Sheaves, Cosheaves and Applications

Justin M. Curry

March 13, 2013

Abstract

This note advertises the theory of cellular sheaves and cosheaves, which are devices for conducting linear algebra parametrized by a cell complex. The theory is presented in a way that is meant to be read and appreciated by a broad audience, including those who hope to use the theory in applications across science and engineering disciplines. We relay two approaches to cellular cosheaves. One relies on heavy stratification theory and MacPherson's entrance path category. The other uses the Alexandrov topology on posets. We develop applications to persistent homology, network coding, and sensor networks to illustrate the utility of the theory. The driving computational force is cellular cosheaf homology and sheaf cohomology. However, to interpret this computational theory, we make use of the Remak decomposition into indecomposable representations of the cell category. The computational formula for cellular cosheaf homology is put on the firm ground of derived categories. This leads to an internal development of the derived perspective for cell complexes. We prove a conjecture of MacPherson that says cellular sheaves and cosheaves are derived equivalent. Although it turns out to be an old result, our proof is more explicit than other proofs and we make clear that Poincaré-Verdier duality should be viewed as an exchange of sheaves and cosheaves. The existence of enough projectives allows us to define a new homology theory for sheaves on posets and we establish some classical duality results in this setting. Finally, we make use of coends as a generalized tensor product to phrase compactly supported sheaf cohomology as the pairing with the image of the constant sheaf through the derived equivalence.

Contents

1 Introduction

The utility of linear algebra in mathematics, science and engineering is manifest. Linear algebra exhibits a surprising fidelity for applications with discrete structures possessing continuous qualities. When wedded with geometry, it provides true local approximations to highly complex phenomena. The linear reflects nonlinear properties that can be effectively computed and, today, put on a computer. The embedding of the continuous world into the discrete matrix of the computable world is a remarkable force of technological innovation.

Classical theorems of 17th, 18th and 19th century mathematics describing the flow and flux of heat, fluid and abstract electromagnetic fields were framed continuously via The Calculus. This form usually came after first thinking in terms of discrete particles continuously moving through small boxes. The passage from discrete to continuous and again to discrete – or model to theory to applications – by programming these renowned equations on computers, represents a cycle of innovation of unparalleled productivity. The theorist and aesthete may lament the degradation of her ideal models, but the engineer and pragmatic judiciously apprehends the differences and the similarities.

To say that for algebraic topologists the situation described is most familiar is an understatement more than mild. The whole enterprise of algebraic topology consists of the principled study of pushing the continuous through the mesh of algebra. However, this perspective betrays the history of the subject, since up until a hundred years ago, the prototype of this study was called "combinatorial topology." Without algebra, numbers were the receptacle for discretization. To perform an alternating count of vertices, edges and faces and discover that for all the Platonic solids, the result is two, would usher in history's first topological invariant in the 1750s – the Euler characteristic.

The complexity of the objects studied by topology rose, the linear giving way to the subtle twists of holomorphic maps and geometry, the grounding in discrete computable invariants continued. It was Gauss – that father figure of so much mathematics – who then understood that averaging the curvature of a surface yielded Euler characteristic. One of the many disciples of Gauss, Johann Benedict Listing, coined the very word "topology" in 1836 and in attempting to define its subtle nature said:

> "By topology we mean the doctrine of the modal features of objects, or of the laws of connection, of relative position and of succession of points, lines, surfaces, bodies and their parts, or aggregates in space, always without regard to matters of measure or quantity." [50]

Listing, together with Gauss, would influence that next great hero-figure of mathematics – Bernhard Riemann – by communicating some notions of topology. Alas, that coarse art of counting required serious strengthening in order to organize the inchoate world of high-dimensional manifolds that Riemann summoned from the heavens. Riemann, who lived a short life constantly in poor health, spent his last years in Italy where he returned a visit to Enrico Betti. Shortly after Riemann's death, Betti produced – without knowing it – the first lifting of Euler characteristic in 1870, by defining a sequence of numbers whose purpose was to count "holes" in a higher dimensional manifold. To do so, Betti introduced the notion of a boundary, thus anticipating the algebraization of combinatorial topology.

The person credited with inventing algebraic topology is Henri Poincaré. In addition to Riemann and Betti, he was influenced by his own investigations in differential equations, celestial mechanics and discontinuous group actions. In his famous treatise "On Analysis Situs" (the preferred name for topology before Lefschetz revived Listing's term), Poincaré tried to assuage an oddly prescient concern of modern times:

> "Persons who recoil from geometry of more than three dimensions may believe this result to be useless and view it as a futile game, if they have not been informed of their error by the use made of Betti numbers by our colleague M. Picard in pure analysis and ordinary geometry." [69]

Having cursorily addressed the criticism of a lack of applications of analysis situs, Poincaré went on to push Betti's notion of a boundary and defined what is now known as homology and the fundamental group. Poincaré then showed that the alternating sum of the Betti numbers yields Euler characteristic. This latter result – after Emmy Noether's formalization of Poincaré's homology in the language of groups – initiated the theme of **categorification** in mathematics.

Poincaré ushered in a theory most complex and sublime that blossomed into a major branch of modern-day mathematics. This branch landed on one side of a rapid bisection of mathematics into "pure" and "applied" parts in the second half of the 20th century. The division unnaturally separated topology, which has been called "the ultimate non-linear data analysis toolkit," from the aforementioned innovative cycle of theory and applications.[1]

In the 21st century an excavation and adaptation of the past 100 years of topology is currently underway. Each time the applicationist comes to another artifact – some advanced technology left behind by a great civilization – they ponder on what purpose it might have served or will serve. Having reached the bedrock which unites these cultures, the layers must be traveled up again.

The basest stratum has revealed the relevance of Poincaré's instruments and in our attempt to adapt these tools we have come across unexpected ways of using them. Consider, for example, the problem of describing the shape of a discrete set of points $\{x_i\}_{i=1}^n$ embedded inside some Euclidean space $\mathbb{R}^N$. The points, not having any geometry of their own, will make their relative positions known by considering the union of the neighborhoods of some radius r about each point, which can be studied as a space in its own right $X_r := \cup_{i=1}^n B(x_i; r) \subset \mathbb{R}^N$. In light of our Cartesian skepticism of the source, we do not put faith in measure and quantity, but rather we satisfy ourselves with the homology $H_*(X_r)$ for each value of the radius r and study those algebraic invariants which *persist* over the base parameter $r \in \mathbb{R}$. Alternatively said, we want to understand how the space $X = \{(X_r, r) \subset \mathbb{R}^{N+1}\}$ fibers over

[1] This quote has been attributed to Gunnar Carlsson.

$\mathbb{R}$ and extract the features of the following assignment of data:

$$\begin{tikzcd} \mathsf{X} \arrow[d, "\pi"'] & \mathsf{H}_n(\pi^{-1}(t)) \\ \mathbb{R} & t \arrow[u, squiggly] \end{tikzcd}$$

This is the exemplar par excellence of a **cellular cosheaf**, which is one of the main attractions of this paper.

Further archaeology reveals that it was Jean Leray who in 1946 said that we should adapt the study of topology from spaces to *representations* or maps [42]:

> "Nous nous proposons d'indiquer sommairement comment les méthodes par lesquelles nous avons etudié la topologie d'un espace peuvent être adaptées à l'étude de la topologie d'une représentation."[2]

Leray, whose own life reminds us of the schism of pure and applied mathematics, was an analyst who had used topology in his work on fluid dynamics. However, when he was captured by German soldiers he denied his competency as a *mécanicien* and sought refuge in the then un-applicable area of algebraic topology [64]. While in residence in a prisoner-of-war camp for officers, he spent five years devoted to generalizing homology to avoid the use of triangulations or smooth structure. His two developments – the **sheaf** and the **spectral sequence** – allowed him to study the topology of maps by assigning data to *closed* subsets of a space X. Upon Leray's release in 1945, these two powerful instruments were taken up by Henri Cartan and Jean-Pierre Serre, among others, where they were stripped down, modified and built up again. By the 1950s sheaves were viewed as a pair of spaces and a map $\pi : E \to X$ whose sections over *open* sets defined the data assigned there. These ideas were replaced, extended and generalized by Alexander Grothendieck, who is perhaps responsible for the largest expansion of the aegis of sheaf theory.

Standing deep in the mines of time we pick up these tools and ask of them "What is their use?" The generality of these ideas provides us with a powerful solvent in which we soak our problems; however, to extract out the desired elements we make use of a discrete mesh to crystallize our solutions. For us this mesh is offered by something that Leray deliberately wanted to avoid, the use of simplicial and cellular complexes, so we have had to wait for his trend of generality to be reversed.

It was only in the 1980's that a more explicit description of sheaves adapted to cell complexes came from the two independent sources of Bob MacPherson and Masaki Kashiwara. Although both were predated by the 1955 thesis work of Sir Christopher Zeeman, the most serious and concrete exposition of the theory of **cellular sheaves** was laid out in the 1984 unpublished thesis of Allen Shepard [80], who worked under MacPherson's direction. The theory has languished, despite the work of Maxim Vybornov, who developed more of the

[2] An attempted translation: "We propose to state briefly how the methods by which we have studied the topology of a space can be adapted to the study of the topology of maps."

theory of cellular sheaves and cosheaves in his 1999 thesis [92]. Perhaps in the sea of vast generality, this island of concreteness has been disregarded by the pure mathematical culture.

However, the technological transfer from algebraic topology to data analysis, sensor networks, and dynamical systems in recent years has led to an infusion of local-to-global ideas. In 2008, Robert Ghrist initiated a call to bring sheaf theory, specifically sheaf cohomology, to bear on a variety of applied problems. **Euler calculus** – a de-categorification of constructible sheaf theory – had already made inroads towards this goal [23]. Heuristically, sheaf cohomology would provide calculable summaries of the topology of data and programs, even if initially there was no topology in sight. The first attack was to model various systems as small categories and to put sheaves on them in the spirit of Grothendieck. However, Tony Pantev brought to our attention the work of Shepard, Vybornov, and Alexander Polishchuck, which allowed us to start developing applications more readily. If this fortuitous interaction had not occurred, we would have found ourselves in Joel Friedman's situation, who carried out the Grothendieck perspective in his work on the Hanna Neumann conjecture in graph theory [32].

In response to rising interest, Bob MacPherson organized an informal seminar at the Institute for Advanced Study to develop applied sheaf (and cosheaf) theory. During 2011-2012, MacPherson gave four lectures and the author gave six lectures where part of this document was first relayed.

1.1 Outline of the Document

This document is intended for a wide readership. As such, there is much more background material than one might expect, to help the neophyte through the text. Occasionally there are remarks and asides intended for the expert. We hope neither of these things discourages readers.

In section 2, we briefly review categories, along with the notion of a representation of a category. The main focus is on reviewing the notion of limits and colimits, which are integral to the abstract definition of sheaves and cosheaves, respectively.

Cosheaves, defined as covariant assignments of data to open sets, we treat in abstract form in section 3. This is intended to put the theory on firm foundations. Unlike with sheaves, one cannot just define the notion of a cosheaf of sets and then use that to define cosheaves of groups, vector spaces, and so on. Thus, we record the definition of a cosheaf valued in an arbitrary data category $\mathcal{D}$. Some serious time is spent on digesting the sheaf and cosheaf axioms, but the benefit is that we learn that $\mathsf{H}_0(-;\mathsf{k})$ is a cosheaf, by making use of the Mayer-Vietoris property. This provides the motivation for using homology with local coefficients to prove the equivalence of locally constant cosheaves with representations of the fundamental groupoid in section 5.1.4. This observation has not been made explicit before.

We have hoped to make the theory a stage filled with live actors. Section 4 develops the view that sheaves come about fundamentally as the sections of a map, thus making contact with Cartan's redevelopment, whereas cosheaves come about as the connected components of the fiber of a map. By taming the maps and spaces through the use of Reeb graphs in section 4.3 we give empirical evidence that one should be able to assign data directly to cells in lieu of open sets.

The very beginning of section 5 requires the least amount of background to start. One can begin there, but be prepared to be see no reference to open sets in the definition of a cellular sheaf or a cellular cosheaf. This is the situation in Shepard's thesis and he does not explain how these gadgets define (co)sheaves in the sense of section 3. This is remedied in two ways:

- Cellular sheaves and cosheaves are instances of constructible sheaves and cosheaves. The vision for this approach is due to Bob MacPherson, who left it to us to develop precise statements and proofs. Consequently, if section 5.1 appears to be technically more challenging than most of the paper, then it is only because we have hoped to provide a service to the mathematical community at large. The reader should feel free to skip to section 5.2 as this is the model of cellular sheaves and cosheaves used throughout the paper. The contributions of section 5.1 are as follows:
  - In 5.1.1 we provide a brief treatment of Whitney and Thom-Mather stratified spaces. Proposition 5.20 contains, apparently for the first time in published form, a modified proof of Mark Goresky's that closed unions of strata have regular neighborhoods.
  - In 5.1.2 we introduce stratified maps, which include Morse functions as a special case. These will be the fundamental source of examples of constructible cosheaves. Not all stratified maps can be triangulated, but the ones that can satisfy an extra condition called Thom's condition $\mathfrak{a}_{\mathsf{f}}$. Lemma 5.34, whose proof is joint with Goresky, provides a technical guarantee based on dimension alone to show when a stratified map satisfies this extra condition.
  - Section 5.1.3 provides an introduction to tame topology and o-minimal structures. Although MacPherson did not use this, we find this extra structure necessary to make certain proofs and constructions go through. Lemma 5.39 contains an easy proof that definable sets and maps are closed under pullback. As a motivating case study, we bring point-cloud data - the starting point for persistent homology - under the umbrella of semi-algebraic geometry and hence o-minimal topology. The Whitney stratifiability of definable sets and maps is integral to our approach in later sections. This is foreshadowed by lemma 5.43, which constructs geometrically a cellular cosheaf from a stratified map $\mathsf{f} : \mathsf{X} \to \mathbb{R}$. This is generalized in all dimensions in theorem 5.72.
  - In section 5.1.4, we begin the development of constructible cosheaves by introducing locally constant cosheaves, which use the abstract open set definition. We then show that these gadgets are equivalent to ones that assign data to points and maps to homotopy classes of paths.
  - Section 5.1.5 is the culmination of section 5.1. We introduce the definable entrance path category, which is our modification of MacPherson's original definition. Borrowing some work of Jon Woolf and David Miller, we then show in proposition

5.62 how all this complexity falls out when X is stratified as a cell complex. To show that representations of the entrance path category define cosheaves, we give a more algorithmic proof of the van Kampen theorem in the stratified setting. Theorem 5.65, the supporting lemma 5.68 and proposition 5.69 should be considered joint with David Lipsky. Theorem 5.72 provides, in essence, a geometric construction of the higher direct image images of the constant cosheaf along a stratified map. This technical tool provides one way of applying cosheaf theory to multi-dimensional persistent homology.

- The approach taken in 5.2 to realizing cellular sheaves and cosheaves as bona fide sheaves and cosheaves is a comparatively more simple theory. Whereas the first approach generalizes to arbitrary stratified spaces, this theory generalizes to arbitrary posets. Here the Alexandrov topology makes a crucial appearance in section 5.2.1. By viewing a cell complex as a set of cells that are partially ordered by the face relation, we can then specialize sheaves on posets to cellular sheaves. Theorem 5.87 states that sheaves on a poset $(\mathsf{X}, \leqslant)$ are equivalent to functors modeled on X as a poset. Although an old result, we give a novel proof via Kan extensions, which although not strictly necessary, may delight the advanced student of category theory.

In section 6 we develop the functoriality of sheaf and cosheaf theory using the working model of section 5.2. Since the partial order in a poset can always be turned around to define a new poset, Kan extensions clarify the difference between sheaves and cosheaves in light of these extra symmetries. We introduced three functors associated to pushing forward sheaves along a map: $\mathsf{f}_*, \mathsf{f}_\dagger$, and $\mathsf{f}_!$. The first two are well-defined for all posets, but the third is a cellular model for the pushforward with compact supports. Confusingly, many esteemed mathematicians will use $\mathsf{f}_!$ to mean what we call $\mathsf{f}_\dagger$. This confusion is understandable because normally the left adjoint of the pullback functor f^* is called $\mathsf{f}_!$, which does not always exist in the situations where the pushforward with compact supports functor is defined. We give the functor $\mathsf{f}_\dagger$ a topological interpretation by interchanging open sets with closed sets and sheaves with cosheaves. This is described in section 6.1.4.

The real computational simplicity of cellular sheaves and cosheaves is finally made apparent in section 7. The reader familiar with ordinary cellular homology and cohomology will find no trouble in adapting that definition to the one here. Section 7.1 will be of interest to the person wanting to compute sheaf cohomology and cosheaf homology by hand.

The visually-minded reader or the person coming from a background in persistent homology may like the use of representation theory in section 7.2. The decomposition of cellular sheaves and cosheaves into indecomposable ones provides a distinguished basis for interpreting the linear algebra computations from section 7.1. The indecomposables that arise in persistent homology are often called "barcodes" and it is the hope of the author that thinking of generalized barcodes will be helpful in understanding cellular sheaves and cosheaves, as well as their homologies. This will require careful attention because one then must think of these generalized barcodes as having topology that is sensitive to its embedding in the surrounding space. This is exemplified in the later theorems in section 8, where Borel-Moore homology makes a strong appearance.

The representation theory perspective is further leveraged to illustrate Ghrist and Hiroaka's work on sheaves in network coding – specifically, their use of a "decoding wire" – in section 9. We also give a combinatorial proof as to why $\mathsf{H}^0(\mathsf{X};\mathsf{F}) \cong \mathsf{H}^1(\mathsf{X};\mathsf{F})$.

The application of sheaves and cosheaves to the study of multi-modal sensor networks is outlined in section 10. It was here that the author first realized the necessity of using representation theory to interpret the topological meaning of cosheaf homology. A deeper examination of the act of sensing makes it clear how sheaves and cosheaves need to be used simultaneously.

Section 11 introduces more of the standard material of derived categories for sheaves and cosheaves. Much of the material is presented for cellular cosheaves and is easily dualized from Shepard's thesis. Here we prove that the computational definitions found in section 7.1 agree with the traditional use of injective and projective resolutions. Invariance under subdivision is then proved in this setting.

However, in the poset setting extra symmetries and new directions for exploration emerge. In particular, the existence of both enough projectives and injectives allows us to define **sheaf homology** and **cosheaf cohomology**. These theories, presented in section 11.4, appear to be another original contribution. Conjecturally, the sheaf homology defined here is the same as the one defined by Bredon [16]. We use a standard trick of Poincaré's to show how in the case of a manifold the author's definition of sheaf homology is meaningful. This theory is invariant under subdivision in the domain of a cellular map, but it is not invariant under subdivision in the target. In this sense the definition of sheaf homology is not topological, but rather is sensitive to the cell structure.

Section 12 contains the author's proof of a conjecture by MacPherson on the derived equivalence of cellular sheaves and cosheaves. The equivalence seems to have been known, but we make clear both its topological origin and provide an explicit formula for turning any sheaf into a complex of projective cosheaves. Verdier duality is then a consequence of this equivalence. We prove, using the formula for the derived equivalence, that the classical Poincaré duality can be phrased using sheaf cohomology *and* sheaf homology, as defined in section 11.4.

We synthesize this equivalence with the process of tensoring a sheaf and a cosheaf together via the use of coends in section 13. This extends an observation from the topos theory community that "cosheaves are valuations on sheaves," i.e. the category of colimit-preserving functors on sheaves is equivalent to the category of cosheaves. Here we observe that compactly supported sheaf cohomology can be defined as the valuation determined by the image of the constant sheaf through the derived equivalence defined in section 12. We conclude with some speculations and hints at further work under way.

1.2 How to Get Straight to the Applications

We advise the reader interested primarily in applications to proceed directly to the definition of a cellular sheaf and cosheaf at the start of section 5. One should skip sections 5.1 and 5.2 and then move onto section 7.1 for the formulas used to compute cellular sheaf cohomology and cosheaf homology. One can then read sections 8, 9, 10 for the applications to persistence, network coding, and sensor networks.

1.3 Acknowledgements

The author owes much to the stewardship and mathematical aesthetic of Robert Ghrist and Bob MacPherson, the former being the author's PhD advisor. The author is indebted to David Lipsky for spending countless hours listening to, and clarifying, many of the ideas and arguments in this work. Appreciation goes to Henry Adams, Steve Awodey, Jonathan Block, Gunnar Carlsson, Mark Goresky, Yasuaki Hiraoka, Sefi Ladkani, Sanjeevi Krishnan, Jacob Lurie, Michael Robinson, Aaron Royer, Hiro Lee Tanaka, David Treumann, and Jon Woolf for the valuable conversations directly concerning the topics in this paper. Additional support and encouragement during the writing of this work was provided by Shiying Dong, Greg Henselman, Michael Lesnick, Vidit Nanda, Amit Patel, Mikael Vejdemo-Johansson, and the author's family.

This work was supported by federal contracts FA9550-09-1-0643, FA9550-12-1-0416, HQ0034-12-C-0027, and a Benjamin Franklin Fellowship from the University of Pennsylvania. The author would also like to thank the hospitality of Princeton University and the Institute for Advanced Study, where much of this work was written.

2 Categories: Limits and Colimits

Categories emerged out of the study of functors, which were originally conceived as a principled way of assigning algebraic invariants to topological spaces. Thus, category theory is part and parcel of the study of algebraic topology. However, from its conception in Samuel Eilenberg and Saunders Mac Lane's 1945 paper on a "General Theory of Natural Equivalence" [28], it was realized that the language of categories provides a way of identifying formal similarities throughout mathematics. The success of this perspective is largely due to the fact that category theory – as opposed to set theory – emphasizes understanding the relationships between objects rather than the objects themselves.

In this section, we provide a brief review of the parts of category theory needed to understand the abstract definitions of a sheaf and cosheaf in section 3. Most importantly, the reader should be able to do the following before moving onto that section:

- Think of the set of open sets of a topological space X as a category.
- Understand how to summarize the behavior of various functors via limits and colimits.

We have tried to provide a self-contained introduction to category theory, but the reader is urged to consult Mac Lane's "Categories for the Working Mathematician" [55] for a more thorough introduction.

2.1 Categories

One should visualize categories as graphs with objects corresponding to vertices and maps as edges between vertices, subject to relations that specify when following one sequence of edges is equivalent to another sequence. One can think of some of the axioms of a category as gluing in triangles and tetrahedra to witness these relations.

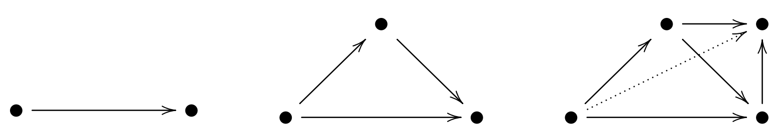

Definition 2.1 (Category)**.** A **category** $\mathcal{C}$ consists of a class of objects denoted $\mathbf{obj}(\mathcal{C})$ and a set of morphisms $\mathbf{Hom}_{\mathcal{C}}(a, b)$ between any two objects $a, b \in \mathbf{obj}(\mathcal{C})$. An individual morphism $f : a \to b$ is also called an arrow since it points (maps) from a to b. We require that the following axioms hold:

- Two morphisms $f \in \mathbf{Hom}_{\mathcal{C}}(a, b)$ and $g \in \mathbf{Hom}_{\mathcal{C}}(b, c)$ can be composed to get another morphism $g \circ f \in \mathbf{Hom}_{\mathcal{C}}(a, c)$.
- Composition is associative, i.e. if $h \in \mathbf{Hom}(c, d)$, then $(h \circ g) \circ f = h \circ (g \circ f)$.
- For each object x there is an identity morphism $\mathrm{id}_x \in \mathbf{Hom}_{\mathcal{C}}(x, x)$ that satisfies $f \circ \mathrm{id}_a = f$ and $\mathrm{id}_b \circ f = f$.

When the category $\mathcal{C}$ is understood, we will sometimes write $\mathbf{Hom}(a, b)$ to mean $\mathbf{Hom}_{\mathcal{C}}(a, b)$.

One can usually ignore the technicality that the collection of objects forms a class rather than a set. A class is a collection of sets that one can refuse to quantify over in a logical sense. This prohibits Russell-type paradoxes gotten by considering the category of all categories that do not contain themselves. Colloquially, one says a proper class is "bigger" than a set. In order to avoid certain machinery that accompanies the use of classes, we will often consider categories that are "small" in a precise sense.[3]

Definition 2.2 (Small Category). A category is **small** if its class of objects is actually a set.

Example 2.3 (Discrete Category). Any set X can be regarded as a **discrete category** $\bar{X}$ with only the identity morphism id_x sitting over each object. There are no non-trivial morphisms.

Recall that a **relation** R on a set X is a subset of the product set $X \times X$. If two elements are related by R, one writes xRy to mean that $(x, y) \in R$. We now give an example of some relations on a set that endow that set with the structure of a category.

Example 2.4 (Posets and Preorders). A **preordered set** is a set X along with a relation $\leqslant$ that satisfies the following two axioms:

- **Reflexivity** – $x \leqslant x$ for all $x \in X$
- **Transitivity** – $x \leqslant y$ and $y \leqslant z$ implies $x \leqslant z$

A **partially ordered set**, or **poset** for short, is a preordered set that additionally satisfies the following third axiom:

- **Anti-Symmetry** – $x \leqslant y$ and $y \leqslant x$ implies $x = y$

Any preordered set $(X, \leqslant)$ defines a category by letting the objects be the elements of X and by declaring each hom set $\mathbf{Hom}(x, y)$ to either have a unique morphism if $x \leqslant y$ or to be empty if $x \nleqslant y$.

We now reach an example of fundamental importance.

Example 2.5. The **open set category** associated to a topological space X, denoted $\mathbf{Open}(X)$, has as objects the open sets of X and a unique morphism $U \to V$ for each pair related by inclusion $U \subseteq V$.

The above examples of categories are quite small when compared to the categories that Eilenberg and Mac Lane first introduced. The categories considered there correspond to **data types** and we will usually refer to them with the letter $\mathcal{D}$. For this paper $\mathcal{D}$ will usually mean one of the following:

[3] The machinery we are referring to is that of Grothendieck universes.

- $\mathcal{Set}$ – the category whose objects are sets and whose morphisms are all set maps (multi-valued maps are prohibited as are partially defined maps)
- $\mathcal{Ab}$ – the category whose objects are abelian groups and whose morphisms are group homomorphisms
- $\mathcal{Vect}$ – the category whose objects are vector spaces and whose morphisms are linear transformations
- $\mathcal{vect}$ – the category whose objects are *finite-dimensional* vector spaces and linear transformations
- $\mathcal{Top}$ – the category whose objects are topological spaces and whose morphisms are continuous maps

The category $\mathcal{vect}$ is an example of a subcategory, which we now define.

Definition 2.6 (Subcategories)**.** Let $\mathcal{C}$ be a category. A **subcategory** $\mathcal{B}$ of $\mathcal{C}$ consists of a collection of objects from $\mathcal{C}$ and a choice of subset of the morphism set $\mathbf{Hom}_{\mathcal{C}}(x, y)$ for each pair $x, y \in \mathbf{obj}(\mathcal{B})$. We require that these morphism sets have the identity and be closed under composition so as to guarantee that $\mathcal{B}$ is a category and that the inclusion $\mathcal{B} \hookrightarrow \mathcal{C}$ is a functor. We say that a subcategory is **full** if $\mathbf{Hom}_{\mathcal{B}}(x, y) = \mathbf{Hom}_{\mathcal{C}}(x, y)$.

Categories have a built-in notion of directionality. For example, in $\mathcal{Set}$ every object X has a unique map from the empty set $\emptyset$, but there are no maps to the empty set. We can abstract out this property, so as to make it apply in other situations.

Definition 2.7 (Initial and Terminal Objects)**.** An object $x \in \mathbf{obj}(\mathcal{C})$ is said to be **initial** if for any other object $y \in \mathbf{obj}(\mathcal{C})$ there is a unique morphism from x to y. Dually, an object y is said to be **terminal** if for any object x there is a unique morphism from x to y.

As already mentioned, in $\mathcal{Set}$ the empty set is initial, but it is not terminal. On the contrary, the terminal object is the one point set $\{\star\}$ since there is only one constant map. Similarly, for $\mathcal{Open}(X)$ the empty set is initial, but the whole space X is terminal. In $\mathcal{Vect}$ the initial and terminal objects coincide with the zero vector space. In some sense, the difference between the initial and terminal objects in a category measure how different it is from its reflection. We now say what we mean by a category's reflection.

Example 2.8 (Opposite Category)**.** For any category $\mathcal{C}$ there is an **opposite category** $\mathcal{C}^{op}$ where all the arrows have been turned around, i.e. $\mathbf{Hom}_{\mathcal{C}^{op}}(x, y) = \mathbf{Hom}_{\mathcal{C}}(y, x)$.

Remark 2.9 (Duality and Terminology). Because one can always perform a general categorical construction in $\mathcal{C}$ or $\mathcal{C}^{op}$ every concept is really two concepts. As we shall see, this causes a proliferation of ideas and is sometimes referred to as the **mirror principle**. The way this affects terminology is that a construction that is dualized is named by placing a "co" in front of the name of the un-dualized construction. Thus, we will have limits and colimits, products and coproducts, equalizers and coequalizers, among other things.

Now we introduce the fundamental device that assigns objects and morphisms in one category to objects and morphisms in another category. Historically, this device was introduced first and categories were summoned into existence to provide a domain and range for this assignment.

Definition 2.10 (Functor)**.** A **functor** $F : \mathcal{C} \to \mathcal{D}$ consists of the following data: To each object $a \in \mathcal{C}$ an object $F(a) \in \mathcal{D}$ is associated, i.e. $a \rightsquigarrow F(a)$. To each morphism $f : a \to b$ a morphism $F(f) : F(a) \to F(b)$ is likewise associated. We require that the functor respect composition and preserve identity morphisms, i.e. $F(f \circ g) = F(f) \circ F(g)$ and $F(\mathrm{id}_a) = \mathrm{id}_{F(a)}$. For such a functor F, we say $\mathcal{C}$ is the **domain** and $\mathcal{D}$ is the **codomain** of F.

Remark 2.11. We can phrase the definition of a functor differently by saying that we have a function $F : \mathbf{obj}(\mathcal{C}) \to \mathbf{obj}(\mathcal{D})$ and functions $F(a, b) : \mathbf{Hom}_{\mathcal{C}}(a, b) \to \mathbf{Hom}_{\mathcal{D}}(F(a), F(b))$ for every pair of objects $a, b \in \mathbf{obj}(\mathcal{C})$. We require that these functions preserve identities and composition. When $F(a, b) : \mathbf{Hom}_{\mathcal{C}}(a, b) \to \mathbf{Hom}_{\mathcal{D}}(F(a), F(b))$ is injective for every pair of objects we say F is **faithful**. When $F(a, b)$ is surjective for every pair of objects we say F is **full**. When a functor is both full and faithful, we say it is **fully faithful**.

An example familiar to every topologist is that of homology and cohomology with field coefficients. In every non-negative degree i, these invariants define functors

$$H_i(-; k) : \mathbf{Top} \to \mathbf{Vect} \qquad \text{and} \qquad H^i(-; k) : \mathbf{Top}^{op} \to \mathbf{Vect}$$

respectively. Here we have used the opposite category as an alternative way of saying cohomology is **contravariant**.

As the reader may well be aware, there are different types of homology theories (Čech, cellular, singular, etc.) and understanding the precise relationships between these motivated the notion of a map between functors.

Definition 2.12 (Natural Transformation)**.** Given two functors $F, G : \mathcal{C} \to \mathcal{D}$ a **natural transformation**, sometimes written $\eta : F \Rightarrow G$, consists of the following information: to each object $a \in \mathcal{C}$, a morphism $\eta(a) : F(a) \to G(a)$ is assigned such that for every morphism $f : a \to b$ in $\mathcal{C}$ the following diagram **commutes**:

$$\begin{tikzcd} F(a) \arrow[r, "\eta(a)"] \arrow[d, "F(f)"'] & G(a) \arrow[d, "G(f)"] \\ F(b) \arrow[r, "\eta(b)"] & G(b) \end{tikzcd}$$

By commutes, we mean $G(f) \circ \eta(a) = \eta(b) \circ F(f)$.

Definition 2.13. Two functors $F, G : \mathcal{C} \to \mathcal{D}$ are said to be **naturally isomorphic** if there is a natural transformation $\eta : F \Rightarrow G$ such that for every object $a \in \mathcal{C}$ the morphism $\eta(a)$ is an isomorphism, i.e. it is invertible. These inverse maps $\eta(a)^{-1}$ define an inverse natural transformation $\eta^{-1} : G \Rightarrow F$.

Functors and natural transformations assemble themselves into a category in their own right. As such, we will sometimes use the notation $\mathsf{F} \to \mathsf{G}$, instead of $\mathsf{F} \Rightarrow \mathsf{G}$, for a natural transformation. We do this to be consistent with our notation since every category has arrows and using different notations for different categories is confusing. In the functor category, we will see that naturally isomorphic functors are isomorphic objects. This demonstrates again the linguistic efficiency of category theory.

Example 2.14 (Functor Category)**.** $\mathbf{Fun}(\mathcal{C}, \mathcal{D})$ denotes the category whose objects are functors from C to D and whose morphisms are natural transformations.

Certain functors deserve special attention. These are the ones that allow us to identify two different categories. One approach to identifying categories is to say that two categories $\mathcal{C}$ and $\mathcal{D}$ are **isomorphic** if there are functors $\mathsf{F} : \mathcal{C} \to \mathcal{D}$ and $\mathsf{G} : \mathcal{D} \to \mathcal{C}$ such that $\mathsf{G} \circ \mathsf{F} = \mathrm{id}_{\mathcal{C}}$ and $\mathsf{F} \circ \mathsf{G} = \mathrm{id}_{\mathcal{D}}$. This definition is so restrictive that it rarely occurs. Thus, we have a looser notion that includes isomorphism as a special case. Instead of asking that $\mathsf{F} \circ \mathsf{G}$ be equal to $\mathrm{id}_{\mathcal{D}}$, we only require that they be isomorphic as objects in $\mathbf{Fun}(\mathcal{D}, \mathcal{D})$ and similarly for $\mathsf{G} \circ \mathsf{F}$ and $\mathrm{id}_{\mathcal{C}}$ in $\mathbf{Fun}(\mathcal{C}, \mathcal{C})$. The reader should compare this with the notion of homotopy equivalence. We phrase this idea in a simpler way that doesn't require us to construct G as a "weak inverse" of F.

Definition 2.15 (Equivalence)**.** A functor $\mathsf{F} : \mathcal{C} \to \mathcal{D}$ induces an **equivalence** of categories if it is bijective on **Hom** sets (fully faithful) and is **essentially surjective**. This last property means that for every object $\mathsf{d} \in \mathcal{D}$ there is an object $\mathsf{c} \in \mathcal{C}$ such that $\mathsf{F}(\mathsf{c})$ is isomorphic to d, i.e. F is bijective on isomorphism classes of $\mathcal{C}$ and $\mathcal{D}$.

Finally, let's analyze how working in the opposite category impacts functors and natural transformations. Observe, first and foremost, that formality allows us to take a functor $\mathsf{F} : \mathcal{C} \to \mathcal{D}$ and define a functor $\mathsf{F}^{\mathsf{op}} : \mathcal{C}^{\mathsf{op}} \to \mathcal{D}^{\mathsf{op}}$. Moreover, a natural transformation $\eta : \mathsf{F} \Rightarrow \mathsf{G}$ translates to a natural transformation $\eta^{\mathsf{op}} : \mathsf{G}^{\mathsf{op}} \Rightarrow \mathsf{F}^{\mathsf{op}}$. This observation allows us to state the equalities

$$\mathbf{Fun}(\mathcal{C}^{\mathsf{op}}, \mathcal{D}^{\mathsf{op}}) = \mathbf{Fun}(\mathcal{C}, \mathcal{D})^{\mathsf{op}} \qquad \text{or} \qquad \mathbf{Fun}(\mathcal{C}^{\mathsf{op}}, \mathcal{D}^{\mathsf{op}})^{\mathsf{op}} = \mathbf{Fun}(\mathcal{C}, \mathcal{D})$$

since $(\mathcal{C}^{\mathsf{op}})^{\mathsf{op}}$ is isomorphic to $\mathcal{C}$ (not just equivalent). See the wonderful work "Abstract and Concrete Categories: The Joy of Cats" [2] for more on duality and category theory more generally.

2.2 Diagrams and Representations

Functors and categories allow us to develop algebra modeled on certain shapes governed by the domain category. For example, we will be interested in studying data arranged in the following forms:

$$\begin{tikzcd} \bullet \arrow[r] \arrow[d] & \bullet \\ \bullet & \end{tikzcd} \qquad \begin{tikzcd} \bullet \arrow[dr, shift left] \arrow[dr, shift right] & \bullet \\ \bullet & \bullet \end{tikzcd} \qquad \begin{tikzcd} & \bullet \arrow[d] \\ \bullet \arrow[r] & \bullet \end{tikzcd}$$

If we imagine the identity arrows in a category as being the vertices themselves, and thus not drawn independently of the objects, each of these shapes gives an example of a finite category.

Definition 2.16 (Diagram)**.** Suppose I is a small category and $\mathcal{C}$ is an arbitrary category. A **diagram** is simply a functor $F : I \to \mathcal{C}$.

Example 2.17 (Constant Diagram)**.** For any category I there is always a diagram for each object $O \in \mathcal{C}$, called the **constant diagram**, $\mathrm{const}_O : I \to \mathcal{C}$ where $\mathrm{const}_O(x) = \mathrm{const}_O(y) = O$ for all objects $x, y \in I$. Every morphism in I goes to the identity morphism.

Definition 2.18 (Representation)**.** A **representation** of a category $\mathcal{C}$ is a functor $F : \mathcal{C} \to$ ***Vect***.

One should note that this definition generalizes the notion of a representation of a group. Every group, say $\mathbb{Z}$ for example, can be considered as a small category with a single object $\star$ and $\mathbf{Hom}(\star, \star) = \mathbb{Z}$. A representation of $\mathbb{Z}$ then corresponds to picking a vector space V and assigning an endomorphism of V for each element of $\mathbb{Z}$, i.e. it is a functor.

$$\begin{tikzcd}
\star \arrow[r, squiggly] \arrow[d, "g"'] & V \arrow[d, "\rho(g)"] \\
\star \arrow[r, squiggly] & V
\end{tikzcd}$$

Maps of representations correspond precisely with natural transformations of such functors. Isomorphic representations are naturally isomorphic functors.[4] These basic notions carry over to the representation theory of arbitrary categories, which allows us to compare different situations in one language.

2.3 Cones and Limits

The next two sections are devoted to studying one way (and a dual way) of summarizing a functor's behavior. This gives a way of compressing the data of a functor into a single object. These concepts are fundamental to the study of sheaves and cosheaves.

Definition 2.19 (Cone)**.** Suppose $F : I \to \mathcal{C}$ is a diagram. A **cone** on F is a natural transformation from a constant diagram to F. Specifically, it is a choice of object $L \in \mathcal{C}$ and a collection of morphisms $\psi_x : L \to F(x)$, one for each x, such that if $g : x \to y$ is a morphism in I, then $F(g) \circ \psi_x = \psi_y$, i.e. the following diagram commutes:

$$\begin{tikzcd}
F(x) \arrow[rr, "F(g)"] & & F(y) \\
& L \arrow[ul, "\psi_x"] \arrow[ur, "\psi_y"'] &
\end{tikzcd}$$

In other words, $\psi_y = F(g) \circ \psi_x$.

[4]Confusingly, the term "equivalent representations" is often used.

Definition 2.20. The collection of cones on a diagram F form a category, which we will call $\mathbf{Cone}(F)$. The objects are cones (L, ψ_x) and a morphism between two cones (L', ψ'_x) and (L, ψ_x) consists of a map $u : L' \to L$ such that $\psi'_x = \psi_x \circ u$ for all x

A limit is simply a distinguished or universal object in the category of cones on F.

Definition 2.21 (Limit)**.** The **limit** of a diagram $F : I \to \mathcal{C}$, denoted $\varprojlim F$ is the terminal object in $\mathbf{Cone}(F)$. This means that a limit is an object $\varprojlim F \in \mathcal{C}$ along with a collection of morphisms $\psi_x : L \to F(x)$ that commute with arrows in the diagram such that whenever there is another object L' and morphisms ψ'_x that also commute there then exists is a unique morphism $u : L' \to \varprojlim F$ that additionally commutes with everything in sight, i.e. $\psi'_x = \psi_x \circ u$ for all x.

$$\begin{tikzcd}
F(x) \arrow[rr, "F(g)"] & & F(y) \\
& \varprojlim F \arrow[ul, "\psi_x"'] \arrow[ur, "\psi_y"] & \\
& L' \arrow[u, dotted, "\exists!"', "u"] \arrow[uul, "\psi'_x"] \arrow[uur, "\psi'_y"'] &
\end{tikzcd}$$

Remark 2.22 (Glossary)*.* Quite confusingly, the following terms are synonyms for limits: inverse limits, projective limits, left roots, lim and $\varprojlim$ are all common.

We now consider some examples of limits over discrete categories.

Example 2.23 (Products)**.** Consider the following index category and diagram:

$$\bullet \qquad \bullet \qquad\qquad\qquad F(i) \qquad F(j)$$

The limit of this diagram is called the **product** and is usually written

$$F(i) \prod F(j).$$

More generally, we define the product to be the limit of any diagram $F : I \to \mathcal{C}$ indexed by a discrete category and write $\prod_i F(i)$. Sometimes one writes $\times_i F(i)$ for the product.

We give an unusual example of a product that will prepare the reader for thinking about the category of open sets.

Example 2.24 (Open Sets: Limits are Intersections)**.** Suppose $\Lambda = \{1, \dots, n\}$ is a finite discrete category, i.e. it has n objects and the only morphisms are the identity morphisms. Now let X be a topological space and let $\mathcal{C} = \mathbf{Open}(X)$ be the category of open sets in X. This is a category that has an object for each open set and a single morphism $U \to V$ if $U \subset V$. A functor $F : \Lambda \to \mathbf{Open}(X)$ is nothing more than a choice of n not necessarily distinct open sets. A cone to F is an open set that includes into all the open sets picked out by F. The limit of F is the largest possible open set that includes into all the open sets picked out by F, i.e.

$$\varprojlim F = \cap_{i=1}^{n} F(i).$$

Example 2.25. Consider the following small category I along with some representation $F : I \to \boldsymbol{Vect}$.

$$\begin{tikzcd} \bullet \arrow[r] \arrow[d] & \bullet \\ \bullet & \end{tikzcd} \qquad \begin{tikzcd} U \arrow[r, "A"] \arrow[d, "B"'] & V \\ W & \end{tikzcd}$$

By thinking about the definition, one can see that

$$\varprojlim F \cong U.$$

Example 2.26 (Pullbacks)**.** Consider the category $J = I^{op}$ and a representation $F : J \to \boldsymbol{Vect}$.

$$\begin{tikzcd} & \bullet \arrow[d] \\ \bullet \arrow[r] & \bullet \end{tikzcd} \qquad \begin{tikzcd} & V \arrow[d, "A"] \\ W \arrow[r, "B"'] & U \end{tikzcd}$$

With some thought one can describe the limit set-theoretically as

$$\varprojlim F \cong \{(v, w) \in V \times W | Av = Bw\},$$

which is called the **pullback**. If $U = 0$, then we re-obtain the product of V and W and one usually writes $V \times W$.

Example 2.27 (Equalizers and Kernels)**.** Consider the following category K and an arbitrary functor $F : K \to \mathcal{D}$.

$$\begin{tikzcd} \bullet \arrow[r, shift left] \arrow[r, shift right] & \bullet \end{tikzcd} \qquad \begin{tikzcd} X \arrow[r, shift left, "f"] \arrow[r, shift right, "g"'] & Y \end{tikzcd}$$

The limit of this diagram, which is also called the **equalizer**, is an object E along with a map h that satisfies $f \circ h = g \circ h$.

$$\begin{tikzcd} E \arrow[r, "h"] & X \arrow[r, shift left, "f"] \arrow[r, shift right, "g"'] & Y \end{tikzcd}$$

If $\mathcal{D} = \boldsymbol{Vect}$ and one sets $g = 0$, then the equalizer is the **kernel**. Thus, if one wants to mimic kernels in data types lacking of zero maps and objects, equalizers can be substituted.

Finally, we finish with an example from representation theory.

Example 2.28 (Invariants)**.** Suppose that V is a vector space with an endomorphism $T : V \to V$, i.e. a $k[x]$-module. Just as a group can be viewed as a category with one object, a ring can be viewed as a category with multiplication corresponding to composition of morphisms and addition corresponding to addition of morphisms, thus such a category has extra structure. Thus the $k[x]$-module determined by V and T is equivalent to a functor $k[x] \to \boldsymbol{Vect}$ that sends the unique object $\star$ to V and sends x to T. The limit of such a functor is called the **invariants** of the action, i.e.

$$I = \{v \in V \,|\, T(v) = v\}.$$

2.4 Co-Cones and Colimits

Here we invoke the mirror principle to dualize the theory of cones and limits. In accordance with usual terminology, we refer to these as *co*cones and *co*limits.

Definition 2.29 (Co-Cone)**.** Given a diagram $F : I \to \mathcal{C}$, a **cocone** is a natural transformation from F to a constant diagram. In other words, it consists of an object $C \in \mathcal{C}$ along with a collection of maps $\phi_x : F(x) \to C$ such that these maps commute with the ones internal to the diagram.

$$\begin{tikzcd}
& C & \\
F(x) \arrow[ur, "\phi_x"] \arrow[rr, "F(g)"'] & & F(y) \arrow[ul, "\phi_y"']
\end{tikzcd}$$

Similarly, there is a category of cocones to a diagram F, denoted $\boldsymbol{CoCone}(F)$. A colimit is a distinguished object in this category.

Definition 2.30 (Colimit)**.** The **colimit** of a diagram F is the initial object in the category $\boldsymbol{CoCone}(F)$. One should practice dualizing the explicit description of the limit in order to understand the following diagram:

$$\begin{tikzcd}
& C' & \\
& \varinjlim F \arrow[u, dotted, "\exists!\ u"] & \\
F(x) \arrow[uur, "\phi'_x"] \arrow[ur, "\phi_x"] \arrow[rr, "F(g)"'] & & F(y) \arrow[uul, "\phi'_y"'] \arrow[ul, "\phi_y"']
\end{tikzcd}$$

Remark 2.31 (Glossary). The following terms are synonyms for colimits: direct limits, inductive/injective limits, right roots, colim and $\varinjlim$ are all used.

To better understand the similarities and differences between limits and colimits, let us re-examine the same examples in the previous section.

Example 2.32 (Coproducts)**.** Consider the following index category and diagram:

$$\bullet \qquad \bullet \qquad\qquad\qquad F(i) \qquad F(j)$$

The colimit of this diagram is called the **coproduct** and is usually written

$$F(i) \coprod F(j).$$

More generally, we define the product to be the limit of any diagram $F : I \to \mathcal{C}$ indexed by a discrete category and write $\coprod_i F(i)$. Alternative notations for the coproduct, depending usually on whether the target category is $\boldsymbol{Set}$, $\boldsymbol{Vect}$ or $\boldsymbol{Ab}$, include

$$\bigoplus_i F(i) \qquad \text{and} \qquad \sum_i F(i).$$

Example 2.33 (Open Sets: Colimits are Unions)**.** Suppose $\Lambda = \{1, \ldots, n\}$ is a finite discrete category. Let $\mathcal{C} = \mathbf{Open}(X)$ be the category of open sets in X. A functor $F : \Lambda \to \mathbf{Open}(X)$ is a choice of n not necessarily distinct open sets. A cocone to F is an open set that contains all the open sets picked out by F. The colimit of F is the smallest possible open set containing all the open sets picked out by F, i.e. the union:

$$\varinjlim F = \cup_{i=1}^{n} F(i)$$

One should note that since the arbitrary union of open sets is still open one could have worked over a larger indexing category Λ.

Example 2.34 (Pushouts)**.** Consider the following small category I and a representation $F : I \to \textbf{\textit{Vect}}$.

$$\begin{tikzcd} \bullet \arrow[r] \arrow[d] & \bullet \\ \bullet & \end{tikzcd} \qquad\qquad \begin{tikzcd} U \arrow[r, "A"] \arrow[d, "B"'] & V \\ W & \end{tikzcd}$$

Contrary to the case of the limit, this one requires a bit more thought. Let's start with something that is *not* a cocone, but is nevertheless naturally built out of pieces of the diagram.

$$\begin{tikzcd} U \arrow[r, "A"] \arrow[d, "B"'] \arrow[dr, "B\oplus A"] & V \arrow[d, "\iota_V"] \\ W \arrow[r, "\iota_W"'] & W \oplus V \end{tikzcd}$$

This is not a cocone because the diagram does not commute since $(Bu, 0) \neq (Bu, Au) \neq (0, Au)$. We can force commutativity by forcing the equivalence relation $[(Bu, 0)] \sim [(0, Au)]$ or equivalently $[(Bu, -Au)] \sim [(0, 0)]$. We thus conclude that

$$\varinjlim F = W \oplus V / \operatorname{im}(B \oplus -A) \quad \phi_U = q \circ \iota_W B = q \circ \iota_V A \quad \phi_W = q \circ \iota_W \quad \phi_V = q \circ \iota_W$$

where q is the quotient map. One should note that this is clearly dual to the limit computation in 2.26 with the added complication that whereas the limit is a sub-object, the colimit is a quotient object.

Like before, if $U = 0$ then the pushout reduces to the coproduct of V and W and one writes it as $V \oplus W$.

Example 2.35. Consider the example $J = I^{op}$ and corresponding representation $F : J \to \textbf{\textit{Vect}}$.

$$\begin{tikzcd} & \bullet \arrow[d] \\ \bullet \arrow[r] & \bullet \end{tikzcd} \qquad\qquad \begin{tikzcd} & V \arrow[d, "A"] \\ W \arrow[r, "B"'] & U \end{tikzcd}$$

One can see that

$$\varinjlim F \cong U.$$

Example 2.36 (Coequalizers and Cokernels)**.** Consider the same category K as before and a functor $\mathsf{F} : \mathsf{K} \to \mathcal{D}$.

$$\bullet \rightrightarrows \bullet \qquad\qquad\qquad \mathsf{X} \underset{\mathsf{g}}{\overset{\mathsf{f}}{\rightrightarrows}} \mathsf{Y}$$

The colimit, which is called the **coequalizer**, is an object E and map h such that $\mathsf{h} \circ \mathsf{f} = \mathsf{h} \circ \mathsf{g}$.

$$\mathsf{X} \underset{\mathsf{g}}{\overset{\mathsf{f}}{\rightrightarrows}} \mathsf{Y} \xrightarrow{\mathsf{h}} \mathsf{E}$$

If $\mathcal{D} = \boldsymbol{Vect}$ and one sets $\mathsf{g} = 0$, then the coequalizer is the **cokernel**. Thus if one wants to mimic cokernels in data types lacking of zero maps and objects, coequalizers can be substituted.

Example 2.37 (Co-invariants)**.** As previously described, a vector space V with an endomorphism T is equivalent to a functor $\mathsf{k}[\mathsf{x}] \to \boldsymbol{Vect}$. The colimit of this functor is called the **coinvariants** of T, i.e.

$$\mathsf{C} = \mathsf{V}/ < \mathsf{T}\nu - \nu > .$$

3 Sheaves and Cosheaves

In its most abstract form, the subject of this article involves the assignment of data to subsets of a space X. This should sound like a very useful thing to do. After all, we have in both pure and applied mathematics many an occasion to record data or solutions in a local, spatially distributed way. Immediate questions arise: To which subsets should we assign data? What should these assignments be used for? What are they to be called?

The author believes such assignments are to be called sheaves or cosheaves depending on whether it is natural to restrict the data from larger spaces to smaller spaces or by extending data from smaller spaces to larger ones. The evolution of these ideas deserves some discussion and the eager historian should consult John Gray's "Fragments of the History of Sheaf Theory," [42] for a more thorough account. However, we outline three basic opinions on what a sheaf (or cosheaf) is really:

- A sheaf is a **system of coefficients** for computing cohomology that weighs and measures parts of the space differently. A cosheaf, in like manner, is a system of coefficients for homology that varies throughout the space.

- A sheaf is an **étalé space** E along with a local homeomorphism $\pi : E \to X$. Analogously, a cosheaf is a locally-connected space D, called the **display locale**, that maps to X [33].

- A sheaf (or a cosheaf) is an **abstract assignment of data** – a functor – that further satisfies a gluing axiom expressed by limits (or colimits).

Historically, the system of coefficients perspective came first. In a 1943 paper Norman Steenrod defined a new homology theory determined by assigning abelian groups directly to *points* of a space X and group isomorphisms to (homotopy classes of) paths between points [85]. This theory was vastly generalized in 1946 by Jean Leray where a **faisceau** (or sheaf) was defined to be a way of assigning modules to *closed sets* in an inclusion-reversing way.

$$\begin{tikzcd} V \arrow[rr] \arrow[dr] & & X \\ & W \arrow[ur] & \end{tikzcd} \qquad \begin{tikzcd} F(V) & & F(X) \arrow[ll] \arrow[dl] \\ & F(W) \arrow[ul] & \end{tikzcd}$$

Although this strengthened the abstract assignment perspective, Leray was still concerned with the cohomological ideas developed by Georges de Rham, Kurt Reidemeister and Hassler Whitney.

By the early 1950s, Henri Cartan and his seminar revised Leray's definition of a sheaf to consist of a local homeomorphism $\pi : E \to X$. One could re-obtain the assignment perspective by attaching to *open sets* U the set of sections of this map over U.

$$U \rightsquigarrow \{s : U \to E \mid \pi \circ s(x) = x\}$$

One plausible explanation for using open sets is provided by the **open pasting lemma**, which states[5] that if $X = \cup U_i$ is a (potentially infinite) union of open sets equipped with continuous sections $s_i : U_i \to E$ that agree on overlaps, then the set-theoretically defined section $s : X \to E$ will also be continuous. If closed sets are used, then this gluing argument only works for covers consisting of finitely many closed sets.

Finally, the Weil conjectures in algebraic geometry motivated the introduction of a more general notion of a topology and cohomology. Following suggestions of Jean-Pierre Serre, the domain of a sheaf was abstracted by Alexander Grothendieck from subsets $U \subseteq X$ to collections of mappings $U \to X$ that satisfy certain conditions reminiscent of an open cover [48]. Defining a sheaf on a Grothendieck topology ushered in the abstract formulation of sheaves using categories, functors and equalizers (limits) found in Michael Artin's 1962 Harvard notes on the subject [8].

All three of these models are useful for thinking about sheaves and cosheaves, but the abstract assignment model is powerful and elegant enough to capture the other two. Moreover, whereas the étalé space perspective can be adapted from sheaves of sets to sheaves of more general data types, the display space perspective on cosheaves appears to only be valid for set-valued cosheaves and cannot be adapted more generally. In particular, since homology requires working with abelian groups or vector spaces, the display space model and the homology perspective describe different types of cosheaves. Thus, the only vantage point capable of reasoning about cosheaves in a unified way is the functorial perspective, where the dualities of category theory can be employed.

In this section, we provide the abstract definition of sheaves and cosheaves, but restrict ourselves to considering open sets and covers in a topological space. We phrase things using limits and colimits that take the shape of a simplicial complex: the nerve of a cover. The sheaf or cosheaf condition says that the value of this limit or colimit is independent of the cover chosen. To make the limits and colimits over covers more computable, we reduce to equalizers and coequalizers. We then specialize to the data type of vector spaces, where Čech homology for a cover is introduced. This evolves into a discussion of why singular zeroth homology defines a cosheaf. As set up for the discussion on general differences between sheaves and cosheaves, we consider how refinement of covers plays interacts with the sheaf and cosheaf property.

3.1 The Abstract Definition

In elementary mathematics one learns that functions are devices for assigning points in one set to points in another. Motivated by differential calculus, one learns properties of functions on metric and topological spaces such as continuity. In its simplest form, continuity of a function states that if $f : X \to Y$ is a function and $\{x_n\}_{n=1}^{\infty}$ is a sequence of points in X converging to some point x, then

$$\lim_{n\to\infty} f(x_n) = f(\lim_{n\to\infty} x_n) = f(x),$$

[5] Munkres calls this the "local formulation of continuity" in theorem 18.2(f) [65]. Munkres reserves the term "pasting lemma" for the closed set version, which is stated directly afterwards as theorem 18.3.

i.e. f commutes with the limits one learns in analysis. Moreover, there is an independence result: The value $f(x)$ is independent of which sequence one used to approximate the point x.

The exact analogous situation occurs in category theory. A functor assigns objects and morphisms of one category to objects and morphisms in another. If a functor commutes with the categorical notion of a limit, then we also say that the functor is continuous. However, since there are so many different shapes of limits in arbitrary categories, this notion is too restrictive. A sheaf is a functor that commutes with limits coming from open covers. Applying the duality principle in category theory, a cosheaf is a functor that preserves colimits coming from open covers.

Pavel Alexandrov introduced in 1928 a method[6] for associating to every open cover an abstract simplicial complex [4]. We will use these shapes to model our limits and colimits of interest.

Definition 3.1. Suppose $\mathcal{U}$ is a set of open sets $\{U_i\}_{i\in\Lambda}$ contained in U. We can take the **nerve** of the cover to get an abstract simplicial complex $N(\mathcal{U})$, whose elements are subsets $I = \{i_0, \ldots, i_n\}$ for which $U_I := U_{i_0} \cap \cdots \cap U_{i_n} \neq \emptyset$. We can regard $N(\mathcal{U})$ as a category whose objects are the finite subsets I such that $U_I \neq \emptyset$ with a unique arrow from $I \to J$ if $J \subseteq I$. Since our intersections are only finite, and the finite intersection of open sets is open, we get natural functors

$$\iota_{\mathcal{U}} : N(\mathcal{U}) \to \mathbf{Open}(X) \qquad \text{or} \qquad \iota_{\mathcal{U}}^{op} : N(\mathcal{U})^{op} \to \mathbf{Open}(X)^{op}.$$

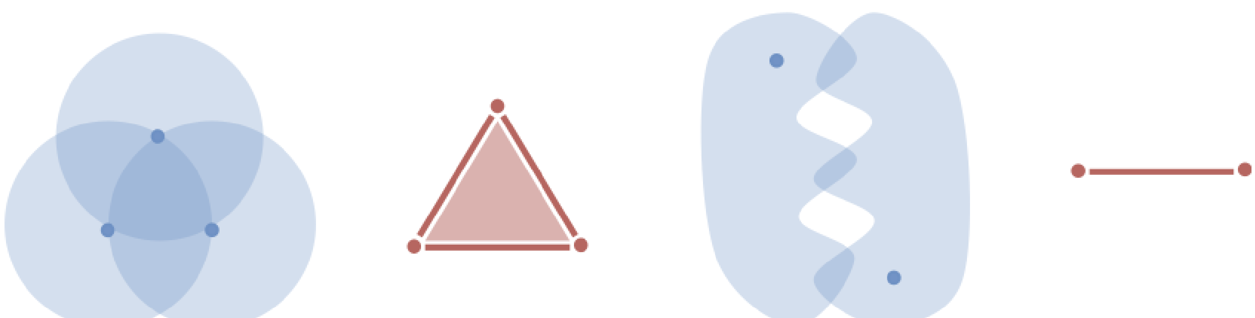

Figure 1: Covers and Their Nerves

In figure 1 we have drawn two different arrangements of open sets and their corresponding nerves, which we have represented graphically to the right. We have added points to each open set to make it clear how many open sets are in the cover. Note that in general, there is nothing to prevent a disconnected open set from being marked by a single label.

The nerve is purely an algebraic and combinatorial model for the cover – it need not respect the topology of the union. However, the **nerve lemma** of Karol Borsuk [14] states

[6] The eager reader is urged to consult [78, 26] for more modern devices associated to an open cover.

that if the intersections are contractible then the nerve and the union have the same homotopy type. The example on the left in figure 1 gives a positive example of the nerve lemma, whereas the example on the right gives a negative one.

The definition of a sheaf or cosheaf requires the synthesis of covers and data. We now introduce the functor that assigns data to open sets.

Definition 3.2 (Pre-Sheaf and Pre-Cosheaf)**.** A **pre-sheaf** is a functor $\mathsf{F} : \mathbf{Open}(\mathsf{X})^{\mathsf{op}} \to \mathcal{D}$ and a **pre-cosheaf** is a functor $\hat{\mathsf{F}} : \mathbf{Open}(\mathsf{X}) \to \mathcal{D}$. If $\mathsf{V} \subset \mathsf{U}$, then we usually write the **restriction map** as $\rho^{\mathsf{F}}_{\mathsf{V},\mathsf{U}} : \mathsf{F}(\mathsf{U}) \to \mathsf{F}(\mathsf{V})$ and the **extension map** as $r^{\hat{\mathsf{F}}}_{\mathsf{U},\mathsf{V}} : \hat{\mathsf{F}}(\mathsf{V}) \to \hat{\mathsf{F}}(\mathsf{U})$. Often we omit the superscript F or $\hat{\mathsf{F}}$.

If one imagines the pre-cosheaf that associates a copy of the field k to every connected component of an open set, then the following diagrams of vector spaces emerge from figure 1:

$$\begin{tikzcd}
 & \mathsf{k} & \\
\mathsf{k} \arrow[ur] \arrow[dd] & & \mathsf{k} \arrow[ul] \arrow[dd] \\
 & \mathsf{k} \arrow[uu] \arrow[ul] \arrow[ur] \arrow[dl] \arrow[dr] \arrow[dd] & \\
\mathsf{k} & & \mathsf{k} \\
 & \mathsf{k} \arrow[ul] \arrow[ur] &
\end{tikzcd}
\qquad\qquad
\mathsf{k} \longleftarrow \mathsf{k}^3 \longrightarrow \mathsf{k}$$

We will examine various ways for computing the colimits of these diagrams explicitly. Since the colimits occur over simplicial complexes, we introduce a structure theorem that allows us to use coequalizers. In the vector space case, this reduces to linear algebra – the colimit will be H_0 of a suitable chain complex.

We want to express the fact that since the colimit of a cover $\mathsf{N}(\mathcal{U}) \to \mathbf{Open}(\mathsf{X})$ is just the union $\mathsf{U} = \cup \mathsf{U}_i$, the data associated to U should be expressible as the colimit of data assigned to the nerve. Moreover, this should be independent of which cover we take.

Definition 3.3 (Sheaves and Cosheaves)**.** Suppose F is a pre-sheaf and $\hat{\mathsf{F}}$ is a pre-cosheaf, both of which are valued in $\mathcal{D}$. Suppose $\mathcal{U} = \{\mathsf{U}_i\}$ is an open cover of U. We say that F is a **sheaf on** $\mathcal{U}$ if the unique map from $\mathsf{F}(\mathsf{U})$ to the limit of $\mathsf{F} \circ \iota_{\mathcal{U}}^{\mathsf{op}}$, written

$$\mathsf{F}(\mathsf{U}) \to \varprojlim_{\mathsf{I} \in \mathsf{N}(\mathcal{U})} \mathsf{F}(\mathsf{U}_{\mathsf{I}}) =: \mathsf{F}[\mathcal{U}],$$

is an isomorphism. Similarly, we say $\hat{\mathsf{F}}$ is a **cosheaf on** $\mathcal{U}$ if the unique map from the colimit of $\hat{\mathsf{F}} \circ \iota_{\mathcal{U}}$ to $\hat{\mathsf{F}}(\mathsf{U})$, written

$$\hat{\mathsf{F}}[\mathcal{U}] := \varinjlim_{\mathsf{I} \in \mathsf{N}(\mathcal{U})} \hat{\mathsf{F}}(\mathsf{U}_{\mathsf{I}}) \to \hat{\mathsf{F}}(\mathsf{U}),$$

is an isomorphism. We say that F is a **sheaf** or $\hat{\mathsf{F}}$ is a **cosheaf** if for every open set U and every open cover $\mathcal{U}$ of U, $\mathsf{F}(\mathsf{U}) \to \mathsf{F}[\mathcal{U}]$ or $\hat{\mathsf{F}}[\mathcal{U}] \to \hat{\mathsf{F}}(\mathsf{U})$ is an isomorphism. For a catchy slogan, we say

On an open set (co)sheaves turn different covers into isomorphic (co)limits.

Remark 3.4 (Stable Under Finite Intersection). Most authors do not introduce the nerve as any part of the definition of a sheaf or cosheaf. Instead, some will require that the cover $\mathcal{U}$ is "stable under finite intersection," i.e. if $U_i, U_j \in \mathcal{U}$, then $U_i \cap U_j \in \mathcal{U}$. This allows those authors to just consider the limit or colimit over the cover and not over some auxiliary construction, like we have done. This works because one can take any cover and then add the intersections after the fact, but this tends to be done unconsciously and without any warning to the reader. Our approach is equivalent to that approach, but we believe it has some added benefits.

We have not stated any requirements on the data category $\mathcal{D}$, but in order to even parse the statement of the (co)sheaf axiom we require that the (co)limits coming from such covers exist. For the most part, we will work in categories where all limits and colimits exist. In analogy with analysis, a category where the limit of any diagram $F : I \to \mathcal{D}$ exists is called **complete**. Similarly, if the colimit of an arbitrary diagram exists, we say $\mathcal{D}$ is **co-complete**. The category $\boldsymbol{Vect}$ is both complete and co-complete.

A particular consequence of the axiom is that for a sheaf, $F(\emptyset)$ must be the limit over covers of the empty set, but since there there are no such covers, this is the limit over the empty diagram, i.e. $\boldsymbol{Cone}(\emptyset) = \mathcal{D}$, whose terminal object is the terminal object of $\mathcal{D}$. Similarly, for a cosheaf $\hat{F}(\emptyset)$ must be the initial object in $\mathcal{D}$. For $\mathcal{D} = \boldsymbol{Vect}$ the initial and terminal objects coincide with the zero vector space.

It is true that if $\mathcal{D}$ has pullbacks (see example 2.26 in section 2) and a terminal object then it has all finite limits. The dual statement that having an initial object and pushouts (see example 2.34) implies finitely co-complete is also true. Thus, if one focuses on sheaves and cosheaves valued in $\boldsymbol{vect}$ – the category of finite dimensional vector spaces and linear maps – then which covers $\mathcal{U}$ one considers needs to be modified. In particular, if the sheaf or cosheaf axiom holds for open covers with two sets, then we can only guarantee that it holds for covers with finitely many open sets. As a purely philosophical point, one wonders whether working with the cover of the complement of the Cantor set given by

$$\mathcal{U} = \{(\frac{3k+1}{3^n}, \frac{3k+2}{3^n}) \subset [0,1] | 0 \leqslant k \leqslant 3^{n-1} - 1,\ 0 \leqslant n < \infty\}$$

would ever be computationally tractable. One might wish to systematically revise the notion of a "cover," and this would lead to the notion of a **Grothendieck site**, which we do not address here.

We now examine the axioms just for covers with only two open sets.

