SHEAVES, COSHEAVES AND APPLICATIONS

Justin Michael Curry

A DISSERTATION in

Mathematics

Presented to the Faculties of The University of Pennsylvania in Partial Fulfillment of the Requirements for the Degree of Doctor of Philosophy

2014

Robert W. Ghrist, Andrea Mitchell University Professor Professor of Mathematics Professor of Electrical and Systems Engineering Supervisor of Dissertation

David Harbater, Christopher H. Browne Distinguished Professor Professor of Mathematics Graduate Group Chairperson

Dissertation Committee:

Robert Ghrist, Andrea Mitchell University Professor of Mathematics Robert MacPherson, Hermann Weyl Professor of Mathematics Tony Pantev, Professor of Mathematics SHEAVES, COSHEAVES AND APPLICATIONS © 2014 Justin Michael Curry

Dedicated to Michael Loyd Curry

ACKNOWLEDGEMENTS

At the current moment in time, a PhD is the highest academic degree awarded in the United States. As such, this thesis reflects over two decades of formal education and schooling across multiple institutions. It also reflects the author's life experience to date, which is formed in many informal and non-academic ways. Accounting for all of these influences and giving credit where credit is due is an impossible task; however, I would like to take some time to thank the many hands which helped this thesis come to be. Given the public nature of this document, I will not always name names, but I will make clear the contributions of my colleagues and teachers.

First and foremost I must thank my parents for bringing me into the world. While my father was in the Navy, my mother had the strongest influence on my education. Instilling a love of reading is, besides giving me life itself, the greatest gift she has given to me. I remember distinctly being told that, given our socio-economic status, receiving a scholarship was the only way I would make it to a university one day, and that reading would take me there. My sense of reverence for reading, among other things, is entirely due to my mother. In contrast, my father engaged me in philosophical dialog at a young age, which is how I gained my first experience with critical thinking. He was never much of a reader; he preferred to sort things out for himself. My entire family — aunts, uncles, cousins and grandparents included — have supported me every step of the way and they know I owe them a great deal.

If anyone thinks that obtaining a PhD comes after an endless stream of successes, they are mistaken. I failed many times, and fortunately I was given many second chances. The educators in the Virginia Beach public school system gave me my first second chance by letting me retake a placement exam for the gifted and talented program. Mr. Ausberry, at Thomas Harrison Middle School, requested that I be accelerated a year in mathematics. Mr. Frutuozo made being a scientist seem fashionable, by being a rock star himself. At Harrisonburg High School, I had many excellent instructors, but I felt the strongest direction and guidance from Henry Buhl, Myron Blosser, Andrew Jackson, Patrick Lintner and David Loughran. Without these hardworking and underpaid teachers, I don't think I would have gotten to go to MIT.

Attending MIT as an undergrad was one of the most formative experiences of my life. It certainly tested and broke the mental toughness that I thought I had. Sitting as a freshman in Denis Auroux's 18.100B and getting my first taste of point-set topology was like stepping in to another dimension. It was too much, too fast, and for a moment I thought that the gate of mathematics was closed to me. Gerald Sussman helped steer me back towards mathematics by preaching the value of the MIT quadrivium: logic/programming, analysis, algebra, geometry, topology, relativity and quantum mechanics. Haynes Miller gave me my second second chance by overlooking my shabby mathematical preparation and letting me study for the Part 1B tripos at Churchill College, as part of the Cambridge-MIT exchange program. Cambridge exposed me to one of the greatest mathematical cultures to ever exist. The integrated nature of the classes and the year-long preparation for the tripos helped me gain independence and synthesize my lessons into a unified whole. It was in the Churchill buttery, where Part II and III students waxed poetic about Riemann surfaces and topoi before I even knew what a ring was, that I decided I had to pursue mathematics for graduate school. Returning to MIT, Haynes exposed me to even more advanced mathematics through summer projects and an IAP project with Aliaa Barakat on integrable systems. Working with Aliaa and, later, Victor Guillemin gave me lots of practice with writing mathematics. All of this has served me well for graduate school.

The University of Pennsylvania appealed to my theory-building nature, but it was having to retake the preliminary exams that helped me become a better problem-solver. While drudging through the *Berkeley Problems in Mathematics* [dSSo4] book, my classes gave me something to look forward to. Tony Pantev made the first-year algebra sequence geometric for me, by introducing us to the Serre-Swan correspondence, categories, simplicial sets, spectra and sheaves. Jonathan Block balanced the algebraic and the geometric in Penn's lengthy topology sequence and introduced us to "Brave New Algebra." The graduate student body at Penn helped contextualize my mathematical lessons, while my roommate, Elaine So, gave me lessons in how to be a better human.

My advisor, Robert Ghrist, believed in me when I did not believe in myself. He taught me to have good taste in mathematics and introduced me to Morse theory, Euler calculus, integral geometry and much more. When I first became his student, the idea that no mathematical object is too abstract to be incarnate resonated deeply with me then, as it does today. Rob outlined a beautiful vision for applied mathematics and worked very hard to realize his ambitious plan. By bringing Yasu Hiraoka, Sanjeevi Krishnan, David Lipsky, Michael Robinson and Radmila Sazdanovic together, Rob augmented my graduate training in profound ways. Given this investment, Rob was extremely generous to let me wander geographically and intellectually. Because of him and Penn's Exchange Scholar program, I was able to live in Princeton for the last few years of my graduate career.

At Princeton, I approached Bob MacPherson in person, who luckily was thinking about applied sheaf theory because of my advisor and Amit Patel, and he agreed to organize a seminar at the Institute for Advanced Study. Listening and watching Bob lecture was like getting to peer through a telescope into the far reaches of the mathematical kingdom. The attendees of this seminar were a motley crew of thinkers and Bob was our shepherd. Bob never said more than was necessary, never wanted his own perspective or understanding to crowd out a newly forming one, and did his best to cultivate each individual's diverse set of mental connections, life experiences and accompanying insights. Many people helped me directly and indirectly while finishing my thesis. Mark Goresky taught me the subtleties of stratification theory, set a high standard for mathematical precision and was enthusiastic and supportive of all my efforts. David Treumann and Jon Woolf both clarified details concerning this work via email. Greg Henselman, Sefi Ladkani, Michael Lesnick, and Jim McClure all provided editorial comments on early drafts of this thesis. Vin de Silva, Matthew Kahle, Dmitriy Morozov, Vidit Nanda, Primoz Skraba and Mikael Vejdemo-Johansson all provided moral support. Ryan and Cate Hodgen kept me sane during my frequent trips to Virginia, where I helped my Dad through the painful process of fighting, and losing to, bladder cancer. My fiancée, Sasha Rahlin, encouraged me to pursue a math major when we first started dating as sophomores, made my junior year abroad doubly wonderful, navigated the stressful two-body aspect of picking a graduate school as a senior, helped me through all of the ups and downs of graduate school along with losing my father, and continues to dazzle me with her focus, drive, beauty and brains. You and Simone are the best.

SHEAVES, COSHEAVES AND APPLICATIONS

Justin Michael Curry

Robert W. Ghrist

This thesis develops the theory of sheaves and cosheaves with an eye towards applications in science and engineering. To provide a theory that is computable, we focus on a combinatorial version of sheaves and cosheaves called **cellular sheaves and cosheaves**, which are finite families of vector spaces and maps parametrized by a cell complex. We develop cellular (co)sheaves as a new tool for topological data analysis, network coding and sensor networks. We utilize the barcode descriptor from persistent homology to interpret cellular cosheaf homology in terms of Borel-Moore homology of the barcode. We associate barcodes to network coding sheaves and prove a duality theorem there. A new approach to multi-modal sensing is introduced, where sheaves and cosheaves model detection and evasion sets. A foundation for multi-dimensional level-set persistent homology is laid via constructible cosheaves, which are equivalent to representations of MacPherson's entrance path category. By proving a van Kampen theorem, we give a direct proof of this equivalence. A cosheaf version of the ith derived pushforward of the constant sheaf along a definable map is constructed directly as a representation of this category. We go on to clarify the relationship of cellular sheaves to cosheaves by providing a *formula* that takes a cellular sheaf and produces a complex of cellular cosheaves. This formula lifts to a derived equivalence, which in turn recovers Verdier duality. Compactly-supported sheaf cohomology is expressed as the coend with the image of the constant sheaf through this equivalence. The equivalence is further used to establish relations between sheaf cohomology and a herein newly introduced theory of *cellular sheaf homology*. Inspired to provide fast algorithms for persistence, we prove that the derived category of cellular sheaves over a 1D cell complex is equivalent to a category of graded sheaves. Finally, we introduce the interleaving distance as an extended metric on the category of sheaves. We prove that global sections partition the space of sheaves into connected components. We conclude with an investigation into the geometry of the space of constructible sheaves over the real line, which we relate to the bottleneck distance in persistence.

CONTENTS

ACKNOWLEDGEMENTS iv vii ABSTRACT LIST OF FIGURES xii xiv PREFACE Т A MATHEMATICAL INTRODUCTION 1 A PRIMER ON CATEGORY THEORY 1 2 Categories 1.1 3 **Diagrams and Representations** 1.2 9 Cones and Limits 10 1.3 **Co-Cones and Colimits** 1.4 13 Adjunctions 16 1.5 THE THEORY OF SHEAVES AND COSHEAVES 18 2 The General Definition 2.1 20 Limits and Colimits over Covers: a Structure Theorem 2.2 25 Rephrased as Equalizers or Co-equalizers 2.2.1 26 2.2.2 Rephrased as Exactness 27 Čech Homology and Cosheaves 28 2.3 **Refinement of Covers** 2.4 32 Generalities on Sheaves and Cosheaves 2.5 34 Pre-sheaves and their Associated Sheaves 2.5.1 35 2.5.2 Grothendieck's Operations 38 Failures to Commute 2.5.3 40 The Existence of Cosheafification 2.5.4 44 PRELIMINARY EXAMPLES 3 49 Sheaves Model Sections 3.1 49 Local Systems: A Bridge Between Sheaves and Cosheaves 3.2 Cosheaves Model Topology 3.3 60 Taming of the Sheaf... and Cosheaf 3.4 63 LINEAR ALGEBRA OVER CELL COMPLEXES 66 Π CELLULAR SHEAVES AND COSHEAVES 67 4 Cell Complexes and the Face-Relation Poset 4.1 67 Partially Ordered Sets: Finite Spaces and Functors 4.2 72 The Alexandrov Topology 4.2.1 72 Functors on Posets 4.2.2 74

53

- 5 FUNCTORS ASSOCIATED TO MAPS 79
 - 5.1 Maps of Posets and Associated Functors 81
 - 5.1.1 Pullback or Inverse Image 82
 - 5.1.2 Application: Subdivision 83
 - 5.1.3 Pushforward or Direct Image
 - 5.1.4 f[†], Pushforwards and Closed Sets 84
 - 5.1.5 f₁: Pushforward with Compact Supports on Cell Complexes 87

83

- 5.2 Calculated Examples 88
 - 5.2.1 Projection to a point 88
 - 5.2.2 Inclusion into a Closed Interval 89
 - 5.2.3 Map to a Circle 90
- 5.3 The Push-Pull Adjunctions 90
- 6 HOMOLOGY AND COHOMOLOGY 92
 - 6.1 Chain Complexes and Homology 93
 - 6.1.1 The Combinatorics of Cell Complexes and Homology 94
 - 6.2 Computational Sheaf Cohomology and Cosheaf Homology 95
 - 6.2.1 Cellular Sheaf Cohomology 95
 - 6.2.2 Cellular Cosheaf Homology 97
 - 6.3 Explaining Homology and Cohomology via Indecomposables 98
 - 6.3.1 Persistence Modules and Barcodes 99
 - 6.3.2 Representation Theory of Categories and the Abelian Structure 102
 - 6.3.3 Quiver Representations and Gabriel's Theorem 106
 - 6.3.4 A Remark on Quivers and Perverse Sheaves 108
- 7 THE DERIVED PERSPECTIVE 110
 - 7.1 Taylor Series for Sheaves 112
 - 7.1.1 Elementary Injectives and Projectives 113
 - 7.1.2 Injective and Projective Resolutions 117
 - 7.2 The Derived Category and Homotopy Theory of Chain Complexes 118
 - 7.3 The Derived Definition of Cosheaf Homology and Sheaf Cohomology 123
 - 7.3.1 Borel-Moore Cosheaf Homology 126
 - 7.3.2 Invariance under Subdivision 126
 - 7.4 Sheaf Homology and Cosheaf Cohomology 128
 - 7.4.1 Invariance under Subdivision 130

III APPLICATIONS TO SCIENCE AND ENGINEERING 131

- 8 TOPOLOGICAL DATA ANALYSIS 132
 - 8.1 Point Clouds and Persistent Homology 132
 - 8.1.1 Level Set and Zigzag Persistence 135
 - 8.2 Approaching Persistence with Sheaves and Cosheaves 136

	8.3	 8.2.1 Cellular Maps and Absolute Homology Cosheaves 136 8.2.2 Local-to-Global Computations via Cellular Sheaves 144 8.2.3 Level Set Persistence Determines Sub-level set Persistence 150 Multidimensional Persistence 154 8.3.1 Generalized Barcodes 157
9	NET 9.1 9.2 9.3	WORK CODING AND ROUTING SHEAVES 159 Duality and Routing Sheaves 161 Counting Paths Cohomologically, or Failures thereof 163 Network Coding Sheaf Homology 164
10	SHE 10.1 10.2 10.3	AVES AND COSHEAVES IN SENSOR NETWORKS 168 A Brief Introduction to Sensors 169 The Coverage Problem: Static and Mobile 169 Intruders and Barcodes 172 10.3.1 Tracking the Topology over Time 172 10.3.2 Linearizing the Sheaf of Sections 175 Multi-Modal Sensing 179 10.4.1 A Deeper Look at Sensing 180 10.4.2 Indecomposables, Evasion Sets, Generalized Barcodes 184
IV	NOV	EL MATHEMATICAL CONTRIBUTIONS 188
11	THE	DEFINABLE ENTRANCE PATH CATEGORY 189
	11.1	Stratification Theory and Tame Topology 190 11.1.1 Whitney Stratified Spaces 193 11.1.2 Stratified Maps and a Counterexample 195 11.1.3 O-minimal Structures 198 11.1.4 Thom-Mather Stratifications 202 11.1.5 Thom Mappings 207 11.1.6 Stratified Maps to the Real Line 210 Representations of the Entrance Path Category 212 11.2.1 Homotopy Links 214 11.2.2 Van Kampen Theorem for Entrance Paths 215 11.2.3 The Equivalence 222 11.2.4 Representations from Stratified Maps 224
12	DUA 12.1 12.2	LITY: EXCHANGE OF SHEAVES AND COSHEAVES231Taking Closures and Classical Dualities Re-Obtained231Derived Equivalence of Sheaves and Cosheaves23512.2.1Linear Duality23712.2.2Verdier Dual Anti-Involution238
13	COSI	HEAVES AS VALUATIONS ON SHEAVES 239
	13.1	Left and Right Modules and Tensor Products 239
	13.2	Compactly-Supported Cohomology 241
		Х

- 13.3 Sheaf Homology and Future Directions 244
- 14 GRADED DESCRIPTIONS OF THE DERIVED CATEGORY 245
 - 14.1 The Derived Category for Complexes of Vector Spaces 246
 - 14.2 Derived Complexes of Cellular Sheaves 248
 - 14.2.1 Counterexample to the Naïve Approach 248
 - 14.2.2 Using the Calculus of Fractions Formulation 250
 - 14.2.3 The Equivalence 250
- 15 A METRIC ON THE CATEGORY OF SHEAVES 254
 - 15.1 Interleavings for Pre-Sheaves 255
 - 15.1.1 Easy Stability 258
 - 15.1.2 Global Sections Obstruct Interleavings 260
 - 15.2 Interleavings for Sheaves 261
 - 15.2.1 The Effect of Sheafification 262
 - 15.2.2 Thickening Global Sections 264
 - 15.2.3 Metric on Sheaves 264
 - 15.3 The Space of Constructible Sheaves over \mathbb{R} 265
 - 15.3.1 Interleavings and Dynamics on Indecomposable Sheaves 268
 - 15.3.2 Coordinates for the Category of Sheaves 270
 - 15.3.3 Towards a Bottleneck Distance for Sheaves 271
- BIBLIOGRAPHY 275
- INDEX 286

LIST OF FIGURES

Figure 1	Covers and Their Nerves 21
Figure 2	Sheaves and Cosheaves of Functions 24
Figure 3	Cohomology Pre-sheaf is Non-Local 36
Figure 4	Is There a Section? 51
Figure 5	Trivial Circle Bundle over the Circle 54
Figure 6	Identification Spaces for the Torus and Klein Bottle 55
Figure 7	Trivial Action with the Torus Map 56
Figure 8	Non-trivial Action from the Klein Bottle 56
Figure 9	Topological Model for Skyscraper Cosheaf61
Figure 10	Barycentric Subdivision of a Singular Chain 62
Figure 11	Sheaf of Sections 65
Figure 12	Cosheaf of Components 65
Figure 13	Bott's Height Function on the Torus 69
Figure 14	Cellular Cosheaf by Taking H_1 of the Star pre-images 70
Figure 15	Cellular Description of j_*k_Y and $j_!k_Y$ respectively 71
Figure 16	Alexandrov Space Associated to the Unit Interval 74
Figure 17	Projection to a Point 89
Figure 18	Inclusion into a Closed Interval 89
Figure 19	Map to a Circle 90
Figure 20	Minimal Cell Structure on an Open Interval 97
Figure 21	Barcodes associated to $T : \mathbb{R}^3 \to \mathbb{R}^2$ for rank(T) = 0, 1, 2 101
Figure 22	Barcodes for the filtration of a Torus 101
Figure 23	Barcodes for a Chain Complex 102
Figure 24	Simply Laced Dynkin Diagrams 107
Figure 25	MacPherson's Motivating Example for the Derived Category 111
Figure 26	Point Cloud Data 133
Figure 27	Height Function on the Circle 136
Figure 28	Homology Cosheaf for a circle, with barcodes 137
Figure 29	Barcodes and the Two Sphere 140
Figure 30	Barcodes for the Cone 141
Figure 31	Barcodes for Bott's torus 143
Figure 32	Determining Sub-level Set from Level Set Persistence 151
Figure 33	A Shape Described with Multi-D Persistence 156
Figure 34	Two Dimensional Barcodes for the Sphere 158
Figure 35	Graph with Decoding Wire 160
Figure 36	No Decoding Edge 163

Figure 37 Decoding Edge 163 Figure 38 Network Coding Sheaf 164 Figure 39 Network Coding Sheaf with Two Decoding Wires 166 Figure 40 Sensors Distributed in a Plane 170 Figure 41 The Space-Time Perspective 171 Figure 42 Mobile Sensor Network 171 Figure 43 Mobile Sensor Network 173 Figure 44 Sheaf of All Possible Evasion Paths 176 Linearized Sheaf of Sections Figure 45 177 Figure 46 Counterexample for the Linearized Sheaf of Sections 178 Multi-Modal Sensors Distributed in a Plane Figure 47 180 Two Multi-Modal Sensors Figure 48 181 Figure 49 Examining the Short Exact Sequence 185 Figure 50 Short Exact Sequence 186 Figure 51 Blowing up at a Point 192 Figure 52 Topologist's Sine Curve 193 Figure 53 Diagram for Whitney Condition (b) 194 Figure 54 Preimage of the Spiral is Not Stratifiable 197 Figure 55 Point Cloud Data 201 Figure 56 A System of Control Data 203 Figure 57 Two Sets of Tubular Neighborhoods 206 Figure 58 Not a Thom Mapping 208 Figure 59 Two Entrance Paths and a Homotopy Between Them 212 Figure 60 Modifying a Homotopy to Stay Inside an Open Set 215 Figure 61 Forcing Move (1c) to Apply 219 Argument for Homotopy Invariance Figure 62 226 Figure 63 A Chain Complex for the 2-Sphere as Barcodes 246 Figure 64 The Counterexample 249 Figure 65 Replacing the Counterexample with a Quasi-Isomorphic Sheaf Figure 66 Do Close Maps Give Rise to "Close" Sheaves? 255 Figure 67 Diagram for the Proof of Lemma 15.1.6 258 Two Maps and their associated Leray Sheaves Figure 68 263

251

The motivation behind this thesis is that sheaves are useful for science and engineering applications. In their most impressionistic form, sheaves are nothing more than a way of tethering data to a space. Passing messages over a network, gathering intelligence in temporal or spatial domains, and characterizing the shape of data as a function of parameter value are all applications where the theory of sheaves and cosheaves are welladapted. There are hopefully many other, yet to be discovered, applications of sheaves that are waiting for the right person to come along and flesh them out. However, the author is of the opinion that applied sheaf theory must confront three issues:

- FOUNDATIONS for technology transfer and communication: The standard formulations of sheaves and derived functors are difficult to communicate to researchers outside of "pure" mathematics. Simpler, alternative descriptions of the sheaves and cosheaves must be presented in an easily accessed format.
- COMPUTATIONS in practice, on a computer, and theoretically: Any application of sheaf theory must be programmable on a computer in an efficient manner.
- **PERTURBATIONS** by noise and approximation: Any technique for modeling the world must be able to account for noise. Being able to test or reject hypotheses formed in a model is essential for interfacing with data from experiments.

In addition to making a first pass at applications of sheaf theory to topological data analysis (TDA) in Chapter 8, network coding in Chapter 9 and sensor networks in Chapter 10, this thesis addresses the above three issues.

A great deal of hard work has already been done to address the first two issues. Combinatorial descriptions of sheaves have been discovered independently by Masaki Kashiwara [Kas84], Robert MacPherson and Christopher Zeeman [Zee62a]. The notion of a cellular sheaf, developed by Allen Shepard [She85] under MacPherson's direction, requires only linear algebra to understand: given a cell complex, a cellular sheaf consists of a choice of vector space for each cell and a choice of linear map from each cell to each of its cofaces, compatible with composition. This notion of a sheaf is easily programmable on a computer, as is its cohomology, which can be further simplified via discrete Morse theory [CGN13]. Unfortunately, Shepard's thesis was never published and the other contributions of Kashiwara and Zeeman overshadowed these modest-looking discoveries; they were developed no further.

In this thesis we attempt to revive the theory of cellular sheaves by developing them for applications and embedding them into a larger mathematical context as well. Although Shepard never explained this, cellular sheaves are actual sheaves when viewed through the Alexandrov topology, as we explain in Chapter 4. This perspective emphasizes the view that cellular sheaves are functors from the face-relation poset. The relationship between functors modeled on posets to sheaves has been explored in many works [Bac75, Yuz81, Yan01, BBR07, Lad08] as well as many others. Cellular sheaves can also be viewed as a special instance of constructible sheaves, which are equivalent to representations of MacPherson's exit path category [Tre09], which we develop in its cosheaf version in Chapter 11 for the first time. To connect these perspectives to a more standard presentation of sheaves, we explain the general theory of sheaves familiar to most mathematicians in Chapter 2. The length of the thesis is in part due to the fact that it attempts to speak in multiple languages to multiple specialists and non-specialists alike.

The use of cosheaves is initially motivated for one simple reason: homology. The applied topology community has made headway into convincing scientists that homology is a useful bit of linear algebra. One can visualize circles in data [Caro9] or holes in sensor networks [dSG06a] and nod when the theory agrees, however cohomology is the theory best suited to sheaves and it is not easily visualized in the absence of Poincaré duality. To continue to use homology when speaking to researchers in other fields, one must work with *cosheaves* instead. However, cosheaves have gained prominence in current research mathematics as well. Costello and Gwilliam's work on factorization algebras in quantum field theory [CG]; Lurie and Salvatore's work on nonabelian Poincaré duality; Ayala, Francis and Tanaka's work on factorization homology and manifold calculus [AFT12]; have all made use of variations on cosheaves. Consequently, setting down a general theory for cosheaves seemed to be well-timed. This thesis provides a proof of the existence of cosheafification for Vect-valued pre-cosheaves, which is non-obvious since cofiltered limits and finite limits do not commute in Vect; one cannot simply dualize a pre-cosheaf into a pre-sheaf and use Grothendieck's sheafification procedure in the opposite category. Fortunately, cellular cosheaves never need to be cosheafified because of their strong finiteness properties. To clarify their relationship with Shepard's theory, we prove that cellular sheaves and cosheaves are derived equivalent in Chapter 12. There is another deeper reason why cosheaves should be studied, which is explained in Chapter 13. Just as tensors take in vectors and spit out numbers, cosheaves act on sheaves and produce vector spaces.

Theoretical computability is one of sheaf theory's greatest strengths. The gluing axiom provides a form of algorithmic compression: if one wants to query the data lying over a large space, it suffices to pass to a cover, compute each piece separately and then glue together the results via a limit (kernel of a matrix). Indeed, the classic Mayer-Vietoris long exact sequence can be viewed as a special case of cellular sheaf cohomology, as Section 8.2.2 explains. Unfortunately, for higher-order stitching together of data, a spectral sequence is required, where Leray differentials frustrate the compression of data offered by sheaves and cohomology. In Chapter 14 we prove in a precise sense that these differentials can be ignored when working over graphs.

The thesis concludes with a first attack on developing a perturbation theory for sheaves. By borrowing the idea of interleavings [CCSG⁺09], in Chapter 15 we introduce an extended metric on the category of sheaves over a metric space. Using this metric we can prove that the assignment from maps $f : Y \to X$ to sheaves f_*k_Y is 1-Lipschitz in these metrics. The broader problem of developing a theory of statistics for sheaves is wide open.

Part I

A MATHEMATICAL INTRODUCTION

This part serves multiple groups of people and can be used in different ways:

- For those who are category theory neophytes, a reading of Chapter 1 is advised, after which they should move on to Part II, with particular emphasis on the beginning of Chapter 4 and Chapter 6.
- Chapter 2 is designed for those who want a general definition of sheaves and cosheaves on a topological space. After looking at the definition, one should proceed as quickly as possible to Chapter 3 to get some simple examples.
- Section 2.2 is meant for people who have always found the expression of the sheaf axiom as an exact sequence a little opaque. Such people are usually frustrated by the notation used in Čech homology, which is the subject of Section 2.3.
- Sections 2.5 and 2.5.4 are for those who think of cosheaves simply as sheaves valued in the opposite category.

1

A PRIMER ON CATEGORY THEORY

"A healthy new seed was planted some twenty odd years ago in the well fertilized soil of the mathematical periodical literature — the notion of a category. It sprouted, took root, flowered, attracted bees, and by now the landscape is dotted with its progeny. It is a beneficent plant: mathematical gardeners have come to appreciate its usefulness in holding down the topsoil and preventing dust storms; indeed, some half dozen books have appeared within the past dozen years putting it to this use. It is a beautiful plant too, whose rapid proliferation has produced many unique and exotic variants; but, perhaps because of its increasingly multiform variety, the book extolling all its loveliness has not yet been written."

— F.E.J. Linton [Lin65]

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" [EM45], 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 Chapter 2. 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" [Mac98] for a book that very well may be the book anticipated by the quote above.

1.1 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 1.1.1 (Category). A **category** C consists of a class of objects denoted obj(C) and a set of morphisms $Hom_C(a, b)$ between any two objects $a, b \in obj(C)$. An individual morphism $f : a \rightarrow b$ is also called an arrow since it points (maps) from a to b. We require that the following axioms hold:

- Two morphisms $f \in Hom_C(a, b)$ and $g \in Hom_C(b, c)$ can be composed to get another morphism $g \circ f \in Hom_C(a, c)$.
- Composition is associative, i.e. if $h \in Hom(c, d)$, then $(h \circ g) \circ f = h \circ (g \circ f)$.
- For each object x there is an identity morphism $id_x \in Hom_C(x, x)$ that satisfies $f \circ id_a = f$ and $id_b \circ f = f$.

When the category **C** is understood, we will sometimes write Hom(a, b) to mean $Hom_C(a, b)$.

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.¹

Definition 1.1.2 (Small and Finite Categories). A category is **small** if its class of objects is actually a set. A category is **finite** if its set of objects has finite cardinality.

¹ The machinery we are referring to is that of Grothendieck universes.

Example 1.1.3 (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-identity 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 1.1.4 (Posets and Preorders). A **preordered set** is a set X along with a relation \leq that satisfies the following two axioms:

Reflexivity — $x \leq x$ for all $x \in X$

TRANSITIVITY — $x \leq y$ and $y \leq z$ implies $x \leq z$

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

ANTI-SYMMETRY — $x \leq y$ and $y \leq x$ implies x = y

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

We now reach our example of fundamental importance.

Example 1.1.5 (Open Set Category). The **open set category** associated to a topological space X, denoted **Open**(X), has as objects the open sets of X and a unique morphism $U \rightarrow V$ for each pair related by inclusion $U \subseteq V$.

There is an example very closely aligned with the category of open sets that is allegedly due to Raoul Bott, who gave it as an example of a topological category [Bot72, LHM⁺10].

Example 1.1.6 (Pointed Open Set Category). The **pointed open set category Open**^{*}(X) associated to a topological space X has pairs (U, x), where U is an open set and x is a point in U, for objects and a unique morphism $(U, x) \rightarrow (V, y)$ if $U \subset V$ and x = y.

This pointed open set category takes us nicely over to a category whose objects are points of a topological space. First, we introduce some terminology.

Definition 1.1.7 (Groupoid). A **groupoid** is a category where every morphisms is invertible. In other words if **G** is a groupoid, then for every pair of objects $x, y \in obj(\mathbf{G})$ and every morphism $\alpha \in Hom_{\mathbf{G}}(x, y)$ there exists a morphism $\beta \in Hom_{\mathbf{G}}(y, x)$ such that $\alpha \circ \beta = id_{u}$ and $\beta \circ \alpha = id_{x}$. **Exercise 1.1.8.** Let **G** be a groupoid with only one object. Show that the structure axioms of a category along with the property of being a groupoid guarantees that **G** is a group. Observe that this gives us a way of treating every group as a category, where multiplication in the group corresponds to composition of morphisms.

Definition 1.1.9 (The Fundamental Groupoid). Let X be a topological space. The **fun-damental groupoid** $\pi_1(X)$ has points $x \in X$ for objects and homotopy classes of paths relative endpoints for morphisms. Specifically,

$$Hom_{\pi_1(X)}(x,y) := \{\gamma : [0,1] \to X \,|\, \gamma(0) = x, \gamma(1) = y\} / \sim$$

where $\gamma \sim \gamma'$ if there exists a third continuous map $h : [0, 1]^2 \to X$ such that $h(0, t) = \gamma(t)$, $h(1, t) = \gamma'(t)$, h(s, 0) = x and h(s, 1) = y.

Remark 1.1.10 (Poincaré ∞ -Groupoid). To a topological space X, one can consider a generalization of the fundamental groupoid, called the **Poincaré** ∞ -**groupoid** $\pi_{\infty}(X)$, which has an object for each point of X, a morphism for every path $\gamma : [0, 1] \rightarrow X$, a "2-morphism" for every continuous map $\sigma : \Delta^2 \rightarrow X$, 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 currently vying to replace ordinary category theory as the foundation for mathematics [Luro9a].

