Eitan Sharon and David Mumford. "2d-shape analysis using conformal mapping". In: *International Journal of Computer Vision* 70.1 (2006), pp. 55–75.

[SMS18] Vin de Silva, Elizabeth Munch, and Anastasios Stefanou. "Theory of interleavings on categories with a flow". In: *Theory and Applications of Categories* 33.21 (2018), pp. 583–607.

[Sri+10] Anuj Srivastava et al. "Shape analysis of elastic curves in euclidean spaces". In: *IEEE transactions on pattern analysis and machine intelligence* 33.7 (2010), pp. 1415–1428.

[ST20] Primoz Skraba and Katharine Turner. "Wasserstein stability for persistence diagrams". In: *arXiv preprint arXiv:2006.16824* (2020).

[Ste18] Anastasios Stefanou. *Dynamics on categories and applications*. State University of New York at Albany, 2018.

[Stu98] K. T. Sturm. "Diffusion processes and heat kernels on metric spaces". In: *The Annals of Probability* 26.1 (1998), pp. 1–55. DOI: 10.1214/aop/1022855410. URL: https://doi.org/10.1214/aop/1022855410.

[Tan+22] Wai Shing Tang et al. “A topological data analytic approach for discovering biophysical signatures in protein dynamics”. In: *PLoS computational biology* 18.5 (2022), e1010045.

[TMB14] Katharine Turner, Sayan Mukherjee, and Doug M Boyer. “Persistent homology transform for modeling shapes and surfaces”. In: *Information and Inference: A Journal of the IMA* 3.4 (2014), pp. 310–344.

[Van98a] Lou Van Den Dries. “O-minimal structures and real analytic geometry”. In: *Current developments in mathematics* 1998.1 (1998), pp. 105–152.

[Van98b] Lou Van den Dries. *Tame topology and o-minimal structures*. Vol. 248. Cambridge university press, 1998.

[Van98c] Lou Van den Dries. *Tame topology and o-minimal structures*. Vol. 248. Cambridge university press, 1998.

[Ver77] Jean-Louis Verdier. *Caégories dérivées, état 0, SGA* $4\frac{1}{2}$. Vol. 569. Lecture Notes in Mathematics. Springer Verlag, 1977.

[Ver96] Jean-Louis Verdier. “Des catégories dérivées des catégories abéliennes”. In: (1996).

[VG05] Marc Vaillant and Joan Glaunes. “Surface matching via currents”. In: *Biennial international conference on information processing in medical imaging*. Springer. 2005, pp. 381–392.

[Vir06] Oleg Yanovich Viro. “Some integral calculus based on Euler characteristic”. In: *Topology and geometry—Rohlin seminar*. Springer. 2006, pp. 127–138.

[Wan+07] Sen Wang et al. “Conformal geometry and its applications on 3D shape matching, recognition, and stitching”. In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* 29.7 (2007), pp. 1209–1220.

[Wan+21] Bruce Wang et al. “A statistical pipeline for identifying physical features that differentiate classes of 3D shapes”. In: *The Annals of Applied Statistics* 15.2 (2021), pp. 638–661.

[Wei99] Charles A Weibel. *History of homological algebra*. 1999.

[You98] Laurent Younes. “Computable elastic distances between shapes”. In: *SIAM Journal on Applied Mathematics* 58.2 (1998), pp. 565–586.