Example 3.5 (Cover by Two Sets). Suppose $\mathcal{D} = \boldsymbol{Set}$, and suppose $\mathcal{U} = \{U_1, U_2\}$ is a cover of U. The sheaf condition says that

$$F(U) \cong \{(s_1, s_2) \in F(U_1) \prod F(U_2) | \rho_{U_{12},U_1}(s_1) = \rho_{U_{12},U_2}(s_2)\} =: F[\mathcal{U}],$$

i.e. $F(U)$ lists the set of consistent choices of elements from $F(U_1)$ and $F(U_2)$. In particular, $F[\mathcal{U}]$ is a sub-object of the product of $F(U_1)$ and $F(U_2)$. For an example, one can let F be

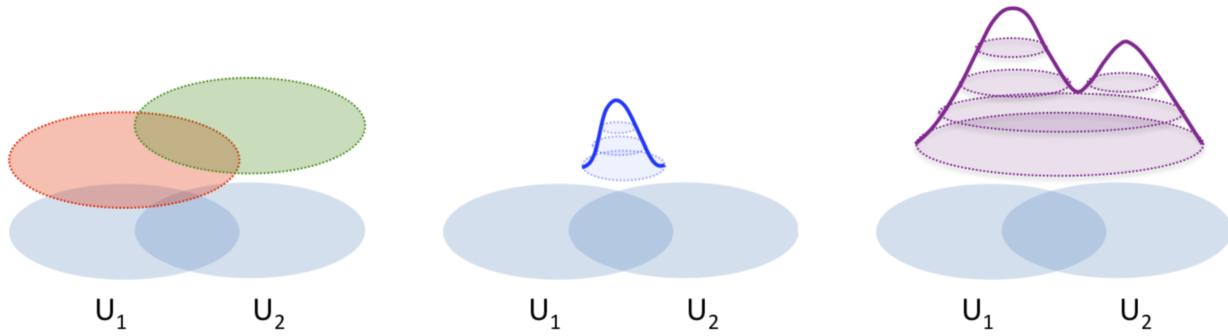

the assignment

$$U \rightsquigarrow \{f : U \to \mathbb{R} \mid \text{continuous}\}.$$

The sheaf axiom then says in order for two functions (or sections) $s_1 = f_1 : U_1 \to \mathbb{R}$ and $s_2 = f_2 : U_2 \to \mathbb{R}$ to determine an element in $U = U_1 \cup U_2$ it is necessary and sufficient that the functions $f_1(x)$ and $f_2(x)$ agree on the overlap $U_{12} = U_1 \cap U_2$.

The cosheaf condition for $\mathcal{D} = \mathbf{Set}$ is slightly strange. It says that

$$\hat{F}(U) \cong (\hat{F}(U_1) \coprod \hat{F}(U_2))/\sim \quad \text{where} \quad s_1 \sim s_2 \Leftrightarrow \exists s_{12} \quad s_1 = r_{U_1,U_{12}}(s_{12}) \quad s_2 = r_{U_2,U_{12}}(s_{12}).$$

In contrast to the sheaf case, the notion of consistent choices no longer applies for cosheaves, because it requires thinking in terms of quotient objects – something human beings are not accustomed to. However, a useful analogy is that one must subtract out or identify those elements that might be counted twice because they come from the intersection. For an example similar in spirit to the sheaf of real-valued functions, we begin by considering the pre-cosheaf of compactly supported functions gotten by assigning

$$U \rightsquigarrow \{f : U \to \mathbb{R} \mid \text{continuous and compactly supported}\}.$$

Extending by zero provides the extension map and identifying the two copies of a function whose support is contained in $U_{12} = U_1 \cap U_2$ prevents double counting on U. However, this is not all that the cosheaf axiom requires. Any compactly supported function should appear as one supported in U_1 or U_2, but this is not always true. Some compactly supported functions are not compact when restricted to any particular open set in a cover. Thus, this pre-cosheaf is not a cosheaf.

The reader familiar with partitions of unity will realize that if X a paracompact Hausdorff space then we can express any compactly supported function $f(x)$ defined on all of U as a *sum* of compactly supported functions on U_1 and U_2. By taking a partition of unity subordinate to the cover $\mathcal{U}$ we get two functions $\lambda_1(x)$ and $\lambda_2(x)$ such that

$$f(x) = f_1(x) + f_2(x) \qquad \text{where} \qquad f_1(x) := \lambda_1(x)f(x) \qquad \text{and} \qquad f_2(x) := \lambda_2(x)f(x).$$

By carrying out the colimit in a data category equipped with sums, such as $\mathcal{D} = \mathbf{Vect}$ of $\mathbf{Ab}$, then compactly supported functions do define a cosheaf valued there. Specifically, If $\mathcal{D} = \mathbf{Vect}$, then the cosheaf axiom for the cover says the sequence

$$\hat{F}(U_{12}) \to \hat{F}(U_1) \oplus \hat{F}(U_2) \to \hat{F}(U) \to 0$$

is exact, where the maps are $(-r_{U_1,U_{12}}, r_{U_2,U_{12}})$ and $r_{U,U_2} + r_{U,U_1}$. Dually, the sheaf axiom says the dual sequence

$$0 \to F(U) \to F(U_1) \times F(U_2) \to F(U_{12})$$

is exact, where the second map is $\rho_{U_{12},U_2} + \rho_{U_{12},U_1}$ and the first map is $(-\rho_{U_1,U}, \rho_{U_2,U})$.

3.2 Limits and Colimits over Covers: a Structure Theorem

The sheaf and cosheaf axioms as stated are meant to emphasize that if one is comfortable with the operations of limits and colimits, then one is already comfortable with sheaves and cosheaves. However, the limits and colimits considered in Definition 3.3 have a special structure. This structure comes from the fact that the indexing category – the nerve – is a simplicial complex.

The first observation one can make is that for any functor $F : N(\mathcal{U})^{op} \to \mathcal{D}$ the limit can be thought of as "sitting inside" the product over the vertices – the vertices corresponding to the elements of the cover through the nerve construction. Dually, the colimit of a functor $\hat{F} : N(\mathcal{U}) \to \mathcal{D}$ can be thought of as a quotient of the coproduct of the functor over the vertices. Said using formulas, this is

$$\varprojlim F \to \prod F(i) \qquad \coprod \hat{F}(i) \to \varinjlim \hat{F}\,.$$

The way to see this is to note that any cone or cocone's morphism must factor through a vertex. However, the difference between the limit or colimit from the functor's aggregate value on vertices is measured by edges in the nerve. This is a reflection of a more general theorem, which we now state.

Theorem 3.6. *A category $\mathcal{D}$ has all (co)limits of an appropriate size if it has all (co)products and (co)equalizers of same such size. Here "size" corresponds to the cardinality of the indexing category of the (co)limit in question.*

Proof Idea. One should consult [13] Prop. 5.22 and 5.23 for a complete proof. To give the reader the idea, one can compute the limit of $F : I \to \mathcal{D}$ by taking the product over all the objects $x \in I$ and separately the product over all morphisms in the indexing category I. The limit is isomorphic to the equalizer going from the first product to the latter, i.e.

$$\varprojlim F \longrightarrow \prod_{x \in I} F(x) \rightrightarrows \prod_{x \to x'} F(x')\,.$$

By dualizing, one can prove the analogous result for colimits. □

This theorem gives us effective means for computing limits and colimits for general data types. We now specialize this result to the limits and colimits pertinent to sheaves and cosheaves.

3.2.1 Rephrased as Equalizers or Co-equalizers

The method outlined in theorem 3.6 for computing limits and colimits contains too much redundant information for the case $I = N(\mathcal{U})^{op}$. As such, we state the precise, simplified formulation here. The sheaf and cosheaf axioms can be rephrased as saying that the following sequences

$$F(U) \xrightarrow{e} \prod F(U_i) \overset{f^-}{\underset{f^+}{\rightrightarrows}} \prod_{i<j} F(U_i \cap U_j) \qquad \coprod_{i<j} \hat{F}(U_i \cap U_j) \overset{g^+}{\underset{g^-}{\rightrightarrows}} \coprod \hat{F}(U_i) \xrightarrow{u} \hat{F}(U)$$

are an equalizer and a co-equalizer respectively.

To describe the maps explicitly requires some work. First, we choose an ordering of the indexing set of the cover $\mathcal{U} = \{U_i\}_{i\in\Lambda}$. To specify a map to a product it suffices to specify maps to each factor of the product. Similarly, maps from a coproduct are specified by maps from each factor. This is summarized by the identities

$$\mathbf{Hom}(X, \prod_i Y_i) \cong \prod_i \mathbf{Hom}(X, Y_i) \qquad \text{and} \qquad \mathbf{Hom}(\coprod_i X_i, Y) \cong \prod_i \mathbf{Hom}_i(X_i, Y).$$

To define the maps e and u we declare $e_i := \rho_{U_i,U}$ and $u_i := r_{U,U_i}$. For the maps $f^{\pm}$ and $g^{\pm}$ we define for each pair $i < j$ the maps

$$f^+_{ij} := \rho_{ij,j} \circ \pi_j \quad f^-_{ij} := \rho_{ij,i} \circ \pi_i \qquad g^+_{ij} := r_{j,ij} \circ \iota_{ij} \quad g^-_{ij} := r_{i,ij} \circ \iota_{ij}$$

where $\pi_i : \prod F(U_i) \to F(U_i)$ is the natural projection and $\iota_{ij} : \hat{F}(U_{ij}) \to \coprod \hat{F}(U_{ij})$ is the natural inclusion.

The reader might find it helpful to think of the maps in between the products as being represented by matrices. In the case of a cover with three elements $\mathcal{U} = \{U_1, U_2, U_3\}$ all of whose pairwise intersections are non-empty, we can write

$$f^+ = \begin{bmatrix} * & \rho_{12,2} & * \\ * & * & \rho_{13,3} \\ * & * & \rho_{23,3} \end{bmatrix} \quad f^- = \begin{bmatrix} \rho_{12,1} & * & * \\ \rho_{13,1} & * & * \\ * & \rho_{23,2} & * \end{bmatrix}.$$

The equalizer condition now reads that $f^+(s_1, s_2, s_3) = f^-(s_1, s_2, s_3)$, i.e.

$$(\rho_{12,2}(s_2), \rho_{13,3}(s_3), \rho_{23,3}(s_3)) = (\rho_{12,1}(s_1), \rho_{13,1}(s_1), \rho_{23,2}(s_2)).$$

3.2.2 Rephrased as Exactness

If $\mathcal{D} = \mathbf{Vect}$, then we can add and subtract maps and look for kernels and cokernels instead of equalizers and co-equalizers. The sheaf and cosheaf axioms then reduce to linear algebra. The modified axioms now read as

$$0 \longrightarrow F(U) \longrightarrow \prod F(U_i) \xrightarrow{d^0} \prod_{i<j} F(U_i \cap U_j)$$

$$\bigoplus_{i<j} \hat{F}(U_i \cap U_j) \xrightarrow{\partial_1} \bigoplus \hat{F}(U_i) \longrightarrow \hat{F}(U) \longrightarrow 0$$

where d^0 is the matrix whose rows are parametrized by pairs $i < j$ and whose columns are parametrized by k with entries given by $d^0_{ij,k} = [k : ij]\rho_{ij,k}$ where

$$[k : ij] = \begin{cases} 0 & \text{if } k \neq i \neq j \\ 1 & \text{if } k = j \\ -1 & \text{if } k = i \end{cases}$$

The matrix ∂_1 is similarly defined except that the rows are indexed by k and columns are indexed by pairs $i < j$ with entries $(\partial_1)_{k,ij} = [k : ij]r_{k,ij}$. Thus the sheaf axiom says that $F(U) \cong \ker(d^0)$ and the cosheaf axiom says that $\hat{F}(U) \cong \operatorname{coker}(\partial_1)$.

In our example of a three set cover $\mathcal{U} = \{U_1, U_2, U_3\}$ all of whose pairwise intersections are non-empty, the definition of d^0 corresponds to taking $f^+ - f^-$, i.e.

$$d^0 = f^+ - f^- = \begin{bmatrix} -\rho_{12,1} & \rho_{12,2} & 0 \\ -\rho_{13,1} & 0 & \rho_{13,3} \\ 0 & -\rho_{23,2} & \rho_{23,3} \end{bmatrix}$$

where each of the $\rho_{ij,k}$'s need to be filled in with some matrix representative of that linear map. The kernel is then identified with $F[\mathcal{U}]$.

3.3 Čech Homology and Cosheaves

In section 3.2.1 we rephrased the limits and colimits coming from covers as equalizers and coequalizers. For the data category $\mathcal{D} = \mathbf{Vect}$ we showed how to reinterpret this as an exact sequence. This perspective is indicative of a deeper and more computational idea, namely that of homology. We now show how to associate to any pre-cosheaf[7] of vector spaces $\hat{F}$ and an open cover $\mathcal{U} = \{U_i\}_{i\in\Lambda}$ a complex of vector spaces whose zeroth homology computes $\hat{F}[\mathcal{U}]$. This allows us to compute the homology of data.

Definition 3.7 (Čech Homology)**.** Given a pre-cosheaf of vector spaces $\hat{F}$ and an open cover $\mathcal{U} = \{U_i\}_{i\in\Lambda}$, we define the **Čech homology on** $\mathcal{U}$ to be the homology of the complex

$$(\check{C}_\bullet(\mathcal{U};\hat{F}), \partial_\bullet) \qquad \text{where} \qquad \check{C}_p(\mathcal{U};\hat{F}) := \bigoplus_{|I|=p+1} \hat{F}(U_I) \qquad \text{for} \qquad I \in N(\mathcal{U}).$$

By choosing an ordering on the index set Λ, we define the differential by extending the formula defined on elements $s_I \in \hat{F}(U_I)$ by linearity, i.e.

$$\partial_p : C_p(\mathcal{U};\hat{F}) \to C_{p-1}(\mathcal{U};\hat{F}) \qquad \partial_p(s_I) := \sum_{k=0}^{p} (-1)^k r_{U_I^{(k)},U_I}(s_I),$$

[7] Or pre-sheaf, but we'll leave it to the reader to dualize.

where the symbol $U_I^{(k)} = U_{i_0} \cap \ldots \cap U_{i_{k-1}} \cap U_{i_{k+1}} \cap \ldots U_{i_p}$ indicates the intersection that omits the kth open set. Thus we can define by the usual formula the pth **Čech homology group**

$$\check{H}_p(\mathcal{U}; \hat{F}) := \frac{\ker \partial_p}{\operatorname{im} \partial_{p+1}} \qquad \text{i.e.} \qquad H_p(\check{C}_\bullet(\mathcal{U}; \hat{F})).$$

To guarantee that Čech homology is well-defined we verify the following lemma:

Lemma 3.8. *The differential ∂ in the Čech complex for a cover $\mathcal{U}$ and a pre-cosheaf $\hat{F}$ of vector spaces satisfies $\partial_p \circ \partial_{p+1} = 0$.*

Proof. The combinatorial nature of the nerve of a cover guarantees that $\partial^2 = 0$. Specifically, there are two ways of going between incident simplices of dimension differing by two. Thus, we get the following diagram of open sets and data:

$$\begin{tikzcd}
& U_I^{(j,k)} & \\
U_I^{(j)} \arrow[ur] & & U_I^{(k)} \arrow[ul] \\
& U_I \arrow[ul] \arrow[ur] \arrow[uu] &
\end{tikzcd}
\qquad\qquad
\begin{tikzcd}
& \hat{F}(U_I^{(j,k)}) & \\
\hat{F}(U_I^{(j)}) \arrow[ur] & & \hat{F}(U_I^{(k)}) \arrow[ul] \\
& \hat{F}(U_I) \arrow[ul] \arrow[ur] \arrow[uu] &
\end{tikzcd}$$

Let's follow a typical element $s_I \in \hat{F}(U_I)$ through the diagram on the right upon applying the formula $\partial \circ \partial$. First note that the fact that $\hat{F}$ is a pre-cosheaf implies that the square commutes, i.e.

$$r_{U_I^{(j,k)}, U_I^{(j)}} \circ r_{U_I^{(j)}, U_I}(s_I) = r_{U_I^{(j,k)}, U_I^{(k)}} \circ r_{U_I^{(k)}, U_I}(s_I) = r_{U_I^{(j,k)}, U_I}(s_I).$$

The first application of ∂ yields $(-1)^j r_{U_I^{(j)}, U_I}(s_I)$ and $(-1)^k r_{U_I^{(k)}, U_I}(s_I)$ as just two components in the formula for $\partial(s_I)$. Assuming $j < k$ and applying the definition of the boundary map to elements in $\hat{F}(U_I^{(j)})$ implies that we must actually delete the $k-1$st entry of $I - \{j\}$ since removing j has caused everything above j to shift down in the ordered list. Thus the image of $\partial^2(s_I)$ in $\hat{F}(U_I^{(j,k)})$ is

$$(-1)^{k-1}(-1)^j r_{U_I^{(j,k)}, U_I}(s_I) + (-1)^k(-1)^j r_{U_I^{(j,k)}, U_I}(s_I) = 0.$$

$\square$

Example 3.9. Consider the covers in figure 1. The pre-cosheaf we described there assigned to each connected component of an open set a copy of the field k. First we consider the cover on the left of figure 1 with three open sets. We label the three vertices of the nerve, starting with the bottom left one and working counter-clockwise, x, y and z respectively. The Čech complex takes the form

$$k_{xyz} \xrightarrow{\partial_2} k_{xy} \oplus k_{xz} \oplus k_{yz} \xrightarrow{\partial_1} k_x \oplus k_y \oplus k_z \longrightarrow 0$$

where, using the lexicographic ordering for a basis, the matrix representatives for ∂_2 and ∂_1 take the following form:

$$\partial_2 = \begin{bmatrix} 1 \\ -1 \\ 1 \end{bmatrix} \qquad \partial_1 = \begin{bmatrix} -1 & -1 & 0 \\ 1 & 0 & -1 \\ 0 & 1 & 1 \end{bmatrix}$$

One can easily verify that the $\ker \partial_1 = \operatorname{im} \partial_2$ and consequently $\check{H}_1 = 0$. Furthermore, $\check{H}_0 \cong k$, which happens to reflect that the union has one connected component. Similarly, one can consider the cover at the right of figure 1. The Čech complex for this cover and the same pre-cosheaf is as follows:

$$k^3 \xrightarrow{\partial_1} k^2 \longrightarrow 0 \qquad \text{where} \qquad \partial_1 = \begin{bmatrix} -1 & -1 & -1 \\ 1 & 1 & 1 \end{bmatrix}$$

Clearly $\check{H}_0 \cong k$, whose dimension agrees with the number of connected components of the union, but also $\check{H}_1 \cong k^2$, which witnesses the presence of two holes in the union.

One can dually define Čech cohomology with coefficients valued in a pre-sheaf F. The discussion of section 3.2.2, along with the examples just presented, can be interpreted as saying a pre-sheaf F or pre-cosheaf $\hat{F}$ is a sheaf or cosheaf if and only if

$$F(U) \cong \check{H}^0(\mathcal{U}; F) \qquad \text{or} \qquad \check{H}_0(\mathcal{U}; \hat{F}) \cong \hat{F}(U).$$

for any choice of cover $\mathcal{U}$ of U.

We would like to use this isomorphism to supply examples of sheaves and cosheaves from standard machinery in algebraic topology. Suppose one has an independent notion of homology, such as singular homology, and one can show it is isomorphic to Čech homology for suitably fine covers (see section 3.4 to see why fineness matters) on suitably nice spaces, then one could define a cosheaf using those values. To make this rigorous, and to also provide a useful criterion for proving when a pre-cosheaf is a cosheaf, we recall a theorem:

Theorem 3.10. *Suppose $\hat{F}$ is a pre-cosheaf, then $\hat{F}$ is a cosheaf if and only if the following the following two properties hold*

- *For all open sets U and V the following sequence is exact*

$$\hat{F}(U \cap V) \to \hat{F}(U) \oplus \hat{F}(V) \to \hat{F}(U \cup V) \to 0.$$

  The first morphism is $(-r_{U,U\cap V}, r_{V,U\cap V})$ and the second is $r_{U\cup V,U} + r_{U\cup V,V}$.

- *If $\{U_\alpha\}$ is directed upwards by inclusion, i.e. for for every pair U_α and U_β there exists U_γ containing both, then the canonical map*

$$\varinjlim_{\alpha} \hat{F}(U_\alpha) \to \hat{F}(\cup U_\alpha)$$

  is an isomorphism.

Dually, turning arrows around and using inverse limits gives a useful criterion for determining when a pre-sheaf is a sheaf.

Proof. Using induction one can prove that the cosheaf property for two sets implies the cosheaf property for finitely many sets (see [16] p. 418 for a proof). We now show that this implies the cosheaf axiom for arbitrary covers. Suppose $\{U_\alpha\}_{\alpha\in\Lambda}$ is a cover indexed by a potentially large, but ordered set Λ. For each finite subset $I \subset \Lambda$ we know that

$$\bigoplus_{\alpha<\beta\in I} \hat{F}(U_{\alpha,\beta}) \to \bigoplus_{\alpha\in I} \hat{F}(U_\alpha) \to \hat{F}(\bigcup_{\alpha\in I} U_\alpha) \to 0$$

is exact. We know that the collection of finite subsets I forms a directed system and that in **Vect** direct limits preserve exactness. As such we have that

$$\varinjlim_{I} \bigoplus_{\alpha<\beta\in I} \hat{F}(U_{\alpha,\beta}) \to \varinjlim_{I} \bigoplus_{\alpha\in I} \hat{F}(U_\alpha) \to \varinjlim_{I} \hat{F}(\bigcup_{\alpha\in I} U_\alpha) \to 0$$

is exact as well, but by using the second property and the fact that the direct limit of the I's is Λ we have

$$\bigoplus_{\alpha<\beta\in\Lambda} \hat{F}(U_{\alpha,\beta}) \to \bigoplus_{\alpha\in\Lambda} \hat{F}(U_\alpha) \to \hat{F}(\bigcup_{\alpha\in\Lambda} U_\alpha) \to 0$$

is exact. This proves the reverse direction. The other direction is clear. □

This theorem then provides us with a useful example of a cosheaf that we have implicitly used to generate examples. We now make this example explicit.

Example 3.11. The assignment to an open set U the 0th singular homology of U

$$U \rightsquigarrow H_0(U;k)$$

is a cosheaf. This follows from the fact that the singular chain complex (see later for a definition) $C_\bullet(-;k)$ can be defined for any subset U of X and homology commutes with direct limits, thus the second property of the theorem holds. The first property in the theorem follows from exactness at the last two spots in the Mayer-Vietoris sequence:

$$\begin{tikzcd} H_1(U\cap V;k) \arrow[r] & H_1(U;k)\oplus H_1(V;k) \arrow[r] & H_1(U\cup V;k) \arrow[dll, out=0, in=180] & \\ H_0(U\cap V;F) \arrow[r] & H_0(U;k)\oplus H_0(V;k) \arrow[r] & H_0(U\cup V;k) \arrow[r] & 0 \end{tikzcd}$$

The moral from this example is that, in essence,

Any functor that satisfies Mayer-Vietoris is a cosheaf.

3.4 Refinement of Covers

We have defined the sheaf and cosheaf axioms for a cover $\mathcal{U}$. The coarsest possible cover of an open set U is the cover with one element $\{U\}$. Thus, one way of interpreting the sheaf and cosheaf axiom is that $F[\mathcal{U}]$ and $\hat{F}[\mathcal{U}]$ are independent of the cover chosen. A logical question to ask is if the axiom holds for some cover, but not all, then for what other covers does the axiom hold? To answer this question, we review some relevant concepts.

Definition 3.12 (Refinement of Covers)**.** Suppose $\mathcal{U}$ and $\mathcal{U}'$ are covers of U, then we say that $\mathcal{U}'$ **refines** $\mathcal{U}$ if for every $U_i' \in \mathcal{U}'$ there is a $U_j \in \mathcal{U}$ and an inclusion $U_i' \to U_j$. Note that every cover refines the trivial cover $\{U\}$.

Definition 3.13. The refinement relation endows the collection of covers of U with the structure of a category $\mathbf{Cov}(U)$, whose objects are covers $\mathcal{U}$ with a unique morphism $\mathcal{U}' \to \mathcal{U}$ if the former refines the latter.

Note that if $U_{i_1}' \to U_{j_1}$ and $U_{i_2}' \to U_{j_2}$, then $U_{i_1}' \cap U_{i_2}' \to U_{j_1} \cap U_{j_2}$. So a refinement induces a functor between nerves, but it depends on which inclusions were chosen.

$$\begin{tikzcd} & \mathbf{Open}(X) & \\ N(\mathcal{U}') \arrow[ur] \arrow[rr] & & N(\mathcal{U}) \arrow[ul] \end{tikzcd} \qquad \begin{tikzcd} & \mathbf{Open}(X)^{op} & \\ N(\mathcal{U}')^{op} \arrow[ur] & & N(\mathcal{U})^{op} \arrow[ul] \arrow[ll] \end{tikzcd}$$

The next lemma shows that these choices don't matter on the level of limits and colimits for pre-sheaves and pre-cosheaves.

Lemma 3.14. *Let $\hat{F}$ and F be a pre-cosheaf and a pre-sheaf respectively. Suppose $\mathcal{U}'$ refines another cover $\mathcal{U}$ of an open set U. Then there are well-defined maps*

$$\hat{F}[\mathcal{U}'] \to \hat{F}[\mathcal{U}] \qquad \text{and} \qquad F[\mathcal{U}] \to F[\mathcal{U}'],$$

i.e. we get functors $\hat{F} : \mathbf{Cov}(U) \to \mathcal{D}$ and $F : \mathbf{Cov}(U)^{op} \to \mathcal{D}$.

Proof. We'll detail the proof for a pre-cosheaf $\hat{F}$ since the case for pre-sheaves can be found in the literature or obtained here via dualizing appropriately. A refinement $\mathcal{U}' \to \mathcal{U}$ defines a natural transformation $\hat{F} \circ \iota_{\mathcal{U}'} \Rightarrow \hat{F} \circ \iota_{\mathcal{U}}$. The colimit defines a natural transformation from $\hat{F} \circ \iota_{\mathcal{U}}$ to the constant diagram $\hat{F}[\mathcal{U}]$. Since the composition of natural transformations is a natural transformation, this induces a cocone $\hat{F} \circ \iota_{\mathcal{U}'} \Rightarrow F[\mathcal{U}]$ which, by the universal property of the colimit, defines a unique induced map there, i.e.

$$\hat{F} \circ \iota_{\mathcal{U}'} \Rightarrow \hat{F} \circ \iota_{\mathcal{U}} \Rightarrow F[\mathcal{U}] \qquad \text{implies} \qquad \exists!\, \hat{F}[\mathcal{U}'] \to \hat{F}[\mathcal{U}].$$

However, if in choosing the inclusions for the refinement we had made a different set of choices, $U_i' \to U_k$ rather than U_j, then a priori we might have expected different maps $\hat{F}[\mathcal{U}'] \to \hat{F}[\mathcal{U}]$. Let us show this choice does not matter. If there is a choice, then we can take $U_i' \to U_j \cap U_k$ as a common refinement. As a consequence of $\hat{F}$ being a functor from the open set category, the different maps to the colimit must agree, as they factor through whatever is assigned on the intersection, i.e. the following diagram commutes:

$$\begin{tikzcd}
 & & \hat{F}(U_j) \arrow[d] \\
\hat{F}(U_i') \arrow[r] \arrow[urr] \arrow[drr] & \hat{F}(U_j \cap U_k) \arrow[r] \arrow[ur] \arrow[dr] & \hat{F}[\mathcal{U}] \\
 & & \hat{F}(U_k) \arrow[u]
\end{tikzcd}$$

□

Corollary 3.15. *If $\hat{F}$ is a cosheaf or F is a sheaf for the cover $\mathcal{U}'$, then it is a cosheaf or sheaf for every cover it refines.*

Proof. Suppose we have a series of refinements

$$\mathcal{U}' \to \mathcal{U} \to \{U\}.$$

To say that $\hat{F}$ or F is cosheaf or sheaf for $\mathcal{U}'$ is to say that the following induced maps are isomorphisms:

$$\begin{tikzcd}
\hat{F}[\mathcal{U}'] \arrow[r] \arrow[rr, bend left, "\cong"] & \hat{F}[\mathcal{U}] \arrow[r] & \hat{F}(U)
\end{tikzcd}
\qquad
\begin{tikzcd}
F(U) \arrow[r] \arrow[rr, bend left, "\cong"] & \hat{F}[\mathcal{U}] \arrow[r] & \hat{F}[\mathcal{U}']
\end{tikzcd}$$

However, by functoriality, the factored maps must themselves be isomorphisms, i.e.

$$\hat{F}[\mathcal{U}] \xrightarrow{\cong} \hat{F}(U) \qquad F(U) \xrightarrow{\cong} F[\mathcal{U}] \; .$$

□

We will make use of this corollary as we begin to consider sheaves and cosheaves on spaces where there is a finest cover. Checking the sheaf or cosheaf axiom there then guarantees it for all covers.

3.5 Generalities on Sheaves and Cosheaves

Sheaves have proved to be highly successful tools in pure mathematics over the past 60-70 years. However, this cannot be said for cosheaves per se. One reason is that some have dualized the concept of a cosheaf to bring it under the technical umbrella of sheaf theory. In particular, we have noted that the data of specifying a functor $\hat{F} : \mathbf{Open}(X) \to \mathcal{D}$ is exactly equivalent to specifying a functor $F : \mathbf{Open}(X)^{op} \to \mathcal{D}^{op}$. As such, one will sometimes see cosheaves defined as sheaves valued in the opposite category. This perspective has certain deficiencies that are worth pointing out. By highlighting these deficiencies, we attempt to better justify the independent study of cosheaves.

First and foremost one should recall that most data categories $\mathcal{D}$, such as $\mathbf{Set}, \mathbf{Vect}$ or $\mathbf{Ab}$, are not equivalent to their opposite categories. Thus the topological simplification of reducing cosheaves to sheaves comes at the cost of making the algebraic thinking more difficult. In particular, certain properties of $\mathcal{D} = \mathbf{Set}, \mathbf{Vect}$ or $\mathbf{Ab}$ are used in the development of sheaf theory, which do not necessarily hold in $\mathcal{D}^{op}$. The centerpiece of this discussion will be understanding that filtered colimits commute with finite limits in $\mathcal{D}$, but cofiltered limits do not necessarily commute with finite colimits in $\mathcal{D}$. Let us now relay the necessary definitions.

Definition 3.16. A non-empty category $\mathcal{C}$ is called **filtered** if the following two properties are satisfied:

- For every pair of objects $x, y \in \mathcal{C}$ there is a third object $z \in \mathcal{C}$ with $x \to z$ and $y \to z$.
- For every pair of parallel morphisms $f, g : x \to y$ there is a third object and morphism $h : y \to z$ such that $h \circ f = h \circ g$.

A category $\mathcal{C}$ is called **cofiltered** if $\mathcal{C}^{op}$ is filtered. Sometimes, when the category is especially simple, we will simply call a cofiltered category filtered.

Example 3.17. In section 3.4 we considered the category of covers $\mathbf{Cov}(U)$ of an open set U. By noting that any two covers have a common refinement, one sees that this is an example of a cofiltered (or *cofiltrant*) category.

Definition 3.18. Suppose I is a filtered indexing category with $F : I \to \mathcal{D}$ and $G : I^{op} \to \mathcal{D}$ diagrams in some category. We will call the colimit of F a **filtered colimit** and the limit of G a **cofiltered limit**.

Now we give an example naturally occurring in sheaf theory.

Example 3.19 ((Co)Stalks)**.** Suppose X is a topological space and x is a point in X. The set of open sets containing x defines a cofiltered subcategory of $\mathbf{Open}(X)$ or a filtered subcategory of $\mathbf{Open}(X)^{op}$. Consequently for a pre-sheaf F or a pre-cosheaf $\hat{F}$, the following

$$F_x := \varinjlim_{U \ni x} F(U) \qquad \varprojlim_{U \ni x} \hat{F}(U) =: \hat{F}_x$$

define a filtered colimit and cofiltered limit, respectively. These are called the **stalk at x** of a pre-sheaf F and the **costalk at x** of a pre-cosheaf $\hat{F}$. Of course, to make such a statement meaningful, one needs to assume the data category $\mathcal{D}$ has the relevant limits and colimits.

The following theorem illustrates one of the fundamental differences between sheaves and cosheaves. It is expressed through the following algebraic fact, which the reader might like to compare with the Fubini theorem.

Theorem 3.20. *Let* I *be a filtered indexing category and* J *a finite category. Then any functor* $\alpha : I \times J \to \mathcal{D}$ *where* $\mathcal{D} = \mathbf{Set}, \mathbf{Vect}$, *or* $\mathbf{Ab}$, *has the property that the natural map*

$$\varinjlim_I \varprojlim_J \alpha(i,j) \to \varprojlim_J \varinjlim_I \alpha(i,j)$$

is an isomorphism. We say for short that "filtered colimits and finite limits commute" in these categories.

Proof. For a proof of this statement, we refer the reader to theorem 3.1.6 in [46]. First note that the product category $I \times J$ is just a product in $\mathbf{Cat}$ - the category of all categories. We can describe the objects of $I \times J$ as pairs of objects, one from each category, and the morphisms as tuples of morphisms, one for each object. One can take α and then define a new functor $\varprojlim_J \alpha : I \to \mathcal{D}$ gotten by assigning to each object $i \in \mathbf{obj}(I)$ the limit over J of $\alpha(i,-) : J \to \mathcal{D}$. Taking the colimit over I defines the first expression. Similar reasoning defines the second. $\square$

For an application of this theorem we introduce some ideas to the world of pre-sheaves valued in $\mathbf{Vect}$. Recall that a complex of vector spaces

$$\cdots \to V_1 \to V_2 \to V_3 \to \cdots$$

is exact at a term in a sequence if the image of the incoming map coincides with the kernel of the outgoing. A sequence of pre-sheaves[8] is exact if and only if for each open set U the associated sequence of vector spaces is exact, i.e.

$$0 \to E \to F \to G \to 0 \qquad \text{iff} \qquad 0 \to E(U) \to F(U) \to G(U) \to 0$$

Theorem 3.20 then implies that for any point $x \in X$ the induced sequence of stalks

$$0 \to E_x \to F_x \to G_x \to 0$$

is exact. Intuitively this is because we can view E as a kernel of the pre-sheaf map $F \to G$ and as already demonstrated, kernels are examples of finite limits. Thus taking the kernel of the stalk map $F_x \to G_x$ is the same as taking the stalk of the kernel of $F \to G$.

Proposition 3.21. *For* $\mathcal{D} = \mathbf{Set}, \mathbf{Vect}$ *or* $\mathbf{Ab}$ *it is not true that cofiltered limits and finite colimits commute. Consequently, if* $A, B, C : \mathbb{N}^{op} \to \mathbf{Ab}$ *(or* $\mathbf{Vect}$*) are functors from the category of natural numbers equipped with the opposite ordering, with natural transformations* $A \to B \to C$ *such that*

$$0 \longrightarrow A_i \longrightarrow B_i \longrightarrow C_i \longrightarrow 0$$

is exact for every i*, then it is* not always *the case that the induced sequence on limits is exact.*

$$0 \longrightarrow \varprojlim A \longrightarrow \varprojlim B \longrightarrow \varprojlim C \longrightarrow 0$$

[8]The corresponding statement for sheaves is *not* true.

Proof. We borrow an example from Jason McCarthy's notes [61]. Consider the following system of short exact sequences of groups:

$$\begin{tikzcd}
\vdots \arrow[d] & \vdots \arrow[d] & \vdots \arrow[d] & \vdots \arrow[d] & \vdots \arrow[d] \\
0 \arrow[r] \arrow[d] & \mathbb{Z} \arrow[r, "n"] \arrow[d, "n+1"'] & \mathbb{Z} \arrow[r] \arrow[d, "n+1"'] & \mathbb{Z}/n \arrow[r] \arrow[d, "\mathrm{id}"'] & 0 \arrow[d] \\
0 \arrow[r] \arrow[d] & \mathbb{Z} \arrow[r, "n"] \arrow[d, "n+1"'] & \mathbb{Z} \arrow[r] \arrow[d, "n+1"'] & \mathbb{Z}/n \arrow[r] \arrow[d, "\mathrm{id}"'] & 0 \arrow[d] \\
0 \arrow[r] & \mathbb{Z} \arrow[r, "n"] & \mathbb{Z} \arrow[r] & \mathbb{Z}/n \arrow[r] & 0
\end{tikzcd}$$

The inverse limit of the first (and second) column with non-zero entries must be zero. To see why, note that the inverse limit can be described as

$$\varprojlim_{\mathbb{N}^{op}} \mathbb{Z}_i = \{(x_i) \in \prod_i \mathbb{Z}_i \mid x_i = (n+1)^{j-i} x_j \; \forall j \geqslant i\},$$

where we have viewed the indexing category as the natural numbers with the opposite ordering. Any non-zero element of the limit would have some non-zero factor x_i and consequently all other factors would be non-zero (since the map $a \to (n+1)a$ is injective). In particular, all higher x_j must be equal to $x_i/(n+1)^{j-i}$, but letting j be suitably large would imply that x_j must be less than one – an impossibility. Thus the induced map of inverse limits is

$$0 \longrightarrow 0 \longrightarrow 0 \longrightarrow \mathbb{Z}/n\mathbb{Z} \longrightarrow 0$$

which is not exact. If the reader prefers an example in the category of vector spaces, one should see Schapira's example 4.2.5 in his notes [72]. □

Thus the statement that short exact sequences of pre-cosheaves induces a short exact sequence on costalks cannot be guaranteed. There is a subtle work-around that says under suitable hypotheses[9] exactness can be guaranteed. This holds for categories like $\boldsymbol{vect}$, the category of finite-dimensional vector spaces, and **ab**, the category of finite abelian groups, because this is where the descending chain condition holds [10].

This last comment about $\boldsymbol{vect}$ provides justification for performing some dualization to obtain results about cosheaves from sheaves. After all, for finite-dimensional vector spaces it is true that

$$\mathbf{Hom}_{\boldsymbol{vect}}(-, k) : \boldsymbol{vect}^{op} \to \boldsymbol{vect}$$

establishes an equivalence of categories.[10] However, issues of stalks versus costalks is not the primary obstacle that the asymmetry of theorem 3.20 presents. That obstacle has to

[9] i.e. the Mittag-Leffler condition. See [45] section 1.12. or [12] pp. 211-214 for more details.

[10] This does *not* extend to an equivalence between $\boldsymbol{Vect}$ and its opposite category. In fact, $\boldsymbol{Vect}^{op}$ is equivalent to the category pro $-\boldsymbol{vect}$, cf. Remark 6.2 of [44].

do with a process known as **sheafification**, which provides a universal tool for turning any pre-sheaf into a sheaf. For most texts on sheaf theory it is presented before almost any other theory is developed.

The most general sheafification process outlined by Grothendieck takes a pre-sheaf F and defines a new pre-sheaf F^+ that assigns to each open set U the filtered colimit of $\mathsf{F} : \mathbf{Cov}(\mathsf{U})^{\mathrm{op}} \to \mathcal{D}$, see [46] section 17.4 for a modern exposition. Applying this construction twice defines a sheaf. However, in order to guarantee that this F^{++} is a sheaf one uses the properties of theorem 3.20. This now gets us to the more fundamental reason why the study of cosheaves may be so obscure: The non-exactness of $\varprojlim$ thwarts the Grothendieck prescription for cosheafification. *For pre-cosheaves valued in* $\mathbf{Set}, \mathbf{Ab}$ *or* $\mathbf{Vect}$ *there is simply no hope in using the standard, most general, cosheafification.*

There are a very small handful of approaches that have been used to circumvent this problem:

1. **Čech Homology and Smoothness**: One approach developed by Bredon [15, 16] is to define an equivalence relation on pre-cosheaves, more nuanced than isomorphism, which is constructed through zig-zag diagrams of *local* isomorphisms. Bredon develops an operation which uses Čech homology to take in one pre-cosheaf and produce another. In the event that the starting pre-cosheaf was equivalent to a cosheaf (Bredon calls such a pre-cosheaf *smooth*), he proves that his construction yields a cosheaf.

2. **Pro-Objects**: Another notable approach is to use pro-objects, i.e. functors $\mathsf{P} : \mathrm{I}^{\mathrm{op}} \to \mathcal{C}$ where I is filtrant. This theory is engineered in such a way that all the desired algebraic properties exist. This approach was perhaps first used by Jean-Pierre Schneiders [76] to develop a rich theory of cosheaves. The problem with pro-objects is its conceptual and algebraic difficulty. For the visually minded, cosheaves of pro-objects are infinite diagrams of infinite diagrams, which obscure the many natural examples of pre-cosheaves and cosheaves that one might want to capture. More recent work [86, 70], has also used this setup for cosheaves.

3. **Topology:** Here, one eschews full generality and works only with certain cosheaves known as **constructible cosheaves**, which can be thought of as cosheaves on particular finite spaces. Cosheafification in this setting exists and is natural. Often one does not even think about needing to cosheafify, because the diagrams are modeled on the points of the space. This school of thought, motivated by the vision and unpublished ideas of Bob MacPherson, has some recent trace in the literature, see [98, 59].

This paper sidesteps the issues of sheafification and cosheafification by focusing on the third approach. We believe that this provides a better way of learning sheaf theory as it removes the ever-present phrase "let *blank* be the sheafification of *blank*" and focuses on the more important technical machinery first.

4 Preliminary Examples

Theories should be motivated by examples. In this section we delay the development of the theory by introducing a few motivating examples. Broadly speaking, all sheaves are realized via local sections associated to a particular map. This principle is rigorously embodied by the *étalé perspective*. There is a similar concept for cosheaves of sets, but we will focus on giving lots of examples rather than emphasizing the duality between sheaves and cosheaves.

4.1 Sheaves Model Sections

Recall that if $f : Y \to X$ is a continuous map then a **section** is a continuous map $g : X \to Y$ such that $f(g(x)) = x$ for all x. This requires, in particular, that f be surjective. Sometimes a map admits a locally-defined section over a subset $U \subset X$, but not a global one. There is a sheaf that tracks this data.

Definition 4.1 (Sheaf of Sections of a Map)**.** Suppose $\pi : E \to X$ is a continuous map. Then we can associate a **sheaf of sections** to this map as follows:

$$U \rightsquigarrow F(U) := \{s : U \to \pi^{-1}(U) \text{ continuous} \,|\pi(s(x)) = x\}.$$

Clearly, if $F(X) \neq \emptyset$, then we can answer positively the question "Does $\pi : E \to X$ have a section?"

$$\begin{tikzcd} E \arrow[d, "\pi"'] \\ X \arrow[u, dotted, bend right, "?"'] \end{tikzcd}$$

To see why this is a pre-sheaf valued in $\mathcal{D} = \boldsymbol{Set}$ note that what is assigned to an open set U is a set of maps. A map whose domain of definition is U can always be restricted to a smaller open subset $V \subset U$ to define a map on V. This process of restricting the domain of definition we write as $\rho_{V,U}(s) := s|_V$, which is what makes this assignment a pre-sheaf.

Let us prove this defines a sheaf. Suppose $U \subset X$ and $\mathcal{U} = \{U_i\}_{i \in \Lambda}$ is an arbitrary open cover of U. We must prove that the map

$$F(U) \to F[\mathcal{U}] := \varprojlim(N(\mathcal{U})^{op} \to \boldsymbol{Open}(X)^{op} \to \boldsymbol{Set})$$

is an isomorphism. Recall that the limit can be described in terms of products and equalizers. As such, every element of the limit is described by a collection of continuous sections $s_i : U_i \to \pi^{-1}(U_i)$, one for each element of the cover, such that on intersections $\rho_{ij,i}(s_i) = \rho_{ij,j}(s_j)$.[11] The natural map from $F(U)$ to $F[\mathcal{U}]$ simply takes a section $s \in F(U)$ to a restricted section $s_i := s|_{U_i}$ and this map is clearly injective. To check surjectivity, note that an element in the limit defines a section over U by setting $s(x) = s_i(x)$ if $x \in U_i$ and this will be continuous by the pasting lemma described at the beginning of section 3.

[11] Here we have adopted the shorthand of referring to open sets via elements of the nerve.

Example 4.2. For a simple example, consider the projection onto the first coordinate $\pi_t : [0,1] \times [0,1] \to [0,1]$, which we regard as taking a time-space coordinate (t, x) to its time coordinate t. There are lots of sections of this map. The map that assigns to each time t a fixed position $p \in [0,1]$ defines a section, so there are uncountably many sections.

Now consider a different map that comes from restricting the time projection map to a subset $E \subseteq [0,1] \times [0,1]$, i.e. $\pi := \pi_t|_E : E \to [0,1]$ is the restricted map. A drawing can be found in figure 2 where E is the region bound between the two curves. Does it have any global sections, i.e. is $F(X) \neq \emptyset$?

Figure 2: Is There a Section?

The answer is clearly no. The example in figure 2 illustrates a concept central to sheaf theory. Although about each point in time t there is some $\epsilon > 0$ such that on the open set $(t - \epsilon, t + \epsilon)$ a continuous section can be defined, there is no globally defined section. Thus local sections (local solutions) exist, but they do not always glue together to define a global section (global solution). This is why we say

"Sheaves mediate the passage from local to global."

Example 4.3 (Square Map)**.** Suppose $f : \mathbb{C} \to \mathbb{C}$ is the map sending a complex number z to z^2. For a point $w = re^{i\theta}$ with $r \neq 0$ there are two points in the fiber: $z = \sqrt{r}e^{i\theta/2}$ and $z' = \sqrt{r}e^{i\theta/2+\pi}$. Consequently, for a small connected neighborhood about w there are two corresponding continuous sections. There is no global section because the square root map is necessarily multi-valued when considered all the whole complex plane.

Lists of similar examples abound in geometry and topology, most of which are concerned with the following mathematical structure.

Definition 4.4. A **fiber bundle** over X consists of total space E equipped with a continuous surjective map $\pi : E \to X$ satisfying the property that for each point $x \in X$ there exists an open neighborhood U such that the following diagram commutes.

$$\begin{tikzcd}
\pi^{-1}(U) \arrow[rr, "h_U"] \arrow[dr, "\pi"'] & & U \times F \arrow[dl, "\pi_1"] \\
& U &
\end{tikzcd}$$

Here F is the fiber space, h_U is a homeomorphism and π_1 is projection onto the first factor. If F is a discrete space then we usually write $\tilde{X}$ instead of E and write that $\pi : \tilde{X} \to X$ is a **covering space**. If each fiber $\pi^{-1}(x)$ is endowed with the structure of a group, i.e. $F = G$ with the discrete topology, so that h_U induces a group isomorphism between $\pi^{-1}(x)$ and G, then E is called a **bundle of groups**. Analogous definitions hold for fiber a ring or a module.

The map $\pi : M \to S^1$ where $M := S^1 \times \mathbb{R} / \sim$ with $(x, y) \sim (x + 2\pi, -y)$ is an example of a fiber bundle over S^1. Restricting the domain of π to the subspace $S^1 \times [-1, 1]$ allows one to think of this map as projecting the Möbius bundle to its core circle. The projection π has a section that embeds S^1 as the zero section, but there are no sections which avoid $S^1 \times \{0\}$. The "hairy ball" theorem is the analogous statement except for the tangent bundle to the two sphere S^2. Sheaf theory is the *lingua franca* for bundle theory and category theory. Thus even the most trivial example of a product bundle, $E = X \times k \to X$ where k is a field, is of interest.

Definition 4.5 (Constant Sheaf)**.** Suppose M is an R-module equipped with the discrete topology and $E = X \times M \to X$ is the product bundle. Then the sheaf of sections of this map is called the **constant sheaf with value** M. If $M = R$ is a field k or the ring $\mathbb{Z}$ we will just say the constant sheaf and write k_X or $\mathbb{Z}_X$.

If $\pi : E \to X$ is not necessarily the product bundle, but has fiber M, then we call the sheaf of sections of π the **locally constant sheaf with value** M.

We can characterize what the constant sheaf k_X assigns to any open set U as the product of the field k for as many connected components as U has. If one recalls the connection between Čech cohomology and the sheaf axiom outlined in 3.3, then one sees that

$$k_X(U) \cong \check{H}^0(U; k_X) \cong H^0(U).$$

If we permit ourselves to imagine what the higher Čech cohomology of a sheaf would yield, we might expect that it would tell you the higher cohomology of the open set U, which could be arrived at by studying **derived sections** of the constant sheaf. Studying the cohomology of a sheaf that is not constant, but only locally constant, then one would get cohomology of a space with "twisted" coefficients.

We introduce one last example of a sheaf that is of paramount importance for this paper.

Definition 4.6 (Constructible Sheaf)**.** Let F be a sheaf on a topological space X. One says that F is **constructible** if there exists a filtration by closed subsets

$$\emptyset = X_{-1} \subset X_0 \subset X_1 \subset \cdots \subset X_n = X$$

such that on the each connected component of the space $X^k = X_k - X_{k-1}$, the **restricted** sheaf $F|_{X^k}$ is locally constant (with value a finite dimensional vector space). Instead of asking for a filtration one can ask for a decomposition of X into disjoint pieces X_σ over each the restricted sheaf is locally constant. Here we define the restricted sheaf $F|_{X^k}(V) := F(U)$ for $V = X^k \cap U$ where U is open in X and V is an open in the subspace topology on X^k.

Remark 4.7. We will need to impose further conditions on the nature of the filtration so that we get nice properties. As stated, there is nothing to prevent us from using a one step filtration of the Cantor set. Expressing precisely these extra conditions will require a short introduction to stratification theory. These extra conditions will allow us to extract the strong geometric and topological content of constructible sheaves.

4.2 Cosheaves Model Topology

The omnipresence of sheaves in geometry and topology should come with no surprise to many researchers in the algebraic cousins of these fields. Remarkably, cosheaves are just as abundant, but this fact is less well appreciated. This might stem from a desire to avoid excessive terminology as very classical constructions in topology might be called cosheaves, but we will briefly reverse this wisdom to provide ourselves with lots of examples.

Perhaps the closest parallel to the sheaf of sections is the cosheaf of pre-images, but the presence of topology makes it a richer object of study.

Definition 4.8 (Cosheaf of Pre-images)**.** Suppose $f : Y \to X$ is a continuous map. We can define the pre-cosheaf of topological spaces $\hat{F} : \mathbf{Open}(X) \to \mathbf{Top}$ by assigning to an open subset the pre-image $f^{-1}(U) \in \mathbf{Open}(Y)$ endowed with the subspace topology, i.e.

$$U \rightsquigarrow f^{-1}(U).$$

Since colimits in the open set category are just unions and $f^{-1}(\cup_i U_i) = \cup_i f^{-1}(U_i)$, this defines a cosheaf.

Example 4.9 (Feature Function)**.** Suppose we have a topological space X, populated with features of interest, expressed as a function $P : \{1, \ldots, n\} \to X$. We get a cosheaf of sets via $\hat{F}(U) = P^{-1}(U)$. A slightly different cosheaf is gotten by letting $\hat{G}(U) = U \cap \mathrm{im}(P)$, which cannot distinguish points with identical images.

In the case $n = 1$ we can linearize this last example to define an example analogous to an example commonly encountered when studying sheaves.

Definition 4.10 (Skyscraper Cosheaf)**.** Suppose $x \in X$ and M is an R-module. Let's define the **skyscraper cosheaf** at x with value M to be

$$\hat{S}_x^M(U) = \begin{cases} M & \text{if } x \in U \\ 0 & \text{otherwise.} \end{cases}$$

When $M = R$, we drop the superscript for notational convenience.

We could adopt the perspective of cosheaves of pre-images as an alternative to continuous functions. This has been suggested in the past by John von Neumann and his derisively-named **pointless topology**, where in place of topological spaces one uses the poset of open sets as a primary notion – an example of a **locale** – and one observes that every continuous map of spaces $f : Y \to X$ induces a functor between categories $f^\circ : \mathbf{Open}(X) \to \mathbf{Open}(Y)$. This perspective will be of use later as we introduce operations on sheaves and cosheaves.

The cosheaf of pre-images will provide us with lots of examples of cosheaves pertinent to topology. However, viewing the entire information of the fiber (pre-image) is often too much to consider. Instead, one can consider invariants of the fiber and get a sometimes simpler, but still content-rich cosheaf (pending certain properties of the invariant).

Example 4.11 (Connected Components of the Fiber)**.** Given a continuous map of spaces $f : Y \to X$, one can define a pre-cosheaf of the components of the pre-image (not path components) $\hat{F} : \mathbf{Open}(X) \to \mathbf{Set}$. This is done via the assignment

$$U \rightsquigarrow \pi_0(f^{-1}(U)).$$

This is not always a cosheaf. However, if Y happens to be locally connected, i.e. the connected components of an open set are open, then it is. Alternatively, one can observe that the functor $\pi_0 : \mathbf{Top}_{lc} \to \mathbf{Set}$ is left adjoint to the discrete space functor and so it preserves colimits [98]. For a concrete example, consider the cosheaf of connected components associated to the map $\pi : E \to X$ found in figure 2.

For spaces where the notion of components and path components agree, one could ask if studying higher invariants of the pre-image is of interest. After π_0 one might consider the fundamental group $\pi_1(-; x_0)$, but we must free ourselves from picking basepoints. As such, we use groupoids instead of groups.

Definition 4.12 (Poincaré and Fundamental Groupoid)**.** To a topological space X, one can consider the gadget $\pi_\infty(X)$, called the **Poincaré ∞-groupoid**, which has an object for each point of X, a morphism for every path $\gamma : [0, 1] \to X$ with $\gamma(0) = x$ indicating the start and $\gamma(1) = y$ the destination, a "2-morphism" for every continuous map $\sigma : \Delta^2 \to X$, and so on for higher Δ^n. The 2-morphisms should be regarded as providing a homotopy between $\sigma|_{0,2}$ and $\sigma|_{1,2} \circ \sigma|_{0,1}$, i.e. a morphism between morphisms. Here $\sigma|_{i,j}$ is the restriction of the map σ to the edge going from vertex i to j. As stated, this is an example of an ∞-category, which is different than a category. However, if we declare the morphisms between x and y to be *homotopy classes* of paths between x and y, where the homotopies fix the endpoints, then we get an ordinary category $\pi_1(X)$ called the **fundamental groupoid**. A **groupoid** is a category where all the morphisms are invertible. We denote the category of groupoids by $\mathbf{Grpd}$.

Van Kampen's theorem in its more sophisticated form [17, 60] tells us that the fundamental groupoid can be built up by gluing together local fundamental groupoids.

Example 4.13 (van Kampen's Theorem)**.** Suppose X is a locally connected topological space, and suppose $\mathcal{U} = \{U_i\}$ is a cover of X by path-connected open subsets, then the van Kampen theorem states that

$$\pi_1(X) \cong \varinjlim_{I \in N(\mathcal{U})} \pi_1(U_I),$$

i.e. the functor $\pi_1 : \mathbf{Open}(X) \to \mathbf{Grpd}$ is a cosheaf for the cover $\mathcal{U}$. However, since any cover is refined by its connected components, which are open by assuming local connectivity, the arguments of section 3.4 imply that the fundamental groupoid is a cosheaf.

Heuristically, we can also consider a sort of linearization and grading of the Poincaré ∞-groupoid by considering for each p the R-module freely generated by the maps of the form $\sigma : \Delta^p \to X$. This is the module of singular p-chains.

Example 4.14 (Singular p-chains). Fix X a topological space and an open subset U. A singular p-chain on U is nothing more than a R-linear combination of maps of the form $\sigma : \Delta^p \to U$. Since we can always post-compose a p-chain on U with an inclusion $U \hookrightarrow V$, this defines a pre-cosheaf

$$C_p(U) = \{\sum_\sigma \lambda_\sigma \sigma | \lambda_\sigma \in R, \sigma : \Delta^p \to U\}.$$

This is, however, not a cosheaf as defined. Try writing down a chain on a union of two open sets as a linear combination of chains on the two sets. A chain needs be sub-divided into pieces coming from each open set, each piece being represented as a map from a fixed simplex. As such, if we define

$$\hat{C}_p(U) := \varinjlim C_p(U)$$

where the colimit is being performed over iterated subdivision, then we obtain a cosheaf [16].

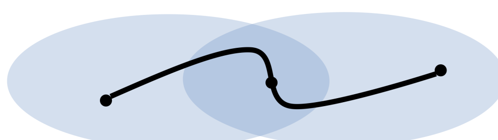

Figure 3: Barycentric Subdivision of a Singular Chain

Remark 4.15 (Mayer-Vietoris and Cosheaves). Another way of seeing that singular p-chains do not define a cosheaf is to recall that the proof of the **Mayer-Vietoris** theorem starts with the observation that the sequence

$$0 \to C_p(U \cap V) \to C_p(U) \oplus C_p(V) \to C_p(U + V) \to 0$$

is exact. Here the $C_p(U + V)$ is just notation for the cokernel of the previous map, thus the sequence is by definition exact. The elements of the cokernel are linear combinations of singular chains strictly contained in either U or V. One then uses barycentric subdivision to show that the *complexes* $C_\bullet(U + V)$ and $C_\bullet(U \cup V)$ are chain homotopy equivalent. Letting $R = k$ be a field, this motivates defining a cosheaf valued in $\mathcal{D} = K^b(\mathbf{Vect}_k)$ by assigning

$$U \rightsquigarrow C_\bullet(U; k)$$

and this will be a cosheaf.[12] The category $K^b(\mathbf{Vect}_k)$ will be discussed later in the paper where it plays a more important role, but briefly stated it is the category whose objects are chain complexes of vector spaces of finite length and whose morphisms consist of equivalence classes of maps where we have identified those that are chain homotopic. This makes

$$C_\bullet(U+V) \cong C_\bullet(U \cup V)$$

thereby forcing the cosheaf axiom to hold. Of course the way this isomorphism is proven is via the use of barycentric subdivision, so we can avoid using cosheaves of chain complexes by working with the cosheaf $\hat{C}_p$ directly.

The cosheaves of singular chains serve a role precisely dual to the sheaves of co-chains commonly encountered in the literature. Consequently, homology is most naturally associated with cosheaf theory and cohomology is naturally associated with sheaf theory. However, there is a deeper duality between sheaves and cosheaves. When considering compactly-supported cohomology or closed (Borel-Moore) homology the natural habitats reverse. The kernel of this idea is present in the following example.

Example 4.16 (Compactly Supported Functions)**.** Suppose X is a locally compact Hausdorff space. Consider the following assignment:

$$\Omega^0_c : U \rightsquigarrow \{f : U \to \mathbb{R} \,|\, \mathrm{supp}(f) \text{ compact}\}$$

Compactly supported functions defined locally can always be extended to larger open sets via extension by zero. If X is a manifold, then we get more cosheaves of compactly supported differential p-forms Ω^p_c for $p \geqslant 0$.

4.3 Taming of the Sheaf... and Cosheaf

As argued, the canonical example of a sheaf is the sheaf of sections of a map. This stands in contrast with the cosheaf of pre-images. However, a legitimate concern of both examples is its lack of computability. This concern is heightened given that the digital computer is becoming an increasingly common tool for modern mathematics.

A natural question might then be "Can we store the sheaf of sections on a computer?" Even in the example depicted in figure 2, it seems unlikely. On a small open set the sheaf of sections is in bijection with the set $\{f : (x-\epsilon, x+\epsilon) \to (a,b) \,|\, \text{continuous}\}$, which is uncountable. Moreover, for simple spaces like the closed unit interval with its Euclidean topology, there are uncountably many open sets that we need to assign data to.

To handle the first problem of "too many sections" in a somewhat ad hoc manner, we can conduct some pre-processing on the input data $\pi : E \to X$. As a motivating example, we can consider a construction normally defined when $X = \mathbb{R}$.

Definition 4.17 (Reeb Graph)**.** Suppose Y is a topological space and $f : Y \to \mathbb{R}$ is a continuous map. The **Reeb graph** [71] is defined to be the quotient space $R(f) := Y/\sim$

[12] The author has recently learned that Jacob Lurie calls this a **homotopy cosheaf**.

where $y \sim y'$ if and only if y and y' belong to the same connected component of the fiber $f^{-1}(t)$.

$$\begin{tikzcd} Y \arrow[rr, "q"] \arrow[dr, "f"'] & & R(f) \arrow[dl, "\pi"] \\ & \mathbb{R} & \end{tikzcd}$$

Observe that $R(f)$ still possesses a map to $\mathbb{R}$. There is clearly a direct generalization for arbitrary base spaces X.

For an example of the Reeb graph, consider our zig-zag in figure 2. Now let's work out what the sheaf of sections for $R(\pi)$ is and what the cosheaf of connected components is as well. Observe that we can probe the sheaf or the cosheaf on $[0, 1] \subset \mathbb{R}$ by asking what it assigns to open sets of the form $(x - \epsilon, x + \epsilon)$. Clearly it is constant except when the open set intersects a "critical value." We express this observation by assigning values directly to cells in the visible decomposition of the codomain of the function. The data over incident edges and vertices are related, but the direction of that relation is dependent on whether we are considering a sheaf or a cosheaf. Making this observation rigorous has tremendous pay off because it allows us to avoid storing infinitely many open sets by instead working with finitely many cells.

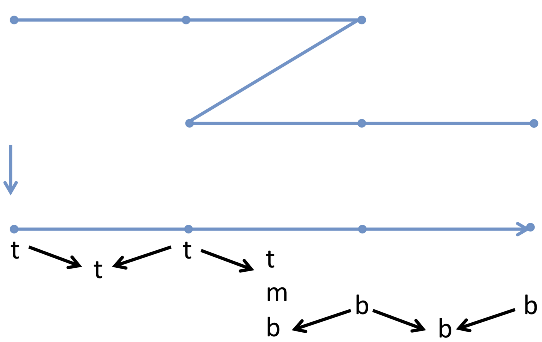

Figure 4: Sheaf of Sections of the Associated Reeb Graph

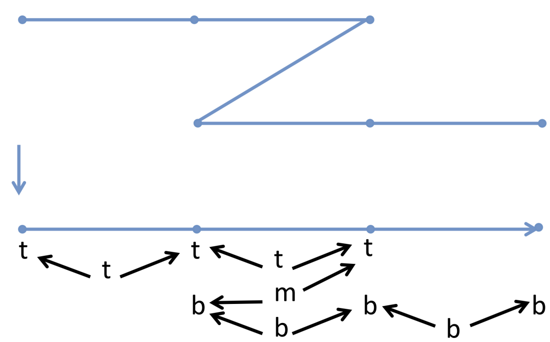

Figure 5: Cosheaf of Connected Components of the Reeb Graph

5 Cellular Sheaves and Cosheaves

We can take it as an experimental observation from figures 4 and 5 that in certain situations a sheaf or a cosheaf can be described as assigning data directly to the cells of a cell complex. Since cell complexes will be objects of primary importance to us, we review some definitions that might be non-standard.

Definition 5.1 (Regular Cell Complex). A **regular cell complex** X is a space equipped with a partition into pieces $\{X_\sigma\}_{\sigma\in P_X}$ such that the following properties are satisfied:

1. **Locally Finite:** Each point $x \in X$ has an open neighborhood U intersecting only finitely many X_σ.
2. X_σ is homeomorphic to $\mathbb{R}^k$ for some k (where $\mathbb{R}^0$ is one point).
3. **Axiom of the Frontier**:[13] If $\bar{X}_\tau \cap X_\sigma$ is non-empty, then $X_\sigma \subseteq \bar{X}_\tau$. When this occurs we say the pair are **incident** or that X_σ is a face of X_τ. The face relation makes the indexing set P_X into a poset by declaring $\sigma \leqslant \tau$.
4. The pair $X_\sigma \subset \bar{X}_\sigma$ is homeomorphic to the pair $\mathrm{int}(B^k) \subset B^k$, i.e. there is a homeomorphism from the closed ball $\varphi : B^k \to \bar{X}_\sigma$ that sends the interior of the ball to X_σ.

Remark 5.2 (Notation). Another common way of notating a cell complex is as a pair $(|X|, X)$ where X is the set of cells and $|X|$ is the topological space being partitioned. To each cell $\sigma \in X$ there is a corresponding topological subspace $|\sigma| \subseteq |X|$. Our definition's notation says that (X, P_X) is a cell complex. Our correspondence between cells and subspaces is $\sigma \rightsquigarrow X_\sigma$. However, we will have occasion to use both of these notations, and will sometimes use all three symbols $\sigma, |\sigma|$ and X_σ to mean the same thing.

It is true that every regular cell complex can be further decomposed so that the resulting space is the homeomorphic image of a simplicial complex. However, for ease of computations we want to work with a class of spaces more general and natural than regular cell complexes. As such, we work with cell complexes, adopting the same convention as in Allen Shepard's thesis [80].

Definition 5.3 (Cell Complex). A **cell complex** is a space X with a partition into pieces $\{X_\sigma\}$ that satisfies the first three axioms of a regular cell complex. Moreover, we require that when we take the one-point compactification of X, then the cells $\{X_\sigma\} \cup \{\infty\}$ are the cells of a regular cell complex structure on $X \cup \{\infty\}$.

Example 5.4. The open interval $(0, 1)$ decomposed with only one open cell is *not* a cell complex. Its one-point compactification is the circle decomposed with one vertex $\{\infty\}$ and one edge $(0, 1)$ whose attaching map is not an embedding, thus contradicting the fourth axiom.

[13] The **frontier** of a subspace A is the complement of A in its closure, i.e. $\mathrm{fr}(A) := \bar{A} - A$. In some forms this axiom reads: if $X_\sigma \neq X_\tau$ and $X_\sigma \cap \bar{X}_\tau \neq \emptyset$ then X_σ is contained in the frontier of X_τ.

Definition 5.5 (Cell category)**.** To a cell complex $(X, \{X_\sigma\}_{\sigma \in P_X})$ we can associate a category $\mathbf{Cell}(X; \{X_\sigma\})$, which is the indexing poset P_X viewed as a category. This means that there is one object σ for each X_σ and a unique morphism $\sigma \to \tau$ for each incident pair $X_\sigma \subseteq \bar{X}_\tau$. When there is no risk of confusion, or a cell structure is specified at the beginning, then we will suppress the extra notation and just use $\mathbf{Cell}(X)$ or X.

We now introduce diagrams indexed by the cell category. These were defined in Shepard's 1985 thesis [80] on p. 6, but were known as **stacks** in the first published volume of Zeeman's 1954 thesis [100] p. 626. In light of subsequent decades of research, Zeeman's thesis provides a discretized model of structures even more general than sheaves and cosheaves. To keep the presentation simple, we give Shepard's definition and its appropriate dualization.

Definition 5.6 (Cellular Sheaves and Cosheaves)**.** A **cellular sheaf** F valued in $\mathcal{D}$ on X is a functor $F : \mathbf{Cell}(X) \to \mathcal{D}$, i.e. it is

- an assignment to each cell X_σ in X an object $F(\sigma)$,
- and to every pair of incident cells $X_\sigma \subset \bar{X}_\tau$ a **restriction map**[14] $\rho^F_{\sigma,\tau} : F(\sigma) \to F(\tau)$.

Dually, a **cellular cosheaf** $\hat{F}$ valued in $\mathcal{D}$ on X is a functor $\hat{F} : \mathbf{Cell}(X)^{op} \to \mathcal{D}$, i.e. an assignment of an object $\hat{F}(\sigma)$ for each cell, and an **extension map** $r_{\sigma,\tau} : \hat{F}(\tau) \to \hat{F}(\sigma)$ for every pair of incident cells $X_\sigma \subset \bar{X}_\tau$.