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 **D**. For this paper **D** will usually mean one of the following:

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

The category vect is an example of a subcategory, which we now define.

Definition 1.1.11 (Subcategories). Let C be a category. A **subcategory B** of C consists of a subcollection of objects from C and a choice of subset of the morphism set $Hom_C(x, y)$ for each pair $x, y \in obj(B)$. We require that these morphism sets have the identity and be closed under composition so as to guarantee that **B** is a category. We say that a subcategory is **full** if $Hom_B(x, y) = Hom_C(x, y)$.

Categories have a built-in notion of directionality. For example, in **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 1.1.12 (Initial and Terminal Objects). An object $x \in obj(C)$ is said to be **initial** if for any other object $y \in obj(C)$ 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 **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 **Open**(X) the empty set is initial, but the whole space X is terminal. In **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 1.1.13 (Opposite Category). For any category C there is an **opposite category** C^{op} where all the arrows have been turned around, i.e. $Hom_{C^{op}}(x, y) = Hom_{C}(y, x)$.

Remark 1.1.14 (Duality and Terminology). Because one can always perform a general categorical construction in C or 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, as we will see shortly, there are 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 1.1.15 (Functor). A functor $F : C \to D$ consists of the following data: To each object $a \in C$ an object $F(a) \in D$ 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 $F(id_{\alpha}) = id_{F(\alpha)}$. For such a functor F, we say C is the **domain** and D is the **codomain** of F.

Remark 1.1.16. We can phrase the definition of a functor differently by saying that we have a *function* $F : obj(C) \rightarrow obj(D)$ and *functions* $F(a,b) : Hom_C(a,b) \rightarrow Hom_D(F(a),F(b))$ for every pair of objects $a, b \in obj(C)$. We require that these functions preserve identities and composition. When $F(a,b) : Hom_C(a,b) \rightarrow Hom_D(F(a),F(b))$ 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 **faithful**. When a functor is both full and faithful, we say it is **fully faithful**.

Exercise 1.1.17. Check that the definition of a subcategory guarantees that the inclusion $B \hookrightarrow C$ is a functor.

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

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

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

Historically, there was a plethora of different homology theories — simplicial, singular, Čech, Vietoris, Alexander, et al — and every time one was introduced a long repetition of the basic properties of that homology theory ensued. Understanding the precise relationships between these motivated the notion of a map between functors, which led in turn to the Eilenberg-Steenrod axioms [Mac89, p.335].

Definition 1.1.18 (Natural Transformation). Given two functors $F, G : C \to D$ a **natural transformation**, sometimes written $\eta : F \Rightarrow G$, consists of the following information: to each object $a \in C$, a morphism $\eta(a) : F(a) \to G(a)$ is assigned such that for every morphism $f : a \to b$ in C the following diagram **commutes**:

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

Definition 1.1.19. Two functors $F, G : C \to D$ are said to be **naturally isomorphic** if there is a natural transformation $\eta : F \Rightarrow G$ such that for every object $a \in C$ 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. Since an arrow is an arrow by any other symbol, we will sometimes use the notation $F \rightarrow G$ to denote a natural transformation, instead of $F \Rightarrow G$. In the functor category, we will see that naturally isomorphic functors are isomorphic objects. This demonstrates again the linguistic efficiency of category theory.

Example 1.1.20 (Functor Category). Fun(C, 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 C and D are **isomorphic** if there are functors $F : C \to D$ and $G : D \to C$ such that $G \circ F = id_C$ and $F \circ G = id_D$. 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 $F \circ G$ be equal to id_D , we only require that they be isomorphic as objects in Fun(D,D) and similarly for $G \circ F$ and id_C in Fun(C,C). The reader should compare this with the notion of homotopy equivalence.

Definition 1.1.21. A pair of functors $F : C \to D$ and $G : D \to C$ together define an **adjoint equivalence** of categories if there are two natural isomorphisms of functors $\varepsilon : F \circ G \to id_D$ and $\eta : id_C \to G \circ F$.

We will see that this notion of a equivalence is a special instance of an adjunction, which is taken up in Section 1.5

Equivalence can also be phrased in a way that doesn't require us to construct G as a "weak inverse" of F.

Definition 1.1.22 (Fully Faithful and Essentially Surjective). A functor $F : C \to D$ 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 $d \in D$ there is an object $c \in C$ such that F(c) is isomorphic to d, i.e. F is bijective on isomorphism classes of C and D.

The notion of equivalence allows us to find compressed presentations of a category.

Definition 1.1.23 (Skeletal Subcategory). Suppose **C** is a category, then a subcategory *S* is **skeletal** if the inclusion functor is an equivalence, and no two objects of *S* are isomorphic.

If **C** is small, then we can describe explicitly how to construct a skeletal subcategory S. On the objects of **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 $\text{Hom}_{S}(\bar{x}, \bar{y}) := \text{Hom}_{C}(x, y)$. **Exercise 1.1.24** (Fundamental groupoid). Suppose X is a path connected space. Show that for any point $x_0 \in X$, the fundamental group $\pi_1(X, x_0)$ is a skeletal subcategory of $\pi_1(X)$.

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 $F : C \to D$ and define a functor $F^{op} : C^{op} \to D^{op}$. Moreover, a natural transformation $\eta : F \Rightarrow G$ translates to a natural transformation $\eta^{op} : G^{op} \Rightarrow F^{op}$. This observation allows us to state the equalities

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

since (C^{op})^{op} is isomorphic to C (not just equivalent). See the wonderful work "Abstract and Concrete Categories: The Joy of Cats" [AHS09a] for more on duality and category theory more generally.

1.2 DIAGRAMS AND REPRESENTATIONS

Categories and functors allow us to develop an *algebra of shape*, the shapes being modeled on the domain category of a functor. For example, we will be interested in studying data arranged in the following forms:



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 1.2.1 (Diagram). Suppose I is a small category and C is an arbitrary category. A **diagram** is simply a functor $F : I \rightarrow C$.

Example 1.2.2 (Constant Diagram). For any category I there is always a diagram for each object $O \in C$, called the **constant diagram**, $const_O : I \rightarrow C$ where $const_O(x) = const_O(y) = O$ for all objects $x, y \in I$. Every morphism in I goes to the identity morphism.

Definition 1.2.3 (Representation). A **representation** of a category C is a functor $F : C \rightarrow$ 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 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.



Maps of representations correspond precisely with natural transformations of such functors. Isomorphic representations are naturally isomorphic functors.² These basic notions carry over to the representation theory of arbitrary categories, which allows us to compare different situations in one language.

1.3 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 1.3.1 (Cone). Suppose $F : I \to C$ is a diagram. A **cone** on F is a natural transformation from a constant diagram to F. Specifically, it is a choice of object $L \in C$ 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:



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

Definition 1.3.2. The collection of cones on a diagram F form a category, which we will call **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 1.3.3 (Limit). The **limit** of a diagram $F : I \to C$, denoted $\varprojlim F$ is the terminal object in **Cone**(F). This means that a limit is an object $\varprojlim F \in C$ along with a collection of morphisms $\psi_x : L \to F(x)$ that commute with arrows in the diagram such that whenever

² Confusingly, the term "equivalent representations" is often used.

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.



Remark 1.3.4 (Glossary). Quite confusingly, the following terms are synonyms for limits: inverse limits, projective limits, left roots, lim and <u>lim</u> are all common.

We now consider some examples of limits over discrete categories.

Example 1.3.5 (Products). Consider the following index category and diagram:

• • F(i) = F(j)

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

```
F(\mathfrak{i})\prod F(\mathfrak{j}).
```

More generally, we define the product to be the limit of any diagram $F : I \to C$ 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 1.3.6 (Open Sets: Limits are Intersections). Suppose $\Lambda = \{1, ..., n\}$ 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 $\mathbf{C} = \mathbf{Open}(X)$ be the category of open sets in X. This is a category that has an object for each open set and a single morphism $U \rightarrow V$ if $U \subset V$. A functor $F : \Lambda \rightarrow \mathbf{Open}(X)$ 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 1.3.7. Consider the following small category I along with some representation $F : I \rightarrow Vect$.



By thinking about the definition, one can see that

$$\underline{\lim} F \cong U.$$

Example 1.3.8 (Pullbacks). Consider the category $J = I^{op}$ and a representation $F : J \rightarrow$ Vect.

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

$$\lim 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 1.3.9 (Equalizers and Kernels). Consider the following category K and an arbitrary functor $F : K \rightarrow D$.

•
$$\longrightarrow$$
• $X \xrightarrow{f} Y$

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

$$E \xrightarrow{h} X \xrightarrow{f} Y$$

If D = 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 1.3.10 (Invariants). Suppose that V is a vector space with an endomorphism $T : V \rightarrow V$, 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] \rightarrow 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.

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

1.4 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 1.4.1 (Co-Cone). Given a diagram $F : I \to C$, a **cocone** is a natural transformation from F to a constant diagram. In other words, it consists of an object $C \in C$ along with a collection of maps $\phi_x : F(x) \to C$ such that these maps commute with the ones internal to the diagram.



Similarly, there is a category of cocones to a diagram F, denoted **CoCone**(F). A colimit is a distinguished object in this category.

Definition 1.4.2 (Colimit). The **colimit** of a diagram F is the initial object in the category **CoCone**(F). One should practice dualizing the explicit description of the limit in order to understand the following diagram:



Remark 1.4.3 (Glossary). The following terms are synonyms for colimits: direct limits, inductive/injective limits, right roots, colim and lim 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 1.4.4 (Coproducts). Consider the following index category and diagram:

• • F(i) 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 $F : I \to C$ indexed by a discrete category and write $\coprod_i F(i)$. Alternative notations for the coproduct, depending usually on whether the target category is **Set**, **Vect**, **Ab** or **Top** include

$$\bigoplus_{i} F(i) \quad \text{and} \quad \sum_{i} F(i) \quad \text{and} \quad \bigsqcup_{i} F(i).$$

Example 1.4.5 (Open Sets: Colimits are Unions). Suppose $\Lambda = \{1, ..., n\}$ is a finite discrete category. Let C = Open(X) be the category of open sets in X. A functor $F : \Lambda \rightarrow Open(X)$ 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:

$$\underline{\lim} F = \bigcup_{i=1}^{n} F(i)$$

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 1.4.6 (Pushouts). Consider the following small category I and a representation $F : I \rightarrow Vect$.



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.



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

$$\varinjlim F = W \oplus V/\operatorname{im}(B \oplus -A) \quad \varphi_{U} = q \circ \iota_{W}B = q \circ \iota_{V}A \quad \varphi_{W} = q \circ \iota_{W} \quad \varphi_{V} = q \circ \iota_{W}$$

where q is the quotient map. One should note that this is clearly dual to the limit computation in 1.3.8 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 1.4.7. Consider the example $J = I^{op}$ and corresponding representation $F : J \rightarrow$ Vect.



One can see that

$$\lim F \cong U$$

Example 1.4.8 (Coequalizers and Cokernels). Consider the same category K as before and a functor $F : K \to D$.

•
$$\Longrightarrow$$
• $X \xrightarrow{t} Y$

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

$$X \xrightarrow{f} Y \xrightarrow{h} E$$

If D = Vect and one sets 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 1.4.9 (Co-invariants). As described in Example 1.3.10, a vector space V with an endomorphism T is equivalent to a functor $k[x] \rightarrow \text{Vect}$. The colimit of this functor is called the **coinvariants** of T, i.e.

$$C = V / < Tv - v > .$$

1.5 ADJUNCTIONS

Adjunctions allow us to derive interesting relationships with almost no effort; they are in essence dualities. For the individual interested in using category theory to model the world, facile manipulations of adjunctions is essential. One often can transform a complicated problem into a simpler one via an adjunction, thereby gaining a computational payoff at the cost of abstraction. This is why using adjunctions between the functors defined in Section 5.1 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 the functor $f_!$ was one of the primary reasons that the notion of a derived category was invented. Only by enlarging the domain could a new, adjoint functor $f^!$ be defined. Here we introduce the general theory.

Definition 1.5.1. Suppose $F : C \to D$ and $G : D \to C$ 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 $\eta : id_C \to G \circ F$ and a natural transformation to $\varepsilon : F \circ G \to id_D$ such that

$$G \xrightarrow{\eta G} GFG \xrightarrow{Ge} G$$
, $F \xrightarrow{F\eta} FGF \xrightarrow{eF} F$

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 [Mac98, p.81] for a list. One can just specify η and ask that it is universal,³ i.e. for each $x \in C$ and for every $y \in D$ 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 $G(f') \circ \eta_x = f$.



Of course we could have just defined ϵ and asked that it is universal in a dual sense.⁴ The point is this: an adjunction is equivalent to specifying for every $x \in C$ and $y \in D$ a natural bijection $\varphi_{x,y}$

 $\operatorname{Hom}_{\mathsf{D}}(\mathsf{F}(x), y) \cong \operatorname{Hom}_{\mathsf{C}}(x, \mathsf{G}(y)).$

The following theorem gives us an abstract criterion for determining when a functor has an adjoint.

³ In other words, initial in a particular comma category; see [Mac98, p.56]

⁴ It is final in a different comma category.

Theorem 1.5.2 (Freyd's Adjoint Functor Theorem). Let **D** be a complete category and $G : \mathbf{D} \to \mathbf{C}$ 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 $x \in \mathbf{C}$ 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 $t : a_i \to a$.

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.

THE THEORY OF SHEAVES AND COSHEAVES

"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."¹

— Jean Leray [Ler46]

In its most general form, the subject of this thesis 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 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," [Gra79] 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 π : E → X. Analogously, a cosheaf is a locally-connected space D, called the display locale, that maps to X [Fun95].
- A sheaf (or a cosheaf) is an **abstract assignment of data** a functor that further satisfies a gluing axiom expressed by limits (or colimits).

^{1 &}quot;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."

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 [Ste43]. 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.



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 each *open set* U the set of sections of this map over U:

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

One plausible explanation for using open sets is provided by the **open pasting lemma**, which states² that if $X = \bigcup U_i$ is a (potentially infinite) union of open sets equipped with continuous sections $s_i : U_i \rightarrow E$ that agree on overlaps, then the set-theoretically defined section $s : X \rightarrow E$ 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 $U \rightarrow X$ that satisfy certain conditions reminiscent of an open cover [MM92]. 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 [AoM62].

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

² Munkres calls this the "local formulation of continuity" in theorem 18.2(f) [Munoo]. Munkres reserves the term "pasting lemma" for the closed set version, which is stated directly afterwards as theorem 18.3.

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 general 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 with the sheaf and cosheaf property.

2.1 THE GENERAL 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

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

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.

Definition 2.1.1. Let X be a topological space and U an open set in X. An **open cover** of U is a collection of open sets $\mathcal{U} := \{U_i\}_{i \in \Lambda}$ whose union is U.
Pavel Alexandrov introduced in 1928 a method³ for associating to every open cover an abstract simplicial complex [Ale28]. We will use these shapes to model our limits and colimits of interest.

Definition 2.1.2. Suppose $\mathcal{U} := \{U_i\}_{i \in \Lambda}$ is an open cover of 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 $I \rightarrow J$ if $J \subseteq I$. Since our intersections are only finite, and the finite intersection of open sets is open, we get natural functors

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

Remark 2.1.3. Sometimes we will use the notation N(U), N_U and N interchangeably, depending on the context.



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 theorem** of Leray and Borsuk [Ler45, Bor48] states 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.

³ In Definition 8.2.16 we consider the "correct" generalization of the nerve.

Definition 2.1.4 (Pre-Sheaf and Pre-Cosheaf). A **pre-sheaf** is a functor $F : \mathbf{Open}(X)^{op} \to \mathbf{D}$ and a **pre-cosheaf** is a functor $\widehat{F} : \mathbf{Open}(X) \to \mathbf{D}$. If $V \subset U$, then we usually write the **restriction map** as $\rho_{V,U}^{F} : F(U) \to F(V)$ and the **extension map** as $r_{U,V}^{\widehat{F}} : \widehat{F}(V) \to \widehat{F}(U)$. Often we omit the superscript F or \widehat{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:



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 $N(U) \rightarrow Open(X)$ is just the union $U = \cup U_i$, 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. Examples where this does not occur are given in Example 2.5.1 and Example 2.5.2.

Definition 2.1.5 (Sheaves and Cosheaves). Suppose F is a pre-sheaf and \widehat{F} is a pre-cosheaf, both of which are valued in **D**. Suppose $\mathcal{U} = \{U_i\}$ is an open cover of **U**. We say that F is a **sheaf on** \mathcal{U} if the unique map from F(U) to the limit of $F \circ \iota_U^{op}$, written

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

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

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

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

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

Remark 2.1.6 (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 **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 $F : I \rightarrow D$ exists is called **complete**. Similarly, if the colimit of an arbitrary diagram exists, we say **D** is **co-complete**. The category **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 are no such covers⁴ this is the limit over the empty diagram, i.e. **Cone**(\emptyset) = **D**, whose terminal object is the terminal object of **D**. Similarly, for a cosheaf $\widehat{F}(\emptyset)$ must be the initial object in **D**. For **D** = **Vect** the initial and terminal objects coincide with the zero vector space.

It is true that if **D** has pullbacks (see Example 1.3.8 in Section 1) and a terminal object then it has all finite limits. The dual statement that having an initial object and pushouts (see Example 1.4.6) implies finitely co-complete is also true. Thus, if one focuses on sheaves and cosheaves valued in vect (the category of finite dimensional vector spaces and linear maps), then the collection of covers of U one can consider must be restricted. 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

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

⁴ Alternatively one argues that the empty set covers itself and hence the value there is chosen to be the initial/terminal object of the category **D**.

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.



FIGURE 2: Sheaves and Cosheaves of Functions

Example 2.1.7 (Cover by Two Sets). Suppose D = Set, and suppose $U = \{U_1, U_2\}$ is a cover of U. The sheaf condition says that

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

i.e. F(U) lists the set of consistent choices of elements from $F(U_1)$ and $F(U_2)$. In particular, F[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 D = Set is slightly strange. It says that

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

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 the assignment

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

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 U₁ and U₂. By taking a partition of unity subordinate to the cover U we get two functions $\lambda_1(x)$ and $\lambda_2(x)$ such that

$$f(x) = f_1(x) + f_2(x)$$
 where $f_1(x) := \lambda_1(x)f(x)$ and $f_2(x) := \lambda_2(x)f(x)$

By carrying out the colimit in a data category equipped with sums, such as D = Vect of Ab, then compactly supported functions do define a cosheaf valued there.

More generally, if D = Vect, then the cosheaf axiom for the cover says the sequence

$$\widehat{\mathsf{F}}(\mathsf{U}_{12}) \to \widehat{\mathsf{F}}(\mathsf{U}_1) \oplus \widehat{\mathsf{F}}(\mathsf{U}_2) \to \widehat{\mathsf{F}}(\mathsf{U}) \to 0$$

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

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

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

2.2 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 2.1.5 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 $F : N(U)^{op} \to D$ 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 $\widehat{F} : N(U) \to D$ can be thought of as a quotient of the coproduct of the functor over the vertices. Said using formulas, this is

$$\varprojlim F \hookrightarrow \prod F(\mathfrak{i}) \qquad \coprod \widehat{F}(\mathfrak{i}) \twoheadrightarrow \varinjlim \widehat{F}.$$

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 2.2.1. A category **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 (Sketch). One should consult [Awo10, Prop. 5.22-3] for a complete proof. To give the reader the idea, one can compute the limit of $F : I \rightarrow D$ 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.

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

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.

2.2.1 Rephrased as Equalizers or Co-equalizers

The method outlined in Theorem 2.2.1 for computing limits and colimits contains too much redundant information for the case $I = N(U)^{op}$. As such, we state the precise, simplified formulation here. The sheaf and cosheaf axioms can be rephrased as saying that the following sequences

$$F(\mathbf{U}) \xrightarrow{e} \prod F(\mathbf{U}_i) \xrightarrow{f^-} \prod_{i < j} F(\mathbf{U}_i \cap \mathbf{U}_j)$$
$$\coprod_{i < j} \widehat{F}(\mathbf{U}_i \cap \mathbf{U}_j) \xrightarrow{g^+} \coprod \widehat{F}(\mathbf{U}_i) \xrightarrow{\mathbf{u}} \widehat{F}(\mathbf{U})$$

are an equalizer and a co-equalizer respectively.

Exercise 2.2.2. Prove that the limit or colimit over the nerve of a cover can be determined after considering only the elements of the cover and their pairwise intersections. Do this by observing that the limit or colimit over the 1-skeleton⁵ of the nerve defines a cone

⁵ In higher homotopy analogs of sheaves and cosheaves one works over the whole Čech tower of a cover [DI01, D0u07, GG].

or cocone over the whole nerve and employing universal properties. Then apply the equation from Theorem 2.2.1 and its dual version to prove that the re-written axioms of Section 2.2.1 and Section 2.2.2 are correct.

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

$$\operatorname{Hom}(X,\prod_{i}Y_{i})\cong\prod_{i}\operatorname{Hom}(X,Y_{i})\quad \text{and}\quad \operatorname{Hom}(\coprod_{i}X_{i},Y)\cong\prod_{i}\operatorname{Hom}_{i}(X_{i},Y)$$

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

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

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

	*	ρ _{12,2}	*		ρ _{12,1}	*	*	
$f^+ =$	*	*	ρ _{13,3}	$f^- =$	ρ _{13,1}	*	*	•
	*	*	ρ _{23,3}		*	ρ _{23,2}	*	

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

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

2.2.2 *Rephrased as Exactness*

If D = 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

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

$$\bigoplus_{i < j} \widehat{F}(U_i \cap U_j) \xrightarrow{a_1} \bigoplus \widehat{F}(U_i) \longrightarrow \widehat{F}(U) \longrightarrow 0$$

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 $F(U) \cong \text{ker}(d^0)$ and the cosheaf axiom says that $\widehat{F}(U) \cong \text{coker}(\partial_1)$.

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[U].

2.3 ČECH HOMOLOGY AND COSHEAVES

In Section 2.2.1 we rephrased the limits and colimits coming from covers as equalizers and coequalizers. For the data category $\mathbf{D} = \mathbf{Vect}$ 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-cosheaf⁶ 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 2.3.1 (Čech Homology). Given a pre-cosheaf of vector spaces \widehat{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

$$(\check{C}_{\bullet}(\mathfrak{U};\widehat{F}), \mathfrak{d}_{\bullet})$$
 where $\check{C}_{p}(\mathfrak{U};\widehat{F}) := \bigoplus_{|I|=p+1} \widehat{F}(U_{I})$ for $I \in N(\mathfrak{U})$.

⁶ Or pre-sheaf, but we'll leave it to the reader to dualize.

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

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

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**

$$\check{H}_{p}(\mathcal{U};\widehat{F}) := \frac{\ker \partial_{p}}{\operatorname{im} \partial_{p+1}}$$
 i.e. $H_{p}(\check{C}_{\bullet}(\mathcal{U};\widehat{F}))$

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

Lemma 2.3.2. The differential ϑ in the Čech complex for a cover \mathcal{U} and a pre-cosheaf F of vector spaces satisfies $\vartheta_p \circ \vartheta_{p+1} = \emptyset$.

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:



Let's follow a typical element $s_I \in \widehat{F}(U_I)$ through the diagram on the right upon applying the formula $\partial \circ \partial$. First note that the fact that \widehat{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 ϑ 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 $\vartheta(s_{I})$. Assuming j < k and applying the definition of the boundary map to elements in $\widehat{F}(U_{I}^{(j)})$ implies that we must actually delete the k-1st en-

try 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 $\widehat{F}(U_I^{(j,k)})$ is

$$(-1)^{k-1}(-1)^{j}r_{u_{I}^{(j,k)},u_{I}}(s_{I}) + (-1)^{k}(-1)^{j}r_{u_{I}^{(j,k)},u_{I}}(s_{I}) = 0.$$

Example 2.3.3. 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:

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

One can easily verify that the ker $\partial_1 = im\partial_2$ 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:

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

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 2.2.2, 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

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

for any choice of cover \mathcal{U} of \mathcal{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 2.4 to see why fineness matters) on nice spaces, then one can also 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 2.3.4. 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

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

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

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

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 [Bre97, 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

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

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

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

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

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

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

$$\mathbf{U} \rightsquigarrow \mathbf{H}_0(\mathbf{U}; \mathbf{k})$$

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:

$$H_{1}(U \cap V; k) \longrightarrow H_{1}(U; k) \oplus H_{1}(V; k) \longrightarrow H_{1}(U \cup V; k) \longrightarrow H_{0}(U \cap V; F) \longrightarrow H_{0}(U; k) \oplus H_{0}(V; k) \longrightarrow H_{0}(U \cup V; k) \longrightarrow 0$$

The moral from this example is that, in essence,

Any functor that satisfies Mayer-Vietoris is a cosheaf.

2.4 REFINEMENT OF COVERS

We have defined the sheaf and cosheaf axioms for a cover \mathcal{U} . The coarsest possible cover of an open set \mathcal{U} is the cover with one element { \mathcal{U} }. Thus, one way of interpreting the sheaf and cosheaf axiom is that $F[\mathcal{U}]$ and $\widehat{F}[\mathcal{U}]$ 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 2.4.1 (Refinement of Covers). Suppose \mathcal{U} and \mathcal{U}' are covers of \mathcal{U} , then we say that \mathcal{U}' refines \mathcal{U} if for every $\mathcal{U}'_i \in \mathcal{U}'$ there is a $\mathcal{U}_j \in \mathcal{U}$ and an inclusion $\mathcal{U}'_i \to \mathcal{U}_j$. Note that every cover refines the trivial cover {U}.

Definition 2.4.2. The refinement relation endows the collection of covers of U with the structure of a category Cov(U), whose objects are covers \mathcal{U} with a unique morphism $\mathcal{U}' \rightarrow \mathcal{U}$ 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.



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

Lemma 2.4.3. Let \hat{F} and F be a pre-cosheaf and a pre-sheaf respectively. Suppose U' refines another cover U of an open set U. Then there are well-defined maps

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

i.e. we get functors \widehat{F} : $Cov(U) \rightarrow D$ and F: $Cov(U)^{op} \rightarrow D$.

Proof. We'll detail the proof for a pre-cosheaf $\widehat{\mathsf{F}}$ 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 $\widehat{\mathsf{F}} \circ \iota_{\mathcal{U}'} \Rightarrow \widehat{\mathsf{F}} \circ \iota_{\mathcal{U}}$. The colimit defines a natural transformation from $\widehat{\mathsf{F}} \circ \iota_{\mathcal{U}}$ to the constant diagram $\widehat{\mathsf{F}}[\mathcal{U}]$. Since the composition of natural transformations is a natural transformation, this induces a cocone $\widehat{\mathsf{F}} \circ \iota_{\mathcal{U}'} \Rightarrow \widehat{\mathsf{F}}[\mathcal{U}]$ which, by the universal property of the colimit, defines a unique induced map there, i.e.

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

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 $\widehat{F}[\mathcal{U}'] \to \widehat{F}[\mathcal{U}]$. 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 \widehat{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:



Corollary 2.4.4. 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

$$\mathfrak{U}' \to \mathfrak{U} \to \{\mathbf{U}\}.$$

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



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

$$\widehat{\mathsf{F}}[\mathcal{U}] \xrightarrow{\cong} \widehat{\mathsf{F}}(\mathcal{U}) \qquad \mathsf{F}(\mathcal{U}) \xrightarrow{\cong} \mathsf{F}[\mathcal{U}] \ .$$

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.

2.5 GENERALITIES ON SHEAVES AND COSHEAVES

Sheaves have proved to be highly successful tools in pure mathematics over the past 60-70 years. This is largely because sheaves provide precise mechanisms for determin-

ing global unknowns from local knowns. These mechanisms are greatly enhanced by considering the operations such as Hom and \otimes on sheaves, as well as the pushforward and pullback of a sheaf along a map, which we define in Section 2.5.2. Some of these operations are only available after applying a certain repair to turn a presheaf into a sheaf, also known as **sheafification**, which we define in Section 2.5.1 after some preliminary discussion of examples.

One would like to know if a similar story is true for cosheaves. After all, any functor $\widehat{F} : \operatorname{Open}(X) \to D$ is exactly equivalent to specifying a functor $F : \operatorname{Open}(X)^{\operatorname{op}} \to D^{\operatorname{op}}$. However, certain asymmetries prevent such an observation from being as useful as one might hope. These asymmetries are outlined in Section 2.5.3 and obstruct the use of Grothendieck's version of sheafification to define an analogous **cosheafification**. However, we prove that such a device must exist in Section 2.5.4 without knowing a particular construction.

2.5.1 Pre-sheaves and their Associated Sheaves

Sheaves are fundamentally local structures. Informally stated, a pre-sheaf can fail to be a sheaf in two independent ways:

- Non-Local: If a pre-sheaf has a section s ∈ F(U) that cannot be constructed from sections over smaller open sets in U a cover of U, for example then F fails to be a sheaf.
- *Inconsistent*: If a pre-sheaf has a pair of sections s ≠ t ∈ F(U) such that when restricted to every smaller open set they define the same section, then F fails to be a sheaf.

Let's illustrate both of these failures with two examples.

Example 2.5.1 (Non-local). Let F be a presheaf of vector spaces over the real line \mathbb{R} , defined as follows:

$$F(\mathbf{U}) = \begin{cases} \mathbf{k} & \text{if } (-1,1) \subset \mathbf{U} \\ \mathbf{0} & \text{o.w.} \end{cases}$$

In particular, F assigns the zero vector space to every open ball $B_r(x)$ centered at $x \in \mathbb{R}$ with $r \leq 1/2$. This collection of balls covers the real line thus if F were a sheaf, then $F(\mathbb{R}) = 0$, but it is the vector space k instead.

An incarnation of this example, depicted in Figure 3, is the pre-sheaf that assigns to every open set U, the first cohomology of the inverse image of U under a map $f : S^1 \to \mathbb{R}$, i.e.

$$\mathbf{U} \rightsquigarrow \mathbf{H}_1(\mathbf{f}^{-1}(\mathbf{U}); \mathbf{k}).$$



FIGURE 3: Cohomology Pre-sheaf is Non-Local

Example 2.5.2 (Inconsistent). Let F be a presheaf of sets over the real line \mathbb{R} , defined as follows:

$$F(U) = \begin{cases} \{a, b\} & \text{if } U = X \\ \{\star\} & \text{o.w.} \end{cases}$$

This presheaf is like two friends that agree on every possible political issue, but still belong to two different political parties.

Fortunately, there is a general method of repair that can make any presheaf F into a sheaf \tilde{F} . Although this method modifies the value of F on open sets, it leaves at least one feature of the presheaf unchanged. This is the stalk of the presheaf.

Definition 2.5.3. Let F be a pre-sheaf on a topological space X and $x \in X$ a point. The **stalk** of F at x is the direct limit of F over open sets U containing x:

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

The stalk is the "local value" of a presheaf at x. Notice that every element $t \in F(U)$ with $x \in U$ has an associated value t_x , which is the image of t in the direct limit.

Remark 2.5.4. Since F only assigns data to open sets, one often uses the direct limit construction to assign data to arbitrary sets of X; the stalk is just a special example of this principle.

Now we introduce the procedure for turning an arbitrary presheaf of sets, vector spaces or groups into a sheaf of the same type.