Since functors between categories assemble themselves into a category of their own, we get categories of cellular sheaves and cosheaves.

Definition 5.7. We denote the category of cellular sheaves on X by

$$\mathbf{Shv}(X; \mathcal{D}) := \mathbf{Fun}(\mathbf{Cell}(X), \mathcal{D})$$

and the category of cellular cosheaves by

$$\mathbf{CoShv}(X; \mathcal{D}) := \mathbf{Fun}(\mathbf{Cell}(X)^{op}, \mathcal{D}).$$

Morphisms are natural transformations of functors. If $\mathcal{D} = \mathbf{Vect}$, then we will omit the notation after the semicolon and write $\mathbf{Shv}(X)$ and $\mathbf{CoShv}(X)$ instead.

The notation deliberately coincides with the notation used for categories of sheaves and cosheaves on an arbitrary topological space, i.e. functors out of the open set category that satisfy the appropriate axiom. This conflict will be resolved using the two different approaches outlined in sections 5.1 and 5.2.

[14] Shepard calls these co-restriction maps since they point from faces to co-faces, but we will see they are restriction maps in the Alexandrov topology.

5.1 Stratifications: A Forge of Theory and Examples

A workshop where many subtle examples are forged is the world of stratified spaces and maps. This workshop contains cell complexes which are the molecules that arise as the simple bonding of atomic copies of Euclidean space. However, the more complicated fusing of arbitrary manifolds, whose interactions are more organic, offer at once a simpler and subtler class of spaces.[15] The insights here are mostly due to Bob MacPherson, who saw through the smoke and fire of algebra and topology to gain a most elegant description of the sorts of sheaves and cosheaves we are interested in. In this section, we try to convey some of his ideas, but the proofs, technical modifications and mistakes are the author's own unless otherwise stated.

Definition 5.8 (Decomposition)**.** A **decomposition** of a space X is a locally finite partition of X into locally closed subsets (sets of the form $U \cap Z$ for U open and Z closed) $\{X_\sigma\}_{\sigma \in P_X}$ called **pieces**, which satisfy the axiom of the frontier. Consequently, P_X is a poset. When the pieces have the additional structure of being manifolds, we call them **strata**.

Remark 5.9. A **stratum** is sometimes used to mean either a union of strata of a fixed dimension or a single connected component in a decomposition. We usually prefer the latter meaning.

We have already encountered an example of a decomposition of a space X, namely a cell complex. Here each piece is homeomorphic to $\mathbb{R}^k$ for some k, which can vary from stratum to stratum. A graph is naturally decomposed into its vertices and open edges. For a decomposition that is not a cell complex, consider the complex numbers $\mathbb{C}$ partitioned into the sets $\{0\}$ and $\mathbb{C} - \{0\}$.

Definition 5.10. Suppose (X, P_X) and (Y, P_Y) are decomposed spaces, then a **decomposition-preserving** map is a continuous map $f : X \to Y$ that sends pieces to pieces, i.e. we have a commutative square

$$\begin{tikzcd} X \arrow[r, "f"] \arrow[d] & Y \arrow[d] \\ P_X \arrow[r, "P_f"] & P_Y \end{tikzcd}$$

In the case where the pieces are strata we call such a map a **stratum-preserving** map.

Much like how the notion of a category emerged through the study of functors, in some sense the necessity for decompositions more general then simplicial or cell complexes came about because not all maps preserved the pieces of those decompositions. We give an example of such a map.

Example 5.11 (Blow-Ups)**.** Consider the map

$$f : \mathbb{R}^2 \to \mathbb{R}^2 \qquad f(x, y) = (x, xy).$$

[15]According to Maxim Vybornov, Kari Vilonen once compared complex algebraic strata to biological cells and simplices to atoms to illustrate the difference in scale.

This map is not triangulable, see [82] page 305. This map is related to the operation in algebraic geometry known as "blowing up at a point." The blow-up map is an endless source of interesting geometry and counter-examples, so it is worth describing. Recall that the space of lines in $\mathbb{R}^2$, written $\mathbb{R}\mathbb{P}^1$ is defined to be the quotient of $\mathbb{R}^2 - \{(0,0)\}$ by the relation that $(x, y) \sim (\lambda x, \lambda y)$ for any $\lambda \neq 0$. Topologically, this quotient is the circle S^1. Tracing the image of the top arc of a circle from 0 to π through the quotient map one gets the complete circle in $\mathbb{R}\mathbb{P}^1$.

The blow-up B of $\mathbb{R}^2$ at the origin is defined to be the closure of the image of the map

$$\mathbb{R}^2 - \{(0,0)\} \hookrightarrow \mathbb{R}^2 \times \mathbb{R}\mathbb{P}^1$$

where the map to the first coordinate is the inclusion and the map to the second coordinate is the quotient map. The blow-up map $\pi : B \to \mathbb{R}^2$ is the projection back from the closure of the graph of this map to the closure of the domain, i.e. $\mathbb{R}^2$. Thus the fiber over $(0,0)$ is a circle, but the fiber over any other point is a single point. One can visualize this by restricting the map to a closed disk centered at the origin. The image is contained in a solid torus and the closure of the image will assign the core circle to the origin. The image of $\mathbb{D}^2 - \{(0,0)\}$ is commonly visualized as a spiral staircase, whose boundary traces out a torus knot. See [7] for a treatment of different constructions of real blow-ups and their functorial properties.

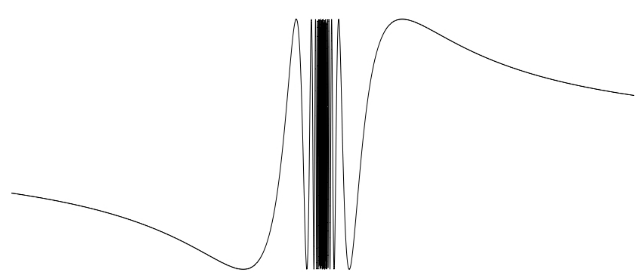

Figure 6: Topologist's Sine Curve

Decomposing spaces and maps gives some control over how these things are built up out of pieces, but it is not quite strong enough to tame the geometry of interest. In particular, the topologist's sine curve drawn in figure 6 can be decomposed into the two pieces

$$X_\tau := \{(x, \sin(1/x)) \mid x \neq 0\} \cup \{(0, y) \mid y \in [-1, 1]\} =: X_\sigma$$

that satisfy the axiom of the frontier $X_\sigma \subset \bar{X}_\tau$, but it does not have the intuitively desired property that

$$\dim X_\sigma < \dim X_\tau$$

([53], p. 131).

Further regularity conditions must be imposed to capture this property and other desired features that hold for piecewise-linear, algebraic, semi-algebraic, sub-analytic and other geometries. Systematic overviews of these different regularity conditions are overwhelming and highly technical. For a taste, one should consult Jörg Schürmann's remarkable service in writing down 14 different regularity conditions and their corresponding implications in Remark 4.1.9 of [77]. To keep the exposition light we focus on a geometric condition and its topological generalization as they have historically had a strong influence on stratification theory.

5.1.1 Whitney and Thom-Mather Stratified Spaces

In this section we relay two ways of fusing manifold pieces into non-manifold wholes. The champions of this section are Hassler Whitney and René Thom.[16] In 1965, Whitney, whose approach relies on the geometry of tangent planes and secant lines, defined two properties that a stratified space should possess [95, 96]. Thom, who proposed in a 1962 paper [88] a definition of a stratified space using tubular neighborhoods, later extracted the topological consequences of Whitney's definition and outlined a more general definition of a stratified space [87]. Thom's definition was first articulated carefully by John Mather in his famous 1970 Harvard "Notes on Topological Stability" [58], which went unpublished for 42 years and are to this day an excellent resource for learning the theory.

Proving that any Whitney stratified space admits the structure of a Thom-Mather stratified space requires substantial work. Thus, we present them below as separate definitions, beginning with Whitney's. We outline the properties that make Whitney stratified spaces nice as motivation for Thom's definition. After introducing both of these definitions we will present a proof[17] that says that closed unions of strata in a Thom-Mather space have regular neighborhoods, i.e. an open neighborhood and a weak deformation retraction. This result plays a key role in theorem 5.72.

Definition 5.12 (Whitney Stratified Spaces)**.** A **Whitney stratified space** is a tuple $(X, M, \{X_\sigma\}_{\sigma\in P_X})$ where X is a closed subset of a smooth manifold M along with a decomposition into pieces $\{X_\sigma\}_{\sigma\in P_X}$ such that

- each piece X_σ is a locally closed smooth submanifold of M, and
- whenever $X_\sigma \leqslant X_\tau$ the pair satisfies **condition (b)**. This condition says if $\{y_i\}$ is a sequence in X_τ and $\{x_i\}$ is a sequence in X_σ converging to $p \in X_\sigma$ and the tangent spaces $T_{y_i}X_\tau$ converges to some plane T at p, and the secant lines ℓ_i connecting x_i and y_i converge to some line ℓ at p, then $\ell \subseteq T$.

[16] For more historical context of these approaches, written by experts, we recommend the recent article [38] and Part I, Section 1 of [39].

[17] A proof appears in Mark Goresky's thesis [41] that was never published and which he graciously provided to the author. We have since modified that proof to suit our purposes.

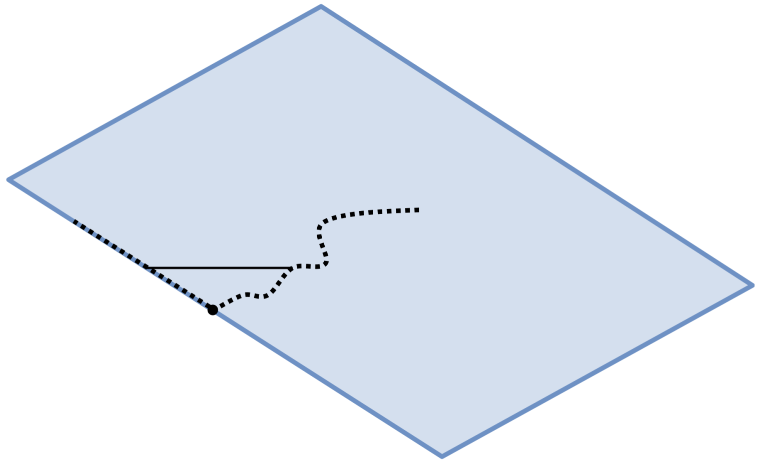

Figure 7: Diagram for Whitney Condition (b)

Remark 5.13. We have omitted **condition (a)** because it is implied by condition (b) ([58] Prop. 2.4). Condition (a) states that if we only consider a sequence y_i in X_τ converging to p such that the tangent planes $T_{y_i}X_\tau$ converge to some plane T, then the tangent plane to p in X_σ must be contained inside T.

The Whitney conditions are important because so many types of spaces admit Whitney stratifications, the most important being semi-algebraic and sub-analytic spaces. Remarkably, these conditions about limits of tangent spaces and secant lines imply strong structural properties of the space. To give the reader a taste for the properties enjoyed by Whitney stratified spaces, we provide a brief list:[18]

- **Dimension is Well-Behaved:** If $X_\sigma \subseteq \mathrm{fr}(X_\tau) := \bar{X}_\tau - X_\tau$, then $\dim X_\sigma < \dim X_\tau$. See Proposition 2.7 of [58] for a proof. This rules out the topologist's sine curve in figure 6 from being Whitney stratified.

- **Good Group of Self-Homeomorphisms:** If x and y belong to the same connected component of a stratum X_σ, then there is a homeomorphism $h : M \to M$ preserving X and other strata such that $h(x) = y$ ([58] pp. 480-481).

- **Local Bundle Structure:** Every stratum X_σ has an open tubular neighborhood T_σ and a projection map $\pi_\sigma : T_\sigma \to X_\sigma$ making it into a fiber bundle. This bundle is equipped with a "distance from the stratum" function $d_\sigma : T_\sigma \to \mathbb{R}_{\geqslant 0}$. If we define $S_\sigma(\epsilon)$ to be $d_\sigma^{-1}(\epsilon)$, then we can identify the map $\pi_\sigma : T_\sigma \to X_\sigma$ with the mapping cylinder of the restricted map $\pi : S_\sigma(\epsilon) \to X_\sigma$ ([40] p. 194). Moreover, the fiber of the bundle has the stratification of a cone on a link.

- **Triangulability:** Every Whitney stratified space can be triangulated [40].

[18]Here we follow part of MacPherson's summary in the appendix of his 1991 Colloquium notes [57].

The third property is historically the most important. It guarantees that a Whitney stratification "looks the same" along all points in a stratum. The tubular neighborhoods exhibit this local triviality. Mather extracts a compatibility condition among these neighborhoods that serves as a segue to Thom-Mather stratifications.

Definition 5.14 (Control Data)**.** Let $(X, M, \{X_\sigma\}_{\sigma \in P_X})$ be a Whitney stratified space and $\{(T_\sigma, \pi_\sigma, d_\sigma)\}$ a family of tubular neighborhoods. We call this family a system of **control data** if the following commutation relations are satisfied: if $X_\sigma \leqslant X_\tau$, then

$$\begin{aligned} \pi_\sigma \circ \pi_\tau &= \pi_\sigma \\ d_\sigma \circ \pi_\tau &= d_\sigma \end{aligned}$$

whenever both sides of the equations are defined.

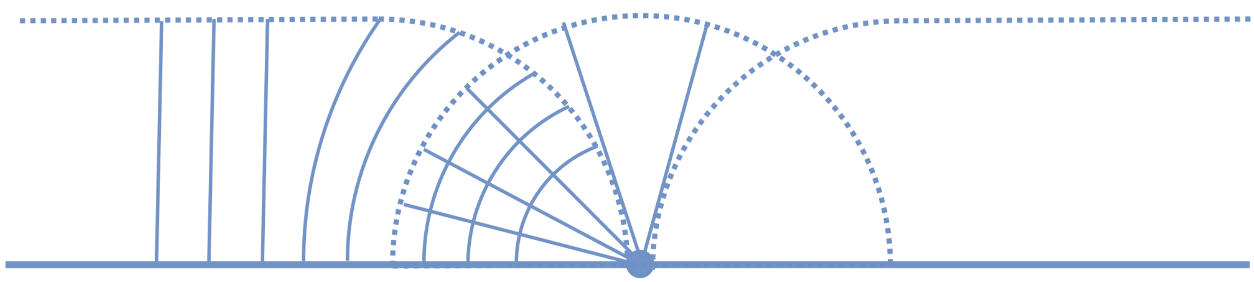

Figure 8: A System of Control Data

Remark 5.15. In figure 8 we have drawn some fibers of the retraction maps for two incident strata. Notice how the fibers must bend in order for the second compatibility condition to hold.

Mather proves that every Whitney stratified space admits a system of control data. The following definition axiomatizes the properties enjoyed by a Whitney space with a system of control data.

Definition 5.16 (Thom-Mather Stratified Spaces)**.** A **Thom-Mather stratified space** consists of a Hausdorff, locally compact topological space X with countable basis for its topology with some smooth structure; a decomposition into topological manifolds $\{X_\sigma\}_{\sigma \in P_X}$; and a family of control data $\{(T_\sigma, \pi_\sigma, d_\sigma)\}_{\sigma \in P_X}$, where T_σ is an open tubular neighborhood of X_σ, $\pi_\sigma : T_\sigma \to X_\sigma$ is a continuous retraction, and $d_\sigma : X_\sigma \to [0, \infty)$ is a continuous distance function. We require that the following conditions hold:

- $X_\sigma = d_\sigma^{-1}(0)$ for all σ.
- For any pair of strata X_σ, X_τ, define $T_{\sigma,\tau} := T_\sigma \cap X_\tau$, $\pi_{\sigma,\tau} := \pi_\sigma|_{T_{\sigma,\tau}}$ and $d_{\sigma,\tau} := d_\sigma|_{T_{\sigma,\tau}}$. We require that
$$(\pi_{\sigma,\tau}, d_{\sigma,\tau}) : T_{\sigma,\tau} \to X_\sigma \times (0, \infty)$$
is a smooth submersion. When $T_{\sigma,\tau} \neq \emptyset$, i.e. when $X_\sigma \leqslant X_\tau$, this implies $\dim X_\sigma < \dim X_\tau$.

- For any trio of strata X_σ, X_τ and X_λ we have

$$\begin{aligned}\pi_{\sigma,\tau} \circ \pi_{\tau,\lambda} &= \pi_{\sigma,\lambda} \\ d_{\sigma,\tau} \circ \pi_{\tau,\lambda} &= d_{\sigma,\lambda}\end{aligned}$$

  whenever both sides of the equation are defined.

Remark 5.17. One should observe that the definition does not require an embedding into an ambient space. Thus Thom-Mather stratified spaces allow us to treat Whitney stratified spaces intrinsically. Any Whitney stratified space (X, M) equipped with a system of control data $\{(T_\sigma, \pi_\sigma, d_\sigma)\}_{\sigma \in P_X}$ defines a Thom-Mather stratified space by intersecting each T_σ, which is open in M, with X.

Thom-Mather stratified spaces exhibit most of the good properties of Whitney stratified spaces. The proof that Thom-Mather spaces can be triangulated was carried out by Goresky [40], among others. His proof views the lines of figure 8 not as fibers of the retraction map π_σ, but rather[19] as fibers of a radial projection map to the boundary of a tubular neighborhood.

Definition 5.18 (Family of Lines). A **family of lines** on a Thom-Mather stratified space is a system of radial projections

$$r_\sigma(\epsilon) : T_\sigma - X_\sigma \to S_\sigma(\epsilon) := d_\sigma^{-1}(\epsilon)$$

one for each stratum X_σ, and a positive number δ, such that whenever $0 < \epsilon < \delta$ and $X_\sigma \leqslant X_\tau$, the following commutation relations hold:

1. $r_\sigma(\epsilon) \circ r_\tau(\epsilon') = r_\tau(\epsilon') \circ r_\sigma(\epsilon) \in S_\sigma(\epsilon) \cap S_\tau(\epsilon')$
2. $d_\sigma \circ r_\tau(\epsilon) = d_\sigma$
3. $d_\tau \circ r_\sigma(\epsilon) = d_\tau$
4. $\pi_\sigma \circ r_\tau(\epsilon) = \pi_\sigma$
5. If $0 < \epsilon < \epsilon' < \delta$, then $r_\sigma(\epsilon') \circ r_\sigma(\epsilon) = r_\sigma(\epsilon')$
6. $\pi_\sigma \circ r_\sigma(\epsilon) = \pi_\sigma$
7. $r_\sigma(\epsilon)|_{T_\sigma(\epsilon) \cap X_\tau} : T_\sigma(\epsilon) \cap X_\tau \to S_\sigma(\epsilon) \cap X_\tau$ is smooth

Remark 5.19. Every Thom-Mather stratified space admits a family of lines. This is the first proposition of [40].

[19]This description does not hold in higher dimensions.

Any family of lines can be used to identify a tubular neighborhood T_σ as a mapping cylinder for the restricted projection map $\pi_\sigma : S_\sigma(\epsilon) \to X_\sigma$. To do so, one defines a stratum-preserving homeomorphism

$$h_\sigma : T_\sigma - X_\sigma \to S_\sigma(\epsilon) \times (0, \infty) \qquad h_\sigma(p) := (r_\sigma(\epsilon)(p), d_\sigma(p))$$

and then extends the map in a suitable way, i.e. one takes $S_\sigma(\epsilon) \times [0, \infty) \sqcup X_\sigma$ and identifies $S_\sigma(\epsilon) \times \{0\}$ with its image under $\pi_\sigma : S_\sigma(\epsilon) \to X_\sigma$. One can check that this allows us to extend our map h_σ to T_σ, which we do so without changing notation. One should interpret this extended homeomorphism h_σ as providing a system of coordinates that is convenient for analyzing neighborhoods of strata. We use this system of coordinates in the following theorem.

Proposition 5.20 (Regular Neighborhoods of Closed Unions of Strata)**.** *Let* X *be a Thom-Mather stratified space and* $W = \cup_{\sigma \in P_W} X_\sigma$ *be a closed union of strata of* X*. The inclusion*

$$W \hookrightarrow U_W(\epsilon/2) := \bigcup_{\sigma \in P_W} T_\sigma(\epsilon/2) \subseteq X$$

is a homotopy equivalence.

Proof. Given a Thom-Mather stratified space, we can equip it with a family of lines [40]. We are going to use the family of lines to construct a weak deformation retraction of $U_W(\epsilon/2)$ inside a larger open neighborhood $U_W(\epsilon) := \cup_{\sigma \in P_W} T_\sigma(\epsilon)$. The idea is to shrink each tubular neighborhood $T_\sigma(\epsilon/2)$ to X_σ in such a way that a line connecting a point $p \in S_\sigma(\epsilon/2)$ and $r_\sigma(\epsilon)(p) \in S_\sigma(\epsilon)$ is stretched to connect $\pi_\sigma(p)$ and $r_\sigma(\epsilon)(p)$ after the homotopy. Figure 9 indicates which neighborhoods are to be collapsed.

To accomplish this stretching, let $f : \mathbb{R} \to [0, 1]$ be any smooth function with the following properties:

$$\begin{aligned} f(x) = 0 \quad &\text{if} \quad x \leqslant \tfrac{1}{2} \\ f(x) = 1 \quad &\text{if} \quad x \geqslant \tfrac{3}{4} \\ f'(x) > 0 \quad &\text{if} \quad x \in (\tfrac{1}{2}, \tfrac{3}{4}) \end{aligned}$$

The homotopy $H_\sigma : U \times [0, 1] \to U$ defined below shrinks $T_\sigma(\epsilon/2)$ to X_σ:

$$H_\sigma(p, t) := \begin{cases} p & \text{if} \quad p \notin T_\sigma(\epsilon) \\ h_\sigma^{-1}(r_\sigma(\epsilon)(p), d_\sigma(p)[(1-t)f(d_\sigma(p)/\epsilon) + t]) & \text{if} \quad p \in T_\sigma(\epsilon) \end{cases}$$

The homotopy is just a straight-line homotopy between the usual distance function d_σ and the shrunken one $d_\sigma(p)f(d_\sigma(p)/\epsilon)$. Moreover, since the homotopy only affects the distance coordinate, properties two and three of definition 5.18 imply that if $X_\sigma \leqslant X_\tau$ then

$$d_\tau(H_\sigma(p, t)) = d_\tau(p) \qquad \text{and} \qquad d_\sigma(H_\tau(p, t)) = d_\sigma(p).$$

As such, the shrinking homotopies can be applied in any order, i.e.

$$H_\tau(H_\sigma(p, t), s) = H_\sigma(H_\tau(p, s), t).$$

Observe that in a Thom-Mather stratified space, if $X_\sigma \neq X_{\sigma'}$ are two strata of the same dimension, then $T_\sigma \cap T_{\sigma'} = \emptyset$. Consequently, the definition for H_σ extends to a homotopy H_i that shrinks all the neighborhoods of strata of dimension i at the same time; one just defines $H_i(p,t) = H_\sigma(p,t)$ if $p \in T_\sigma(\epsilon)$. The commutation relation now extends to the statement that for any i and j

$$H_i(H_j(p,t),s) = H_j(H_i(p,s),t).$$

Thus, our desired homotopy can be defined to be

$$H(p,t) := H_0(H_1(H_2(\cdots(H_m(p,t),t)\cdots,t),t),t)$$

where the order of the composition doesn't matter and m is the maximum dimension of a stratum appearing in W. If we let $r_t(p) = H|_{U(\epsilon/2)\times I}$, then r_t defines a weak deformation retract of $U_W(\epsilon/2)$ to W, that is, $r_t(W) \subseteq W$ for all t, $r_0(U_W(\epsilon/2)) \subseteq W$ and $r_1 = \text{id}$. It is easy to show that this implies that $W \hookrightarrow U_W(\epsilon/2)$ is a homotopy equivalence. $\square$

Remark 5.21. One could imagine performing these homotopies at separate times by letting the homotopy parameter in dimension i be a function $s_i(t) = f(t-i)$ where the shrinking homotopy in dimension i is performed in the interval $(i+1/2, i+3/4)$. This is how it is done in Goresky's thesis [41]. This makes his homotopy

$$H_i^G(p,t) := \begin{cases} p & \text{if} \quad p \notin T_\sigma(\epsilon) \\ h_\sigma^{-1}(r_\sigma(\epsilon)(p), d_\sigma(p)[(1-s_i(t))f(d_\sigma(p)/\epsilon)+s_i(t)]) & \text{if} \quad p \in T_\sigma(\epsilon) \end{cases}$$

easier to visualize. However, the advantage of choosing $s_i(t) = t$ is that the homotopy is stratum preserving up until $t = 0$.

Remark 5.22. Of course, for a given stratum X_σ, away from its frontier the retraction map r_0 coincides with the tubular projection π_σ.

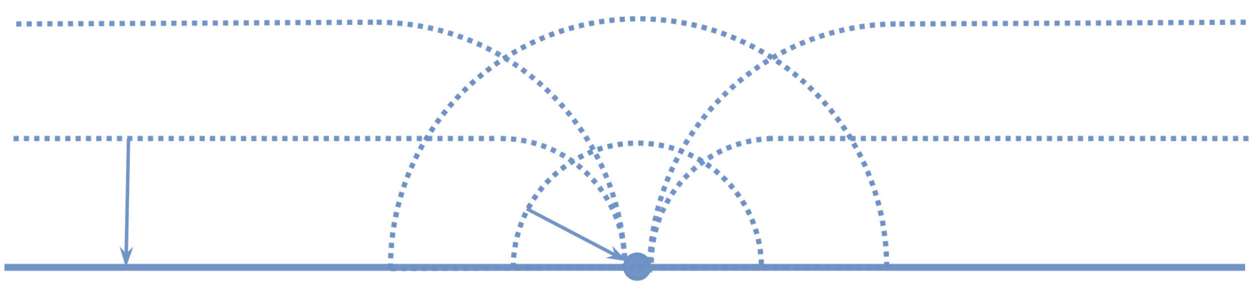

Figure 9: Two Sets of Tubular Neighborhoods

As the above proposition shows, control data is essential for providing Whitney stratified spaces with good neighborhoods. Not only do they endow Whitney stratified spaces with the structure of a Thom-Mather stratified space, they allow us to construct Goresky's family of lines to carry out these retractions. These retractions are instrumental to the cosheaves that we will construct in lemma 5.43 and theorem 5.72. There is another technical tool that we need that can only be developed in the presence of control data.

Definition 5.23. A **stratified vector field** η on $(X, \{X_\sigma\}_{\sigma \in P_X})$ is a collection of vector fields $\{\eta_\sigma\}_{\sigma \in P_X}$ with one smooth vector field on each stratum.

When it is meaningful to compare these vector fields, it is remarkable to note that this collection need not be continuous. Nevertheless, in the presence of control data, the flow generated by such a discontinuous vector field is continuous.

Definition 5.24. A stratified vector field η on X is said to be **controlled** by $\{T_\sigma, \pi_\sigma, d_\sigma\}$ if the following compatibility conditions are satisfied for any pair of strata $X_\sigma \leqslant X_\tau$:

$$\begin{aligned} \eta_\tau(d_{\sigma,\tau}(p)) &= 0 \\ d(\pi_{\sigma,\tau})(\eta_\tau(p)) &= \eta_\sigma(\pi_{\sigma,\tau}(p)) \end{aligned}$$

where ever both sides of the equation are defined.

5.1.2 Stratified Maps: Glued Together Fiber Bundles

Our main purpose for considering Whitney (and hence Thom-Mather) stratified spaces is to understand stratified maps. Such maps include Morse functions as a special case and are a good model for understanding moduli problems that commonly arise in applications. Over a given stratum, a stratified map looks like a fiber bundle and all fibers are homeomorphic in a stratum-preserving way. However, as we try to compare a fiber over one stratum with a fiber over that stratum's frontier, the blow-up map of example 5.11 frustrates our intuition. Thus, we introduce a more restrictive class of stratified maps called Thom maps. Finally, we illustrate that such general stratified maps are not necessarily closed under pullback. This motivates the move to tame topology in section 5.1.3.

Definition 5.25 (Whitney Stratified Map)**.** Suppose $f : M \to N$ is a smooth map between manifolds that contain stratified spaces $(X, \{X_\sigma\}_{\sigma \in P_X})$ and $(Y, \{Y_\sigma\}_{\sigma \in P_Y})$ such that $f(X) \subset Y$ with $f|_X$ *proper*. We say f is a Whitney **stratified map** if the pre-image of each stratum Y_σ is a union of connected components of strata of X and f takes these components submersively onto Y_σ.

Remark 5.26. To say that a map is (Whitney) **stratifiable** is to say there exists a stratification of X and Y such that the map is stratified. Often we will neglect to include the ambient manifolds and will say "Let $f : X \to Y$ be a stratified map."

Remark 5.27. When $N = Y$ is stratified as a single stratum, we say that f is a **stratified submersion**, i.e. $f|_X$ is proper and for each stratum X_σ $f|_{X_\sigma}$ is a submersion.

Recall that Ehresmann's theorem states that proper submersions are fiber bundles. Thus, over each stratum a stratified map is a fiber bundle. However, Ehresmann's theorem does not say that the local trivializations can be chosen to respect the stratification. This stratified analog of Ehresmann's theorem is expressed in Thom's first isotopy lemma ([39] p. 41).

Lemma 5.28 (Thom's First Isotopy Lemma)**.** *Let* $f : M \to \mathbb{R}^n$ *be a (proper) stratified submersion for* $X \subseteq M$ *a Whitney stratified subset. Then there is a stratum-preserving homeomorphism*

$$h : X \to \mathbb{R}^n \times (f^{-1}(0) \cap X)$$

which is smooth on each stratum and commutes with the projection to $\mathbb{R}^n$*. In particular, the fibers of* $f|_X$ *are homeomorphic by a stratum preserving homeomorphism.*

Remark 5.29. Of course, this implies that for a general stratified map, for every stratum of the codomain Y_σ, the fibers of $f|_{f^{-1}(Y_\sigma)} : f^{-1}(Y_\sigma) \to Y_\sigma$ are homeomorphic in a stratum-preserving way. This lemma will be used implicitly throughout the section. It expresses the idea that stratified maps are "glued together fiber bundles."

As one can imagine, there is a second isotopy lemma, which applies to a more restrictive class of stratified maps. We will not state the second isotopy lemma, rather we will use some of the theory leading up to it.

Definition 5.30. A **Thom mapping** is a stratified map $f : (X, M) \to (Y, N)$ that satisfies **condition** a_f for every pair of strata $X_\tau \geqslant X_\sigma$: let x_i be a sequence of points in X_τ converging to a point $p \in X_\sigma$. Suppose $\ker d(f|_{X_\tau})_{x_i} \subseteq T_{x_i}M$ converges to a plane $K \subseteq T_pM$, then $\ker d(f|_{X_\sigma})_p \subseteq K$.

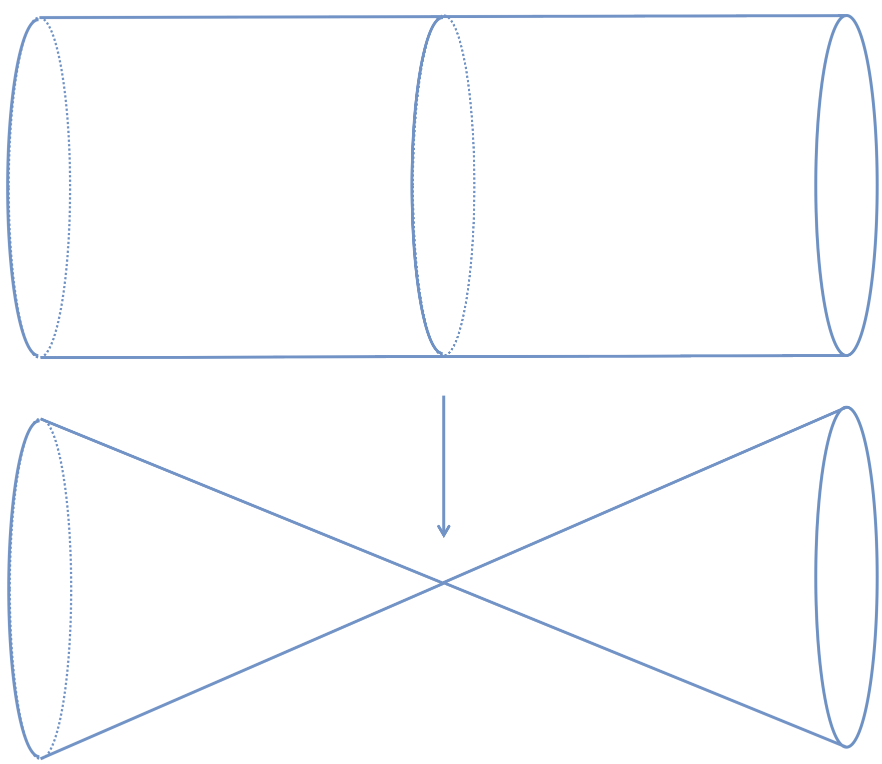

Figure 10: Not a Thom Mapping

In figure 10, we have drawn an example of a mapping that is not a Thom mapping.[20] Other non-examples include the blow-up map discussed in example 5.11. Any map that is triangulable satisfies Thom's condition a_f for that triangulation viewed as a stratification. It has been a long standing conjecture that every smooth Thom mapping is triangulable. Masahiro Shiota appears to have proven this conjecture in the C^∞ case [83], but we have chosen not to rely on this conjecture. Instead, we only need the following proposition of Mather's (Proposition 11.3 of [58]).

Proposition 5.31. *Suppose $f : X \to Y$ is a Thom mapping and a system of control data $\{T\}$ for Y is given. There exists a family of tubular neighborhoods $\{T'\}$ for X over $\{T\}$, which satisfies the following compatibility conditions:*

(a) If $X_\sigma \leqslant X_\tau$, then $\pi'_\sigma \circ \pi'_\tau = \pi'_\sigma$ for points in $T'_\sigma \cap T'_\tau$ in M. Furthermore, if $f(X_\sigma)$ and $f(X_\tau)$ lie in the same stratum of Y, then $d'_\sigma \circ \pi'_\tau = d'_\sigma$ where both sides are defined.

(b) If Y_σ is a stratum that contains X_σ, then

$$f(\pi'_\sigma(p)) = \pi_\sigma(f(p))$$

for all $p \in T'_\sigma \cap f^{-1}(T_\sigma)$.

Remark 5.32. The first condition is weaker than the usual definition of control data when the strata are not mapped to the same stratum. Consequently, the above notion of a **system $\{T'\}$ of control data over** $\{T\}$ is not the same as two systems of control data.

Just as the notion of control data generalizes to control data over control data, controlled vector fields generalize to controlled vector fields over controlled vector fields.

Definition 5.33. Suppose $f : X \to Y$ is a Thom mapping and $\{T'\}$ is a system of control data over $\{T\}$. If $\eta = \{\eta_\sigma\}$ is a controlled vector field on $\{Y_\sigma\}$ controlled by $\{T\}$, then there exists a stratified vector field $\eta' = \{\eta'_\sigma\}$ on $\{X_\sigma\}$ satisfying the following compatibility conditions:

(a) For any X_σ and $p \in X_\sigma$, we have

$$(df|_{X_\sigma})(\eta'_\sigma(p)) = \eta_\sigma(f(p))$$

where Y_σ is the stratum of Y that contains $f(p)$.

(b) For any $X_\sigma \leqslant X_\tau$, there is a neighborhood N'_σ in T'_σ such that for $p \in T'_\sigma \cap X_\tau$ we have

$$d(\pi'_{\sigma,\tau})(\eta'_\tau(p)) = \eta'_\sigma(\pi'_{\sigma,\tau}(p))$$

and if X_σ and X_τ are carried to the same stratum of Y, then we have further the condition that

$$\eta'_\tau(d'_{\sigma,\tau}(p)) = 0.$$

[20]This example is borrowed from [53].

Thus, the notion of a **controlled vector field η' over η** is a weaker one than a pair of controlled vector fields on X and Y that commute with the Thom mapping f.

The following result, proven with help from Mark Goresky, gives a useful criterion for determining when a stratified map is a Thom mapping, so as to make the above constructions possible there. It rests on the observation that all the classical examples of stratified maps $f : X \to Y$ that aren't Thom maps require considering a pair of strata $Y_\sigma < Y_\tau$ in Y whose codimension is at least two. Combinatorially, this allows us to have a pair of strata $X_\sigma < X_\tau$ in X such that $\dim X_\sigma \cap f^{-1}(p) > \dim X_\tau \cap f^{-1}(x_i)$ even though $\dim X_\sigma < \dim X_\tau$. In the following lemma we show that if the codomain only has strata of codimension 1, then the map is a Thom mapping.

Lemma 5.34. *Suppose $f : (X, M) \to (Y, N)$ is a Whitney stratified map that is C^1 on the ambient manifold M. Let $Y' = Y_\sigma \cup Y_\tau$ be the union of two strata whose difference in dimension is one. The restricted map $f' : (X', M) \to (Y', N)$ where $X' := f^{-1}(Y')$ is a Thom map.*

Proof. The proof is local, so we consider the following setup instead: Suppose $f : (X, M) \to (Y, \mathbb{R}^{k+1})$ is a Whitney stratified map where Y is the upper half plane in $\mathbb{R}^{k+1}$, i.e. $Y := \{(y_1, \cdots, y_{k+1}) \,|\, y_{k+1} \geqslant 0\}$. We assume that the stratification of the map stratifies Y as $Y_\tau := \{y_{k+1} > 0\} \cong \mathbb{R}^{k+1}$ and $Y_\sigma := \{y_{k+1} = 0\} \cong \mathbb{R}^k$. Let X_σ be a stratum of X that is mapped to Y_σ and X_τ a stratum mapped to Y_τ. Suppose $\{x_i\}$ is a sequence in X_τ and $\ker df|_{X_\tau}(x_i) =: K_i$ converges to a subspace $K_\infty \subseteq T_pM$ where $p \in X_\sigma$. We want to show that $K_p := \ker df|_{X_\sigma}(p) \subseteq K_\infty$. By passing to a subsequence we can further assume that the tangent planes $T_{x_i}X_\tau =: T_i$ converges to $T_\infty \subseteq T_pM$. By Whitney's condition (a), $T_pX_\sigma \subset T_\infty$.

Denote by $\rho_Y(y) := \pi_{k+1}(y_1, \ldots, y_{k+1}) = y_{k+1}$ the "distance from the stratum" function on Y. By pre-composing with f, this defines a function $\rho_X(x) := \rho_Y(f(x))$. Any vector $v \in T_i$ with $d\rho_X(x_i)(v) \neq 0$ must also have $df|_{X_\tau}(x_i)(v) \neq 0$ since the chain rule implies that $d\rho_X(x_i) = d\rho_Y(f(x_i)) \circ df|_{X_\tau}(x_i)$ and thus $v \notin K_i$.

Let $\pi_\sigma : \mathbb{R}^{k+1} \to Y_\sigma$ be the projection onto the first k coordinates. The restriction of π_σ to Y_τ, written $\pi_{\sigma,\tau}$, is a submersion. By virtue of $\pi_{\sigma,\tau} \circ f|_{X_\tau}$ being a submersion, any vector $w \in T_{\pi_\sigma(f(x_i))}Y_\sigma$ has a lift $w_{f(x_i)} \in T_{f(x_i)}Y_\tau$ so that $w_{f(x_i)} \in \ker d\rho_Y(f(x_i))$, which in turn has a lift $\tilde{w}_i \in T_i$. Consequently, $df|_{X_\tau}(x_i)(\tilde{w}_i) \neq 0$ and thus $\tilde{w}_i \notin K_i$. Moreover, $\tilde{w}_i$ is orthogonal to $\nabla\rho_X(x_i)$ since any lift of w is chosen to factor through the kernel of $d\rho_Y(f(x_i)) = \pi_{k+1}$.

Thus, each T_i can be written as $\tilde{T}_{\pi_\sigma(f(x_i))}Y_\sigma \oplus K_i \oplus \nabla\rho_X(x_i)$. Since $T_pX_\sigma \subset T_\infty$ the isomorphism $T_\infty \cong T_pX_\sigma \oplus (T_pX_\sigma)^\perp$ can be further refined as $T_\infty \cong \tilde{T}_{f(p)}Y_\sigma \oplus K_p \oplus (T_pX_\sigma)^\perp$. We have assumed that f is C^1 on the ambient manifold M so that the lifts $\tilde{T}_{\pi_\sigma(f(x_i))}Y_\sigma$ must converge (perhaps after passing again to a subsequence) to $\tilde{T}_{f(p)}Y_\sigma$. Additionally, $\nabla\rho_X(x_i)$ converges to a subspace of $(T_pX_\sigma)^\perp$. Finally, since $\dim X_\sigma < \dim X_\tau$, dimension constraints force $K_p \subseteq K_\infty$. This proves the lemma. $\square$

This lemma is instrumental for our proof of theorem 5.72. On it's own, it has a useful corollary.

Corollary 5.35. *Any stratified map* $f : (X, M) \to (Y, \mathbb{R})$ *that is* C^1 *on the ambient manifold is a Thom map.*

One would like to say that given a general (not necessarily Thom) stratified map $f : X \to Y$, one could take a path $\gamma : [0, 1] \to Y$ so that the pullback $\gamma^* f : f^{-1}(\gamma) \to I$ is stratified and hence, by the above corollary, a Thom mapping. However, as the next example shows, the pullback need not be stratifiable, so the hypothesis for the corollary fails.[21]

Example 5.36. The blow-up map $\pi : B \to \mathbb{R}^2$ is a Whitney stratified map, that is not a Thom mapping. The closure S of the "quick spiral"

$$S := \mathrm{cl}\{(r, \theta) \in \mathbb{R}^2 \mid r = e^{-\theta^2}\}$$

is also Whitney stratified despite wrapping around the origin infinitely many times ([67] Example 1.4.8). However, the inverse image $\pi^{-1}(S)$ cannot be stratified because the inverse image of $(0, 0)$ is S^1, which is of the same dimension as the inverse image of the spiral, despite the fact that the former is in the frontier of the latter. Since being Whitney stratified implies a drop in dimension of the frontier, contraposition shows that the inverse image cannot be Whitney stratified.

In theorem 5.72 we will give a direct geometric construction of several cosheaves associated to a stratified map. To do so we will need to consider a class of subsets and maps that have all the geometric properties of stratified spaces as well as being preserved under inverse images. This is provided in section 5.1.3.

5.1.3 Persistence and Tame Topology

Although stratification theory provides a first pass at taming geometry, it is unsuitable from our perspective because pathologies can still creep in via the inverse image, as example 5.36 showed. General stratified spaces and maps are still not tame enough. However, most sets and maps encountered in nature have extra structure. For instance, computer scientists commonly work with piecewise-linear (PL) spaces, which are describable in terms of affine spaces and matrix inequalities. Some algebraic geometers work with semialgebraic spaces, which use zeros and inequalities of polynomials to define their spaces. Analysts tend to use analytic or subanalytic spaces, because the theory is well behaved. Traditionally, one has had to make a choice, once and for all, to speak only of PL geometry, or only of algebraic geometry, or only of analytic geometry. The curse of Babel has confused and separated these domains for a hundred years.

In 1984, Grothendieck declared that an axiomatic "tame topology" or "topologie modérée" should be developed by extracting out precisely those properties that make these classes of spaces good ones [43]. MacPherson put forth in his lecture notes for the 1991 AMS colloquium lectures a definition of what should constitute a "good" class of subsets of a manifold M [57]. Namely, a subset S is good if there is a Whitney stratification of M such that S is a

[21] We are indebted to Mark Goresky for suggesting the key ideas of this example.

union of strata. These subsets should be closed under the finite set-theoretic operations of unions, intersections and differences. Additionally, the closure of any good subset should be good.

In 1996, Lou van den Dries and his student Chris Miller set forth a most satisfactory definition in their paper "Geometric Categories and O-minimal Structures" [90]. Taking requests from sheaf theorists [73] and other working geometers, their paper is a valuable service to the community. It globalized a local solution to Grothendieck's program known as **o-minimal topology**. The theory of o-minimal topology is grounded in model theory and logic, but it has left almost no trace from those fields. All the logical operations of $\forall, \exists, \vee, \wedge$ are converted into familiar operations in geometry. Each of the above languages (PL, semialgebraic, subanalytic) are instances of an o-minimal structure. The common fundamental theorems, each expressed in their own language, can be reduced to universal logical operations, and hence geometric ones. We will start by examining o-minimal structures as they form the local models of Miller and van den Dries definition. The reader is urged to consult the textbook "Tame Topology and O-minimal Structures" [91] as it is an excellent introduction that requires virtually no pre-requisites.

Definition 5.37 ([91], p. 2)**.** An **o-minimal structure on** $\mathbb{R}$ is a sequence $\mathcal{O} = \{\mathcal{O}_n\}_{n\geqslant 0}$ satisfying

1. $\mathcal{O}_n$ is a boolean algebra of subsets of $\mathbb{R}^n$, i.e. it is a collection of subsets of $\mathbb{R}^n$ closed under unions and complements, with $\emptyset \in \mathcal{O}_n$;
2. If $A \in \mathcal{O}_n$ then $A \times \mathbb{R}$ and $\mathbb{R} \times A$ are both in $\mathcal{O}_{n+1}$;
3. The sets $\{(x_1, \ldots, x_n) \in \mathbb{R}^n \,|x_i = x_j\}$ for varying $i \leqslant j$ are in $\mathcal{O}_n$;
4. If $A \in \mathcal{O}_{n+1}$ then $\pi(A) \in \mathcal{O}_n$ where $\pi : \mathbb{R}^{n+1} \to \mathbb{R}^n$ is projection onto the first n factors;
5. For each $x \in \mathbb{R}$ we require $\{x\} \in \mathcal{O}_1$ and $\{(x, y) \in \mathbb{R}^2 \,|x < y\} \in \mathcal{O}_2$;
6. The only sets in $\mathcal{O}_1$ are the finite unions of open intervals and points.

When working with a fixed o-minimal structure $\mathcal{O}$ on $\mathbb{R}$ we say a subset of $\mathbb{R}^n$ is **definable** if it belongs to $\mathcal{O}_n$. A map is definable if its graph is definable.

Remark 5.38. One should note that the third and sixth property together prohibit any spiral that wraps infinitely many times around the origin from being part of an o-minimal structure. Thus, the quick spiral in example 5.36 is not definable.

Now we prove that definable sets and maps are closed under pullbacks.

Lemma 5.39. *Suppose* $f : X \to Z$ *and* $g : Y \to Z$ *are definable maps, then the pullback* $X \times_Z Y := \{(x, y) \in X \times Y \,|\, f(x) = g(y)\}$ *is a definable set and the restrictions of the projection maps are definable as well.*

Proof. First note that if $X \in \mathcal{O}_n$ and $Y \in \mathcal{O}_m$, then $X \times Y = (X \times \mathbb{R}^m) \cap (\mathbb{R}^n \times Y)$ is in $\mathcal{O}_{n+m}$. Since Γ_f and Γ_g are definable, we know that $\Gamma_f \times Y = \{(x, y, f(x))\}$ and $\Gamma_g \times X = \{(x', y', g(y')\}$ are both definable subsets of $X \times Y \times Z$. Since the intersection is definable, and a point in the intersection has $(x, y, f(x)) = (x', y', g(y'))$, the image of the projection to $X \times Y$ is the pullback. One can then use B.3 of [90] to conclude that the restriction to the pullback of the projection maps to X and Y is definable as well. □

There are surprising facts that follow from the axioms of an o-minimal structure. For example, if $A \in \mathcal{O}$, then the closure $\bar{A}$ is in $\mathcal{O}$ ([91] Ch. 1, 3.4). Another surprising fact is that definable sets can be Whitney stratified [52]. Thus, these sets meet the requirements of MacPherson to form a good class of subsets. Perhaps even better than MacPherson's sets, definable sets can be given finite cell decompositions, where "cell" has its own special meaning ([91] Ch. 3).

The prototypical o-minimal structure is the class of semialgebraic sets, which has become increasingly relevant in applied mathematics.

Definition 5.40. A **semialgebraic** subset of $\mathbb{R}^n$ is a subset of the form

$$X = \bigcup_{i=1}^{p} \bigcap_{j=1}^{q} X_{ij}$$

where the sets X_{ij} are of the form $\{f_{ij}(x) = 0\}$ or $\{f_{ij} > 0\}$ with f_{ij} a polynomial in n variables.

The only semi-algebraic subsets of $\mathbb{R}$ are finite unions of points and open intervals. From the definition, one sees that the class of semialgebraic sets is closed under finite unions and complements. The **Tarski-Seidenberg** theorem states that the projection onto the first m factors $\mathbb{R}^{m+n} \to \mathbb{R}^m$ sends semialgebraic subsets to semialgebraic subsets [22]. We can deduce from this theorem all of the conditions of o-minimality.

Semialgebraic maps are defined to be those maps $f : \mathbb{R}^k \to \mathbb{R}^n$ whose graphs are semialgebraic subsets of the product. It is a fact that semi-algebraic sets and maps can be Whitney stratified [81]. This allows us to consider the following example of a semi-algebraic family of sets:

Example 5.41 (Point-Cloud Data)**.** Suppose Z is a finite set of points in $\mathbb{R}^n$. For each $z \in Z$, consider the square of the distance function

$$f_z(x_1, \ldots, x_n) = \sum_{i=1}^{n} (x_i - z_i)^2.$$

By the previously stated facts we know that the sets

$$B_z := \{x \in \mathbb{R}^{n+1} \mid f_z(x_1, \ldots, x_n) \leqslant x_{n+1}\}$$

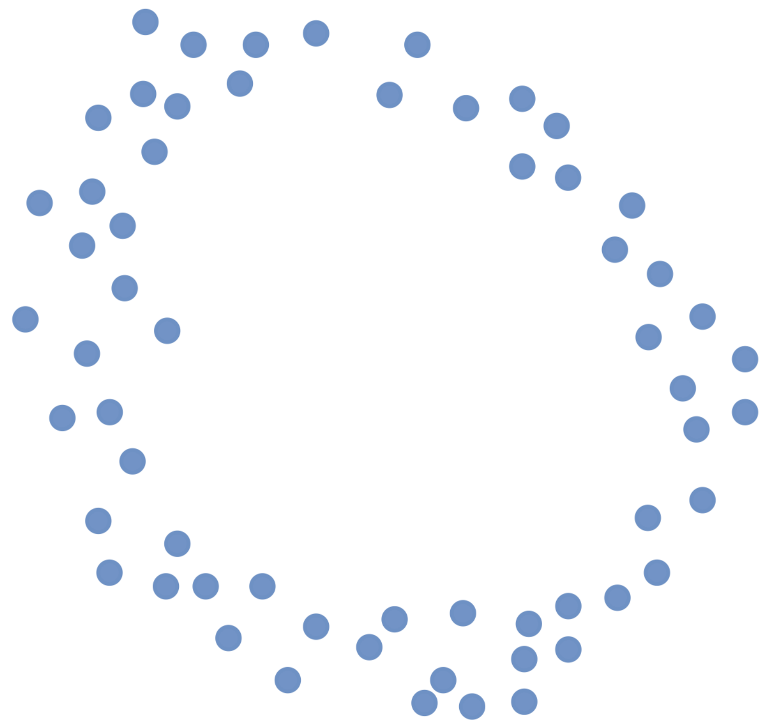

Figure 11: Point Cloud Data

are semialgebraic along with their unions and intersections. Denote by X the union of the B_z. The Tarski-Seidenberg theorem implies that the map

$$f : X \to \mathbb{R} \qquad f^{-1}(r) := \cup_{z \in Z} B(z, \sqrt{r}) = \{x \in \mathbb{R}^n \mid \exists z \in Z \,\text{s.t.}\, f_z(x) \leqslant r\}$$

is semialgebraic. In particular the topology of the fiber (of the union of the closed balls) can only change finitely many times.

Point-Cloud data is the standard entrée to the area of **persistent homology**. Although the philosophy of persistence can be applied whenever there is a filtration of some space, the study of point-cloud data illustrates most of the important ideas.[22] However, the treatment above is non-standard; o-minimal topology is almost never invoked. The way point-cloud data is normally treated is to start with some finite set of points $Z \subset \mathbb{R}^n$ and view the union of balls of a fixed radius $\cup_{z \in Z} B(z, r) =: X_r$ as a space in its own right. Then one observes that there are natural inclusions

$$X_{r_0} \hookrightarrow X_{r_1} \hookrightarrow X_{r_2} \hookrightarrow X_{r_3} \cdots$$

whenever $r_0 \leqslant r_1 \leqslant r_2 \leqslant \cdots$ and so on. Applying the ith homology functor $H_i(-; k)$ turns this diagram of spaces into a diagram of vector spaces.

$$H_i(X_{r_0}; k) \to H_i(X_{r_1}; k) \to H_i(X_{r_2}; k) \to H_i(X_{r_3}; k) \to \cdots \tag{5.42}$$

[22] For an excellent survey article see Robert Ghrist's Chauvenet prize-winning survey article "Barcodes: The Persistent Topology of Data" [35].

Notice that this perspective relies heavily on the fact that the radii $r_i \in \mathbb{R}$ can be compared by using the standard order on $\mathbb{R}$. Our perspective[23] is slightly different and is prototypical for the general perspective of theorem 5.72.

Lemma 5.43. *Any stratified map* $f : X \to \mathbb{R}$ *defines, for each* i*, a cellular cosheaf.*

Proof. The map $f : X \to \mathbb{R}$ defined above has as fibers the spaces X_r. Because it is stratifiable with finitely many strata, we have the following decomposition of the codomain:

$$(-\infty, 0) \leftarrow \{0\} \to (0, t_1) \leftarrow \{t_1\} \to (t_1, t_2) \leftarrow \{t_2\} \to (t_2, t_3) \cdots$$

The points t_i indicate the radii (the "times") where the topology of the union of the balls changes. Since the fiber $X_{t_i} := f^{-1}(t_i)$ is a closed union of strata, proposition 5.20 implies (after first choosing a system of control data and then regarding X as Thom-Mather stratified) that we can fix an $\epsilon > 0$ such that the neighborhood $U_{t_i}(\epsilon) = \cup_{\sigma \in X_{t_i}} T_\sigma(\epsilon/2)$ contains X_{t_i} as a weak deformation retract. Since f is proper, we claim that there exists a point $s_i^- \in (t_{i-1}, t_i)$ such that $X_{s_i^-}$ is contained in $U_{t_i}(\epsilon)$. Suppose for contradiction that for all $n >> 0$ there exists a point $x_n \in f^{-1}([t_i - \frac{1}{n}, t_i - \frac{1}{n+1})) \cap U_{t_i}(\epsilon)^c$. If this is possible, then $\{x_n\}$ defines a sequence with no convergent subsequence, which contradicts the fact that $f^{-1}([t_{i-1}, t_i])$ is compact. Consequently, there exists an n such that if $s_i^- := t_i - \frac{1}{n}$, then $f^{-1}([s_i^-, t_i]) \subseteq U_{t_i}(\epsilon)$. The composition of the inclusion followed by the retraction

$$\begin{tikzcd} & U_{t_i}(\epsilon) & \\ X_{s_i^-} \arrow[ur, hook] & & X_{t_i} \arrow[ul, hook', "\cong"'] \end{tikzcd}$$

allows us to define maps between the homology of the typical fiber over (t_{i-1}, t_i) to the homology of the fiber X_{t_i}.

$$H_i(X_{s_i^-}; k) \to H_i(X_{t_i}; k)$$

An analogous argument allows us to find an $s_i^+ \in (t_i, t_{i+1})$ such that $X_{s_i^+} \subset U_{t_i}(\epsilon)$. We can construct a vector field on (t_{i-1}, t_i) that flows from the point s_{i-1}^+ to s_i^-. Lifting this vector field to a controlled one over this one, allows us to flow the fiber over s_{i-1}^+ to the fiber over s_i^-, thus realizing the homeomorphisms $X_{s_{i-1}^+} \cong X_{s_i^-}$ explicitly. For convenience, we drop the decorations and choose any point $s_i \in (t_{i-1}, t_i)$ to get our modified version of equation 5.42:

$$\cdots \leftarrow H_i(X_{s_i}; k) \to H_i(X_{t_i}; k) \leftarrow H_i(X_{s_{i+1}}; k) \to H_i(X_{t_{i+1}}) \leftarrow \cdots$$

One should note that this diagram is contravariant with respect to the poset indexing the stratification of $\mathbb{R}$, thus we have constructed geometrically a cellular cosheaf. □

[23] We are indebted to Bob MacPherson for first suggesting to us how the apparatus of stratification theory is applicable to point-cloud data.

Corollary 5.44. *The semialgebraic function* $f : X \to \mathbb{R}$ *in example 5.41 defines, for each* i*, a cellular cosheaf.*

Although the above construction may appear convoluted, it is geometrically natural. Instead of using the order on $\mathbb{R}$ to get a diagram of vector spaces and maps, we have a diagram indexed by the pieces of a stratification of $\mathbb{R}$. This new diagram is specifically adapted to the topological changes in the family $\{X_r\}$.

In **multi-dimensional persistence** we imagine the need for more than one parameter to distinguish features in a point cloud. The traditional story of persistence no longer applies since $\mathbb{R}^n$ for $n \geqslant 2$ has no natural (partial) order. In contrast, every situation where multi-dimensional persistence can be treated as a stratified map (which is effectively always), the partial order of the pieces in a stratification presents itself as a most natural candidate.

However, the geometry of stratified spaces in more than one dimension is subtle and a poset will not always suffice. In section 5.1.5, we will introduce a small category (usually equivalent to a finite one) that allows us to track persistent features in a more careful way. The proof of lemma 5.43 contains the essential ideas of this more general picture. By considering certain definable paths in the parameter space, and analyzing their inverse images, which will be definable, we can try to reduce a multi-dimensional problem to a one-dimensional one. This is the high-level outline of how theorem 5.72 associates a constructible cosheaf to a general definable map.

We conclude with the definition Miller and van den Dries proposed in section 1 of [90]. This definition allows us to verify definability locally, and allows us to work inside manifolds other than $\mathbb{R}^n$.

Definition 5.45 (Analytic-Geometric Categories)**.** A **analytic-geometric category** $\mathcal{G}$ is given by assigning to each analytic manifold M a collection of subsets $\mathcal{G}(M)$ such that following conditions are satisfied:

1. $\mathcal{G}(M)$ is a boolean algebra of subsets of M, with $M \in \mathcal{G}(M)$.
2. If $A \in \mathcal{G}(M)$, then $A \times \mathbb{R} \in \mathcal{G}(M)$.
3. If $f : M \to N$ is a proper analytic map and $A \in \mathcal{G}(M)$, then $f(A) \in \mathcal{G}(N)$.
4. If $A \subseteq M$ and $\{U_i\}_{i \in \Lambda}$ is an open covering of M, then $A \in \mathcal{G}(M)$ if and only if $A \cap U_i \in \mathcal{G}(U_i)$ for all $i \in \Lambda$.
5. Every bounded set in $\mathcal{G}(\mathbb{R})$ has finite boundary.

Remark 5.46. This defines a category in the usual sense. An object of $\mathcal{G}$ is a pair (A, M) with $A \in \mathcal{G}(M)$. A morphism $f : (A, M) \to (B, N)$ is a continuous map $f : A \to B$ whose graph

$$\Gamma(f) := \{(a, f(a)) \in M \times N \mid a \in A\}$$

is an element of $\mathcal{G}(M \times N)$.

The category of $\mathcal{G}$-sets and $\mathcal{G}$-maps, although we will prefer to use the term "definable," has all the properties one could desire, including being closed under inverse images (as long as the domain is closed, see [90], D.7) and Whitney stratifiability ([90] D.16).[24]

5.1.4 Local Systems and Constructibility

In this section we develop the algebraic groundwork for understanding constructible cosheaves more generally. This section is mostly free of stratification theory. We only assume for this section that X is a paracompact Hausdorff space that is locally path connected and locally simply connected. Here we develop the theory of locally constant cosheaves and show that these are equivalent to representations of the fundamental groupoid $\pi_1(X)$, which we call local systems. We finish the section by giving examples of local systems coming from fiber bundles. Since stratified maps are "glued together fiber bundles," this motivates the study of constructible sheaves and sheaves.

Definition 5.47. The sheaf A_X or cosheaf $\hat{A}_X$ on X valued in $\mathbf{Vect}$ is **constant** with value A if for every open set U they make the following assignments:

$$A_X : U \rightsquigarrow A^{\times \pi_0(U)} \qquad \hat{A}_X : U \rightsquigarrow A^{\oplus \pi_0(U)}.$$

A sheaf F or cosheaf $\hat{F}$ is **locally constant** if for each point x there is an open neighborhood U such that F or $\hat{F}$ is constant, i.e. there is a vector space A such that $F|_U \cong A_X$ or $\hat{F}|_U \cong \hat{A}_X$.

As a consequence of this definition and the topological assumptions on X, a locally constant sheaf or cosheaf possesses for each point $x \in X$ a collection of connected neighborhoods containing x all of which take identical values. As a consequence $F(U) \to F_x$ or $\hat{F}_x \to \hat{F}(U)$ is an isomorphism. Moreover, for any other point x' contained in U, the stalk or costalk at x' can be chosen to be isomorphic to $F(U)$ or $\hat{F}(U)$ respectively. By chaining together these sorts of isomorphisms, one can show the following classical theorem:

Theorem 5.48. *Suppose X is a locally path connected, locally simply-connected paracompact Hausdorff space. A locally constant sheaf determines a a local system, where a local system is defined to be a representation of the fundamental groupoid of X, i.e.*

$$\mathcal{L} : \pi_1(X) \to \mathbf{Vect}.$$

Similarly, any locally constant cosheaf valued in $\mathbf{Vect}$ determines a local system.

Proof. By taking stalks or costalks we can define the functor $\mathcal{L}$ on objects $x \in X$ to be F_x or $\hat{F}_x$, respectively. Since the theorem is well known (see [1] for a proof, which we follow here) for sheaves we present the cosheaf-theoretic proof instead.

[24] The authors of [90] acknowledge that there is a gap in the proof of Whitney stratifiability of $\mathcal{G}$-maps, but Ta Lê Loi, among others, has since filled in this gap [51].

Call a subset K of X *fine*[25] for a cosheaf $\hat{F}$ if it is connected and is contained in a connected open set V such that $\hat{F}|_V$ is a constant cosheaf. For any set of points $\{x_i\}$ in a fine set K we have a collection of isomorphisms

$$\begin{tikzcd} & \hat{F}(V) & \\ \hat{F}_{x_i} \arrow[ur, "\pi_i"] & \hat{F}_{x_j} \arrow[u, "\pi_j"] & \hat{F}_{x_k} \arrow[ul, "\pi_k"'] \end{tikzcd}$$

that when composed together allows us to define an invertible map from $\hat{F}_{x_i} \to \hat{F}_{x_k}$ via $\pi_k^{-1}\pi_i$. Of course this map agrees with the composition of the analogously defined map

$$\hat{F}_{x_i} \to \hat{F}_{x_j} \to \hat{F}_{x_k}$$

because $\pi_k^{-1}\pi_j\pi_j^{-1}\pi_i = \pi_k^{-1}\pi_i$.

Now we claim that given a path $\gamma : [0,1] \to X$ there exists a sequence of points $0 = a_0 < a_1 < \cdots < a_n = 1$ so that for all i the set $\gamma([a_i, a_{i+1}])$ is fine. This is the case because every point $\gamma(t)$ possesses a fine neighborhood and by continuity there are open intervals V_t of t such that $\gamma(V_t)$ is fine. If we choose intervals $[a, a']$ contained in each V_t, the interiors of these intervals will form an open cover of $[0,1]$. By compactness, finitely many of these intervals will do. Choosing such a finite list, merging and ordering the endpoints, gives the requested sequence.

From the sequence we can define the map $\rho(\gamma) : \hat{F}_{\gamma(0)} \to \hat{F}_{\gamma(1)}$ to be the composite

$$\hat{F}_{\gamma(a_0)} \to \hat{F}_{\gamma(a_1)} \to \cdots \to \hat{F}_{\gamma(a_n)}.$$

This map is well defined by virtue of the fact that it is invariant under the addition of extra points a' to the sequence above. Consequently, if any different sequence was chosen we could have merged it with this one and deduced that these maps were the same.

A similar argument can be used to show that for homotopies $H : [0,1] \times [0,1] \to X$ there are sequences $\{a_i\}_{i=1}^n$ and $\{b_j\}_{j=1}^m$ so that the sets $H([a_i, a_{i+1}] \times [b_j, b_{j+1}])$ are fine. Using the same concatenation of isomorphisms proves that if γ and γ' are homotopic relative endpoints, then the above defined maps $\hat{F}_{\gamma(0)} \to \hat{F}_{\gamma(1)}$ and $\hat{F}_{\gamma'(0)} \to \hat{F}_{\gamma'(1)}$ are the same. $\square$

This theorem is one direction in an interesting equivalence.

Theorem 5.49. *A representation of a fundamental groupoid $\mathcal{L} : \pi_1(X) \to \mathbf{Vect}$ determines a locally constant sheaf and a locally constant cosheaf respectively.*

Proof. Again the sheaf theoretic version of this statement is well known (see [1]), so we carry out the cosheaf version. Assume we have a local system $\mathcal{L}$, then we define the associated cosheaf to be

$$\hat{L} : U \rightsquigarrow H_0(U; \mathcal{L}) := \varinjlim_{x \in U} \mathcal{L}|_U,$$

[25] In [1] they use the word "good."

which is a cosheaf on account of the fact that colimits commute with colimits. The fact that $\hat{L}$ is locally constant comes from the fact that each point x has a simply connected neighborhood U for which the local system $H_0(U;\mathcal{L}) \cong \mathcal{L}(x)$ for any $x \in U$.

Although it is not pointed out in the literature, the classical proof for sheaves follows by making the exact dual assignment.

$$L : U \rightsquigarrow H^0(U;\mathcal{L}) := \varprojlim_{x\in U} \mathcal{L}|_U$$

□

Remark 5.50 (Alternative Proof). The introduction of an apparently superfluous $H_0(-;\mathcal{L})$ is an invocation of the principle that H_0 is a cosheaf. This principle, expressed in theorem 3.10, actually states that "H_0 for any homology theory that satisfies Mayer-Vietoris is a cosheaf." This is true again for this case, but it requires that the reader know that local systems allow us to define a homology theory with "twisted coefficients." This theory, which uses singular chains with coefficients determined by $\mathcal{L}$, satisfies the Eilenberg-Steenrod axioms ([94] Ch. 6) and thus Mayer-Vietoris ([84] Ch. 4.6). To complete our alternative proof of the above theorem, we check one more hypothesis of theorem 3.10. We observe that for an upward increasing sequence of open sets $\{U_i\}$ we have a directed system of chain complexes of twisted singular chains whose right most end point looks like the (non-exact) sequence

$$C_1(U_i;\mathcal{L}) \to C_0(U_i;\mathcal{L}) \to 0.$$

Taking H_0 involves only taking a cokernel, which commutes with direct limits. This proves that $H_0(-;\mathcal{L})$ is a cosheaf.

We leave it to the reader to prove that this defines an equivalence between local systems and locally constant sheaves and cosheaves. Moreover it is interesting to note that local systems define an abstract equivalence between locally constant sheaves and cosheaves gotten by taking H^0 and H_0 respectively.

One of the flexibilities offered by working with paths and points is that we get get more examples of locally constant cosheaves.

Proposition 5.51 (Fiber Bundles Give Local Systems). *Suppose* $\pi : E \to B$ *is a fiber bundle, then for each* i *the homology of the fiber* $H_i(\pi^{-1}(b);k)$ *defines a local system.*

Proof. This is easily seen because any path $\gamma : [0,1] \to X$ determines a pullback bundle γ^*E which is trivial, so there is an isomorphism $H_i(\pi^{-1}(\gamma(0));k) \to H_i(\pi^{-1}(\gamma(1));k)$. Moreover, any homotopy of paths $H : [0,1]^2 \to X$ determines a trivial pullback bundle H^*E. □

Example 5.52 (Torus). Using the equivalence between locally constant cosheaves and local systems one can use the above proposition to create numerous examples of cosheaves. For example, the torus $T = S^1 \times S^1$ defines a trivial circle bundle over S^1 by projection onto the first factor. Taking H_1 of the fiber gives a topological model for the constant cosheaf on S^1. See figures 12 and 13.

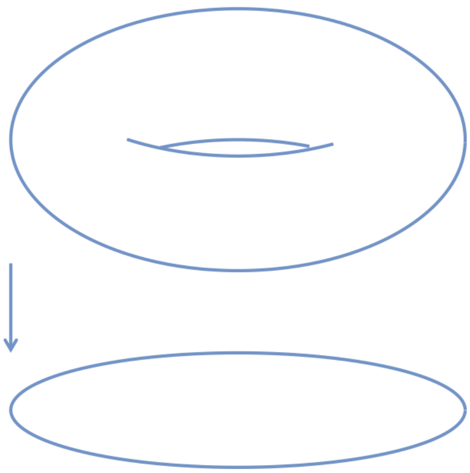

Figure 12: Trivial Circle Bundle

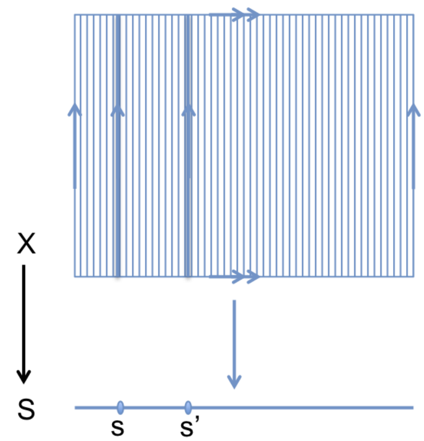

Figure 13: Identification Space Model

Example 5.53 (Klein Bottle)**.** Let's modify the previous example to see what happens with a non-trivial circle bundle over the circle. Written using the identification space model, we see explicitly how the fiber must twist in the Klein bottle. Specifically, imagine that we consider two points s and s' in the base. If we consider any path between s and s' that does not wind around the circle, then the action on the fiber can be chosen to be trivial; see figure 14. However, if we consider a path from s to itself that laps once around the circle, then the action on H_1 is given by multiplying by -1 as it reverses the orientation of the fiber; see figure 15.

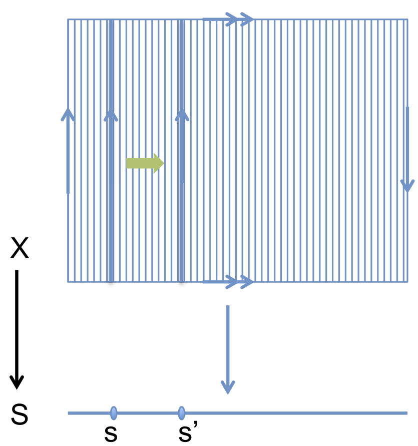

Figure 14: Trivial Action on Fiber

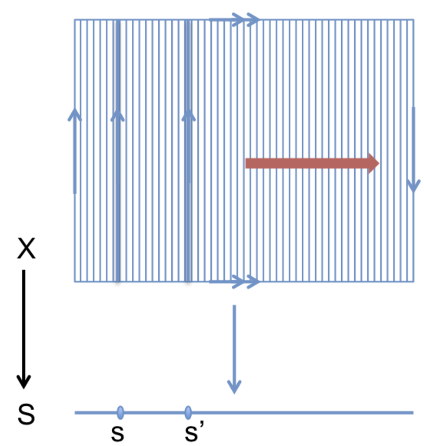

Figure 15: Non-trivial Action on Fiber

As already described, stratified maps $f : X \to Y$ consist of "glued together fiber bundles," i.e. there is a partition of the codomain Y into pieces over which the map looks like a fiber bundle. By proposition 5.51, if we take the ith homology of the fiber $f^{-1}(y)$, we get the structure of several disjoint local systems over Y. We'd like to synthesize these disjoint local systems into a single device. This device is a constructible cosheaf.

Definition 5.54. A cosheaf $\hat{F}$ on X is **constructible** if there exists a decomposition $X_\bullet$ of X into pieces X_σ such that $\hat{F}|_{X_\sigma}$ is locally constant. We say that $\hat{F}$ is constructible with respect to the decomposition $X_\bullet$. The same definition holds for sheaves F.

By the definition of a stratified map, the codomain X can be decomposed into pieces over which the map f defines a fiber bundle. By the earlier example and theorem, taking the homology of the fiber over these pieces defines a local system on each piece and consequently, via the equivalence, a locally constant cosheaf on each piece. This would define a constructible cosheaf with the one annoying detail that we have not defined a cosheaf $\hat{F}$ on the whole space. There are two possible resolutions to this problem:

- For sheaves, this result is well known, but it requires developing the derived category and proving essential "base change theorems"; see [77] section 4.2.2. One could try to re-develop this machinery for cosheaves, but it has not yet been done.
- One could prove geometrically that a stratified map induces a representation of a category associated to X analogous to the fundamental groupoid. In similar spirit to the proof above, one could show that such a representation defines a cosheaf in an open set sense.

We will take this second approach. When Bob MacPherson relayed (without proof) the observation that taking ith homology of a stratified map defines a constructible cosheaf in a lecture at the IAS, the open set version of cosheaves was not under consideration. A completely different definition that was only suited for the constructible setting was used. MacPherson's vantage point was more geometric and relied on a suitably refined stratified analog of the fundamental groupoid, which we decided to call the **entrance path category**. Constructible cosheaves were *defined* as representations of the entrance path category instead. In the next section we develop this perspective and also use the open set version of cosheaves. This will complete one way of proving that

Cellular sheaves and cosheaves are bona fide sheaves and cosheaves constructible with respect to a cell structure.

5.1.5 Representations of the Entrance Path Category

The utility of cosheaves is limited by its abstract quality and by a lack of computability. Section 5.1.4 demonstrated that locally constant cosheaves can be described by assigning vector spaces to points and linear maps to homotopy classes of paths. This description allows us to compress the amount of information that encodes a locally constant cosheaf. As described below, for a connected space X, a representation of the fundamental groupoid $\mathcal{L} : \pi_1(X) \to \mathbf{Vect}$ is equivalent to a representation of the fundamental group of X.

$$\pi_1(X; x_0) \to \mathbf{Vect}$$

In other words, a locally constant cosheaf is equivalent to a $\pi_1(X, x_0)$-module, which is more easily programmed on a computer. For constructible cosheaves, we will introduce the entrance path category as the stratified analog of $\pi_1(X)$. To gain an analogous compression of information, we will outline how to pass to an equivalent, yet smaller, version of the entrance path category. We will then prove that constructible cosheaves are equivalent to representations of the entrance path category.

In addition to compression of information, there are geometric considerations motivating this equivalence. Just as fiber bundles give rise to locally constant cosheaves, we'd like to prove that stratified maps give rise to constructible cosheaves. If we are successful in this task, we bring cosheaf theory in contact with a large area of mathematics:

- **Morse theory:** Morse functions are just particular instances of stratified maps $f : M \to \mathbb{R}$.
- **Picard-Lefschetz theory:** The complex analog of Morse theory studies algebraic maps $\pi : X \to \mathbb{C}$, which are necessarily stratified.
- **Point Cloud Data and Persistence:** Semialgebraic families are described by semi-algebraic maps, which are stratified.

In the case of Morse theory, or any situation with a Whitney stratified map to a one-dimensional space, the proof of lemma 5.43 goes through and thus we do get a cosheaf. In the case of maps to higher-dimensional spaces, we can only treat definable ones. Thus, we get representations of a **definable entrance path category**, which also defines a cosheaf. It is unclear to us whether this is equivalent to the full entrance path category.

Given a Whitney (or Thom-Mather) stratified space, the entrance path category looks very much like the fundamental groupoid. It has objects that are points and morphisms that are paths. However, the paths and homotopies must respect the stratification. A path may wind around in a given stratum and it may enter deeper levels of the stratification, but upon doing so, it may never return to its higher level.[26]

Definition 5.55. Let $(X, \{X_\sigma\}_{\sigma \in P_X})$ be a Whitney (or Thom-Mather) stratified space. We define the **entrance path category** $\mathbf{Entr}(X, \{X_\sigma\})$ to be the category whose objects are points of X and whose morphisms are homotopy classes of paths $\gamma(t)$ whose ambient dimension (the pure dimension of the containing stratum) is non-increasing with t. We call such paths **entrance paths**. Moreover we require the homotopies $h(s, t)$ to be entrance paths for every fixed s. We write $\mathbf{Entr}(X)$ when a given stratification is understood.

Opposite to the entrance path category is the **exit path category**, written $\mathbf{Exit}(X)$ whose objects are the same, but whose paths ascend into higher dimensional strata.

Remark 5.56 ("Tame" Homotopies). David Treumann's thesis [89], which was written under MacPherson's direction, contains one of the first published accounts of the exit path category.

[26] "Fácilis descensus Averno; noctes atque dies patet atri ianua Ditis; sed revocare gradum superasque evadere ad auras, hoc opus, hic labor est."

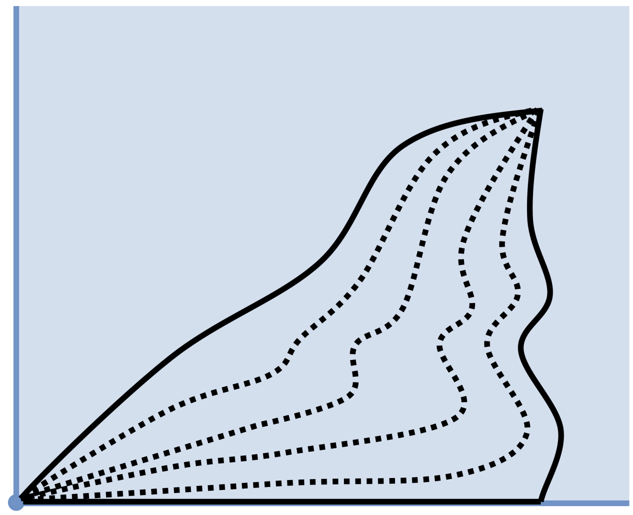

Figure 16: Two Entrance Paths and a Homotopy Between Them

However, he added an additional hypothesis that the homotopies should be "tame," which he defines by saying that $h : [0,1]^2 \to X$ should admit a triangulation of $[0,1]^2$ such that the interior of each simplex in the triangulation is contained in some stratum of X. Jon Woolf [98] uses a version of the exit and entrance path category based on Quinn's theory of homotopically stratified spaces and does not require Treumann's tameness assumption. Homotopically stratified spaces are more general than Whitney or Thom-Mather stratified spaces, so we may invoke some of Woolf's results. Nevertheless, Treumann's modification foreshadows our own.

Definition 5.57 (Definable Entrance Path Category)**.** For a fixed analytic-geometric category $\mathcal{G}$ we can consider the **definable entrance path category** to have the same objects as before, but whose morphisms are definable entrance paths, where identify entrance paths related by definable homotopies $h : I^2 \to X$. There should be a triangulation of I^2, so that the image of every open cell is contained in some stratum of X. This category will be written $\mathbf{Entr}_{\mathcal{G}}(X, \{X_\sigma\})$. Dually, we have a definable exit path category $\mathbf{Exit}_{\mathcal{G}}(X)$.

Remark 5.58. We will not need to use the definable entrance path category until theorem 5.72, so one may temporarily ignore this restrictive definition.

From the perspective of a computer, the entrance path category definition is entirely too unwieldy to be useful. Storing the points of any space we are accustomed to thinking about (circles, tori, Klein bottles, etc.) is simply too much data to consider. Fortunately, these categories are equivalent to much simpler categories. After all, if two points x and x' lie in the same stratum, then there is a path from x to x'. Following this path and then following back along it defines a contractible loop based at x, which corresponds to the identity morphism id_x. Thus x and x' are isomorphic as objects in this category. This argument is the same one that shows that if a space is connected, then for any choice of basepoint $x_0 \in X$ we have

$\pi_1(X) \cong \pi_1(X, x_0)$. We offer the analogous statement for more general stratified spaces by explaining a general principle first.

Definition 5.59 (Skeletal Subcategory)**.** Suppose $\mathcal{C}$ is a category, then a subcategory $\mathcal{S}$ is **skeletal** if the inclusion functor is an equivalence, and no two objects of $\mathcal{S}$ are isomorphic.

If $\mathcal{C}$ is small, then we can describe explicitly how to construct a skeletal subcategory $\mathcal{S}$. On the objects of $\mathcal{C}$ we define an equivalence relation that says $x \sim x'$ if and only if x and x' are isomorphic. To define a skeletal subcategory we pick one object $x \in \bar{x}$ from each equivalence class and define the morphisms to be $\mathbf{Hom}_{\mathcal{S}}(\bar{x}, \bar{y}) := \mathbf{Hom}_{\mathcal{C}}(x, y)$.

Since a skeletal subcategory is *equivalent* to its surrounding category, we get a compressed representation of the entrance path category by choosing a single point from each connected component in the stratification.

Example 5.60 (Entrance Path Category for S^1)**.** Now consider the circle S^1 stratified as a single pure stratum. The argument above shows that we can view the entrance path category of S^1 as equivalent to the fundamental group $\pi_1(S^1, x_0)$. This is a category with a single object $\star$ whose Hom-set corresponds to a loop for each homotopy class of path, i.e. $\mathbf{Hom}(\star, \star) \cong \mathbb{Z}$.

Example 5.61 (Manifolds)**.** More generally, if the space X is a manifold, stratified as a single pure stratum, then the entrance path category is equivalent to the fundamental group.

If we believe MacPherson's characterization of constructible (co)sheaves, then we can reach our much sought after explanation of why cellular sheaves and cosheaves are actually sheaves and cosheaves. Part of the explanation rests on the following characterization of the entrance path category for cell complexes.

Proposition 5.62 (Entrance Path Category for Cell Complexes)**.** *If $(X, \{X_\sigma\}_{\sigma \in P_X})$ is stratified as a cell complex, then each stratum is contractible and there is only one homotopy class of entrance paths between any two incident cells. As such*

$$\mathbf{Entr}(X) \cong \mathbf{Cell}(X)^{op} = P_X^{op} \qquad \text{and} \qquad \mathbf{Exit}(X) \cong \mathbf{Cell}(X) = P_X.$$

To prove this proposition, we need a better understanding of the entrance path category. To do so, we pick out a distinguished class of entrance paths.

Definition 5.63 (Homotopy Link)**.** Suppose X is a decomposed space and $X_\sigma \leqslant X_\tau$ are two incident pieces. The **homotopy link** of X_σ in X_τ is defined to be the space of paths $\gamma : I \to X_\sigma \cup X_\tau$ such that $\gamma([0,1)) \subset X_\tau$ and $\gamma(1) \in X_\sigma$, i.e. it is the space of paths that enter X_σ at the last possible moment.

We now adapt a proof of Jon Woolf's ([98], Lemma 3.2) to our situation.

Lemma 5.64. *Let $(X, \{X_\sigma\})$ be a Thom-Mather stratified space. Any entrance path is homotopic through entrance paths to an element of the homotopy link.*

Proof. Suppose $\gamma : [0,1] \to X$ is an entrance path. By compactness, it can only intersect finitely many pieces in the stratification of X. We write X^j to denote the union of all dimension j pieces. For any $i \leqslant j$, we have that $X^i \leqslant X^j$.

We claim that one can show that every entrance path γ starting in a stratum X^k and ending in a stratum X^i that intersects potentially every stratum in between

$$X^k \geqslant X^{k-1} \geqslant \cdots \geqslant X^i$$

is homotopic to a path $\gamma'(t)$ which sends every $t \in [0,1)$ to X^k and then enters X^i at the last possible moment.

To define the homotopy, one focuses on pulling the path off the last stratum X^j that γ enters before entering X^i, i.e. there is a partition $0 < t_1 < \cdots < t_n < 1$ of $[0,1]$ such that $\gamma(0) \in X^k$, $\gamma([t_n,1]) \subseteq X^i$ and $\gamma([t_{n-1},t_n]) \subseteq X^j$. First, we show that we can pull the path off X^i into X^j so that it enters only at $t=1$. The schematic uses the fundamental observation that stratified spaces can be treated locally as a system of fiber bundles.

Pick a point $x_j := \gamma(t_n - \epsilon) \in X^j$ and consider its homeomorphic image (which we call x_j') in the fiber over $\gamma(t_n)$. There is a homotopy from the path γ relative the end points $x_j = \gamma(t_n - \epsilon)$ and $\gamma(t_n)$ to the piece-wise path that is constant in the fiber, connects x_j to x_j', and then heads straight to $\gamma(t_n)$ while staying in the fiber over that point. By the path lifting property for fiber bundles, we can consider a lift of $\gamma([t_n,1])$ starting with x_j' which ends at x_j'' in the fiber over $\gamma(1)$. Repeating the same argument, we can then consider a path that heads from x_j'' to $\gamma(1)$ while staying in the fiber. Now the path enters the stratum X^i at the last possible moment.

Repeating this argument and using the conical structure of the fiber, allows us to lift the path out of the X^j stratum and into higher ones. □

This result allows us to take representative entrance paths that are easy to understand. Every element of the homotopy link is an entrance path, but not every entrance path is an element of the homotopy link. Moreover, it is not clear that two paths that are homotopic as entrance paths are homotopic as entrance paths (after we have moved them into the link as in the above proof). Fortunately, David Miller has recently shown this is the case [63]. At a high level, this provides a proof of proposition 5.62, which can also be seen using easier methods.

Proof of 5.62. Since the pieces in a cell structure on X are all contractible, each cell X_σ has a single path component in its homotopy link in X_τ. Thus the skeleton of the entrance path category for a cell complex is

$$\mathbf{Entr}(X) \cong X^{op},$$

which was wanted. □

If we can show that the entrance path category can be built up locally, then we can prove that representations of this category define cosheaves. The ability to build up locally the entrance path category is the van Kampen theorem adapted to stratified spaces. Ostensibly, David Treumann's published version of his thesis [89] proves the van Kampen theorem for

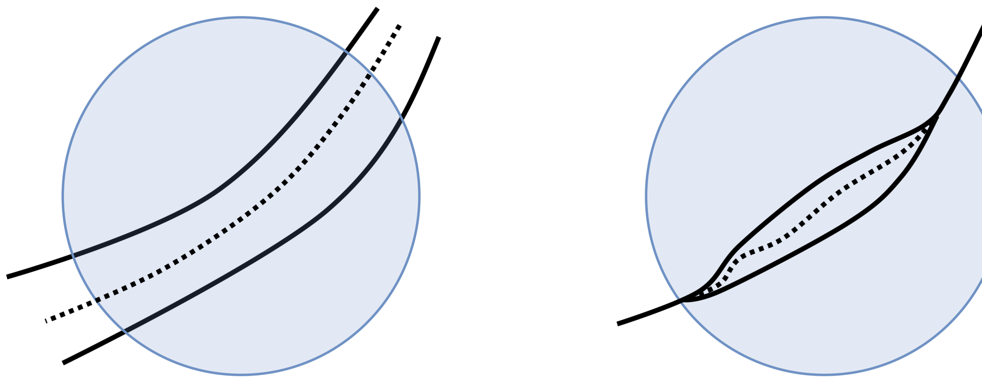

Figure 17: Modifying a Homotopy to Stay Inside an Open Set

the exit path 2-category, but the elegant inductive argument in proposition 5.9 appears to have an error.[27] Jacob Lurie has a proof for the ∞-category case [54]. Jon Woolf has outlined another argument [97] based on his classification of $\mathbf{Set}$-valued representations of the entrance path category as branched covers. The following proof, joint with Dave Lipsky, is more direct and algorithmic, but less elegant in many respects.

The main difficulty in proving the van Kampen theorem is that given a cover, a homotopy of entrance paths restricts to a free homotopy between entrance paths and not a homotopy relative endpoints; see figure 17. In contrast to the fundamental groupoid, we cannot freely add paths to make this homotopy respect endpoints, the entrance path property must be preserved and this significantly complicates the proof. We borrow Treumann's idea of using a triangulation T of I^2 such that $h : I^2 \to X$ sends open cells of T to strata of X. We then, after sufficient refinement, define a homeomorphism of I^2 that allows us to treat the triangulation as a piecewise linear one. For a piecewise linear triangulation, we outline an explicit algorithm for replacing the homotopy h with a composition of homotopies preserving endpoints, each of which is supported on a triangle in I^2. Let us state our desired version of the van Kampen theorem and give the first step of the proof.

Theorem 5.65 (Van Kampen Theorem for Entrance Paths)**.** *If X is a Whitney stratified space and $\mathcal{U} = \{U_i\}$ is a cover, then*

$$\mathbf{Entr}(X) \cong \varinjlim_{I \in N(\mathcal{U})} \mathbf{Entr}(U_I).$$

Each open set is given the induced stratification from the whole space. We assume that every homotopy $h : I^2 \to X$ admits a triangulation of the domain so that for each open cell in the triangulation there is a stratum of X that contains its image. Moreover, the same result holds for the definable entrance path category.

[27] The argument inducts on the number of triangles in a triangulation of I^2. The statement that the closure of the complement of a single triangle is homeomorphic to I^2 is not true if, for example, the triangle has two vertices on one side of the square and the third on another side.

Proof. The colimit is an ordinary colimit in the category of all categories. The diagram that sends each $I \in N(\mathcal{U})$ to $\mathbf{Entr}(U_I)$ we will call V. We already know that the inclusions of the open sets $U_I \hookrightarrow X$ induce functors $\phi_I : \mathbf{Entr}(U_I) \to \mathbf{Entr}(X)$ and that these define a cocone $\phi : V \Rightarrow \mathbf{Entr}(X)$, i.e. a natural transformation from V to the constant diagram on $N(\mathcal{U})$ with value $\mathbf{Entr}(X)$. Now suppose $\phi' : V \Rightarrow \mathcal{C}$ is another cocone. We need to check that there exists a unique map $u : \mathbf{Entr}(X) \to \mathcal{C}$ that makes all the functors commute, i.e. $u \circ \phi_I = \phi'_I$.

On objects, $u(x) := \phi'_i(x)$ for whatever open set U_i contains x. The choice doesn't matter since if U_j also contains x, then the functor defined on the intersection causes $\phi_i \circ \phi_{ij}(x) = \phi_j \circ \phi_{ij}(x)$. Now we must define $u(\gamma)$ for γ an entrance path in X. By compactness, we can pass to a finite subcover of $\{U_i\}$ to cover the path γ. We can break up γ into shorter paths $\gamma_{i_1}, \ldots, \gamma_{i_n}$, each of which lie in some element of the cover. We define $u(\gamma) := \phi'_{i_n}(\gamma_{i_n}) \circ \cdots \circ \phi'_{i_1}(\gamma_{i_1})$. We must show that this definition is invariant under homotopy to complete the proof. This is accomplished by lemma 5.68 together with proposition 5.69. □

Definition 5.66. Call a homotopy **U-elementary** if there is an interval $[a, b] \subset I$ such that $h(s, t)$ is independent of s so long as $t \notin [a, b]$ and the image of $I \times [a, b]$ under h is contained in U. See figure 17 for an illustrative cartoon.

Remark 5.67. We will use a slightly strange way of orienting the unit square $I^2 = [0, 1] \times [0, 1] \ni (s, t)$. The "top edge" is the edge where $t = 0$ and the "bottom edge" is the edge $t = 1$. We will use this language because an entrance path enters "deeper" levels of a stratification.

Lemma 5.68. *Let X be a Whitney stratified space along with a cover $\mathcal{U}$ and let $\alpha(t)$ and $\beta(t)$ be entrance paths with the same start and end points. Let $h : I^2 \to X$ be a homotopy (relative endpoints) through entrance paths connecting $\alpha(t) = h(0, t)$ to $\beta(t) = h(1, t)$. If I^2 admits a piecewise-linear triangulation T such that every open cell in T is mapped to a stratum of X, then we may define a sequence of new homotopies $h_1, \ldots, h_n : I^2 \to X$, each of which are elementary for some element of the cover, so that the composite connects $\alpha \simeq \beta$. Informally speaking, each homotopy h_i will be supported on a single triangle in the barycentric subdivision of the triangulation T.*

Proof. Since the image of I^2 is compact, a finite subcover of $\mathcal{U}$ will do. After sufficient refinement, we can assume that each triangle in T is contained in some element of the subcover. By taking the barycentric subdivision T', we can refer to the vertices of any triangle in T via barycentric labels v, e, f depending on whether the vertex is at the barycenter of a vertex, edge or face in the original triangulation. Since each open cell in T is mapped to a stratum of X, the triangles in T' satisfy the following fundamental property: $h(\sigma_f)$, where $\sigma_f := [v, e, f] - [v, e]$, is contained in some stratum X_f; $h(\sigma_e)$, where $[v, e] - v = \sigma_e$, is contained in X_e; $h(v)$ is contained in X_v and $X_f \geqslant X_e \geqslant X_v$. We will refer to the dimension of these containing strata as the "dimension" of f, v and e, respectively.

By the fundamental property of triangles in T' we know, for example, that the path parameterized by going from f to e to v along the boundary of a triangle is a valid entrance path and this is homotopic through entrance paths to one that goes from f directly to v. This is the prototypical "move" that we will use to define a given h_i in our new homotopy between α to β. By reparameterizing the triangle, this move defines an elementary homotopy of entrance paths.

As a preparatory step we replace the entrance path $\alpha(t) := h(0,t)$ with the path that starts at $(s,t) = (1,0)$ and goes along the top edge of the square to $(0,0)$, then to $(0,1)$ and finally to $(1,1)$. Because the homotopy is constant along the top and bottom edges, this only affects the parameterization of the path, but now our modified path and $\beta(t) := h(1,t)$ share the same endpoints in I^2. We will now refer to our intermediate paths γ by a sequences of vertices in T', written $w_1 \cdots w_n$, which taken two at a time define edges $\gamma_i = w_a w_b$ labelled by a pair of letters fv or vf, fe or ef, ev or ve. Observe that if the image of vf under h is a valid entrance path, then this implies that $\dim v = \dim e = \dim f$ for the triangle containing that particular edge $v_i f_i$.

If an entrance path γ ever has a vertex appear twice in its list, then this indicates a loop that must be contained in the same stratum. By virtue of the fact that I^2 is simply connected, the portion of the path between the repeated vertices can be homotopically reduced to the constant path via the argument used to prove the van Kampen theorem for the fundamental groupoid. We will avail ourselves of this operation, which we call the **fundamental groupoid sweep** F. For example, if γ contains $\cdots e_i v_i e_i \cdots$ in its list of visited vertices, then $F(\gamma)$ will replace the portion $e_i v_i e_i$ with just e_i. Of course, $F^2 = F$.

We retain the (s,t) coordinates to determine valid moves in our homotopy. We do this because, by assumption, for all s, $h(s,t)$ is an entrance path in t and thus the dimension decreases in that direction. Now we can describe our algorithm:

If $F(\gamma) = \gamma$ and $s(w_i) = s(w_j)$ for all $i \neq j$, then we are done. Otherwise, apply F and starting with γ_1, ask of γ_i if there is a triangle to the left (with respect to the induced orientation of following the path) and apply one of following rules:

(1a) If $\gamma_i = vf$ or fv, then replace γ_i with $\gamma_i' := vef$ or fev, where e belongs to the triangle to the left.

(1b) If $\gamma_i = ev$ and $s(e) > s(v)$ or if $\gamma_i = ve$ and $s(v) > s(e)$, then $\gamma_i' := efv$ or $\gamma_i' := vfe$ where f belongs to the triangle to the left.

(1c) If $\gamma_i = fe$ and $s(f) \leqslant s(v) \leqslant s(e)$ or if $\gamma_i = ef$ and $s(e) \leqslant s(v) \leqslant s(f)$ where v belongs to the triangle to the left, then $\gamma_i' := fve$ or $\gamma_i' := evf$.

If none of the above apply, consider adjacent paths $\gamma_i * \gamma_{i+1}$ two at a time and ask if the following rule is applicable:

(2) If $\gamma_i * \gamma_{i+1} = fev$ or vef where the interior of the triangle is kept to the left, then $(\gamma_i * \gamma_{i+1})' = fv$ or vf.

After each application of a rule, one must check whether $F(\gamma) = \gamma$ and $s(w_i) = s(w_j)$ and repeat as many times as necessary. The algorithm must terminate by virtue of the fact that each step reduces the number of triangles to the left.

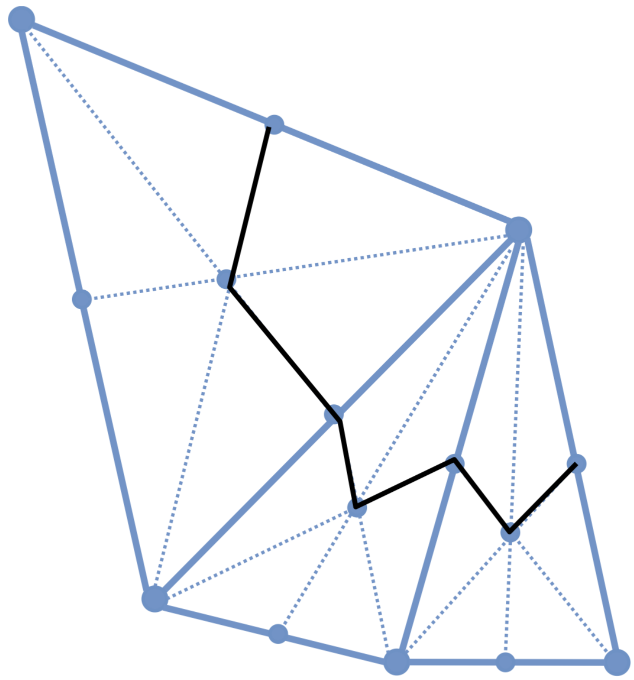

Figure 18: Forcing Move (1c) to Apply

Observe that the only way for a path γ_i not to have a triangle to its left is if it lies on the boundary of I^2 and it is following the boundary clockwise. If γ_i does not belong to the $s = 1$ edge, then that contradicts the assumption that $F(\gamma) = \gamma$ as the total path must return to the point $(s,t) = (1,1)$. If the edge does lie on $s = 1$, then part of the desired homotopy has been achieved and it need not be moved. If there are no triangles to the left and $F(\gamma) = \gamma$, then the algorithm has finished.

The rationale for rule (1b) is that any point p on the edge ev determines an entrance path $h(s(p),t)$, which drops into the interior σ_f of the triangle to the left, thus bounding the dimension of f by the dimension of e. The rule (1c) uses similar reasoning. If $s(f) \leqslant s(v) \leqslant s(e)$ where v is the triangle to the left, then the entrance path determined by v $h(s(v),t)$ flows into σ_f or σ_e thus bounding the dimension of e by the dimension of v. Let us now prove the correctness of the algorithm.

Suppose γ has a vf of fv in sequence. Since a f vertex cannot belong to the boundary of I^2, this implies that there is a triangle to the left and that rule (1a) can be applied. Thus, to show that at least one move can be applied up until the algorithm finishes, we assume that no vf's or fv's appear in γ. Suppose γ consists of only e's and v's. Since the start of γ has $s(v) = 1$, having s non-decreasing would imply that γ is contained in $s = 1$ and the algorithm would be finished. Otherwise there is a pair such that (1b) can be applied. Now assume that our path has e's, v's and f's, with no fv/vf pairs and such that for all ev/ve pairs, s is increasing.

Because the value of s must go from 1 back to 1, if s is not constant along γ, then there must be at least one s decreasing to increasing turning point. Because γ is piecewise linear, by turning point we mean the shortest adjacent collection of edges $\gamma_i * \cdots * \gamma_k$ where the s value goes from strictly decreasing to strictly increasing, i.e. there is an edge along which s is strictly decreasing, then potentially several edges where s is constant and finally an edge which increases in s. To determine the "handedness" of these turning points we must further specify the t behavior. If the turning point consists of only two edges, then we can ask if the difference in t of the first and last vertex is positive or negative. If the turning point has at least one constant s value edge, then we can use the difference in t along the edge to determine if the turning point is t positive or t negative.

Suppose we have a t negative s decreasing-to-increasing turning point. If the minimal s value is obtained on this turning point, then since t must go from 0 to 1, we can conclude that the path must intersect itself at some point, contradicting the assumption that $F(\gamma) = \gamma$. To avoid self-intersection, there must be at least one t positive turning point. Since there are no decreasing ev/ve pairs, the decreasing edge must be either ef or fe. If the next term is a v, then either rule (1a) or (2) would have to apply, respectively. Thus, we can conclude that the next term is either an f or an e. Inducting on the length of the s constant portion of the turning point and using the fact that f, e and v cannot be collinear, we can show so long as rules (1a), (1b) or (2) cannot be applied, that the s increasing edge has to be an fe/ef edge. Consequently the last two edges in such a turning point is either fef or efe.

Now we visit the last t positive s decreasing-to-increasing turning point. Since the reasoning is so similar, assume that the last two edges are efe. We aim to show that if no other rules are applicable, then the rule (1c) must be applicable for some ef/fe edge. Let us refer to the vertices in the original triangulation T containing these barycenters as v_1, v_2 and v_3, whose s coordinates are s_1, s_2 and s_3 respectively. Since all the vertices cannot be collinear, we let s_2 have the largest s value. The vertex f is the centroid of $\{v_1, v_2, v_3\} \subset I^2$, the first $e = e_{12}$ is the centroid of $\{v_1, v_2\}$ and the next $e = e_{23}$ is the centroid of $\{v_2, v_3\}$. By assumption $1/3(s_1 + s_2 + s_3) = s(f) < s(e_{23}) = 1/2(s_2 + s_3)$, thus if the next vertex visited is v_3, then we can apply rule (1b) to $e_{23}v_3$. If the next vertex visited is v_2, then we can apply rule (2). Thus, we must assume that the next vertex is $f' = 1/3(v_2 + v_3 + v_4)$. Now we reason on $e_{23}f'$. The next vertex cannot be a v, otherwise (1a) could be applied. If the next edge visited is e_{34}, then there are two possibilities. Either $s(e_{34}) < s(f')$, which would contradict the fact that we are at the last t positive turning point, or $s(f') \leqslant s(e_{34})$, which would imply that $2s_2 \leqslant s_3 + s_4$, but this would imply that $s_2 = s(v_2) \leqslant s(f')$ and consequently rule (1c) would apply. Thus, assuming $s(e_{34}) < s(f')$, we must remain in the link of v_2 and proceed to e_{24}. Repeating inductively, and using the fact that there are only finitely many triangles, there must be a point where rule (1c) is applicable; see figure 18. This completes the proof. □

Now we must select out a class of triangulations of the unit square I^2 that can be deformed in an entrance-path preserving way to a PL triangulation.

Proposition 5.69. *Suppose that a (definable) triangulation $\varphi : |K| \to I^2$ by a finite simplicial complex is C^2 when restricted to the edges of $|K|$. There is a (definable) homeomorphism g of I^2 so that after suitable refinement, the triangulation is piecewise-linear.*

Proof. The strategy of the proof is to add additional vertices $\{w_i\}$ to the image $\varphi(e)$ of each edge in I^2 so that the line segment connecting any two adjacent vertices w_0, w_1 is to one side of the curve $\varphi(e)$ between $s(w_0)$ and $s(w_1)$. We will then locally scale the t value in such a way as to push that part of $\varphi(e)$ to the line segment.

To add in these vertices in a principled way, we first consider the critical set of the s value of $\varphi(e)$ for every edge e in $|K|$. We remove the entire critical set from e. If the critical set contains an interval, then we know that portion of the edge is already linear and need not consider it. Now we use the implicit function theorem to write the remainder of the edge $\varphi(e) - \{ds|_{\varphi(e)} = 0\}$ as a function of s. For each of these functions we find the critical set of its first derivative ("inflection points") and remove these as well. What is remaining of $\varphi(e)$ is a collection of open concave and convex arcs, each of which have boundary points w_i, w_{i+1} in the various critical sets we have removed. Write $\ell_i(s)$ for the equation of the t coordinate of the line connecting w_i to w_{i+1}, i.e. the graph of the line is $(s, \ell_i(s))$. We also write the portion of $\varphi(e)$ between w_i and w_{i+1} as $\varphi_i(s)$.

Possibly after further removal of points, we assume that each arc $\varphi_i(s)$ has a tubular neighborhood T_i that contains $\ell_i(s)$ and each of these neighborhoods are pairwise disjoint. We are now going to define a homeomorphism that is the identity outside of T_i. To do so we need one more pair of functions.

$$\begin{aligned} \mu_\pm(x) &= x \quad \text{if} \quad |x| \geqslant 1 \\ \mu_\pm(x) &= 2x \pm 1 \quad \text{if} \quad -1 \leqslant x \leqslant \tfrac{-1}{2} \\ \mu_\pm(x) &= \tfrac{2}{3}x \pm \tfrac{1}{3} \quad \text{if} \quad \tfrac{-1}{2} \leqslant 1 \end{aligned}$$

Now we can define a homeomorphism g_i on T_i using $+$ if the function $\varphi_i(s)$ is concave and $-$ if the function $\varphi_i(s)$ is convex.

$$g_i(s,t) := (s, 2 \cdot |\varphi_i(s) - \ell_i(s)| \cdot \mu_\pm(\frac{t - \ell_i(s)}{2|\varphi_i(s) - \ell_i(s)|}) + \ell_i(s))$$

Since the domains of each T_i are disjoint we can define the homeomorphism g to be g_i when in T_i and the identity otherwise. This straightens each of the $\varphi_i(s)$. The portions of $\varphi(e)$ removed are already piecewise-linear. This makes each $g \circ \varphi(\bar{\sigma})$ into a piecewise-linear polyhedron, the interior of which is mapped via h to a single stratum of X. By adding edges and vertices, we can refine the stratification of $g \circ \varphi(\bar{\sigma})$ to be a piece-wise linear triangulation. $\square$

We will break our proof of MacPherson's characterization into two parts. The first shows that any representation defines a constructible cosheaf. The second shows that any constructible cosheaf defines a representation of the entrance path category.

Theorem 5.70 (Representations are Cosheaves)**.** *Let X be a Thom-Mather stratified space. Any representation of the entrance path category*

$$\mathbf{Entr}(X, \{X_\alpha\}) \to \mathbf{Vect}$$

defines a constructible cosheaf.

Proof. To produce a cosheaf from a representation $\hat{F} : \mathbf{Entr}(X) \to \mathbf{Vect}$ we take colimits over the restriction of $\hat{F}$ to the entrance path category of U (with its induced stratification):

$$\hat{F}(U) := \varinjlim_{\mathbf{Entr}(U)} \hat{F}|_U$$

This is clearly a pre-cosheaf since if $U \hookrightarrow V$, the colimit over $\hat{F}|_V$ defines by restriction a cocone over $\hat{F}|_U$ and thus a unique map $\hat{F}(U) \to \hat{F}(V)$. We can describe more explicitly this colimit as follows:

Given a point $x \in X_\sigma$ in a stratum of dimension i, there is a basis of conical neighborhoods $U_x \cong \mathbb{R}^i \times C(L)$ where L is the stratified fiber of the retraction map π_σ and $C(L)$ is its open cone. For such a neighborhood, x is the terminal object in $\mathbf{Entr}(U_x)$, thus the colimit returns the value of $\hat{F}(x)$. Moreover, this shows that the costalks of the pre-cosheaf defined stabilize for small contractible sets containing x.

To show this is actually a cosheaf we use the version of the van Kampen theorem adapted to the entrance path category just proved in theorem 5.65:

$$\varinjlim_{I \in N(\mathcal{U})} \mathbf{Entr}(U_I) \cong \mathbf{Entr}(X)$$

As a consequence of colimits commuting with colimits we get that for a representation of the entrance path category $\hat{F}$

$$\varinjlim_{I \in N(\mathcal{U})} \hat{F}(U_I) := \varinjlim_{I \in N(\mathcal{U})} \varinjlim_{\mathbf{Entr}(U_I)} \hat{F}|_{U_I} \cong \varinjlim_{\mathbf{Entr}(U_I)} \varinjlim_{I \in N(\mathcal{U})} \hat{F}|_{U_I} \cong \varinjlim_{\mathbf{Entr}(U)} \hat{F}|_U.$$

This establishes the cosheaf axiom. □

Theorem 5.71 (Representations of the Entrance Path Category). *Every cosheaf $\hat{F}$ with finite-dimensional costalks that is constructible with respect to a Thom-Mather stratification of X determines a representation of the entrance path category.*

$$\mathbf{Entr}(X, \{X_\alpha\}) \to \boldsymbol{vect}$$

Proof. If X is a Thom-Mather stratified space, then we know that every point $x \in X_\sigma$ in a stratum of dimension i has a neighborhood $U_x \cong \mathbb{R}^i \times C(L)$, where L is the fiber of π_σ, and $C(L)$ is the open cone. Now suppose $\hat{F}$ is a constructible cosheaf, which we assume has finite-dimensional costalks. We claim that for U_x suitably small we can show that

$$\hat{F}_x \cong \hat{F}(U_x).$$

This is not so easy to see and a proof would require substantial more development of cosheaf theory. Heuristically, if the value of $\hat{F}$ on a sequence of conical neighborhoods never stabilized then this would contradict the constancy of the cosheaf on sets of the form $U_x \cap X_\tau$. For a rigorous proof, one dualizes a constructible cosheaf to a constructible sheaf by post-composing with $\mathbf{Hom}_{\boldsymbol{vect}}(-, k)$, which is an equivalence, and we can apply the proof for constructible sheaves found on p. 84 of [37].

Consequently, if $y \in U_x \cap X_\tau$ is a point in a nearby stratum, then there is an analogous neighborhood U_y contained in U_x. Repeating the same argument, we can then use the maps present in a cosheaf to define a map from the costalk at y to the costalk at x:

$$\hat{F}_x \xrightarrow{\cong} \hat{F}(U_x) \longleftarrow \hat{F}(U_y) \xleftarrow{\cong} \hat{F}_y$$

Recall that the restriction of a constructible cosheaf to any stratum defines a locally constant cosheaf. For arbitrary points y' in the stratum X_τ we can consider a path $y' \rightsquigarrow y$ and use theorem 5.48 to define a map from $\hat{F}_{y'}$ to $\hat{F}_y$. Postcomposing with the above map defines the map $\hat{F}_{y'} \to \hat{F}_x$. This explains why constructible cosheaves map naturally define ways of specializing a costalk over one stratum to a costalk in its frontier.

To show homotopy invariance, we appeal to the van Kampen theorem 5.65 to reduce the argument to elementary homotopies of a particular form. Assume $\alpha(t)$ goes from $z \in X_\lambda$ directly to $x \in X_\sigma$ and that $\beta(t)$ goes from z to y and then x. Since restriction to any stratum defines a locally constant cosheaf, we can appeal to the homotopy invariance of theorem 5.48 to position these paths and points to be inside U_x and so that $U_x \supset U_y \supset U_z$. By choosing a similar set of isomorphisms, we get two commutative diagrams, which followed along the top edge corresponds to the action of $\alpha(t)$ and followed along the bottom edges corresponds to $\beta(t)$.

$$\begin{tikzcd} \hat{F}(U_z) \arrow[rr] \arrow[dr] & & \hat{F}(U_x) \\ & \hat{F}(U_y) \arrow[ur] & \end{tikzcd} \qquad \begin{tikzcd} \hat{F}_z \arrow[rr] \arrow[dr] & & \hat{F}_y \\ & \hat{F}_x \arrow[ur] & \end{tikzcd}$$

□

We want to show that stratified maps induce representations of the entrance path category, which, by the first part of our equivalence, defines a constructible cosheaf.

Theorem 5.72 (Cosheaves from Stratified Maps). *Fix an analytic-geometric category $\mathcal{G}$. If Y is a closed set in $\mathcal{G}(N)$ and $f : (Y, N) \to (X, M)$ is a C^1 proper definable map, then for each i, the assignment*

$$x \in X \rightsquigarrow H_i(f^{-1}(x); k)$$

defines a representation of the definable entrance path category of X, where the stratification is gotten by the stratification induced by f.

Proof. Let $\gamma : I \to X$ be a definable map that satisfies the entrance path condition, i.e. as t increases the dimension of the ambient stratum is non-increasing. Thus $\gamma(0)$ is in a stratum of dimension greater than or equal to $\gamma(1)$. By lemma 5.39, we know that the pullback $Y_\gamma := I \times_X Y$ is definable, as is the pullback of f along γ, written $\gamma^* f$. Since definable sets can be Whitney stratified, Y_γ admits a system of control data, and may be regarded as a Thom-Mather stratified space.

The argument from lemma 5.43 provides us with the prototype for getting a diagram of spaces for every path. We will repeat it here for convenience and make any necessary modifications. By definable triviality (4.11 of [90]), there exists a finite partition of $[0,1]$ such that over each interval the inverse image is homeomorphic to the product:

$$\begin{tikzcd} f^{-1}((t_i,t_{i+1})) \arrow[rr, "\cong"] \arrow[dr, "f"'] & & F\times(t_i,t_{i+1}) \arrow[dl] \\ & (t_i,t_{i+1}) & \end{tikzcd}$$

By properness we can, for any fixed $\epsilon > 0$, find an s_i^+ such that $f^{-1}([t_i, s_i^+])$ is contained in $U_i(\epsilon) := \cup T_\sigma(\epsilon/2)$ for $Y_\sigma \subseteq f^{-1}(t_i) := Y_i$. The retraction we constructed in proposition 5.20 gives a retraction map $r_i^+ := H(p, 0)$ from $U_i(\epsilon) \to Y_i$. This allows us to define a map on fibers

$$f^{-1}(s_i^+) \hookrightarrow U_i(\epsilon) \to Y_i.$$

Applying some homology functor $H_n(-; k)$ defines the representation locally on the path. Of course, we must show that this representation is independent of the point s_i taken. If $s_i^{+'} \in [t_i, s_i^+]$ is another point, then the composition of the trivialization with the retraction witnesses the homotopy between these two choices.

$$F \times [s_i^{+'}, s_i^+] \cong f^{-1}([s_i^{+'}, s_i^+]) \hookrightarrow U_i(\epsilon) \to Y_i$$

Similarly, one can find a point s_{i+1}^- so that it's fiber is contained in $U_{i+1}(\epsilon)$ and the retraction r_{i+1}^- defines a map from that fiber to the fiber $Y_{i+1} = f^{-1}(t_{i+1})$. By Thom's first isotopy lemma there is a homeomorphism $\varphi_{i+1,i}$ taking the fiber over s_i^+ to the fiber over s_{i+1}^-. This homeomorphism is gotten by constructing a vector field that flows from s_i^+ to s_{i+1}^- and lifting it to a controlled vector field on $f^{-1}((t_i, t_{i+1}))$ via proposition 9.1 of [58]. Finally, one must observe that the filtration of X by strata of a given dimension or less, the restriction of γ to the half-open interval $[t_i, t_{i+1})$ is contained inside a single stratum of X and thus the retraction r_i^+ induces a homotopy equivalence between the fiber over s_i^+ and the fiber over t_i. Applying our homology functor to the following composition defines the total action associated to this path:

$$\cdots (r_{i+1})_*^- \circ (\varphi_{i+1,i})_* \circ (r_i^+)_*^{-1} \cdots$$

It remains to be seen that this map is invariant under definable homotopies of entrance paths. Suppose $h : I \times I \to X$ is a definable homotopy. Again, the pullback $Y_h := I^2 \times_X Y$ is definable, as is the map h^*f, and both can be stratified. Thus, we have reduced everything to considering a stratified map to the square I^2. By the van Kampen theorem 5.65, it suffices to check homotopy invariance on an elementary homotopy, such as the one depicted in figure 16. Let us assume that h is a homotopy between an entrance path $\alpha(t) = h(0, t)$, which goes from a stratum X_λ and enters a stratum X_σ at the last possible moment $t = 1$, and an entrance path $\beta(t) = h(1, t)$, which enters X_τ at $t = 1/2$ and then goes to X_σ at $t = 1$. Moreover, we assume that h takes the complement of $\{t = 1\} \cup \{(1, t) \mid t \geqslant 1/2\}$ to the stratum

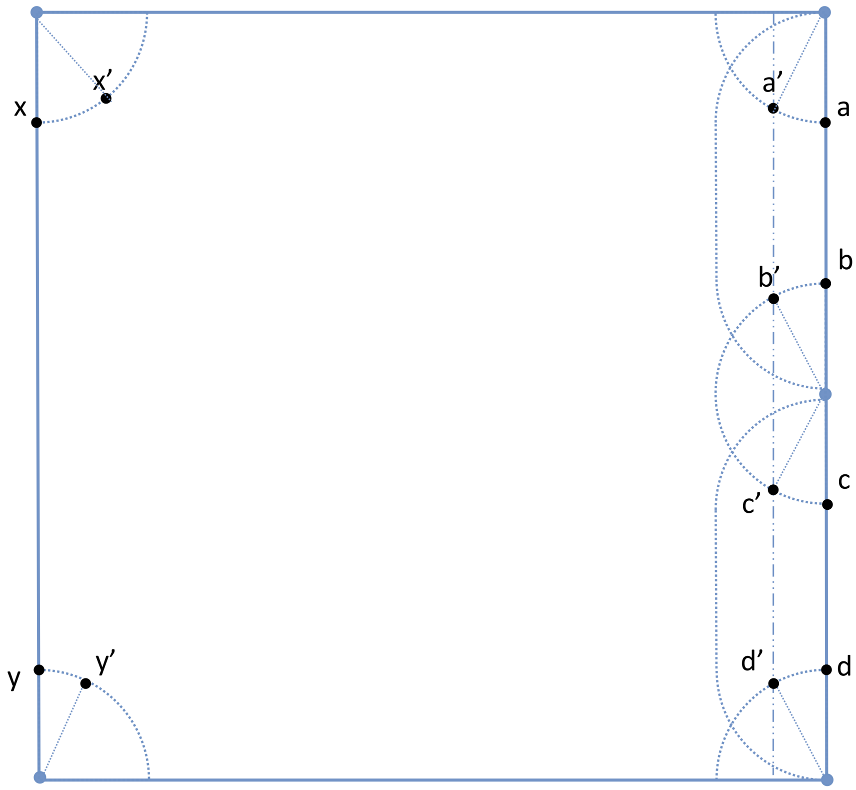

Figure 19: Argument for Homotopy Invariance

X_λ. This guarantees that the fibers over $x, x', y, y', a, a', b, b', c'$ and d' in figure 19 can all be identified.

Let T be a system of control data for Y_h, obtained in a specific way. By restricting to the strata over $s = 0$ and $s = 1$ respectively, we get control data for $Y_{s=0}$ and $Y_{s=1}$, both of which are inside $I^2 \times_X Y \subset \mathbb{R}^2 \times N$. The spaces $Y_{s=0}$ and $Y_{s=1}$ can be identified with the inclusions of Y_α or Y_β, which are contained in $\mathbb{R} \times N$. The manner in which Mather constructs control data in proposition 7.1 of [58] can be used to extend the control data for Y_α and Y_β to control data for $Y_{s=0}$ and $Y_{s=1}$ inside $\mathbb{R}^2 \times N$ respectively. This is how we obtain those tubular neighborhoods in $\{T\}$ and the rest can be constructed to be compatible with those. This allows us to use the control data T to meaningfully compare the construction above for $\alpha(t)$ and $\beta(t)$.

We can describe the maps associated to $\alpha(t)$ and $\beta(t)$ as follows: By properness, we assume the fiber over x is contained in a regular neighborhood, which retracts via r_x to the fiber over $(0, 0)$. There is a homeomorphism $\varphi_{y,x}$ from the fiber over x to the fiber over y. Finally, we can assume that the fiber over y retracts via r_y to the fiber over $(0, 1)$. Thus the action associated to $\alpha(t)$ is the map

$$(r_y)_* \circ (\varphi_{y,x})_* \circ (r_x)_*^{-1} : H_n(Y_{(0,0)}) \to H_n(Y_{(0,1)})$$

where we have implicitly pre-composed r_x with the inclusion of the fiber.

For $\beta(t)$, the action is similar:

$$(r_d)_* \circ (\varphi_{d,c})_* \circ (r_c)_*^{-1} \circ (r_b)_* \circ (\varphi_{b,a})_* \circ (r_a)_*^{-1} : H_n(Y_{(1,0)}) \to H_n(Y_{(1,1/2)}) \to H_n(Y_{(1,1)})$$

The strategy of the proof is to pick a path $\gamma(t)$ that interpolates $\alpha(t)$ and $\beta(t)$ and show that the associated map on homology agrees with both $\alpha(t)$ and $\beta(t)$. This path is indicated by the dotted-and-dashed line passing through a', b', c' and d' in figure 19. The representation associated to $\gamma(t)$ is

$$(r_1)_* \circ (\varphi_{d',c'})_* \circ (i_{c'})_*^{-1} \circ (i_{b'})_* \circ (\varphi_{b',a'})_* \circ (r_0)_*^{-1}.$$

Here the maps $i_{b'}$ and $i_{c'}$ denote the inclusion of $Y_{b'}$ and $Y_{c'}$ into the inverse image of the interval $[b', c']$. The action on homology of $(i_{c'})_*^{-1} \circ (i_{b'})_*$ agrees with an analogously constructed homeomorphism $\varphi_{c',b'}$, but we will find it easier to equate the map associated to $\beta(t)$ and $\gamma(t)$ as written above.

Because the control data $\{T\}$ extends the control data for Y_α and Y_β, the retraction map r_a can be taken to be the restriction of a retraction map $r_0 : U_{Y_{(1,0)}}(\epsilon) \to Y_{(1,0)}$ constructed in proposition 5.20. This in turn can be taken to be the restriction of the tubular projections used to define a retraction map $r_{s=1} : U_{s=1}(\epsilon) \to Y_{s=1}$. The commutation relations for control data allow us to imagine first taking the fiber $Y_{a'}$ over a' and retracting to the strata over the edge $e_0 := \{(1,t) \,|\, 0 < t < 1/2\}$, and then retracting to the fiber over $(1,0)$. This allows us to factor r_0 as

$$r_0 = r_a \circ r_{e_0},$$

but the image of $Y_{a'}$ under r_{e_0} may not be contained in Y_a or any single fiber. This would be true if, for example, the control data defining the retraction to Y_{e_0} satisfied

$$\pi_{e_0}(f(p)) = f(\pi_\sigma(p))$$

for each stratum Y_σ that f carried to e_0, but in general it does not. This is what necessitates the use of the Thom properties given by lemma 5.34 and property (b) of proposition 5.31.

By lemma 5.34, we know that restricting the codomain to the complement of the vertices, the mapping h^*f is a Thom mapping. Consequently, if we pick a tubular neighborhood T_{e_0} for the edge $e_0 := \{(1,t) \,|\, 0 < t < 1/2\}$, there exists a system of control data $\{T'\}$ over T_{e_0} and the interior of I^2 by proposition 5.31. If we restrict to those tubular neighborhoods coming from strata in Y_{e_0}, then property (a) of proposition 5.31 implies that this restricted collection of tubular neighborhoods defines actual control data for Y_{e_0}, which we call $\{T'\}_{Y_{e_0}}$. A priori, the analogous restriction of $\{T\}$ to Y_{e_0} defines a different system of control data. However, by Mather's uniqueness result,[28] there is a homeomorphism ψ_{e_0} of Y_{e_0} that takes $\{T'\}_{Y_{e_0}}$ to

[28] Mather mentions at the bottom of page 492 of [58] that any Whitney stratified subset Z of a manifold M has a unique, up to isomorphism, structure as a Thom-Mather stratified set. This is not explicitly proved, but it follows from Mather's cor. 10.3 as explained by Goresky: Suppose Z is given two different structures of control data $\{T\}$ and $\{T'\}$. If we consider $Z \times \mathbb{R}$ as a Whitney stratified subset of $M \times \mathbb{R}$, then $\{T\}$ and $\{T'\}$ can be extended to control data on $Z \times (-\epsilon, \epsilon)$ and $Z \times (1-\epsilon, 1+\epsilon)$, respectively. Then, using the proof of prop. 7.1, one can find control data on all of $Z \times \mathbb{R}$ that agrees with the ϵ extensions of $\{T\}$ and $\{T'\}$. This space, now viewed as a Thom-Mather stratified set, is then isomorphic via cor. 10.3 to the set where just $Z \times \mathbb{R}$ is given the extension of just the control data $\{T\}$.

$\{T\}_{Y_{e_0}}$. This implies that if Y_σ is a stratum that is mapped to $(1,0)$ and Y_τ is mapped to e_0, then

$$\pi_\sigma = \pi_\sigma \circ \psi_{e_0} \circ \pi_{\tau'} \qquad \text{since} \qquad \pi_\tau = \psi_{e_0} \circ \pi_{\sigma'}$$

By repeating the construction of a retraction outlined in proposition 5.20, but using the control data $\{T'\}_{Y_{e_0}}$ instead to construct the family of lines, we get a map r'_{e_0} that carries the fiber over a' to the fiber over a. Post-composing r'_{e_0} with ψ_{e_0} gives the equality $r_{e_0} = \psi_{e_0} \circ r'_{e_0}$. This construction gives the left most triangle in the following commutative diagram:

$$\begin{tikzcd}
Y_{(1,0)} & Y_a \arrow[l, "r_a\psi_{e_0}"'] \arrow[r, "\varphi_{b,a}"] & Y_b \arrow[r, "r_b\psi_{e_0}"] & Y_{(1,1/2)} & Y_c \arrow[l, "r_c\psi_{e_1}"'] \arrow[r, "\varphi_{d,c}"] & Y_d \arrow[r, "r_d\psi_{e_1}"] & Y_{(1,1)} \\
& Y_{a'} \arrow[ul, "r_0"] \arrow[u, "r'_{e_0}"'] \arrow[r, "\varphi_{b',a'}"'] & Y_{b'} \arrow[u, "r'_{e_0}"] \arrow[ur, "r_{1/2}"'] \arrow[r, "i_{b'}"'] & Y_{[b',c']} \arrow[u, "r_{1/2}"'] & Y_{c'} \arrow[ul, "r_{1/2}"'] \arrow[u, "r'_{e_1}"'] \arrow[l, "i_{c'}"] \arrow[r, "\varphi_{d',c'}"'] & Y_{d'} \arrow[u, "r'_{e_1}"] \arrow[ur, "r_1"'] &
\end{tikzcd}$$

Now we explain the other maps in this diagram. The homeomorphisms $\varphi_{b,a}$ and $\varphi_{b',a'}$ are constructed by taking a controlled vector field $\{\eta_f, \eta_{e_0}, \eta_{e_1}\}$ in I^2 minus the vertices using the control data $\{T'\}$ over $\{T_f, T_{e_0}, T_{e_1}\}$. Since $d\pi_{e_0}(\eta_f(s,t)) = \eta_{e_0}(\pi_{e_0}(s,t))$ the controlled vector field over this one commutes with f and gives

$$\varphi_{b,a} \circ r_{e'_0} = r_{e'_0} \circ \varphi_{b',a'}.$$

Again, the commutation relations in proposition 5.20 allows us to, using the control data $\{T\}$ to factor $r_{1/2} = r_b \circ r_{e_0}$. However, the uniqueness theorem tells us that $r_{e_0} = \psi_{e_o} \circ r'_{e_0}$ where ψ_{e_0}. A simple diagram chase now completes the argument. Comparing the maps associated to $\alpha(t)$ and $\gamma(t)$ is much simpler and uses the same ideas. We leave it to the reader. $\square$

Remark 5.73 (Alternative Idea for a Proof). An alternative approach makes use of the properties of o-minimal structures. The generic triviality theorem 4.11 of [90] guarantees that we have a definable trivialization of the map over $(0,1)$.

$$\begin{tikzcd}
f^{-1}(\gamma((0,1))) \arrow[rr, "\cong", "h"'] \arrow[dr, "\gamma^* f"'] & & F \times (0,1) \arrow[dl] \\
& (0,1) &
\end{tikzcd}$$

For each point x in the fiber F, we get a lift $\{x\} \times (0,1)$ of the open interval. Applying the inverse homeomorphism, $h^{-1}(\{x\} \times (0,1))$ defines a definable path $\alpha_x : (0,1) \to f^{-1}(\gamma([0,1]))$.

Mário Edmundo and Luca Prelli, in their recent note [27] reworking the six basic Grothendieck operations for sheaves in the o-minimal setting, have given a tantalizing reformulation of what characterizes a definable proper map. They use an idea of Ya'acov Peterzil and Charles Steinhorn [66] that shows that being definably compact (equivalently, closed and bounded) is equivalent to being able to to complete curves. A map $f : Y \to X$ is definably proper if for every definable curve $\alpha : (0,1) \to Y$ and every definable map $[0,1] \to X$ there is at least one way to complete the diagram:

$$\begin{tikzcd} (0,1) \arrow[r, "\alpha"] \arrow[d] & Y \arrow[d, "f"] \\ {[0,1]} \arrow[r] \arrow[ur, "\bar{\alpha}"] & X \end{tikzcd}$$

If one assumes all the maps are continuous as well as definable then the completion in the diagram above is unique.[29]

In our situation, the hypotheses guarantee that for each point $x \in F$, we can complete the curve $\alpha_x : (0,1) \to I \times_X Y$ to a curve $\bar{\alpha}_x : [0,1]$. By associating endpoints over 0 to endpoints over 1 we define a set-theoretic map $g : f^{-1}(\gamma(0)) \to f^{-1}(\gamma(1))$. The hard work to show that this map g is continuous and is invariant under homotopy.

The utility of this theorem is that it gives a much more geometric and combinatorial description of constructible cosheaves. As an example, it allows one to study Morse theory using the induced stratification of $\mathbb{R}$ given by critical points and intervals determined by a Morse function $f : M \to \mathbb{R}$.

Example 5.74 (Morse Torus)**.** The example that many people use to first learn Morse theory is the torus, stood up on end, mapped to the real line using a height function. Since we want to envision the height function as giving a stratified mapping down to the real line. We rotate the picture 90° to obtain the following family of spaces. Now the theorem tells

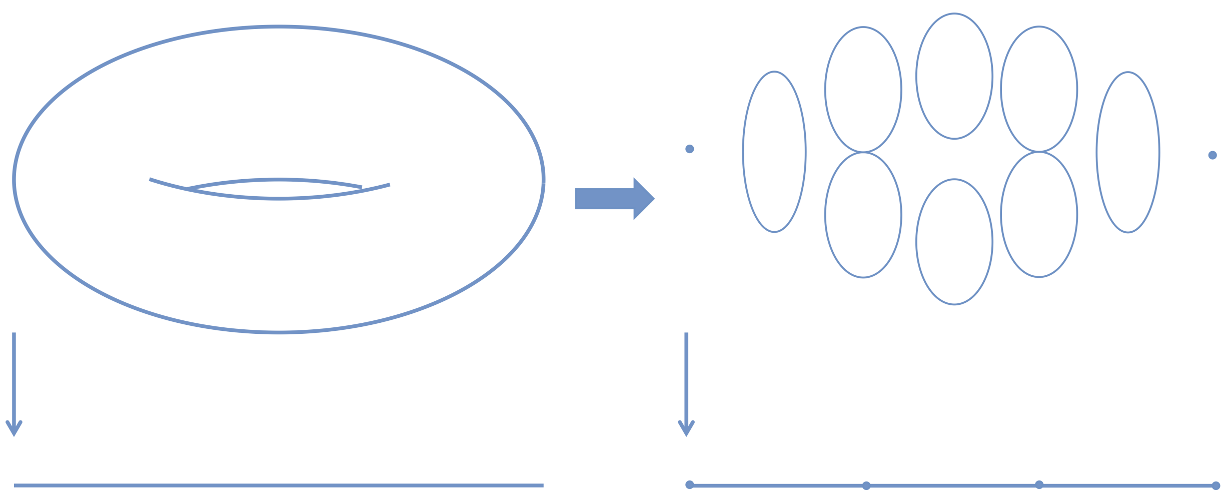

Figure 20: Morse Functions are Stratified Maps

us that we can pick a non-zero integer i and take ith homology of the fiber with coefficients in k and this will define for us a cellular cosheaf of vector spaces. We sketch the diagram of vector spaces so obtained: Once we've obtained this diagram of vector spaces we can begin

[29]One of the unusual features of o-minimal topology is that definable maps need not always be continuous, thus the added hypothesis. Even discontinuous maps can have triangulable graphs.

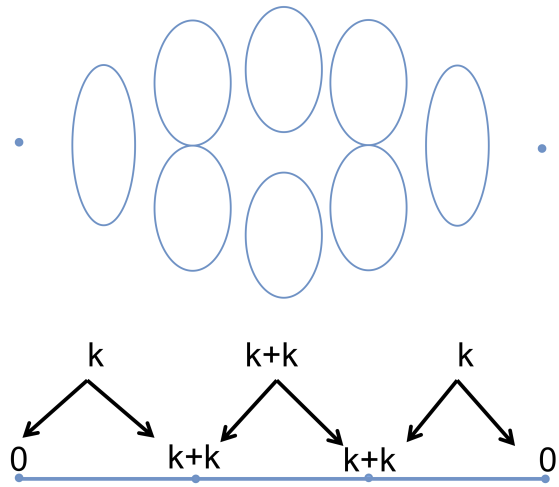

Figure 21: Cellular Cosheaf Obtained by Taking H_1 of the Fiber

to probe what it tells us about the original space. This question will be taken up in due time.

One of the major pay-offs of using cellular stratifications is that we get a strong compression of the data that specifies a sheaf or cosheaf. With general stratifications having non-simply connected strata the process of recording a map for each non-trivial homotopy class can require recording a countably infinity of maps (although finite generation simplifies things). In cellular stratifications we only need to record a single map for each incidence relation.

Example 5.75 (Klein Bottle Revisited)**.** As already seen, a Klein bottle can be viewed as a non-trivial S^1 bundle over the circle. We've already seen how the this leads to a representation of the fundamental group $\pi_1(S^1; x_0) \cong \mathbb{Z}$. We can concoct a cellular cosheaf that describes this bundle in a different way. Let s and s' denote two vertices in a cell decomposition which includes two edges a and b. We can imagine calling the edge a the short edge between s and s' and let b be the long edge. To each cell $\hat{F}$ assigns the homology of the fiber to that cell. The actions are encoded using maps between the cells. This gives us a diagram of vector spaces in the shape of the entrance path category for this different stratification of S^1:

$$\begin{tikzcd}
& \hat{F}(b) \arrow[dl, "-1"'] \arrow[dr, "1"] & \\
\hat{F}(s) & & \hat{F}(s') \\
& \hat{F}(b) \arrow[ul, "1"] \arrow[ur, "1"'] &
\end{tikzcd}$$

Observe that instead of storing a matrix for each integer, corresponding to each homotopy class of map, we only need to record a map for each pair of incident cells.