Sheafification

We begin our introduction to the sheaf associated to a presheaf with a careful introduction to products and disjoint unions, following [DL12].

Definition 2.5.5 (Disjoint Unions and Products in Set). Suppose $\{X_s\}_{s \in S}$ is a family of sets indexed by S. The **disjoint union** is a union that tracks the indexing set:

$$\bigsqcup_{s\in S} X_s := \bigcup_{s\in S} X_s \times \{s\}$$

The **product** can be written as a set of maps:

$$\prod_{s \in S} X_s := \{f : S \to \bigsqcup_{s \in S} X_s \,|\, f(s) \in X_s \,\forall s \in S\}$$

The projection maps

$$\pi_{s'}:\prod_{s\in S}X_s\to X_{s'}$$

are defined by evaluation $\pi_{s'}(f) = f(s')$.

Definition 2.5.6 (Sheafification). Let $F : \mathbf{Open}(X)^{op} \to \mathbf{Set}$ be a presheaf of sets and let F_x denote the stalk of F at x. Now for each open set U, form the product $\prod_{x \in U} F_x$. The **sheafification** \widetilde{F} of F assigns to every open set U the functions in $\prod_{x \in U} F_x$ that "locally extend," i.e.

$$\widetilde{F}(U) := \{ s \in \prod_{x \in U} F_x \mid \forall x \in U \ s(x) \in F_x \text{, } \exists V \ni x \ V \subset U \ t \in F(V) \ s.t. \ t_y = s(y) \ \forall y \in V \}$$

There is a natural transformation $\theta : F \to \widetilde{F}$ that takes every element $s \in F(U)$ to the map $s : x \in U \mapsto s_x \in F_x$. In particular, $\theta_x : F_x \to \widetilde{F}_x$ is an isomorphism.

One can summarize sheafification more elegantly in the language of categories. Since every sheaf is also a pre-sheaf, we have an inclusion functor

$$\iota : Shv(X; Set) \hookrightarrow Fun(Open(X)^{op}, Set) =: PreShv(X; Set)$$

that has a left adjoint, i.e. there is a universal natural transformation θ : id_{PreShv} $\rightarrow \iota \circ (-)$, see Section 1.5 for a reminder. Such a subcategory is called **reflective**. This guarantees,

for example, that if F is an arbitrary pre-sheaf and G is a sheaf regarded as a pre-sheaf $G = \iota(G)$, then we have the following universal property:



Pulling back along $\theta_{\rm F}$ induces the natural isomorphism of Hom-sets:

$$Hom_{Shv}(F, G) \cong Hom_{PreShv}(F, \iota(G))$$

2.5.2 Grothendieck's Operations

What makes sheaf theory such a powerful machine is that there are many natural operations on sheaves and well understood adjunctions between these operations. However, many of these operations only exist with the aid of sheafification. In particular, there are the following six operations, grouped into three adjoint pairs, the third of which exists only in a suitable enlargement of the category of sheaves.

$$(f^*, f_*)$$
 (\otimes, \mathcal{Hom}) $(f_!, f^!)$

Here we will consider only four out of the six in order to forego this extra difficulty of "enlarging" the category of sheaves.

Definition 2.5.7 (Pushforward Sheaf). Let $f : Y \to X$ be a continuous map and let G be a sheaf on Y. The **pushforward sheaf** is defined by the formula:

$$f_*G(U) := G(f^{-1}(U))$$

There should be an inverse operation that takes a sheaf F on X and pulls back along $f: Y \rightarrow X$. After all, if $i: W \hookrightarrow X$ is the inclusion of an open set, then a natural candidate for the pullback sheaf i^*F would be the restriction of the domain of definition of F to only those open sets contained in W.

$$F|_W(U) = F(U)$$

However, if $f : Y \rightarrow X$ is not an open map, then there is no hope for an easy definition. Sheafification, however, comes to the rescue.

Definition 2.5.8 (Pullback Sheaf). Let $f : Y \to X$ be a continuous map and F a sheaf on X. The **pullback sheaf**, written f^*F is the sheafification of the pre-sheaf

$$U \rightsquigarrow \varinjlim_{V \supset f(U)} F(V)$$

Example 2.5.9 (The Stalk). Let $i : \{x\} \hookrightarrow X$ be the inclusion of a point into a space with a sheaf F defined on it. The sheaf $i^*F \cong F_x$ is the stalk at x.

Exercise 2.5.10. Verify for $i : W \hookrightarrow X$ that $i^*F = F|_W$

The next pair of interest is the middle pair.

Definition 2.5.11 (Sheaf Hom). Suppose F and G are sheaves of abelian groups over a single space X. The **sheaf hom** Hom(F, G) assigns to every open set

$$\mathcal{H}om(F,G)(U) := \operatorname{Hom}_{\mathbf{Shv}(U)}(F|_{U},G|_{U})$$

For the algebraically minded, there should be a knee-jerk response for an associated tensor sheaf, however the naïve assignment needs to be sheafified.

Definition 2.5.12. Suppose F and G are sheaves of abelian groups over a single space X. The **tensor product of sheaves** $F \otimes G$ is defined to be the sheafification of the assignment

$$U \rightsquigarrow F(U) \otimes G(U)$$

The reader is encouraged to work through the following exercise, borrowed from [Acho7] with a few extra hints.

Exercise 2.5.13. Let Q be the sheaf of sections of the map $f : S^1 \to S^1$ defined via complex coordinates as $f(z) = z^2$, i.e.

$$\mathbf{Q}(\mathbf{U}) := \{ \mathbf{s} : \mathbf{U} \to \mathbf{S}^1 \mid \mathbf{f} \circ \mathbf{s}(z) = z \}.$$

Check that this sheaf has no global sections. Now let Q_k be the sheaf which assigns to each open set U the k vector space freely generated by the set Q(U). Show by taking a carefully chosen cover of S¹ that

$$F: U \rightsquigarrow Q_k(U) \otimes Q_k(U)$$

is not a sheaf. Observe that we have a natural method for tensoring elements of $Q_k(U)$ together via pointwise multiplication. Any element $s \in Q_k(U)$ satisfies $(s \otimes s)(z) = s(z)^2 = z$, but there are interesting cross-multiple terms.

Although working out the above exercise is rewarding, the category theorist knows that since tensor products are colimit constructions and the sheaf axiom involves limits, one should instantly be suspicious of such a construction defining a sheaf. However, one can construct cosheaf-theoretic analogs of the above functors and there the tensor cosheaf is naturally a cosheaf, but cosheaf Hom needs to be "cosheafified," if such a thing exists.

2.5.3 Failures to Commute

Unfortunately, the universe appears to have a sort of handedness that makes certain constructions for sheaves natural, but not so for cosheaves. This is because most data categories **D**, such as **Set**, **Vect** or **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 **D** = **Set**, **Vect** or **Ab** are used in the development of sheaf theory, which do not necessarily hold in **D**^{op}. The centerpiece of this discussion will be understanding that filtered colimits commute with finite limits in **D**, but cofiltered limits do not necessarily commute with finite colimits in **D**. Let us now relay the necessary definitions.

Definition 2.5.14. A non-empty category **C** is called **filtered** if the following two properties are satisfied:

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

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

Example 2.5.15. In Section 2.4 we considered the category of covers 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 2.5.16. Suppose I is a filtered indexing category with $F : I \rightarrow D$ and $G : I^{op} \rightarrow D$ 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 already introduced in the context of sheafification.

Example 2.5.17 ((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 Open(X) or a filtered

subcategory of **Open**(X)^{op}. Consequently for a pre-sheaf F or a pre-cosheaf F, the following

$$F_{x} := \varinjlim_{U \ni x} F(U) \qquad \lim_{U \ni x} \widehat{F}(U) =: \widehat{F}_{x}$$

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 \widehat{F} . Of course, to make such a statement meaningful, one needs to assume the data category **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 2.5.18. Let I be a filtered indexing category and J a finite category. Then any functor α : I × J → D where D = Set, Vect, or Ab, has the property that the natural map

$$\varinjlim_{\mathrm{I}} \varprojlim_{\mathrm{J}} \alpha(\mathfrak{i},\mathfrak{j}) \to \varprojlim_{\mathrm{J}} \varinjlim_{\mathrm{I}} \alpha(\mathfrak{i},\mathfrak{j})$$

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 [KSo6]. First note that the product category $I \times J$ is just a product in **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 $\varprojlim_J \alpha : I \to D$ gotten by assigning to each object $i \in obj(I)$ the limit over J of $\alpha(i, -) : J \to D$. Taking the colimit over I defines the first expression. Similar reasoning defines the second.

For an application of this theorem we introduce some ideas to the world of pre-sheaves valued in **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 map. A sequence of pre-sheaves⁷ is exact if and only if for each open set U the associated sequence of vector spaces is exact, i.e.

 $0 \to E \to F \to G \to 0$ iff $0 \to E(U) \to F(U) \to G(U) \to 0$

⁷ The corresponding statement for sheaves is not true.

Theorem 2.5.18 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 2.5.19. For D = Set, Vect or Ab it is not true that cofiltered limits and finite colimits commute. Consequently, if A, B, C : $\mathbb{N}^{op} \rightarrow Ab$ (or Vect) are functors from the category of natural numbers equipped with the opposite ordering, with natural transformations $A \rightarrow B \rightarrow C$ 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.

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

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



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

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

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 $a \rightarrow (n + 1)a$ 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 [Sch]. \Box

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 hypotheses⁸ exactness can be guaranteed. This holds for categories like **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 [AM69].

This last comment about 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

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

establishes an equivalence of categories.⁹ However, issues of stalks versus costalks is not the primary obstacle that the asymmetry of Theorem 2.5.18 presents. That obstacle has to 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 F : $Cov(U)^{op} \rightarrow D$, see [KSo6, Sec. 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 2.5.18. This now gets us to the more fundamental reason why the study of cosheaves may be so obscure: The non-exactness of lim thwarts the Grothendieck prescription for cosheafification. *For pre-cosheaves valued in* Set, Ab or 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 [Bre68, Bre97] is to define an equivalence relation on pre-cosheaves, more nuanced than iso-

⁸ i.e. the Mittag-Leffler condition. See [KS02, Sec. 1.12] or [AHS09b, pp. 211-214] for more details.

⁹ This does *not* extend to an equivalence between Vect and its opposite category. In fact, Vect^{op} is equivalent to the category pro – vect, cf. [Isao2, Rmk.6.2].

morphism, which is constructed through zig-zag diagrams of *local* isomorphisms. Bredon develops an operation which uses Čech homology to take in one precosheaf 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 $P : I^{op} \rightarrow C$ 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 [Sch87] 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 [Sug01, Pra11], 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 [Woo09, MT12].

For the most part, we choose to sidestep 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. However, we provide a proof that cosheafification does exist in the next section.

2.5.4 The Existence of Cosheafification

Grothendieck gave us a general construction of the sheafification functor (-) that works in more general data categories **D**, which comes from applying a certain functor $(-)^+$ twice. The requirement on the category **D** includes, among other things, that filtered colimits and finite limits commute. If we were to regard a pre-cosheaf $\hat{F} : \mathbf{Open}(X) \to \mathbf{D}$ as a pre-sheaf valued in \mathbf{D}^{op} , then the condition that filtered colimits and finite limits commute in \mathbf{D}^{op} would become the condition that cofiltered limits and finite colimits commute in **D**, which is patently false when $\mathbf{D} = \mathbf{Set}$, **Vect** or **Ab**, as the example in Proposition 2.5.19 showed. Consequently, we do not have a clear answer to the question: Does the inclusion functor ι have a *right* adjoint $(-)_{\#}$?

$$\iota : CoShv(X; D) \hookrightarrow Fun(Open(X), D) \Longrightarrow PreCoshv(X; D)$$

so that the dual universal property is satisfied, i.e. if \widehat{F} is a pre-cosheaf and \widehat{G} is a cosheaf with a morphism $\widehat{G} \to \widehat{F}$, then there is a unique way of completing the diagram.



In the case where D = Set, Jon Woolf's paper [Wooo9] contains a construction of cosheafification. Unfortunately, this cannot be adapted to categories like Vect or Ab. For a high-level reason why, Mac Lane and Moerdijk explain on page 95 of [MM92] that a sheaf of abelian groups can be identified with an abelian group object in the category of sheaves. Since sheafification preserves finite products, sheafification of pre-sheaves of sets lifts to a functor between abelian group objects. Moreover, since the forgetful functor for : $Ab \rightarrow Set$ preserves limits (but not colimits), any sheaf of groups defines a sheaf of sets. Trying to repeat this last line of reasoning for cosheaves of groups fails, i.e. a cosheaf of groups does not, by forgetting, define a cosheaf of sets.

Our approach is to verify abstractly whether cosheafification exists without constructing it. Of course, one would like to use Freyd's Adjoint Functor Theorem 1.5.2, but we will use a different theorem [AR94, Thm 6.28] that is easier to check.

Theorem 2.5.20. Assuming Vopenka's principle (a large cardinal axiom), every full subcategory **B** of a locally presentable category **C**, where **B** is closed under colimits, is coreflective, i.e. the inclusion functor $\iota : \mathbf{B} \hookrightarrow \mathbf{C}$ has a right adjoint (a cofree functor).

We will leave Vopenka's principle as a black box and assume it, even though many category theorists cringe at its very name. We prove that cosheafification exists by verifying the hypotheses of the above theorem for our case of interest.

Corollary 2.5.21. The category of cosheaves of vector spaces is a coreflective subcategory of Fun(Open(X), Vect), i.e. cosheafification exists.

Proof. It is clear that the category of cosheaves is closed under colimits, since we can define the colimit to be the pre-cosheaf, which open-by-open assigns the colimit of vector spaces over that open set. This pre-cosheaf is a cosheaf, because for a fixed cover, each vector space in the diagram is expressed as a colimit and colimits commute with colimits.

It remains to be seen that the category of pre-cosheaves is locally presentable. this means that the category is locally small, has small colimits, has a small set of objects S that generates **PreCoshv**(*X*; **Vect**) in the sense that every pre-cosheaf is a colimit of objects in S, and every object is small. The first two statements are easily addressed. **Open**(*X*) is a small category and **Vect** is locally small, so the functor category is locally small. Colimits of pre-cosheaves are defined open-by-open. Since **Vect** is cocomplete, pre-cosheaves valued in **Vect** is also cocomplete. Now we address the existence of a generating set.

Define, for each open set $U \in Open(X)$, the following pre-cosheaf:

$$\hat{h}_{U}(V) = \left\{ \begin{array}{ll} k & \text{if } U \subset V \\ 0 & o.w. \end{array} \right.$$

We'd like to say that every pre-cosheaf is a colimit of pre-cosheaves of the above form. The corresponding statement for pre-sheaves is proved in pages 41-42 of [MM92]. We will go ahead and repeat the argument here. Note that if \hat{G} is an arbitrary pre-cosheaf, then

Hom_{PreCoshv}
$$(\hat{h}_{U}, \widehat{G}) \cong \widehat{G}(U).$$

Observe that if $U \subset U'$, then we get a map of pre-cosheaves (a natural transformation) $\hat{h}_{U'} \rightarrow \hat{h}_{U}$. This in turn induces a map

$$\operatorname{Hom}_{\operatorname{PreCoshv}}(\hat{h}_{U}, G) \to \operatorname{Hom}_{\operatorname{PreCoshv}}(\hat{h}_{U'}, G)$$

which coincides with the internal extension map of \widehat{G} , that is $r_{U',U} : \widehat{G}(U) \to \widehat{G}(U')$. In other words, the functor

$$\begin{split} &\mathsf{R}: \mathsf{PreCoshv}(X;\mathsf{Vect}) \to \mathsf{PreCoshv}(X;\mathsf{Vect}) \\ &\mathsf{G} \rightsquigarrow (\mathsf{U} \mapsto \mathrm{Hom}_{\mathsf{PreCoshv}}(\hat{\mathfrak{h}}_{\mathsf{U}},\widehat{\mathsf{G}}) \cong \widehat{\mathsf{G}}(\mathsf{U})) \end{split}$$

is isomorphic to the identity functor. Since adjoints are unique up to isomorphism, then we can conclude that its (left) adjoint must also be isomorphic to the identity functor. However, we will construct explicitly the adjoint, which, combined with the fact that it must be the identity functor, exhibits \hat{G} as the colimit of pre-cosheaves of the form \hat{h}_{U} .

For each pre-cosheaf \hat{G} , define the following **category of elements** $J_{\hat{G}}$.¹⁰ The objects of $J_{\hat{G}}$ are pairs

 (\mathbf{U}, \mathbf{x}) where $\mathbf{U} \in \mathbf{Open}(\mathbf{X})$ $\mathbf{x} \in \widehat{\mathbf{G}}(\mathbf{U})$.

¹⁰ It should be noted that in [MM92], the category of elements is written $\int \widehat{G}$.

A morphism $(U, x) \to (U', x')$ is defined if $U \subset U'$ and $x' = r_{U',U}(x)$. Clearly, there is a projection functor $\pi_{\widehat{G}} : J_{\widehat{G}} \to \mathbf{Open}(X)$ and by formality, there is a dual functor $\pi_{\widehat{G}}^{op} : J_{\widehat{G}}^{op} \to \mathbf{Open}(X)^{op}$.

Denote by \mathcal{Y} the functor

$$\mathcal{Y}: \mathbf{Open}(X)^{\mathrm{op}} \to \mathbf{PreCoshv}(X; \mathbf{Vect}) \qquad \qquad \mathcal{U} \rightsquigarrow \hat{\mathfrak{h}}_{\mathcal{U}}.$$

We claim that the left adjoint L to the functor R considered above can be constructed object-wise as follows: for each pre-cosheaf \widehat{G} define

$$\mathcal{L}(\widehat{\mathcal{G}}) := \varinjlim \mathcal{Y} \circ \pi_{\widehat{\mathcal{G}}}^{\operatorname{op}}.$$

We claim that \widehat{G} is the colimit. This is already given from the fact that L must be isomorphic to the identity functor, but let's at least check how \widehat{G} is a co-cone, to make the statement more plausible. For each object (U, x) in $J_{\widehat{G}}$ the map to \widehat{G} is defined by

$$x \in G(U) \cong \operatorname{Hom}_{\operatorname{\mathsf{PreCoshv}}}(\hat{h}_U, G) \ni \psi_{U,x}$$

where $\psi_{U,x}$ is the natural transformation that sends $1 \in \hat{h}_{U}(U)$ to $x \in \hat{G}(U)$ and then for any larger open set $U \subset U'$ sends $1 \in \hat{h}_{U}(U')$ to $r_{U',U}(x)$. Observe that if $U \subset U'$ and $(U', x') \to (U, x)$ is a morphism in $J_{\hat{G}}^{op}$, so $r_{U',U}(x) = x'$, then we have the following commutative diagram:



At the risk of demonstrating the obvious, the above diagram commutes if the the following diagram commutes for an arbitrary triple of open sets $V \subset V' \subset V''$. We will check it for the interesting boundary case $U \subset U' \subset U''$.

$$\begin{split} \hat{h}_{U'}(U'') &= k \xrightarrow{1} \hat{h}_{U}(U'') \stackrel{r_{U'',U'}(x)}{=} \widehat{G}(U'') \\ \hat{h}_{U'}(U') &= k \xrightarrow{1} \hat{h}_{U}(U') \stackrel{r_{U}}{=} k \xrightarrow{(x) = x'} \widehat{G}(U') \\ \hat{h}_{U'}(U) &= 0 \longrightarrow \hat{h}_{U}(U) = k \xrightarrow{x} \widehat{G}(U) \end{split}$$

This completes the plausibility check. We use the observation that L is isomorphic to the identity functor to conclude that \widehat{G} is actually the colimit. The conclusion is that

$$\widehat{\mathsf{G}} \cong \varinjlim^{\mathsf{y}} {}^{\mathsf{y}} \circ \pi^{\mathsf{op}}_{\widehat{\mathsf{G}}}$$

i.e. \widehat{G} is expressible as a small colimit of pre-cosheaves of the form \widehat{h}_{U} where the size of the indexing set is bounded by the product of the cardinality of **Open**(X) and the maximum cardinality of $\widehat{G}(U)$ over varying U.

Now it remains to check the smallness of objects in PreCoshv(X; Vect). An object \widehat{G} is small if there exists a regular cardinal κ such that $Hom(\widehat{G}, -)$ commutes with directed colimits of diagrams indexed by categories of cardinality at most κ .

Observe that for one of our pre-cosheaves \hat{h}_{U} is compact since if (F_i) is a direct system of pre-cosheaves, then

$$\varinjlim \operatorname{Hom}(\hat{h}_{U}, \widehat{F}_{i}) \cong \varinjlim \widehat{F}_{i} \cong \operatorname{Hom}(\hat{h}_{U}, \varinjlim \widehat{F}_{i}).$$

As already shown, for every pre-cosheaf \widehat{G} there exists a diagram whose cardinality is the cardinality of $J_{\widehat{G}}$, which we will call κ_J . Now, we know how to express \widehat{G} as a colimit of \hat{h}_{U} 's. Thus,

$$\begin{aligned} \operatorname{Hom}(\widehat{\mathsf{G}}, \varinjlim_{i} \widehat{\mathsf{F}}_{i}) &\cong \operatorname{Hom}(\varinjlim_{U} \widehat{\mathsf{h}}_{U}, \varinjlim_{i} \widehat{\mathsf{F}}_{i}) \\ &\cong \varprojlim_{U} \operatorname{Hom}(\widehat{\mathsf{h}}_{U}, \varprojlim_{i} \widehat{\mathsf{F}}_{i}) \\ &\cong \varprojlim_{U} \varinjlim_{i} \operatorname{Hom}(\widehat{\mathsf{h}}_{U}, \widehat{\mathsf{F}}_{i}) \\ &\cong \varprojlim_{i} \varinjlim_{U} \operatorname{Hom}(\widehat{\mathsf{h}}_{U}, \widehat{\mathsf{F}}_{i}) \\ &\cong \varinjlim_{i} \operatorname{Hom}(\varinjlim_{U} \widehat{\mathsf{h}}_{U}, \widehat{\mathsf{F}}_{i}) \\ &\cong \varinjlim_{i} \operatorname{Hom}(\varinjlim_{U} \widehat{\mathsf{h}}_{U}, \widehat{\mathsf{F}}_{i}) \\ &\cong \varinjlim_{i} \operatorname{Hom}(\widehat{\mathsf{G}}, \widehat{\mathsf{F}}_{i}) \end{aligned}$$

The third line follows from compactness of \hat{h}_{U} . The fourth line follows from the fact that in **Set**, κ_{J} small colimits commute with κ_{J} -filtered colimits. This completes the proof.

PRELIMINARY EXAMPLES

"The content of a mathematical theory is never larger than the set of examples that are thoroughly understood."

— Vladmir Arnol'd [Arno4]

Theories should be motivated by examples. In this chapter we develop the common themes these examples share. Broadly speaking, all sheaves are realized via local sections associated to a particular map. This principle is rigorously embodied by the *étalé perspective*. Similarly, all cosheaves of sets are realized by connected components of the fiber of a map, embodied by the *display perspective*, which is a generalization of the Reeb graph construction outlined in Definition 3.4.1.

3.1 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 definition has the property that f is 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 3.1.1 (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?"



To see why this is a pre-sheaf valued in D = 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

$$F(U) \rightarrow F[U] := \lim(N(U)^{op} \rightarrow Open(X)^{op} \rightarrow Set)$$

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 $\rho_{ij,i}(s_i) = \rho_{ij,j}(s_j)$.¹ The natural map from F(U) to F[U] simply takes a section $s \in F(U)$ to the collection of restricted sections $\{s_i := s|_{U_i}\}$. If two sections over U differ at a point x, then they will define different sections over $U_i \ni x$, thus the natural map is 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 Chapter 2.

Example 3.1.2. For a simple example, consider the projection onto the first coordinate $\pi_t : [0,1] \times [0,1] \rightarrow [0,1]$, 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 4 where E is the region bound between the two curves. Does it have any global sections, i.e. is $F(X) \neq \emptyset$?

The answer is clearly no. The example in Figure 4 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 3.1.3 (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

¹ Here we have adopted the shorthand of referring to open sets via elements of the nerve.



FIGURE 4: Is There a Section?

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 3.1.4. 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.



Here F is the fiber space, h_U is a homeomorphism and π_U is projection onto the first factor. If F is a discrete space then we usually write \tilde{X} instead of E and say 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 3.1.5 (Constant Sheaf). Suppose A is an R-module equipped with the discrete topology and $E = X \times A \rightarrow X$ is the product bundle. The sheaf of sections of this map is called the **constant sheaf** A_X . If A = 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 . We will almost always work over a field k.

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

For locally connected spaces the constant sheaf k_X assigns to any open set U the product of the field k for as many connected components as U has. To see this, let us be more precise about how the algebraic structure of the module/vector space interacts with the topological structure of a fiber bundle [Mac91, Sec. 7.2.1].

Definition 3.1.6 (Local Systems). Let k be a field viewed as a topological space with the discrete topology. A **local system** is a covering space $\pi : L \to X$ equipped with the structure of a k-vector space on each fiber. Specifically, an n-dimensional local **system** on a topological space X is a topological space L, a map of spaces $\pi : L \to X$, and, for each point $p \in X$, a k-vector space structure on $\pi^{-1}(p)$ with the following property: Every point $p \in X$ has a neighborhood U such that there is a homeomorphism $h_U : \pi^{-1}(U) \to U \times k^n$ such that $\pi_U \circ h_U = \pi$ and for each $x \in U$ the vector space structure on $\pi^{-1}(x)$ is induced by h from the one on $\{x\} \times k^n$.

Our definition of the locally constant sheaf F in Definition 3.1.5 is more accurately defined as an n-dimensional local system L, at least when A is a k-vector space. Consider one of the distinguished open neighborhoods U of a point $x \in X$ provided by the definition. Here we have a commutative diagram:



For a locally connected space we can assume U is connected by replacing U with whatever connected component of U contains x. Consider a section of $\pi : L \to X$ over U. Since k^n has the discrete topology any section s over U has to be constant since the image of a connected set is always connected. This implies that every section s has the form $h_U \circ s(y) = (y, (v_1, ..., v_n))$ for every $y \in U$ and for some fixed vector $\bar{v} \in k^n$. Consequently, for the distinguished neighborhood U, the set of sections

$$F(U) = A_X(U) \cong k^n \cong A.$$

Moreover, by the local system condition, we can form any linear combination of sections $s_1, s_2 \in F(U)$ to obtain a third section $\alpha s_1 + \beta s_2 \in F(U)$. This implies that the locally

constant sheaf is actually a sheaf valued in Vect — the category of vector spaces. Arguing in the same way for each connected component tells us that over a union U' of disjoint, connected, distinguished neighborhoods, a locally constant sheaf has the value

$$F(\mathbf{U}') \cong A^{\pi_0(\mathbf{U}')} \cong H^0(\mathbf{U}'; A) \cong H^0(\mathbf{U}'; k)^n.$$

This illustrates that a locally constant sheaf can be thought of as taking H^0 of a space with "twisted" coefficients.

3.2 LOCAL SYSTEMS: A BRIDGE BETWEEN SHEAVES AND COSHEAVES

Formulating locally constant sheaves as a twisted H^0 presents an obvious dualization in terms of H_0 . Moreover this duality reaches higher by considering higher cohomology and homology as well. To understand this, we will need to understand local systems better.

Definition 3.2.1. The collection of local systems over X forms a category Loc(X). A morphism of between two local systems $\pi : L \to X$ and $\pi' : L' \to X$ is a map $\varphi : L \to L'$ such that $\pi' \circ \varphi = \pi$ and the restricted map $\varphi_x : \pi^{-1}(x) \to \pi'^{-1}(x)$ is a linear transformation.

The following theorem is classical and allows us to use two definition of local systems interchangeably.

Theorem 3.2.2. If X is a locally connected and locally simply connected space, then the category of local systems is equivalent to the category of representations of the fundamental groupoid of X, i.e.

 $Loc(X) \simeq Rep(\pi_1(X)).$

Recall that the objects of $\operatorname{Rep}(\pi_1(X))$ are functors $\mathcal{L} : \pi_1(X) \to \operatorname{Vect}$.

Remark 3.2.3. For a connected space X fixing a base point x_0 provides a skeletal subcategory $\pi_1(X; x_0) \hookrightarrow \pi_1(X)$. Precomposing \mathcal{L} with this inclusion defines a representation of the fundamental group $\pi_1(X; x_0)$.

Proof (Idea). The functor that realizes this equivalence is very easy to describe. Given a local system L one defines a representation $\mathcal{L} : \pi_1(X) \to \text{Vect}$ by assigning to points $x \in X$, the vector space $\pi^{-1}(x) =: \mathcal{L}(x)$. Now suppose $\gamma : [0, 1] \to X$ is a path connecting x to y. Since $\pi : L \to X$ is a covering space, there is a unique lift $\tilde{\gamma}$ connecting any element of $\pi^{-1}(x)$ to an element in $\pi^{-1}(y)$. These lifts piece together to define a **monodromy map** $\mu_{\gamma} : \mathcal{L}(x) \to \mathcal{L}(y)$. Since local systems are fiber bundles, a homotopy of paths pulls back



FIGURE 5: Trivial Circle Bundle over the Circle

to a trivial bundle, which shows that the map μ_{γ} is invariant under homotopy classes of paths rel endpoints.

Moreover, one can construct a space associated to a functor $\mathcal{L} : \pi_1(X) \to \text{Vect}$ by considering the product $L := \prod_{x \in X} \mathcal{L}(x)$ and topologizing suitably. For example, one could consider a basis of open sets around a point $v \in \mathcal{L}(x)$ given by the collection of elements { $w \in L \mid \exists \gamma, \text{s.t.} \mu_{\gamma}(v) = w$ }. This construction mirrors the usual construction of a classifying space given in Hatcher [Hato2, Sec. 1.3] or Munkres [Munoo, Ch. 13].

This equivalence allows us to define plenty of examples of local systems coming from fiber bundles.

Proposition 3.2.4 (Fiber Bundles Give Local Systems). Suppose $\pi : E \to X$ is a fiber bundle, then for each i the homology of the fiber $H_i(\pi^{-1}(x);k)$ defines a local system. Dually, for each i the cohomology of the fiber $H^i(\pi^{-1}(x);k)$ defines a representation of the fundamental groupoid and consequently a local system.

Proof. This is easily seen because any path $\gamma : [0, 1] \to X$ determines a pullback bundle $\gamma^* E$ 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 $H^* E$.



FIGURE 6: Identification Spaces for the Torus and Klein Bottle

Let us now consider two examples. Firstly, in Figure 5 we drew a map from the torus to a circle. To define this map one considers the torus as the space $S^1 \times S^1$ and defines the map to be projection onto the first factor. Secondly, consider the analogous projection map from the Klein bottle to the circle. An identification space model is depicted for both of these maps are drawn in Figure 6 with the left hand side being the torus with its map and the right hand side being the Klein bottle with its map.