5.2 Partially Ordered Sets: Finite Spaces and Functors

The preceding section explained one way of thinking of cellular sheaves and cosheaves as more traditional gadgets that involve the assignment of data to open sets. In summary, by passing through an equivalence between (co)sheaves constructible with respect to a cellular stratification and representations of a particular category we obtained cellular (co)sheaves. This equivalence is useful for its geometric and topological content, but it obfuscates the algebraic properties and symmetries constructible cosheaves possess.

In this section we take a second vantage point. Instead of working on a space X that is decomposed into cells $\{X_\sigma\}_{\sigma \in P_X}$, we work over the indexing poset P_X. Partially ordered sets can be endowed with a topology making cellular sheaves and cosheaves into actual sheaves and cosheaves on this topology.

Here one can illuminate all of the general machinery of classical sheaf theory, but with a combinatorial finiteness that bends the theory to direct computation and understanding. Some of the explicit treatment of sheaves on posets is contained in the clear and concise work of Sefi Ladkani [47], but we streamline the discussion by using Kan extensions, which clarifies how cosheaves on a poset X differ from sheaves on X^{op}.

5.2.1 The Alexandrov Topology

In this section we consider how partially ordered sets (posets) arise naturally for spaces of "finite type" [57] and how these posets can be treated as spaces in their own right.

Definition 5.76. A **pre-order** consists of a set P and a relation $\leqslant$ that is reflexive and transitive. A **poset** is a pre-order where the relation is also anti-symmetric, i.e. $x \leqslant y$ and $y \leqslant x$ implies $x = y$. A map f of pre-orders is one that respects $\leqslant$. That is if $x \leqslant y$ then $f(x) \leqslant f(y)$. Pre-orders and order preserving maps form a category $\mathbf{Preorder}$. The collection of all posets form a subcategory of this category.

Any pre-order P has an associated poset. This poset is gotten by defining an equivalence relation on P via $x \sim y$ if and only if $x \leqslant y$ and $y \leqslant x$. One can check that this surjection is order-preserving. This construction defines a right adjoint to the inclusion of posets into pre-orders [98].

Example 5.77. The exit and entrance path categories are strictly more sophisticated structures than preorders. Nevertheless, one can associate a pre-order to a stratified space by declaring $x \leqslant y$ if and only if there is an exit path from x to y. For a general stratified space this preorder loses the information of how (the homotopy class of path, for example) the two points are related. When X is stratified as a cell complex, the exit path category is a preorder, with no loss of information. When we passed to the skeletal subcategory of the exit (or entrance) path category we foreshadowed the process of taking the associated poset to a preorder.

Every pre-order can be equipped with a topology. However, it was first defined for finite posets by Pavel Alexandrov [3, 5] and the general definition carries his name.

Definition 5.78. On a pre-order $(\mathsf{P}, \leqslant)$ define the **Alexandrov topology** to be the topology whose open sets are the sets that satisfy the following property:

$$x \in U \qquad x \leqslant y \qquad \Rightarrow \qquad y \in U$$

A basis is given by the sets of the form $U_x := \{y \in \mathsf{P} | x \leqslant y\}$ – what we will call the **open star at x**. Similarly, we define the **closure of x** by $\bar{x} := \{y \in \mathsf{P} | y \leqslant x\}$. When P is a finite poset, then a basis of closed sets is given by the $\bar{x}$'s.

Remark 5.79 (P will mean a poset). Although spaces equipped with a pre-order are an interesting class of structures to consider, we will now work exclusively with posets. We do this to prevent closed loops from occurring in chains of related elements, as this would complicate our story.

Example 5.80. Consider $(\mathbb{R}, \leqslant)$ with the usual partial order. The open sets are all those open or half open intervals such that the righthand endpoint is $+\infty$. Observe that the closed set $(-\infty, 0)$ cannot be written as an intersection of closed sets of the form $\bar{t}$. Thus the closures at t do not form a basis.

The dictionary between cellular complexes and Alexandrov spaces is easily described. First we introduce another definition.

Definition 5.81 (Star)**.** Let $(X, \{X_\sigma\})_{\sigma \in \mathsf{P}_X}$ be a cell complex. Every cell X_σ has a **star**, which is a set that consists of all those cells X_τ such that $X_\sigma \leqslant X_\tau$.

$$\mathrm{star}(X_\sigma) := \{X_\tau \,|\, X_\sigma \leqslant X_\tau\}$$

Since this definition only depends on the incidence relation of cells, we often drop the distinction between X_σ and its label σ. Thus the star is also described as a subset of the poset P_X consisting of those labels τ such that $\sigma \leqslant \tau$.

The Alexandrov topology on the indexing poset P_X of a cell complex allows us to define a continuous surjective map that comes from sending each cell X_σ to its label σ. This continuous surjective map gives an alternative way of describing how the Alexandrov topology arises. It is the quotient space where we identify two points x and y if and only if they belong to the same cell.

$$\begin{tikzcd} X \arrow[d, "q"] \\ \mathsf{P}_X := X/\sim \end{tikzcd}$$

The inverse image of the star of σ is an open union of cells, which is open. Thus this map is continuous and the topology that makes this map continuous is the Alexandrov topology.

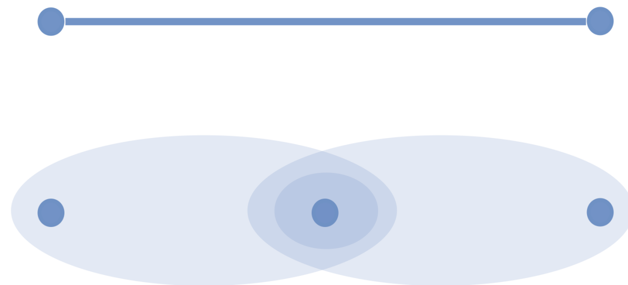

Figure 22: Alexandrov Space Associated to the Unit Interval

Example 5.82 (The Interval)**.** Suppose $X = [0, 1]$ is the unit interval given a cell complex structure with two vertices and one open interval. The face relation poset P_X takes the following form:

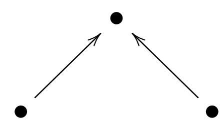

The Alexandrov topology has basic open sets corresponding to the star of each cell. The stars of the two vertices intersect each other. In figure 22, we have drawn the basic open sets.

The reader should note that although this topology is non-Hausdorff, it is highly relevant to concepts in algebraic topology.

Remark 5.83 (Weak Homotopy Equivalence)*.* There is a remarkable theorem due to Michael McCord [62] that states that every finite simplicial complex is weakly homotopy equivalence to an Alexandrov space. Thus, if one is interested in the topological properties of simplicial complexes, one should care about (non-Hausdorff) Alexandrov spaces. McCord gives constructions of classical operations in algebraic topology, including suspension, in the Alexandrov setting.

5.2.2 Functors on Posets

We now want to understand how data modeled on posets can be treated as a sheaf or cosheaf on the Alexandrov topology. To do so we take adopt an the elegant approach of Kan extensions. To motivate this concept we will consider the relationship between a poset and its topology.

Observe that the correspondence between the relation internal to the poset P and the containment relation for the open sets in the Alexandrov topology is order-reversing. Said more succinctly, we have an inclusion functor that is contravariant, i.e.

$$\iota : P \to \mathbf{Open}(P)^{op} \qquad p \mapsto U_p.$$

A natural question to ask is

"Given a functor $F : P \to \mathcal{D}$, is there a consistent way of extending F to a functor $R : \mathbf{Open}(P)^{op} \to \mathcal{D}$?"

One can hope that since the image of the inclusion $\iota : P \to \mathbf{Open}(P)^{op}$ is a basis for the topology. Consequently we can express arbitrary open sets as unions (colimits or limits in the opposite category) of basic open sets $\iota(p) = U_p$.

A candidate extension would be to define

$$F(U) := \varprojlim_{U_p \subset U} F(p)$$

or as the colimit of F over $U_p \subset U$. However, we should have some consistency. If one views $U_p = \{p' | p \leqslant p'\}$ as a subcategory of the category P, then it has an initial object p and thus the limit of the diagram $F|_{U_p}$ is $F(p)$, i.e.

$$\varprojlim_{p \leqslant p'} F(p') \cong F(p).$$

This guides us to the following possible extension.

$$\begin{tikzcd} P \arrow[r, "F"] \arrow[d, "\iota"'] & \mathcal{D} \\ \mathbf{Open}(P)^{op} \arrow[ur, "\lim_{U_p \subset U} F(p) =: F(U)"'] & \end{tikzcd}$$

This extension is nice for many reasons. By using limits to define data on larger open sets we have forced the sheaf axiom to hold, so this extension is in fact a sheaf. Moreover it illustrates through example a more general concept, which we now define.

Remark 5.84 (Caveat). We will make use of Kan extensions at a few points throughout the paper, but its immediate application is a theorem that says functors out of posets can be identified with sheaves. The proof of this theorem is described casually without the language of Kan extensions in [47], but adopting this language will be powerful and will make certain categorical properties transparent.

Definition 5.85 (Kan Extensions). Suppose $\mathcal{B}, \mathcal{C}$ and $\mathcal{D}$ are categories, $F : \mathcal{B} \to \mathcal{D}$ and $E : \mathcal{B} \to \mathcal{C}$ are functors, then the **right Kan extension of** F **along** E written $R = \mathrm{Ran}_E F : \mathcal{C} \to \mathcal{D}$ is a functor and a natural transformation $\epsilon : RE \to F$ that is universal in the following sense. For every functor $H : \mathcal{C} \to \mathcal{D}$ with a natural transformation $\alpha : H \circ E \to F$ there exists a unique natural transformation $\sigma : H \to R$, i.e. $\mathbf{Nat}(H, R) \cong \mathbf{Nat}(H \circ E, F)$.

$$\begin{tikzcd} \mathcal{B} \arrow[r, "F"] \arrow[d, "E"'] & \mathcal{D} \\ \mathcal{C} \arrow[ur, "R = \mathrm{Ran}_E F"'] & \end{tikzcd}$$

The **left Kan extension of F along E** written $L = \mathrm{Lan}_E F : \mathcal{C} \to \mathcal{D}$ is a functor with a natural transformation $\eta : F \to L \circ E$ that is universal as well. If $H : \mathcal{C} \to \mathcal{D}$ is a functor with a natural transformation $\omega : F \to H \circ E$, then there exists a unique $\tau : L \to H$, i.e. $\mathbf{Nat}(L, H) \cong \mathbf{Nat}(F, H \circ E)$.

$$\begin{tikzcd} \mathcal{B} \arrow[r, "F"] \arrow[d, "E"'] & \mathcal{D} \\ \mathcal{C} \arrow[ur, "L=\mathrm{Lan}_E F"'] & \end{tikzcd}$$

Remark 5.86 (Existence of Kan Extensions). Kan extensions do not always exist, but we have already alluded to a situation where they do. If $\mathcal{D}$ has enough limits and colimits, then we can give pointwise formulae for the left and right Kan extensions respectively:

$$\mathrm{Lan}_E F(c) := \varinjlim_{E(b)\to c} F(b) \qquad \mathrm{Ran}_E F(c) := \varprojlim_{c \to E(b)} F(b)$$

One of the reasons that sheaves and cosheaves on Alexandrov spaces are so well-behaved is that every open set has a finest cover, so in particular, by Corollary 3.15, we only need to check the (co)sheaf axiom on this cover, and it will be guaranteed for all others. Furthermore, every point in an Alexandrov space has a smallest open neighborhood, and the (co)stalks are just the values on these minimal open sets. This is how we can use Kan extensions to create a dictionary between (co)sheaves on Alexandrov spaces and functors out of posets.

Theorem 5.87. *Let P be a poset and $\mathcal{D}$ a category that is both complete and co-complete. Then the following categories are equivalent*

$$\mathbf{Fun}(P, \mathcal{D}) \cong \mathbf{Shv}(P; \mathcal{D}) \qquad \mathbf{Fun}(P^{op}, \mathcal{D}) \cong \mathbf{CoShv}(P; \mathcal{D})$$

Proof. We claim that taking the right Kan extension of $F : P \to \mathcal{D}$ along the inclusion $\iota : P \to \mathbf{Open}(P)^{op}$ produces a sheaf. Suppose U is an open set in the Alexandrov topology, i.e. one for which $p \in U$ and $p \leqslant p' \Rightarrow p' \in U$. It is true that every open set can be expressed as a union $U = \cup_{p \in U} U_p$ and thus the finest possible cover is $\{U_p\}_{p \in U}$. The right Kan extension then defines $F(U) := F[\{U_p\}_{p\in U}]$ so the sheaf axiom holds for that cover, but by Corollary 3.15, this means that F is a sheaf. To go from a sheaf to the diagram, one simply takes stalks at every point. Since the smallest neighborhood containing p is U_p, we get that $F_p = F(U_p) = F(p)$.

The dual argument for cosheaves is completely analogous: we take the left Kan extension of $\hat{F} : P^{op} \to \mathcal{D}$ along the inclusion $\iota : P^{op} \to \mathbf{Open}(P)$ to get a cosheaf. Taking costalks returns a diagram from a cosheaf. □

Remark 5.88 (Stalks and Costalks on Posets). To elaborate on the proof, let us compute some invariants. Recall that the stalk and costalk at a point $p \in P$ for a sheaf and cosheaf respectively is described via the use of filtered colimits and limits.

$$F_p := \varinjlim_{U \ni p} F(U) \qquad \text{and} \qquad \hat{F}_p := \varprojlim_{U \ni p} \hat{F}(U)$$

In both cases when P is a poset with the Alexandrov topology there is a smallest open set containing p, namely $U_p = \{q | p \leqslant q\}$, so $F_p \cong F(U_p) = F(p)$ and $\hat{F}_p = \hat{F}(U_p) = \hat{F}(p)$.

Definition 5.89 (Sections)**.** Let $(\mathsf{P}, \leqslant)$ be a poset and $\mathsf{F} : \mathsf{P} \to \mathcal{D}$ a sheaf and $\hat{\mathrm{F}} : \mathsf{P}^{\mathsf{op}} \to \mathcal{D}$. let $\mathsf{Z} \subset \mathsf{P}$ be any subset. We define the **sections over** Z to be

$$\Gamma(\mathsf{Z}; \mathsf{F}) := \varprojlim \mathsf{F}|_{\mathsf{Z}} \qquad \text{and} \qquad \varinjlim \hat{\mathrm{F}}\,|_{\mathsf{Z}} =: \Gamma(\mathsf{Z}; \hat{\mathrm{F}}).$$

When $\mathsf{Z} = \mathsf{P}$, we call these **global sections**. Note that $\Gamma(\mathsf{Z}; -)$ is context dependent: different definitions are used pending whether a sheaf or cosheaf is used.

The above theorem provides the simplest explanation of why cellular sheaves and cosheaves deserve to be called sheaves and cosheaves. When theorem 5.87 is specialized to the face relation poset P_{X} of a cell complex, also called the cell category $\mathsf{P}_{\mathsf{X}} = \mathbf{Cell}(\mathsf{X})$ in definition 5.3, we get that the category of sheaves in definition 5.7. We summarize these observations in the following corollary.

Corollary 5.90. *Let* $(\mathsf{X}, \mathsf{P}_{\mathsf{X}})$ *be a cell complex. A* ***cellular sheaf on*** X *is a sheaf on* P_{X} *equipped with the Alexandrov topology. Such a sheaf is uniquely determined by a functor* $\mathsf{F} : \mathsf{P}_{\mathsf{X}} \to \mathcal{D}$*. A* ***cellular cosheaf on*** X *is a cosheaf on* P_{X} *with the Alexandrov topology. Such a cosheaf is uniquely determined by a functor* $\hat{\mathrm{F}} : \mathsf{P}_{\mathsf{X}}^{\mathsf{op}} \to \mathcal{D}$*.*

To close, we point out one of the symmetries that Alexandrov spaces possess.

Claim 5.91. *In the Alexandrov topology, arbitrary intersections of open sets are open and arbitrary unions of closed sets are closed. Thus, every Alexandrov space possesses a dual topology by exchanging open sets with closed sets.*

This observation would have pleased Leray. It demonstrates that one can also think of a functor $\mathsf{F} : \mathsf{P} \to \mathcal{D}$ on a poset as either a sheaf or as a "cosheaf on closed sets." What distinguishes these two though is whether we use limits or colimits to extend to larger sets. We will consider this perspective in greater detail in section .

6 Functoriality Under Maps

Since sheaves and cosheaves as defined here assign data to open sets, maps between spaces should only make reference to open sets. In the case where our spaces are partially ordered sets endowed with the Alexandrov topology, it suffices to work directly with points since they are in bijection with a basis for the topology. However, playing these perspectives off of each other adds depth to the theory. In particular, by restricting our attention to these spaces, and using Kan extensions, we define the basic functors on (co)sheaves without making use of (co)sheafification. Pedagogically this is advantageous because the operation of sheafification tends to obfuscate the underlying ideas of sheaf theory. The lack of an explicit cosheafification[30] process has historically, been a stumbling block for the theory.

Recall that the definition of a continuous map $f : X \to Y$ says that the inverse image of an open set of Y is an open set of X. This observation can be expressed by saying that we have a functor

$$\mathring{f} : \mathbf{Open}(Y) \to \mathbf{Open}(X) \qquad U \subseteq Y \rightsquigarrow f^{-1}(U) \subseteq X.$$

By formality, we also have a functor from the corresponding opposite categories

$$\mathring{f} : \mathbf{Open}(Y)^{op} \to \mathbf{Open}(X)^{op}.$$

We have purposely suppressed the op superscript on $\mathring{f}$ for legibility. Now suppose we are given a pre-sheaf G on X, then we get a naturally associated pre-sheaf on Y by observing that the diagram

$$\begin{tikzcd}
\mathbf{Open}(X)^{op} \arrow[r, "G"] & \mathcal{D} \\
\mathbf{Open}(Y)^{op} \arrow[u, "\mathring{f}"] \arrow[ur, dashed] &
\end{tikzcd}$$

has a natural completion given by pre-composition.

Definition 6.1 (Pushforward Sheaf and Cosheaf)**.** Suppose $f : X \to Y$ is a continuous map of spaces, G and $\hat{G}$ a pre-sheaf and pre-cosheaf respectively, then define the **pushforward** or **direct image** pre-sheaf and pre-cosheaf via $f_*G(U) := G(f^{-1}(U))$ and $f_* \hat{G}(U) := \hat{G}(f^{-1}(U))$. Because f^{-1} commutes with unions, we get that if G or $\hat{G}$ is a sheaf or cosheaf, then so is the pushforward. Moreover, this operation is functorial with respect to maps between (co)sheaves, so we get functors

$$f_* : \mathbf{Shv}(X; \mathcal{D}) \to \mathbf{Shv}(Y; \mathcal{D}) \qquad f_* : \mathbf{CoShv}(X; \mathcal{D}) \to \mathbf{CoShv}(Y; \mathcal{D}).$$

There is also a pull-back functor associated to a continuous map $f : X \to Y$, but it's construction is less obvious. Namely, if G is a pre-sheaf on Y, then there is no clear way to define a pre-sheaf on X because for an open set on X, $f(U)$ may not be open. The solution

[30] It can be checked abstractly that cosheafification does exist for most data categories.

usually used is to take a system of approximations of $f(U)$ by open sets and to define the pullback sheaf as the limit of these approximations.

$$f^*G(U) := \varinjlim_{V \supset f(U)} F(V).$$

Thinking categorically, the problem of "approximation" has been encountered before. Namely, how can we complete the following diagram?

$$\begin{tikzcd} \mathbf{Open}(Y)^{op} \arrow[r, "G"] \arrow[d, "\mathring{f}"'] & \mathcal{D} \\ \mathbf{Open}(X)^{op} \arrow[ur, dashed, "?"'] & \end{tikzcd}$$

Again, by assuming that $\mathcal{D}$ has sufficient colimits, we can could fill in the diagram by taking the left Kan extension of G along $\mathring{f}$ and that will yield the candidate formula for the pullback just presented.

Unfortunately, this definition for the pull-back of a sheaf does not always define a sheaf. In classical sheaf theory this defect is circumvented by a process called **sheafification**, which can take a pre-sheaf and produce a sheaf. In standard books, one would define the pullback sheaf as the sheafification of the pre-sheaf defined here. However, this process is a bit mysterious. After programming your pre-sheaf to take certain values, it is hard to say what the values of the associated sheaf are, but the working mechanics of sheaf theory allow one to talk about pre-sheaves knowing that all these statements can be reflected into sheaves.

6.1 Maps of Posets and Associated Functors

As already noted, sheaves and cosheaves on posets are easier to manipulate. The functors associated to a map of posets are explicitly defined without extra processing. Since posets can be made into a topological space the functors which exist for sheaves on general spaces can be studied here is a more tightly controlled laboratory. Moreover, since Alexandrov spaces have extra symmetries new functors not normally encountered exist here.

Definition 6.2 (Map of Posets)**.** Suppose $(X, \leqslant_X)$ and $(Y, \leqslant_Y)$ are posets. A map of posets is a map of sets $f : X \to Y$ that is order-preserving, i.e. if $x \leqslant_X x'$ then $f(x) \leqslant_Y f(x')$. Alternatively, since a poset can be viewed as a category, a map of posets is just a functor. When it is clear from context we will abbreviate $(X, \leqslant_X)$ by just X.

Remark 6.3 (Notation and Cell Complexes). In our effort to treat posets as spaces, we have used X and Y to denote partially ordered sets equipped with the Alexandrov topology. This might cause confusion since our canonical example of a poset will be the indexing poset of a cell complex (X, P_X). Note that cell complexes consist of a pair of spaces, one is X, the Hausdorff space that is partitioned into pieces X_σ, the other is P_X, the poset of labels σ. From here on out we will work primarily with the poset P_X as this is the combinatorial approximation to X. Thus, keeping in line with Shepard [80], we change our notation from (X, P_X) to $(|X|, X)$. Thus we have the following dictionary:

Dictionary	Old Notation	New Notation
Underlying Hausdorff Space	X	$\|X\|$
Underlying Alexandrov Space	P_X	X
Set of Points in a "Cell"	X_σ	$\|\sigma\|$
"Cell" viewed as a point	σ	σ
Cellular Sheaf	$F : P_X \to \mathcal{D}$	$F : X \to \mathcal{D}$

6.1.1 Pullback or Inverse Image

Recall that a sheaf on a poset $(Y, \leqslant_Y)$ is a functor $G : Y \to \mathcal{D}$. Similarly, a cosheaf is a functor $\hat{F} : Y^{op} \to \mathcal{D}$. For both structures the pull-back functor f^* is easily described. It is the obvious pre-composition that completes the following diagram.

$$\begin{tikzcd} Y \arrow[r, "G"] & \mathcal{D} \\ X \arrow[u, "f"] \arrow[ur, dashed] & \end{tikzcd}$$

Definition 6.4 (Pullback for Poset Maps)**.** Given a sheaf G on Y and a map of posets $f : X \to Y$, we can define the **pullback** or **inverse image** sheaf f^*G on X as follows:

- $f^*G(x) = G(f(x))$
- If $x \leqslant x'$, then let $\rho^{f^*G}_{x',x} = \rho^G_{f(x'),f(x)}$
- If $\eta : G \to H$ is a morphism in $\mathbf{Shv}(Y)$, i.e. a natural transformation of diagrams over Y, then $f^*\eta : f^*G \to f^*H$ is a morphism in $\mathbf{Shv}(X)$ defined by declaring $f^*\eta(x) : f^*G(x) \to f^*H(x)$ to be equal to $\eta(f(x)) : G(f(x)) \to H(f(x))$.

The same definition and arguments go through for a cosheaf on Y with suitable modification, i.e $r^{f^*G}_{x,x'} = r^G_{f(x),f(x')}$. Thus, we get functors

$$f^* : \mathbf{Shv}(Y; \mathcal{D}) \to \mathbf{Shv}(X; \mathcal{D}) \qquad f^* : \mathbf{CoShv}(Y; \mathcal{D}) \to \mathbf{CoShv}(X; \mathcal{D}).$$

The definition of the pullback seems almost too good to be true, but one can check that the pre-sheaf description we outlined earlier agrees with this definition. Observe that if one applies that definition then

$$f^*F(U_x) := \varinjlim_{V \supset f(U_x)} F(V) \cong F(V_{f(x)}) = F(f(x)),$$

where we have used the fact that the smallest open set containing $f(U_x) = f(\{x' | x \leqslant x'\})$ is $V_{f(x)} = \{y | f(x) \leqslant y\}$.

Example 6.5 (Constant Sheaf and Cosheaf)**.** Consider the constant map $p : X \to \star$. A sheaf G on $\star$ consists of a single vector space W and the identity morphism so we'll just call G by the name W. We define the constant sheaf on X with value W to be $W_X := p^*W$. One sees that it is a sheaf that assigns W to every cell with all the restriction maps being the identity. Similarly, the constant cosheaf with value W is $\hat{W}_X := p^*W$.

6.1.2 Application: Subdivision

In the case where the poset is the face relation of a cell complex certain natural maps present themselves, such as subdivision.

Definition 6.6 ([80] 1.5, p.29)**.** A **subdivision** of a cell complex X is a cell complex X' with $|X'| = |X|$ and where every cell of X is a union of cells of X'.

Untangling the definition a bit we see that if σ is a cell of X, then there is a collection of cells $\{\sigma'_i\}$ such that $\cup_i |\sigma'_i| = |\sigma|$. As such, we can define a surjective map of posets $s : X' \to X$ defined by making $s(\sigma') = \sigma$ if $|\sigma'| \subseteq |\sigma|$.

Claim 6.7. *Subdivision of a cell complex* X *induces an order preserving map* $s : X' \to X$ *of the corresponding face-relation posets.*

Proof. The ordering on X' is given by the face relation. Suppose $\sigma' \leqslant \tau'$, then either $s(\sigma') = s(\tau')$ or not. If not, then σ' and τ' belong to the subdivision of two cells $\sigma \leqslant \tau$. $\square$

We are going to use this fact to define the subdivision of a sheaf in a cleaner manner than is found in [80].

Definition 6.8. Suppose F is a sheaf on X and $s : X' \to X$ is a subdivision of X, then we define the subdivided sheaf $F' := s^* F$.

6.1.3 Pushforward or Direct Image

By adopting a point-theoretic picture rather than an open set-theoretic picture of sheaves and cosheaves over posets, we got an easy definition for the pullback functor. In the introduction we outlined a general definition for the pushforward functor f_* on sheaves and cosheaves on an arbitrary topological space. Interestingly enough, although f_* had a simple description using open sets, the point-level description requires thought.

Definition 6.9 (Pushforward for Poset Maps)**.** Given a sheaf F on X and a map of posets $f : X \to Y$ we can define a sheaf on Y as follows:

- The **pushforward** of a sheaf is the right Kan extension of F along f, i.e. $\mathrm{Ran}_f F$.

$$f_*F(y) = \varprojlim_{f(x) \geqslant y} F(x)$$

- Suppose $y \leqslant y'$, then $\{x | f(x) \geqslant y'\} \subseteq \{x | f(x) \geqslant y\}$. Any limit over the bigger set defines a cone over the smaller set by restriction, thus the universal property of limits guarantees the existence of a unique map $f_*F(y) \to f_*F(y')$ that we will call $\rho^{f_*F}_{y',y}$.
- Suppose $\eta : F \to G$ is a map of sheaves, i.e. a natural transformation of diagrams over X. Then for any sub-poset U of X, post-composing the limit over U of F with the

arrows in the natural transformation defines a cone over G restricted to U. By the universal property of limits there must be an induced map.

$$\varprojlim_{x \in U} F \to \varprojlim_{x \in U} G$$

For cosheaves, the dual arguments go through with the slight modification that we use the left Kan extension along $f^{op} : X^{op} \to Y^{op}$.

$$f_* \hat{F}(y) := \varinjlim_{f(x) \geqslant y} \hat{F}(x).$$

Since both of these constructions are functorial, we have redefined two functors:

$$f_* : \mathbf{Shv}(X; \mathcal{D}) \to \mathbf{Shv}(Y; \mathcal{D}) \qquad f_* : \mathbf{CoShv}(X; \mathcal{D}) \to \mathbf{CoShv}(Y; \mathcal{D})$$

Example 6.10 (Global Sections)**.** This functor is extremely useful as it gives us a way of defining the global sections of a sheaf or a cosheaf. For the constant map $p : X \to \star$ we offer the following definitions:

$$p_* F(\star) \cong F(X) = \Gamma(X; F) = H^0(X; F) \qquad p_* \hat{F}(\star) \cong \hat{F}(X) = \Gamma(X; \hat{F}) = H_0(X; \hat{F})$$

In section 11 we will use this definition as the prototype for defining "higher" pushforward or direct image functors.

6.1.4 $f_\dagger$, Pushforwards and Closed Sets

One of the advantages of describing the standard functors of sheaf theory in the setting of posets is the presence of extra symmetries. Abstract definitions lend themselves to being dualized. In particular, in our point-theoretic definition of the pushforward we made use of Kan extensions, which come in two variants: left and right. In this section we consider the other variant and give a topological explanation for its origin.

Remark 6.11 (Caveat). When we discuss adjunctions in section 6.3 between these functors, we will see that for sheaves the functor $f_\dagger$ defined below is the left adjoint to f^* (for cosheaves it will be the right adjoint). There seems to be a strong trend to call the left adjoint of f^* by a different name: $f_!$. According to Joel Friedman ([32] p. 22) the tradition goes back to Grothendieck in [9] SGA Exposé I, Proposition 5.1. The same notation is used by Ladkani [47], Lurie, Beilinson, Bernstein and others.

This is unfortunate, since the notation $f_!$ is perhaps even more firmly established for the pushforward with compact supports functor used in classical sheaf theory. The reason seems to be that for general sheaves, there is no left adjoint to f^*, so it would be clear from context which was meant. However, for cellular sheaves, both functors exist and are useful.

Definition 6.12 (Pushforward with Open Supports)**.** Given a sheaf F on X and a map of posets $f : X \to Y$ we can define a sheaf on Y as follows:

- The **pushforward with open supports** of a sheaf is the left Kan extension of F along f, i.e. $\mathrm{Lan}_f F$.
$$f_\dagger F(y) = \varinjlim_{f(x)\leqslant y} F(x)$$

- If $y \leqslant y'$, then $\{x|f(x) \leqslant y\} \subseteq \{x|f(x) \leqslant y'\}$ and since any colimit over the bigger set defines a cocone over the smaller set by restriction, we get a unique map $\rho^{f_\dagger F}_{y',y} : f_\dagger F(y) \to f_\dagger F(y')$.

- If we have a map of sheaves $\eta : F \to G$, then pre-composing the arrows for colim G with η defines a co-cone over F. By universal properties we get an induced map
$$\varinjlim_{x\in V} F \to \varinjlim_{x\in V} G.$$

Dually, for cosheaves we use the right Kan extension along f^{op}.

$$f_\dagger \hat{F}(y) := \varprojlim_{f(x)\leqslant y} \hat{F}(x)$$

Both of these constructions are functorial and thus we have defined two functors:

$$f_\dagger : \mathbf{Shv}(X;\mathcal{D}) \to \mathbf{Shv}(Y;\mathcal{D}) \qquad f_\dagger : \mathbf{CoShv}(X;\mathcal{D}) \to \mathbf{CoShv}(Y;\mathcal{D})$$

This functor appears to be quite unusual, despite its naturality from the categorical perspective. To explain its topological origin, we revisit some of the original ideas of Alexandrov.

When Alexandrov first defined his topology he did two things differently:

1. He only defined the topology for *finite* posets.

2. He defined the *closed* sets to have the property that if $x \in V$ and $x' \leqslant x$, then $x' \in V$.

Let us repeat the initial analysis of sheaves and diagrams indexed over posets, where we now put closed sets on equal footing with open sets. Observe that as before we have an inclusion functor:

$$j : (X, \leqslant) \to \mathbf{Closed}(X) \qquad x \rightsquigarrow \bar{x} := \{x'|x' \leqslant x\}$$

Consequently, we have a similar diagram for a functor $F : X \to \mathcal{D}$ as before.

$$\begin{tikzcd} X \arrow[r, "F"] \arrow[d, "j"'] & \mathcal{D} \\ \mathbf{Closed}(X) \arrow[ur, dotted, "?"'] & \end{tikzcd}$$

If we choose the left Kan extension, we'd like to say the extended functor is a cosheaf on closed sets, i.e. use the definition of a cosheaf but replace open sets with closed sets.

Unfortunately, this concept is not well defined for general topological spaces because the arbitrary union (colimit) of closed sets is not always closed. For Alexandrov spaces this property does hold and this illustrates one of the extra symmetries this theory possesses.

However, in order for the Kan extension to take a diagram and make it into a cosheaf, we need to know whether the image of the inclusion functor defines a basis for the closed sets. In example 5.80 we showed that this is not always the case. The topology generated by the image of this functor is called the **specialization** topology and it suffers from certain technical deficiencies. In particular, order-preserving maps are not necessarily continuous in this topology, thus it fails to give a functorial theory. Fortunately, for finite posets these topologies agree and we can talk about cosheaves on closed sets without any trouble.

We now can give a topological explanation for the existence of the functor $f_\dagger$. It is the functor analogous to ordinary pushforward where we have adopted closed sets as the indexing category for cosheaves and sheaves. If $f : X \to Y$ is a map of posets, then f^c is the induced map between closed sets. The dagger pushforward is then the obvious completion of the diagram.

$$\begin{tikzcd} \mathbf{Closed}(X) \arrow[r, "F"] & \mathcal{D} \\ \mathbf{Closed}(Y) \arrow[u, "f^c"] \arrow[ur, dashed, "f_\dagger F"'] & \end{tikzcd}$$

In section 11.4 this functor provides the foundation for defining **sheaf homology** and **cosheaf cohomology** – theories that don't exist for general spaces.

6.1.5 $f_!$: Pushforward with Compact Supports on Cell Complexes

The three functors f^*, f_* and $f_\dagger$ induced by a map of posets are well defined for any poset and any diagram. However, when Shepard wrote his thesis the only posets that he considered were posets coming from cell complexes. By working in this smaller class and imitating the theory of constructible sheaves, Shepard described another functor that is not defined for arbitrary Alexandrov spaces: the pushforward with compact supports $f_!$.

This fourth functor is meant to provide a cellular (constructible) analog of a functor naturally defined for sheaves on more general topological spaces and the name is borrowed from there. The reader must keep this in mind since every set in a finite Alexandrov space is compact. Thus, when we say "pushforward with compact supports" we mean a discrete model for the pushforward with compact supports functor defined for locally compact Hausdorff spaces.

Following Shepard, this functor $f_!$ will only be defined for cellular maps, which are stratified (or even definable) maps naturally adapted to cell complexes.

Definition 6.13 (Cellular Map [80] pg. 32)**.** Let X and Y be cell complexes. A **cellular map** $(|f|, f)$ consists of a map of posets $f : X \to Y$ and a continuous "geometric" map $|f| : |X| \to |Y|$ satisfying the following compatibility conditions:

1. For every $\sigma \in X$, $|f|(|\sigma|)$ is the cell $|f(\sigma)|$.

2. The restricted map $|f|_{|\sigma|} : |\sigma| \to |f(\sigma)|$ is the projection $\mathbb{R}^{n+k} \to \mathbb{R}^n$ onto the first n coordinates.

3. Given $\sigma \in X$ and $y, z \in |f(\sigma)|$, $|f|^{-1}(y) \cap |\bar{\sigma}|$ is compact if and only if $|f|^{-1}(z) \cap |\bar{\sigma}|$ is.

Remark 6.14. The first and second conditions clearly restrict the types of maps of posets that can be considered. It appears that the third condition is redundant given the first two, but this is how it is recorded in Shepard's thesis.

Example 6.15. Let $X = [0,1)$ be given the cell structure $x = 0$ and $b = (0,1)$. Let $Y = [0,1) \times [0,1)$ be given the simplest possible cell structure. The underlying posets for these spaces are as follows:

$$\begin{tikzcd} & b \\ x \arrow[ur] & \end{tikzcd} \qquad \begin{tikzcd} & \sigma & \\ a \arrow[ur] & & b \arrow[ul] \\ & x \arrow[ul] \arrow[ur] \arrow[uu] & \end{tikzcd}$$

Here x refers to the vertex, a and b the open edges, and σ is the open face. Clearly, $f(x) = a$ and $f(b) = \sigma$ would be a map of these posets, but it is not a cellular map.

The definition of $f_!$ uses kernels and other standard linear algebra operations. As such, we now assume $\mathcal{D} = \mathbf{Vect}$ and suppress it from our notation.

Definition 6.16 (Pushforward with Compact Supports)**.** Given a sheaf F on X and a cellular map $f : X \to Y$, we can define the **pushforward with compact supports** sheaf on Y as follows:

- $f_!F(\tau) = \{s \in \Gamma(f^{-1}(\tau); F) \mid s(\sigma) = 0$ if $|\bar{\sigma}| \cap f^{-1}(y)$ not compact for $y \in |\tau|\}$

- Let $\gamma \leqslant \tau$ be cells in Y, and let $s \in f_!F(\gamma)$ and $t \in f_!F(\tau)$. We define $\rho^{f_!F}_{\tau,\gamma}(s) = t$ if for every $\sigma \in f^{-1}(\tau)$ and every $\lambda \in f^{-1}(\gamma)$ such that $\lambda \leqslant \sigma$ $\rho^F_{\sigma,\lambda}(s(\lambda)) = t(\sigma)$. If there is no such $t \in f_!F(\tau)$ then we define $\rho^{f_!F}_{\tau,\gamma}(s) = 0$.

The notation $\Gamma(-; F)$ for sections is explained in definition 5.89. The verification that $f_!F$ is actually a sheaf and that it is functorial, is much more drawn out and is done in detail in [80] pp. 35-38. As such we have defined a functor

$$f_! : \mathbf{Shv}(X) \to \mathbf{Shv}(Y)$$

Remark 6.17 (Compact Supports for Cosheaves). The definition for cosheaves cannot be written so simply because the vector space of "compactly supported" sections of a cosheaf, is a quotient of the space of all sections. The simplest definition would be, assuming $\hat{F} : X^{op} \to \mathbf{vect}$, to take transposes and turn $\hat{F}$ into a sheaf F and apply the definition above. We will not make use of the cosheaf version of this functor.

6.2 Calculated Examples

In this section we compute explicit examples of the functors defined above. To avoid clutter, we consider only sheaves and leave it the reader to dualize and check the corresponding functors on cellular cosheaves. We further assume that $\mathcal{D} = \textbf{Vect}$ and leave it as implicit that all operations are to be performed in vector spaces.

The notation $\square$ will be a place holder for any one of the three symbols $*, \dagger, !$.

6.2.1 Projection to a point

We consider the constant map $p : X \to \star$. The output of $p_\square F$ is a single vector space, namely $p_\square F(\star)$.

Figure 23: Projection to a Point

Without too much effort we compute the following:

- $p_* F = \varprojlim\{F(x') \to F(a')\} \cong F(x')$
- $p_\dagger = \varinjlim\{F(x') \to F(a')\} \cong F(a')$

For the pushforward with compact supports, we will be extra careful. Recall the definition states that $p_! F(\tau) = \{s \in \Gamma(p^{-1}(\tau); F) | s(\sigma) = 0 \text{ if } |\bar{\sigma}| \cap p^{-1}(y) \text{ not compact for } y \in |\tau|\}$.

In our example y can be the only point $\star$ and $p^{-1}(\star) = X$. Thus we have only two cells to check whether their closures are compact or not. Clearly $\bar{x}' = x'$ is compact, but $\bar{a}' = X$ is not compact. The definition then says that we only allow sections whose value on a' is zero.

- $p_! F = \ker(\rho_{x',a'} : F(x') \to F(a'))$

6.2.2 Inclusion into a Closed Interval

Here we encounter an open inclusion $j : X \to Y$. The first thing to note is that in this case, the value of $j_! F$ is not going to change since either $j^{-1}(y) = \{x\}$ or it is empty. Since points are closed and bounded, the compactness condition on $|\bar{\sigma}| \cap \{j^{-1}(y)\}$ is always satisfied.

We see in this example that

- $j_* F(x) = \varprojlim\{F(x') \to F(a')\} \cong F(x')$, $j_* F(a) = F(a')$, and less intuitively, $j_* F(y) \cong F(a')$.
- $j_\dagger F(x) = F(x')$, $j_\dagger F(a) \cong F(a')$, and $j_\dagger F(y) = \varinjlim\{\emptyset\} = 0$.
- $j_! F \cong j_\dagger F$.

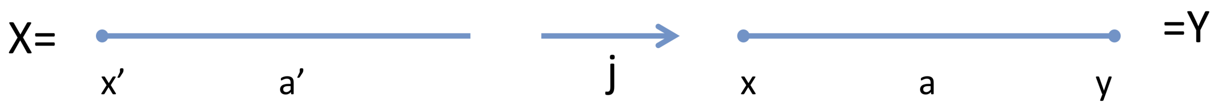

Figure 24: Inclusion into a Closed Interval

6.2.3 Map to a Circle

Here is an example where the function is bijective and continuous (in both topologies), but not an embedding, i.e. the domain is not homeomorphic with its image.

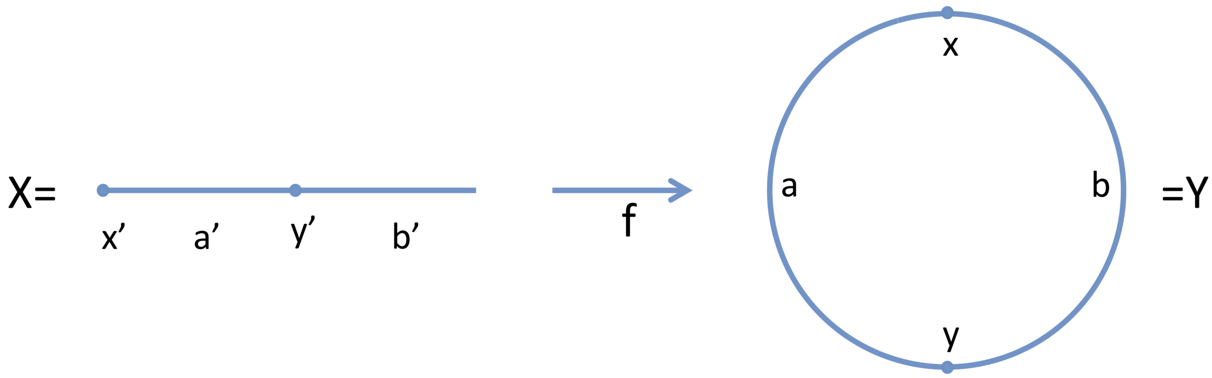

Figure 25: Map to a Circle

All three sheaves agree on the values and the restriction maps $f_\square F(y) \cong F(y') \to F(a') \cong f_\square F(a)$. We concentrate on the other two cells.

- Here diagram we are taking the limit over is disconnected because the inverse image of the star of x in Y is disconnected. Consequently, $f_* F(x) = \varprojlim\{F(x') \to F(a') \quad F(b')\} \cong F(x') \oplus F(b')$ and $f_* F(b) = F(b')$.

- Here $f_\dagger F(x) = F(x')$, but for similar reasons as before $f_\dagger F(b) = \varinjlim\{F(x') \quad F(y') \to F(b')\} = F(x') \oplus F(b')$.

- As noted, f is injective thus the value of $f_! F$ on any cell in the image is un-changed. However, we need to pay careful attention to how the restriction map is defined. The map f is injective so the fiber over a is a' and over x is x', but $x' \not\leq a'$ so the restriction map must be zero.

6.3 Adjunctions

Adjunctions allow us to derive interesting relationships with almost no effort. For the individual interested in using category theory to model the world, facile manipulations of adjunctions is essential. One of the tricks that one might hope to exploit is to transform a complicated problem into a simpler one via an adjunction, thus we get a computational payoff at the cost of abstraction. This is why using adjunctions between the functors defined in section 6.1 is one of the key technical skills every sheaf theorist must master.

Adjunctions also have played an essential role in the development of sheaf theory. Finding an adjoint to $f_!$ was one of the primary reasons that the derived categories were invented. Only by enlarging the domain could a new functor $f^!$ be defined.

Before addressing specifically the functors for sheaves and cosheaves, we review the general theory.

Definition 6.18. Suppose $F : \mathcal{C} \to \mathcal{D}$ and $G : \mathcal{D} \to \mathcal{C}$ are functors. We say that (F, G) is an **adjoint pair** or that F **is left adjoint to** G (or equivalently G is right adjoint to F) if we have a natural transformation $\eta : \mathrm{id}_{\mathcal{C}} \to G \circ F$ and a natural transformation to $\epsilon : F \circ G \to \mathrm{id}_{\mathcal{D}}$ such that

$$G \xrightarrow{\eta G} GFG \xrightarrow{G\epsilon} G, \qquad F \xrightarrow{F\eta} FGF \xrightarrow{\epsilon F} F$$

We call η the **unit** of the adjunction and ϵ the **counit** of the adjunction.

There are about a half-dozen different, but equivalent, ways of defining an adjunction; see [56] p. 81 for a list. One can just specify η and ask that it is universal,[31] i.e. for each $x \in \mathcal{C}$ and for every $y \in \mathcal{D}$ there is a map $\eta_x : x \to GF(x)$ such that if we have $f : x \to G(y)$, then there exists a unique map $f' : F(x) \to y$ with $G(f') \circ \eta_x = f$.

$$\begin{tikzcd} x \arrow[r, "\eta_x"] \arrow[dr, "f"'] & GF(x) \arrow[d, dotted] \\ & G(y) \end{tikzcd}$$

Of course we could have just defined ϵ and asked that it is universal in a dual sense.[32] The point is this – an adjunction is equivalent to specifying for every $x \in \mathcal{C}$ and $y \in \mathcal{D}$ a natural bijection $\varphi_{x,y}$

$$\mathbf{Hom}_{\mathcal{D}}(F(x), y) \cong \mathbf{Hom}_{\mathcal{C}}(x, G(y)).$$

The following theorem gives us an abstract criterion for determining when a functor has an adjoint. Coupled with our use of Kan extensions, we will be able to derive adjunctions with no additional work.

Theorem 6.19 (Freyd's Adjoint Functor Theorem)**.** *Let $\mathcal{D}$ be a complete category and $G : \mathcal{D} \to \mathcal{C}$ a functor, then G has a left adjoint F if and only if G preserves all limits and satisfies the **solution set condition**. This condition states that for each object $x \in \mathcal{C}$ there is a set I and an I-indexed family of arrows $f_i : c \to G(a_i)$ such that every arrow $f : x \to G(a)$ can be factored as $x \to G(a_i) \to G(a)$, where the first map is $f_i : x \to G(a_i)$ and the second is G applied to to some $t : a_i \to a$.*

[31]In other words, initial in a particular comma category; see [56] p. 56.

[32]It is final in a different comma category.

The solution set condition holds nearly all the time, so in practice one only needs to check that G preserves limits, in which case G is a right adjoint (has a left adjoint). Dually, for a functor to be a left adjoint it needs to preserve colimits.

Since in our construction of the functors associated to a map, we made explicit use of limits and colimits, corresponding to the right and left Kan extensions respectively, and (co)limits commute with (co)limits, the following theorems are automatic. However, we check them explicitly for sheaves and leave the dual proof for the reader to fill out on their own.

Theorem 6.20. *The functors* $f^* : \mathbf{Shv}(Y) \to \mathbf{Shv}(X)$ *and* $f_* : \mathbf{Shv}(X) \to \mathbf{Shv}(Y)$ *form an adjoint pair* (f^*, f_*) *and thus*

$$\mathbf{Hom}_{\mathbf{Shv}(X)}(f^*G, F) \cong \mathbf{Hom}_{\mathbf{Shv}(Y)}(G, f_*F).$$

Dually, the functors for cosheaves satisfy the opposite adjunction (f_*, f^*)

$$\mathbf{Hom}_{\mathbf{CoShv}(Y)}(f_* \hat{F}, \hat{G}) \cong \mathbf{Hom}_{\mathbf{CoShv}(X)}(\hat{F}, f^* \hat{G}).$$

Proof. Recall that $f^*(f_*F)(x) = (f_*F)(f(x))$. Using the fact that $(f_*F)(f(x)) = \varprojlim\{F(z)|f(z) \geqslant f(x)\}$, we get a map to $F(x)$ since $x \in f^{-1}(f(x))$ and this morphism is final for each x. This implies there is a natural transformation of functors $f^* \circ f_* \to \mathrm{id}$, which is universal (final).

Similarly, $f_*(f^*G)(y) = \varprojlim\{f^*G(x) = G(f(x))|f(x) \geqslant y\}$ and since $y \leqslant f(x)$ we can use the restriction map $\rho^G_{f(x),y} : G(y) \to G(f(x))$. The universal property of the limit guarantees a map $G(y) \to \varprojlim G(f(x)) = f_*f^*G(y)$ and thus a natural transformation of functors $\mathrm{id} \to f_*f^*$. $\square$

Theorem 6.21. *The functors* $f_\dagger : \mathbf{Shv}(X) \to \mathbf{Shv}(Y)$ *and* $f^* : \mathbf{Shv}(Y) \to \mathbf{Shv}(X)$ *form an adjoint pair* $(f_\dagger, f^*)$ *and thus*

$$\mathbf{Hom}_{\mathbf{Shv}(Y)}(f_\dagger F, G) \cong \mathbf{Hom}_{\mathbf{Shv}(X)}(F, f^*G).$$

Dually, the functors for cosheaves satisfy the opposite adjunction $(f^*, f_\dagger)$

$$\mathbf{Hom}_{\mathbf{CoShv}(X)}(f^* \hat{G}, \hat{F}) \cong \mathbf{Hom}_{\mathbf{CoShv}(Y)}(\hat{G}, f_\dagger \hat{G}).$$

Proof. $f_\dagger(f^*G)(y) = \varinjlim\{G(f(x))|f(x) \leqslant y\}$ so again we can use the restriction maps to define maps to $G(y)$. The universal property of colimits gives a map $f_\dagger f^*G(y) \to G(y)$ and thus a map of functors $f_\dagger f^* \to \mathrm{id}$. Similar arguments give a map $\mathrm{id} \to f^*f_\dagger$ $\square$

To conclude, we derive the first interesting consequence of an adjunction. In effect it reduces all the possible natural transformations between a certain pair of functors to a single vector space.

Proposition 6.22. *If* $F : X \to \mathbf{Vect}$ *is a sheaf and* $p : X \to \star$ *is the constant map, then*

$$\mathbf{Hom}_{\mathbf{Shv}(X)}(p^*k, F) \cong \mathbf{Hom}_{\mathbf{Vect}}(k, p_*F) \cong F(X) = H^0(X; F).$$

Proof. The first isomorphism is the adjunction (p^*, p_*). The second isomorphism is simply the observation that every linear map is determined by where it sends 1, i.e. $\mathbf{Hom}_{\mathbf{Vect}}(k, W) \cong W$. $\square$

7 Homology and Cohomology

In sections 5.2 and 6 we worked over arbitrary posets. We did this because it was natural and some applications may need this level of generality. In this section, we eschew this generality and restrict ourselves to posets arising as the face relation of a finite cell complex. This is beneficial not only because cell complexes are of great interest, but because sheaves and cosheaves over them have easily defined cohomology and homology theories.

We will start by describing a simple generalization of cellular cohomology and homology where we have augmented the coefficients by placing vector spaces over individual cells and linear maps between incident cells. This is a generalization in the sense that if one restricts to the case where every cell is assigned the one-dimensional vector space k and all the incident linear maps are the identity, we recover classical cellular (co)homology. However interesting this special case may be, it misses a theory general enough to compute homological invariants of data varying over a cell complex.

The theory presented is combinatorial and computable. One needs only a good working knowledge of linear algebra to be able to use it. However, one can compute cellular sheaf cohomology without understanding it. To clarify the meaning of these computations we adopt a representation-theoretic perspective. This allows us to break up sheaves and cosheaves into the basic building blocks of indecomposable representations of the cell category. Thus, borrowing terminology from the persistent homology community, we use "generalized barcodes" to see the topology of data in a wider world of applications. These ideas are be put into practice in sections 8, 9, and 10, where many examples are considered.

7.1 Computational Sheaf Cohomology and Cosheaf Homology

To introduce cellular sheaf cohomology, we analyze some of the combinatorial properties that cell complexes have. These combinatorial properties are not shared with all posets. This is why sheaves over cell complexes have a computationally tractable cohomology theory.

Definition 7.1. We write $\sigma \leqslant_i \tau$ if the difference in dimension of the cells is i.

Lemma 7.2. *If $\sigma \leqslant_2 \tau$, then there are exactly two cells λ_1, λ_2 where $\sigma \leqslant_1 \lambda_i \leqslant_1 \tau$.*

We want to invent a sign condition that distinguishes these two different sequences of incidence relations.

$$\begin{tikzcd}
 & \tau & \\
\lambda_1 \arrow[ur] & & \lambda_2 \arrow[ul] \\
 & \sigma \arrow[ul] \arrow[ur] &
\end{tikzcd}$$

Definition 7.3 (Signed Incidence Relation)**.** A **signed incidence relation** is an assignment to any pair of cells $\sigma, \tau \in X$ a number $[\sigma : \tau]$ that is zero if $\sigma \not\leqslant \tau$ and is otherwise ± 1.

One way to get a signed incidence relation is to just choose a **local orientation** (via the homeomorphism of each cell $|\sigma|$ with $\mathbb{R}^k$) for each cell without regard to global consistency. Then for every pair of incident cells $\sigma \leqslant \tau$ we have a number $[\sigma : \tau] = \pm 1$ given by $+1$ if the orientations agree and -1 otherwise.

Another way is motivated by working with regular cell complexes, where we can subdivide so that we have a simplicial complex. We can refer to any cell by a list of its vertices. If we order the set of vertices, then we have a procedure for orienting the cells. A local orientation of a cell $\sigma \in X$ consists of divvying up the set of ordered lists representing σ into classes each of which are invariant under even permutations. We can then pick the class with the list of vertices in increasing order as "the" orientation. Either method allows us to define a complex of vector spaces associated to either a cellular sheaf or a cosheaf.

7.1.1 Cellular Sheaf Cohomology

Definition 7.4 ([100, 80])**.** Given a cellular sheaf $F : X \to \textbf{\textit{Vect}}$ we define its **compactly supported** k **co-chains** to be the product[33] of the vector spaces residing over all the k-dimensional cells.

$$C_c^k(X; F) = \bigoplus_{\sigma^k} F(\sigma^k)$$

These vector spaces are graded components in a complex of vector spaces $C_c^\bullet(X; F)$. The differentials are defined by

$$\delta_c^k = \sum_{\sigma \leqslant \tau} [\sigma^k : \tau^{k+1}] \rho_{\tau,\sigma}.$$

The cohomology of this complex

$$0 \longrightarrow \oplus F(\text{vertices}) \xrightarrow{\delta_c^0} \oplus F(\text{edges}) \xrightarrow{\delta_c^1} \oplus F(\text{faces}) \longrightarrow \cdots \qquad = C_c^\bullet(X; F)$$

is defined to be the **compactly supported cohomology** of F, i.e. $H_c^k(X; F) = \ker \delta_c^k / \operatorname{im} \delta_c^{k-1}$.

Lemma 7.5. $(C_c^\bullet(X; F), \delta_c^\bullet)$ *is a chain complex.*

Proof. To see why the chain complex condition $\delta_c^{k+1}\delta_c^k = 0$ is assured, lemma 7.2 is crucial. This is the very same lemma that proves that ordinary cellular homology is computed via a chain complex. One must now observe that varying data over the cells does not change the

[33]Here we implicitly assume that X has finitely many cells in a given dimension so products and direct sums agree.

result.

$$\begin{aligned}
\delta_c\delta_c &= \sum_{\sigma\leqslant_1\tau}[\sigma:\tau]\rho_{\tau,\sigma}(\delta_c)\\
&= \sum_{\sigma\leqslant_1\tau}[\sigma:\tau]\rho_{\tau,\sigma}(\sum_{\gamma\leqslant_1\sigma}[\gamma:\sigma]\rho_{\sigma,\gamma})\\
&= \sum_{\gamma\leqslant_1\sigma\leqslant_1\tau}[\gamma:\tau]\rho_{\tau,\sigma}\rho_{\sigma,\gamma}\\
&= \sum_{\gamma\leqslant_1\sigma\leqslant_1\tau}[\gamma:\tau]\rho_{\tau,\gamma}\\
&= \sum_{\gamma\leqslant_1\sigma\leqslant_1\tau}([\gamma:\sigma_1][\sigma_1:\tau]+[\gamma:\sigma_2][\sigma_2:\tau])\rho_{\tau,\gamma}\\
&= 0
\end{aligned}$$

□

To define the arbitrarily-supported cochain complex associated to a cellular sheaf F on X, we simply remove all the cells from X without compact closures and apply the same formula.

Definition 7.6 (Ordinary Cohomology)**.** Let X be a cell complex and $F : X \to \textbf{\textit{Vect}}$ a cellular sheaf. Let $j : X' \to X$ be the subcomplex consisting of cells that do not have vertices in the one-point compactification of X. Define the ordinary cochains and cohomology by

$$C^\bullet(X;F) = C^\bullet_c(X';j^*F) \qquad H^i(X;F) := H^i_c(X';j^*F)$$

The situation may seem a bit unusual. The naturally defined chain complex computes a more restrictive type of cohomology. To get the standard cohomology, one needs to remove non-compact cells. When we define cohomology via the derived perspective of section 11, this quirk of linear algebra disappears. Ordinary cohomology will fall out naturally using limits and injective resolutions, and compactly-supported sheaf cohomology will require some finesse.

Example 7.7 (Compactly Supported vs. Ordinary Cohomology)**.** To see why the naïve chain complex computes compactly supported cohomology, consider the example of the half-open interval $X = [0, 1)$ decomposed as $x = \{0\}$ and $a = (0, 1)$. Now consider the constant sheaf k_X. To compute compactly supported cohomology, we must first pick a local orientation of our space. By choosing the orientation that points to the right, we get that $[x : a] = -1$. The cohomology of our sheaf is computed via the complex

$$k \xrightarrow{-1} k\,,$$

which yields $H^0_c = H^1_c = 0$. If we follow the prescription for computing ordinary cellular sheaf cohomology, then we must remove the vector space sitting over a in our computation. The resulting complex is simply the vector space k placed in degree 0, so $H^0(X;k_X) = k$ and is zero in higher degrees.

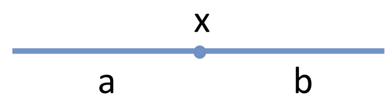

Figure 26: Minimal Cell Structure on an Open Interval

Example 7.8 (Open Interval)**.** If we pretended for a moment that the pure stratum $Y = (0,1)$ is a cell complex[34] with no other cells, then computing the compactly supported cohomology of the constant sheaf would yield a vector space in degree one and nowhere else, hence $H^1_c(Y; k_Y) = k$.

To make this example a legitimate example, as in figure 26, we place a vertex at $x = 1/2$. We call our new cells $a = (0, 1/2)$ and $b = (1/2, 1)$. If we orient our 1-cells to point to the right, then $[x : a] = 1$ and $[x : b] = -1$. Using the lexicographic ordering on our cells to get a basis for $C^1_c(Y; k_Y)$ we can compute explicitly the compactly supported cohomology.

$$\delta^0_c = \begin{bmatrix} 1 \\ -1 \end{bmatrix} : k_x \to k_a \oplus k_b \qquad \Rightarrow \qquad H^0_c = 0 \quad H^1_c = k$$

7.1.2 Cellular Cosheaf Homology

For cellular cosheaves the exact dual construction works, but the terminology is slightly different.

Definition 7.9 (Borel-Moore Cosheaf Homology)**.** Let X be a cell complex and let $\hat{F} : X^{op} \to \mathbf{Vect}$ be a cellular cosheaf. Define the **Borel-Moore homology** $H^{BM}_\bullet(X; \hat{F})$ to be the homology of the following complex:

$$C^{BM}_\bullet(X; \hat{F}) = \qquad \cdots \longrightarrow \oplus \hat{F}(\text{faces}) \xrightarrow{[e:\sigma]r_{e,\sigma}} \oplus \hat{F}(\text{edges}) \xrightarrow{[v:e]r_{v,e}} \oplus \hat{F}(\text{vertices}) \longrightarrow 0$$

Definition 7.10 (Ordinary Cosheaf Homology)**.** Let X be a cell complex and let $\hat{F} : X^{op} \to \mathbf{Vect}$ be a cellular cosheaf. By discarding all the cells without compact closure, we obtain the maximal compact subcomplex X'. If we write $j : X' \hookrightarrow X$ for the inclusion, then we can define the ordinary chain complex to be

$$C_\bullet(X; \hat{F}) = C^{BM}_\bullet(X'; j^* \hat{F}).$$

Applying the definition above gives the ordinary **cosheaf homology** $H_\bullet(X; \hat{F})$ of a co-sheaf.

All of the examples of cellular sheaf cohomology dualize to give interesting examples of cellular cosheaf homology. Let us define the functor that performs this operation.

[34] Recall that we require a cell complex to have a one point compactification that is a regular cell complex.

Definition 7.11 (Linear Duality)**.** Let $\widehat{V} : \mathbf{Shv}(X; \mathbf{vect}_k)^{op} \to \mathbf{CoShv}(X; \mathbf{vect}_k)$ be the contravariant equivalence from sheaves to cosheaves, both valued in *finite* dimensional vector spaces, defined as follows:

$$\begin{tikzcd} F(\tau) \arrow[r, squiggly] & F(\tau)^* \arrow[d, "\rho^*_{\tau,\sigma}"] & \widehat{V}(F)(\tau) \arrow[d, "r_{\sigma,\tau}"] \\ F(\sigma) \arrow[u, "\rho_{\tau,\sigma}"] \arrow[r, squiggly] & F(\sigma)^* & \widehat{V}(F)(\sigma) \end{tikzcd}$$

Lemma 7.12. *Taking linear duals preserves cohomology, i.e.* $H_i(X; \widehat{V}(F)) \cong H^i(X; F)$ *and* $H_i^{BM}(X; \widehat{V}(F)) \cong H_c^i(X; F)$.

7.2 Explaining Homology and Cohomology via Indecomposables

Sheaf cohomology is notoriously difficult to interpret. Every time a successful interpretation is discovered, a cornerstone of a theory waiting to be fleshed out is put into place. For example, the Cousin problems of complex analysis asks whether a meromorphic function with a given divisor (zeros and poles) exists or not. When Cartan and Serre interpreted this problem in terms of sheaf theory, sheaf cohomology groups gave a complete classification and obstruction theory; see [42] p. 17. The narrative that falls out of those historical successes is that sheaf cohomology gives calculable obstructions to finding solutions.

However, when the above interpretation fails, we need to compute examples and extract new interpretations. When computing sheaf cohomology, one encounters a plethora of choices that obfuscate the natural meaning of the vector spaces $H^i(X; F)$: picking ordered bases for each $F(\sigma)$, choosing local orientations, computing kernels and quotients, taking representative elements of cohomology or homology, etc. Each of these lead one farther from a workable interpretation of the topology of data.

The experience of the author in computing examples of sheaf cohomology has led him to believe that the best way of circumventing these issues is to borrow an idea from the representation theory of quivers. Specifically, if one knows the direct sum decomposition of a sheaf into indecomposable sheaves, then one gets a distinguished basis for sheaf cohomology. These indecomposables allow one to see how data travels through a space.

7.2.1 Representation Theory of Categories and the Abelian Structure

For the purposes of this section, there is no real difference between cellular sheaves and cosheaves – they are both representations of the cell category $\mathbf{Cell}(X)$. Recall, for any category $\mathcal{C}$ the **category of representations** is defined to be the category of functors to $\mathbf{Vect}$:

$$\mathbf{Rep}(\mathcal{C}) := \mathbf{Fun}(\mathcal{C}, \mathbf{Vect})$$

This category has the structure of an abelian category, which we explain in this section. In effect, this means we can do everything in $\mathbf{Rep}(\mathcal{C})$ that we can do in $\mathbf{Vect}$: take kernels and cokernels of maps between representations, talk about images of maps, add maps and so on. We will introduce these properties as we need them.

Claim 7.13. *For $\mathcal{C}$ a category, $\mathbf{Rep}(\mathcal{C})$ is an **exact category**. This means we can talk about exact sequences. Specifically:*

- *There is a **zero representation** given by sending all objects and morphisms to the zero object and the zero morphism.*
- *Between any two representations F and G there is a **zero map**, which can be factored through the zero representation.*
- *Since any morphism $\eta : \mathsf{F} \to \mathsf{G}$ is a natural transformation occurring inside $\mathbf{Vect}$, there are associated **kernel and cokernel representations** denoted $\ker(\eta)$ and $\operatorname{coker}(\eta)$ defined by taking kernels and cokernels object wise:*

$$\begin{tikzcd}
\ker(\eta(c')) \arrow[r] & \mathsf{F}(c') \arrow[r, "\eta(c')"] & \mathsf{G}(c') \arrow[r] & \operatorname{coker}(\eta(c')) \\
\ker(\eta(c)) \arrow[r] \arrow[u, dotted] & \mathsf{F}(c) \arrow[r, "\eta(c)"] \arrow[u, "\mathsf{F}(f)"'] & \mathsf{G}(c) \arrow[r] \arrow[u, "\mathsf{G}(f)"'] & \operatorname{coker}(\eta(c)) \arrow[u, dotted]
\end{tikzcd}$$

- *There is an **image representation** $\operatorname{im}(\eta)$ defined as the object-wise image.*

As usual, we say that a sequence of representations $\mathsf{A} \to \mathsf{B} \to \mathsf{C}$ is **exact** at B if the kernel of the outgoing morphism is equal to the image of the incoming morphism. A longer sequence is exact if it is exact at each place with an incoming and outgoing morphism.

We are going to do a brief sketch of some representation theory for categories, using the terminology introduced.

Definition 7.14. A **subrepresentation** E consists of a choice of subspace $\mathsf{E}(c) \to \mathsf{F}(c)$ for each object that is invariant under all the linear maps $\mathsf{F}(f)$. Restriction of $\mathsf{F}(f)|_{\mathsf{E}(c)} =: \mathsf{E}(f)$ makes E into a representation of its own right. Said more succinctly, $\mathsf{E} \to \mathsf{F}$ is a natural transformation of functors that is object wise an inclusion, i.e.

$$0 \to \mathsf{E} \to \mathsf{F}$$

is an exact sequence. Dually, we can say G is a quotient representation by saying $\mathsf{F} \to \mathsf{G} \to 0$ is an exact sequence.

Definition 7.15. Suppose $\mathsf{F} : \mathcal{C} \to \mathbf{Vect}$ and $\mathsf{G} : \mathcal{C} \to \mathbf{Vect}$ are two representations of a small category $\mathcal{C}$, then we can define the **direct sum** of these two representations $\mathsf{H} = \mathsf{F} \oplus \mathsf{G}$ by defining on objects $\mathsf{H}(c) := \mathsf{F}(c) \oplus \mathsf{G}(c)$ and on morphisms $\mathsf{H}(f) = \mathsf{F}(f) \oplus \mathsf{G}(f) : \mathsf{H}(c) \to \mathsf{H}(c')$.