Consider the local system gotten by taking $H_1(-;k)$ of the fiber $\pi_T : T \to S^1$. Between any two points s and s' there are two homotopy classes of paths connecting them: one that in the identification space model proceeds directly from s to s' and one that wraps around the circle using the implied identification. Choosing a basis for the vector space $H_1(\pi_T^{-1}(-);k)$ involves choosing a cycle along with an orientation. If one considers the monodromy map associated to either path, one can see that in the bases indicated for $H_1(\pi_T^{-1}(s);k)$ and $H_1(\pi_T^{-1}(s');k)$ in Figure 7 both monodromies are trivial (i.e. the identity map $k \to k$) as indicated by the green arrows.

Now consider the local system gotten by taking $H_1(-;k)$ of the fiber for $\pi_K : K \to S^1$. Choosing the same bases as before the monodromy associated to the longer path that wraps around the identification space is non-trivial

$$H_1(\pi_K^{-1}(s');k) = k \xrightarrow{-1} k = H_1(\pi_K^{-1}(s);k)$$

as indicated by the red arrows in Figure 8.



FIGURE 7: Trivial Action with the Torus Map



FIGURE 8: Non-trivial Action from the Klein Bottle
We now will show that these examples actually provide examples of locally constant cosheaves, or sheaves if cohomology is taken. First, we will need some alternative definitions.

Definition 3.2.5. Let X be a locally connected space. A sheaf A_X or cosheaf \hat{A}_X on X valued in **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 \mathcal{F} is **locally constant** if for each point x there is an open neighborhood U such that F or \mathcal{F} is constant, i.e. there is a vector space A such that $F|_U \cong A_X$ or $\mathcal{F}|_U \cong \hat{A}_X$.

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 $F(U) \rightarrow F_x$ or $\mathcal{F}_x \rightarrow \mathcal{F}(U)$ 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 $\mathcal{F}(U)$ respectively. By chaining together these sorts of isomorphisms, one can show the following theorem:

Theorem 3.2.6. 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.

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

Similarly, any locally constant cosheaf valued in 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 \mathcal{F}_x , respectively. Since the theorem is well known (see [Acho7] for a proof, which we follow here) for sheaves we present the cosheaf-theoretic proof instead.

Call a subset K of X *fine*² for a cosheaf \mathcal{F} if it is connected and is contained in a connected open set V such that $\mathcal{F}|_V$ is a constant cosheaf. For any set of points $\{x_i\}$ in a fine set K we have a collection of isomorphisms



² In [Acho7] they use the word "good."

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

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

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

Now we claim that given a path $\gamma : [0,1] \rightarrow X$ 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 $\rho(\gamma) : \mathcal{F}_{\gamma(0)} \to \mathcal{F}_{\gamma(1)}$ to be the composite

$$\mathcal{F}_{\gamma(\mathfrak{a}_0)} \to \mathcal{F}_{\gamma(\mathfrak{a}_1)} \to \cdots \to \mathcal{F}_{\gamma(\mathfrak{a}_n)}$$

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 $H : [0,1] \times [0,1] \rightarrow X$ 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 $\mathcal{F}_{\gamma(0)} \rightarrow \mathcal{F}_{\gamma(1)}$ and $\mathcal{F}_{\gamma'(0)} \rightarrow \mathcal{F}_{\gamma'(1)}$ are the same.

Moreover, one can show that representations of the fundamental groupoid give rise to locally constant sheaves and cosheaves. This will require a slightly more sophisticated version of van Kampen's theorem found in [Bro67, May99].

Proposition 3.2.7 (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 π_1 : **Open**(X) \rightarrow **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 2.4 imply that the fundamental groupoid is a cosheaf.

Theorem 3.2.8. A representation of a fundamental groupoid $\mathcal{L} : \pi_1(X) \to \text{Vect}$ determines a locally constant sheaf and a locally constant cosheaf respectively.

Proof. Again the sheaf theoretic version of this statement is well known (see [Acho7]), so we carry out the cosheaf version. Assume we have a local system \mathcal{L} , then we define the associated cosheaf to be

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

which is a cosheaf on account of the fact that colimits commute with colimits. The fact that \widehat{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 3.2.9 (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 2.3.4, 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 [Whi78, Ch. 6] and thus Mayer-Vietoris [Spa94, Ch. 4.6]. To complete our alternative proof of the above theorem, we check one more hypothesis of Theorem 2.3.4. 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.

This establishes the following corollary, justifying that local systems do indeed form a bridge between locally constant sheaves and cosheaves.

Corollary 3.2.10. With the above topological assumptions on X, the category of locally constant sheaves and locally constant cosheaves are equivalent.

We leave the details to the reader, who should note that any local system \mathcal{L} defines a sheaf by taking $H^0(-;\mathcal{L})$ and a cosheaf by taking $H_0(-;\mathcal{L})$.

3.3 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 3.3.1 (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.

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

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 3.3.2 (Feature Function). Suppose we have a topological space X, populated with features of interest, expressed as a function $P : \{1, ..., n\} \rightarrow X$. We get a cosheaf of sets via $\widehat{F}(U) = P^{-1}(U)$. A slightly different cosheaf is gotten by letting $\widehat{G}(U) = U \cap im(P)$, 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 3.3.3 (Skyscraper Cosheaf). Suppose $x \in X$ and V is an k-vector space. Let's define the **skyscraper cosheaf** at x with value V to be

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

When V = k, we drop the superscript for notational convenience.



FIGURE 9: Topological Model for Skyscraper Cosheaf

A topological incarnation for the skyscraper is depicted in Figure 9. Here one makes the assignment

$$\mathbf{U} \rightsquigarrow \mathbf{H}_{0}(\mathbf{f}^{-1}(\mathbf{U});\mathbf{k})$$

where f is the map that maps the circle to the point x, i.e. the constant map with value x.

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 $f^{\circ} : Open(X) \to Open(Y)$. 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).

Definition 3.3.4 (Cosheaf of Connected Components). 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

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

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 π_0 : **Top**_{lc} \rightarrow **Set** is left adjoint to the discrete space functor and so it preserves colimits [Woo09].

Example 3.3.5 (The Square Map, Again). Consider the cosheaf of connected components associated to the map $f : S^1 \to S^1$ defined in complex coordinates as $f(z) = z^2$. For every point $p \in S^1$ there are two connected components in the fiber over p. However, there is only one connected component over the whole of S^1 . This illustrates a sort of "twisted" H_0 already alluded to.

Exercise 3.3.6. Work out the cosheaf of connected components associated to the map $\pi: E \to X$ found in Figure 4.

The next example provides a derived version of the principle that H_0 is a cosheaf.

Example 3.3.7 (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

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

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

$$\hat{C}_{p}(U) := \lim C_{p}(U)$$

where the colimit is being performed over iterated subdivision, then we obtain a cosheaf [Bre97].



FIGURE 10: Barycentric Subdivision of a Singular Chain

Remark 3.3.8 (Mayer-Vietoris and Cosheaves). Another way of seeing that singular pchains do not define a cosheaf is to recall that the proof of the **Mayer-Vietoris** theorem starts with the observation that the sequence

$$0 \to C_{\mathfrak{p}}(U \cap V) \to C_{\mathfrak{p}}(U) \oplus C_{\mathfrak{p}}(V) \to C_{\mathfrak{p}}(U+V) \to 0$$

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 $D = K^b(\text{Vect}_k)$ by assigning

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

and this will be a cosheaf.³ The category $K^b(\text{Vect}_k)$ 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 3.3.9 (Compactly Supported Functions). Suppose X is a locally compact Hausdorff space. Consider the following assignment:

$$\Omega_{c}^{0}: U \rightsquigarrow \{f: U \to \mathbb{R} | \operatorname{supp}(f) \text{ compact} \}$$

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_c^p for $p \ge 0$.

3.4 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

³ The author has recently learned that Jacob Lurie calls this a homotopy cosheaf.

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 4, it seems unlikely. On a small open set the sheaf of sections is in bijection with the set

{f: $(x - \epsilon, x + \epsilon) \rightarrow (a, b)$ | 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 3.4.1 (Reeb Graph). Suppose Y is a topological space and $f : Y \to \mathbb{R}$ is a continuous map. The **Reeb graph** [Ree46] is defined to be the quotient space $R(f) := Y/ \sim$ where $y \sim y'$ if and only if y and y' belong to the same connected component of the fiber $f^{-1}(t)$.



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 from Figure 4. 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 11: Sheaf of Sections

FIGURE 12: Cosheaf of Components

→ •

Part II

LINEAR ALGEBRA OVER CELL COMPLEXES

In this part we emphasize that cellular sheaves and cosheaves are nothing more than linear algebra parametrized by a cell complex. The use of the term "sheaf" is justified by the Alexandrov topology, which makes functors out of posets into sheaves or cosheaves. Explicit proofs are presented since the primary reference of Shepard [She85] is unpublished and not easily accessed. Cellular sheaf cohomology and cosheaf homology are presented computationally in Chapter 6 and then put on the firm foundation of derived categories in Chapter 7. The novel contributions from this part, aside from working explicitly with cosheaves, are the introduction of "barcodes" to interpret cellular sheaf cohomology and cosheaf homology and exploiting the existence of enough projectives for cellular sheaves to define sheaf *homology*.

CELLULAR SHEAVES AND COSHEAVES

"Sheaf theory is [where] you do topology horizontally and algebra vertically."

- attributed to Maurice Auslander by [Gra79]

We can take it as an experimental observation from Figures 11 and 12 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 may be non-standard.

4.1 CELL COMPLEXES AND THE FACE-RELATION POSET

Definition 4.1.1 (Regular Cell Complex [Mac14b]). 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:¹ 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 $\sigma \leq \tau$.
- 4. The pair $X_{\sigma} \subset \bar{X}_{\sigma}$ is homeomorphic to the pair $int(B^k) \subset B^k$, 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 4.1.2 (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

¹ The **frontier** of a subspace A is the complement of A in its closure, i.e. $fr(A) := \overline{A} - A$. In some forms this axiom reads: if $X_{\sigma} \neq X_{\tau}$ and $X_{\sigma} \cap \overline{X}_{\tau} \neq \emptyset$ then X_{σ} is contained in the frontier of X_{τ} .

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.

Definition 4.1.3 (Cell Complex [She85, Mac14b]). 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 4.1.4. 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.

Definition 4.1.5 (Cell category). To a cell complex $(X, \{X_{\sigma}\}_{\sigma \in P_X})$ we can associate a category **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 **Cell**(X) or X.

We now introduce diagrams indexed by the cell category. These were defined in Shepard's 1985 thesis [She85, p. 6], but were known as "stacks" in the first published volume of Zeeman's 1954 thesis [Zee62a, p. 626]. However, this term came to be used for an entirely different construction in the landmark paper of Pierre Deligne and David Mumford [DM69], which introduced the fundamental concept of algebraic and moduli stacks for algebraic geometry. Zeeman's usage of the term is now extinct, but his work anticipates MacPherson's cellular perverse sheaves (cf. Definition 6.3.25) although MacPherson was unaware [Mac14a] of the content of Zeeman's thesis [Zee62a, Zee62b, Zee63]. To keep the presentation simple, we give Shepard's definition and its appropriate dualization.

Definition 4.1.6 (Cellular Sheaves and Cosheaves). A **cellular sheaf** F valued in D on X is a functor $F : Cell(X) \rightarrow D$, i.e. it is

• an assignment to each cell X_{σ} in X an object $F(\sigma)$,



FIGURE 13: Bott's Height Function on the Torus

• and to every pair of incident cells $X_{\sigma} \subset \bar{X}_{\tau}$ a **restriction map**² $\rho_{\sigma,\tau}^{F} : F(\sigma) \to F(\tau)$.

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

Let us consider a few natural examples.

Example 4.1.7 (Bott's Torus). The following example was first popularized by Raoul Bott in his book on Morse theory [Mac14a]. Consider the height function on the torus, rotated by 90^{circ} so that the real line is underneath the torus, as shown in Figure 13. By taking the pre-image of the star of each cell, one obtains a diagram of spaces $\hat{F} : X^{op} \rightarrow \text{Top}$. By post-composing this diagram with $H_1(-;k)$, one obtains the cellular cosheaf indicated in Figure 14.

Example 4.1.8 (Klein Bottle Revisited). As seen in Section 3.2, a Klein bottle can be viewed as a non-trivial S¹ bundle over the circle. We've already seen how the this leads to

² 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.



FIGURE 14: Cellular Cosheaf by Taking H1 of the Star pre-images

a representation of the fundamental groupoid $\pi_1(S^1)$. 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, as in Figure 8. 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 a cell structure on S¹:



Example 4.1.9. Let Y = (0, 1) be the open unit interval in \mathbb{R} . Denote the inclusion of Y into \mathbb{R} by j. To the constant sheaf on Y, written k_Y , we can associate two sheaves on \mathbb{R} :



FIGURE 15: Cellular Description of j_*k_Y and $j_!k_Y$ respectively

 j_*k_Y and $j_!k_Y$. To describe these sheaves, we can think cellularly. For the first sheaf j_*k_Y the cellular sheaf is simply the diagram of vector spaces

$$0 \leftarrow k \rightarrow k \leftarrow k \rightarrow 0.$$

For j_1k_Y the cellular sheaf is the diagram of vector spaces

$$0 \leftarrow 0 \rightarrow k \leftarrow 0 \rightarrow 0.$$

The key difference being that the latter sheaf is zero on the endpoints $\{0\}$ and $\{1\}$. However, one can recover the classical, open set description of the sheaves j_*k_Y and $j_!k_Y$ by considering any ordinary (Hausdorff) open set on the real line and then computing the limit of the diagram that lies over the open set. In Figure 15 we have drawn the two cellular sheaves of interest and the value of the limit over each open set. Of course, one can dualize the discussion and consider cosheaves instead and use colimits to get functors from the open set category as defined in Chapter 2. This perspective is developed more fully in Section 11.2.3.

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

Definition 4.1.10. We denote the category of cellular sheaves on X by

$$Shv(X; D) := Fun(Cell(X), D)$$

and the category of cellular cosheaves by

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

Morphisms are natural transformations of functors. If D = Vect, then we will omit the notation after the semicolon and write Shv(X) and 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 in Section 4.2 in one way, and in Chapter 11 in an entirely different way.

4.2 PARTIALLY ORDERED SETS: FINITE SPACES AND FUNCTORS

Cellular sheaves and cosheaves earn their finiteness by assigning data directly to cells, rather than open sets. This turns out to not be entirely true; Cellular sheaves and cosheaves are simply operating on a different topology than the one we are accustomed to. 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 [Lado8], but we streamline the discussion by using Kan extensions, which clarifies how cosheaves on a poset X differ from sheaves on X^{op}.

4.2.1 The Alexandrov Topology

In this section we introduce a class of non-Hausdorff spaces called **Alexandrov spaces**. The reader should note that although this topology is non-Hausdorff, it is highly relevant to concepts in algebraic topology. There is a remarkable theorem due to Michael McCord [McC66] 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 even gives constructions of classical operations in algebraic topology, including suspension, in the Alexandrov setting. However, our ambitions for this section are far more limited. Let us begin with the necessary definitions.

Definition 4.2.1. A **pre-order** consists of a set P and a relation \leq that is reflexive and transitive. A **poset** is a pre-order where the relation is also anti-symmetric, i.e. $x \leq y$ and $y \leq x$ implies x = y. A map f of pre-orders is one that respects \leq . That is if $x \leq y$ then $f(x) \leq f(y)$. Pre-orders and order preserving maps form a category **Preorder**. The collection of all posets form a subcategory of this category.

Every pre-order can be equipped with a topology. However, it was first defined for finite posets by Pavel Alexandrov [Ale₃₇, Ale₄₇] and the general definition carries his name.

Definition 4.2.2. On a pre-order (P, \leq) define the **Alexandrov topology** to be the topology whose open sets are the sets that satisfy the following property:

$$x \in U$$
 $x \leqslant y$ \Rightarrow $y \in U$

A basis is given by the sets of the form $U_x := \{y \in P | x \leq y\}$ — what we will call the **open star at** *x***.** Similarly, we define the **closure of** *x* by $\bar{x} := \{y \in P | y \leq x\}$. When P is a finite poset, then a basis of closed sets is given by the \bar{x} 's.

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 $x \leq y$ and $y \leq x$. One can check that this surjection is order-preserving. This construction defines a right adjoint to the inclusion of posets into pre-orders [Woo09].

Remark 4.2.3 (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 4.2.4. Consider (\mathbb{R}, \leq) with the usual partial order. The open sets are all those open or half open intervals such that the right-hand endpoint is $+\infty$. Observe that the closed set $(-\infty, 0)$ cannot be written as an intersection of closed sets of the form \overline{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 4.2.5 (Star). Let $(X, \{X_{\sigma}\})_{\sigma \in P_X}$ 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 $X_{\sigma} \leq X_{\tau}$.

$$star(X_{\sigma}) := \{X_{\tau} \mid X_{\sigma} \leq X_{\tau}\}$$

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 $\sigma \leq \tau$.

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.



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 16: Alexandrov Space Associated to the Unit Interval

Example 4.2.6 (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 16, we have drawn the basic open sets.

4.2.2 Functors on Posets

We want to understand how data modeled on posets can be treated as a sheaf or cosheaf on the Alexandrov topology. To do so we use the elegant, but sophisticated, 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.

$$\iota: P \to \textbf{Open}(P)^{op} \qquad p \mapsto U_p.$$

A natural question to ask is

"Given a functor $F : P \to D$, is there a consistent way of extending F to a functor $R : Open(P)^{op} \to D$?"

One can hope to perform this extension since the image of the inclusion $\iota : P \rightarrow Open(P)^{op}$ 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 $U_p = \{p' | p \leq p'\}$ 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.

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 4.2.7 (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 [Lado8], but adopting this language will be powerful and will make certain categorical properties transparent.

Definition 4.2.8 (Kan Extensions). Suppose **B**, **C** and **D** are categories, $F : \mathbf{B} \to \mathbf{D}$ and $E : \mathbf{B} \to \mathbf{C}$ are functors, then the **right Kan extension of** F **along** E written $R = \operatorname{Ran}_E F : \mathbf{C} \to \mathbf{D}$ is a functor and a natural transformation $\epsilon : RE \to F$ that is universal in the following sense. For every functor $H : \mathbf{C} \to \mathbf{D}$ with a natural transformation $\alpha : H \circ E \to F$ there exists a unique natural transformation $\sigma : H \to R$, i.e. $\operatorname{Nat}(H, R) \cong \operatorname{Nat}(H \circ E, F)$.



The **left Kan extension of F along E** written $L = Lan_EF : C \rightarrow D$ is a functor with a natural transformation $\eta : F \rightarrow L \circ E$ that is universal as well. If $H : C \rightarrow D$ is a functor with a natural transformation $\omega : F \rightarrow H \circ E$, then there exists a unique $\tau : L \rightarrow H$, i.e. $Nat(L, H) \cong Nat(F, H \circ E)$.



Remark 4.2.9 (Existence of Kan Extensions). Kan extensions do not always exist, but we have already alluded to a situation where they do. If **D** has enough limits and colimits, then we can give point-wise formulae for the left and right Kan extensions respectively:

 $Lan_{E}F(c) := \varinjlim_{E(b) \to c} F(b) \qquad Ran_{E}F(c) := \varprojlim_{c \to E(b)} F(b)$

One of the reasons that sheaves and cosheaves on Alexandrov spaces are so wellbehaved is that every open set has a finest cover, so in particular, by Corollary 2.4.4, 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 4.2.10. Let P be a poset and **D** a category that is both complete and co-complete. Then the following categories are equivalent

$$\operatorname{Fun}(\mathsf{P},\mathsf{D}) \cong \operatorname{Shv}(\mathsf{P};\mathsf{D}) \qquad \operatorname{Fun}(\mathsf{P}^{\operatorname{op}},\mathsf{D}) \cong \operatorname{CoShv}(\mathsf{P};\mathsf{D})$$

Proof. We claim that taking the right Kan extension of $F : P \to D$ along the inclusion $\iota : P \to Open(P)^{op}$ produces a sheaf. Suppose U is an open set in the Alexandrov

topology, i.e. one for which $p \in U$ and $p \leq p' \Rightarrow p' \in U$. It is true that every open set can be expressed as a union $U = \bigcup_{p \in U} U_p$ 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 2.4.4, 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 $\hat{F} : P^{op} \to D$ along the inclusion $\iota : P^{op} \to Open(P)$ to get a cosheaf. Taking costalks returns a diagram from a cosheaf.

Remark 4.2.11 (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.

$$F_p := \varinjlim_{U \ni p} F(U)$$
 and $\widehat{F}_p := \varinjlim_{U \ni p} \widehat{F}(U)$

In both cases when P is a poset with the Alexandrov topology there is a smallest open set containing p, namely $U_p = \{q | p \leq q\}$, so $F_p \cong F(U_p) = F(p)$ and $\widehat{F}_p = \widehat{F}(U_p) = \widehat{F}(p)$.

Definition 4.2.12 (Sections). Let (P, \leq) be a poset and $F : P \to D$ a sheaf and $\widehat{F} : P^{op} \to D$. let $Z \subset P$ be any subset. We define the **sections over** Z to be

$$\Gamma(Z;F) := \lim_{Z \to 0} F|_Z$$
 and $\lim_{Z \to 0} \widehat{F}|_Z =: \Gamma(Z;\widehat{F}).$

When Z = P, we call these **global sections**. Note that $\Gamma(Z; -)$ 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 4.2.10 is specialized to the face relation poset P_X of a cell complex, also called the cell category $P_X = \text{Cell}(X)$ in Definition 4.1.3, we get that the category of sheaves in Definition 4.1.10. We summarize these observations in the following corollary.

Corollary 4.2.13. Let (X, P_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 $F : P_X \to D$. A **cellular cosheaf on** X is a cosheaf on P_X with the Alexandrov topology. Such a cosheaf is uniquely determined by a functor $\widehat{F} : P_X^{op} \to D$.

To close, we point out one of the symmetries that Alexandrov spaces possess.

Claim 4.2.14. 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 $F : P \to D$ 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 5.1.4.

FUNCTORS ASSOCIATED TO MAPS

"Qui sème le foncteur récolte la structure."1

— Bourbaki

Since sheaves and cosheaves as defined here assign data to open sets, maps between spaces should only make reference to open sets. In Section 2.5.2 we briefly introduced how to pushforward or pullback a sheaf along a map between spaces. 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 process has historically been a stumbling block for the theory.

Recall that the definition of a continuous map $f : X \rightarrow Y$ 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

$$\operatorname{\check{f}}: \operatorname{\mathbf{Open}}(\operatorname{Y}) \to \operatorname{\mathbf{Open}}(\operatorname{X}) \qquad \operatorname{U} \subseteq \operatorname{Y} \rightsquigarrow \operatorname{f}^{-1}(\operatorname{U}) \subseteq \operatorname{X}.$$

By formality, we also have a functor from the corresponding opposite categories

$$\mathbf{\mathring{f}}: \mathbf{Open}(\mathbf{Y})^{\mathrm{op}} \to \mathbf{Open}(\mathbf{X})^{\mathrm{op}}.$$

^{1 &}quot;Who sows the functor reaps the structure."

We have purposely suppressed the op superscript on 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



has a natural completion given by pre-composition.

Definition 5.0.15 (Pushforward Sheaf and Cosheaf). Suppose $f : X \to Y$ is a continuous map of spaces, G and \widehat{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 $f_*\widehat{G}(U) := \widehat{G}(f^{-1}(U))$. Because f^{-1} commutes with unions, we get that if G or \widehat{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

 $f_*: \textbf{Shv}(X; \textbf{D}) \rightarrow \textbf{Shv}(Y; \textbf{D}) \qquad f_*: \textbf{CoShv}(X; \textbf{D}) \rightarrow \textbf{CoShv}(Y; \textbf{D}).$

There is also a pullback 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 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?

$$Open(Y)^{op} \xrightarrow{G} D$$

$$\stackrel{f}{\downarrow} \xrightarrow{?} Open(X)^{op}$$

Again, by assuming that **D** has sufficient colimits, we can could fill in the diagram by taking the left Kan extension of G along f and that will yield the candidate formula for the pullback just presented. Unfortunately, this definition for the pullback of a sheaf does not always define a sheaf. In sheaf theory over general spaces, this defect is circum-

vented by sheafification, as in Section 2.5.2. Fortunately, for the Alexandrov topology, this circumvention is unnecessary.

5.1 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 5.1.1 (Map of Posets). Suppose (X, \leq_X) and (Y, \leq_Y) are posets. A map of posets is a map of sets $f : X \to Y$ that is order-preserving, i.e. if $x \leq_X x'$ then $f(x) \leq_Y f(x')$. 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, \leq_X) by just X.

Remark 5.1.2 (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 [She85], 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	Х	X
Underlying Alexandrov Space	P _X	Х
Set of Points in a "Cell"	X_{σ}	$ \sigma $
"Cell" viewed as a point	σ	σ
Cellular Sheaf	$F:P_X\to \boldsymbol{D}$	$F:X\to \mathbf{D}$

5.1.1 Pullback or Inverse Image

Recall that a sheaf on a poset (Y, \leq_Y) is a functor $G : Y \to D$. Similarly, a cosheaf is a functor $\widehat{F} : Y^{op} \to D$. For both structures the pull-back functor f^* is easily described. It is the obvious pre-composition that completes the following diagram.



Definition 5.1.3 (Pullback for Poset Maps). Given a sheaf G on Y and a map of posets $f : X \rightarrow Y$, 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 $\rho_{x',x}^{f^*G} = \rho_{f(x'),f(x)}^G$
- If $\eta : G \to H$ is a morphism in Shv(Y), i.e. a natural transformation of diagrams over Y, then $f^*\eta : f^*G \to f^*H$ is a morphism in 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 $r_{x,x'}^{f^*G} = r_{f(x),f(x')}^G$. Thus, we get functors

$$f^*: \mathsf{Shv}(Y; \mathsf{D}) \to \mathsf{Shv}(X; \mathsf{D}) \qquad f^*: \mathsf{CoShv}(Y; \mathsf{D}) \to \mathsf{CoShv}(X; \mathsf{D}).$$

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

$$f^*F(U_x) := \lim_{V \supset f(U_x)} F(V) \cong F(V_{f(x)}) = F(f(x)),$$

where we have used the fact that the smallest open set containing $f(U_x) = f(\{x' | x \leq x'\})$ is $V_{f(x)} = \{y | f(x) \leq y\}$.

Example 5.1.4 (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 $\hat{W}_X := p^*W$.

5.1.2 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 5.1.5 ([She85] 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 5.1.6. 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 $\sigma' \leq \tau'$, then either $s(\sigma') = s(\tau')$ or not. If not, then σ' and τ' belong to the subdivision of two cells $\sigma \leq \tau$. \Box

We are going to use this fact to define the subdivision of a sheaf in a cleaner manner than is found in [She85].

Definition 5.1.7. Suppose F is a sheaf on X and $s : X' \to X$ is a subdivision of X, then we define the subdivided sheaf $F' := s^*F$.

5.1.3 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 5.1.8 (Pushforward for Poset Maps). Given a sheaf F on X and a map of posets $f : X \rightarrow Y$ 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. Ran_fF.

$$f_*F(y) = \varprojlim_{f(x) \ge y} F(x)$$

• Suppose $y \leq y'$, then $\{x|f(x) \geq y'\} \subseteq \{x|f(x) \geq y\}$. 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 $f_*F(y) \rightarrow f_*F(y')$ that we will call $\rho_{y',y}^{f_*F}$.

Suppose η : F → 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 $f^{op} : X^{op} \to Y^{op}$.

$$f_*\widehat{F}(y) := \varinjlim_{f(x) \ge y} \widehat{F}(x).$$

Since both of these constructions are functorial, we have redefined two functors:

$$f_*: Shv(X; D) \rightarrow Shv(Y; D)$$
 $f_*: CoShv(X; D) \rightarrow CoShv(Y; D)$

Example 5.1.9 (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 $p : X \to *$ we offer the following definitions:

$$p_*F(\star) \cong F(X) = \Gamma(X;F) = H^0(X;F)$$
 $p_*\widehat{F}(\star) \cong \widehat{F}(X) = \Gamma(X;\widehat{F}) = H_0(X;\widehat{F})$

In Section 7 we will use this definition as the prototype for defining "higher" pushforward or direct image functors.

5.1.4 f⁺, 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 5.1.10 (Caveat). 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 ([Fri11] p. 22) the tradition goes back to Grothendieck in [AGV72] SGA Exposé I, Proposition 5.1. The same notation is used by Ladkani [Lado8], 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 5.1.11 (Pushforward with Open Supports). Given a sheaf F on X and a map of posets $f : X \rightarrow Y$ 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. Lan_fF.

$$f_{\dagger}F(y) = \varinjlim_{f(x) \leqslant y} F(x)$$

- If $y \leq y'$, then $\{x|f(x) \leq y\} \subseteq \{x|f(x) \leq y'\}$ and since any colimit over the bigger set defines a cocone over the smaller set by restriction, we get a unique map $\rho_{y',y}^{f_{\dagger}F}$: $f_{\dagger}F(y) \rightarrow f_{\dagger}F(y')$.
- If we have a map of sheaves $\eta : F \to G$, then pre-composing the arrows for colimG 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}.

$$f_{\dagger}\widehat{F}(y) := \varprojlim_{f(x) \leqslant y} \widehat{F}(x)$$

Both of these constructions are functorial and thus we have defined two functors:

$$f_{\dagger}: \textbf{Shv}(X; \textbf{D}) \rightarrow \textbf{Shv}(Y; \textbf{D}) \qquad f_{\dagger}: \textbf{CoShv}(X; \textbf{D}) \rightarrow \textbf{CoShv}(Y; \textbf{D})$$

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 $x' \leq x$, 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:

$$j: (X, \leq) \rightarrow \mathbf{Closed}(X) \qquad x \rightsquigarrow \bar{x} := \{x' | x' \leq x\}$$

Consequently, we have a similar diagram for a functor $F : X \to D$ as before.



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 4.2.4 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.

$$\frac{\text{Closed}(X) \xrightarrow{F} D}{f^{c} \uparrow f^{\dagger}F}$$

$$\frac{f^{c}}{\text{Closed}(Y)}$$

In Section 7.4 this functor provides the foundation for defining **sheaf homology** and **cosheaf cohomology** — theories that don't exist for general spaces.

5.1.5 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 f₁ will only be defined for cellular maps, which are stratified (or even definable) maps naturally adapted to cell complexes.

Definition 5.1.12 (Cellular Map [She85] pg. 32). Let X and Y be cell complexes. A **cellular map** (|f|, f) consists of a map of posets f : X \rightarrow Y and a continuous "geometric" map $|f| : |X| \rightarrow |Y|$ 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 $y, z \in |f(\sigma)|, |f|^{-1}(y) \cap |\overline{\sigma}|$ is compact if and only if $|f|^{-1}(z) \cap |\overline{\sigma}|$ is.

Remark 5.1.13. 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 5.1.14. 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:



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_1 uses kernels and other standard linear algebra operations. As such, we now assume D =Vect and suppress it from our notation.

Definition 5.1.15 (Pushforward with Compact Supports). Given a sheaf F on X and a cellular map $f : X \rightarrow Y$, we can define the **pushforward with compact supports** sheaf on Y as follows:

- $f_!F(\tau) = \{s \in \Gamma(f^{-1}(\tau); F) \mid s(\sigma) = 0 \text{ if } |\overline{\sigma}| \cap f^{-1}(y) \text{ not compact for } y \in |\tau| \}$
- Let $\gamma \leq \tau$ be cells in Y, and let $s \in f_!F(\gamma)$ and $t \in f_!F(\tau)$. We define $\rho_{\tau,\gamma}^{f_!F}(s) = t$ if for every $\sigma \in f^{-1}(\tau)$ and every $\lambda \in f^{-1}(\gamma)$ such that $\lambda \leq \sigma \rho_{\sigma,\lambda}^{F}(s(\lambda)) = t(\sigma)$. If there is no such $t \in f_!F(\tau)$ then we define $\rho_{\tau,\gamma}^{f_!F}(s) = 0$.

The notation $\Gamma(-; F)$ for sections is explained in Definition 4.2.12. 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 [She85] pp. 35-38. As such we have defined a functor

$$f_!: \mathbf{Shv}(X) \to \mathbf{Shv}(Y)$$

Remark 5.1.16 (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 $\hat{F} : X^{op} \rightarrow \text{vect}$, 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.

5.2 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 D =Vect and leave it as implicit that all operations are to be performed in vector spaces.

The notation \Box will be a place holder for any one of the three symbols *, †, !.

5.2.1 *Projection to a point*

We consider the constant map $p : X \to \star$. The output of $p_{\Box}F$ is a single vector space, namely $p_{\Box}F(\star)$.

Without too much effort we compute the following:



FIGURE 17: Projection to a Point

- $p_*F = \underline{\lim} \{F(x') \to F(a')\} \cong F(x')$
- $p_{\dagger} = \lim \{F(x') \to F(a')\} \cong F(a')$

For the pushforward with compact supports, we will be extra careful. Recall the definition states that $p_!F(\tau) = \{s \in \Gamma(p^{-1}(\tau); F) | s(\sigma) = 0 \text{ if } |\overline{\sigma}| \cap p^{-1}(y) \text{ not compact for } y \in |\tau|\}.$

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.

• $p_!F = \ker(\rho_{x',a'}: F(x') \to F(a'))$

5.2.2 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.



FIGURE 18: Inclusion into a Closed Interval

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, $j_*F(y) \cong F(a')$.
- $j_{\dagger}F(x) = F(x'), j_{\dagger}F(a) \cong F(a'), \text{ and } j_{\dagger}F(y) = \varinjlim\{\emptyset\} = 0.$

•
$$j_!F \cong j_{\dagger}F$$
.

5.2.3 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 19: Map to a Circle

All three sheaves agree on the values and the restriction maps $f_{\Box}F(y) \cong F(y') \rightarrow F(a') \cong f_{\Box}F(a)$. 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, $f_*F(x) = \varprojlim \{F(x') \rightarrow F(a') | F(b')\} \cong F(x') \oplus F(b')$ and $f_*F(b) = F(b')$.
- Here $f_{\dagger}F(x) = F(x')$, but for similar reasons as before $f_{\dagger}F(b) = \varinjlim\{F(x') | F(y') \rightarrow F(b')\} = F(x') \oplus F(b')$.
- As noted, f is injective thus the value of f_1F 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 $x' \nleq a'$ so the restriction map must be zero.

5.3 THE PUSH-PULL ADJUNCTIONS

Recall from Section 1.5 that adjunctions allow us transform a complicated problem into an easy one. To derive these adjunctions, we can take two approaches: Use Freyd's adjoint functor theorem 1.5.2, or explicitly construct the adjunction. 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 5.3.1.** The functors $f^* : Shv(Y) \to Shv(X)$ and $f_* : Shv(X) \to Shv(Y)$ form an adjoint pair (f^*, f_*) and thus

$$\operatorname{Hom}_{\operatorname{Shv}(X)}(f^*G, F) \cong \operatorname{Hom}_{\operatorname{Shv}(Y)}(G, f_*F).$$

Dually, the functors for cosheaves satisfy the opposite adjunction (f_*, f^*)

$$\operatorname{Hom}_{\operatorname{\mathsf{CoShv}}(Y)}(f_*\widehat{F},\widehat{G})\cong\operatorname{Hom}_{\operatorname{\mathsf{CoShv}}(X)}(\widehat{F},f^*\widehat{G}).$$

Proof. Recall that $f^*(f_*F)(x) = (f_*F)(f(x))$. Using the fact that $(f_*F)(f(x)) = \varprojlim \{F(z)|f(z) \ge f(x)\}$, 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 $f^* \circ f_* \to id$, which is universal (final).

Similarly, $f_*(f^*G)(y) = \lim_{G \to G} \{f^*G(x) = G(f(x)) | f(x) \ge y\}$ and since $y \le f(x)$ we can use the restriction map $\rho_{f(x),y}^G : G(y) \to G(f(x))$. The universal property of the limit guarantees a map $G(y) \to \lim_{G \to G} G(f(x)) = f_*f^*G(y)$ and thus a natural transformation of functors $id \to f_*f^*$.

Theorem 5.3.2. 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

$$\operatorname{Hom}_{\operatorname{Shv}(Y)}(f_{\dagger}F,G) \cong \operatorname{Hom}_{\operatorname{Shv}(X)}(F,f^{*}G).$$

Dually, the functors for cosheaves satisfy the opposite adjunction (f^*, f_{\dagger})

$$Hom_{\mathsf{CoShv}(X)}(f^*\widehat{G},\widehat{F}) \cong Hom_{\mathsf{CoShv}(Y)}(\widehat{G},f_{\dagger}\widehat{G}).$$

Proof. $f_{\dagger}(f^*G)(y) = \varinjlim\{G(f(x))|f(x) \leq y\}$ so again we can use the restriction maps to define maps to G(y). The universal property of colimits gives a map $f_{\dagger}f^*G(y) \rightarrow G(y)$ and thus a map of functors $f_{\dagger}f^* \rightarrow id$. Similar arguments give a map $id \rightarrow f^*f_{\dagger}$

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 5.3.3. If $F : X \rightarrow Vect$ is a sheaf and $p : X \rightarrow \star$ is the constant map, then

$$\operatorname{Hom}_{\operatorname{Shv}(X)}(p^*k, F) \cong \operatorname{Hom}_{\operatorname{Vect}}(k, p_*F) \cong F(X) = H^0(X; F).$$

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. Hom_{Vect} $(k, W) \cong W$.

6

HOMOLOGY AND COHOMOLOGY

"The de Rham complex may be viewed as a God-given set of differential equations, whose solutions are the closed forms.... A measure of the size of the space of 'interesting' solutions is the definition of the de Rham cohomology."

— Raoul Bott and Loring Tu [BT82, p. 15]

In Section 4.2 and Chapter 5 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 Chapters 8, 9, and 10, where many examples are considered.
6.1 CHAIN COMPLEXES AND HOMOLOGY

Definition 6.1.1. A **Z-graded vector space** V^* is a collection of vector spaces $\{V^i\}_{i \in \mathbb{Z}}$ with one for each integer. A **graded map** is a collection of linear maps $f : V^i \to W^i$. The **category of Z-graded vector spaces**, **grVect**, has graded vector spaces V^* for objects and graded maps for morphisms.

A (co)chain complex is a graded vector space with extra structure.

Definition 6.1.2. A cochain complex consists of a collection of vector spaces called cochain groups $\{V^i\}_{i \in \mathbb{Z}}$ and a collection of linear maps called differentials $d^i : V^i \to V^{i+1}$ that satisfy $d^{i+1} \circ d^i = 0$ for every $i \in \mathbb{Z}$. We denote a cochain complex by $(V^{\bullet}, d^{\bullet})$. Alternatively said, a cochain complex is a graded vector space equipped with a degree one increasing map that when composed twice gives the zero map.

A chain complex is a cochain complex with different notation. The chain groups $\{V_i\}_{i \in \mathbb{Z}}$ and boundary maps $\partial_i : V_i \to V_{i-1}$ are decorated with subscripts; this is the only difference. The maps satisfy $\partial_{i-1} \circ \partial_i = 0$ for every $i \in \mathbb{Z}$. We denote a chain complex by $(V_{\bullet}, \partial_{\bullet})$.

Remark 6.1.3. Since chain complexes and cochain complexes are the same thing, merely dressed up in different notation, we will usually just say "Let (V^{\bullet} , d^{\bullet}) be a chain complex" and let the mathematical notation be precise. As an aside, one can also say that a chain complex is *homologically indexed* if it is written as (V_{\bullet} , ∂_{\bullet}) or *cohomologically indexed* if it is written as (V^{\bullet} , d^{\bullet}).

Remark 6.1.4. Sometimes we drop the subscript or superscript • and write (V, ∂) or (V, d) to refer to a chain complex or cochain complex. Dropping the superscript can lead to overloaded notation. For example, the expression $d^2 = 0$ is a synonym for "d is a differential," i.e. $d^{i+1} \circ d^i = 0$ for ever $i \in \mathbb{Z}$, but it could also mean that the map $V^2 \rightarrow V^3$ is zero. This is one of the perils of cohomological indexing for chain complexes, but the ambiguity is resolved by scoping the context. If we are speaking at the high-level of viewing a chain complex as a different sort of structure, then the former interpretation is intended. If we are talking about the particulars of a given chain complex, then the latter is meant.

Definition 6.1.5. The **category of chain complexes**, $Ch^{\bullet}(Vect)$, has chain complexes for objects, and chain maps $f^{\bullet} : (V, d_V) \to (W, d_W)$ for morphisms, i.e. a collection of maps $f^i : V^i \to W^i$ such that $f^{i+1} \circ d_V^i = d_W^i \circ f^i$.

There is a natural functor ι : grVect \rightarrow Ch(Vect) that treats a \mathbb{Z} -graded vector space as a chain complex with zero differentials, i.e.

$$\{V^i\}_{i\in\mathbb{Z}}\rightsquigarrow (V^\bullet, 0)$$

Taking cohomology of a chain complex defines a functor going the other way.

Definition 6.1.6. Cohomology is a functor H^* : Ch(Vect) \rightarrow grVect, which takes a chain complex (V, d) and places the quotient vector space

$$H^{i}(V, d) := \ker(d^{i}) / \operatorname{im}(d^{i-1})$$

in degree i. Without too much work one can show that a chain map f^{\bullet} induces maps of the associated cohomology spaces $H^{i}(f) : H^{i}(V, d_{v}) \to H^{i}(W, d_{w})$, making H^{*} into a functor.

6.1.1 The Combinatorics of Cell Complexes and Homology

The motivation for chain complexes and homology comes from computing invariants of topological spaces. As already indicated, posets can be regarded as topological spaces, but not every poset has the nice structure that the face-relation poset of a cell complex has. This nice structure is what determines whether certain sequence of vector spaces and maps defines a chain complex.

Definition 6.1.7. We write $\sigma \leq_i \tau$ if the difference in dimension of the cells is i.

Lemma 6.1.8. If $\sigma \leq_2 \tau$, then there are exactly two cells λ_1, λ_2 where $\sigma \leq_1 \lambda_i \leq_1 \tau$.

We want to invent a sign condition that distinguishes these two different sequences of incidence relations.



Definition 6.1.9 (Signed Incidence Relation). A **signed incidence relation** is an assignment to any pair of cells σ , $\tau \in X$ a number [$\sigma : \tau$] $\in \{0, \pm 1\}$ such that

- if $[\sigma : \tau] \neq 0$, then $\sigma \leq_1 \leq_1 \tau$, and
- if γ and τ are any pair of cells, the sum $\sum_{\sigma} [\gamma : \sigma][\sigma : \tau] = 0$.

One way to get a signed incidence relation is to 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 $\sigma \leq \tau$ 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 enables us to define a chain complex associated to a cell complex.

Proposition 6.1.10 (Cellular Cohomology). Let X be a cell complex equipped with a sign relation. Let $C^n(X;k)$ be the vector space spanned by the n-dimensional calls of X. We define a map $\delta : C^n \to C^{n+1}$ on the basis by defining $\delta(\sigma) = \sum_{\tau} [\sigma : \tau] \tau$. Clearly $\delta^2 = 0$.

6.2 COMPUTATIONAL SHEAF COHOMOLOGY AND COSHEAF HOMOLOGY

We now provide formulae for computing cellular sheaf cohomology and cellular cosheaf homology that is completely analogous to cellular cohomology.

6.2.1 Cellular Sheaf Cohomology

Definition 6.2.1 ([Zee62a, She85]). Given a cellular sheaf $F : X \rightarrow$ Vect we define its **compactly supported** k **co-chains** to be the product¹ 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

$$\delta^k_c = \sum_{\sigma \leqslant \tau} [\sigma^k : \tau^{k+1}] \rho_{\tau,\sigma}.$$

The cohomology of this complex

$$0 \longrightarrow \oplus F(\text{vertices}) \xrightarrow{\delta_c^0} \oplus F(\text{edges}) \xrightarrow{\delta_c^1} \oplus F(\text{faces}) \longrightarrow \cdots = C_c^{\bullet}(X;F)$$

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 6.2.2. $(C_c^{\bullet}(X; F), \delta_c^{\bullet})$ is a chain complex.

¹ Here we implicitly assume that X has finitely many cells in a given dimension so products and direct sums agree.

Proof. To see why the chain complex condition $\delta_c^{k+1}\delta_c^k = 0$ is assured, Lemma 6.1.8 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 result.

$$\begin{split} \delta_{c}\delta_{c} &= \sum_{\sigma \leqslant 1\tau} [\sigma:\tau]\rho_{\tau,\sigma}(\delta_{c}) \\ &= \sum_{\sigma \leqslant 1\tau} [\sigma:\tau]\rho_{\tau,\sigma}(\sum_{\gamma \leqslant 1\sigma} [\gamma:\sigma]\rho_{\sigma,\gamma}) \\ &= \sum_{\gamma \leqslant 1\sigma \leqslant 1\tau} [\gamma:\tau]\rho_{\tau,\sigma}\rho_{\sigma,\gamma} \\ &= \sum_{\gamma \leqslant 1\sigma \leqslant 1\tau} [\gamma:\tau]\rho_{\tau,\gamma} \\ &= \sum_{\gamma \leqslant 1\sigma \leqslant 1\tau} ([\gamma:\sigma_{1}][\sigma_{1}:\tau] + [\gamma:\sigma_{2}][\sigma_{2}:\tau])\rho_{\tau,\gamma} \\ &= 0 \end{split}$$

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 6.2.3 (Ordinary Cohomology). Let X be a cell complex and $F : X \rightarrow$ Vect a cellular sheaf. Let $j : X' \rightarrow X$ 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 7, 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 6.2.4 (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 $H_c^0 = H_c^1 = 0$. 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 o, so $H^0(X; k_X) = k$ and is zero in higher degrees.



FIGURE 20: Minimal Cell Structure on an Open Interval

Example 6.2.5 (Open Interval). If we pretended for a moment that the pure stratum Y = (0, 1) is a cell complex² 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 $H_c^1(Y; k_Y) = k$.

To make this example a legitimate example, as in Figure 20, 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 $C_c^1(Y; k_Y)$ we can compute explicitly the compactly supported cohomology.

$$\delta_{c}^{0} = \begin{bmatrix} 1 \\ -1 \end{bmatrix} : k_{x} \to k_{a} \oplus k_{b} \qquad \Rightarrow \qquad H_{c}^{0} = 0 \quad H_{c}^{1} = k$$

6.2.2 Cellular Cosheaf Homology

For cellular cosheaves the exact dual construction works, but the terminology is slightly different.

² Recall that we require a cell complex to have a one point compactification that is a regular cell complex.

Definition 6.2.6 (Borel-Moore Cosheaf Homology). Let X be a cell complex and let \widehat{F} : $X^{op} \rightarrow \text{Vect}$ be a cellular cosheaf. Define the **Borel-Moore homology of** $H^{BM}_{\bullet}(X;\widehat{F})$ to be the homology of the following complex:

$$C^{BM}_{\bullet}(X;\widehat{F}) = \longrightarrow \oplus \widehat{F}(faces) \xrightarrow{[e:\sigma]r_{e,\sigma}} \oplus \widehat{F}(edges) \xrightarrow{[v:e]r_{v,e}} \oplus \widehat{F}(vertices) \longrightarrow 0$$

Definition 6.2.7 (Ordinary Cosheaf Homology). Let X be a cell complex and let $\widehat{F} : X^{op} \rightarrow$ **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

$$C_{\bullet}(X;\widehat{F}) = C_{\bullet}^{BM}(X';j^*\widehat{F}).$$

Applying the definition above gives the ordinary **cosheaf homology** $H_{\bullet}(X; \widehat{F})$ of a cosheaf.

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.

Definition 6.2.8 (Linear Duality). Let \widehat{V} : Shv(X;vect_k)^{op} \rightarrow CoShv(X;vect_k) be the contravariant equivalence from sheaves to cosheaves, both valued in *finite* dimensional vector spaces, defined as follows:

$F(\tau) \longrightarrow I$	$F(au)^*$	$\widehat{V}(F)(\tau)$
ρ _{τ,σ}	$\rho_{\tau,\sigma}^*$	r _{o,τ}
$F(\sigma) \longrightarrow F(\sigma)^*$		$\widehat{V}(F)(\sigma)$

Lemma 6.2.9. Taking linear duals preserves cohomology, i.e. $H_i(X; \hat{V}(F)) \cong H^i(X; F)$ and $H_i^{BM}(X; \hat{V}(F)) \cong H_c^i(X; F)$.

6.3 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 [Gra79] 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.

6.3.1 Persistence Modules and Barcodes

To begin the introduction of representation theory gently, we will describe a convenient visualization technique called a **barcode**, which was first described by Carlsson, Zomorodian, Collins and Guibas [CZCGo4]. The motivation for those authors was to provide a simple shape descriptor for data that could be used by scientists not trained in representation theory, but we will adapt it for understanding sheaves and cosheaves.

To begin, let us recast a chain complex as a special instance of the following structure:

Definition 6.3.1. A **persistence module** consists of a collection of vector spaces $\{V^i\}_{i \in \mathbb{Z}}$, one for each integer, and a collection of linear maps $\rho_V^i : V^i \to V^{i+1}$. If $i \leq j$, then we define $\rho_{i,j}^V := \rho_V^{j-1} \circ \cdots \circ \rho_V^i$ to be the uniquely determined map from $\rho_{i,j}^V : V^i \to V^j$. We denote a persistence module by (V, ρ_V) , but we may suppress the V in ρ_V or even drop the ρ_V all together.

Observe that one can add two persistence modules to create a third persistence module, i.e. if (V^{\bullet}, ρ_V) and (W, ρ_W) are two persistence modules, then one obtains a third persistence module (U, ρ_U) by defining $U^i := V^i \oplus W^i$ and $\rho_U^i := \rho_V^i \oplus \rho_W^i$. We denote the sum by $(V \oplus W, \rho_V \oplus \rho_W)$ or more simply by $V \oplus W$.

There is a fundamental structure theorem for persistence modules, due to Crawley-Boevey [CB12], that explains how any persistence module can be written as a direct sum of simpler persistence modules. We now introduce these simpler persistence modules.

Definition 6.3.2. Recall that an **interval** in (\mathbb{Z}, \leq) is a subset $I \subset \mathbb{Z}$ having the property that if $i, k \in I$ and if there is a $j \in I$ such that $i \leq j \leq k$, then $j \in I$. An **interval module** k_I assigns to each element $i \in I$ the vector space k and assigns the zero vector space to

elements in $\mathbb{Z} \setminus I$. All maps $\rho_{i,j}$ are the zero map, unless $i, j \in I$ and $i \leq j$, in which case $\rho_{i,j}$ is the identity map.

Since interval modules are completely determined by the interval where they assign non-zero vector spaces, we can draw a **bar** to represent an interval module. The following structure theorem shows that any persistence module can be represented by a collection of bars, called a **barcode**.

Theorem 6.3.3 (Decomposition for Pointwise-Finite Persistence Modules). If (V, ρ_V) is a persistence module for which every vector space V^i is finite-dimensional, then the module is isomorphic to a direct sum of interval modules, i.e.

$$V \cong \bigoplus_{I \in D} k_I.$$

Here D is a multi-set of intervals. A multi-set is a set allowing repetitions, i.e. a set equipped with a function μ indicating the multiplicity of each given element.

Remark 6.3.4. We will refer to the length of the bar as $\ell(I) = j - i$

This theorem summarizes a great deal of elementary linear algebra and quiver representation theory. For linear algebra, it has the fundamental theorem of linear algebra as a consequence [Str93], i.e. any map of vector spaces $T : V \rightarrow W$ has a matrix representation that is diagonal with 0 and 1 entries, the number of 1s corresponding to the rank of the matrix, cf. [Art91] Chapter 4, Proposition 2.9. Said differently, there are vector space isomorphisms making the following diagram commute:



Here im(T), ker(T), and cok(T) refer to the image, kernel and cokernel of T respectively. Although the image of T is properly a subspace of W, the first isomorphism theorem identifies it with V modulo the kernel.

Example 6.3.5 (Barcodes for Linear Algebra). Consider any linear map $T : \mathbb{R}^3 \to \mathbb{R}^2$ as a persistence module by extending by zero vector spaces and maps. There are three isomorphism classes of such persistence modules determined by the rank of T. The associated barcodes are depicted in Figure 21.



FIGURE 21: Barcodes associated to $T:\mathbb{R}^3\to\mathbb{R}^2$ for rank(T)=0,1,2

Example 6.3.6 (Barcodes for Filtrations). Barcodes with longer bars appear in filtrations of topological spaces. For example, consider the standard height function on the torus $h : X \to \mathbb{R}$. By choosing a discrete set of points $\{t_0 < t_1, \ldots\}$ to sample the function h at, we get a sequence of spaces $\{X_{\leq t_i} = h^{-1}(-\infty, t_i]\}$, which after taking homology in some degree $i \ge 0$ defines a persistence module. Such an example is depicted in Figure 22.



FIGURE 22: Barcodes for the filtration of a Torus

Now we reach an example that will be useful when we undertake the derived category of chain complexes.

Example 6.3.7 (Barcodes for Chain Complexes). As already remarked, a chain complex of vector spaces is a special example of a persistence module and, consequently, has a barcode decomposition. With a moment's reflection one can see that any chain complex

can be written as the direct sum of two types of modules: the length zero interval modules

$$S_i: \cdots \to 0 \to k \to 0 \to \cdots$$

and the length one interval modules.

 $P_i: \cdots \to 0 \to k \to k \to 0 \to \cdots$

Figure 23 gives a visual depiction of such a barcode decomposition.



FIGURE 23: Barcodes for a Chain Complex

6.3.2 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 Cell(X). Recall, for any category C the **category of representations** is defined to be the category of functors to Vect:

$$\operatorname{Rep}(\mathbf{C}) := \operatorname{Fun}(\mathbf{C}, \operatorname{Vect})$$

This category has the structure of an abelian category, which we explain in this section. In effect, this means we can do everything in Rep(C) that we can do in 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 6.3.8. For **C** a category, **Rep**(**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 η : F → G is a natural transformation occurring inside Vect, there are associated kernel and cokernel representations denoted ker(η) and coker(η) defined by taking kernels and cokernels object wise:

$$\ker(\eta(c')) \longrightarrow F(c') \xrightarrow{\eta(c')} G(c') \longrightarrow \operatorname{coker}(\eta(c'))$$

$$\uparrow^{F(f)} \qquad \uparrow^{G(f)} \qquad \downarrow^{G(f)} \qquad \uparrow^{G(f)} \qquad \downarrow^{G(f)} \qquad \uparrow^{G(f)} \qquad \downarrow^{G(f)} \qquad \downarrow^{G(f)}$$

• There is an **image representation** $im(\eta)$ defined as the object-wise image.

As usual, we say that a sequence of representations $A \rightarrow B \rightarrow C$ 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 6.3.9. A **subrepresentation** E consists of a choice of subspace $E(c) \rightarrow F(c)$ for each object that is invariant under all the linear maps F(f). Restriction of $F(f)|_{E(c)} =: E(f)$ makes E into a representation of its own right. Said more succinctly, $E \rightarrow F$ is a natural transformation of functors that is object wise an inclusion, i.e.

$$0 \rightarrow E \rightarrow F$$

is an exact sequence. Dually, we can say G is a quotient representation by saying $F \rightarrow G \rightarrow 0$ is an exact sequence.

Definition 6.3.10. Suppose $F : C \to Vect$ and $G : C \to Vect$ are two representations of a small category C, then we can define the **direct sum** of these two representations $H = F \oplus G$ by defining on objects $H(c) := F(c) \oplus G(c)$ and on morphisms $H(f) = F(f) \oplus G(f) : H(c) \to H(c')$.

The above definition further clarifies the structure of Rep(C).

Claim 6.3.11. For **C** a category, **Rep**(**C**) is both an exact and an **additive category**. This latter definition requires the following:

- For any two representations F and G the set Hom_{Rep(C)}(F, 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 **Rep**(**C**) is an abelian category.

Remark 6.3.12. 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 6.3.13 (Indecomposable). A representation $F : C \rightarrow Vect$ 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

$$0 \to E \to F \to G \to 0$$

with neither E nor G the zero representation, then $F \ncong E \oplus G$, i.e. the sequence does not split.

Exercise 6.3.14. Verify that the interval modules in Definition 6.3.2 are indecomposable representations. What is the underlying category that these modules represent?

Remark 6.3.15 (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 $E \rightarrow F$ yields the identity, or having a map back from G to F, which post-composed with the surjection is the identity on G.

Definition 6.3.16 (Remak Decomposition). A direct sum decomposition of an object $F \in \text{Rep}(C)$

$$F \cong F_1 \oplus \cdots \oplus F_n$$

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. The structure theorem 6.3.3 provides an example where this is the case. Sir Michael Atiyah considered such a question in the very general setting of abelian categories [Ati₅₆]. 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 6.3.17** (Krull-Schmidt Theorem for Representations [Ati₅₆]). Suppose A is an abelian category, further satisfying

- 1. For every pair of objects $Hom_A(A, B)$ 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

$$\operatorname{Hom}_{A}(B, B') \to \operatorname{Hom}_{A}(A, A') \qquad \eta \mapsto \psi \circ \eta \circ \varphi$$

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 A = Rep(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.

$$\operatorname{\mathsf{Rep}}_{f}(\mathsf{C}) := \operatorname{\mathsf{Fun}}(\mathsf{C}, \operatorname{\mathsf{vect}}) \subset \operatorname{\mathsf{Fun}}(\mathsf{C}, \operatorname{\mathsf{Vect}}) =: \operatorname{\mathsf{Rep}}(\mathsf{C})$$

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.

$$\operatorname{Hom}(F,G) \subseteq \prod_{c \in C} \operatorname{Hom}_{\operatorname{vect}}(F(c),G(c))$$

One severe restriction one can make to insure that Atiyah's first condition holds is to assume that the category C has finitely many objects. This is not strictly necessary, but it does provide us with the following corollary:

Corollary 6.3.18 (Sheaves and Cosheaves on Finite Posets have Remak Decompositions). Suppose (X, \leq) is a finite poset, then Shv(X) and 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 6.3.19 (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

$$\sigma \subset \text{supp}(F)$$
 and $\tau \subset \overline{\sigma}$ or $\sigma \subset \overline{\tau}$ then $\tau \subset \text{supp}(F)$

which by connectedness of X would imply that supp(F) = X; a contradiction to the supposition that neither F nor G was the zero sheaf. (To see why 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.

$$k \rightarrow k \neq (k \rightarrow 0) \oplus (0 \rightarrow k).$$

6.3.3 *Quiver Representations and Gabriel's Theorem*

"These graphs arise in a multitude of classification problems in mathematics, such as classification of simple Lie algebras, singularities, platonic solids, reflection groups, etc. In fact, if we needed to make contact with an alien civilization and show them how sophisticated our civilization is, perhaps showing them Dynkin diagrams would be the best choice!" [EGH⁺11]

There are natural examples of representations of categories where these ideas and their consequences have been studied. One such example is the category associated to a directed graph.

Definition 6.3.20. A **quiver** or **directed graph** is defined by a pair of sets consisting of "edges" E and "vertices" V along with a pair of functions h, t : $E \rightarrow V$, 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, but there is a natural category associated to a directed graph, which we now define.



FIGURE 24: Simply Laced Dynkin Diagrams

Definition 6.3.21. To a quiver we can associate a category called the **free category** or **path category** written 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.

Definition 6.3.22. A **quiver representation** is thus nothing more than a functor F : Free(X) \rightarrow 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.

Every finite dimensional quiver representation can be decomposed into a direct sum of indecomposable representations. However, this list can be very unwieldy. Gabriel's theorem provides an precise description of which quivers admit a finite list of indecomposable representations.

Theorem 6.3.23 (Gabriel's Theorem [DW05]). Let Q denote a quiver. The category of representations Rep(Q) has finitely many indecomposables if and only if the underlying undirected graph of Q is a union of Dynkin graphs of type A_n , D_n , E_6 , E_7 or E_8 . These are depicted in Figure 24.

6.3.4 A Remark on Quivers and Perverse Sheaves

From a quiver representation, one can always construct a cellular sheaf or cosheaf over a one-dimensional base space. This is done by turning every map

$$F(s(e)) \xrightarrow{\rho_{t,s}} F(t(e))$$

into one of the following diagrams:

$$F(s(e)) \xrightarrow{\rho_{t,s}} F(t(e)) = F(e) \xrightarrow{id} F(t(e)) \qquad F(s(e)) \xrightarrow{id} F(t(e)) = F(e) \xrightarrow{\rho_{t,s}} F(t(e))$$

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 $s(e) \leq t(e)$. For example, consider the poset coming from the face relation of the cell complex $Y = [0, 1) \times [0, 1)$:



A quiver representation produces a diagram of vector spaces



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.

A relaxed version of this observation generates a combinatorial model for perverse sheaves, which was invented by Bob MacPherson [GMV96] and explored by Maxim Vybornov [Vyb97, Vyb98].

Definition 6.3.24. A **perversity** $p : \mathbb{Z}_{\geq 0} \to \mathbb{Z}$ is a function from the non-negative integers to the integers such that p(0) = 0 and p takes every interval $\{0, ..., k\}$ bijectively to an interval $\{a, ..., a + k\}$ where $a \in \mathbb{Z}_{\leq 0}$.

Definition 6.3.25 (Cellular Perverse Sheaves). Let X be a cell complex and P_X its associated poset. Let $p : \mathbb{Z}_{\geq 0} \to \mathbb{Z}$ be a perversity. Define a quiver Q_X whose vertices are the elements of P_X and whose edges have the property that if $\tau \to \sigma$ is an edge, then σ is incident to τ and $p(\dim \tau) = p(\dim \sigma) + 1$. A **cellular perverse sheaf** assigns to each vertex of Q_X a vector space $\mathcal{P}(\nu)$ and to each edge from τ to σ a linear map $\kappa_{\sigma,\tau} : \mathcal{P}(\tau) \to \mathcal{P}(\sigma)$. These maps satisfy the chain complex condition for any pair of vertices γ, τ

$$\sum_{\sigma} \kappa_{\gamma,\sigma} \circ \kappa_{\sigma,\tau} = 0$$

where σ ranges over all vertices containing with an edge from τ and to λ .