The above definition further clarifies the structure of $\mathbf{Rep}(\mathcal{C})$.

Claim 7.16. *For $\mathcal{C}$ a category, $\mathbf{Rep}(\mathcal{C})$ is both an exact and an **additive category**. This latter definition requires the following:*

- *For any two representations F and G the set $\mathbf{Hom}_{\mathbf{Rep}(\mathcal{C})}(\mathsf{F}, \mathsf{G})$ has the structure of an abelian group (with the zero map being the additive identity) making composition bilinear.*

- *The direct sum of two representations is a representation.*

*A category that is exact and additive is defined to be **abelian**. Thus $\mathbf{Rep}(\mathcal{C})$ is an abelian category.*

Remark 7.17. In any additive category, it can be shown that having finite direct sums (finite coproducts) implies the existence of finite direct products (finite products) and these are isomorphic.

Definition 7.18 (Indecomposable)**.** A representation $\mathsf{F} : \mathcal{C} \to \mathbf{Vect}$ is called **indecomposable** if whenever F is written as a direct sum of representations one of the representations is the zero one; i.e. every direct sum decomposition is trivial.

Said using sequences, a representation F is indecomposable if whenever we have a short exact sequence of representations

$$0 \to \mathsf{E} \to \mathsf{F} \to \mathsf{G} \to 0$$

with neither E nor G the zero representation, then $\mathsf{F} \ncong \mathsf{E} \oplus \mathsf{G}$, i.e. the sequence does not split.

Remark 7.19 (Splitting Lemma). There is a general lemma called the **splitting lemma**, which provides equivalent ways of saying that F is indecomposable. It states that writing F as a direct sum is equivalent to either having a map back from F to E, which precomposed with the inclusion $\mathsf{E} \to \mathsf{F}$ yields the identity, or having a map back from G to F, which post-composed with the surjection is the identity on G.

Definition 7.20 (Remak Decomposition)**.** A direct sum decomposition of an object $\mathsf{F} \in \mathbf{Rep}(\mathcal{C})$

$$\mathsf{F} \cong \mathsf{F}_1 \oplus \cdots \oplus \mathsf{F}_n$$

where each F_i is indecomposable and non-zero is called a **Remak decomposition**.

A fact that we would very much like to know is whether every representation admits a Remak decomposition. Sir Michael Atiyah considered such a question in the very general setting of abelian categories [11]. He developed a bi-chain condition and proved that under this condition every non-zero object admitted a Remak decomposition. We use a stronger condition of finite-dimensionality that Atiyah showed implied his bi-chain condition.

Theorem 7.21 (Krull-Schmidt Theorem for Representations [11])**.** *Suppose $\mathcal{A}$ is an abelian category, further satisfying*

1. *For every pair of objects* $\mathbf{Hom}_{\mathcal{A}}(A, B)$ *is a finite dimensional vector space, and*

2. *Conjugation is linear, i.e. for every pair of morphisms* $\varphi : A \to B$ *and* $\psi : B' \to A'$ *the following map is linear*

$$\mathbf{Hom}_{\mathcal{A}}(B, B') \to \mathbf{Hom}_{\mathcal{A}}(A, A') \qquad \eta \mapsto \psi \circ \eta \circ \varphi$$

then the Krull-Schmidt theorem holds. This says that every non-zero object A *has a Remak decomposition and for any two such decompositions*

$$A \cong A_1 \oplus \cdots \oplus A_n \qquad A \cong A'_1 \oplus \cdots \oplus A'_m$$

$n = m$ *and after re-ordering* $A_i \cong A'_i$.

For $\mathcal{A} = \mathbf{Rep}(\mathcal{C})$ the second condition is certainly satisfied. The first condition imposes significantly stronger conditions. First of all, we must restrict to the full subcategory of finite dimensional representations.

$$\mathbf{Rep}_f(\mathcal{C}) := \mathbf{Fun}(\mathcal{C}, \boldsymbol{vect}) \subset \mathbf{Fun}(\mathcal{C}, \mathbf{Vect}) =: \mathbf{Rep}(\mathcal{C})$$

Secondly, one must observe that for any two representations F and G the space of natural transformations is a subspace of a potentially infinite product of finite dimensional spaces.

$$\mathbf{Hom}(F, G) \subseteq \prod_{c \in \mathcal{C}} \mathbf{Hom}_{\boldsymbol{vect}}(F(c), G(c))$$

One severe restriction one can make to insure that Atiyah's first condition holds is to assume that the category $\mathcal{C}$ has finitely many objects. This is not strictly necessary, but it does provide us with the following corollary:

Corollary 7.22 (Sheaves and Cosheaves on Finite Posets have Remak Decompositions)**.** *Suppose* $(X, \leqslant)$ *is a finite poset, then* $\mathbf{Shv}(X)$ *and* $\mathbf{CoShv}(X)$ *satisfy the Krull-Schmidt theorem.*

The example that we have in mind, of course, is the poset associated to a a cell complex X. In this situation, one can recognize a large set of examples of indecomposable representations.

Lemma 7.23 (Constant (Co)Sheaves are Indecomposable)**.** *Suppose* X *is a connected cell complex, then the constant sheaf* k_X *and the constant cosheaf* $\hat{k}_X$ *are indecomposable.*

Proof. We'll state the proof for sheaves and leave it to the reader to dualize for cosheaves. Suppose for contradiction that $k_X \cong F \oplus G$ where neither F nor G is the zero sheaf. Now as a consequence of neither F nor G being zero, and k_X being one dimensional on each cell, there must be a pair of cells σ and τ such that one is in the support of F and the other is in the support of G. We argue that we can choose σ and τ such that one is the face of the other. If not, then the support of F (or G) would be closed under the following operations

$$\sigma \subset \text{supp}(F) \quad \text{and} \quad \tau \subset \bar{\sigma} \quad \text{or} \quad \sigma \subset \bar{\tau} \qquad \text{then} \qquad \tau \subset \text{supp}(F)$$

which by connectedness of X would imply that $\text{supp}(F) = X$; a contradiction to the supposition that neither F nor G was the zero sheaf. (To see why $\text{supp}(F) = X$, one can imagine drawing the Hasse diagram of the poset X and realizing that connectedness means that the diagram is connected.) Thus we have such a pair $\sigma \subset \tau$ with one in the support of F and the other in the support of G, but this also can not occur since the identity cannot be written as a sum of zero maps.

$$k \to k \neq (k \to 0) \oplus (0 \to k).$$

□

7.2.2 Quiver Representations and Cellular Sheaves

There are natural examples of representations of categories where these ideas and their consequences have been studied. One such example is the category of quiver representations.

A **quiver** or **directed graph** is nothing more than a pair of sets consisting of "edges" E and "vertices" V along with a pair of functions $h, t : E \to V$, which we think of as denoting the head and tail of a directed edge respectively. Alternatively, a quiver can be topologically regarded as a one-dimensional cell complex equipped with a local orientation of its edges.

One should be careful to note that a directed graph is not a category in and of itself; rather, there is a natural category associated to a directed graph, which we now define.

Definition 7.24. To a quiver we can associate a category called the **free category** or **path category** written $\mathbf{Free}(X)$. The objects are vertices and the morphisms are directed paths between vertices. Since paths are just concatenated edges, we think of the morphisms as being freely generated by the edges. We must consider simply sitting at a vertex as the identity directed path connecting the vertex to itself.

A **quiver representation** is thus nothing more than a functor $F : \mathbf{Free}(X) \to \mathbf{Vect}$. Because a general path is simply a sequence of edges, such a functor is equivalent to specifying a vector space for each vertex in V and a linear map for each edge in E that goes from the source to the target.

From a quiver representation, one can always construct a cellular sheaf or cosheaf. This comes from treating the quiver as a one-dimensional cell complex and then turning every map

$$F(s(e)) \xrightarrow{\rho_{t,s}} F(t(e))$$

into one of the following diagrams:

$$F(s(e)) \xrightarrow{\rho_{t,s}} F(t(e)) = F(e) \xleftarrow{\text{id}} F(t(e)) \qquad F(s(e)) \xleftarrow{\text{id}} F(t(e)) = F(e) \xrightarrow{\rho_{t,s}} F(t(e))$$

The former choice would make a quiver representation into a cellular sheaf, the latter into a cellular cosheaf.

There are dangers in trying to use quiver theory as a substitute for cellular sheaf or cosheaf theory. One might try to think of a poset as a certain type of quiver with vertices corresponding to elements and an edge between two elements if $s(e) \leqslant t(e)$. For example, consider the poset coming from the face relation of the cell complex $Y = [0, 1) \times [0, 1)$:

$$\begin{tikzcd}
 & \sigma & \\
a \arrow[ur] & & b \arrow[ul] \\
 & x \arrow[ul] \arrow[ur] \arrow[uu] &
\end{tikzcd}$$

A quiver representation produces a diagram of vector spaces

$$\begin{tikzcd}
 & F(\sigma) & \\
F(a) \arrow[ur] & & F(b) \arrow[ul] \\
 & F(x) \arrow[ul] \arrow[ur] \arrow[uu] &
\end{tikzcd}$$

that does not commute. In contrast, if F were a cellular sheaf, then the two triangles would commute. If we were to impose "relations" on the quiver representation by identifying different paths, then we could recover cellular sheaves. See [34] for more related to this viewpoint.

8 Barcodes: Persistent Homology Cosheaves

In section 7.2 we proposed the adoption of representation theory as a principled way for interpreting cellular sheaf cohomology and cellular cosheaf homology. In this section, we put that idea to the test by considering cosheaves that are naturally associated to stratified spaces and maps. In so doing, we synthesize the combinatorial theory of sections 5.2, 6 and 7 with the geometric theory of section 5.1.

We observed in theorem 5.72 that stratified maps give rise to cellular cosheaves by taking ith homology of the fiber. That is, given a stratified map $f : Y \to X$ with X stratified as a cell complex, we get a cellular cosheaf as follows:

$$\hat{F}_i(\sigma) = H_i(f^{-1}(t); k) \qquad t \in \sigma$$

These examples provide excellent test cases for interpreting cellular cosheaf homology because it should be directly related to the homology of the domain space Y.

To keep our examples simple, in section 8.1 we focus on the case where $X = \mathbb{R}$ or $[0, 1]$ with some cell structure. Since this case incorporates Morse theory, level and sub-level set persistence [21], we begin with an introduction to the language of "barcodes," which is a catchy name for certain one-dimensional indecomposable representations. We also compute cosheaf homology using the formulas of section 7.1 and show how barcodes provide a visual proxy for these computations. Finally, in section 8.2 we show how one can use a stratified map and its associated barcodes to compute the total homology of the domain space. We do this via a spectral sequence argument, which generalizes to higher dimensional base spaces. This generality suggests that using cosheaves associated to stratified maps provides a natural setting for multi-dimensional persistence.

8.1 Barcodes and One-Dimensional Persistence

For any cell structure on $[0, 1]$ a cellular cosheaf $\hat{F}$ is nothing other than a diagram of vector spaces and maps of the following form:

$$V_1 \leftarrow V_2 \to V_3 \leftarrow \cdots \to V_{n-2} \leftarrow V_{n-1} \to V_n$$

Such diagrams are called **zigzag modules** in [20] and they are built up from indecomposable representations of a particularly simple form. The indecomposables consist of an interval of one-dimensional vector spaces connected by identity maps, with zero maps and spaces on either side.

$$0 \leftarrow 0 \to k \leftarrow \cdots \to k \leftarrow 0 \to 0$$

Since indecomposables look like solid stretches of one-dimensional vector spaces, a convenient way of depicting them visually is to simply draw a solid black line between the first and last non-zero vector space in the indecomposable. By the Krull-Schmidt theorem, we can use this set of solid black lines as a unique visual identifier of the Remak decomposition of a representation. One need only scan this set of lines to be able to uniquely identify the representation. As such, they are colloquially referred to as **barcodes**.[35]

[35]The author is unsure if this line of reasoning went into the adoption of this term.

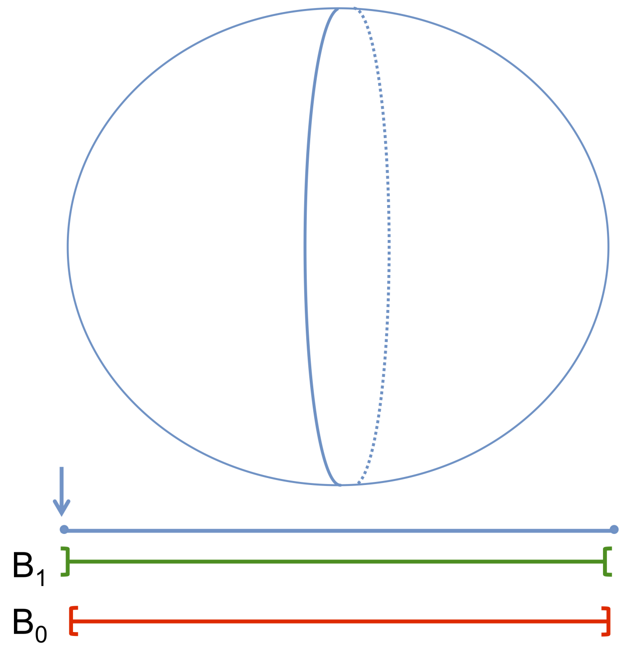

Figure 27: Barcodes and the Two Sphere

Remark 8.1 (Barcode Notation). We will use intervals to represent barcodes and these will be sensitive to whether the first or last vector space in an indecomposable representation falls on a vertex or an open interval in some stratification of $[0,1]$ or $\mathbb{R}$. For visual clarity, we will adopt the convention that a turned around square bracket is equivalent to a round one, i.e. $(x_i, x_{i+1}) \rightsquigarrow]x_i, x_{i+1}[$ and $[x_i, x_{i+1}) \rightsquigarrow [x_i, x_{i+1}[$ and so on.

Example 8.2 (Height function on the Two Sphere)**.** Let $h : S^2 \to \mathbb{R}$ be the standard height function on the two sphere. This is an example of a stratified map that is also a Morse function with two critical values. The Morse function property is nice because it allows us to flow the fiber over non-critical values to critical ones by using the gradient vector field (swapping the sign when appropriate). In figure 27 we have drawn the map and the associated barcodes. The barcode decomposition for the cosheaf associated to taking H_0 of the fiber is trivial because it is already an indecomposable cosheaf.

$$\hat{F}_0 : \qquad k \longleftarrow k \longrightarrow k$$

Similarly, taking H_1 of the fiber also yields an indecomposable cosheaf

$$\hat{F}_1 : \qquad 0 \longleftarrow k \longrightarrow 0$$

Now let us compute cosheaf homology. Since the space $X = [0,1]$ is compact, ordinary and compactly supported cosheaf homology agree. We label our cells as $x = 0$, $a = (0,1)$ and $y = 1$. To get an ordered basis and matrix representatives for our homology computation, we choose the local orientation pointing to the right and use the lexicographic ordering on the cells. For $\hat{F}_0$ we get the following boundary matrix and homology groups:

$$\partial_1 = \begin{bmatrix} -1 \\ 1 \end{bmatrix} : k_a \to k_x \oplus k_y \qquad \Rightarrow \qquad H_1(X;\hat{F}_0) = 0 \quad H_0(X;\hat{F}_0) = k.$$

For $\hat{F}_1$ the computation is even easier:

$$\partial_1 : k_a \to 0 \qquad \Rightarrow \qquad H_1(X;\hat{F}_1) = k \quad H_0(X;\hat{F}_1) = 0$$

It appears that whether the barcode lives over an open or closed interval impacts its cosheaf homology. Let us consider a map that leads naturally to a half-open barcode.

Example 8.3 (Height function on a Cone)**.** The height function on a cone is not a Morse function because differentiability breaks down at the cone point; see figure 28. One could use stratified Morse theory as a substitute, but we'll use cosheaf homology. Here the cosheaf $\hat{F}_0$ is the same as the previous example; we will not repeat the computation. The cosheaf $\hat{F}_1$

$$\hat{F}_1 : \qquad 0 \longleftarrow k \longrightarrow k$$

exhibits different behavior. The cosheaf homology computation for this cosheaf reveals that the half-open barcode, embedded inside a compact interval, has no non-zero homology groups.

$$\partial_1 = \mathrm{id} : k_a \to k_y \qquad \Rightarrow \qquad H_1(X;\hat{F}_1) = 0 \quad H_0(X;\hat{F}_1) = 0$$

The examples above lead themselves naturally to a complete classification of the cosheaf homology of barcodes over a stratified compact subset of the real line.

Claim 8.4. *Suppose X is a compact interval stratified into points x_i and intervals a_i. To each indecomposable sheaf or cosheaf on X we can associate one of the following intervals:*

$$[\text{—}] \qquad [\text{—}[\qquad]\text{—}] \qquad]\text{—}[$$

Moreover, the cosheaf homology of an indecomposable can be summarized as follows:

$$H_0^{BM}([\text{—}]) = k \qquad H_1^{BM}(]\text{—}[) = k \qquad H_i^{BM}([\text{—}[) = H_i^{BM}(]\text{—}]) = 0$$

All other Borel-Moore homology groups are zero. Moreover, since cosheaf homology commutes with finite direct sums, cellular cosheaf homology of $\hat{F}$ on X can be computed using the disjoint union of bar-codes $B_{\hat{F}}$ associated to the Remak decomposition of $\hat{F}$.

$$H_i(X;\hat{F}) \cong \oplus H_i^{BM}(B_{\hat{F}}) \qquad i = 0, 1$$

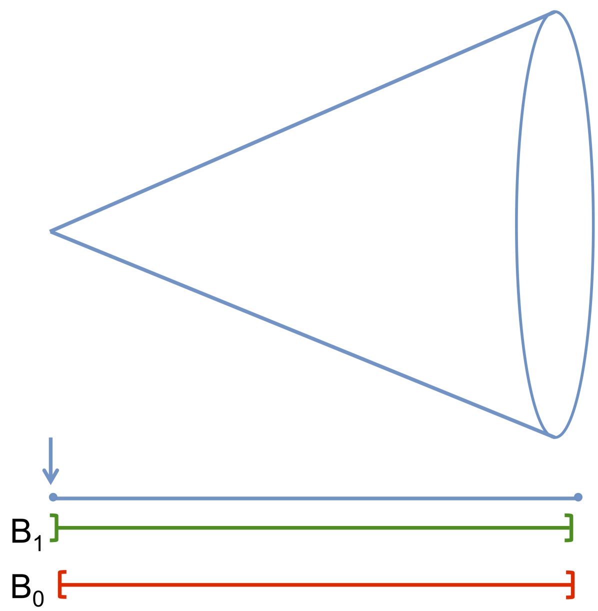

Figure 28: Barcodes for the Cone

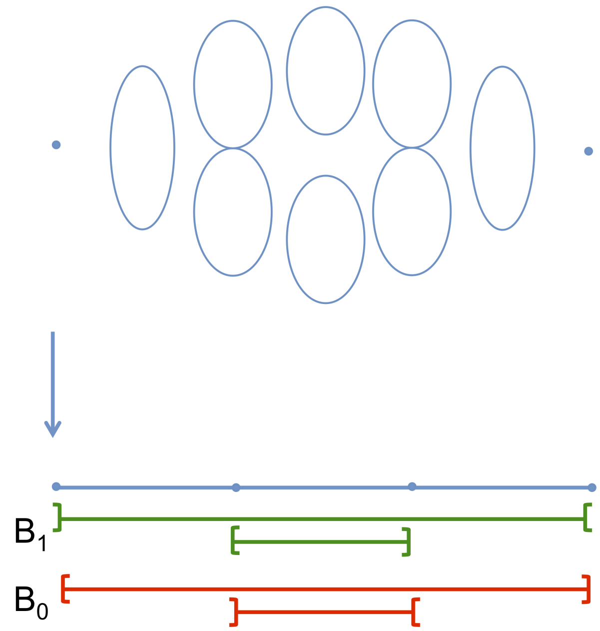

Figure 29: Barcodes for Bott's torus

Proof. The proof of the claim uses the previous computations and an invariance under subdivision argument, which is proved in theorem 11.27. The true sticking point is why ordinary cosheaf homology becomes Borel-Moore (cosheaf) homology. A more complete answer is developed in section **??**, but we sketch some of the argument now.

If we imagine that we are extending the constant cosheaf supported on a barcode to a cosheaf defined on all of $[0,1]$, then the natural way of extending by zero is to use the pushforward with open supports functor $j_\dagger$, where

$$j : B \hookrightarrow [0,1] \qquad \text{and} \qquad \hat{k}_B \mapsto j_\dagger \hat{k}_B.$$

The process of taking cosheaf homology is to then push forward this cosheaf to a point. However, using lemma 11.25 from section **??**, we see the following is true:

$$p_! \hat{k}_B \cong p_* j_\dagger \hat{k}_B \qquad H_i^{BM}(B; \hat{k}_B) \cong H_i([0,1]; j_\dagger \hat{k}_B)$$

When it is clear that we are working on $[0,1]$ we may write $\hat{k}_B$ instead of $j_\dagger \hat{k}_B$. □

Let's illustrate the utility of the claim by computing cosheaf homology over X using two different methods:

- Using the computational formulae of section 7.1
- Determining the barcode decomposition and applying claim 8.4.

Example 8.5 (Height function on the Torus)**.** The standard introductory example of Morse theory, first popularized by Raoul Bott, is the height function on the torus. In figure 29 we have drawn the behavior of the fibers over the critical values and the non-critical intervals. For the sake of brevity, let us write out only the cosheaf $\hat{F}_1$:

$$0 \longleftarrow k_a \longrightarrow k_y^2 \longleftarrow k_b^2 \longrightarrow k_z^2 \longleftarrow k_c \longrightarrow 0$$

Here the maps from k_a to k_y^2 and k_c to k_z^2 are the diagonal maps

$$r_{z,a} = \begin{bmatrix} 1 \\ 1 \end{bmatrix} = r_{z,c}$$

and the other maps are the identity. Choosing the orientation that points to the right, we get the follow matrix representation for the boundary map:

$$\partial_1 = \begin{bmatrix} 1 & -1 & 0 & 0 \\ 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & -1 \\ 0 & 0 & 1 & -1 \end{bmatrix} \qquad H_1(X; \hat{F}_1) =< \begin{bmatrix} 1 \\ 1 \\ 1 \\ 1 \end{bmatrix} > \qquad H_0(X; \hat{F}_1) \cong k$$

However, if we change our bases as follows

$$\begin{bmatrix} y_1' = y_1 \\ y_2' = y_1 + y_2 \end{bmatrix} \qquad \begin{bmatrix} b_1' = b_1 \\ b_2' = b_1 + b_2 \end{bmatrix} \qquad \begin{bmatrix} z_1' = z_1 \\ z_2' = z_1 + z_2 \end{bmatrix}$$

then our cosheaf $\hat{F}_1$ can then be written as the direct sum of two indecomposable cosheaves:

$$0 \longleftarrow 0 \longrightarrow k_{y_1'} \longleftarrow k_{b_1'} \longrightarrow k_{z_1'} \longleftarrow 0 \longrightarrow 0$$

$$0 \longleftarrow k_a \longrightarrow k_{y_2'} \longleftarrow k_{b_2'} \longrightarrow k_{z_2'} \longleftarrow k_c \longrightarrow 0$$

Hence it is apparent that

$$H_i(X; \hat{F}_1) \cong H_i^{BM}([\text{—}]) \oplus H_i^{BM}(]\text{—}[).$$

8.2 Barcode Homology and Multi-Dimensional Persistence

The astute reader may have noticed a pattern in the cosheaf homology computations of the previous section. Although no single cosheaf $\hat{F}_i$ captures the homology of the domain of a stratified map $f : Y \to X$, together the collection of cosheaves do.

$$\hat{F}_0 \qquad \hat{F}_1 \qquad \hat{F}_2 \qquad \cdots$$

In this section we point out a simple, visual way of computing the homology of Y using the barcodes of each of the $\hat{F}_i$. This result, first explicitly observed by Dan Burghalea and Tamal Dey in 2011 [19], was obtained independently by the author as a corollary of a more general theorem.

Corollary 8.6. *Assume* Y *is compact and* $\mathsf{f} : \mathsf{Y} \to \mathsf{X}$ *is a (Whitney) stratified map with* $\mathsf{X} \subset \mathbb{R}$*. For each* i*, let* $\mathsf{B_i}$ *denote the disjoint union of the barcode intervals associated to taking* i*th homology of the fiber. The following is true:*

$$\mathsf{H_i}(\mathsf{Y}; \mathsf{k}) \cong \mathsf{H_0}(\mathsf{X}; \hat{\mathrm{F}}_\mathsf{i}) \oplus \mathsf{H_1}(\mathsf{X}; \hat{\mathrm{F}}_{\mathsf{i}-1}) \cong \mathsf{H}_0^{\mathsf{BM}}(\mathsf{B_i}) \oplus \mathsf{H}_1^{\mathsf{BM}}(\mathsf{B}_{\mathsf{i}-1})$$

The reader should verify this corollary in the few examples already computed. The degree-shift should indicate that a spectral sequence is involved. This is indeed the case and it has a well-established history.

Leray defined a spectral sequence specifically adapted to understanding sheaf cohomology of a series of pushforward sheaves along an *arbitrary* continuous map $\mathsf{f} : \mathsf{Y} \to \mathsf{X}$. Serre, who wanted to use Leray's ideas for explicitly the case of singular homology, made the restriction to fiber bundles familiar to many mathematicians today. Thus, the reader might be surprised to see a Leray-Serre type spectral sequence in the stratified setting.

The explicit application of spectral sequences to stratified maps appears to have been pioneered[36] by Sir Christopher Zeeman. However, Zeeman's work occurred before most stratification theory really developed, thus the attribution is somewhat anachronistic. In "Dihomology II" [101], Zeeman expressed a prescient concern in finding a class of maps of finite type that are general enough to include simplicial maps, fiber bundles, ramified coverings, orbifolds $\mathsf{M} \to \mathsf{M}/\mathsf{G}$, projections to symmetric products, and many others. He called such maps "polyfibre maps" to express the idea that the fibers could change from point to point. By restricting to these maps, Zeeman obtained a class of spectral sequences that are finitely computable. Moreover, he framed his theory to work over cell complexes since one could use fewer cells than in a general triangulation.

Perhaps even more shocking is the fact that Zeeman worked explicitly with cellular cosheaves[37] to organize his computations. Zeeman proves that the cellular cosheaves $\hat{\mathrm{F}}_\mathsf{i}$ that arise from maps of this kind can be knit together in a spectral sequence to compute the homology of Y. However, some of his requirements are unnecessary. We state a more general result and a sketch of the proof.

Theorem 8.7. *Assume* Y *is compact and* $\mathsf{f} : \mathsf{Y} \to \mathsf{X}$ *is a (Whitney) stratified map such that* X *is stratified as a regular cell complex. There is a spectral sequence stitching together all the cellular cosheaf homologies of* $\hat{\mathrm{F}}_\mathsf{i}$*, which computes the total homology of* Y*:*

$$\mathsf{H_i}(\mathsf{X}; \hat{\mathrm{F}}_\mathsf{j}) \Rightarrow \mathsf{H}_{\mathsf{i}+\mathsf{j}}(\mathsf{Y}; \hat{\mathsf{k}}_\mathsf{Y}).$$

Remark 8.8. We have assumed compactness of Y so that we can use the usual formula for computing cosheaf homology on X without having to discard cells.

[36] Much to his dismay, the author only learned of Zeeman's work after developing independently a similar set of ideas.

[37] Zeeman used the term **stacks** - a term which now refers to something entirely different. Such a confusion of terminology made a literature search impossible.

Proof Sketch. We begin with a cover of X by the open stars of the vertices $\{U_x\}$. Now we observe that we have the following double complex involving the singular chain cosheaves defined in our preliminary examples section 4.

$$\begin{tikzcd}
\hat{C}_2(Y) \arrow[d] & \oplus_x \hat{C}_2(f^{-1}(U_x)) \arrow[l] \arrow[d] & \oplus_{x_1 \neq x_2} \hat{C}_2(f^{-1}(U_{x_1} \cap U_{x_2})) \arrow[l] \arrow[d] & \cdots \arrow[l] \\
\hat{C}_1(Y) \arrow[d] & \oplus_x \hat{C}_1(f^{-1}(U_x)) \arrow[l] \arrow[d] & \oplus_{x_1 \neq x_2} \hat{C}_1(f^{-1}(U_{x_1} \cap U_{x_2})) \arrow[l] \arrow[d] & \cdots \arrow[l] \\
\hat{C}_0(Y) & \oplus_x \hat{C}_0(f^{-1}(U_x)) \arrow[l] & \oplus_{x_1 \neq x_2} \hat{C}_0(f^{-1}(U_{x_1} \cap U_{x_2})) \arrow[l] & \cdots \arrow[l]
\end{tikzcd}$$

It is a general fact that the rows of of this double complex involving $\hat{C}_k$ are exact. This has to do with the fact that

$$0 \leftarrow \hat{C}_k(U \cup V) \leftarrow \hat{C}_k(U) \oplus \hat{C}_k(V) \leftarrow \hat{C}_k(U \cap V) \leftarrow 0$$

and consequently the Čech homology of $\hat{C}_k$ on any cover $\mathcal{U}$ of Y is zero except in degree zero where it is isomorphic with $\hat{C}_k(Y)$. The acyclic assembly lemma from homological algebra says that the homology of the double complex is isomorphic to the homology of the left most column, which is wanted.

Now we are going to apply the deformation retraction argument and additionally subdivide the cell structure on X to get a new cell complex X' so that there is a bijection between the nerve of the cover and the cells in X'. This allows us to refer to elements of the double intersection as edges and so on.

$$\begin{tikzcd}
\oplus_x \hat{C}_2(f^{-1}(x)) \arrow[d] & \oplus_{x_1 \neq x_2} \hat{C}_2(f^{-1}(e)) \arrow[l] \arrow[d] & \cdots \arrow[l] \\
\oplus_x \hat{C}_1(f^{-1}(x)) \arrow[d] & \oplus_{x_1 \neq x_2} \hat{C}_1(f^{-1}(e)) \arrow[l] \arrow[d] & \cdots \arrow[l] \\
\oplus_x \hat{C}_0(f^{-1}(x)) & \oplus_{x_1 \neq x_2} \hat{C}_0(f^{-1}(e)) \arrow[l] & \cdots \arrow[l]
\end{tikzcd}$$

Taking homology in the vertical direction produces the cosheaves of interest

$$\begin{tikzcd}
\oplus_x \hat{F}_2(x) & \oplus_e \hat{F}_2(e) \arrow[l] & \cdots \arrow[l] \\
\oplus_x \hat{F}_1(x) & \oplus_e \hat{F}_1(e) \arrow[l] & \cdots \arrow[l] \\
\oplus_x \hat{F}_0(x) & \oplus_e \hat{F}_0(e) \arrow[l] & \cdots \arrow[l]
\end{tikzcd}$$

which one can recognize is the preparatory step in taking cellular cosheaf homology. Taking homology horizontally produces the E^2 page in a spectral sequence, which collapses in the case that X is one-dimensional, otherwise we get another complex

$$\begin{tikzcd}
\mathsf{H}_0(\mathsf{X}';\hat{\mathrm{F}}_2) & \mathsf{H}_1(\mathsf{X}';\hat{\mathrm{F}}_2) & \mathsf{H}_2(\mathsf{X}';\hat{\mathrm{F}}_2) \\
\mathsf{H}_0(\mathsf{X}';\hat{\mathrm{F}}_1) & \mathsf{H}_1(\mathsf{X}';\hat{\mathrm{F}}_1) & \mathsf{H}_2(\mathsf{X}';\hat{\mathrm{F}}_1) \arrow[ull] \\
\mathsf{H}_0(\mathsf{X}';\hat{\mathrm{F}}_0) & \mathsf{H}_1(\mathsf{X}';\hat{\mathrm{F}}_0) & \mathsf{H}_2(\mathsf{X}';\hat{\mathrm{F}}_0) \arrow[ull]
\end{tikzcd}$$

and the standard spectral sequence argument converges to the total homology of Y. Technically, the cosheaf homology groups are over the new cell structure X', but an invariance under subdivision argument that will be presented later on shows that these groups agree with the ones computed over X. □

Corollary 8.6 indicates that barcodes give an alternative way of computing homology; although it can be more work to compute a barcode than to compute homology. The fact that theorem 8.7 works over spaces X of greater than one dimension is enticing: the cosheaf perspective carries over without modification. It is the belief of the author that cosheaves provide an alternative approach to multi-dimensional persistence, which has been of recent interest. We give a sketch of what a multi-dimensional barcode might look like.

Example 8.9 (Shadow of the Sphere)**.** Consider the standard Euclidean sphere S^2 embedded in $\mathbb{R}^3$. Let $\mathsf{f}:\mathsf{S}^2\to\mathbb{R}^2$ be the projection onto the first two factors of $\mathbb{R}^3$. The image $\mathsf{X}=\mathsf{f}(\mathsf{S}^2)$ is the closed unit disk. This map is definable in any o-minimal structure on $\mathbb{R}$ and hence stratified. Since the fiber over any point in the disk is zero-dimensional, we only get the cosheaf $\hat{\mathrm{F}}_0$, which can be viewed as a representation of the entrance path category of the natural stratification on X or as a cellular cosheaf in some stratification.

Applying the Krull-Schmidt theorem, we see that this representation splits as the direct sum of constant cosheaves supported on the open and closed disk respectively. The open disk has $\mathsf{H}_0^{\mathsf{BM}}=\mathsf{H}_1^{\mathsf{BM}}\cong 0$ and $\mathsf{H}_2^{\mathsf{BM}}\cong\mathsf{k}$. The closed disk only has $\mathsf{H}_0^{\mathsf{BM}}\cong\mathsf{k}$ as it is compact. Applying theorem 8.7, we see that the spectral sequence collapses immediately, so we can read off the homology of the two sphere from these generalized barcodes:

$$\mathsf{H}_0(\mathsf{S}^2)\cong\mathsf{H}_0(\mathsf{X};\hat{\mathrm{F}}_0)\cong\mathsf{k}\qquad \mathsf{H}_1(\mathsf{S}^2)\cong\mathsf{H}_1(\mathsf{X};\hat{\mathrm{F}}_0)\cong 0\qquad \mathsf{H}_2(\mathsf{S}^2)\cong\mathsf{H}_2(\mathsf{X};\hat{\mathrm{F}}_0)\cong\mathsf{k}$$

To date, the main approaches to multi-dimensional persistence include representations of quivers or standard poset structures on $\mathbb{R}^n$ [49]. One of the primary obstacles to any of these approaches is that there is no nice class of indecomposables to play the role of barcodes. For quivers, a strict classification of when a quiver has a finite or tame set of indecomposables is provided by Gabriel's theorem; see [25] for a nice survey article. Constructible sheaves and cosheaves, which are not quivers, but rather quivers with relations, suffer from similar pathologies. Nevertheless, we give a loose definition of a multi-dimensional barcode:

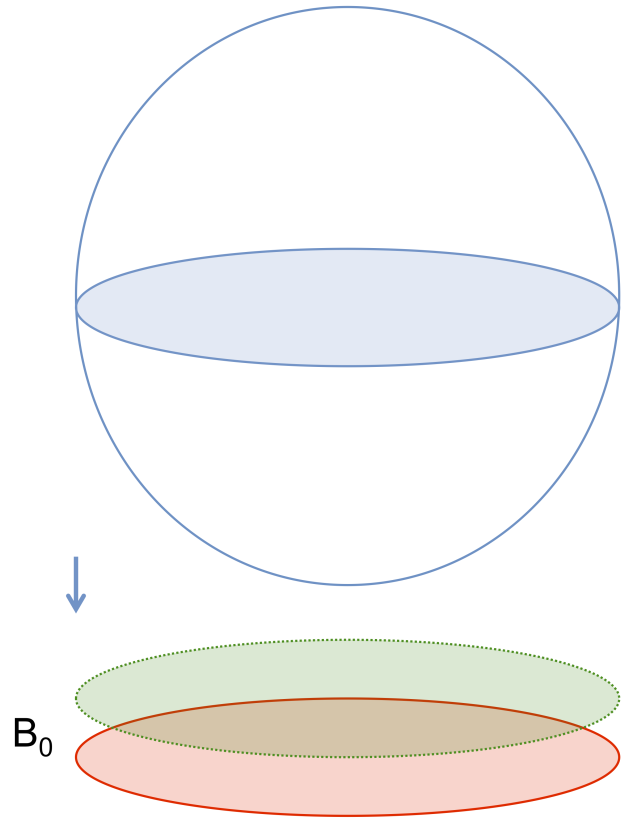

Figure 30: Two Dimensional Barcodes for the Sphere

Definition 8.10 (Generalized Barcodes)**.** A constant sheaf or cosheaf supported on a subspace of X will sometimes be referred to as a **generalized barcode**.

One should note that not every indecomposable sheaf or cosheaf defines a generalized barcode. For example, the cosheaf $\hat{\mathrm{F}}_0$ associated to the map

$$\mathsf{f} : \mathbb{C}^* \to \mathbb{C}^* \qquad\qquad z \mapsto z^2$$

is an indecomposable, but it is only locally constant and not constant. The presence of holonomy means that it cannot simply be "scanned" and have its homological type determined.

Even in the absence of barcodes, it must be stressed that sheaves and cosheaves are powerful data management tools that are endowed with a rich computational framework. We believe that they will play an increasingly important role in multi-dimensional persistence

9 Network Coding and Routing Sheaves

In the previous section, we introduced the language of barcodes and integrated them with a cosheaf-theoretic perspective on Morse theory and persistent homology. The fundamental idea was that by decomposing a cosheaf into indecomposables, we were able to understand cosheaf homology via the Borel-Moore homology of the barcode. In this section, we attempt to do the same thing for cellular sheaves on graphs. We apply the barcode perspective, wherever possible, to a class of sheaves introduced by Robert Ghrist and Yasuaki Hiraoka [36]. These sheaves were specifically designed to model the flow of information over graphs and the generalized barcode decomposition can aid in visualizing this flow.

First, we review some basic definitions for graphs.

Definition 9.1. Let X be a directed graph consists of a pair of sets E and V of edges and vertices and a pair of functions $h, t : E \to V$ that return the **head** and **tail** of an edge respectively. A directed edge goes from its tail to its head. The set of **incoming edges** to a vertex v, written $\mathrm{in}(v)$, is the set of edges whose heads are v, i.e. $h^{-1}(v)$. The set of **outgoing edges** at v is the set of edges whose tails are at v, i.e. $t^{-1}(v) = v$.

Definition 9.2. Let X be a directed graph with vertex set V and edge set E. A **capacity function** is a function c from the edge set to either the non-negative reals $\mathbb{R}_{\geqslant 0}$ or the non-negative integers $\mathbb{Z}_{\geqslant 0}$.

Definition 9.3 (Network Coding Cell Sheaf)**.** Suppose X is a directed graph with a capacity function c. A **network coding sheaf** on X is a cellular sheaf $F : X \to \mathbf{Vect}$ constructed as follows:

- To an edge $e \in X$ we let $F(e) = k^{c(e)}$, a vector space of dimension equal to the capacity.
- To a vertex v we let $F(v) = k^{c(v)} \cong \oplus_{e_i \in \mathrm{in}(v)} k^{c(e)}$.
- The restriction maps are given by ordinary projections for the incoming edges, i.e. $\rho_{e_i,v} := \mathrm{proj}_{F(e_i)}$ for $e_i \in \mathrm{in}(v)$, but for the outgoing edges some non-trivial coding may be performed, i.e. any linear map $\Phi_{e_k,v} : F(v) \to F(e_k)$ for $e_k \in \mathrm{out}(v)$ will do. We write $\Phi_v = \oplus_{e_k \in \mathrm{out}(v)} \Phi_{v,e_k}$ for the **total coding** through v.

Remark 9.4. It should be noted that in [36], they do not use cellular sheaves. This was primarily due to the lack of a good reference.

In [36] they do not define network coding sheaves for arbitrary directed graphs. Instead, they consider a graph with a distinguished set of sources and targets (senders and receivers) and they augment the graph by adding **decoding wires** directed to go from a target vertex back to a subset of source vertices. Heuristically for Ghrist and Hiraoka, the purpose of these edges is to make global sections of this sheaf correspond to closed loops through the graph. This topological reasoning is correct, but oversimplifies how network codings can produce counterintuitive weavings and splittings of data.

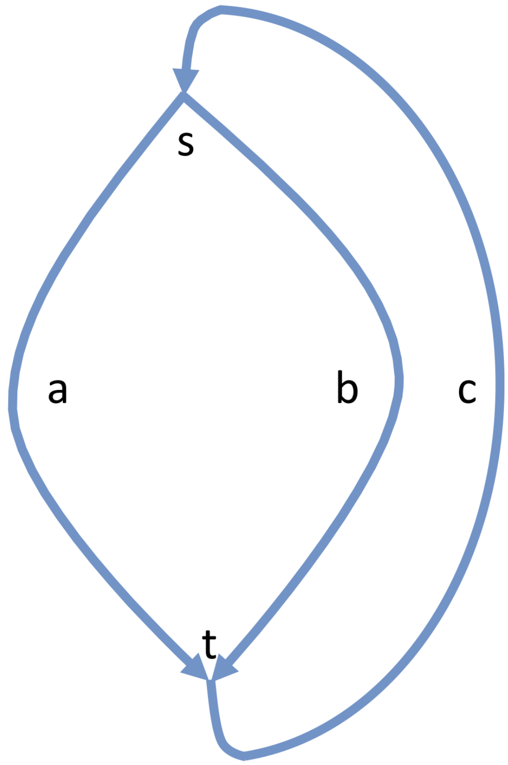

Figure 31: Graph Augmented with Decoding Wire

Example 9.5. Consider the graph in figure 47 with constant capacity function $\mathsf{c} = 1$. Consequently, all edges and vertices get a one dimensional vector space $\mathsf{k} = \mathbb{R}$ with the exception of $\mathsf{F}(\mathsf{t}) \cong \mathsf{k}^2$. Define the coding maps $\rho_{\mathsf{a},\mathsf{s}} = \mathrm{id} = \rho_{\mathsf{b},\mathsf{s}}$ and

$$\rho_{\mathsf{c},\mathsf{t}} = \begin{bmatrix} \frac{1}{2} & \frac{1}{2} \end{bmatrix}.$$

We pick a local orientation implied by the source and target vertices. The one and only coboundary matrix can be written as follows:

$$\delta^0 := \begin{bmatrix} -1 & 1 & 0 \\ -1 & 0 & 1 \\ 1 & -1/2 & -1/2 \end{bmatrix}$$

Consequently, $\mathsf{H}^0(\mathsf{X};\mathsf{F}) \cong \mathsf{H}^1(\mathsf{X};\mathsf{F}) \cong \mathsf{k}$. The one global section is supported over the entire graph; it is not simply a loop through the graph.

The previous example of a network coding sheaf is an example of an indecomposable sheaf that is not a generalized barcode in the sense of definition 8.10. To better understand the flow of data over graphs, as well as the utility of the barcode perspective, we consider a simpler class of network coding sheaves.

9.1 Duality and Routing Sheaves

Definition 9.6 (Routing Sheaf)**.** A particular type of network coding sheaf is a **routing sheaf**. Here we assume that the capacity function is constant $c = 1$, and the coding maps Φ_v can be written as a binary matrix with at most one 1 in each column and row. Said another way, at a vertex v the total coding map maps to zero as many incoming edges as desired, so long as there is a bijection of the remaining incoming edges and a subset of the outgoing edges. The total coding map through v is then a matrix representation of this bijection.

The advantage of routing sheaves is that they are simple to visualize: Start at a source and use a color pen to track how an edge emanating from that source gets bounced around under the routing directions at each subsequent vertex. If at any point in your drawing you run into a vertex that sends your edge's data to zero, stop on that vertex with your pen. This argument essentially establishes the following proposition.

Proposition 9.7 (Structure Theorem for Routing Sheaves)**.** *Suppose* $X = (V, E, h, t)$ *is a directed graph, then every routing sheaf* $F : X \to \mathbf{Vect}$ *can be realized as the pushforward with compact support of a disjoint union of half-open intervals* [—[*or circles, whose images can intersect only at vertices.*

A consequence of this result combined with Poincaré duality is the following corollary.

Corollary 9.8 (Duality)**.** *For any routing sheaf* F *one has*

$$H^0(X; F) \cong H^1(X; F).$$

Proof. By proposition 9.7, every network coding sheaf can be written as a direct sum of constant sheaves supported on half-open intervals or circles. Half-open intervals embedded into compact spaces (extending by zero using $j_!$) have trivial cohomology in both degrees. Poincaré duality for S^1 establishes the corollary. □

One can also prove this duality in the more general setting of network coding sheaves via a simple combinatorial argument.

Proposition 9.9. *For an any network coding sheaf* F*, we have the following isomorphisms:*

$$\bigoplus_v F(v) \cong \bigoplus_e F(e) \qquad H^0(X; F) \cong H^1(X; F)$$

Proof. By construction of a network coding sheaf there is a bijection between the sum of the vector spaces over the edges $e \in in(v)$ and the vector space over the vertex v.

$$\bigoplus_{e \in in(v)} F(e) = F(v).$$

By definition of a graph, every edge is the incoming edge for a unique vertex. Thus, by summing over all vertices, we sum over all edges without double-counting. This proves the first isomorphism. The second isomorphism follows by the rank-nullity theorem. □

Ideally, one could interpret these cohomology groups as something meaningful to obtain a useful duality result, but this is still missing. Ghrist and Hiraoka interpret H^0 as a vector space spanned by independent information flows, but in the case of routing sheaves, H^1 is the group that counts closed trajectories of information flow. For routing sheaves, one could say that the Poincaré dual of the fundamental class of an information loop would yield a point whose removal would cease the flow of information. One might call this a "cut equals flow" theorem. This is only a pale shade of the greater "Max-Cut Min-Flow" theorem [29, 31, 79], which compares the maximum possible flow with the minimum capacity cut required to disconnect the graph.

9.2 Counting Paths Cohomologically, or Failures thereof

Regardless of network coding sheaves connections with duality, one would like to know what information cellular sheaf cohomology can capture over a graph. Is it possible, for example, to build a sheaf that encodes source-to-target paths cohomologically? Suppose we allow the source to have its own independent capacity, without regard for the number of incoming edges. As the next example shows such a "pseudo network coding" (NC) sheaf cannot encode source-to-target paths cohomologically.

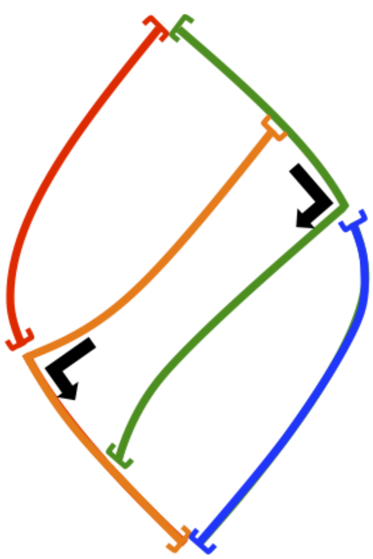

Figure 32: Pseudo-NC Sheaf with No Decoding Edge

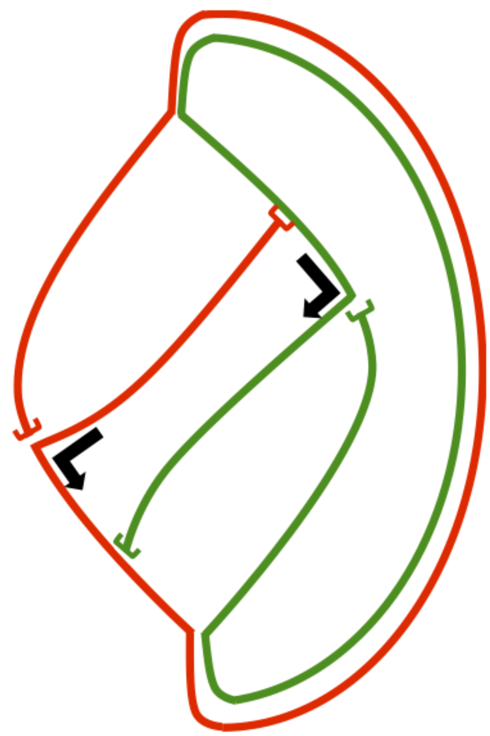

Figure 33: Network Coding Sheaf with Decoding Edge

Example 9.10 (Decoding Edge and Barcodes)**.** In figures 32 and 33, the barcode decomposition of a network coding sheaf is drawn with and without a decoding edge. With the particular choices made there is no flow from source to target. Yet the sheaf in figure 32 decomposes as the constant sheaf on two half open intervals and two closed intervals:

$$F_{no} \cong (j_o)_! k_{[0,1)} \oplus (j_b)_! k_{[0,1)} \oplus (i_r)_* k_{[0,1]} \oplus (i_g)_* k_{[0,1]} \Rightarrow H^0(X; F_{no}) \cong k^2.$$

This is bad if we want our sheaf to encode cohomologically the presence of source-to-target information paths.

However, with the use of a decoding edge (decoding edges) the network decomposes into only two half open intervals:

$$F_{de} \cong (j_r)_! k_{[0,1)} \oplus (j_g)_! S k_{[0,1)} \Rightarrow H^0(X; F_{de}) \cong 0$$

which was wanted.

10 Sheaves and Cosheaves in Sensor Networks

In this section, we consider a candidate application of sheaves and cosheaves to problems in sensor networks. Section 10.1 outlines some real-world sensors as well as their mathematical abstraction. With this abstraction in hand, we consider in section 10.2 the classic problem of determining when a sensor network has completely covered a region. The introduction of time-dependent sensor networks necessitates the sheaf-theoretic approach, despite the fact that it is unwieldy in its most general form.

In section 10.3 we attempt to "linearize" the sheaves and cosheaves used in studying sensor networks in the hope that sheaf cohomology and cosheaf homology will give us an obstruction-theoretic approach to sensing. An approach of Henry Adams is considered in section 10.3.1, as well as his counter-example to that approach. By using cosheaf-theoretic reasoning, we give a principled explanation for why this approach fails in proposition 10.7. An approach of the author and Robert Ghrist is then considered in 10.3.2. This approach succeeds where the previous approach fails, but it too suffers from giving false positives, as the example in proposition 10.14 shows. The example constructed there, which is joint with David Lipsky, uses one of the 12 indecomposable representations of the Dynkin diagram D_4.

Finally, a linear model for multi-modal sensing is presented in section 10.4. It was there that the author realized the necessity of using indecomposables to interpret sheaf cohomology computations. A delightful examination of the act of sensing in section 10.4.1 shows how sheaves and cosheaves work in tandem. Theorem 10.19 uses a long exact sequence in sheaf cohomology to obtain a forcing result in multi-modal sensing. Finally, the role of higher-dimensional barcodes in multi-modal sensing is considered in section 10.4.2.

10.1 A Brief Introduction to Sensors

Sensors are devices with delimited purview. They can measure certain properties and interact with occupants of a particular part of space-time. Examples abound in our world and they operate via differing modalities. Here are a few examples:

Example 10.1 (Sight)**.** Our eyes are highly tuned sensors that can detect photons with certain frequencies (visible light and colors) and their spatial range can be on the order of kilometers. Some man-made satellites orbiting the Earth have cameras with a greater spatial resolution and frequency response – they help us navigate by providing detailed pictures of roads, weather and climate. Eyes and satellites have a large scale and are very expensive. Cheaper sensors which can read only very coarse changes in light levels are found in our traffic lights, door ways and bathrooms.

Example 10.2 (Weight and Pressure)**.** Buried in roads or placed under door mats are sensors designed to respond to pressure. These open doors or gates or initiate changes in traffic signals. Some are more passive and merely collect data. A cable as thick as a thumb can be laid across a road and will record when something heavy (like a car) drives over it. Two spikes in pressure close in time indicate when a car's front and back wheels respectively drove over the cable. From this city officials can measure how fast cars are going as well as density and total volume of traffic.

Example 10.3 (Radio Frequency ID)**.** Some readers probably have a university card, or building card, that grants them access through locked doors merely by tapping on a sensor. Commuters drive cars equipped with sensors that allow them to pass through tolls without stopping. Some scientists tag animals to study a species' habits and movements. In all these cases, the sensor or the tag emits an electromagnetic field with limited spatial range (a few centimeters, meters, or kilometers) and only when inside this range is a tuned circuit thereby completed, connecting the sensors (card reader, toll booth, etc.) with the things being sensed (ID card, tag, etc.).

Although the physical mechanisms that allow each of these sensors to sense is different, there are some broad commonalities: spatially localized sensors return data in the presence of certain occupants, which we call **intruders**.

10.2 The Coverage Problem: Static and Mobile

The way we model sensors is to first identify the physical domain where the sensing is taking place – a two-dimensional Euclidean plane could represent the floor of a building – and we represent the sensors spatially via their support – a door mat with pressure sensors would be a rectangle in the plane. Or we could think of the field of view of a camera in a ceiling pointed directly down as a disk in the plane.

For the moment we ignore the type of data a sensor reports (we'll take that up later when we work with sheaves and cosheaves) and instead we consider the **coverage problem**: Given a collection of sensors distributed in a physical domain D, can we monitor the entire region without gaps?

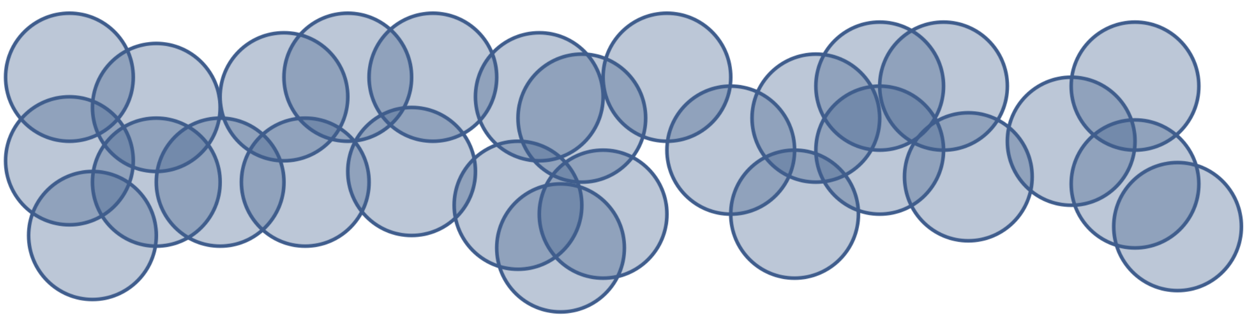

Figure 34: Sensors Distributed in a Plane

If we have good knowledge of the sensors which live on the boundary of our region, then we can, following the work of Vin de Silva and Robert Ghrist [24], give a certificate of coverage using relative homology. However, we frame this question using sheaves of sets instead, so as to better handle the time-dependent scenario.[38]

[38] The author would like to thank Gunnar Carlsson and Rob Ghrist for their insights here.

Definition 10.4. Let D be a spatial region of interest and denote by $\mathsf{D} \times [0,1]$ a region of space-time. This carries with it a map that keeps track of time via projection onto the second factor, i.e. $\pi_2 : \mathsf{D} \times [0,1] \to [0,1]$. We assume that D can be given a cell structure so that the sensors' **coverage region** $\mathsf{S} \subset \mathsf{D} \times [0,1]$ and the **evasion region** $\mathsf{E} := \mathsf{S}^c$ can be written as the union of cells. To study the **intruder problem** is to analyze the associated sheaf of sections of the map $\pi := \pi_2|_{\mathsf{E}} : \mathsf{E} \to [0,1]$, which we assume can be made cellular. Saying that there is an **evasion path** is to say there is a global section of this map, i.e. a $\mathsf{s} : [0,1] \to \mathsf{E}$ such that $\pi \circ \mathsf{s} = \mathrm{id}$.

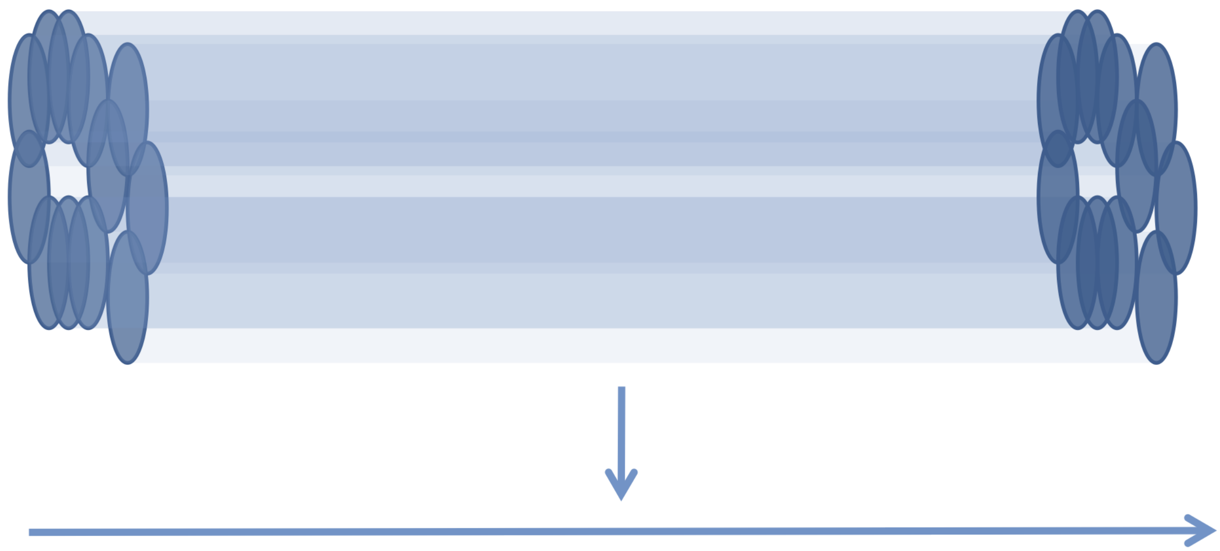

Figure 35: The Space-Time Perspective

Example 10.5. For the situation depicted in figure 35, the intruder problem has a clear answer. An intruder can evade detection by residing in either one of the two holes present. Picking a point and then resting there for all time determines a global section of the time projection map.

It should be clear that our sheaf-theoretic question is equivalent to a much simpler one: "Is the complement of the sensed region (the uncovered region) in D non-empty?" Thinking in terms of sheaves, at this point, buys us nothing.

Where sheaves begin to offer a hint of leverage is in the time-dependent scenario. Here we imagine the sensors can move around in our domain D. Now it is possible that the sensed region S does not look like a product of space and time.

Example 10.6. In figure 36 we imagine that there is a one-dimensional environment of interest that sits vertically over each point on the time axis. Between the black lines is a region that is currently being unmonitored. To begin there is only one connected component of the unmonitored region. As time marches forward to the right a second connected component

Figure 36: Mobile Sensor Network

of the unmonitored region opens up, followed shortly by a third. Two of these three merge and then disappear leaving only one component of unmonitored territory.

In this case the non-existence of an evasion path is clear: no intruder could have gone undetected without time-traveling. This corresponds to the ready-seen fact that this map has no section, i.e. there does not exist a map $s : [0, 1] \to E$ such that $\pi \circ s = \mathrm{id}$.

What is the purpose of considering sheaves at all? If we can stare at the drawing and detect whether a section exists or not, why bother with high-flown machinery? However, what is easily seen in toy examples, can quickly become unmanageable. The only mathematics that formalizes intuition about sections is sheaf theory and moreover, once formalized using cellular sheaves, it can be programmed on a computer.

However, there is a disadvantage with using sheaves of sets. We'd like to be able to calculate an obstruction that would certify whether a global section exists or not. One of the stated purposes of using sheaf cohomology is to provide such a calculable obstruction. Unfortunately, cohomology requires the linear structure of vector spaces, which we do not have here. In the next section we consider what happens when we naively "linearize" the sheaf of sections of a map.

10.3 Intruders and Barcodes

In this section, we use cellular sheaves and cosheaves to analyze the intruder problem in the time-dependent case. We assume that the time projection map π is cellular in order to take advantage of the functors in section 6. By putting a sheaf or cosheaf on the evasion region and pushing forward along π, we reduce the intruder problem to one dimension where we can use the barcode perspective of section 8. There are two main approaches, both of which have their drawbacks:

- One approach is to study the homology of the evasion region at each moment in time $\pi^{-1}(t)$. By theorem 5.72, this determines a cellular cosheaf.

- The second approach is to linearize the space of sections of the map π. To make the space of sections finite, we pass to the Reeb graph of the evasion region. This determines a cellular sheaf and stays true to the original intruder problem.

10.3.1 Tracking the Topology over Time

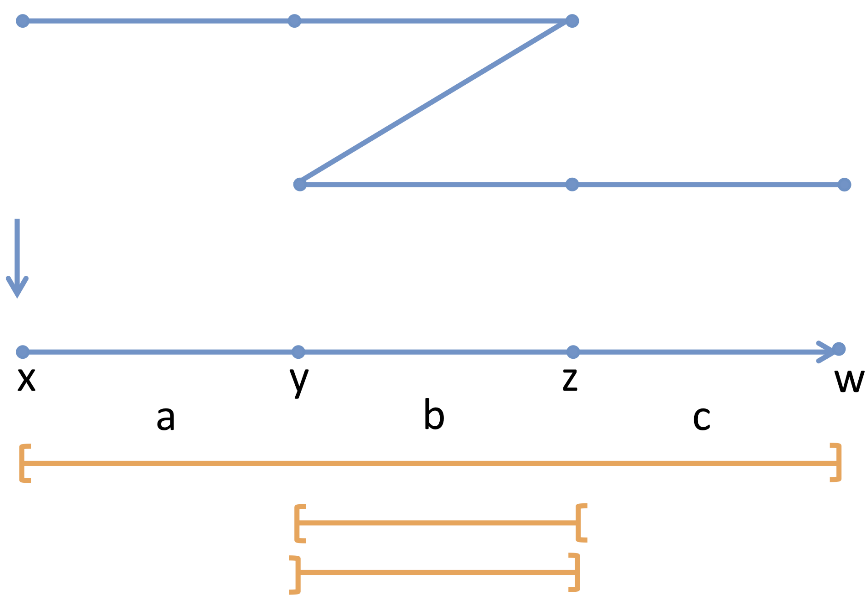

Figure 37: Mobile Sensor Network

To simplify the topology, we focus on the Reeb graph version of figure 36. This is drawn and labelled in figure 37. Since everything is occurring in two-dimensional space-time, the only interesting homological invariant of the fiber is H_0. Studying this is equivalent to studying the pushforward cosheaf $\hat{F} := \pi_* \hat{k}_E$. In the parlance of [20], this is simply a zigzag module of the following form:

$$\hat{F}(x) \xleftarrow{r_{x,a}} \hat{F}(a) \xrightarrow{r_{y,a}} \hat{F}(y) \xleftarrow{r_{y,b}} \hat{F}(b) \xrightarrow{r_{z,b}} \hat{F}(z) \xleftarrow{r_{z,c}} \hat{F}(c) \xrightarrow{r_{w,c}} \hat{F}(w)$$

$$k_x \longleftarrow k_a \longrightarrow k_y^2 \longleftarrow k_b^3 \longrightarrow k_z^2 \longleftarrow k_c \longrightarrow k_w$$

If we choose for each cell in $[0,1]$ the ordered basis given by the top down ordering on the page of the cells in the fiber we get the following matrix representations of the extension maps:

$$r_{y,a} = \begin{bmatrix} 1 \\ 0 \end{bmatrix} \quad r_{y,b} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 1 \end{bmatrix} \quad r_{z,b} = \begin{bmatrix} 1 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \quad r_{x,c} = \begin{bmatrix} 0 \\ 1 \end{bmatrix}$$

We can decompose this cosheaf into indecomposables simply by performing the correct change of basis:

$$\begin{bmatrix} y_1' = y_1 \\ y_2' = y_1 - y_2 \end{bmatrix} \qquad \begin{bmatrix} b_1' = b_1 - b_2 + b_3 \\ b_2' = b_1 - b_2 \\ b_3' = b_2 - b_3 \end{bmatrix} \qquad \begin{bmatrix} z_1' = z_2 \\ z_2' = z_1 - z_2 \end{bmatrix}$$

The reader should match the resulting indecomposables with the barcodes drawn in figure 37.

$$\begin{tikzcd} k_x & k_a \arrow[l] \arrow[r] & k_{y_1'} & k_{b_1'} \arrow[l] \arrow[r] & k_{z_1'} & k_c \arrow[l] \arrow[r] & k_w \end{tikzcd}$$

$$\begin{tikzcd} 0 & 0 \arrow[l] \arrow[r] & k_{y_2'} & k_{b_2'} \arrow[l] \arrow[r] & 0 & 0 \arrow[l] \arrow[r] & 0 \end{tikzcd}$$

$$\begin{tikzcd} 0 & 0 \arrow[l] \arrow[r] & 0 & k_{b_3'} \arrow[l] \arrow[r] & k_{z_2'} & 0 \arrow[l] \arrow[r] & 0 \end{tikzcd}$$

The presence of a long barcode may seem surprising. It indicates that there is a connected component of the evasion region that persists for all time. The following proposition explains why this long barcode must exist.

Proposition 10.7. *Suppose* $E \subset D \times [0,1]$ *is a compact connected evasion region such that* $\pi = \pi_2|E$ *is surjective, i.e. there is at each point in time somewhere an intruder can evade detection, then the Remak decomposition of* $\pi_* \hat{k}_E$ *must have a barcode that stretches the length of* $[0,1]$.

Proof. The proof starts with the easy observation that if $f : Y \to X$ is a continuous map and $\hat{G}$ is a cosheaf on Y, then we have that $H_0(Y; \hat{G}) \cong H_0(X; f_* \hat{G})$. This follows from the commutativity of the following diagrams and functoriality of pushforward.

$$\begin{tikzcd} Y \arrow[rr, "f"] \arrow[dr, "p"'] & & X \arrow[dl, "p"] \\ & \star & \end{tikzcd} \qquad \begin{tikzcd} \hat{G} \arrow[rr, "f_*"] \arrow[dr, "p_*"'] & & f_* \hat{G} \arrow[dl, "p_*"] \\ & p_* \hat{G} \cong (p \circ f)_* \hat{G} & \end{tikzcd}$$

Setting $Y = E$, $X = [0,1]$, $f = \pi$ and $\hat{G} = \hat{k}_E$, we can use the fact that E is connected to get that $p_* \hat{k}_E \cong H_0(E; \hat{k}_E) \cong k$. We know that any (co)sheaf over $[0,1]$ can be written as a direct sum of constant (co)sheaves supported on barcodes.

$$\pi_* \hat{k}_E \cong \hat{k}_{B_1} \oplus \cdots \oplus \hat{k}_{B_n} \qquad \text{and}$$

Now we combine this with the fact that homology commutes with direct sums.

$$k \cong H_0(E; \hat{k}_E) \cong H_0([0,1]; \pi_* \hat{k}_E) \cong \bigoplus_i H_0([0,1]; \hat{k}_{B_i}) \cong \bigoplus_i H_0^{BM}(B_i).$$

Consequently, there can be only one closed barcode. We argue that this unique closed barcode must have support on all of $[0,1]$. Since we know that the constant section $1_{\mathsf{E}} \in \Gamma(\mathsf{E}; \hat{\mathsf{k}}_{\mathsf{X}})$ has support on all of E, the pushforward section $\pi_* 1_{\mathsf{E}}$ that generates the closed barcode must have support on all of $[0,1]$, since π is surjective. $\square$

Remark 10.8. We have implicitly used sheaf-theoretic reasoning with H^0 taking the place of H_0. The argument about the support of the section is better expressed using stalks.