7

THE DERIVED PERSPECTIVE

"The maker of a sentence launches out into the infinite and builds a road into Chaos and old Night, and is followed by those who hear him with something of wild, creative delight."

— Ralph Waldo Emerson [Eme60, p.59]

The need for a derived perspective can be stated with one picture. In Figure 25 two maps are drawn to the two-sphere S². One is defined on the wedge sum $S^2 \vee S^1$ and maps the S¹ to a point. The other is defined on the closed disk \mathbb{D}^2 and maps the boundary circle to a point. If one is only allowed to look at the homology of the fiber for both of these maps, they will not be able to tell them apart. The derived category is the universal solution to this problem, as well as many others.

In the derived category, one does not consider a sheaf in isolation, but rather one considers complexes of sheaves or, alternatively said, sheaves of complexes. In order to solve the problem presented by Figure 25, one works with the sheaf of cochains on each fiber, along with their differentials. This transition is formally motivated via an analogy with Taylor series in Section 7.1. 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 7.1.1. Injective and projective resolutions are then introduced in Section 7.1.2.

Section 7.2 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 7.2.3 and 7.2.8.

The derived definition of cosheaf homology is given in Section 7.3 and the derived definition of sheaf cohomology can be dualized from there or looked up in Shepard's thesis [She85]. These definitions should be regarded as the true definition of cosheaf



FIGURE 25: MacPherson's Motivating Example for the Derived Category

homology and sheaf cohomology. The compactly supported variant, which we call Borel-Moore cosheaf homology, is defined in Section 7.3.1. The derived functor formalism allows us to resolve the question of invariance under subdivision in Section 7.3.2 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 9.3.

7.1 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:

$$f(\mathbf{x}) \simeq f(\mathbf{a}) + f'(\mathbf{a})(\mathbf{x} - \mathbf{a}) + \frac{f''(\mathbf{a})}{2!}(\mathbf{x} - \mathbf{a})^2 + \cdots$$

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 homology groups:

$$\pi_0(X, x) = \pi_1(X, x) = \pi_2(X, x) = \cdots = H_2(X) = H_1(X) = H_0(X)$$

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 $I^{0}(U)$.¹ 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.

7.1.1 Elementary Injectives and Projectives

In this section we consider the basic building blocks of the derived category, which are injective or projective objects. These objects are characterized by universal mapping properties. For cellular sheaves and cosheaves the injectives and projectives can be described explicitly.

Injectives

Definition 7.1.1. A representation of a small category $I : C \rightarrow Vect$ is **injective** if, for any natural transformation $\eta : A \rightarrow I$ and any injection $\iota : A \hookrightarrow B$, there is an extension $\tilde{\eta} : B \rightarrow I$ such that $\eta = \tilde{\eta} \circ \iota$. Said using diagrams, they are characterized by the usual universal property:



Exercise 7.1.2. Use the universal property of an injective representation to prove the following statements:

¹ There is no guarantee for exactness at higher terms.

Every short exact sequence

$$0 \rightarrow I \rightarrow A \rightarrow Q \rightarrow 0$$
,

where I is injective, splits, i.e. if $\iota : I \hookrightarrow A$ is an inclusion, then $A \cong I \oplus cok(\iota)$.

• $\prod A_i$ is injective if and only if each A_i is injective.

We will make use of these properties in Lemma 7.1.6.

For our first example of an injective representation, we consider an injective cell sheaf. These sheaves are supported on the closures of cells.

Definition 7.1.3. An elementary injective cell sheaf on X concentrated on $\sigma \in X$ with value $W \in$ Vect is given by

$$\left[\sigma\right]^{W}(au) = egin{cases} W & ext{if } au \leqslant \sigma, \ 0 & ext{other wise.} \end{cases}$$

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 7.1.4. 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 $W \in$ **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.

Lemma 7.1.5. 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

$$\operatorname{Hom}_{\operatorname{Shv}(X)}(A, (\mathfrak{i}_x)_*W) \cong \operatorname{Hom}_{\operatorname{Vect}}(A(x), W)$$

 $\operatorname{Hom}_{\operatorname{CoShv}(X)}(\hat{A},(\mathfrak{i}_{x})_{\dagger}\hat{W}) \cong \operatorname{Hom}_{\operatorname{Vect}}(\hat{A}(x),\hat{W})$

and the fact that in the category of vector spaces every object is injective.

There is one last lemma that tells us that the considering sheaves of the form $[x]^W$ suffices for understanding all injective sheaves over a cell complex. The proof is presented in [She85, Thm. 1.3.2, p.19-20], but we will give the direction needed for our derived equivalence proof given in Theorem 12.2.1.

Lemma 7.1.6 (All Injectives are Sums). Let X be the face-relation poset of a cell complex. A sheaf I is injective if and only if it is isomorphic to a one of the form $\bigoplus_{\sigma} [\sigma]^{V_{\sigma}}$.

Proof. One can easily check that the direct sum of elementary injective sheaves is injective, so this proves the 'if' direction.

To prove that every injective is isomorphic to a direct sum of injectives — the 'only if' direction — requires a little work. Assume for induction that every injective sheaf I that is non-zero on at most $k \leq n-1$ cells is isomorphic to $\bigoplus_{\sigma}[\sigma]^{V_{\sigma}}$. Now consider an injective sheaf that is non-zero on exactly n cells. Let σ be a cell of maximal dimension where $I(\sigma) =: V \neq 0$. Since I is zero on all higher cells incidence to σ , there is a non-zero map η from the skyscraper sheaf S_{σ}^{V} to I with $\eta(\sigma) = id_{V}$. There is also a non-zero map $\iota: S_{\sigma}^{V} \to [\sigma]^{V}$. This gives us a diagram



and the implicated existence of a map $\tilde{\eta} : [\sigma]^V \to I$. If $\tau \leq \sigma$, then by the fact that $\tilde{\eta}$ is a sheaf map,

$$\mathrm{id}_V = \tilde{\eta}(\sigma) \circ \rho_{\sigma,\tau}^{\lfloor \sigma \rfloor} = \rho_{\sigma,\tau}^{\mathrm{I}} \circ \tilde{\eta}(\tau)$$

the map $\tilde{\eta}$ is injective. By the second property of Exercise 7.1.2, we can deduce that $I \cong [\sigma]^V \oplus \operatorname{cok}(\tilde{\eta})$. Since $\operatorname{cok}(\tilde{\eta})$ is zero wherever I is and also zero on σ , it is non-zero on at most n - 1 cells and the induction hypothesis applies. The zero sheaf is clearly equal to a direct sum of elementary injectives with the zero vector space, which checks the base case, completing the induction.

Projectives

There is a dual universal object that is called projective.

Definition 7.1.7. 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:



As before, we have some dual consequences:

Every short exact sequence

$$0 \rightarrow A \rightarrow B \rightarrow P \rightarrow 0$$

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 7.1.8. 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 $W \in$ Vect to be $\{x\}^W = (i_x)_{\dagger}W$ and the corresponding elementary projective cosheaf to be $[\hat{x}]^W := (i_x)_{\ast} \hat{W}$.

We leave it to the reader to check that these objects are actually projective.

Mapping Identities

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.

$$\begin{split} & \operatorname{Hom}_{\mathsf{Shv}}([\tau]^{U},[\sigma]^{W}) = \begin{cases} \operatorname{Hom}_{\mathsf{Vect}}(U,W) & \text{if } \sigma \leqslant \tau, \\ 0 & \text{o.w.} \end{cases} \\ & \operatorname{Hom}_{\mathsf{Shv}}(\{\tau\}^{U},\{\sigma\}^{W}) = \begin{cases} \operatorname{Hom}_{\mathsf{Vect}}(U,W) & \text{if } \sigma \leqslant \tau, \\ 0 & \text{o.w.} \end{cases} \\ & \operatorname{Hom}_{\mathsf{CoShv}}([\hat{\sigma}]^{W},[\hat{\tau}]^{U}) = \begin{cases} \operatorname{Hom}_{\mathsf{Vect}}(W,U) & \text{if } \sigma \leqslant \tau, \\ 0 & \text{o.w.} \end{cases} \end{split}$$

$$\operatorname{Hom}_{\mathsf{CoShv}}(\{\hat{\sigma}\}^{W},\{\hat{\tau}\}^{U}) = \begin{cases} \operatorname{Hom}_{\mathsf{Vect}}(W,U) & \text{if } \sigma \leqslant \tau, \\ 0 & \text{o.w.} \end{cases}$$

7.1.2 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 7.1.9. Every sheaf $F : X \rightarrow$ **Vect** on a poset (X, \leq) 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

$$0 \to F \to I := \prod_{x} [x]^{F(x)} = \prod_{x} (\mathfrak{i}_{x})_{*} F(x).$$

The map to I is defined easily using the standard adjunctions

$$\iota \in \operatorname{Hom}(F, \prod_{x} (\mathfrak{i}_{x})_{*}F(x)) \cong \prod_{x} \operatorname{Hom}(F, \prod_{x} (\mathfrak{i}_{x})_{*}F(x)) \cong \prod_{x} \operatorname{Hom}(F(x), F(x)) \ni \prod_{x} \operatorname{id}_{F(x)}.$$

We encourage the reader to describe this map is explicitly, by seeing how a single $id_{F(x)}$ traces through this adjunction, which we'll call $\iota_x \in Hom(F, (i_x)_*F(x))$.

Similarly, for a cosheaf $\hat{F} : X^{op} \to Vect$ on an Alexandrov space we could have built a projective surjection by taking

$$\hat{P}_{0} := \bigoplus_{x} (\mathfrak{i}_{x})_{*} \widehat{F}(x) \to \widehat{F} \to 0$$

where the map $\pi_0 : P_0 \to \widehat{F}$ is gotten through the corresponding adjunction for cosheaves and using the contravariance of Hom in the first slot

$$\operatorname{Hom}(\bigoplus_{x}(\mathfrak{i}_{x})_{*}\widehat{F}(x),\widehat{F})\cong\prod\operatorname{Hom}((\mathfrak{i}_{x})_{*}\widehat{F}(x),\widehat{F})\cong\prod\operatorname{Hom}(\widehat{F}(x),\widehat{F}(x)).$$

Corollary 7.1.10. Every sheaf $F : X \to Vect$ has an injective resolution. Dually, every cosheaf $\widehat{F} : X^{op} \to Vect$ 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:



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{F} :

$$\cdots \hat{P}_2 \to \hat{P}_1 \to \hat{P}_0 \to \widehat{F} \to 0.$$

These exact sequences can be used to replace F or \hat{F} in a suitable sense, defined by the derived category. Before moving onto that discussion, we note one interesting point.

Proposition 7.1.11. The length of injective resolution of any sheaf $F \in Shv(X)$ 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 $cok(id : F(x') \to F(x')) = 0$. So for any maximally ordered chain in X, I^1 is zero on the top-most element. Arguing inductively finishes the proof.

7.2 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 $\phi : \widehat{F} \to \widehat{G}$ is a map of cosheaves, then it can be checked from the universal properties of projective objects, that this induces a map of complexes



where all the squares in sight commute. For a hint on how to see this, consider the composite map $\hat{P}_0 \rightarrow \hat{F} \rightarrow \hat{G}$ 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 $\hat{P}_0 \rightarrow \hat{Q}_0$. To get the next, all important step, one must recognize that having maps from $\hat{P}_0 \rightarrow \hat{Q}_0$ and $\hat{F} \rightarrow \hat{G}$ induces maps between the kernels of the map $\hat{P}_0 \rightarrow \hat{F}$ and $\hat{Q}_0 \rightarrow \hat{G}$. 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 $\phi = id$.



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 7.2.1 (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



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

$$\hat{\mathsf{P}}_{\bullet}$$
: $[\hat{\mathfrak{a}}] \to \hat{k}_X.$

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:

$$\hat{\mathrm{Q}}_ullet: \qquad [\hat{\mathrm{x}}] \oplus [\hat{\mathrm{y}}] o [\hat{\mathrm{a}}] \oplus [\hat{\mathrm{x}}] \oplus [\hat{\mathrm{y}}] o \hat{\mathrm{k}}_X$$

Definition 7.2.2 (Chain Homotopy). Suppose (A^{\bullet}, d_A) and (B^{\bullet}, d_B) are two (cohomological) chain complexes and ϕ^{\bullet} and ψ^{\bullet} are two chain maps, then **a chain homotopy** h^{\bullet} is a chain map $h^i : A^i \to B^{i-1}$

such that

$$\phi^i - \psi^i = h^{i+1} d^i_A + d^{i-1}_B h^i.$$

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

 $\varphi\circ\psi\sim id\qquad and\qquad \psi\circ\varphi\sim id$

then one says A[•] and B[•] are chain homotopy equivalent.

Example 7.2.3 (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



is clearly the identity on \hat{P}_{\bullet} , but the composite



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 7.2.4. Let A be an abelian category, such as the category of sheaves or cosheaves. The **category of chain complexes** in A, written $C^{b}(A)$ has objects that are chain complexes and morphisms that are chain maps.

The **homotopy category of complexes** $K^{b}(A)$ of an abelian category A has the same objects as $C^{b}(A)$, but where we have identified chain homotopic maps.

Definition 7.2.5. For A = Shv(X) we define the **bounded derived category of sheaves** $D^{b}(Shv(X))$ to be $K^{b}(Inj - Shv(X))$ the homotopy category that uses only complexes of injective sheaves.

Similarly, for A = CoShv(X), we define the **bounded derived category of cosheaves** $D^{b}(CoShv(X))$ by $K^{b}(Proj - CoShv(X))$ 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 7.2.6. Suppose we are given a complex of cellular sheaves

 $(F^{\bullet}, d^{\bullet}): \cdots \rightarrow F^{i-1} \rightarrow F^{i} \rightarrow F^{i+1} \rightarrow \cdots,$

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

$$\mathcal{H}^{i}(\mathsf{F}^{\bullet}): \qquad \sigma \rightsquigarrow \mathcal{H}^{i}(\mathsf{F}^{\bullet}(\sigma))$$

which is a cellular sheaf. The restriction maps being defined as the induced map on cohomology for the chain map $F^{\bullet}(\sigma) \rightarrow F^{\bullet}(\tau)$ for $\sigma \leq \tau$.

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)

$$\mathcal{H}^*: C^b(\operatorname{Shv}(X)) \to \operatorname{Shv}(X; \operatorname{grVect}) \qquad F^\bullet \rightsquigarrow \bigoplus_i \mathcal{H}^i(F^\bullet).$$

There are completely dual notions of **homology cosheaves**, where we generally use homological indexing and notation $(\hat{F}_{\bullet}, \partial_{\bullet})$.

Definition 7.2.7 (Quasi-Isomorphisms). A map of complexes of sheaves (or cosheaves) $\alpha^{\bullet} : F^{\bullet} \to G^{\bullet}$ such that the induced map

$$\mathcal{H}(\alpha^{\bullet}): \mathcal{H}^{i}(\mathsf{F}^{\bullet}) \to \mathcal{H}^{i}(\mathsf{G}^{\bullet})$$

is an isomorphisms for every i, is called a **quasi-isomorphism**.

The term "quasi-isomorphism" reflects the fact that if α^{\bullet} : $F^{\bullet} \rightarrow G^{\bullet}$ is a quasiisomorphism, then there does not always exist an inverse map β^{\bullet} that gives the identity, or even chain homotopic to the identity, any map back may simply not exist.

Example 7.2.8. 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

$$0 \longrightarrow S_{a} \longrightarrow 0$$

$$\downarrow \qquad \qquad \downarrow \qquad \qquad \downarrow$$

$$0 \longrightarrow [a] \longrightarrow [x] \oplus [y]$$

which is a quasi-isomorphism. However, there does not exist a map of sheaves $[a] \rightarrow S_a$.

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 **B** that is closed under certain

operations, then we can consider the following universal problem: suppose $L : \mathbf{B} \to \mathbf{C}$ 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 $\mathbf{B}[Q^{-1}]$.



An alternative approach to the derived category of an abelian category A is to define

 $D(A) := K(A)[Q^{-1}]$ $Q = {quasi-isomorphisms}$

where we have removed the boundedness hypothesis.

One then proves the following claim to re-obtain the definition we provided here

Theorem 7.2.9 ([Aluo9] 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.

7.3 THE DERIVED DEFINITION OF COSHEAF HOMOLOGY AND SHEAF COHOMOL-OGY

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 [She85] pp. 28-29. Let's more or less repeat the proof for cellular cosheaves since it is nowhere in the literature.

Definition 7.3.1. Given a cosheaf \widehat{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:

$$Lf_*\widehat{F} := f_*P_{\bullet}.$$

We define the ith derived functor by

$$L_{i}f_{*}\widehat{F} := \mathcal{H}_{i}(f_{*}P_{\bullet}).$$

In the special case where $f = p : X \rightarrow \star$ we write

$$H_i(X; \widehat{F}) := L_i p_* \widehat{F}$$

for the ith cosheaf homology group of \widehat{F} .

We now aim to prove the following theorem.

Theorem 7.3.2. 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 $\hat{P}_{\bullet} \rightarrow \hat{F}$ and then take cellular chains of each cosheaf to obtain the following double complex:



Now we make use of the following two observations, which dualize [She85, Thm. 1.3.10, 1.4.1].

Lemma 7.3.3. For 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 $|\bar{\sigma}|$, has the homeomorphism type of a closed k-simplex. $C_{\bullet}(X; [\hat{\sigma}])$ is the chain complex that computes the cellular homology of $Y = |\{\tau \leq \sigma | \bar{\tau} \text{ is compact} \}|$, 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 $|\bar{\sigma}|$ 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 $|\bar{\sigma}|$ are simply connected. This proves the claim.

Lemma 7.3.4. For any cellular cosheaf \hat{F} on a cell complex X we have that

$$\operatorname{colim} \hat{\mathsf{F}} \cong \operatorname{cok}(\mathsf{C}_1(X; \hat{\mathsf{F}}) \to \mathsf{C}_0(X; \hat{\mathsf{F}})).$$

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 colim \hat{F} is a choice of maps $\psi_{\sigma} : \hat{F}(\sigma) \to \text{colim}\hat{F}$. Let $\Psi = \oplus \psi_{\sigma} : \oplus \hat{F}(\sigma) \to \text{colim}\hat{F}$, now consider the factorization of this map through the image:



Now we can use the im Ψ to define a new co-cone over the diagram \hat{F} simply by precomposing 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 $u \circ j = \operatorname{id}$ since there is only one map $\operatorname{colim} \hat{F} \to \operatorname{colim} \hat{F}$.

Now observe that $C_0(X; \hat{F}) = \bigoplus \hat{F}(v_i)$ 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 $\bigoplus \hat{F}(v_i)$. 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



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 $Tot_{\bullet}(C_i(X; \hat{P}_j))$ induces quasi-isomorphisms between $colim\hat{P}_{\bullet}$ and $C_{\bullet}(X; \hat{F})$. We have thus established the theorem.

7.3.1 Borel-Moore Cosheaf Homology

Definition 7.3.5. Suppose \widehat{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 $\widehat{F}(\infty) = 0$. Alternatively said we look at the inclusion $j : X \to X \cup \{\infty\}$ and define

$$\Gamma^{BM}(X;\widehat{F}) := \mathfrak{p}_*\mathfrak{j}_{\dagger}\widehat{F}.$$

Another possible definition is to dualize a cellular cosheaf of finite-dimensional vector spaces to a cellular sheaf by post-composing $\hat{F} : X \to \text{Vect}$ with $\text{Hom}_{\text{vect}}(-,k)$, apply $p_!$ and then dualize back.

Remark 7.3.6 (Functoriality). The definitions that involve the one-point compactification are deficient in the following way. A map of cell complexes $f : X \rightarrow Y$ does not necessarily extend to a map between the one-point compactifications. It is for this reason that for functoriality, the definition using p_1 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 7.3.7. For any cellular cosheaf \hat{F} on a cell complex X we have that $\Gamma^{BM}(X; \hat{F}) \cong cok(C_1^{BM}(X; \hat{F}) \to C_0^{BM}(X; \hat{F})).$

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



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.

7.3.2 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 7.3.8. Suppose F is a sheaf on X and $s : X' \to X$ is a subdivision of X, then we define the subdivided sheaf $F' := s^*F$.

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)



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 [She85, Thm. 1.5.2]:

Theorem 7.3.9. 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

$$s_*s^*F(y) = \varprojlim\{s^*F(x)|s(x) \ge y\}$$

=
$$\varprojlim\{F(s(x))|s(x) \ge y\}$$

(surjectivity) =
$$\varprojlim\{F(x)|x \ge y\}$$

(sheaf - axiom) =
$$F(U_y)$$

=
$$F(y)$$

So we have that for the subdivision map $s_*s^*F \cong F$ and as a consequence

$$(\mathbf{p}_{\mathbf{X}'})\mathbf{F}'\cong\mathbf{p}_{\mathbf{X}}\mathbf{F}.$$

Now we can just take the associated right derived functors to obtain the result. \Box

7.4 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. **Ab**, **Vect**, etc., the existence of enough projective sheaves is driven by the underlying topology of the space.

Proposition 7.4.1. 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.²

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



whose value on the open set U is



so in particular the surjection must factor through zero — a contradiction.

² The author would like to acknowledge the contributions of Valery Alexeev, David Treumann, and Jon Woolf on mathoverflow in regards to this question.
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.

$$\mathsf{P}^{\mathsf{0}} := \oplus_{\sigma \in \mathsf{X}} \{\sigma\}^{\mathsf{F}(\sigma)} \to \mathsf{F} \to \mathsf{0}$$

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

$$\cdots P^{-3} \rightarrow P^{-2} \rightarrow P^{-1} \rightarrow P^{0} \rightarrow F \rightarrow 0$$

This motivates the following definitions:

Definition 7.4.2. Given a cellular sheaf, we can construct its projective resolution $P_{\bullet} \rightarrow F$, 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 $\bigoplus_{\dim \tau = n} F(\tau)$ and dually for limits of cosheaves.
- 3. This defect, which is measured by the difference of Hⁿ(X; F) and colimF, 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 7.4.3. 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:

So the same abstract diagram of vector spaces $F : X \to \text{Vect}$ defines a diagram over $\widehat{F} : \widetilde{X}^{op} \to \text{Vect}$, i.e. a cosheaf on the dual cell structure. Since they are the same diagrams everything about them is the same.

7.4.1 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 7.4.4. Suppose F is a sheaf on X and X' is a subdivision of X, then

$$(\mathbf{p}_{\mathbf{X}'})_{\dagger}\mathbf{F}' \cong (\mathbf{p}_{\mathbf{X}})_{\dagger}\mathbf{F}$$

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

$$\begin{split} s_{\dagger}s^*\mathsf{F}(y) &= \ colim\{s^*\mathsf{F}(x)|s(x)\leqslant y\}\\ &= \ colim\{\mathsf{F}(s(x))|s(x)\leqslant y\}\\ (surjectivity) &= \ colim\{\mathsf{F}(x)|x\leqslant y\}\\ (check-directly) &= \ \mathsf{F}(y) \end{split}$$

and thus

$$(\mathbf{p}_{X'})_{\dagger} \mathbf{s}^* \mathbf{F} = (\mathbf{p}_X)_{\dagger} \mathbf{s}_{\dagger} \mathbf{s}^* \mathbf{F} \cong (\mathbf{p}_X)_{\dagger} \mathbf{F}.$$

Taking the left derived functors gives the higher result.

Part III

APPLICATIONS TO SCIENCE AND ENGINEERING

This part constitutes a first application of cellular sheaves and cosheaves to problems in science and engineering.

Chapter 8 begins with a short introduction to persistent homology, which we reformulate using sheaves and cosheaves. The advantage of this reformulation is the ability to distribute homology computations and aggregate efficiently, as noted in Section 8.2.2, which is part of joint work [CGN13] where discrete Morse theory is adapted to compute cellular sheaf cohomology. A theorem that uses the apparatus of spectral sequences to connect level set and sub-level set persistence is proved in Section 8.2.3. A motivating example for multi-dimensional persistence and an introduction of "generalized barcodes" concludes the chapter.

Chapter 9 reviews an application of sheaves to network coding introduced first in [GH11]. However, here the language of cellular sheaves is employed and the barcode method is used to visualize the flow of data. A duality theorem for network coding sheaves is proved in two different ways.

Chapter 10 casts various sensor network problems in the language of sheaves. Any attempt to use level set persistence to study the intruder problem is proven to be a "no-go" using the machinery of cosheaves. The main contribution of the chapter is the introduction of a linearized model for multi-modal sensing in Section 10.4. It was this model that first motivated the author to take up the theory of indecomposables as a way of interpreting sheaf cohomology.

TOPOLOGICAL DATA ANALYSIS

"We are enabled to divide into forms, following the objective articulation; we are not to hack off parts like a clumsy butcher."

- Plato's Phaedrus 265e, translated by R. Hackforth

Suppose two patients enter an office with a recent diagnosis of cancer. If they are very lucky, they might have their respective tumors biopsied and sent off for genetic analysis. The genetic analysis consists of measuring the gene expression levels of over a thousand genes in each of the respective tumors. These two sets of expression levels are then compared with a few hundred other tumor samples, where varying therapies were used to varying degrees of success. How should the doctor determine which course of action to take? If we are to carve the universe at its joints, on which side do these patients lie?

A similar, but apparently less dramatic, situation occurs in topology. Given two topological spaces X and Y, how do we discern whether X and Y are essentially the same or different? The entire apparatus of algebraic topology was developed to address this problem. It turns out that these two situations are not just formally similar. Topological data analysis stems from the observation that data of the above form can give rise to topological spaces, which can in turn be discriminated using classical constructions such as homology [LSL⁺13].

8.1 POINT CLOUDS AND PERSISTENT HOMOLOGY

To illustrate why the applicability of homological methods is not so far-fetched, consider the following toy problem, which serves as the standard entrée into **persistent homology**. Consider a finite set of points $\{x_i\} \subset \mathbb{R}^n$. How do we describe the perceived shape of such a set of points, such as the set depicted in Figure 26? The human brain is a patternmaking machine that connects the dots and returns the knee-jerk response that the points in Figure 26 appears to form a circle. However, what do we mean by this and how can we automate this process so as to remove human subjectivity?

Mathematics abounds with rigorous formulations of what does or does not look like a circle. However, only topology provides a definition that is robust with respect to



FIGURE 26: Point Cloud Data

perturbation and noise. One universal definition of a circle is given by the Eilenberg-Mac Lane space $K(\mathbb{Z}, 1)$, which is homotopy equivalent to S¹. However, homology provides a shape descriptor based on linear algebra, which can be efficiently computed. But there is still the problem that the homology of the set of points in Figure 26 is not isomorphic to the homology of the circle. To get around this, we consider the union of Euclidean balls of some radius $r \cup B(x_i, r) =: X_r$, which mirrors the "connecting the dots" procedure that the brain applies. 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 $r_0 \leq r_1 \leq r_2 \leq \cdots$ and so on. Applying the ith homology functor $H_i(-;k)$ turns this diagram of spaces into a diagram of vector spaces, which defines a persistence module, cf. Definition 6.3.1.

$$H_{i}(X_{r_{0}};k) \rightarrow H_{i}(X_{r_{1}};k) \rightarrow H_{i}(X_{r_{2}};k) \rightarrow H_{i}(X_{r_{3}};k) \rightarrow \cdots$$

By applying the structure theorem 6.3.3, we can determine the barcodes of the collection of points. Long bars are considered to be robust topological signals in the data set. For Figure 26, there would be one long bar in the persistence module corresponding to H_0 , indicating that after a certain radius the the space X_r is connected, and another long bar in the module corresponding to H_1 , indicating the apparent circle in the data set.

To summarize, we have the following prototypical pipeline of topological data analysis.

Definition 8.1.1 (Point Cloud Persistence). The **point cloud persistence pipeline** consists of the following ingredients and operations:

- 1. Let X denote a point cloud, i.e. the union of a finite set of points $\{x_i\} \subset \mathbb{R}^n$.
- 2. The union of balls $X_r := \bigcup_{x_i \in X} B(x, r)$ (or the Vietoris-Rips complex [Ghro8]) defines a functor from the real line, viewed as poset, to the category of topological spaces (simplicial complexes) and maps, i.e.

$$\mathsf{G}:(\mathbb{R},\leqslant)\to\mathsf{Top}\qquad r\rightsquigarrow X_r\qquad r\leqslant r'\rightsquigarrow X_r\hookrightarrow X_{r'}$$

3. Postcomposing this functor with homology H_{*}, defines a graded representation of the real line, which is equivalently a graded sheaf on the Alexandrov topology or a graded persistence module:

$$H_*(X_{r_0};k) \to H_*(X_{r_1};k) \to H_*(X_{r_2};k) \to H_*(X_{r_3};k) \to \cdots$$

4. Applying Theorem 6.3.3 produces a Remak decomposition of this representation into a multiset of interval modules, which is visualized as a barcode by the end user.

The first and second steps in this pipeline offer the chance for endless modification and application. Instead of considering a collection of points, one can start with a space X and a function $f: X \to \mathbb{R}$ and consider the family of sub-level sets $X_r := f^{-1}(-\infty, r]$. As long as the function and space are sufficiently nice, we can use Theorem 6.3.3 to produce a barcode.

Exercise 8.1.2. Determine the barcodes associated to the function $f(x) = x^3 - x$.

Exercise 8.1.3. Recast the basic theorems of Morse theory in terms of barcodes. See Figure 22 for inspiration.

8.1.1 Level Set and Zigzag Persistence

Despite their successes, persistence modules are not the end-all, be-all of topological data analysis. In [CdS10] Gunnar Carlsson and Vin de Silva gave three examples where diagrams of vector spaces and maps of the form

$$V_1 \leftrightarrow V_2 \leftrightarrow V_3 \leftrightarrow \cdots \leftrightarrow V_{n-2} \leftrightarrow V_{n-1} \leftrightarrow V_n$$

are of interest. One example comes from estimating the probability density function from which a point-cloud is drawn. One can try to smooth the data by defining a function $\rho_r(x)$ that counts the number of points within a radius r of the point x. If one then tries to take the 25% densest points measured according to a sequence of radii $r_1 < r_2 < \cdots < r_n$, the the only way of comparing features comes from a **zigzag** diagram of the form

$$X^p_{r_1} \to X^p_{r_1} \cup X^p_{r_2} \leftarrow X^p_{r_2} \to \dots \leftarrow X^p_{r_n}$$

where $X_{r_k}^p$ indicates the p% densest points measured according to the function $\rho_{r_k}(x)$. Similarly, if one has a function $f: X \to \mathbb{R}$, but not the computational power to investigate the entire sub-level set $f^{-1}(-\infty, r]$, one could choose a mesh $t_0 < t_1 < t_2 < \cdots < t_n$ and consider the zigzag of pre-images given by

$$f^{-1}(t_0) \to f^{-1}[t_0, t_1] \leftarrow f^{-1}(t_1) \to \cdots \leftarrow f^{-1}(t_n).$$

Applying homology in some degree i gives the traditional definition of **level set persistent** homology. This complicates the usual TDA pipeline because, a priori, the structure theorem 6.3.3 fails to apply. Fortunately, Gabriel's Theorem 6.3.23 tells us that the direction of the arrows doesn't matter for the representations of A_n type quivers, so the decomposition into interval modules still applies; we still have barcodes [CdS10].

The fully general definition of level set persistence usually given adheres to this perspective that the assignment of homology to closed intervals is fundamental.

Definition 8.1.4. The **interval category** of \mathbb{R} , written $Int(\mathbb{R})$, is the category whose objects are closed intervals $[x, y] \subset \mathbb{R}$ and whose morphisms are inclusions $[x, y] \hookrightarrow [x', y']$.

Definition 8.1.5 (Level Set Persistent Homology). Suppose $f : X \to \mathbb{R}$ is a function, not necessarily continuous. The **ith level set persistence of** $f L_i$ is the representation of the interval category given by

$$L_i: Int(\mathbb{R}) \rightarrow Vect$$
 $[x, y] \rightsquigarrow H_i(f^{-1}([x, y]); k).$



FIGURE 27: Height Function on the Circle

Critique of the Definition of Level Set Persistence

The definition of level set persistence suffers in one crucial respect. The definition is *non-local* and consequently requires the storage of an infinite amount of data. Consider the map f depicted in Figure 27. Level set persistence for H_1 will assign the zero vector space to every interval of diameter less than y - x. Only after inspecting intervals large enough will the topological feature of the circle appear.

A sheaf-theoretic approach to level-set persistence offers the advantage of being local. However, the analogous sheaf version of the level set H_1 just examined is the zero sheaf. The trade-off appears to be too great. However, the apparent disadvantage is remedied via the use of sheaf-cohomology, which preserves all the information of the domain space while simultaneously being local.

8.2 APPROACHING PERSISTENCE WITH SHEAVES AND COSHEAVES

8.2.1 Cellular Maps and Absolute Homology Cosheaves

The contravariant nature of sheaves should remind us that they are most naturally associated with cohomology of a space. Cosheaves, being covariant with respect to the



FIGURE 28: Homology Cosheaf for a circle, with barcodes

inclusion of opens, are naturally associated with homology of a space. Since we are working over a field, we can pass back and forth between these perspectives. However, to coincide with the traditional use of homology in persistence, we first introduce our version of persistence using cosheaves.

Definition 8.2.1. Suppose X and Y are cell complexes and $f : Y \to X$ is a *proper* cellular map, see Definition 5.1.12 for a reminder, then for each natural number $i \ge 0$ we have the following ith **absolute homology cosheaf** \widehat{F}_i , which assigns to a cell σ in X the ith homology of the pre-image $f^{-1}(\operatorname{star}(\sigma))$, i.e.

$$\widehat{F}_{i}(\sigma) := H_{i}(f^{-1}(star(\sigma)); k).$$

This is clearly a cellular cosheaf since if $\sigma \leq \tau$, then $star(\tau) \subseteq star(\sigma)$ and thus we have a map

$$r_{\sigma,\tau}: H_i(f^{-1}(\operatorname{star}(\tau)); k) \to H_i(f^{-1}(\operatorname{star}(\sigma)); k).$$

By Corollary 6.3.18 we know that every absolute homology cosheaf can be decomposed into indecomposable sheaves. For cellular maps $f : Y \to X$ where X is a compact subset of \mathbb{R} , the absolute homology cosheaves have the following form:

$$H_{i}(f^{-1}(star(x_{0}));k) \leftarrow H_{i}(f^{-1}((x_{0},x_{1}));k) \rightarrow H_{i}(f^{-1}(star(x_{1}));k) \leftarrow \cdots$$

Hence, by Gabriel's Theorem 6.3.23 absolute homology cosheaves over $X \subset \mathbb{R}$ can be assigned barcodes. However, by inspecting the support of these indecomposable cosheaves we observe that there are four types of bars that make up any barcode:

Example 8.2.2. Let $h : S^1 \to \mathbb{R}$ be the standard height function on the circle, drawn in Figure 27. The only absolute homology cosheaf of interest is \widehat{F}_0 , since the fibers have no higher homology. The associated pre-images, values of the homology cosheaf and barcode are drawn in Figure 28.

Remark 8.2.3 (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. $(x_i, x_{i+1}) \rightarrow]x_i, x_{i+1}[$ and $[x_i, x_{i+1}] \rightarrow [x_i, x_{i+1}[$ and so on.

By viewing these barcodes as cosheaves, which have a homology theory, we can compute barcode homology.

Lemma 8.2.4. Suppose X is a compact subset of \mathbb{R} , equipped with some cell structure. The cosheaf homology of the four types of indecomposable cosheaves coincides with the Borel-Moore homology of the underlying barcode, i.e.

$$H_0^{BM}([-]) = k$$
 $H_1^{BM}(]-[) = k$ $H_i^{BM}([-]) = H_i^{BM}([-]) = 0$

with all other Borel-Moore homology groups being zero. Moreover, since cosheaf homology commutes with finite direct sums, cellular cosheaf homology of \hat{F} on X can be computed using the barcode $B_{\hat{F}}$ associated to the Remak decomposition of \hat{F} .

$$H_{i}(X; \widehat{F}) \cong \oplus H_{i}^{BM}(B_{\widehat{F}}) \qquad i = 0, 1$$

In short,

$$H_0(X; \widehat{F})$$
 counts closed bars and $H_1(X; \widehat{F})$ counts open bars.

Proof. The proof of the claim uses simple computations, illustrated in the examples, and an invariance under subdivision argument, which is proved in Theorem 7.3.9. The true sticking point is why ordinary cosheaf homology becomes Borel-Moore (cosheaf) homology, which is developed in section 7.3.1.

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

$$j: B \hookrightarrow [0, 1]$$
 and $\hat{k}_B \mapsto j_{\dagger} \hat{k}_B$.

The process of taking cosheaf homology is to then push forward this cosheaf to a point. However, using Lemma 7.3.7 we see the following is true:

$$p_!\hat{k}_B \cong p_*j_!\hat{k}_B \qquad \mathsf{H}^{\mathsf{BM}}_{\mathsf{i}}(\mathsf{B};\hat{k}_B) \cong \mathsf{H}_{\mathsf{i}}([0,1];j_!\hat{k}_B)$$

When it is clear that we are working on [0, 1] we may write \hat{k}_B instead of $j_{\dagger}\hat{k}_B$.

One of the advantages of absolute homology cosheaves is that over the real line they can be used to compute the homology of the domain.

Corollary 8.2.5 ("The Barcode Trick"). Assume Y is compact and $f : Y \to X$ is a cellular map with $f(Y) = X \subset \mathbb{R}$. For each i, let B_i denote the barcode associated to the ith absolute homology cosheaf. The following is true:

$$H_i(Y;k) \cong H_0(X;\widehat{F}_i) \oplus H_1(X;\widehat{F}_{i-1}) \cong H_0^{BM}(B_i) \oplus H_1^{BM}(B_{i-1})$$

Proof. This is an immediate corollary of Theorem 8.2.15 and Lemma 8.2.4.

Let $h: S^2 \to \mathbb{R}$ be the standard height function on the two sphere. In Figure 29 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.

$$\widehat{F}_0: \qquad k \longleftarrow k \longrightarrow k$$

Similarly, taking H_1 of the fiber also yields an indecomposable cosheaf

$$\widehat{F}_1: \quad 0 \longleftarrow k \longrightarrow 0$$

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



Example 8.2.6 (Height function on the Two Sphere).

FIGURE 29: Barcodes and the Two Sphere

lexicographic ordering on the cells. For \widehat{F}_0 we get the following boundary matrix and homology groups:

$$\vartheta_1 = \begin{bmatrix} -1 \\ 1 \end{bmatrix} : k_a \to k_x \oplus k_y \qquad \Rightarrow \qquad H_1(X; \widehat{F}_0) = 0 \quad H_0(X; \widehat{F}_0) = k.$$

For \widehat{F}_1 the computation is even easier:

$$\vartheta_1: k_a \to 0 \qquad \Rightarrow \qquad H_1(X; F_1) = k \quad H_0(X; F_1) = 0$$

One can then check 8.2.5 simply as follows:



FIGURE 30: Barcodes for the Cone

Example 8.2.7 (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 30. One could use stratified Morse theory as a substitute, but we'll use cosheaf homology. Here

the cosheaf \widehat{F}_0 is the same as the previous example; we will not repeat the computation. The cosheaf \widehat{F}_1

$$\widehat{F}_1: \quad 0 \longleftarrow k \longrightarrow k$$

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.

$$\vartheta_1 = \mathrm{id} : k_a \to k_y \qquad \Rightarrow \qquad H_1(X; \widehat{F}_1) = 0 \quad H_0(X; \widehat{F}_1) = 0$$

Checking 8.2.5 again gives

$$H_0(X; \hat{F}_1) = 0 \qquad H_1(X; \hat{F}_1) = 0$$
$$H_0(X; \hat{F}_0) = k \qquad H_1(X; \hat{F}_0) = 0$$
$$H_0(C) = k \qquad H_1(C) = 0 \qquad H_2(C) = 0$$

Let's illustrate the utility of the barcode trick by computing cosheaf homology over X using two different methods:

- Using the computational formulae of section 6.2
- Determining the barcode decomposition and applying claim 8.2.4.

Example 8.2.8 (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 31 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 :

$$0 \longleftrightarrow k_a \longrightarrow k_y^2 \longleftrightarrow k_b^2 \longrightarrow k_z^2 \longleftrightarrow k_c \longrightarrow 0$$

Here the maps from k_a to k_u^2 and k_c to k_z^2 are the diagonal maps

$$\mathbf{r}_{z,a} = \begin{bmatrix} 1 \\ 1 \end{bmatrix} = \mathbf{r}_{z,c}$$



FIGURE 31: Barcodes for Bott's torus

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:

$$\partial_{1} = \begin{bmatrix} 1 & -1 & 0 & 0 \\ 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & -1 \\ 0 & 0 & 1 & -1 \end{bmatrix} \qquad H_{1}(X;\widehat{F}_{1}) = < \begin{bmatrix} 1 \\ 1 \\ 1 \\ 1 \\ 1 \end{bmatrix} > \qquad H_{0}(X;\widehat{F}_{1}) \cong k$$

However, if we change our bases as follows

$$\begin{bmatrix} y_1' = y_1 \\ y_2' = y_1 + y_2 \end{bmatrix} \begin{bmatrix} b_1' = b_1 \\ b_2' = b_1 + b_2 \end{bmatrix} \begin{bmatrix} z_1' = z_1 \\ z_2' = z_1 + z_2 \end{bmatrix}$$

then our cosheaf \hat{F}_1 can then be written as the direct sum of two indecomposable cosheaves:

$$0 \longleftrightarrow 0 \longrightarrow k_{y'_{1}} \longleftrightarrow k_{b'_{1}} \longrightarrow k_{z'_{1}} \longleftrightarrow 0 \longrightarrow 0$$
$$0 \longleftrightarrow k_{a} \longrightarrow k_{y'_{2}} \longleftrightarrow k_{b'_{2}} \longrightarrow k_{z'_{2}} \longleftrightarrow k_{c} \longrightarrow 0$$

Hence it is apparent that

$$H_{i}(X; \widehat{F}_{1}) \cong H_{i}^{BM}([--]) \oplus H_{i}^{BM}(]--[).$$

From which the homology of the torus can be directly observed.

$$\begin{split} H_0(X;\widehat{F}_1) &= k & H_1(X;\widehat{F}_1) = k \\ H_0(X;\widehat{F}_0) &= k & H_1(X;\widehat{F}_0) = k \\ H_0(T) &= k & H_1(T) = k^2 & H_2(T) = k \end{split}$$

8.2.2 Local-to-Global Computations via Cellular Sheaves

Aside from the decoration of cosheaves, section 8.2.1 is completely classical and can be stated much more generally using non-cellular sheaves and cosheaves. However, one can still use cellular versions to compute homological information of non-cellular spaces as already noted by the author in [CGN13].

The Čech Approach

Recall that any topological space X equipped with a cover \mathcal{U} has an associated simplicial approximation $N_{\mathcal{U}}$ given by the nerve construction considered in Definition 2.1.2. The nerve theorem tells us when this approximation is "good enough" for the purposes of cohomology.

Theorem 8.2.9 (Nerve Theorem [Ler45, Bor48]). Given a topological space X and a cover \mathcal{U} , if the support $U_{\sigma} \subset X$ of each $\sigma \in N_{\mathcal{U}}$ is acyclic (i.e., the reduced cohomology $\tilde{H}^{\bullet}(U_{\sigma}; R) = 0$ vanishes), then $H^{\bullet}(N_{\mathcal{U}}; R) \cong H^{\bullet}(X; R)$.

Typically, the coarsest covers do not satisfy the acyclicity assumption. One achieves this by refinement of the cover, with the additional cost of more simplices in $N_{\mathfrak{U}}$. However, one can dodge this refinement by recording the cohomology of the intersection as a cellular sheaf on $N_{\mathfrak{U}}$.

Definition 8.2.10 (Čech Sheaves). The **Čech cellular sheaves** \check{C}^n associated to the cover \mathcal{U} of a space X are defined on the nerve $N_{\mathcal{U}}$ by the following data. Each $\sigma \in N_{\mathcal{U}}$ is assigned the vector space $\check{C}^n(\sigma) = H^n(\mathcal{U}_{\sigma}; k)$ and each face relation $\sigma \subset \tau$ is assigned the linear map $\check{C}^n_{\sigma\tau} : H^n(\mathcal{U}_{\sigma}; R) \to H^n(\mathcal{U}_{\tau}; R)$ arising from the inclusion of supports $\mathcal{U}_{\tau} \hookrightarrow \mathcal{U}_{\sigma}$.

If all simplex supports are acyclic, then \check{C}^0 reduces to the constant sheaf on $N_{\mathcal{U}}$ and all other \check{C}^n s are trivial; in the absence of acyclicity assumptions, the following result yields a simple correction.

Proposition 8.2.11. Let X be a topological space and \mathcal{U} a cover whose nerve $N_{\mathcal{U}}$ is at most one-dimensional. Then, for each $n \in \mathbb{N}$,

$$\mathrm{H}^{\mathrm{n}}(\mathrm{X};\mathrm{k}) \cong \mathrm{H}^{0}(\mathrm{N}_{\mathrm{\mathcal{U}}};\check{\mathrm{C}}^{\mathrm{n}}) \oplus \mathrm{H}^{1}(\mathrm{N}_{\mathrm{\mathcal{U}}};\check{\mathrm{C}}^{\mathrm{n}-1}).$$

We note that similar results have been obtained by Burghelea and Dey [BD13], as well as Carlsson, de Silva, and Morozov [CdSM09] in the context of zig-zag persistence. The difference between their results and ours is that their results depend on the decomposition of zig-zag persistence modules into indecomposable modules (barcodes). Our result makes the recognition that these modules are rightly conceived as sheaves over a linear nerve with a cohomology that can be quickly computed using discrete Morse theory [CGN13].

Proposition 8.2.11 generalizes the familiar Mayer-Vietoris long exact sequence, as the next example shows.

Example 8.2.12 (Mayer-Vietoris). The Mayer-Vietoris Theorem states that given an open cover of X by two open sets $\mathcal{U} = \{A, B\}$ we have the following exact sequence of R-modules and maps:

Ideally one can determine the unknown cohomology of $X = A \cup B$ by inspecting the terms on either side. More formally, one uses universal constructions to force $H^n(X)$ into a more manageable short exact sequence:



The dotted maps are defined by the universal property of the cokernel and kernel, respectively. Over an arbitrary coefficient ring R we would have to solve an extension problem in order to infer $H^n(X)$. If we take R to be a field k, then every short exact sequence splits and we can deduce that

$$H^{n}(X) \cong \ker(\delta^{0}_{n}) \oplus \operatorname{cok}(\delta^{0}_{n-1}) \cong H^{0}(N_{A,B}; \check{C}^{n}) \oplus H^{1}(N_{A,B}; \check{C}^{n-1})$$

where $N_{A,B}$ is the unit interval, viewed as the nerve of a two-element cover.

The Leray Approach

As Section 8.2.1 already indicated with examples, one can compute homological information of a space Y with a suitably nice map $f: Y \to X$. We view this construction from a different perspective. By assuming that X comes with a cover \mathcal{V} , having nerve N_{\mathcal{V}}, one can pull-back the associated Čech sheaf on N_{\mathcal{V}} along f to yield local information about Y.

Definition 8.2.13 (Leray Sheaves). The Leray cellular sheaves L^n associated to a map $f : Y \to X$ and a cover \mathcal{V} of $f(Y) \subset X$ are defined over the nerve $N_{\mathcal{V}}$ as follows. Each simplex $\sigma \in N_{\mathcal{V}}$ is assigned the cohomology of the preimage of its support, i.e., $L^n(\sigma) = H^n(f^{-1}(V_{\sigma});k)$; furthermore, each face relation $\sigma \subset \tau$ is assigned the map induced on cohomology by the inclusion $f^{-1}(V_{\tau}) \hookrightarrow f^{-1}(V_{\sigma})$.

Remark 8.2.14. • We will sometimes use the notation F^n in place of L^n to emphasize the association to $f: Y \to X$.

• The absolute homology cosheaf \widehat{F}^n of Definition 8.2.1 is clearly the cosheaf-theoretic version of Lⁿ when the cover is given by the open stars of the cells.

• A More general version of the Leray sheaf is given by the right derived pushforward of the constant sheaf along the map f.

In the special case where X = Y and f is the identity map, the Leray sheaves clearly coincide with the Čech sheaves associated to the cover \mathcal{V} of X. Thus, the following result generalizes Proposition 8.2.11.

Theorem 8.2.15. Let $f : Y \to X$ be continuous. Assume a cover \mathcal{V} of the image $f(Y) \subset X$ whose nerve $N_{\mathcal{V}}$ is at most one-dimensional. Then, for each $n \in \mathbb{N}$,

$$H^{n}(Y;k) \cong H^{0}(N_{\mathcal{V}};L^{n}) \oplus H^{1}(N_{\mathcal{V}};L^{n-1}).$$

Proof. The theorem is a simple consequence of the Leray spectral sequence which packages the cohomology of Y into a coefficient system over the space X from a map $f: Y \rightarrow X$ [McCo1]. The restriction to a one-dimensional nerve forces the spectral sequence to collapse on the second page and hence yield the desired isomorphisms. More precisely, for each open $V \subset f(Y)$, let $C^n(V; R)$ denote the vector space freely generated by the set of all cochains defined on V. Clearly if $V \subset U$, then there is a surjection $C^n(U; R) \rightarrow C^n(V; R)$ defined by restriction of cochains. The sheaf \widetilde{C}^n associated to this presheaf of singular cochains is consequently *flabby* (see [Ramo5, p. 97]).

Consider the following double complex of vector spaces:

$$\begin{array}{c} & & & & & & & & \\ \uparrow & & & & \uparrow & & & \\ C^{2}(Y) \longrightarrow \bigoplus_{\dim \sigma=0} \widetilde{C}^{2}(f^{-1}(V_{\sigma})) \longrightarrow \bigoplus_{\dim \tau=1} \widetilde{C}^{2}(f^{-1}(V_{\tau}) \longrightarrow 0 \\ & & & \uparrow & & \\ C^{1}(Y) \longrightarrow \bigoplus_{\dim \sigma=0} \widetilde{C}^{1}(f^{-1}(V_{\sigma})) \longrightarrow \bigoplus_{\dim \tau=1} \widetilde{C}^{1}(f^{-1}(V_{\tau}) \longrightarrow 0 \\ & & & \uparrow & & \\ C^{0}(Y) \longrightarrow \bigoplus_{\dim \sigma=0} \widetilde{C}^{0}(f^{-1}(V_{\sigma})) \longrightarrow \bigoplus_{\dim \tau=1} \widetilde{C}^{0}(f^{-1}(V_{\tau})) \longrightarrow 0 \end{array}$$

It follows from standard results [Bre97, Thm II.5.5, Thm III.4.13]) that the rows are exact. By the acyclic assembly lemma [Wei94], the spectral sequence converges to the cohomology of the leftmost column, i.e., $H^{\bullet}(Y;k)$. If one takes cohomology in the ver-

tical direction, one obtains the defined cochain groups associated to the Leray cellular sheaves Lⁿ:

Taking cohomology horizontally corresponds precisely to computing separately (in parallel, if one wishes) the cohomology of the Leray sheaves L^n over N_V , thus producing the final stable page of the spectral sequence.



Over a general ring R, these terms prescribe a filtration of the cohomology, giving rise to extension problems; however, over a field k one can read off the cohomology directly. \Box

Note that the proof indicates precisely where we require the one-dimensional nerve restriction. Without this assumption in place, the second page of the spectral sequence may not be stable and the conclusion of the theorem need not hold.

A Unifying Perspective

There is a more sophisticated version of the nerve described originally by Segal [Seg68] which is homotopically faithful to the underlying space independent of the particulars

of the cover. This notion has been used in recent applications [ZCo8] and parallelizations for homology computation [LZ].

Definition 8.2.16 (Mayer Vietoris Blowup). Let X be a topological space equipped with a cover \mathcal{U} with nerve $N_{\mathcal{U}}$. The **Mayer Vietoris blowup** $M_{\mathcal{U}}$ associated to \mathcal{U} is a subset of the product $X \times N_{\mathcal{U}}$ defined as follows. The pair (x, s) lies in $M_{\mathcal{U}}$ if and only if there is some simplex $\sigma \in N_{\mathcal{U}}$ for which $x \in U_{\sigma}$ and $s \in \sigma$.

Remark 8.2.17. The original description given by Segal is that (up to barycentric subdivision) $M_{\mathcal{U}}$ is the classifying space of a topological category, whose objects are pairs (x, U_{σ}) with $x \in U_{\sigma}$ and whose morphisms are pointed inclusions $(x, U_{\sigma}) \rightarrow (y, U_{\tau})$ where $\tau \subset \sigma \subset I$ and I is the indexing set of the cover $\mathcal{U} = \{U_i\}_{i \in I}$.

Segal provides an updated version of the nerve theorem using this construction [Seg68, Prop. 4.1]

Lemma 8.2.18 (Generalized Nerve Theorem). If X is a paracompact Hausdorff space and $\mathcal{U} = \{U_i\}_{i \in I}$ is an open cover of X, then $M_{\mathcal{U}}$ is homotopic to X.

Proof. An explicit proof is provided by Segal using linear homotopies. Here we take a slightly higher-brow approach.

Being a subset of the product, $M_{\mathcal{U}}$ is equipped with natural surjective projection maps



The map ρ_1 has contractible fibers: for any $x \in X$, we have $\rho_1^{-1}(x) = \{x\} \times \sigma_x$ where σ_x is the unique simplex of maximal dimension whose support contains x. Thus, by Quillen's Theorem A [Qui73], the Mayer-Vietoris blowup is homotopy-equivalent to X via ρ_1 in full generality.

On the other hand, it is easy to see that the map ρ_2 fails to have contractible fibers precisely when the simplex supports are not contractible. In fact, given $s \in N_{\mathcal{U}}$, the fiber $\rho_2^{-1}(s)$ has the homotopy type of the support of σ_s , which is the unique simplex of maximal dimension whose realization contains s. Since cohomology is a homotopy invariant, this leads to the following observation which unifies the Čech and Leray approaches.

Proposition 8.2.19. The Leray cellular sheaves L^n associated to the map $\rho_2 : M_{\mathcal{U}} \to N_{\mathcal{U}}$, where $N_{\mathcal{U}}$ is covered by (small neighborhoods of the topological) simplices $\{\sigma\}_{\sigma \in N_{\mathcal{U}}}$, are isomorphic to the Čech cellular sheaves \check{C}^n associated to the cover \mathcal{U} .

- **Remark 8.2.20.** The commonality between the Čech and Leray approaches comes as no surprise to anyone sufficiently familiar with spectral sequences (and would have surprised neither Čech nor Leray).
 - Both strategies are examples of *distributed* cohomology computation because in order to determine the sheaf Čⁿ or Lⁿ, one only needs to compute cohomology locally: of a non-trivial intersection of covering sets in the former case, or of a small neighborhood of the fiber f⁻¹(x) in the latter case. In principle, one can assign each local computation to a different processor, compute the appropriate sheaf cohomology over a decidedly nicer space (either N_u or Y depending on the circumstances), and aggregate this information to compute the desired cohomology of X.
 - By taking the appropriate linear duals and working with cosheaves, all of our results transform to computations of homology rather than cohomology.

8.2.3 Level Set Persistence Determines Sub-level set Persistence

We will use Theorem 8.2.15 to obtain a non-obvious theorem in persistence, namely that level set persistence determines sub-level set persistence. By making use of 8.2.4 we illustrate how one can take the absolute homology cosheaves (or Leray sheaves) equipped with a barcode decomposition and sweep from left to right, applying 8.2.5 to determine the barcodes of the associated sublevel set persistence modules. An example is drawn in Figure 32. Stated formally, we have the following theorem.

Theorem 8.2.21 (Level Set to Sublevel Set Persistence). Let F^k denote the Leray sheaf associated to a proper map $f : X \to \mathbb{R}$, whose stalk at x is the cohomology of the fiber $H^k(f^{-1}(x))$. We can define a functor

$$S^k: (\mathbb{R}, \leqslant)^{op} \to \text{Vect} \qquad S^k(t) := H^0((-\infty, t]; F^k) \oplus H^1((-\infty, t]; F^{k-1}) \cong H^k(f^{-1}(-\infty, t])$$

whose value records the cohomology of the entire sublevel set. The maps

$$S^k(t') o S^k(t) \qquad t \leqslant t'$$

are defined sheaf theoretically by observing that we have maps

$$\mathsf{H}^{0}((-\infty, \mathbf{t}']; \mathsf{F}^{k}) \to \mathsf{H}^{0}((-\infty, \mathbf{t}]; \mathsf{F}^{k}) \qquad \mathsf{H}^{1}((-\infty, \mathbf{t}']; \mathsf{F}^{k-1}) \to \mathsf{H}^{1}((-\infty, \mathbf{t}]; \mathsf{F}^{k-1})$$

the sum of which define the desired map.



FIGURE 32: Determining Sub-level Set from Level Set Persistence

Observe that since $f: X \to \mathbb{R}$ is a proper map, the restriction of f to $X_{\leq t} := f^{-1}((-\infty, t])$ is also a proper map. Consequently, we can apply Theorem 8.2.15 to the space $X_{\leq t}$ instead, but we have to restrict the sheaves to the subspace $(-\infty, t]$. Fortunately, restriction of a sheaf to a subspace is a standard operation in the Grothendieck six-functor formalism, presented in 2.5.2: If $\iota : (-\infty, t] \hookrightarrow \mathbb{R}$ is the inclusion, then the application of Theorem 8.2.15 to the restriction reads

$$\mathsf{H}^k(X_{\leqslant t};k) \cong \mathsf{H}^0((-\infty,t];\iota^*F^k) \oplus \mathsf{H}^1((-\infty,t];\iota^*F^{k-1})$$

The upshot of this formula is that we can define a family of vector spaces, one for each $t \in \mathbb{R}$ that records the cohomology of the sublevel set $X_{\leq t}$

$$S(t) := H^0((-\infty, t]; \iota^* F^k) \oplus H_1((-\infty, t]; \iota^* F^{k-1}))$$

given by computing sheaf cohomology of the restriction of the Leray sheaves to the subspace $(-\infty, t]$. What remains to be shown is that there are maps

$$S(t') \rightarrow S(t)$$
 $t \leq t'$

that can be defined purely sheaf-theoretically. To do this, we will make use of some standard adjunctions in sheaf theory.

Theorem 8.2.22. Let $f : Y \to Z$ be a continuous map. The functors $f^* : Shv(Z) \to Shv(Y)$ and $f_* : Shv(Y) \to Shv(Z)$ form an adjoint pair (f^*, f_*) and thus

$$\operatorname{Hom}_{\operatorname{Shv}(Y)}(f^*G, F) \cong \operatorname{Hom}_{\operatorname{Shv}(Z)}(G, f_*F).$$

In the above adjunction for sheaves, let $Y = (-\infty, t]$, $Z = (-\infty, t']$ and f = j be the inclusion of Y as a closed subspace of Z. Observe that if we set $F = j^*G$ in the above adjunction, then we get an isomorphism

$$\operatorname{Hom}_{\operatorname{Shv}(Y)}(j^*G, j^*G) \cong \operatorname{Hom}_{\operatorname{Shv}(Z)}(G, j_*j^*G).$$

that is natural in G. This defines the unit of the adjunction:

$$\mathrm{id}_{\mathsf{Shv}}(Y) o \mathfrak{j}_*\mathfrak{j}^*$$

We recall a basic theorem about the pushforward sheaf along a closed immersion [Ive86, II.5 p. 102].

Proposition 8.2.23. For $j : Y \hookrightarrow Z$ the inclusion of a closed subspace, the functor $j_* :$ Shv(Y) \rightarrow Shv(Z) is exact, i.e. it sends exact sequences of sheaves to exact sequences of sheaves. Moreover, j^* is always exact.

Lemma 8.2.24. Suppose $j : Y \hookrightarrow Z$ is the inclusion of a closed subspace and F is a sheaf on Z, then there is an induced map from the cohomology of F on Z to the cohomology of j^*F on Y.

$$H^{\iota}(Z;F) \rightarrow H^{\iota}(Y;j^{*}F)$$

Proof. We can read off the proof from the following diagram of spaces.



Sheaf cohomology is defined as the right derived functor of pushforward to a point. If we want to compute sheaf cohomology of F, one takes an injective resolution of F

$$0 \to F \to I^\bullet$$

and applies p_{Z*} to the injective resolution.

$$Rp_{Z*}F := p_{Z*}I^{\bullet}$$

This results in a chain complex of vector spaces, whose cohomology is the sheaf cohomology of F. We usually save this step for last as it takes us out of the category of chain complexes of vector spaces and into the category of graded vector spaces. This is written as follows.

$$R^{i}p_{Z*}F := H^{i}(p_{Z*}I^{\bullet}) =: H^{i}(Z;F)$$

Since j^{*} is exact we will consider an injective resolution of F and pull that back to an injective resolution of j^{*}F. Observe the following string of identities.

$$Rp_{Y*}j^*F := p_{Y*}j^*I^{\bullet}$$
$$= (p_Z \circ j)_*j^*I^{\bullet}$$
$$= p_{Z*}j_*j^*I^{\bullet}$$

The unit of the adjunction defines a map of sheaves

$$F
ightarrow j_* j^* F$$
,

which also defines a map on complexes of sheaves and hence injective resolutions.

$$I^{\bullet} \rightarrow j_* j^* I^{\bullet}$$

Because j^* and j_* are exact and preserve injectives, $j_*j^*I^\bullet$ is an injective resolution of j_*j^*F , thus

$$Rp_{Z*}j_*j^*F = p_{Z*}j_*j^*I^{\bullet} = Rp_{Y*}j^*F$$

and hence the unit of the adjunction defines a map

$$Rp_{Z*}F \to Rp_{Y*}j^*F \implies H^i(Z;F) \to H^i(Y;j^*F.)$$

Remark 8.2.25 (Abuse of Notation). Common practice in the sheaf literature is to suppress the notation j*F and to just write

$$H^{i}(Y;F) := H^{i}(Y;j^{*}F).$$

The reasoning is that F is a sheaf on Z and hence the only way to parse the formula on the left is to realize that the sheaf must be restricted to the subspace Y.