As a consequence, we obtain a negative result, which is almost identical to a result of Henry Adams.

Corollary 10.9. *Having a barcode associated to* $\pi_* \hat{\mathsf{k}}_{\mathsf{E}}$ *whose support is all of* $[0,1]$ *does not indicate the existence of an evasion path.*

Remark 10.10. The above proof gives a cosheaf-theoretic explanation of why we shouldn't expect barcodes to detect the existence of an evasion path. Homology of the evasion region is not sensitive to its embedding, thus a long barcode will appear even if it is embedded in a way that would require an intruder to time travel. In this sense, Corollary 8.6 can be interpreted as a stability result: although half-open barcodes can pop in and out of existence, based on the embedding, there must always be one and only one closed barcode.

10.3.2 Linearizing the Sheaf of Sections

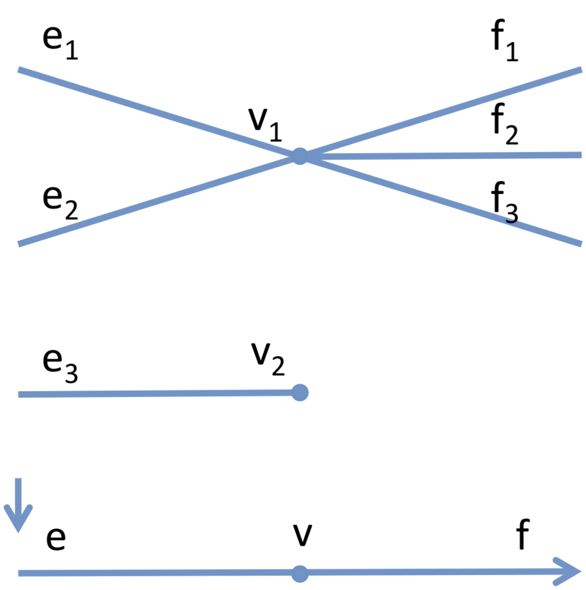

Figure 38: Sheaf of All Possible Evasion Paths

In light of the inability of the pushforward cosheaf $\pi_* \hat{\mathsf{k}}_{\mathsf{E}}$ to distinguish when an evasion path exists or not, we return to the original sheaf-theoretic formulation of the intruder problem. To make the sheaf of sections finite enough to work with, we take the Reeb graph of the map $\pi : E \to [0,1]$. From this setup, we can extract a cellular map of 1-dimensional cell complexes, normally called $\tilde{\pi} : R(\pi) \to [0,1]$, but we will abuse notation and assume that our input $\pi : E \to [0,1]$ is already a Reeb graph.

By picking a directionality of $[0,1]$ we can endow E with the structure of a directed graph. On this directed graph we can define the following cellular sheaf, which is meant to pushforward to a linear model of the sheaf of sections. It is very closely related[39] to the network coding sheaves defined in section 9.

Definition 10.11. Let X be an acyclic directed graph. We define a cellular sheaf G that assigns to an edge the one-dimensional vector space k and assigns to a vertex the space freely generated by all possible directed routings through that vertex.

We allow special treatment to a subset of sources S and sinks T, where we allow $G(v) = k^{\mathrm{out}(v)}$ for $v \in S$ and $G(v) = k^{\mathrm{in}(v)}$ for $v \in T$. All other sources and sinks get the zero vector space. The restriction mappings send a routing to the edges that participate in that routing.

Example 10.12. For a concrete example, where we focus on a small part of a graph, consider the graph in figure 38. The definition of the sheaf G makes

$$G(v_1) = < e_i \otimes f_j \,|\, i = 1,2;\ j = 1,2,3 > \cong k^6 \qquad \rho_{e_i,v_1}(e_j \otimes f_k) = \delta_{ij} \quad \rho_{f_j,v_1}(e_i \otimes f_k) = \delta_{jk}$$

and we choose to make $G(v_2) = 0$. The reason we have decided to set $G(v_2) = 0$ comes from the extra information of the projection map to $[0,1]$. We call such a vertex an **internal** source or sink. In the context of the intruder problem, an internal source represents an impossible entry point for an intruder. If we push the sheaf G along the projection map π we then get the following assignments of data:

$$\pi_* G(e) \cong < e_1, e_2, e_3 > \quad \pi_* G \cong G(v_1) \quad \pi_* G(f) = < f_1, f_2, f_3 >$$

Example 10.13. Let us consider the example drawn in figure 39, but now with the sheaf just defined. We set $F = \pi_* G$, whose values are below:

$$F(x) \xrightarrow{\rho_{a,x}} F(a) \xleftarrow{\rho_{a,y}} F(y) \xrightarrow{\rho_{b,y}} F(b) \xleftarrow{\rho_{b,z}} F(z) \xrightarrow{\rho_{c,z}} F(c) \xleftarrow{\rho_{c,w}} F(w)$$

$$k_x \longrightarrow k_a \longleftarrow k_y \longrightarrow k_b^3 \longleftarrow k_z \longrightarrow k_c \longleftarrow k_w$$

The two restriction maps of any interest include into the top section and the bottom section, respectively.

$$\rho_{b,y} = \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} \qquad \rho_{b,z} = \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}$$

Without a change of basis one can see that this sheaf splits as the direct sum of indecomposables, whose barcodes are drawn in figure 39.

[39] In a sense, the sheaf defined here gives all possible codings. It approximates a "stack" of network coding sheaves.

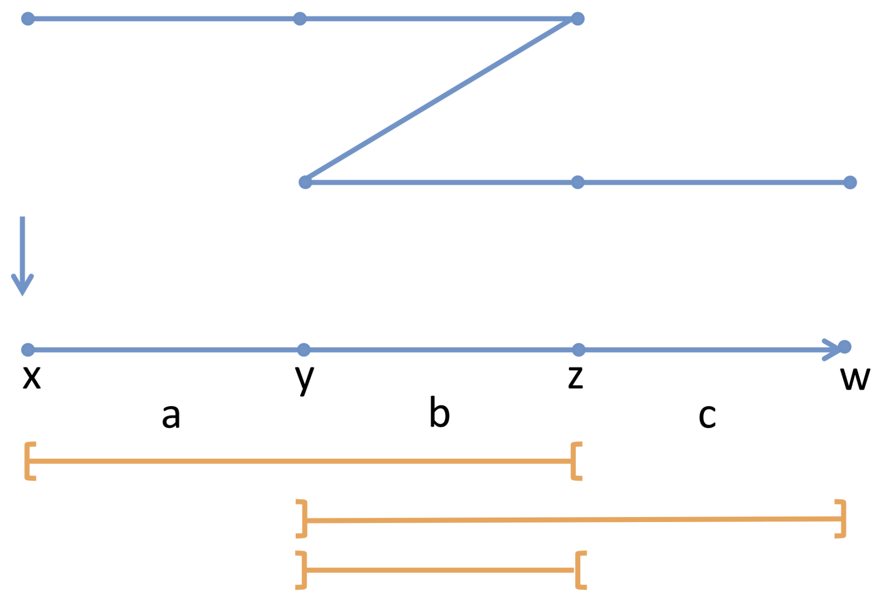

Figure 39: Linearized Sheaf of Sections

The previous example offers a glimmer of hope. No intruder can evade detection and the absence of a long barcode reflects that. Moreover, the sheaf cohomology computation shows $H^0([0,1];F) \cong 0$, which would be a promising shortcut to computing barcodes. Alas, the linearized sheaf of sections fairs no better than the cosheaf of components. Here we provide a counterexample, joint with Dave Lipsky, to either of the hopes that non-zero $H^0([0,1];F)$ or a long barcode provides an if and only if criterion for the existence of an evasion path.

Proposition 10.14. *Although it is true that the existence of an evasion path implies the existence of a long barcode (and thus $H^0([0,1];F) \neq 0$) it is not true that having a long barcode (or $H^0([0,1];F) \neq 0$, which is a weaker condition) implies the existence of an evasion path.*

Proof. In figure 40 we have drawn the counter-example, which we now explain. The component coming into p appears immediately after time 0, so it is impossible for an intruder to enter there. Similarly, there is a component leaving from q that closes up right before time 1. The pushforward sheaf then takes the following form

$$k_x \longrightarrow k_a^2 \longleftarrow k_y^3 \longrightarrow k_b^3 \longleftarrow k_z^3 \longrightarrow k_c^2 \longleftarrow k_w$$

The maps from $F(y)$ and $F(z)$ to $F(b)$ are the identity maps. The two maps that require some inspection are built out of a projection and a trace map.

$$\rho_{a,y} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 1 \end{bmatrix} \qquad \rho_{c,z} = \begin{bmatrix} 1 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}.$$

The change of basis required to obtain the desired Remak decomposition indicated by the barcodes is not so easily seen. The interval decomposition algorithm outlined in [20] provides

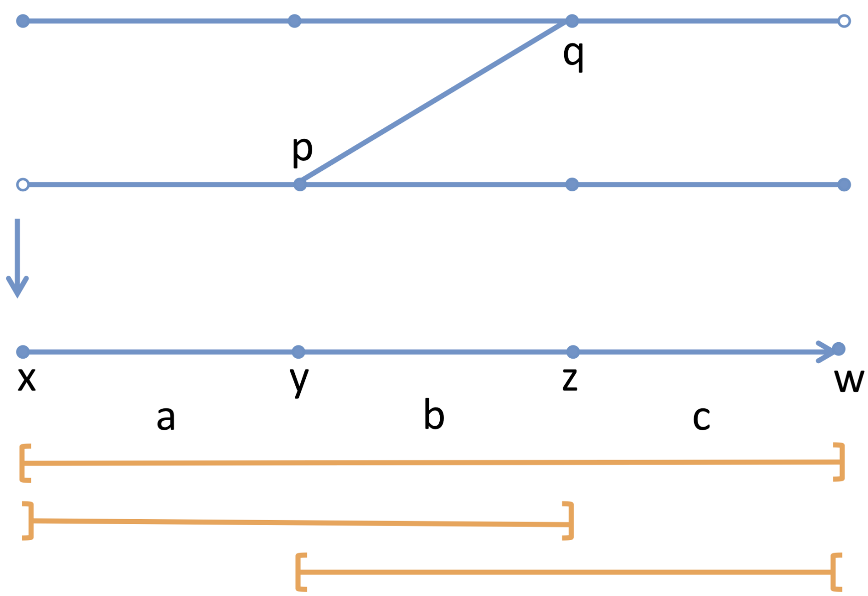

Figure 40: Counterexample for the Linearized Sheaf of Sections

a sure-fire method of obtaining it. It is left to the reader to verify the the barcodes in figure 40.

Instead, we give a sheaf-theoretic justification for the existence of a long barcode. There is a unique non-zero global section of G and it is supported everywhere except on the prong incoming to p and outgoing from q. Explicitly, it comes from choosing a compatible kernel for the restriction matrices $\rho_{a,y}$ and $\rho_{c,z}$. At q the routing through the top path is annihilated by the "negative" of the routing through the bottom path; it is as if two intruders are traveling with opposite charges. As a consequence, its support surjects onto all of $[0,1]$. Since $\mathsf{H}^0([0,1];\mathsf{F}) \cong \mathsf{k}$ we can infer the existence of one closed barcode, and because this section has global support, the barcode must be long. □

Remark 10.15 (Dynkin Diagrams and Stalks). Recall that $\mathsf{F} := \pi_* \mathsf{G}$. Consider the sheaf G implied by figure 40. When restricted to the open stars at p and q separately G is equivalent to one of the 12 indecomposable representations of the Dynkin diagram D_4; see [30], p. 83. Since the open stars intersect, one can show that the entire sheaf G on E is indecomposable. This cannot be used directly to show that a long barcode must exist. The pushforward of an indecomposable representation is not necessarily indecomposable. However, the argument using stalks indicates that some sections (subrepresentations), must have global support.

10.4 Multi-Modal Sensing

In this section we will explore the following cartoon for multi-modal sensing:

- We have a region W thought of as a topological space that is tame enough to be triangulated. This space is populated by agents of interest and sensors.

- There is a vector space of properties k^n, usually $\mathbb{R}^n$ or $\mathbb{C}^n$, and every intruder is tagged with an unchanging **property vector** $v \in k^n$. These property vectors might record colors (which we pretend has a linear structure), sounds, thermal signatures or, in the context of wireless network data, a unique wireless SSID (we imagine scaling corresponds to the strength of the signal). In future applications, k may be a ring that stores data, just as $\mathbb{Z}$ is used to record counts and $\mathbb{Z}^n$ records counts of different types of targets.

- There are sensors who monitor subspaces of k^n and subspaces of X. "Monitors" means explicitly that a sensor i with support $V_i \subset X$ has attached to it a subspace of the vector space dual to property space, i.e. $S_i \subset k^{n*}$. For simplicity, we assume that $S_i = \text{span}\{\xi_i\} =:< \xi_i >$. The act of sensing corresponds to taking a property vector $v \in k^n$ and returning a number $\xi_i(v)$ that records the strength of the detection. Outside of the sensor's support V_i, the sensor must return zero on every vector. In the overlap of two sensors' supports, the vector space that is sensed is the internal direct sum.

This cartoon specifically suggests the use of constructible sheaves and cosheaves as a model. Because the roles of sensors and intruders are formally dual, we will have to use *both* sheaves and cosheaves. Understanding the formal properties of sensing and evasion will lead us naturally to some long-exact sequences in cohomology, which will necessitate the introduction of barcodes to understand these results.

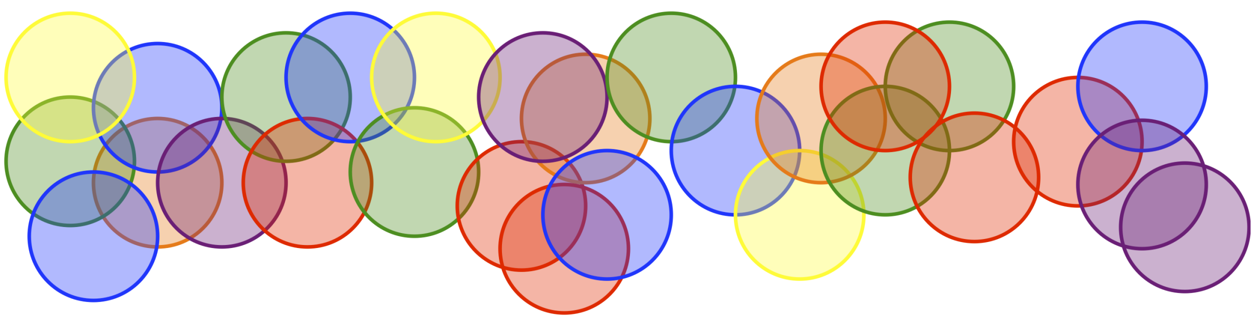

Figure 41: Multi-Modal Sensors Distributed in a Plane

We are going to work with a simplified version of the above cartoon. To detach ourselves from an embedding of the sensors into W, we will use the Čech nerve associated to the sensors supports. This will provide us with a simplicial complex X and this where we will define sheaves and cosheaves. Since we can only analyze the intruder problem inside sensor's support, we call this a **relative intruder problem**. Working strictly inside the coverage region will introduce counter-intuitive results, such as claim 10.24. Nevertheless, this setup is a prototype for future applications of sheaves and cosheaves to multi-modal sensing.

10.4.1 A Deeper Look at Sensing

Let us investigate a little more deeply the picture of multi-modal sensing presented to us in the above cartoon. In figure 42, we consider a situation where we have a sensor capable of detecting "red" properties and a sensor capable of detecting "green" properties.[40]

On the nerve of the sensor cover, the organizing diagram of vector spaces is clear.

$$< r^* > \hookrightarrow < r^*, g^* > \hookleftarrow < g^* > \qquad k \hookrightarrow k^2 \hookleftarrow k$$

The direction of the arrows indicates that a cellular sheaf is best used to collate sensing abilities. However, the diagram of abstract vector spaces on the right has no way of telling whether an individual copy of k should correspond to $< r^* >$ or $< g^* >$ or $< b^* >$. Such a distinction requires that we embed our sensing sheaf into a global system of coordinates $(k^n)^*_X$. This motivates the following definition.

Definition 10.16 (Sensing Sheaf)**.** Suppose we have a multi-modal sensor network distributed in a space W. Form the nerve given by the intersections of the sensors' supports and call this simplicial complex X. We define a **sensing sheaf** F by assigning to each vertex v in X the subspace $S_v \subset (k^n)^*$. Over higher simplices σ we assign the following vector spaces and use the natural inclusions for the maps internal to the sheaf:

$$F(\sigma) = S_{v_0} + \cdots + S_{v_n} \qquad F(\sigma) \hookrightarrow F(\tau) \quad \sigma \leqslant \tau.$$

Here we have used the internal sum of subspaces to reflect the fact there may be dependencies. The internal sum is only defined in the presence of an ambient space, thus part of the data of a sensing sheaf is an embedding into the constant sheaf of all sensing abilities:

$$\iota_F : F \hookrightarrow k^{n*}_X.$$

Now suppose we have an intruder, which we imagine as a point in the union of the red and green sensors in figure 42. The intruder has a property vector $v \in k^n$ that lists its various attributes, its colors in this example. What number does the sensor return while the intruder is in the red sensor's domain? By design, it is $r^*(v)$, the contraction of the red covector and the property vector v. If $k^n = k^3 =< v_r, v_g, v_b >$ is a three dimensional property space spanned by the attributes "red," "green," and "blue," equipping it with the standard Euclidean inner product allows us represent this measurement by the matrix product

$$r^*(v) = \begin{bmatrix} 1 & 0 & 0 \end{bmatrix} \begin{bmatrix} v_r \\ v_g \\ v_b \end{bmatrix} = v_r$$

However, if the sensors can collaborate and share information, then we can store together the observations when the intruder is in the intersection of the red and green sensors' support.

$$\begin{bmatrix} r^*(v) \\ g^*(v) \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix} \begin{bmatrix} v_r \\ v_g \\ v_b \end{bmatrix} = \begin{bmatrix} v_r \\ v_g \end{bmatrix}$$

[40]We use scare quotes to indicate that the terms can be substituted for whatever application is of interest.

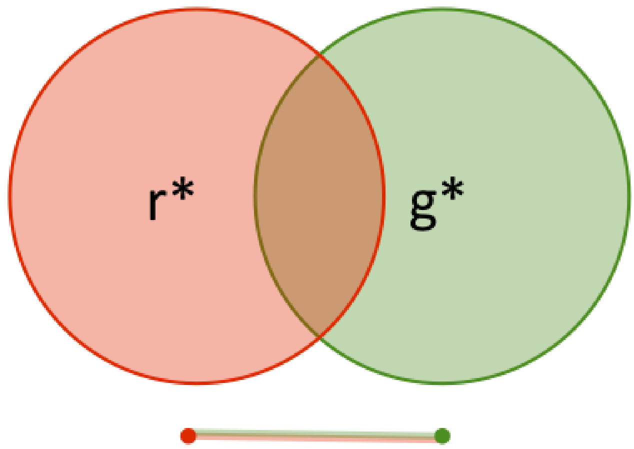

Figure 42: Two Multi-Modal Sensors

We can package these measurements into a cellular cosheaf, where two observations are the same modulo the properties unobserved by the sensors.

$$< r > \leftarrow < r, g > \rightarrow < g > \qquad \frac{k^3}{< g, b >} \leftarrow \frac{k^3}{< b >} \rightarrow \frac{k^3}{< r, b >}$$

One should note that the right side gives an equivalent formulation for the **measurement cosheaf** of the figure 42. We have over each cell passed to the quotient space where the properties that are invisible to each of the sensors is treated as zero. In other words, the vectors produced by the process of measurement must naturally be considered modulo the unknown.

Definition 10.17 (Evasion Co-Sheaf)**.** Given a sheaf F whose restriction maps are inclusions, along with a fixed embedding into a locally constant sheaf of vector spaces G (we take $G = k_X^{n*}$), we define the **annihilator cosheaf** $\widehat{Ann}(F)$ as follows:

- $\widehat{Ann}(F)(\sigma) = \{v^* \in G(\sigma)^* | v^*(\iota(w)) = 0 \quad \forall w \in F(\sigma)\}$
- If $\sigma \subset \bar{\tau}$, then $r_{\sigma,\tau} : \widehat{Ann}(F)(\tau) \to \widehat{Ann}(F)(\sigma)$ is the inclusion.

When using the language of sensing sheaves, we will call $\widehat{Ann}(F) =: \hat{E}$ the **evasion cosheaf**.

Lemma 10.18. *Let F be a sensing sheaf on X, then the evasion cosheaf is canonically identified as the linear dual of the cokernel of the embedding, that is to say that $\hat{E} \cong \hat{V}(\mathrm{cok}(\iota))$ in the diagram below.*

$$\begin{tikzcd}
0 \arrow[r] & F \arrow[r, "\iota"] \arrow[d, squiggly] & G \arrow[r, "q"] \arrow[d, squiggly] & \mathrm{cok}(\iota) \arrow[r] \arrow[d, squiggly] & 0 \\
0 & \hat{V}(F) \arrow[l] & \hat{V}(G) \arrow[l] & \hat{V}(\mathrm{cok}) \arrow[l] & 0 \arrow[l]
\end{tikzcd}$$

Proof. Here we make use of the fact that for cellular sheaves, the cell-by-cell cokernel of the maps $\iota(\sigma) : F(\sigma) \to G(\sigma)$ defines a sheaf. This is not always true for general sheaves. Reducing the argument to a cell-by-cell one, we have a short exact sequence of vector spaces

$$0 \longrightarrow V \xrightarrow{\iota} W \xrightarrow{q} W/V \longrightarrow 0$$

where we can identify

$$\mathrm{Ann}_W(V) = \{\varphi : W \to k \mid \varphi(v) = 0 \, \forall v \in V\} \cong (W/V)^*$$

and of course all the restriction maps get sent to restriction maps

$$\begin{tikzcd}
V_2 \arrow[dr] & & (W/V_2)^* \arrow[dd] \\
& W \arrow[ur] \arrow[dr] & \\
V_1 \arrow[uu] \arrow[ur] & & (W/V_1)^*.
\end{tikzcd}$$

□

This identification of evasion cosheaves with the linear dual of a cokernel means that we can leverage a classical technique in studying the relative intruder problem. After all, to every short exact sequence of sheaves we get an induced long exact sequence of sheaf cohomology. In the context of multi-modal sensing this relates in a precise way the topology of the total covered region and the cohomology of the sensing and evasion sheaves.

Theorem 10.19 (Sensing-Evasion Decomposition)**.** *Given a sensing sheaf of vector spaces* $\iota : F \to G = \tilde{k}_X^{n^*}$ *we obtain a long exact sequence of sheaf cohomology groups*

$$\begin{tikzcd}
0 \arrow[r] & H^0(X;F) \arrow[r] & H^0(X;k)^{\oplus n} \arrow[r] & H^0(X;\mathrm{cok}(\iota)) \arrow[dll, "\delta^0"] \\
& H^1(X;F) \arrow[r] & \cdots \arrow[r] & H^k(X;\mathrm{cok}(\iota)) \arrow[dll, "\delta^k"] \\
& H^{k+1}(X;F) \arrow[r] & H^{k+1}(X;k)^{\oplus n} \arrow[r] & H^{n+1}(X;\mathrm{cok}(\iota)) \arrow[dll, "\delta^{k+1}"] \\
& \cdots & &
\end{tikzcd}$$

Where $H^k(X;\mathrm{cok}(\iota))$ *gets identified with the evasion co-sheaf's homology* $H_k(X;\hat{E})$ *via the linear duality functor, i.e.* $V : \mathrm{cok}(\iota)) \rightsquigarrow E$.

Proof. The proof is immediate from standard homological algebra techniques. □

10.4.2 Indecomposables, Evasion Sets, Generalized Barcodes

One of the drawbacks of theorem 10.19 is that we have no good interpretation of what the sheaf cohomology groups mean. Let's consider again figure 42, but this time let us focus only on the Čech complex and each of the three sheaves that appear in the short exact sequence. This is depicted in figure 43.

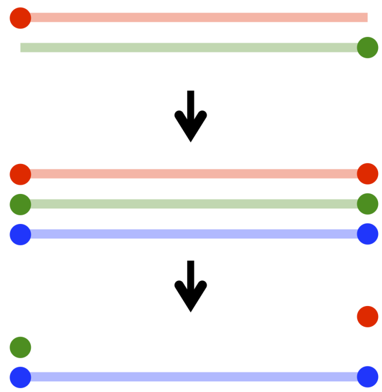

Figure 43: Examining the Short Exact Sequence

As can be clearly seen each sheaf appearing in the sequence is already written as a direct sum of indecomposables, which because the nerve is a one-simplex, look like barcodes. By using the observation $H^i([0,1];F) \cong \oplus H^i_c(B_i)$, which we have already made heavy use of, we can determine all the sheaf cohomology of interest for this example.

Example 10.20 (Red-Green Sensors)**.** By inspection of the indecomposable presentations of the three sheaves F, k^3_X and $\mathrm{cok}(\iota)$ in figure 43 we see that

$$H^i(X;F) \cong 0 \; i = 0,1; \qquad H^0(X;k^3_X) \cong k^3 \qquad H^0(X;\mathrm{cok}(\iota)) \cong H_0(X;\hat{E}) \cong k^3$$

The interpretation of each of the three generators in the evasion cosheaf homology is that there is a connected component where red, green and blue can separately evade.

Definition 10.21 (Evasion and Detection Sets)**.** Let $v \in k^n$ be a property vector and F a sensing sheaf on X. Define the **evasion set** E_v to be the set of points in X where an intruder with property vector v can go without being detected. Dually, call the set of points where v can be detected the **detection set** D_v.

Since we are working with cellular sheaves where individual sensors have support equal to the open star of their designated vertex in the simplicial complex X, thus the detection set D_v is equal to the union of all the stars of the sensors that can see v, hence D_v is an open union of cells. This proves the following lemma.

Lemma 10.22. *For any property vector v, the evasion and detection sets form an open-closed decomposition of X, that is*

$$X = E_v \cup D_v \qquad E_v \cap D_v = \emptyset, \qquad E_v \text{ open.}$$

When X is compact this means that E_v is compact as well.

We record another easy lemma, connected to our desire to get an indecomposable presentation for our evasion cosheaves.

Lemma 10.23. *Suppose that all sensors must pull their sensing capabilities from a fixed orthonormal basis of* k^{n*}*, say* $v_1^*, \dots, v_n^*$*, then the evasion cosheaf splits as a direct sum decomposition of constant cosheaves supported on the evasion sets for* $v_1, \dots, v_n$

$$\hat{E} \cong \hat{k}_{E_{v_1}} \oplus \cdots \oplus \hat{k}_{E_{v_n}}$$

with the further observation that each $\hat{k}_{E_{v_i}}$ *has a Remak decomposition as a sum of constant cosheaves supported on the components of* E_{v_i}*.*

Proof. The fact that the sensor capabilities can only be chosen from a fixed orthonormal basis, implies that we can write the constant sheaf k_X^{n*} as a direct sum of $k_X \oplus \cdots \oplus k_X$ where we think of each copy of k_X as being the constant sheaf generated by $< v_i^* >$. As a consequence we get the following diagram

$$\begin{array}{ccccc} & k_{D_{v_1}} & \hookrightarrow & k_X & \\ {\scriptstyle \pi_{v_1^*}} \nearrow & & & & \searrow {\scriptstyle i_{v_1^*}} \\ F & \vdots & & \vdots & k_X^{n*} \\ {\scriptstyle \pi_{v_n^*}} \searrow & & & & \nearrow {\scriptstyle i_{v_n^*}} \\ & k_{D_{v_n}} & \hookrightarrow & k_X & \end{array}$$

Now we can use for each factor the following standard short exact sequence of sheaves

$$0 \longrightarrow k_{D_{v_i}} \longrightarrow k_X \longrightarrow k_{E_{v_i}} \longrightarrow 0$$

and thus the cokernel splits as a direct sum $\oplus_{i=1}^n k_{E_{v_i}}$. □

The above lemma implies that in certain cases we can interpret the homology of the evasion cosheaf in terms of the topology of the evasion sets.

We have one more observation we'd like to leverage.

Claim 10.24. *If sensor's abilities are pulled from a fixed orthonormal basis* $v_1^*, \dots, v_n^*$ *and moreover the detection sets are not pairwise disjoint, then the sensing sheaf has no global sections.*

Proof. This follows from the fact that if an edge is common to two different detection sets, then there can be no global sections since the following sheaf has no non-zero global sections

$$< v_i^* > \hookrightarrow < v_i^*, v_j^* > \hookleftarrow < v_j^* > .$$

□

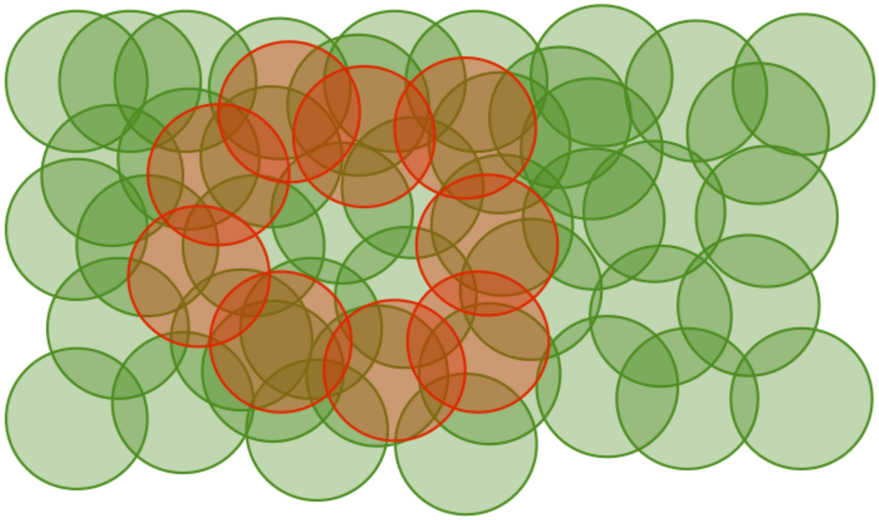

Figure 44: Short Exact Sequence

For the example considered in figure 44 assume that the space of properties is two dimensional, spanned by red and green. Then the theorem 10.19 provides the following forcing result

$$0 \to \mathsf{H}^0_{\mathsf{c}}(\mathsf{X};\mathsf{F}) \cong 0 \to \mathsf{H}^0_{\mathsf{c}}(\mathsf{X};\mathsf{k}^{2*}_{\mathsf{X}}) \cong \mathsf{k}^2 \to \mathsf{H}^0_{\mathsf{c}}(\mathsf{X};\mathrm{cok}) \to \mathsf{H}^1_{\mathsf{c}}(\mathsf{X};\mathsf{F}) \cong \mathsf{k} \to 0$$

which upon careful inspection reveals that the red evasion set must be disconnected.

11 The Derived Perspective

The derived perspective is a very important one for sheaf theory. In this perspective, one does not consider a sheaf in isolation, but rather one considers complexes of sheaves or, alternatively said, sheaves of complexes. This transition is motivated via an analogy with Taylor series in section 11.1. Injective and projective sheaves are introduced as the basic building blocks for the derived category, just as polynomials are the basis for Taylor series. Because the Alexandrov topology is so simple, we can described explicitly the elementary injective and projective (co)sheaves in section 11.1.1. Injective and projective resolutions are then introduced in section 11.1.2.

Section 11.2 gives a high-level introduction to the homological algebra techniques necessary to understanding the derived category. The explicitness of cellular sheaves allows us to give concrete examples of what is usually taken on faith when first learning the subject. The notion that maps are unique up to homotopy and that sheaves can be "quasi"-isomorphic without being isomorphic, are demonstrated in examples 11.12 and 11.17.

The derived definition of cosheaf homology is given in section 11.3 and the derived definition of sheaf cohomology can be dualized from there or looked up in Shepard's thesis [80]. These definitions should be regarded as the true definition of cosheaf homology and sheaf cohomology. The compactly supported variant, which we call Borel-Moore cosheaf homology, is defined in section 11.3.1. The derived functor formalism allows us to resolve the question of invariance under subdivision in section 11.3.2 with considerable ease.

Finally, we exploit the special features of the Alexandrov topology to develop two new theories: sheaf *homology* and cosheaf *cohomology*. Although these theories are invariant under subdivision in the domain of a map, they are not invariant under subdivision in the target of a map. These theories are sensitive to both the cell structure and the embedding. We compute some explicit examples of these theories in section 11.4.2.

11.1 Taylor Series for Sheaves

When first learning about the derived perspective a helpful analogy might be the following. We can approximate suitably nice functions around a point via the use of Taylor series:

$$f(x) \simeq f(a) + f'(a)(x-a) + \frac{f''(a)}{2!}(x-a)^2 + \cdots$$

The working physicist or engineer appreciates deeply how by only using a few terms, one can make serious headway into the analysis of integrals or other problems involving f.

In similar spirit one might start approximating or "taking the Taylor series expansion" of a topological space X via its homotopy or cohomology groups:

$$\pi_0(X,x) \quad \pi_1(X,x) \quad \pi_2(X,x) \quad \cdots \qquad\qquad \cdots \quad H_2(X) \quad H_1(X) \quad H_0(X)$$

One should realize that both of these series expansions arise from more fundamental sequences:

$$X \to \Omega_x X \to \Omega_x^2 X \to \cdots \qquad\qquad \cdots \to C_2(X;k) \to C_1(X;k) \to C_0(X;k)$$

Here $\Omega_x X$ denotes the space of loops in X based at x (and iterated applications thereof) and $C_p(X;k)$ denotes the p-chains.

For a sheaf F on a topological space X one also has a similar process. Namely, there is an exact sequence called a **resolution**

$$0 \to F \to I^0 \to I^1 \to I^2 \to \cdots$$

that when evaluated on an open set $U \subset X$ produces a sequence

$$0 \to F(U) \to I^0(U) \to I^1(U) \to I^2(U) \to \cdots$$

that is exact at $F(U)$ and $I^0(U)$.[41] Like the physicist with their Taylor series, one can discard the original sheaf and work solely with the terms in the sequence $I^\bullet(U)$. This sequence is a chain complex with potentially interesting cohomology. For each i, these cohomologies piece together to provide a pre-sheaf description of the complex of sheaves $I^\bullet$ – the derived replacement for F.

$$U \rightsquigarrow H^i(I^\bullet(U)) =: H^i(U;F).$$

If we specialize to the constant sheaf $F = k_X$, then we obtain another familiar series expansion of the space X: the cohomology. However, this series is much more general, as it encodes the cohomology of each open set in X. Consequently, even if one embeds X into the contractible cone CX, the constant sheaf and its derived replacement will remember the topology on X.

However, just as the reason that Taylor series are amenable to analysis because polynomials have simple properties, for general sheaves we must develop an algebraic analogue of a polynomial, which are the injective sheaves.

11.1.1 Elementary Injectives and Projectives

Definition 11.1. A representation of a small category $I : \mathcal{C} \to \mathbf{Vect}$ is **injective** if, for any natural transformation $\eta : A \to I$ and any injection $\iota : A \hookrightarrow B$, there is an extension $\tilde{\eta} : B \to I$ such that $\eta = \tilde{\eta} \circ \iota$. Said using diagrams, they are characterized by the usual universal property:

$$\begin{tikzcd} 0 \arrow[r] & A \arrow[r, "\iota"] \arrow[d, "\eta"'] & B \arrow[dl, dotted, "\exists\tilde{\eta}"] \\ & I & \end{tikzcd}$$

By playing around with the abstract properties one can easily prove the following statements:

- Every short exact sequence
$$0 \to I \to A \to Q \to 0$$
where I is injective, splits.

[41] There is no guarantee for exactness at higher terms.

- $\prod A_i$ is injective if and only if each A_i is injective.

The first statement is of particular importance. It expresses the idea that *indecomposables do not survive the derived perspective*. The second statement allows us to construct lots of injectives by taking products of perhaps easier to define injectives.

For our first example of an injective representation, we consider an injective cell sheaf. These sheaves are supported on the closures of cells.

Definition 11.2. An elementary injective cell sheaf on X concentrated on $\sigma \in X$ with value $W \in \textbf{Vect}$ is given by

$$[\sigma]^W(\tau) = \begin{cases} W & \text{if } \tau \leqslant \sigma, \\ 0 & \text{other wise.} \end{cases}$$

where the only possible non-zero restriction maps are the identity.

In order to prove that this sheaf is actually injective we introduce an alternative definition of injective sheaves and cosheaves defined on arbitrary posets. This definition makes use of the functors f_* and $f_\dagger$.

Definition 11.3. Let $i_x : \star \to X$ be the map that assigns to the one element poset the value $x \in X$, i.e. $x = i_x(\star)$. Define the elementary injective sheaf on $x \in X$ with value $W \in \textbf{Vect}$ to be $[x]^W = (i_x)_* W$ and the corresponding elementary injective cosheaf to be $\{\hat{x}\}^W := (i_x)_\dagger \hat{W}$.

One can see that for cosheaves, the elementary injectives are concentrated on the open stars of cells. To prove these objects are actually injective we make use of the adjunctions already defined.

Claim 11.4. *The sheaf* $[x]^W = (i_x)_* W$ *and cosheaf* $\{\hat{x}\}^W := (i_x)_\dagger \hat{W}$ *are injective.*

Proof. The proof is immediate from the following adjunctions

$$\mathbf{Hom}_{\mathbf{Shv}(X)}(A, (i_x)_* W) \cong \mathbf{Hom}_{\mathbf{Vect}}(A(x), W)$$

$$\mathbf{Hom}_{\mathbf{CoShv}(X)}(\hat{A}, (i_x)_\dagger \hat{W}) \cong \mathbf{Hom}_{\mathbf{Vect}}(\hat{A}(x), \hat{W})$$

and the fact that in the category of vector spaces every object is injective. $\square$

There is a dual universal object that is called projective.

Definition 11.5. A representation of a small category P is **projective** if for any natural transformation $\epsilon : P \to A$ and any surjection $\pi : B \twoheadrightarrow A$, there is a map $\tilde{\epsilon} : P \to B$ such that $\epsilon = \pi \circ \tilde{\epsilon}$. Said using diagrams:

$$\begin{tikzcd} & P \arrow[dl, dotted, "\exists \tilde{\epsilon}"'] \arrow[d, "\epsilon"] & \\ B \arrow[r, "\pi"'] & A \arrow[r] & 0 \end{tikzcd}$$

As before, we have some dual consequences:

- Every short exact sequence
$$0 \to A \to B \to P \to 0$$
where P is projective, splits.

- $\oplus B_i$ is projective if and only if each B_i is projective.

Since the adjunctions will be our guide we make the following definitions.

Definition 11.6. Let $i_x : \star \to X$ be the map that assigns to the one element poset the value $x \in X$, i.e. $x = i_x(\star)$. Define the elementary projective sheaf on $x \in X$ with value $W \in \mathbf{Vect}$ to be $\{x\}^W = (i_x)_\dagger W$ and the corresponding elementary projective cosheaf to be $[\hat{x}]^W := (i_x)_* \hat{W}$.

We leave it to the reader to check that these objects are actually projective.

Before moving on to the derived definition of sheaf cohomology, we record some useful identities that should be evident from the definition and the adjunctions.

$$\mathbf{Hom}_{\mathbf{Shv}}([\tau]^U, [\sigma]^W) = \begin{cases} \mathbf{Hom}_{\mathbf{Vect}}(U, W) & \text{if } \sigma \leqslant \tau, \\ 0 & \text{o.w.} \end{cases}$$

$$\mathbf{Hom}_{\mathbf{Shv}}(\{\tau\}^U, \{\sigma\}^W) = \begin{cases} \mathbf{Hom}_{\mathbf{Vect}}(U, W) & \text{if } \sigma \leqslant \tau, \\ 0 & \text{o.w.} \end{cases}$$

$$\mathbf{Hom}_{\mathbf{CoShv}}([\hat{\sigma}]^W, [\hat{\tau}]^U) = \begin{cases} \mathbf{Hom}_{\mathbf{Vect}}(W, U) & \text{if } \sigma \leqslant \tau, \\ 0 & \text{o.w.} \end{cases}$$

$$\mathbf{Hom}_{\mathbf{CoShv}}(\{\hat{\sigma}\}^W, \{\hat{\tau}\}^U) = \begin{cases} \mathbf{Hom}_{\mathbf{Vect}}(W, U) & \text{if } \sigma \leqslant \tau, \\ 0 & \text{o.w.} \end{cases}$$

11.1.2 Injective and Projective Resolutions

As promised, we aim to prove every sheaf has a resolution by injective sheaves. This follows from the following claim, which we now prove. Although this theorem is true for general spaces, we work with Alexandrov spaces arising as posets as usual.

Claim 11.7. *Every sheaf $F : X \to \mathbf{Vect}$ on a poset $(X, \leqslant)$ possibly of infinite size, F admits an inclusion into an injective sheaf. Dually, every cosheaf admits is surjected onto by a projective cosheaf.*
$$0 \to F \to I$$

Proof. We construct I explicitly. It is given by

$$0 \to \mathsf{F} \to \mathrm{I} := \prod_x [x]^{\mathsf{F}(x)} = \prod_x (i_x)_* \mathsf{F}(x).$$

The map to I is defined easily using the standard adjunctions

$$\iota \in \mathbf{Hom}(\mathsf{F}, \prod_x (i_x)_* \mathsf{F}(x)) \cong \prod_x \mathbf{Hom}(\mathsf{F}, \prod_x (i_x)_* \mathsf{F}(x)) \cong \prod_x \mathbf{Hom}(\mathsf{F}(x), \mathsf{F}(x)) \ni \prod_x \mathrm{id}_{\mathsf{F}(x)}.$$

We encourage the reader to describe this map is explicitly, by seeing how a single $\mathrm{id}_{\mathsf{F}(x)}$ traces through this adjunction, which we'll call $\iota_x \in \mathbf{Hom}(\mathsf{F}, (i_x)_*\mathsf{F}(x))$.

Similarly, for a cosheaf $\hat{\mathrm{F}} : \mathsf{X}^{op} \to \mathbf{Vect}$ on an Alexandrov space we could have built a projective surjection by taking

$$\hat{\mathsf{P}}_0 := \bigoplus_x (i_x)_* \hat{\mathrm{F}}(x) \to \hat{\mathrm{F}} \to 0$$

where the map $\pi_0 : \mathsf{P}_0 \to \hat{\mathrm{F}}$ is gotten through the corresponding adjunction for cosheaves and using the contravariance of **Hom** in the first slot

$$\mathbf{Hom}(\bigoplus_x (i_x)_* \hat{\mathrm{F}}(x), \hat{\mathrm{F}}) \cong \prod \mathbf{Hom}((i_x)_* \hat{\mathrm{F}}(x), \hat{\mathrm{F}}) \cong \prod \mathbf{Hom}(\hat{\mathrm{F}}(x), \hat{\mathrm{F}}(x)).$$

□

Corollary 11.8. *Every sheaf* $\mathsf{F} : \mathsf{X} \to \mathbf{Vect}$ *has an injective resolution. Dually, every cosheaf* $\hat{\mathrm{F}} : \mathsf{X}^{op} \to \mathbf{Vect}$ *has a projective resolution.*

Proof. Since cokernels exist in the category of sheaves by taking element-by-element quotients and by iteratively applying the claim, we obtain an injective resolution of F:

$$\begin{tikzcd} \mathsf{F} \arrow[r, "\iota^0"] & \mathrm{I}^0 \arrow[rr, "\iota^1 = j_1\pi_0"] \arrow[dr, "\pi_0"'] & & \mathrm{I}^1 \arrow[rr, "\iota^2 = j_2\pi_1"] \arrow[dr, "\pi_1"'] & & \mathrm{I}^2 & \cdots \\ & & \mathrm{cok}(\iota^0) \arrow[ur, "j_1"'] & & \mathrm{cok}(\iota^1) \arrow[ur, "j_2"'] & & \cdots \end{tikzcd}$$

Iterating the analogous process for projective cosheaves, replacing kernels where one sees cokernels above, one obtains an exact sequence of cosheaves called the **projective resolution** of $\hat{\mathrm{F}}$:

$$\cdots \hat{\mathsf{P}}_2 \to \hat{\mathsf{P}}_1 \to \hat{\mathsf{P}}_0 \to \hat{\mathrm{F}} \to 0.$$

□

These exact sequences can be used to replace F or $\hat{\mathrm{F}}$ in a suitable sense, defined by the derived category. Before moving onto that discussion, we note one interesting point.

Proposition 11.9. *The length of injective resolution of any sheaf* $F \in \mathbf{Shv}(X)$ *is bounded by the length of longest chain in the poset. In particular for* X *a cell complex, it is bounded by the dimension.*

Proof. Pick a maximal ordered subset in X and consider its top element, say x', then $I^0(x') = F(x')$ since nothing is larger than x'. The cokernel sheaf of ι^0 evaluated on x' is then $\mathrm{cok}(\mathrm{id} : F(x') \to F(x')) = 0$. So for any maximally ordered chain in X, I^1 is zero on the top-most element. Arguing inductively finishes the proof. $\square$

11.2 The Derived Category and Homotopy Theory of Chain Complexes

The purpose of the derived category is to replace the category of sheaves with a category of complexes where certain operations are more natural. We have already shown that one can replace a sheaf by its injective resolution and a cosheaf by its projective resolution. This will define our derived replacement on the level of objects, but we have not yet shown how a map of sheaves or cosheaves induces a map on the level of resolutions.

If $\phi : \hat{F} \to \hat{G}$ is a map of cosheaves, then it can be checked from the universal properties of projective objects, that this induces a map of complexes

$$\begin{tikzcd} \cdots & \hat{P}_2 \arrow[r] \arrow[d, dotted, "\phi_2"] & \hat{P}_1 \arrow[r] \arrow[d, dotted, "\phi_1"] & \hat{P}_0 \arrow[r] \arrow[d, dotted, "\phi_0"] & \hat{F} \arrow[r] \arrow[d, "\phi"] & 0 \\ \cdots & \hat{Q}_2 \arrow[r] & \hat{Q}_1 \arrow[r] & \hat{Q}_0 \arrow[r] & \hat{G} \arrow[r] & 0 \end{tikzcd}$$

where all the squares in sight commute. For a hint on how to see this, consider the composite map $\hat{P}_0 \to \hat{F} \to \hat{G}$ and let $\hat{G} = A$ and $B = \hat{Q}_0$ in the definition of the universal property defining a projective object. This induces our first map $\hat{P}_0 \to \hat{Q}_0$. To get the next, all important step, one must recognize that having maps from $\hat{P}_0 \to \hat{Q}_0$ and $\hat{F} \to \hat{G}$ induces maps between the kernels of the map $\hat{P}_0 \to \hat{F}$ and $\hat{Q}_0 \to \hat{G}$. Since $\hat{Q}_1$ surjects onto the kernel of the latter map repeating the initial argument provides a map from $\hat{P}_1$ to $\hat{Q}_1$. This shows that the projective replacement of cosheaves is functorial.

Aside from functoriality, there is one more snag that needs to be mentioned: For a sheaf or a cosheaf it is possible that the choice of injective or projective resolution is not unique. If one really wants to use these as replacements for the original sheaf or cosheaf, there must be a strong relationship between these two complexes. This is best seen by specializing the functoriality discussion above to the case $\phi = \mathrm{id}$.

$$\begin{tikzcd} \cdots & \hat{P}_2 \arrow[r] \arrow[d, dotted, "\phi_2"] & \hat{P}_1 \arrow[r] \arrow[d, dotted, "\phi_1"] & \hat{P}_0 \arrow[r] \arrow[d, dotted, "\phi_0"] & \hat{F} \arrow[r] \arrow[d, "\mathrm{id}"] & 0 \\ \cdots & \hat{Q}_2 \arrow[r] & \hat{Q}_1 \arrow[r] & \hat{Q}_0 \arrow[r] & \hat{F} \arrow[r] & 0 \end{tikzcd}$$

The resulting map of complexes need not be a term-by-term isomorphism with all squares in sight commuting, but rather a more general notion must be substituted, namely the definition of chain homotopy. Before giving that, let us give an example.

Example 11.10 (Non-Unique Projective Resolutions)**.** Let us work again over our test space of the closed unit interval $X = [0, 1]$ stratified as $x = 0$, $y = 1$ and $a = (0, 1)$. The constant cosheaf $\hat{k}_X$ is then modeled as

$$\begin{tikzcd} & k \arrow[dl, "1"'] \arrow[dr, "1"] & \\ k & & k \end{tikzcd}$$

Revisiting the definition of the elementary projective cosheaves, there is one obvious projective resolution because the constant cosheaf on this stratification of the unit interval is already projective, so we have the identity map

$$\hat{P}_\bullet : \qquad [\hat{a}] \to \hat{k}_X.$$

On the other hand, following blindly the prescription provided for computing the projective resolution of an arbitrary cellular cosheaf would have lead us to the following "canonical" resolution:

$$\hat{Q}_\bullet : \qquad [\hat{x}] \oplus [\hat{y}] \to [\hat{a}] \oplus [\hat{x}] \oplus [\hat{y}] \to \hat{k}_X$$

Definition 11.11 (Chain Homotopy)**.** Suppose $(A^\bullet, d_A)$ and $(B^\bullet, d_B)$ are two (cohomological) chain complexes and $\phi^\bullet$ and $\psi^\bullet$ are two chain maps, then **a chain homotopy** $h^\bullet$ is a chain map $h^i : A^i \to B^{i-1}$

$$\begin{tikzcd} \cdots \arrow[r] & A^{i-2} \arrow[r] \arrow[d] & A^{i-1} \arrow[r] \arrow[dl] \arrow[d, "\phi^{i-1}"', "\psi^{i-1}"] & A^i \arrow[r] \arrow[dl] \arrow[d, "\phi^i"', "\psi^i"] & A^{i+1} \arrow[r] \arrow[dl] \arrow[d, "\phi^{i+1}"', "\psi^{i+1}"] & \cdots \\ \cdots \arrow[r] & B^{i-2} \arrow[r] & B^{i-1} \arrow[r] & B^i \arrow[r] & B^{i+1} \arrow[r] & \cdots \end{tikzcd}$$

such that

$$\phi^i - \psi^i = h^{i+1} d_A^i + d_B^{i-1} h^i.$$

In which case we say that $\phi \sim \psi$ are **chain homotopic**.

Consider now two chain maps $\phi : A^\bullet \to B^\bullet$ and $\psi : B^\bullet \to A^\bullet$, such that

$$\phi \circ \psi \sim \mathrm{id} \qquad \text{and} \qquad \psi \circ \phi \sim \mathrm{id}$$

then one says $A^\bullet$ and $B^\bullet$ are **chain homotopy equivalent**.

Example 11.12 (Non-unique, but equivalent)**.** Consider again the case of the two different projective resolutions of the constant sheaf $\hat{k}_X$ on the closed unit interval. On the one hand the composite

$$\begin{tikzcd}
0 \arrow[r] \arrow[d] & {[\hat{a}]} \arrow[d] \\
{[\hat{x}] \oplus [\hat{y}]} \arrow[r] \arrow[d] & {[\hat{x}] \oplus [\hat{a}] \oplus [\hat{y}]} \arrow[d] \\
0 \arrow[r] & {[\hat{a}]}
\end{tikzcd}$$

is clearly the identity on $\hat{P}_\bullet$, but the composite

$$\begin{tikzcd}
{[\hat{x}] \oplus [\hat{y}]} \arrow[r] \arrow[d] & {[\hat{x}] \oplus [\hat{a}] \oplus [\hat{y}]} \arrow[d] \\
0 \arrow[r] \arrow[d] & {[\hat{a}]} \arrow[d] \\
{[\hat{x}] \oplus [\hat{y}]} \arrow[r] & {[\hat{x}] \oplus [\hat{a}] \oplus [\hat{y}]}
\end{tikzcd}$$

cannot possibly be the identity because one map factors through zero. However, if we employ a self-homotopy of $\hat{Q}_\bullet$ by defining a homotopy for the only possible degree to be

$$h^0 : [\hat{a}] \oplus [\hat{x}] \oplus [\hat{y}] \to [\hat{x}] \oplus [\hat{y}]$$

which is zero on the a component and the identity elsewhere. One can then check that this defines a homotopy between the identity map and the map indicated in the second composite.

The conclusion from the example should be that although one can use different projective resolutions, the choice is irrelevant up to homotopy. The derived category should not be able to discriminate between them. As such, we make the following definitions.

Definition 11.13. Let $\mathcal{A}$ be an abelian category, such as the category of sheaves or cosheaves. The **category of chain complexes** in $\mathcal{A}$, written $C^b(\mathcal{A})$ has objects that are chain complexes and morphisms that are chain maps.

The **homotopy category of complexes** $K^b(\mathcal{A})$ of an abelian category $\mathcal{A}$ has the same objects as $C^b(\mathcal{A})$, but where we have identified chain homotopic maps.

Definition 11.14. For $\mathcal{A} = \mathbf{Shv}(X)$ we define the **bounded derived category of sheaves** $D^b(\mathbf{Shv}(X))$ to be $K^b(\text{Inj} - \mathbf{Shv}(X))$ the homotopy category that uses only complexes of injective sheaves.

Similarly, for $\mathcal{A} = \mathbf{CoShv}(X)$, we define the **bounded derived category of cosheaves** $D^b(\mathbf{CoShv}(X))$ by $K^b(\text{Proj} - \mathbf{CoShv}(X))$ where complexes of projective cosheaves are used instead.

This definition, is an equivalent reformulation of another definition of the derived category. This other perspective is built on the foundational notion of a quasi-isomorphism, which is in turn built on the idea of a cohomology sheaf or homology cosheaf.

Definition 11.15. Suppose we are given a complex of cellular sheaves

$$(F^\bullet, d^\bullet): \qquad \cdots \to F^{i-1} \to F^i \to F^{i+1} \to \cdots,$$

i.e. for each cell σ we have a complex of vector spaces. For each i we can define the ith **cohomology sheaf** as the assignment

$$\mathcal{H}^i(F^\bullet): \qquad \sigma \rightsquigarrow H^i(F^\bullet(\sigma))$$

which is a cellular sheaf. The restriction maps being defined as the induced map on cohomology for the chain map $F^\bullet(\sigma) \to F^\bullet(\tau)$ for $\sigma \leqslant \tau$.

Considering all i at once defines a functor from the category of complexes of sheaves and the category of graded sheaves (sheaves of graded vector spaces with level preserving restriction maps)

$$\mathcal{H}^*: C^b(\mathbf{Shv}(X)) \to \mathbf{Shv}(X; \mathrm{gr}\,\mathbf{Vect}) \qquad F^\bullet \rightsquigarrow \bigoplus_i \mathcal{H}^i(F^\bullet).$$

There are completely dual notions of **homology cosheaves**, where we generally use homological indexing and notation $(\hat{F}_\bullet, \partial_\bullet)$.

Definition 11.16 (Quasi-Isomorphisms)**.** A map of complexes of sheaves (or cosheaves) $\alpha^\bullet: F^\bullet \to G^\bullet$ such that the induced map

$$\mathcal{H}(\alpha^\bullet): \mathcal{H}^i(F^\bullet) \to \mathcal{H}^i(G^\bullet)$$

is an isomorphisms for every i, is called a **quasi-isomorphism**.

The term "quasi-isomorphism" reflects the fact that if $\alpha^\bullet: F^\bullet \to G^\bullet$ is a quasi-isomorphism, then there does not always exist an inverse map $\beta^\bullet$ that gives the identity, or even chain homotopic to the identity, any map back may simply not exist.

Example 11.17. Consider again the unit interval $X = [0,1]$ decomposed into two vertices x and y and an open interval a. Consider the stalk sheaf S_a that assigns k to a and is zero everywhere else. It's injective resolution defines a chain map

$$\begin{tikzcd} 0 \arrow[r] \arrow[d] & S_a \arrow[r] \arrow[d] & 0 \arrow[d] \\ 0 \arrow[r] & {[a]} \arrow[r] & {[x] \oplus [y]} \end{tikzcd}$$

which is a quasi-isomorphism. However, there does not exist a map of sheaves $[a] \to S_a$.

The slogan most commonly associated with the derived category is that one "formally inverts the quasi-isomorphisms." This is formalized by the process of localizing categories. Namely, if Q is a collection of morphisms in $\mathcal{B}$ that is closed under certain operations, then we can consider the following universal problem: suppose $L : \mathcal{B} \to \mathcal{C}$ is a functor such that if $\alpha \in Q$, then $L(\alpha)$ is an isomorphism, then every such functor factors through the **category localized at** Q, written $\mathcal{B}[Q^{-1}]$.

$$\begin{tikzcd} \mathcal{B} \arrow[rr, "L"] \arrow[dr] & & \mathcal{C} \\ & \mathcal{B}[Q^{-1}] \arrow[ur, dotted, "\exists"'] & \end{tikzcd}$$

An alternative approach to the derived category of an abelian category $\mathcal{A}$ is to define

$$D(\mathcal{A}) := K(\mathcal{A})[Q^{-1}] \qquad Q = \{\text{quasi} - \text{isomorphisms}\}$$

where we have removed the boundedness hypothesis.

One then proves the following claim to re-obtain the definition we provided here

Theorem 11.18 ([6] Thm 6.7)**.** *Suppose* A *is an abelian category with enough projectives, then* $D^-(A) \cong K^-(P)$ *where* P *denotes projective objects of* A*. Similarly, if* A *has enough injectives then* $D^+(A) \cong K^+(I)$ *where* I *denotes injective objects of* A*.*

11.3 The Derived Definition of Cosheaf Homology and Sheaf Cohomology

We are now in a position to give the derived definition of cosheaf homology and show that it agrees with the computational formula provided earlier. This discussion can be dualized and readily found in the literature. The proof that the formula for sheaves computes the cohomology as defined by taking an injective resolution and applying $\Gamma(X; -) = p_*$ can be found in [80] pp. 28-29. Let's more or less repeat the proof for cellular cosheaves since it is nowhere in the literature.

Definition 11.19. Given a cosheaf $\hat{F}$ on X we define the **left derived pushforward** along $f : X \to Y$ by taking a projective resolution and applying pushforward term by term:

$$Lf_* \hat{F} := f_* P_\bullet.$$

We define the ith derived functor by

$$L_i f_* \hat{F} := \mathcal{H}_i(f_* P_\bullet).$$

In the special case where $f = p : X \to \star$ we write

$$H_i(X; \hat{F}) := L_i p_* \hat{F}$$

for the ith cosheaf homology group of $\hat{F}$.

We now aim to prove the following theorem.

Theorem 11.20. *Let $p : X \to \star$ be the constant map and $\hat{F}$ a cellular cosheaf on X. Then the left derived functors of p_* agree with the computational formula for homology, i.e.* $L_i p_* \hat{F} = H_i(X; \hat{F})$.

Proof. Begin with a projective resolution of $\hat{P}_\bullet \to \hat{F}$ and then take cellular chains of each cosheaf to obtain the following double complex:

$$\begin{tikzcd}
 & \vdots & \vdots & \vdots & \\
\cdots & C_1(X;\hat{P}_1) \arrow[r] \arrow[d] & C_1(X;\hat{P}_0) \arrow[r] \arrow[d] & C_1(X;\hat{F}) \arrow[r] \arrow[d] & 0 \\
\cdots & C_0(X;\hat{P}_1) \arrow[r] \arrow[d] & C_0(X;\hat{P}_0) \arrow[r] \arrow[d] & C_0(X;\hat{F}) \arrow[r] \arrow[d] & 0 \\
\cdots & \operatorname{colim} \hat{P}_1 \arrow[r] \arrow[d] & \operatorname{colim} \hat{P}_0 \arrow[r] \arrow[d] & \operatorname{colim} \hat{F} \arrow[r] \arrow[d] & 0 \\
 & 0 & 0 & 0 &
\end{tikzcd}$$

Now we make use of the following two observations.

Lemma 11.21 (cf. [80] Thm 1.3.10)**.** *For $\hat{P}$ a projective cosheaf*

$$H_p(C_\bullet^{BM}(X;\hat{P})) \cong H_p(C_\bullet(X;\hat{P})) \cong 0$$

for $p > 0$.

Proof. Observe that we can assume that $\hat{P}$ is an elementary projective co-sheaf with value $\mathbb{R}$, i.e. $[\hat{\sigma}]$, since $C_\bullet^{BM}(X; \oplus A_i) = \oplus C_\bullet^{BM}(X; A_i)$.

Everything follows from the following consequence of our definition of a cell complex: In the one-point compactification of X, the closure of any cell $\sigma \in X$, call it $|\bar{\bar{\sigma}}|$, has the homeomorphism type of a closed k-simplex.

$C_\bullet(X; [\hat{\sigma}])$ is the chain complex that computes the cellular homology of $Y = |\{\tau \leqslant \sigma | \bar{\tau} \text{ is compact}\}|$, which is a closed k-simplex minus the star of a vertex. On the other hand, $C_\bullet^{BM}(X; [\hat{\sigma}])$ is equal to the chain complex calculating the cellular homology of $|\bar{\bar{\sigma}}|$ except in degree zero if $|\bar{\sigma}|$ is not compact. Notice that H_1 for both of these complexes is the same, as $|\bar{\sigma}|$ and $|\bar{\bar{\sigma}}|$ are simply connected. This proves the claim. ☐

Lemma 11.22 (cf. [80] Thm 1.4.1)**.** *For any cellular cosheaf $\hat{F}$ on a cell complex X we have that*

$$\operatorname{colim} \hat{F} \cong \operatorname{cok}(C_1(X;\hat{F}) \to C_0(X;\hat{F})).$$

Proof. First let us prove that taking the coproduct of $\hat{F}$ over all the cells obtains a vector space that surjects onto the colimit. As part of the definition of $\operatorname{colim}\hat{F}$ is a choice of maps $\psi_\sigma : \hat{F}(\sigma) \to \operatorname{colim}\hat{F}$. Let $\Psi = \oplus\psi_\sigma : \oplus\hat{F}(\sigma) \to \operatorname{colim}\hat{F}$, now consider the factorization of this map through the image:

$$\begin{tikzcd} \oplus\hat{F}(\sigma) \arrow[rr, "\Psi"] \arrow[dr] & & \operatorname{colim}\hat{F} \\ & \operatorname{im}\Psi \arrow[ur, "j"] & \end{tikzcd}$$

Now we can use the $\operatorname{im}\Psi$ to define a new co-cone over the diagram $\hat{F}$ simply by pre-composing the factorized map with the inclusions $i_\sigma : \hat{F}(\sigma) \to \oplus\hat{F}(\sigma)$. Since the colimit is the initial object in the category of co-cones, there must be a map $u : \operatorname{colim}\hat{F} \to \operatorname{im}\Psi$ and thus $u \circ j = \mathrm{id}$ since there is only one map $\operatorname{colim}\hat{F} \to \operatorname{colim}\hat{F}$.

Now observe that $C_0(X; \hat{F}) = \oplus\hat{F}(v_i)$ surjects onto the colimit of $\hat{F}$ by virtue of the fact that since every cell $\sigma \in X$ has at least one vertex as a face, the map Ψ factors through $\oplus\hat{F}(v_i)$. Thus there is a surjection from $\Psi' : C_0(X; \hat{F}) \to \operatorname{colim}\hat{F}$. Notice that by universal properties of the cokernel of $\partial_0 : C_1(X; \hat{F}) \to C_0(X; \hat{F})$ it suffices to check that $\Psi' \circ \partial_0 = 0$. However, this is clear since every e edge has two vertices v_1 and v_2 (we've discarded all those edges without compact closures), then we need only check the claim for each diagram of the form

$$\begin{tikzcd} & \hat{F}(e) \arrow[dl, "r_{e,v_1}"'] \arrow[dr, "r_{e,v_2}"] & \\ \hat{F}(v_1) & & \hat{F}(v_2) \end{tikzcd}$$

where it is clear that the colimit can be written as $\hat{F}(v_1) \oplus \hat{F}(v_2)$ modulo the equivalence relation $(r_{e,v_1}(w), 0) \simeq (0, r_{e,v_2}(w))$, i.e. $\partial_0|_e(w) = (-r_{e,v_1}(w), r_{e,v_2}(w)) \simeq (0,0)$. □

From these two theorems we can conclude that the columns away from the chain complex of $\hat{F}$ are exact and thus $\mathrm{Tot}_\bullet(C_i(X; \hat{P}_j))$ induces quasi-isomorphisms between $\operatorname{colim}\hat{P}_\bullet$ and $C_\bullet(X; \hat{F})$. We have thus established the theorem. □

11.3.1 Borel-Moore Cosheaf Homology

Definition 11.23. Suppose $\hat{F}$ is a cellular cosheaf. Define $\Gamma^{BM}(X; F)$ to be the colimit of the diagram extended over the one-point compactification of X where we define $\hat{F}(\infty) = 0$. Alternatively said we look at the inclusion $j : X \to X \cup \{\infty\}$ and define

$$\Gamma^{BM}(X; \hat{F}) := p_* j_\dagger \hat{F}.$$

Another possible definition is to dualize a cellular cosheaf of finite-dimensional vector spaces to a cellular sheaf by post-composing $\hat{F} : X \to \boldsymbol{Vect}$ with $\mathbf{Hom}_{\boldsymbol{vect}}(-, k)$, apply $p_!$ and then dualize back.

Remark 11.24 (Functoriality). The definitions that involve the one-point compactification are deficient in the following way. A map of cell complexes $f : X \to Y$ does not necessarily extend to a map between the one-point compactifications. It is for this reason that for functoriality, the definition using $p_!$ is preferred.

Now we can prove that the formula provided calculates the Borel-Moore homology of a cosheaf $\hat{F}$ by establishing the following lemma:

Lemma 11.25. *For any cellular cosheaf* $\hat{F}$ *on a cell complex* X *we have that* $\Gamma^{BM}(X;\hat{F}) \cong \operatorname{cok}(C_1^{BM}(X;\hat{F}) \to C_0^{BM}(X;\hat{F}))$.

Proof. The proof above goes through until the last argument. Now we have edges e with only one vertex. However, by extending and zeroing out at infinity to get that the colimit of

$$\begin{tikzcd} & \hat{F}(e) \arrow[dl, "r_{e,v}"'] \arrow[dr, "0"] & \\ \hat{F}(v) & & \hat{F}(\infty) = 0 \end{tikzcd}$$

is exactly equal to the co-equalizer of $r_{e,v} : \hat{F}(e) \to F(v)$ and the zero morphism, i.e. the cokernel. □

11.3.2 Invariance under Subdivision

Now we take up the question of invariance under subdivision by applying the derived perspective. For convenience, we work with sheaves, but the reasoning can be dualized.

Definition 11.26. Suppose F is a sheaf on X and $s : X' \to X$ is a subdivision of X, then we define the subdivided sheaf $F' := s^*F$.

For an example, let X be the unit interval $[0, 1]$ stratified in the obvious way with $x = 0$, $y = 1$ and $a = (0, 1)$. Now consider a sheaf F on X. We will want to investigate what happens to this sheaf as we subdivide the space. In this example, the barycentric subdivision of X produces a space X' with a third vertex $\bar{a}$ and two edges a_x and a_y. The obvious way of defining a subdivided sheaf is to define $F'(\bar{a}) = F'(a_x) = F'(a_y) = F(a)$ where we use the identity map for the two new restriction maps. Observe that if F is the elementary injective sheaf $[a]$, then F' is *not* an injective sheaf, yet nevertheless F' and F have isomorphic cohomology.

More generally we are concerned with the following diagram of spaces (posets)

$$\begin{tikzcd} X' \arrow[rr, "s"] \arrow[dr, "p_{X'}"'] & & X \arrow[dl, "p_X"] \\ & \star & \end{tikzcd}$$

and the induced functors on sheaves. For example, if we analyze the ordinary pushforward functor, then we would obtain the following result, which is a simplified proof of one found in [80]:

Theorem 11.27 (cf. [80], Thm. 1.5.2)**.** *Suppose* F *is a sheaf on* X *and* X' *is a subdivision of* X*, then*

$$H^\bullet(X;F) \cong H^\bullet(X';F')$$

Proof. Observe that since $p_{X'} = p_X \circ s$, then $(p_{X'})_* = (p_X)_* \circ s_*$. Now recall

$$(p_{X'})_*F' = (p_{X'})_*s^*F = (p_X)_* \circ s_*s^*F.$$

The question then boils down to understanding the relationship between s_*s^*F and F. Unraveling the definition reveals

$$\begin{array}{rcl} s_*s^*F(y) & = & \varprojlim\{s^*F(x)|s(x) \geqslant y\} \\ & = & \varprojlim\{F(s(x))|s(x) \geqslant y\} \\ \text{(surjectivity)} & = & \varprojlim\{F(x)|x \geqslant y\} \\ \text{(sheaf} - \text{axiom)} & = & F(U_y) \\ & = & F(y) \end{array}$$

So we have that for the subdivision map $s_*s^*F \cong F$ and as a consequence

$$(p_{X'})F' \cong p_X F.$$

Now we can just take the associated right derived functors to obtain the result. □

11.4 Sheaf Homology and Cosheaf Cohomology

There is a surprising symmetry in the land of cellular sheaves and cosheaves, which is unique to the land of Alexandrov spaces and deserves to be explored. Contrary to the existence of enough injective sheaves, which for general sheaves is gotten as a consequence of the target category, e.g. $\mathbf{Ab}, \mathbf{Vect}$, etc., the existence of enough projective sheaves is driven by the underlying topology of the space.

Proposition 11.28. *Suppose* X *is a topological space with the property that there is a point* $x \in X$ *such that for every open neighborhood* $U \ni x$ *there is a strictly smaller open neighborhood* $V \subset U$*. Then the category of sheaves on* X *does not have enough projectives.*[42]

Proof. Consider the map $i : x \hookrightarrow X$ and the sheaf i_*k. Suppose it has a projective resolution, i.e. a projective sheaf P and a surjection $P \to i_*k$. Now let's examine this map evaluated on an open set $U \ni x$. By assumption there is another open set $V \subset U$ and we can put the constant sheaf extended by zero on V, denote the inclusion by $j : V \hookrightarrow X$. Note that we have the following diagram of sheaves

$$\begin{tikzcd} j_!\tilde{k}_V \arrow[r] & i_*k \arrow[r] & 0 \\ & P \arrow[u] \arrow[ul, dotted] & \end{tikzcd}$$

[42]The author would like to acknowledge the contributions of Valery Alexeev, David Treumann, and Jon Woolf on mathoverflow in regards to this question.

whose value on the open set U is

$$
\begin{tikzcd}
j_!\tilde{k}_V(U) = 0 \arrow[r] & i_*k(U) = k \arrow[r] & 0 \\
 & P(U) \arrow[u] \arrow[ul, dotted] &
\end{tikzcd}
$$

so in particular the surjection must factor through zero – a contradiction. $\square$

Contrary to sheaves on manifolds and other Hausdorff spaces, cellular sheaves are can be viewed as sheaves on finite posets and as such do not suffer from the above argument. In fact, computing a projective resolution is as easy as computing injective resolutions. To see how this goes recall we need to find a projective sheaf that surjects onto our sheaf of interest.

$$P^0 := \oplus_{\sigma\in X}\{\sigma\}^{F(\sigma)} \to F \to 0$$

serves nicely and by finding the kernel sheaf (which is easier to understand than cokernels!) and then iterating this process will obtain a projective resolution

$$\cdots P^{-3} \to P^{-2} \to P^{-1} \to P^0 \to F \to 0.$$

This motivates the following definitions:

Definition 11.29. Given a cellular sheaf, we can construct its projective resolution $P_\bullet \to F$, calculate colimits of $P_\bullet$ and take the cohomology of the resulting complex of vector spaces. Assuming F was in degree zero, this will be concentrated in negative degree and we define the **homology of a cellular sheaf** F to be $H_i(X; F) := H^{-i}(p_\dagger P_\bullet)$.

Similarly we define the **cohomology of a cellular cosheaf** $\hat{F}$ by taking its injective resolution $\hat{I}^\bullet$, and taking limits, i.e. $H^i(X; \hat{F}) = H^i(\hat{p}_*\hat{I}^\bullet)$.

The reasons for it's apocryphal nature are many:

1. Only for (co)sheaves over finite spaces are there enough projectives and enough injectives.

2. Spaces for which there is not a fixed n so that every cell σ contains in its star a cell τ such that $\dim\tau = n$ cannot hope to have the same computational formula for (co)homology because we can't treat the colimit (in the case of a sheaf) as a quotient object of $\oplus_{\dim\tau=n}F(\tau)$ and dually for limits of cosheaves.

3. This defect, which is measured by the difference of $H^n(X; F)$ and $\operatorname{colim} F$, is only the first in a series of obstructions that appear to detect whether X is a cell structure on a manifold.

The evidence for the last two observations is further solidified in view of the following theorem.

Theorem 11.30. *Suppose F is a cellular sheaf on a triangulated closed n-manifold X, then F defines a cellular cosheaf on the dual triangulation and moreover all the homologies and cohomologies of both agree.*

Proof. This is a consequence of the following simple observation:

$$\begin{tikzcd} F(\sigma^i) \arrow[r, "\rho_{\sigma,\tau}"] \arrow[d, squiggly] & F(\tau^{i+1}) \arrow[d, squiggly] \\ \hat{F}(\tilde{\sigma}^{n-i}) \arrow[r, "\rho_{\tilde{\sigma},\tilde{\tau}}"] & \hat{F}(\tilde{\tau}^{n-i-1}) \end{tikzcd}$$

So the same abstract diagram of vector spaces $F : X \to \textbf{\textit{Vect}}$ defines a diagram over $\hat{F} : \tilde{X}^{op} \to \textbf{\textit{Vect}}$, i.e. a co-sheaf on the dual cell structure. Since they are the same diagrams everything about them is the same. $\square$

11.4.1 Invariance under Subdivision

One can ask whether this new invariant is invariant under subdivision. In this section we show that it is invariant for the domain of a map, but is not invariant under subdivision of the target. One can see this latter claim by an earlier example already considered with the pushforward with open supports to a circle.

Theorem 11.31. *Suppose F is a sheaf on X and X' is a subdivision of X, then*

$$(p_{X'})_\dagger F' \cong (p_X)_\dagger F$$

and consequently

$$(Lp_{X'})_\dagger F' \cong (Lp_X)_\dagger F$$

thus sheaf homology is invariant under subdivision. Similarly, the same result should hold for cosheaf cohomology.

Proof. Getting right down to it we see

$$\begin{aligned} s_\dagger s^* F(y) &= \text{colim}\{s^*F(x) | s(x) \leqslant y\} \\ &= \text{colim}\{F(s(x)) | s(x) \leqslant y\} \\ (\text{surjectivity}) &= \text{colim}\{F(x) | x \leqslant y\} \\ (\text{check} - \text{directly}) &= F(y) \end{aligned}$$

and thus

$$(p_{X'})_\dagger s^* F = (p_X)_\dagger s_\dagger s^* F \cong (p_X)_\dagger F.$$

Taking the left derived functors gives the higher result. $\square$

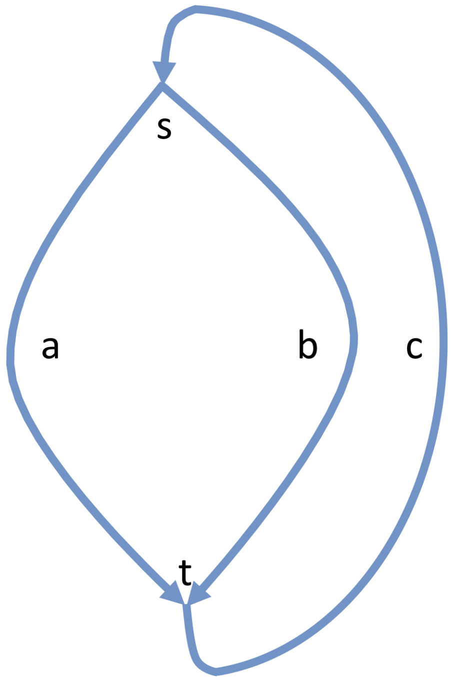

Figure 45: Directed Graph with Labels

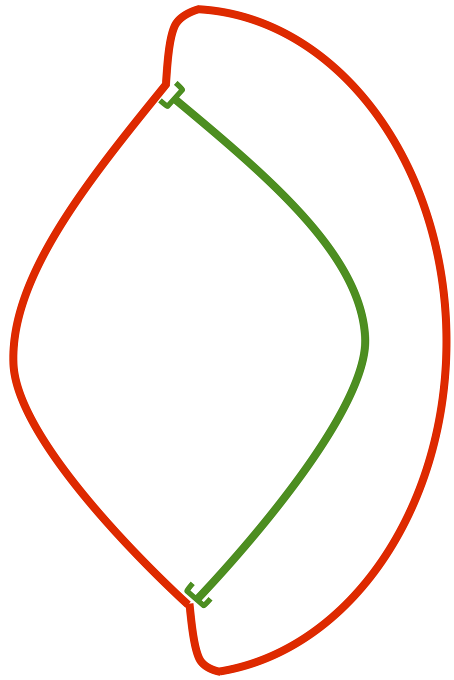

Figure 46: Network Coding Sheaf

11.4.2 Examples of Sheaf Homology over Graphs

One of the virtues of the network coding sheaves in section 9 is that they are easy to construct, have interesting sheaf cohomology, and provide lots of examples. As such, we will use some examples from there to explore this not well understood theory.

Example 11.32. Consider the network coding sheaf implied by figure 45. Viewed as a diagram of vector spaces, it takes the following form:

$$\begin{tikzcd}
 & \mathsf{k}_s \arrow[dl, "1"'] \arrow[dr, "0"] \arrow[drr] & & \\
\mathsf{k}_a & & \mathsf{k}_b & \mathsf{k}_c \\
 & \mathsf{k}_t^2 \arrow[ul, "\pi_1"] \arrow[ur, "\pi_2"'] \arrow[urr, "\pi_1"'] & &
\end{tikzcd}$$

Since the category of complexes of sheaves is additive, we can consider each indecomposable sheaf separately and compute its sheaf homology. If one considers just the red loop as a constant sheaf (barcode) R, it takes the following form:

$$\begin{tikzcd}
& \mathsf{k}_s \arrow[dl, "1"'] \arrow[dr] \arrow[drr] & & \\
\mathsf{k}_a & & 0 & \mathsf{k}_c \\
& \mathsf{k}_t \arrow[ul, "1"] \arrow[ur] \arrow[urr, "1"'] & &
\end{tikzcd}$$

A projective sheaf that surjects onto R is supported on the star of s and t respectively, i.e. $\mathsf{P}_0 := \{\mathsf{s}\} \oplus \{\mathsf{t}\}$:

$$\begin{tikzcd}
& \mathsf{k}_s \arrow[dl, "1"'] \arrow[dr] \arrow[drr] & & \\
\mathsf{k}_s \oplus \mathsf{k}_t & & \mathsf{k}_s \oplus \mathsf{k}_t & \mathsf{k}_s \oplus \mathsf{k}_t \\
& \mathsf{k}_t \arrow[ul, "1"] \arrow[ur] \arrow[urr, "1"'] & &
\end{tikzcd}$$

The kernel sheaf of the natural transformation $\mathsf{P}_0 \Rightarrow \mathsf{R}$ is also projective, which we call P_1 and finishes the projective replacement of the sheaf R.

$$\begin{tikzcd}
& 0 \arrow[dl] \arrow[dr] \arrow[drr] & & \\
{[1\ -1]} & & \mathsf{k}_s \oplus \mathsf{k}_t & {[1\ -1]} \\
& 0 \arrow[ul] \arrow[ur] \arrow[urr] & &
\end{tikzcd}$$

If we take the colimit of P_1 and P_0 separately, the sheaf map $\mathsf{P}_1 \to \mathsf{P}_0$ induces a map on colimits that defines the boundary in the chain complex computing sheaf homology:

$$\partial_1 : \mathsf{k}^4 \to \mathsf{k}^2 \qquad \begin{bmatrix} 1 & 1 & 0 & 1 \\ -1 & 0 & 1 & -1 \end{bmatrix} \quad \Rightarrow \quad \mathsf{H}_0(\mathsf{X};\mathsf{R}) = 0 \quad \mathsf{H}_1(\mathsf{X};\mathsf{R}) \cong \mathsf{k}^2$$

Repeating the same reasoning for the green barcode G yields homology groups $\mathsf{H}_0(\mathsf{X};\mathsf{G}) = 0$ and $\mathsf{H}_1(\mathsf{X};\mathsf{G}) \cong \mathsf{k}$. Since our original network coding sheaf F is a direct sum $\mathsf{R} \oplus \mathsf{G}$ we obtain that the sheaf homology of the sheaf in figure 45 is

$$\mathsf{H}_0(\mathsf{X};\mathsf{F}) = 0 \qquad \mathsf{H}_1(\mathsf{X};\mathsf{F}) \cong \mathsf{k}^3.$$

Exercise 11.33. As an exercise, and to indicate the sensitivity of sheaf homology to its embedding, we ask the reader to verify that the sheaf homology groups of the sheaf in figure 46 are

$$\mathsf{H}_0(\mathsf{X};\mathsf{F}) = 0 \qquad \mathsf{H}_1(\mathsf{X};\mathsf{F}) \cong \mathsf{k}^8.$$

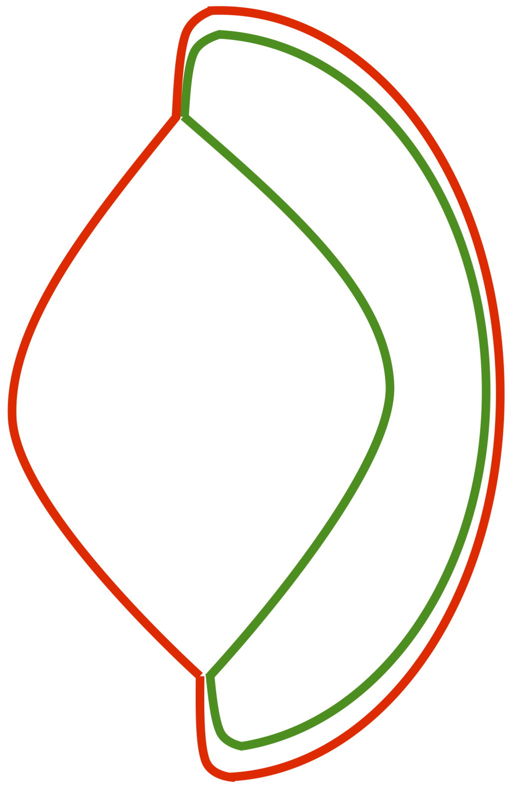

Figure 47: Network Coding Sheaf with Two Decoding Wires

12 Duality: Exchange of Sheaves and Cosheaves

In this section, we are concerned with the derived equivalence of cellular sheaves and cosheaves. In section 12.1, we introduce the functor that establishes this equivalence and try to motivate it topologically via taking the "closure" of the data over an open cell. In the case when X is a manifold, theorem 12.3 gives us a duality result for data that relates sheaf cohomology with our new theory of sheaf homology. Finally, the equivalence is proved in section 12.2.

12.1 Taking Closures and Classical Dualities Re-Obtained

In this section we are going to explain the all-important Poincaré-Verdier duality as an exchange of sheaves and cosheaves. To introduce this duality, we explain an odd, but clean way of going from a cellular sheaf to a complex of cellular cosheaves. This is meant to express the idea that duality is an exchange of open and closed cells.

Suppose we start with a sheaf F on the unit interval $\mathsf{X} = [0, 1]$ stratified with end points

$x = 0$, $y = 1$, and $a = (0, 1)$. Such a sheaf is just a diagram of vector spaces of the form

$$\begin{tikzcd} & F(a) & \\ F(x) \arrow[ur, "\rho_{a,x}"] & & F(y). \arrow[ul, "\rho_{a,y}"'] \end{tikzcd}$$

Now we are going to extend the value of the sheaf on a cell σ to its closure $\bar{\sigma}$ by defining $\hat{F}(\tau) = F(\sigma)$ for every cell $\tau \leqslant \sigma$ and using the identity maps from σ to its faces. This in effect smears the value of the sheaf on an open cell onto all of its faces. However, what should we do to the values of the sheaf already stored on a face τ? This is where we use the different slots in a complex of vector space to store independently the values:

$$\begin{tikzcd} F(a) & F(a) \arrow[l, "\mathrm{id}"'] \arrow[r, "\mathrm{id}"] & F(a) \\ F(x) \arrow[u, "\rho_{a,x}"] & 0 \arrow[l] \arrow[r] \arrow[u] & F(y). \arrow[u, "\rho_{a,y}"'] \end{tikzcd}$$

For dimension reasons, it should be clear that this smearing operation defines an assignment of chain complexes to each open cell with chain maps extending to the faces:

$$\begin{tikzcd} & \widehat{\mathcal{P}}(F)(a) \arrow[dl, "r^{\bullet}_{x,a}"'] \arrow[dr, "r^{\bullet}_{y,a}"] & \\ \widehat{\mathcal{P}}(F)(x) & & \widehat{\mathcal{P}}(F)(y). \end{tikzcd}$$

This motivates the following general definition of a functor $\widehat{\mathcal{P}}$: to a cellular sheaf $F \in \mathbf{Shv}(X)$ we associate the following cosheaf of chain complexes $\widehat{\mathcal{P}}(F)$

$$\widehat{\mathcal{P}}(F): \qquad \sigma \qquad \rightsquigarrow \qquad F(\sigma) \to \bigoplus_{\sigma \leqslant_1 \tau} F(\tau) \to \bigoplus_{\sigma \leqslant_2 \gamma} F(\gamma) \to \cdots .$$

where $F(\sigma)$ is placed in cohomological degree $\dim|\sigma|$ or homological degree $-\dim|\sigma|$ However, in order for this to be a chain complex, following two arrows in sequence should give zero. In order to guarantee this we need to use the fact that X is a cell complex, and as such for any pair of cells $\sigma \leqslant_2 \gamma$ differing in dimension by two, there are precisely two ways τ_1, τ_2 of going between σ and γ. Using the signed incidence relations $[\sigma : \tau_i]$ and the restrictions maps internal to F allows us to define the differentials in this complex by $d^{i+1} := \oplus[\tau : \gamma]\rho^F_{\gamma,\tau}$. Now let's consider a cell λ that is a codimension one face of σ, then the extension map $r^{\bullet}_{\lambda,\sigma}$ is defined to be the chain map

$$\begin{tikzcd} 0 \arrow[r] \arrow[d, "r^{i-1}_{\lambda,\tau}"'] & F(\sigma) \arrow[r, "d^i"] \arrow[d, "r^i_{\lambda,\tau}"'] & \bigoplus_{\sigma \leqslant_1 \tau} F(\tau) \arrow[r, "d^{i+1}"] \arrow[d, "r^{i+1}_{\lambda,\tau}"'] & \bigoplus_{\sigma \leqslant_2 \gamma} F(\gamma) . \arrow[d, "r^{i+2}_{\lambda,\tau}"'] \\ F(\lambda) \arrow[r, "d^{i-1}"'] & \bigoplus_{\lambda \leqslant_1 \sigma} F(\sigma) \arrow[r, "d^i"'] & \bigoplus_{\lambda \leqslant_2 \tau} F(\tau) \arrow[r, "d^{i+1}"'] & \bigoplus_{\lambda \leqslant_3 \gamma} F(\gamma) \end{tikzcd}$$

The reason it is a chain map is clear from the fact that if $\lambda \leqslant \sigma$ then $U_\sigma \subset U_\lambda$ and so the chain complex $\widehat{\mathcal{P}}(F)(\sigma)$ simply includes term by term into the chain complex $\widehat{\mathcal{P}}(F)(\lambda)$.

Although the idea of a cosheaf of chain complexes is perhaps easier to visualize, for actual algebraic manipulation, one uses a chain complex of cosheaves to express the same idea in a different way.

Definition 12.1 (Poincaré-Verdier Equivalence Functor)**.** Let X be a cell complex and let $\mathbf{Shv}(X)$ and $\mathbf{CoShv}(X)$ denote the categories of cellular sheaves and cosheaves respectively. We define the **Poincaré-Verdier Equivalence Functor** $\widehat{\mathcal{P}} : D^b(\mathbf{Shv}(X)) \to D^b(\mathbf{CoShv}(X))$ by the following formula: to a sheaf $F \in \mathbf{Shv}(X)$ we associate the following complex of projective co-sheaves, the cohomological degree corresponding to the dimension of the cell:

$$\cdots \longrightarrow \bigoplus_{\sigma^i \in X} [\hat{\sigma}^i]^{F(\sigma^i)} \xrightarrow{[\sigma:\gamma]\rho^F} \bigoplus_{\gamma^{i+1} \in X} [\hat{\gamma}^{i+1}]^{F(\gamma^{i+1})} \xrightarrow{[\gamma:\tau]\rho^F} \bigoplus_{\tau^{i+2} \in X} [\hat{\tau}^{i+2}]^{F(\tau^{i+2})} \longrightarrow \cdots$$

Here σ^i denotes the i-cells and $[\sigma^i : \gamma^{i+1}] = \{0, \pm 1\}$ records whether the cells are incident and whether orientations agree or disagree. The maps in between are to be understood as the matrix $\oplus[\sigma^i : \gamma^{i+1}]\rho^F_{\sigma,\gamma}$.

For a complex of sheaves

$$\begin{tikzcd}
F^i \arrow[d] & \rightsquigarrow & \cdots \arrow[r] & \bigoplus_{\gamma^{j+1} \in X} [\hat{\gamma}^{j+1}]^{F^i(\gamma^{j+1})} \arrow[r, "{[\gamma:\tau]\rho^{F^i}}"'] \arrow[d] & \bigoplus_{\tau^{j+2} \in X} [\hat{\tau}^{j+2}]^{F^i(\tau^{j+2})} \arrow[r] \arrow[d] & \cdots \\
F^{i+1} \arrow[d] & \rightsquigarrow & \cdots \arrow[r] & \bigoplus_{\gamma^{j+1} \in X} [\hat{\gamma}^{j+1}]^{F^{i+1}(\gamma^{j+1})} \arrow[r, "{[\gamma:\tau]\rho^{F^{i+1}}}"'] \arrow[d] & \bigoplus_{\tau^{j+2} \in X} [\hat{\tau}^{j+2}]^{F^{i+1}(\tau^{j+2})} \arrow[r] \arrow[d] & \cdots \\
F^{i+2} & \rightsquigarrow & \cdots \arrow[r] & \bigoplus_{\gamma^{j+1} \in X} [\hat{\gamma}^{j+1}]^{F^{i+2}(\gamma^{j+1})} \arrow[r, "{[\gamma:\tau]\rho^{F^{i+2}}}"'] & \bigoplus_{\tau^{j+2} \in X} [\hat{\tau}^{j+2}]^{F^{i+2}(\tau^{j+2})} \arrow[r] & \cdots
\end{tikzcd}$$

where we then pass to the totalization.

Before discussing why this functor is an equivalence, let us deduce a few computational consequences of this functor.

Theorem 12.2. *If F is a cell sheaf on a cell complex X, then*

$$H^i_c(X; F) \cong H_{-i}(X; \widehat{\mathcal{P}}(F)).$$

Proof. First we apply the equivalence functor $\widehat{\mathcal{P}}$ to F

$$0 \longrightarrow \bigoplus_{v \in X} [\hat{v}]^{F(v)} \xrightarrow[{[v:e]\rho e,v}]{} \bigoplus_{e \in X} [\hat{e}]^{F(e)} \xrightarrow[{[e:\sigma]\rho_{\sigma,e}}]{} \bigoplus_{\sigma \in X} [\hat{\sigma}]^{F(\sigma)} \longrightarrow \cdots$$

Taking colimits (pushing forward to a point) term by term produces the complex of vector spaces

$$0 \longrightarrow \bigoplus_{v \in X} F(v) \xrightarrow[{[v:e]\rho e,v}]{} \bigoplus_{e \in X} F(e) \xrightarrow[{[e:\sigma]\rho_{\sigma,e}}]{} \bigoplus_{\sigma \in X} F(\sigma) \longrightarrow \cdots$$

which the reader should recognize as being the computational formula for computing compactly supported sheaf cohomology. $\square$

Now let us give a simple proof of the standard Poincaré duality statement on a manifold X with coefficients in an arbitrary cell sheaf F, except this time the sheaf homology groups are used.

Theorem 12.3. *Suppose* F *is a cell sheaf on a cell complex* X *that happens to be a compact manifold (so it has a dual cell structure* $\hat{X}$*), then*

$$H^i(X;F) \cong H_{n-i}(X;F).$$

Where the group on the right is not just notational, but it indicates the left-derived functors of $p_\dagger$ *on sheaves.*

Proof. We repeat the first step of the proof of the previous theorem. By feeding F through the equivalence $\widehat{\mathcal{P}}$ we get a complex of cosheaves. Pushing forward to a point yields a complex whose (co)homology is the compactly supported cohomology of the sheaf F. Now we recognize that the formula yields a formula for the Borel-Moore homology for the cosheaf naturally defined on the dual cell structure.

$$\begin{tikzcd} 0 \arrow[r] & \bigoplus_{v\in X} F(v) \arrow[r, "{[v:e]\rho_{e,v}}"'] \arrow[d, squiggly] & \bigoplus_{e\in X} F(e) \arrow[r, "{[e:\sigma]\rho_{\sigma,e}}"'] \arrow[d, squiggly] & \bigoplus_{\sigma\in X} F(\sigma) \arrow[r] \arrow[d, squiggly] & \cdots \\ 0 \arrow[r] & \bigoplus_{\tilde{v}\in\tilde{X}} \hat{F}(\tilde{v}) \arrow[r, "{[\tilde{v}:\tilde{e}]\rho_{\tilde{e},\tilde{v}}}"'] & \bigoplus_{\tilde{e}\in\tilde{X}} \hat{F}(\tilde{e}) \arrow[r, "{[\tilde{e}:\tilde{\sigma}]\rho_{\tilde{\sigma},\tilde{e}}}"'] & \bigoplus_{\tilde{\sigma}\in\tilde{X}} \hat{F}(\tilde{\sigma}) \arrow[r] & \cdots \end{tikzcd}$$

Taking the homology of the bottom row is the usual formula for the Borel-Moore homology of a cellular cosheaf except the top dimensional cells are place in degree 0, the $n-1$ cells in degree -1, and so on. Everything being shifted by $n = \dim X$ we get the isomorphism

$$H_{-i}(X;\widehat{\mathcal{P}}(F)) \cong H^{BM}_{n-i}(\tilde{X};\hat{F}).$$

However, we already observed in theorem 12.3 that the diagrams $\hat{F}$ on $\tilde{X}$ and F on X are the same in every possible way, so in particular sheaf homology of F must coincide with cosheaf homology of $\hat{F}$. Thus using compactness to drop the Borel-Moore label and chaining together the previous theorem we get

$$H^i(X;F) \cong H_{-i}(X;\widehat{\mathcal{P}}(F)) \cong H_{n-i}(\tilde{X};\hat{F}) \cong H_{n-i}(X;F).$$

□

12.2 Derived Equivalence of Sheaves and Cosheaves

Historically, the derived equivalence of cellular sheaves and cosheaves appears in a few places and is re-discovered again and again. In chronological order, the first published proof appears to be in the 1998 paper of Peter Schneider in "Verdier Duality on the Building" [75], which is a follow-up of a longer paper connecting sheaves, buildings and representation theory [74]. Unfortunately, Schneider uses the term "local coefficient systems" to mean what we mean by cellular cosheaves. At around the same time Maxim Vybornov made explicit mention of the relationship between sheaves and cosheaves, relating them through Koszul duality [92], but it took up until 2005 for Kohji Yanagawa to explicitly state that Vybornov's work implied the derived equivalence of sheaves and cosheaves [99].

However, the perspective presented here was arrived at independently of the above work. In early March 2012, Bob MacPherson gave a lecture (which the author attended) where he conjectured that the derived category of cellular sheaves and cosheaves should be equivalent. Within a few weeks the author produced a proof. After some truly insightful comments from David Lipsky, the equivalence was refined to its current form.

Although the ideas were foreshadowed by many sources, the use of stalk (co)sheaves appears to be a novel way of arguing.

Theorem 12.4 (Equivalence)**.** $\widehat{\mathcal{P}} : \mathsf{D}^b(\mathbf{Shv}(X)) \to \mathsf{D}^b(\mathbf{CoShv}(X))$ *is an equivalence.*

Proof. First let us point out that the functor $\widehat{\mathcal{P}}$ really is a functor. Indeed if $\alpha : F \to G$ is a map of sheaves then we have maps $\alpha(\sigma) : F(\sigma) \to G(\sigma)$ that commute with the respective restriction maps ρ^F and ρ^G. As a result, we get maps $[\hat{\sigma}]^{F(\sigma)} \to [\hat{\sigma}]^{G(\sigma)}$. Moreover, these maps respect the differentials in $\widehat{\mathcal{P}}(F)$ and $\widehat{\mathcal{P}}(G)$, so we get a chain map. It is clearly additive, i.e. for maps $\alpha, \beta : F \to G$ $\widehat{\mathcal{P}}(\alpha+\beta) = P(\alpha)+P(\beta)$. This implies that $\widehat{\mathcal{P}}$ preserves homotopies.

It is also clear that $\widehat{\mathcal{P}}$ preserves quasi-isomorphisms. Note that a sequence of cellular sheaves $A^\bullet$ is exact if and only if $A^\bullet(\sigma)$ is an exact sequence of vector spaces for every $\sigma \in X$. This implies that $\widehat{\mathcal{P}}(A^\bullet)$ is a double-complex with exact rows. By standard results surrounding the theory of spectral sequences or by the acyclic assembly lemma ([93] Lem. 2.7.3) we get that the totalization is exact.

Let us understand what this functor does to an elementary injective sheaf $[\sigma]^V$. Applying the definition we can see that

$$\widehat{\mathcal{P}} : [\sigma]^V \qquad \rightsquigarrow \qquad \bigoplus_{\tau^0 \subset \sigma} [\hat{\tau}^0]^V \longrightarrow \cdots \longrightarrow \bigoplus_{\tau^i \subset \sigma} [\hat{\tau}^i]^V \longrightarrow \cdots \longrightarrow [\hat{\sigma}]^V$$

which is nothing other than the projective cosheaf resolution of the skyscraper (or stalk) cosheaf $\hat{S}^V_\sigma$ supported on σ, i.e.

$$\hat{S}^V_\sigma(\tau) = \begin{cases} V & \sigma = \tau \\ 0 & \text{o.w.} \end{cases}$$

Consequently, there is a quasi-isomorphism $q : \widehat{\mathcal{P}}([\sigma]^V) \to \hat{S}^V_\sigma[-\dim \sigma]$ where $\hat{S}^V_\sigma$ is placed in degree equal to the dimension of σ assuming that $[\sigma]^V$ is initially in degree 0. By abusing notation slightly and letting $\mathcal{P}$ send cosheaves to sheaves, we see that

$$\mathcal{P}(q) : \mathcal{P}\widehat{\mathcal{P}}([\sigma]^V) \to \mathcal{P}(S^V_\sigma) = [\sigma]^V$$

and thus we can define a natural transformation from $\mathcal{P}\widehat{\mathcal{P}}$ to $\mathrm{id}_{\mathsf{D}^b(\mathbf{Shv})}$ when restricted to elementary injectives. However, by lemma **??** we know that every injective looks like such a sum, so this works for injective sheaves concentrated in a single degree. However, it is clear that $\widehat{\mathcal{P}}$ sends a complex of injectives, before taking the totalization of the double complex to the projective resolutions of a complex of skyscraper cosheaves. Applying $\widehat{\mathcal{P}}$ to the quasi-isomorphism relating the double complex of projective cosheaves to the complex of skyscrapers, extends the natural transformation to the whole derived category. However, since $\widehat{\mathcal{P}}$ preserves quasi-isomorphisms, this natural transformation is in fact an equivalence. This shows $\mathcal{P}\widehat{\mathcal{P}} \cong \mathrm{id}_{D^b(\boldsymbol{Shv})}$. Repeating the argument starting from co-sheaves shows that

$$\widehat{\mathcal{P}} : D^b(\boldsymbol{Shv}(X)) \leftrightarrow D^b(\boldsymbol{CoShv}(X)) : \mathcal{P}$$

is an adjoint equivalence of categories. □

The above proof should be taken as the primary duality result from which other dualities spring. This was not always appreciated and the author's first attack on the proof was to chain together two well-known dualities, which we review in the next two sections.

12.2.1 Linear Duality

There is an endofunctor on the category of finite dimensional vector spaces $\boldsymbol{vect}$ given by sending a vector space to its dual $V \rightsquigarrow V^*$. This functor has the effect of taking a cellular sheaf (F, ρ) to a cellular co-sheaf (F^*, ρ^*), since the restriction maps get dualized into extension maps. It is contravariant since a sheaf morphism $F \to G$ gets sent to a co-sheaf morphism in the opposite direction $F^* \leftarrow G^*$ as one can easily check. We can promote this functor to the derived category, using a subscript f to remind the reader when we restrict to the finite dimensional full subcategories.

Definition 12.5 (Linear Duals). Define $\hat{V} : D^b(\boldsymbol{Shv}(X))^{op} \to D^b(\boldsymbol{CoShv}(X))$ as follows

- $\hat{V}(F^\bullet) = (F^*)^{-\bullet}$, i.e. take a sheaf in slot i, dualize its internal restriction maps $\rho^{F^i}_{\sigma,\tau}$ to extension maps $r^{F^{i*}}_{\tau,\sigma}$ to obtain a co-sheaf and then put it in slot $-i$.
- $\hat{V}$ sends differentials between sheaves d^i to their adjoints in negative degree $\partial^{-i-1} := (d^i)^*$

$$\star(\cdots \longrightarrow F^i \xrightarrow{d^i} F^{i+1} \longrightarrow \cdots) \;=\; \cdots \longrightarrow [(F^{i+1})^*]^{-i-1} \xrightarrow[\partial^{-i-1}]{} [(F^i)^*]^{-i} \longrightarrow \cdots$$

We'll adopt the convention that lowering the index increases the degree $\partial^{-i-1} \to \partial_i$.

We will reserve the right to abuse notation and let V map from co-sheaves to sheaves in the obvious manner, i.e. $V : D^b(\boldsymbol{CoShv}(X))^{op} \to D^b(\boldsymbol{Shv}(X))$ or formally equivalent $V : D^b(\boldsymbol{CoShv}(X)) \to D^b(\boldsymbol{Shv}(X))^{op}$.

Lemma 12.6. $\hat{V}_f : D^b(\boldsymbol{Shv}_f(X)) \to D^b(\boldsymbol{CoShv}_f(X))^{op}$ *is an equivalence of categories.*

Proof. It is clear that if $\alpha : I^\bullet \to J^\bullet$ is a map in the category of complexes of sheaves homotopic to zero $\alpha \simeq 0$, i.e. there exists a map $h : I^\bullet \to J^{\bullet-1}$, written $h : I \to J[-1]$ such that $\alpha^n - 0^n = d_J^{n-1}h^n + h^{n+1}d_I^n$. Writing out how $\hat{V}$ acts carefully we see that $\hat{V}(\alpha) : \hat{V}(J) \to \hat{V}(I)$ and $\hat{V}(h) : V(J[-1]) = \hat{V}(J)[+1] \to \hat{V}(I)$ defines a homotopy between $\hat{V}(\alpha)$ and $\hat{V}(0) = 0$ by setting $(h^*)^\bullet = \hat{V}(h)^{\bullet-1}$.

$\hat{V}$ thus sends $K^b(\mathrm{Inj}-S_f)^{op}$ to $K^b(\mathrm{Proj}-C_f)$ and composed twice $V\hat{V} : K^b(\mathrm{Inj}-S_f) \to K^b(\mathrm{Inj}-S_f)$ is naturally isomorphic to the identity functor, so it is an equivalence. We can repeat the arguments for co-sheaves and use formality to put the ^{op} where we want. $\square$

12.2.2 Verdier Dual Anti-Involution

Definition 12.7 (Verdier Dual)**.** The Verdier dual functor $D : D(\mathbf{Shv}_f(X)) \to D(\mathbf{Shv}_f(X))^{op}$ is defined as $D := \mathcal{H}om(-, \omega_X^\bullet)$. Recall that $\mathcal{H}om(F, G)$ is a sheaf whose value on a cell σ is given by $\mathbf{Hom}(F|_{st(\sigma)}, G|_{st(\sigma)})$, i.e. natural transformations between the restrictions to the star of σ.

The complex of injective sheaves $\omega_X^\bullet$ is called the dualizing complex of X. It has in negative degree ω_X^{-i} the sum over the one-dimensional elementary injectives concentrated on i-cells $[\gamma^i]$. The maps between use the orientations on cells to guarantee it is a complex.

$$\cdots \longrightarrow \oplus_{|\tau|=i+1}[\tau] \xrightarrow{\oplus[\gamma:\tau]} \oplus_{|\gamma|=i}[\gamma] \xrightarrow{\oplus[\sigma:\gamma]} \oplus_{|\sigma|=i-1}[\sigma] \longrightarrow \cdots$$

The Verdier dual of F is the complex of sheaves $D^\bullet F := \mathcal{H}om(F, \omega_X^\bullet)$. Written out explicitly it is

$$\cdots \longrightarrow \oplus_{|\tau|=i+1}[\tau]^{F(\tau)^*} \xrightarrow[\oplus[\gamma:\tau]\rho^*]{} \oplus_{|\gamma|=i}[\gamma]^{F(\gamma)^*} \xrightarrow[\oplus[\sigma:\gamma]\rho^*]{} \oplus_{|\sigma|=i-1}[\sigma]^{F(\sigma)^*} \longrightarrow \cdots$$

Proposition 12.8. *The functor* $\widehat{\mathcal{P}} : D^b(\mathbf{Shv}_f(X)) \to D^b(\mathbf{CoShv}_f(X))$ *composed with linear duality* $V : D^b(\mathbf{CoShv}_f(X)) \to D^b(\mathbf{Shv}_f(X))^{op}$ *gives the Verdier dual anti-equivalence, i.e.* $D \cong V\widehat{\mathcal{P}}$.

Proof. Just check by hand. $\square$

Remark 12.9. We could have used well-known facts about Verdier duality to prove a weaker version of our main theorem by restricting to finitely generated stalks.

13 Cosheaves as Valuations on Sheaves

The development of cosheaves as a theory is largely fragmented. Researchers at different points in time have found a use for it here and there, at the service of different purposes and interests. The more strongly categorical and logical community have done some considerable work understanding the relationship between the topos of sheaves and cosheaves. One insight that seems very worthwhile is that cosheaves act on sheaves in a natural way. Although one can use a little bit of category theory to draw this conclusion, we use this to give some surprising reformulations of classical sheaf theory. Namely, the primary observation of this section is that the action of taking compactly supported cohomology of a sheaf can be interpreted as an action of a very particular cosheaf on the category of all sheaves.

13.1 Left and Right Modules and Tensor Products

Suppose R is a ring with unit 1_R. One can think of R as a category with a single object $\star$ whose set of morphisms

$$\mathbf{Hom}_R(\star,\star) \cong R$$

has the structure of an abelian group. The multiplication in the ring plays the role of a composition so $r \cdot s = r \circ s$. The abelian group structure, which corresponds to the ability to add morphisms $r+s$, reflects the fact that rings have an underlying abelian group structure. One says that R is a pre-additive category, or is a category enriched in $\mathbf{Ab}$ – the category of abelian groups.

An additive functor $B : R \to \mathbf{Ab}$ is a functor that preserves the abelian group structure, so it picks out a single abelian group, which we also call B, and satisfies the relation $(r+s)\cdot B = r \cdot B + s \cdot B$ and $(rs)\cdot B = r\cdot(s\cdot B)$ so such a functor is precisely the data of a **left R-module**. Dually, a contravariant functor $A : R^{op} \to \mathbf{Ab}$ prescribes the data of a **right R-module**. Taking the tensor product over $\mathbb{Z}$ of A and B allows us to define a bi-module

$$A \otimes B : R^{op} \times R \to \mathbf{Ab} \qquad (\star,\star) \mapsto A \otimes_{\mathbb{Z}} B.$$

The latter is the group freely generated by pairs of elements from A and B modulo the usual relations $(a+a')\otimes b = a\otimes b + a'\otimes b$ and $a \otimes (b+b') = a\otimes b + a \otimes b'$. However, in the presence of the action of a ring R, there is another tensor product $A \otimes_R B$ that further quotients $A \otimes_{\mathbb{Z}} B$ by the relation $(a\cdot r)\otimes b = a \otimes (r \cdot b)$. Said using diagrams, we require that for each r, the following diagram commutes.

$$\begin{tikzcd}
A \otimes B \arrow[r, "1_A \otimes B(r)"] \arrow[d, "A(r)\otimes 1_B"'] & A \otimes B \arrow[d] \\
A \otimes B \arrow[r] & A \otimes_R B
\end{tikzcd}$$

In other words there is a coequalizer

$$\begin{tikzcd}
A \otimes_{\mathbb{Z}} R \otimes_{\mathbb{Z}} B \arrow[r, shift left, "(a,r,b)\mapsto(a,rb)"] \arrow[r, shift right, "(a,r,b)\mapsto(ar,b)"'] & A \otimes_{\mathbb{Z}} B \arrow[r] & A \otimes_R B
\end{tikzcd}$$

that realizes the tensor product using purely categorical operations. This allows us to work in a greater degree of generality by making use of a special type of colimit called a **coend**, that generalizes the tensor product described above.

Definition 13.1 (Tensoring Sheaves with Cosheaves)**.** Let X be a topological space and let $\hat{G}$ and F be a pre-cosheaf and a pre-sheaf respectively, both valued in $\mathbf{Vect}$. Note that for every pair of objects $U \to V$ in $\mathbf{Open}(X)$ we have a diagram

$$\begin{tikzcd} & \hat{G}(V) \otimes F(V) \\ \hat{G}(U) \otimes F(V) \arrow[ur, "r^{G}_{V,U} \otimes \mathrm{id}"] \arrow[dr, "\mathrm{id} \otimes \rho^{F}_{U,V}"'] & \\ & \hat{G}(U) \otimes F(U) \end{tikzcd}$$

which is the building block in defining the **coend** or **tensor product over** X

$$\bigoplus_{U \to V} \hat{G}(U) \otimes F(V) \rightrightarrows \bigoplus_{W} \hat{G}(W) \otimes F(W) \to \int^{\mathbf{Open}(X)} \hat{G}(W) \otimes F(W) =: \hat{G} \otimes_X F.$$

We illustrate this definition with an immediate example.

Example 13.2 (Stalks and Skyscraper Cosheaf)**.** Recall that we defined the *skyscraper cosheaf* at x to be the cosheaf

$$\hat{S}_x(U) = \begin{cases} k & \text{if } x \in U \\ 0 & \text{other wise} \end{cases}$$

With some thought one can show that the tensor product of any pre-sheaf F with the cosheaf $\hat{S}_x$ yields

$$\hat{S}_x \otimes_X F \cong F_x$$

by treating F as a variable which can range over all pre-sheaves, one gets, in particular, a functor

$$\hat{S}_x \otimes_X - : \mathbf{Shv}(X) \to \mathbf{Vect} \qquad F \rightsquigarrow F_x.$$

The previous example demonstrates an important observation: *The operation of taking stalks is equivalent to the process of tensoring with the skyscraper cosheaf.*

To see how far this observation can be generalized, note that if we fix $\hat{G}$ and let F vary then we get a functor

$$\hat{G} \otimes_X - : \mathbf{Shv}(X) \to \mathbf{Vect}$$

that is defined in terms of colimits and is thus co-continuous (it sends colimits to colimits). Now we are free to take an arbitrary cosheaf and let it act on sheaves. The "one obvious choice" of taking stalks at a point is run over by a veritable slew of valuations, one for each cosheaf. Moreover, it is clear that this description extends to a pairing between the symmetric monoidal categories $\mathbf{CoShv}(X)$ and $\mathbf{Shv}(X)$, i.e.

$$-\otimes_X - : \mathbf{CoShv}(X) \times \mathbf{Shv}(X) \to \mathbf{Vect} \qquad (\hat{G}, F) \mapsto \hat{G} \otimes_X F := \int^{\mathbf{Open}(X)} \hat{G}(U) \otimes F(U),$$

although we haven't used the sheaf or cosheaf axiom anywhere, so the pairing is actually valid for pre-sheaves and pre-cosheaves.

13.2 Compactly-Supported Cohomology

Although the idea of using coends to tensor together sheaves and cosheaves has been independently re-discovered many times, cf. Jean-Pierre Schneider's 1987 work [76], it has not been used to do any serious work. This is a shame in light of the following 1985 theorem of A.M. Pitts [68].

Theorem 13.3. *Let X be any topological space. Every colimit-preserving functor on sheaves arises by tensoring with a cosheaf, i.e.*

$$\mathbf{CoShv}(X; \mathbf{Set}) \cong \mathbf{Fun}^{co-cts}(\mathbf{Shv}(X; \mathbf{Set}), \mathbf{Set}).$$

This theorem is also stated in Marta Bunge and Jonathan Funk's 2006 book "Singular Coverings of Toposes" [18] as theorem 1.4.3, which further surveys some of Lawvere's philosophy of distributions on topoi. The topos community deserves commendation for keeping the study of cosheaves alive during the past few decades, but so far work in the enriched and computable setting of vector spaces is largely missing.

We attempt to partly remedy this gap by establishing a connection between the tensor operation and the cohomology of sheaves. However, instead of establishing an enriched version of Pitt's theorem,[43] we will use it as a guide. For example, in classical sheaf theory, compactly supported cohomology is gotten by taking the constant map $p : X \to \star$ and associating to it the pushforward with compact supports functor $p_! : \mathbf{Shv}(X) \to \mathbf{Shv}(\star) \cong \mathbf{Vect}$. Of course, just applying $p_!$ defines only compactly supported zeroth cohomology of a sheaf $H^0_c(X; F)$. To get the higher compactly supported cohomology groups one takes an injective resolution and applies $p_!$ to the resolution. The result will be a complex of vector spaces, whose cohomology in turn produces the desired groups:

$$F \to I^\bullet \rightsquigarrow Rp_! := p_! I^\bullet \qquad H^i(p_! I^\bullet) := H^i_c(X; F).$$

Historically the first fundamental duality result in sheaf theory was the statement that $Rp_!$ admits a right adjoint on the level of the derived category. This adjunction is sometimes called **global Verdier duality**:

$$\mathbf{Hom}(Rp_! F, G) \cong \mathbf{Hom}(F, p^! G).$$

[43] We delay this for another paper.

By applying the fact that left adjoints are co-continuous one is led to believe, in light of Pitt's theorem, that there should be a cosheaf that realizes the operation of taking derived pushforward with compact supports.

In light of the derived equivalence between cellular sheaves and cosheaves established in this paper, we provide an explicit description of the complex of cosheaves that realizes the derived pushforward.

In preparation, one should note that there are several cosheaves that realize the operation of taking stalks at a point x in the cellular world. One is

$$\hat{\delta}_{\sigma}(\tau) = \begin{cases} k & \text{if } \sigma = \tau \\ 0 & \text{o.w.} \end{cases}$$

The other is the correct formulation of $\hat{S}_{\sigma}$ when using the Alexandrov topology

$$[\hat{\sigma}](\tau) = \begin{cases} k & \text{if } \tau \leqslant \sigma \\ 0 & \text{o.w.} \end{cases}$$

Recall that this is also the elementary projective cosheaf concentrated on σ with value k.

Observe that the first cosheaf returns the value $F(\sigma)$ because every other cell is tensored with zero. The second cosheaf works by restricting the non-zero values of F to the closure of the cell σ, but this restricted diagram has a terminal object given by $F(\sigma)$, so the colimit returns $F(\sigma)$ as well.

This allows us to state the main theorem of this section.

Theorem 13.4. *Let X be a cell complex, then the operation $Rp_! : \mathbf{Shv}(X) \to \mathbf{Vect}$ on cellular sheaves is equivalent to tensoring with the image of the constant sheaf through the derived equivalence, i.e.*

$$\widehat{\mathcal{P}}(k_X) = \bigoplus_{v \in X} [\hat{v}] \to \bigoplus_{e \in X} [\hat{e}] \to \bigoplus_{\sigma \in X} [\hat{\sigma}] \to \cdots .$$

Proof. The proof is immediate given the previous description of taking stalks, i.e. one can check directly the formula

$$\widehat{\mathcal{P}}(k_X) \otimes_X F \cong \bigoplus_{v \in X} F(v) \to \bigoplus_{e \in X} F(e) \to \bigoplus_{\sigma \in X} F(\sigma) \to \cdots$$

whose cohomology is by definition the compactly supported cohomology of a cellular sheaf F. □

This perspective is especially satisfying for the following reason: it makes transparent how the underlying topology of the space X is coupled with the cohomology of a sheaf F. Compactly supported sheaf cohomology arises by tensoring with the complex of cosheaves that computes the Borel-Moore homology of the underlying space.

13.3 Sheaf Homology and Future Directions

The perspective of tensoring sheaves and cosheaves together offers numerous directions for further research both in pure and applied sheaf theory. Just the heuristic that

each cosheaf determines a (co-)continuous valuation on the category of sheaves,

is suggestive of the idea that if we are going to use sheaves to model the world, then cosheaves should allow us to weight different models of the world.

After having recovered some classical operations on sheaves, we are left with many more to consider. For example the constant cosheaf $\hat{\mathsf{k}}_X$ should act on sheaves by returning its colimit, i.e. zeroth sheaf homology

$$\hat{\mathsf{k}}_X \otimes_X - : \boldsymbol{Shv}(X) \to \boldsymbol{Vect} \qquad F \rightsquigarrow H_0(X; F) = p_\dagger F.$$

By taking a projective resolution of the constant cosheaf once and for all, one then gets for free a way of computing **higher sheaf homology**. This yet-to-be-explored theory has only recently found its use in applications, e.g. the work of Sanjeevi Krishnan on max-flow min-cut.

Additionally, the decategorification of the pairing of the categories of constructible sheaves and cosheaves provides an alternative approach to the study of Euler integration and leads in a natural way to the study of **higher Euler calculus** through higher K-theory. More directly the operation of pairing sheaves and cosheaves is reminiscent of a convolution operation. This area is under active research in collaboration with Aaron Royer.

References

[1] P. Achar. Applications of homological algebra: Introduction to perverse sheaves. Course Notes from Louisiana State University, Spring Available at `https://www.math.lsu.edu/~pramod/tc/07s-7280/`, 2007.

[2] Jiri Adamek, Horst Herrlich, and George E. Strecker. *Abstract and Concrete Categories: The Joy of Cats*. Dover Books on Mathematics. Dover Publications, 2009.

[3] P. Alexandroff. Diskrete räume. *Mat. Sb*, 2:501–518, 1937.

[4] Paul Alexandroff. Über den allgemeinen dimensionsbegriff und seine beziehungen zur elementaren geometrischen anschauung. *Mathematische Annalen*, 98:617–635, 1928.

[5] P. S. Alexandrov. *Combinatorial Topology*. Dover Publications, Inc., 31 East 2nd Street, Mineola, NY, 1947.

[6] Paolo Aluffi. *Algebra. Chapter 0*, volume 104. Graduate Studies in Mathematics American Mathematical Society, 2009.

[7] G. Arone and M. Kankaanrinta. On the functoriality of the blow-up construction. *Bulletin of the Belgian Mathematical Society-Simon Stevin*, 17(5):821–832, 2010.

[8] M. Artin and Harvard University. Dept. of Mathematics. *Grothendieck topologies: notes on a seminar*. Harvard University, Dept. of Mathematics, 1962.

[9] Michael Artin, Alexander Grothendieck, and Jean-Louis Verdier. *Théorie des topos et cohomologie étale des schémas. Tome 1: Théorie des topos*. Lecture Notes in Mathematics, Vol. 269. Springer-Verlag, Berlin, 1972. Séminaire de Géométrie Algébrique du Bois-Marie 1963–1964 (SGA 4), Dirigé par M. Artin, A. Grothendieck, et J. L. Verdier. Avec la collaboration de N. Bourbaki, P. Deligne et B. Saint-Donat.

[10] M.F. Atiyah and I.G. MacDonald. *Introduction To Commutative Algebra*. Addison-Wesley series in mathematics. Westview Press, 1969.

[11] Michael Atiyah. On the krull-schmidt theorem with application to sheaves. *Bull. Soc. Math. France*, 84, 1956.

[12] Christian Ausoni, Kathyrn Hess, and Jerome Scherer. *Alpine Perspectives on Algebraic Topology: Third Arolla Conference on Algebraic Topology, August 18-24, 2008, Arolla, Switzerland*, volume 502 of *Contemporary Mathematics*. American Mathematical Society, 2009.

[13] Steve Awodey. *Category Theory*. Oxford University Press, 2nd edition, 2010.

[14] Karol Borsuk. On the imbedding of systems of compacta in simplicial complexes. *Fund. Math.*, 35:217–234, 1948.

[15] Glen Bredon. Cosheaves and homology. *Pacific J. Math.*, 25(1):1–32, 1968.

[16] Glen Bredon. *Sheaf Theory*, volume 170 of *Graduate Texts in Mathematics*. Springer-Verlag, 2nd edition, 1997.

[17] R BROWN. Groupoids and van kampen's theorem. In *Proc. London Math. Soc.(3)*, volume 17, pages 385–401, 1967.

[18] Marta Bunge and Jonathan Funk. *Singular coverings of toposes*, volume 1890 of *Lecture Notes in Mathematics*. Springer-Verlag, 2006.

[19] Dan Burghelea and Tamal Dey. Persistence for circle valued maps. 2011.

[20] Gunnar Carlsson and Vin de Silva. Zigzag persistence. *Foundations of Computational Mathematics*, 10(4):367–405, August 2010.

[21] Gunnar Carlsson, Vin de Silva, and Dmitriy Morozov. Zigzag persistent homology and real-valued functions. *Proceedings of the Annual Symposium on Computational Geometry*, pages 247–256, 2009.

[22] Michel Coste. An introduction to semialgebraic geometry. lecture notes, October 2002.

[23] Justin Curry, Robert Ghrist, and Michael Robinson. Euler calculus with applications to signals and sensing. *Proceedings of Symposia in Applied Mathematics*, 70:75–146, 2012.

[24] Vin de Silva and Robert Ghrist. Coordinate-free coverage in sensor networks with controlled boundaries via homology. *Int. J. Robotics Research*, 25(12):1205–1222, 2006.

[25] Harm Derkden and Jerzy Weyman. Quiver representations. *Notices of the AMS*, 52(2):200–206, Feb 2005.

[26] Daniel Dugger and Daniel Isaksen. Hypercovers in topology. K-theory Preprint Archives, `http://www.math.uiuc.edu/K-theory/0528/`, November 2001.

[27] Mário Edmundo and Luca Prelli. The six grothendieck operations on o-minimal sheaves. Available at `http://www.ciul.ul.pt/~edmundo/6Go_03_10_12.pdf`, 2012.

[28] Samuel Eilenberg and Saunders MacLane. General theory of natural equivalences. *Transactions of the American Mathematical Society*, 58(2):pp. 231–294, 1945.

[29] P. Elias, A. Feinstein, and C. Shannon. A note on the maximum flow through a network. *IRE Transactions on Information Theory Volume*, 2(4):117–119, December 1956.

[30] P. Etingof, O. Golberg, S. Hensel, T. Liu, A. Schwendner, D. Vaintrob, and E. Yudovina. *Introduction to Representation Theory*, volume 59 of *Student Mathematical Library*. AMS, 2011.

[31] L. R. Ford and D. Fulkerson. Maximal flow through a network. *Canadian Journal of Mathematics*, 8:399–404, 1956.

[32] Joel Friedman. Sheaves on graphs, their homological invariants, and a proof of the hanna neumann conjecture. arXiv:1105.0129 [math.CO], 2011.

[33] J. Funk. The display locale of a cosheaf. *Cahiers de Topologie et Géométrie Différentielle Catégoriques*, 36(1):53–93, 1995.

[34] Sergei Gelfand, Robert MacPherson, and Kari Vilonen. Perverse sheaves and quivers. *Duke Mathematical Journal*, 83(3):621–643, June 1996.

[35] Robert Ghrist. Barcodes: The persistent topology of data. *Bulletin Amer. Math. Soc.*, 45:61–75, 2008.

[36] Robert Ghrist and Yasuaki Hiraoka. Applications of sheaf cohomology and exact sequences on network codings. 2011.

[37] M. Goresky and R. MacPherson. Intersection homology ii. *Inventiones Mathematicae*, 72(1):77–129, 1983.

[38] Mark Goresky. Introduction to the papers of r. thom and j. mather. *Bulletin (New Series) of the American Mathematical Society*, 49(4):469–474, October 2012.

[39] Mark Goresky and Robert MacPherson. *Stratified Morse Theory*, volume 14 of *Ergebnisse Der Mathematik Und Ihrer Grenzgebiete. 3. Folge a Series of Modern Surveys in Mathematics*. Springer-Verlag, 1988.

[40] R.M. Goresky. Triangulation of stratified objects. *Proceedings of the American Mathematical Society*, pages 193–200, 1978.

[41] Robert Mark Goresky. *Geometric Cohomology and Homology of Stratified Objects*. PhD thesis, Brown University PhD Thesis, May 1976.

[42] John W. Gray. Fragments of the history of sheaf theory. *Lecture Notes in Mathematics*, 753:1–79, 1979.

[43] A. Grothendieck. Sketch of a programme. 1997. Orginally written in 1984. Translated and available online at `http://matematicas.unex.es/~navarro/res/esquisseeng.pdf`.

[44] Daniel C. Isaksen. Calculating limits and colimits in pro-categories. *Fund. Math.*, 175:175–194, 2002.

[45] Masaki Kashiwara and Pierre Schapira. *Sheaves on Manifolds: With a Short History "Les debuts de la theorie des faisceaux" by Christian Houzel*, volume 292 of *Grundlehren der Mathematischen Wissenschaften*. Springer-Verlag, 2nd printing 1994, 3rd printing 2002 edition, 2002.

[46] Masaki Kashiwara and Pierre Schapira. *Categories and Sheaves*, volume 332 of *Grundlehren der Mathematischen Wissenschaften*. Springer-Verlag, 2006.

[47] Sefi Ladkani. On derived equivalences of categories of sheaves over finite posets. 2007.

[48] S.M. Lane and I. Moerdijk. *Sheaves in Geometry and Logic: A First Introduction to Topos Theory*. Universitext (1979). Springer-Verlag, 1992.

[49] Michael Lesnick. Multidimensional interleavings and applications to topological inference. 2012.

[50] Johann Benedict Listing. The MacTutor History of Mathematics archive.

[51] Ta Lê Loi. Lecture 2: Stratifications in o-minimal structures. In *The Japanese-Australian Workshop on Real and Complex Singularities—JARCS III*, volume 43 of *Proc. Centre Math. Appl. Austral. Nat. Univ.*, pages 31–39. Austral. Nat. Univ., Canberra, 2010.

[52] T.L. Loi. Verdier and strict thom stratifications in o-minimal structures. *Illinois Journal of Mathematics*, 42(2):347–356, 1998.

[53] Y.C. Lu. *Singularity theory and an introduction to catastrophe theory*. Universitext (1979). Springer-Verlag, 1976.

[54] Jacob Lurie. Derived algebraic geometry vi: E_k algebras. *arXiv preprint arXiv:0911.0018*, 2009.

[55] Saunders Mac Lane. *Categories for the Working Mathematician*, volume 5 of *Graduate Texts in Mathematics*. Springer-Verlag, 2 edition, 1998.

[56] Saunders Mac Lane. *Categories for the Working Mathematician*, volume 5 of *Graduate Texts in Mathematics 5*. Springer-Verlag, 2nd edition, 1998.

[57] Robert MacPherson. Intersection homology and perverse sheaves. Colloquium lecture notes distributed by the American Mathematical Society, 1991.

[58] John; Mather. Notes on topological stability. *Bulletin (New Series) of the American Mathematical Society*, 49(4):475–506, October 2012.

[59] Davesh Maulik and David Treumann. Constructible functions and lagrangian cycles on orbifolds. 2012.

[60] J. P. May. *Concise Course in Algebraic Topology*. Chicago Lectures in Mathematics, University of Chicago Press, 1999.

[61] Jason McCarty. Limits and colimits in abelian categories.

[62] M.C. McCord. Singular homology groups and homotopy groups of finite topological spaces. *Duke Mathematical Journal*, 33(3):465–474, 1966.

[63] David Miller. Popaths and holinks. *Journal of Homotopy and Related Structures*, 1(1):1–9, 2006.

[64] Haynes Miller. Leray in oflag xviia: The origins of sheaf theory, sheaf cohomology, and spectral sequences. 2000.

[65] J.R. Munkres. *Topology*. Prentice Hall, Incorporated, 2000.

[66] Y. Peterzil and C. Steinhorn. Definable compactness and definable subgroups of o-minimal groups. *Journal of the London Mathematical Society*, 59(3):769–786, 1999.

[67] M.J. Pflaum. *Analytic and Geometric Study of Stratified Spaces: Contributions to Analytic and Geometric Aspects*, volume 1768. Springer, 2001.

[68] A. M. Pitts. On product and change of base for toposes. *Cahiers de Top et Geom. Diff. Categoriques*, 26(1):43–61, 1985.

[69] Henri Poincare. *Papers on Topology: analysis situs and its five supplements, translated by John Stillwell*, volume 37 of *History of Mathematics*. American Mathematical Society, 2010.

[70] Andrei V. Prasolov. Precoshieves of pro-sets and abelian pro-groups are smooth. May 2011.

[71] Georges Reeb. Sur les points singuliers d' une forme de pfaff complement integrable ou d' une fonction numerique. *Comptes Rendus de L' Academie ses Seances*, 222:847–849, 1946.

[72] Pierre Schapira. Algebra and topology.

[73] W. Schmid and K. Vilonen. Characteristic cycles of constructible sheaves. *Inventiones Mathematicae*, 124(1):451–502, 1996.

[74] Peter Schneider. Representation theory and sheaves on the bruhat-tits building. *Publications mathematiques de l' I.H.E.S.*, 85:97–191, 1997.

[75] Peter Schneider. Verdier duality on the building. *J. reine angew Math.*, 494:205–218, 1998.

[76] Jean-Pierre Schneiders. Cosheaves homology. *Bull. Soc. Math. Belgique*, 39:1–31, 1987.

[77] Jörg Schürmann. *Topology of Singular Spaces and Constructible Sheaves*, volume 63 of *Monografie Matematyczne Instytut Matematyczny New Series*. Birkhäuser Verlag, 2003.

[78] Graeme Segal. Classifying spaces and spectral sequences. *Publications mathèmatiques de l 'I.H.È.S.*, 34, 1968.

[79] P.D. Seymour. The matroids with the max-flow min-cut property. *Journal of Combinatorial Theory, Series B*, 23(2):189–222, 1977.

[80] Allen Shepard. *A Cellular Description of the Derived Category of a Stratified Space.* PhD thesis, Brown University PhD Thesis, May 1985.

[81] M. Shiota. *Geometry of Subanalytic and Semialgebraic Sets.* Progress in Mathematics. Birkhäuser, 1997.

[82] Masahiro Shiota. *Geometry of Subanalytic and Semialgebraic Sets*, volume 150 of *Progress in Mathematics.* Birkhäuser, Boston, MA.

[83] Masahiro Shiota. Thom's conjecture on triangulations of maps. *Topology*, 39:383–399, 2000.

[84] Edwin Spanier. *Algebraic Topology.* Springer-Verlag, 1994.

[85] N. E. Steenrod. Homology with local coefficients. *Annals of Mathematics*, 44(4):pp. 610–627, 1943.

[86] Yuichi Sugiki. The category of cosheaves and laplace transforms. Preprint Series 33, The University of Tokyo Graduate School of Mathematical Sciences, December 2001.

[87] R. Thom. Ensembles et morphismes stratifiés. *Bull. Amer. Math. Soc*, 75(2):240–284, 1969.

[88] René Thom. La stabilité topologique des applications polynomiales. *Enseignement Math. (2)*, 8:24–33, 1962.

[89] David Treumann. Exit paths and constructible stacks. *Compositio Mathematica*, 145:1504–1532, 2009.

[90] L. van den Dries and C. Miller. Geometric categories and o-minimal structures. *Duke Math. J.*, 84(2):497–540, 1996. Updated Version at `http://www.math.osu.edu/~miller.1987/newg.pdf` with list of errors and corrections at `http://www.math.osu.edu/~miller.1987/geofix.pdf`.

[91] L.P.D. van den Dries. *Tame Topology and O-minimal Structures.* London Mathematical Society Lecture Note Series. Cambridge University Press, 1998.

[92] Maxim Vybornov. *Sheaves on Triangulated Spaces and Koszul Duality.* PhD thesis, Yale University, 1999.

[93] Charles A. Weibel. *An Introduction to Homological Algebra*, volume 38 of *Cambridge studies in advanced mathematics.* Cambridge University Press, 1994.

[94] George Whitehead. *Elements of Homotopy Theory*, volume 61 of *Graduate Texts in Mathematics.* Springer-Verlag, 1978.

[95] Hassler Whitney. Local properties of analytic varieties. In *Differential and Combinatorial Topology (A Symposium in Honor of Marston Morse)*, pages 205–244. Princeton Univ. Press, Princeton, N. J., 1965.

[96] Hassler Whitney. Tangents to an analytic variety. *Annals of Mathematics*, 81(3):pp. 496–549, 1965.

[97] Jon Woolf. 2012-2013.

[98] Jonathan Woolf. The fundamental category of a stratified space. *Journal of Homotopy and Related Structures*, 4(1):359–387, 2009.

[99] Kohji Yanagawa. Dualizing complex of the incidence algebra of a finite regular cell complex.

[100] Christopher Zeeman. Dihomology. i. relations between homology theories. *Proc. London Math. Soc.*, 3(12):609–638, 1962.

[101] Christopher Zeeman. Dihomology. ii. *Proc. London Math. Soc.*, 3(12):639–689, 1962.