As a corollary we obtain our desired result, Theorem 8.2.21.

8.3 MULTIDIMENSIONAL PERSISTENCE

One of the single greatest theoretical challenges to topological data analysis is a foundation for multi-dimensional persistence [Les12]. To consider why data analysts might want such a thing, consider the following example.

Suppose X is the shape depicted in Figure 33. A common feature of interest in applications [LSL⁺13] is the presence of *flares* or *tendrils*. Sublevel set persistence provides a method for detecting such features. Consider the p**th eccentricity functional** on X:

$$\mathsf{E}^{\mathsf{p}}(\mathsf{x}) := \left(\int_{\mathsf{y} \in \mathsf{X}} \mathsf{d}(\mathsf{x}, \mathsf{y})^{\mathsf{p}} \mathsf{d} \mathsf{y} \right)^{\mathsf{p}}.$$

If we filter by superlevel sets, the four endpoints of the perceived flares in Figure 33 will come into view. Said using homology, there are a suitable large range of values t for which $E_{\ge t}^p := \{x \in X | E^p(x) \ge t\}$ will have

$$H_0(E^p_{\geqslant t};k) \cong k^4.$$

This formally expresses the four flare-like features we see in the space X.

Now suppose that we are not just interested in the number of eccentric features, but rather we are interested in holes with high eccentricity value, i.e. the persistence module

$$H_1(E_{\geq t}^p;k)$$

is of interest. However, what size of hole is of interest, and what can be regarded as noise? In other words, what is the behavior of the two-parameter family of vector spaces

$$MP_1(t,r) := H_1((E_{\geq t}^p)^r;k)$$

where Y^r denotes the set of points within distance r of a subspace Y. This family of vector spaces defines a functor

$$MP_1: (\mathbb{R}^2, \leq) \rightarrow Vect$$
 where $MP_1(t, r) := H_1((E_{\geq t}^p)^r; k)$

where \mathbb{R}^2 is viewed as a poset under the relation $(t, r) \leq (t', r')$ if and only if t' - t and r' - r are non-negative numbers.

This gives the general definition of a multidimensional persistence module as introduced in [CZ09].

Definition 8.3.1 (Multi-dimensional Sub-Level Set Persistence). Suppose we are given a map $f : X \to \mathbb{R}^n$, with coordinate functions $f(x) = (f_1(x), \dots, f_n(x))$. The **ith multidimensional persistence module** is defined to be the functor

$$MP_i: (\mathbb{R}^n, \leq) \to \text{Vect} \qquad (t_1, \dots, t_n) \rightsquigarrow H_i(\{x \in X \mid f_j(x) \leq t_j, 1 \leq j \leq n\}; k)$$

Critique of the Definition of Multi-D Persistence

There are a few problems that such a definition suffers from:

- Such a definition is strongly dependent on the particular choice of basis given to Rⁿ. If one is given an abstract function f, valued in say a manifold M, then the definition fails.
- The definition is not local.
- There is no decomposition theorem akin to Theorem 6.3.3. There are no barcodes. This is fundamentally due to Gabriel's Theorem 6.3.23, but an explicit example is provided in [CZ09, Sec.5.2], which suggests other invariants as well.



FIGURE 33: A Shape Described with Multi-D Persistence

Sheaves and cosheaves overcome the first problem by putting level set persistence as the primary object of interest. Locality is also provided by using sheaves. Our definition of a multi-dimensional persistence module is simply the Leray sheaf.

Definition 8.3.2 (Multi-D Persistence as the Derived Pushforward). Suppose $f : Y \to X$ is a continuous map. The right derived pushforward of the constant sheaf Rf_*k_Y is gotten by taking a resolution of k_Y by the complex of singular cochains \tilde{C}^{\bullet} and pushing forward along f. The Leray sheaves are the nth cohomology sheaves of this complex, i.e.

$$\mathbf{F}^{\mathbf{n}} := \mathcal{H}^{\mathbf{n}}(\mathbf{f}_{*}\widetilde{\mathbf{C}}^{\bullet}) =: \mathbf{R}^{\mathbf{n}}\mathbf{f}_{*}\mathbf{k}_{\mathbf{Y}}.$$

If f is proper, then the stalks of F^n record the cohomology of the level set, i.e. $F_x^n \cong H^n(f^{-1}(x);k)$. We define n**th multi-dimensional level set persistence** of a *proper* map $f: Y \to X$ to be the Leray sheaf F^n .

Remark 8.3.3. If f is not proper, then one can use the pushforward with compact supports to encode the compactly supported cohomology of the fiber.

One can always obtain the traditional sub-level set definition from this definition by using a multi-dimensional version of Theorem 8.2.21, the Leray spectral sequence.

Sheaves also suffer from the lack of a nice set of indecomposables, but the Grothendieck operations provide one possible approach.

8.3.1 Generalized Barcodes

The need for a generalized notion of a barcode comes from the need to communicate topological summaries to non-topologists. A scientist can easily understand a histogram and the barcode is only subtly different from a histogram.

Definition 8.3.4 (Generalized Barcodes). Suppose F^n is the Leray sheaf associated to a proper map $f : Y \to X$. A **generalized barcode** for F^n is the expression of F as follows.

$$\mathsf{F}^{\mathfrak{n}} \cong \bigoplus_{b \in B} (\mathfrak{j}_b)_! k_{\mathsf{Z}_b}$$

where $j_b : Z_b \to X$ are maps indexed by the barcode B and where each $(j_b)_! k_{Z_b}$ is assumed to be indecomposable.

Remark 8.3.5. It is not clear when or if such a generalized barcode exists for a given $f: Y \rightarrow X$ and n.

An example barcode is provided in Figure 34 and discussed in the next example.

Example 8.3.6 (Shadow of the Sphere). Consider the standard Euclidean sphere S^2 embedded in \mathbb{R}^3 . Let $f : S^2 \to \mathbb{R}^2$ be the projection onto the first two factors of \mathbb{R}^3 . The image $X = f(S^2)$ is the closed unit disk.

One can use a multi-D analog of Corollary 8.2.5 to determine the homology of the two-sphere.

$$H_0(S^2) \cong H_0(X; \widehat{F}_0) \cong k \qquad H_1(S^2) \cong H_1(X; \widehat{F}_0) \cong 0 \qquad H_2(S^2) \cong H_2(X; \widehat{F}_0) \cong k$$



FIGURE 34: Two Dimensional Barcodes for the Sphere

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 [GH11]. 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 9.0.7. 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 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 9.0.8. 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 $\mathbb{R}_{\geq 0}$ or the non-negative integers $\mathbb{Z}_{\geq 0}$.

Definition 9.0.9 (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 $F : X \rightarrow Vect$ constructed as follows:

- To an edge e ∈ X we let F(e) = k^{c(e)}, a vector space of dimension equal to the capacity.
- To a vertex v we let $F(v) = k^{c(v)} \cong \bigoplus_{e_i \in in(v)} k^{c(e)}$.
- The restriction maps are given by ordinary projections for the incoming edges, i.e. ρ_{e_i,ν} := proj_{F(e_i)} for e_i ∈ in(ν), but for the outgoing edges some non-trivial coding may be performed, i.e. any linear map Φ_{e_k,ν} : F(ν) → F(e_k) for e_k ∈ out(ν) will do. We write Φ_ν = ⊕_{e_k∈out(ν)}Φ_{ν,e_k} for the **total coding** through ν.

Remark 9.0.10. It should be noted that in [GH11], they do not use cellular sheaves. This was primarily due to the lack of a good reference.



FIGURE 35: Graph with Decoding Wire

In [GH11] 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. **Example 9.0.11.** Consider the graph in Figure 35 with constant capacity function c = 1. Consequently, all edges and vertices get a one dimensional vector space $k = \mathbb{R}$ with the exception of $F(t) \cong k^2$. Define the coding maps $\rho_{\alpha,s} = id = \rho_{b,s}$ and

$$\rho_{c,t} = \begin{bmatrix} \frac{1}{2} & \frac{1}{2} \end{bmatrix}.$$

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, $H^0(X; F) \cong H^1(X; F) \cong k$. 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 8.3.4. 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 9.1.1 (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_{ν} can be written as a binary matrix with at most one 1 in each column and row. Said another way, at a vertex ν 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 ν 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 9.1.2 (Structure Theorem for Routing Sheaves). Suppose X = (V, E, h, t) is a directed graph, then every routing sheaf $F : X \rightarrow Vect$ 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 9.1.3 (Duality). For any routing sheaf F one has

$$\mathrm{H}^{0}(\mathrm{X};\mathrm{F})\cong\mathrm{H}^{1}(\mathrm{X};\mathrm{F}).$$

Proof. By Proposition 9.1.2, 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 S¹ establishes the corollary.

One can also prove this duality in the more general setting of network coding sheaves via a simple combinatorial argument.

Proposition 9.1.4. For an any network coding sheaf F, we have the following isomorphisms:

$$\bigoplus_{\nu} F(\nu) \cong \bigoplus_{e} F(e) \qquad H^{0}(X;F) \cong H^{1}(X;F)$$

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.

$$\bigoplus_{e \in \mathfrak{in}(\nu)} F(e) = F(\nu).$$

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. \Box

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 H⁰ as a vector space spanned by independent information flows, but in the case of routing sheaves, H¹ 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 [EFS₅₆, FF₅₆, Sey₇₇], 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 36: No Decoding Edge

FIGURE 37: Decoding Edge

Example 9.2.1 (Decoding Edge and Barcodes). In Figures <u>36</u> and <u>37</u>, 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 <u>36</u> decomposes as the constant sheaf on two half open intervals and two closed intervals:

$$F_{no} \cong (j_{0})_{!} k_{[0,1)} \oplus (j_{b})_{!} k_{[0,1)} \oplus (i_{r})_{*} k_{[0,1]} \oplus (i_{g})_{*} k_{[0,1]} \Rightarrow H^{0}(X; F_{no}) \cong k^{2}.$$

This is bad if we want our sheaf to encode cohomologically the presence of source-totarget information paths. However, with the use of a decoding edge (decoding edges) the network decomposes into only two half open intervals:

$$\mathsf{F}_{de} \cong (\mathfrak{j}_r)_! \mathsf{k}_{[0,1)} \oplus (\mathfrak{j}_g)_! \mathsf{Sk}_{[0,1)} \Rightarrow \mathsf{H}^0(\mathsf{X}; \mathsf{F}_{de}) \cong \mathsf{O}$$

which was wanted.

9.3 NETWORK CODING SHEAF HOMOLOGY

One of the virtues of the network coding sheaves is that they are easy to construct, have interesting sheaf cohomology, and provide lots of examples. As such, we will use them as a testing ground for the new theory of sheaf homology developed in Section 7.4.



FIGURE 38: Network Coding Sheaf
Example 9.3.1. Consider the network coding sheaf implied by Figure <u>38</u>. Viewed as a diagram of vector spaces, it takes the following form:



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:



A projective sheaf that surjects onto R is supported on the star of s and t respectively, i.e. $P_0 := \{s\} \oplus \{t\}$:



The kernel sheaf of the natural transformation $P_0 \Rightarrow R$ is also projective, which we call P_1 and finishes the projective replacement of the sheaf R.



If we take the colimit of P_1 and P_0 separately, the sheaf map $P_1 \rightarrow P_0$ induces a map on colimits that defines the boundary in the chain complex computing sheaf homology:

$$\partial_1 : k^4 \to k^2$$
 $\begin{bmatrix}
1 & 1 & 0 & 1 \\
-1 & 0 & 1 & -1
\end{bmatrix}$
 \Rightarrow
 $H_0(X; R) = 0$
 $H_1(X; R) \cong k^2$

Repeating the same reasoning for the green barcode G yields homology groups $H_0(X;G) = 0$ and $H_1(X;G) \cong k$. Since our original network coding sheaf F is a direct sum $R \oplus G$ we obtain that the sheaf homology of the sheaf in Figure 38 is

$$H_0(X;F) = 0$$
 $H_1(X;F) \cong k^3$.



FIGURE 39: Network Coding Sheaf with Two Decoding Wires

Exercise 9.3.2. 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 39 are

 $H_0(X;F)=0 \qquad H_1(X;F)\cong k^8.$

10

SHEAVES AND COSHEAVES IN SENSOR NETWORKS

"What strength, what art can then suffice, or what evasion bear him safe through the strict sentries and stations thick of angels watching round?"

— John Milton, Paradise Lost, Book II, Line 410 [MP08]

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 10.3.1. An approach of the author and Robert Ghrist is then considered in Section 10.3.2. This approach succeeds where the previous approach fails, but it too suffers from giving false positives, as the example in Proposition 10.3.8 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 10.4.4 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 10.1.1 (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 10.1.2 (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 10.1.3 (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



FIGURE 40: Sensors Distributed in a Plane

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?

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 [dSGo6b], 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.¹

Definition 10.2.1. Let D be a spatial region of interest and denote by $D \times [0, 1]$ a region of space-time. This carries with it a map that keeps track of time via projection onto the second factor, i.e. $\pi_2 : D \times [0, 1] \rightarrow [0, 1]$. We assume that D can be given a cell structure so that the sensors' **coverage region** $S \subset D \times [0, 1]$ and the **evasion region** $E := S^c$ can be written as the union of cells. To study the **intruder problem** is to analyze the associated sheaf of sections of the map $\pi := \pi_2|_E : E \rightarrow [0, 1]$, 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 $s : [0, 1] \rightarrow E$ such that $\pi \circ s = id$.

Example 10.2.2. For the situation depicted in Figure 41, 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.

¹ The author would like to thank Gunnar Carlsson and Rob Ghrist for their insights here.



FIGURE 42: Mobile Sensor Network

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 10.2.3. In Figure 42 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 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 $s : [0, 1] \rightarrow E$ such that $\pi \circ s = id$.

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 naïvely "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 5. 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 11.2.17, 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

To simplify the topology, we focus on the Reeb graph version of Figure 42. This is drawn and labeled in Figure 43. Since everything is occurring in two-dimensional space-time, the only interesting homological invariant of the fiber is H₀. Studying this is equivalent



FIGURE 43: Mobile Sensor Network

to studying the pushforward cosheaf $\hat{F} := \pi_* \hat{k}_E$. In the parlance of [CdS10], this is simply a zigzag module of the following form:

$$\widehat{F}(x) \stackrel{r_{x,a}}{\longleftrightarrow} \widehat{F}(a) \stackrel{r_{y,a}}{\longrightarrow} \widehat{F}(y) \stackrel{r_{y,b}}{\longleftarrow} \widehat{F}(b) \stackrel{r_{z,b}}{\longrightarrow} \widehat{F}(z) \stackrel{r_{z,c}}{\longleftarrow} \widehat{F}(c) \stackrel{r_{w,c}}{\longrightarrow} \widehat{F}(w)$$

$$k_x \longleftarrow k_a \longrightarrow k_y^2 \longleftarrow k_b^3 \longrightarrow k_z^2 \longleftarrow k_c \longrightarrow k_w$$

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:

$$\mathbf{r}_{y,a} = \begin{bmatrix} 1 \\ 0 \end{bmatrix} \quad \mathbf{r}_{y,b} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 1 \end{bmatrix} \quad \mathbf{r}_{z,b} = \begin{bmatrix} 1 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \quad \mathbf{r}_{x,c} = \begin{bmatrix} 0 \\ 1 \end{bmatrix}$$

We can decompose this cosheaf into indecomposables simply by performing the correct change of basis:

$$\begin{bmatrix} y_1' = y_1 \\ y_2' = y_1 - y_2 \end{bmatrix} \begin{bmatrix} b_1' = b_1 - b_2 + b_3 \\ b_2' = b_1 - b_2 \\ b_3' = b_2 - b_3 \end{bmatrix} \begin{bmatrix} z_1' = z_2 \\ z_2' = z_1 - z_2 \end{bmatrix}$$

The reader should match the resulting indecomposables with the barcodes drawn in Figure 43.

$$k_{x} \longleftrightarrow k_{a} \longrightarrow k_{y_{1}'} \longleftrightarrow k_{b_{1}'} \longrightarrow k_{z_{1}'} \longleftrightarrow k_{c} \longrightarrow k_{w}$$

$$0 \longleftrightarrow 0 \longrightarrow k_{y_{2}'} \longleftrightarrow k_{b_{2}'} \longrightarrow 0 \longleftrightarrow 0 \longrightarrow 0$$

$$0 \longleftrightarrow 0 \longleftrightarrow k_{b_{2}'} \longrightarrow k_{z_{2}'} \longleftrightarrow 0 \longrightarrow 0$$

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 10.3.1. 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 \widehat{G} is a cosheaf on Y, then we have that $H_0(Y; \widehat{G}) \cong H_0(X; f_*\widehat{G})$. This follows from the commutativity of the following diagrams and functoriality of pushforward.



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.

$$\pi_* \hat{k}_{\mathsf{E}} \cong \hat{k}_{\mathsf{B}_1} \oplus \dots \oplus \hat{k}_{\mathsf{B}_n}$$
 and

Now we combine this with the fact that homology commutes with direct sums.

$$\mathbf{k} \cong \mathbf{H}_0(\mathsf{E}; \hat{\mathbf{k}}_{\mathsf{E}}) \cong \mathbf{H}_0([0, 1]; \pi_* \hat{\mathbf{k}}_{\mathsf{E}}) \cong \bigoplus_{i} \mathbf{H}_0([0, 1]; \hat{\mathbf{k}}_{\mathsf{B}_i}) \cong \bigoplus_{i} \mathbf{H}_0^{\mathsf{BM}}(\mathsf{B}_i)$$

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 $1_E \in \Gamma(E; \hat{k}_X)$ has support on all of E, the pushforward section $\pi_* 1_E$ that generates the closed barcode must have support on all of [0, 1], since π is surjective.

Remark 10.3.2. 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 10.3.3. Having a barcode associated to $\pi_* \hat{k}_E$ whose support is all of [0, 1] does not indicate the existence of an evasion path.

Remark 10.3.4. 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 8.2.5 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*

In light of the inability of the pushforward cosheaf $\pi_* \hat{k}_E$ 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 $\pi : E \rightarrow [0, 1]$. From this setup, we can extract a cellular map of 1-dimensional cell complexes, normally called $\tilde{\pi} : R(\pi) \rightarrow [0, 1]$, but we will abuse notation and assume that our input $\pi : E \rightarrow [0, 1]$ 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 related² to the network coding sheaves defined in Section 9.

² In a sense, the sheaf defined here gives all possible codings. It approximates a "stack" of network coding sheaves.



FIGURE 44: Sheaf of All Possible Evasion Paths

Definition 10.3.5. 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 $G(v) = k^{out(v)}$ for $v \in S$ and $G(v) = k^{in(v)}$ 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 10.3.6. For a concrete example, where we focus on a small part of a graph, consider the graph in Figure 44. The definition of the sheaf G makes

$$G(v_1) = \langle e_i \otimes f_j | i = 1, 2; j = 1, 2, 3 \rangle \cong k^6 \qquad \rho_{e_i,v_1}(e_j \otimes f_k) = \delta_{ij} \quad \rho_{f_i,v_1}(e_i \otimes f_k) = \delta_{jk}$$

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:

$$\pi_* G(e) \cong \langle e_1, e_2, e_3 \rangle$$
 $\pi_* G \cong G(v_1)$ $\pi_* G(f) = \langle f_1, f_2, f_3 \rangle$



FIGURE 45: Linearized Sheaf of Sections

Example 10.3.7. Let us consider the example drawn in Figure 45, but now with the sheaf just defined. We set $F = \pi_*G$, whose values are below:

 $F(x) \xrightarrow{\rho_{a,x}} F(a) \xleftarrow{\rho_{a,y}} F(y) \xrightarrow{\rho_{b,y}} F(b) \xleftarrow{\rho_{b,z}} F(z) \xrightarrow{\rho_{c,z}} F(c) \xleftarrow{\rho_{c,w}} F(w)$

$$k_x \longrightarrow k_a \longleftrightarrow k_y \longrightarrow k_b^3 \longleftrightarrow k_z \longrightarrow k_c \longleftrightarrow k_w$$

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 45.

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-



FIGURE 46: Counterexample for the Linearized Sheaf of Sections

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 10.3.8. 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 46 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

$$k_x \longrightarrow k_a^2 \longleftrightarrow k_y^3 \longrightarrow k_b^3 \longleftrightarrow k_z^3 \longrightarrow k_c^2 \longleftrightarrow k_w$$

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 [CdS10] provides a sure-fire method of obtaining it. It is left to the reader to verify the the barcodes in Figure 46.

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 $H^0([0, 1]; F) \cong k$ we can infer the existence of one closed barcode, and because this section has global support, the barcode must be long.

Remark 10.3.9 (Dynkin Diagrams and Stalks). Recall that $F := \pi_*G$. Consider the sheaf G implied by Figure 46. 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 D₄; see [EGH⁺11], 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 (sub-representations), 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 kⁿ, usually ℝⁿ or ℂⁿ, and every intruder is tagged with an unchanging property vector v ∈ kⁿ. 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 Z is used to record counts and Zⁿ 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



FIGURE 47: Multi-Modal Sensors Distributed in a Plane

 $S_i = \text{span}\{\xi_i\} =: < \xi_i >$. The act of sensing corresponds to taking a property vector $\nu \in k^n$ and returning a number $\xi_i(\nu)$ 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.

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 10.4.9. 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 48, we consider a situation where we have a sensor capable of detecting "red" properties and a sensor capable of detecting "green" properties.³

³ We use scare quotes to indicate that the terms can be substituted for whatever application is of interest.



FIGURE 48: Two Multi-Modal Sensors

On the nerve of the sensor cover, the organizing diagram of vector spaces is clear.

$$< r^* > \hookrightarrow < r^*, g^* > \longleftrightarrow < g^* > \qquad k \hookrightarrow k^2 \leftrightarrow k$$

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 $< r^* > \text{ or } < g^* > \text{ or } < b^* >$. 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 10.4.1 (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:

$$F(\sigma) = S_{\nu_0} + \dots + S_{\nu_n} \qquad F(\sigma) \hookrightarrow F(\tau) \quad \sigma \leqslant \tau.$$

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:

$$\iota_{\mathsf{F}}:\mathsf{F}\hookrightarrow k_X^{\mathfrak{n}*}.$$

Now suppose we have an intruder, which we imagine as a point in the union of the red and green sensors in Figure 48. 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 co-vector and the property vector v. If $k^n = k^3 = \langle v_r, v_g, v_b \rangle$ 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

$$\mathbf{r}^{*}(\mathbf{v}) = \begin{bmatrix} 1 & 0 & 0 \end{bmatrix} \begin{bmatrix} \mathbf{v}_{r} \\ \mathbf{v}_{g} \\ \mathbf{v}_{b} \end{bmatrix} = \mathbf{v}_{r}$$

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.

$$\begin{bmatrix} \mathbf{r}^*(\mathbf{v}) \\ \mathbf{g}^*(\mathbf{v}) \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix} \begin{bmatrix} \mathbf{v}_{\mathbf{r}} \\ \mathbf{v}_{g} \\ \mathbf{v}_{b} \end{bmatrix} = \begin{bmatrix} \mathbf{v}_{\mathbf{r}} \\ \mathbf{v}_{g} \end{bmatrix}$$

We can package these measurements into a cellular cosheaf, where two observations are the same modulo the properties unobserved by the sensors.

$$<$$
r $>$ \leftarrow > \rightarrow > $\frac{k^3}{<$ 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 48. 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 10.4.2 (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:

- $\widehat{\operatorname{Ann}}(F)(\sigma) = \{ v^* \in \operatorname{G}(\sigma)^* | v^*(\iota(w)) = 0 \quad \forall w \in F(\sigma) \}$
- If $\sigma \subset \overline{\tau}$, 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 $Ann(F) =: \hat{E}$ the **evasion cosheaf**.

Lemma 10.4.3. 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 $\hat{E} \cong \hat{V}(cok(\iota))$ in the diagram below.



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

$$\operatorname{Ann}_{W}(V) = \{ \varphi : W \to k \, | \, \varphi(v) = 0 \, \forall v \in V \} \cong (W/V)^{*}$$

and of course all the restriction maps get sent to restriction maps



1	-	-	-	
- 1				
- 1				
- 1				
ι				l

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 10.4.4** (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



Where $H^k(X; cok(\iota))$ gets identified with the evasion co-sheaf's homology $H_k(X; \hat{E})$ via the linear duality functor, i.e. $V : cok(\iota)) \rightsquigarrow E$.

Proof. The proof is immediate from standard homological algebra techniques. \Box

10.4.2 Indecomposables, Evasion Sets, Generalized Barcodes

One of the drawbacks of Theorem 10.4.4 is that we have no good interpretation of what the sheaf cohomology groups mean. Let's consider again Figure 48, 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 49.

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 $H^i([0, 1]; F) \cong \bigoplus H^i_c(B_i)$, which we have already made heavy use of, we can determine all the sheaf cohomology of interest for this example.

Example 10.4.5 (Red-Green Sensors). By inspection of the indecomposable presentations of the three sheaves F, k_X^3 and $cok(\iota)$ in Figure 49 we see that

 $H^{i}(X;F) \cong 0i = 0,1;$ $H^{0}(X;k_{X}^{3}) \cong k^{3}$ $H^{0}(X;cok(\iota)) \cong H_{0}(X;\hat{E}) \cong k^{3}$

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 10.4.6 (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 .



FIGURE 49: Examining the Short Exact Sequence

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 10.4.7. For any property vector v, the evasion and detection sets form an openclosed decomposition of X, that is

$$X = E_v \cup D_v$$
 $E_v \cap D_v = \emptyset$, E_v open.

When X is compact this means that E_{ν} is compact as well.

We record another easy lemma, connected to our desire to get an indecomposable presentation for our evasion cosheaves.

Lemma 10.4.8. Suppose that all sensors must pull their sensing capabilities from a fixed orthonormal basis of k^{n*} , say v_1^*, \ldots, v_n^* , then the evasion cosheaf splits as a direct sum decomposition of constant cosheaves supported on the evasion sets for v_1, \ldots, v_n

$$\hat{E} \cong \hat{k}_{E_{\nu_1}} \oplus \cdots \oplus \hat{k}_{E_{\nu_n}}$$

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 $k_X \oplus \cdots \oplus k_X$



FIGURE 50: Short Exact Sequence

where we think of each copy of k_X as being the constant sheaf generated by $\langle v_i^* \rangle$. As a consequence we get the following diagram



Now we can use for each factor the following standard short exact sequence of sheaves

 $0 \longrightarrow k_{D_{\nu_i}} \longrightarrow k_X \longrightarrow k_{E_{\nu_i}} \longrightarrow 0$

and thus the cokernel splits as a direct sum $\oplus_{i=1}^{n} k_{E_{\nu_i}}$.

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 10.4.9. If sensor's abilities are pulled from a fixed orthonormal basis v_1^*, \ldots, 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

$$<\nu_i^*> \hookrightarrow <\nu_i^*, \nu_j^*> \longleftrightarrow <\nu_j^*>.$$

For the example considered in Figure 50 assume that the space of properties is two dimensional, spanned by red and green. Then the Theorem 10.4.4 provides the following forcing result

$$0 \to H^0_c(X; F) \cong 0 \to H^0_c(X; k_X^{2*}) \cong k^2 \to H^0_c(X; cok) \to H^1_c(X; F) \cong k \to 0$$

which upon careful inspection reveals that the red evasion set must be disconnected.

Part IV

NOVEL MATHEMATICAL CONTRIBUTIONS

This part represents the mathematical heart of the thesis, although many of its results were motivated by the applications considered in Part III.

Chapter 11 is by far the most technically demanding part of the thesis. It takes up and proves an equivalence between constructible cosheaves and representations of MacPherson's entrance path category, which hinges on a proof of the Van Kampen theorem for this category. The full machinery of stratification theory is then used to construct representations of the (definable) entrance path category from a stratified (definable) map. This part also rests on proving a codimension-criterion under which Thom's condition a_f always holds.

Chapter 12 proves that Verdier duality is rightly conceived as an exchange of sheaves and cosheaves. An explicit formula for the derived equivalence of cellular sheaves and cosheaves is presented.

Chapter 13 uses the formula of Chapter 12 to prove that compactly supported cellular sheaf cohomology can be viewed as taking a (derived) coend with the image of the constant sheaf under this formula.

Chapter 14 proves that the derived category of cellular sheaves over a onedimensional base space is equivalent to a graded category. This formalizes the intuition of why spectral sequences over graphs always collapse on the E_2 page.

Chapter 15 introduces the interleaving distance for sheaves defined on a metric space. Although officially an extended pseudo-metric on the category of pre-sheaves, we prove it is an extended metric on the category of sheaves. One of the most fundamental properties of this extended metric is that global sections places sheaves into distinct connected components. To illustrate the theory more concretely, we take up an explicit description of the space of constructible sheaves over the real line.

11

THE DEFINABLE ENTRANCE PATH CATEGORY

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

— Virgil's Aeneid, Book 6, Lines 124-9

Fundamentally, one-dimensional persistent homology tries to understand topological changes in a one-parameter family of spaces. Multi-dimensional persistence tries to understand topological changes in a multi-parameter family of spaces; the leap in complexity from one dimension to two can not be overstated. The model problem of interest is to describe how the homology of the fiber of a map $f : Y \rightarrow X$ changes as one queries points or subsets in X. For general maps, this problem is entirely too unwieldy.

In this chapter we focus on a broad class of maps where this problem has an interesting answer: definably stratified maps. Informally, stratified maps are glued together fiber bundles. Definable maps are ones that can be defined with finitely many logical operations. Every definable map is stratified so we study simply definable maps.

The upshot of this chapter is that the homology of the fibers of a definable map give rise to a representation of a particular quiver with relations — a category in other words — called the definable entrance path category, whose general version was introduced by MacPherson to study general stratified maps. If one considers the opposite of the entrance path category, i.e. the exit path category, one obtains a constructible sheaf, which we now define.

Definition 11.0.10 (Constructible Sheaves and Cosheaves). Let F be a sheaf valued 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 $F|_{X^k}$ (the pullback of F along the inclusion $X^k \hookrightarrow X$) is locally constant. Alternatively, instead of asking for a filtration one can ask for a decomposition of X into disjoint pieces X_{σ} over which the restricted sheaf is locally constant.

Dually, we will call a cosheaf, with costalks valued in **vect**, constructible if its linear dual is a constructible sheaf.

One of the purposes of this chapter is 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 the introduction of stratification theory.

11.1 STRATIFICATION THEORY AND TAME TOPOLOGY

As wonderful as fiber bundles and local systems may be, they still fail to capture the sort of structure we are interested in because the topology of the fiber can never change. In order to bring cosheaves into contact into a larger realm of mathematics, we will need to consider stratified maps. To whet the appetite, stratified maps will allow us to describe in one language:

- 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 semialgebraic maps, which are stratified.

In other words, stratification theory gives a system of geometry for exploring a wealth of examples, appearing in mathematics and nature. Stratification theory does this by breaking up a space or map into regions, over which the usual analysis of manifolds and fiber bundles apply.

Definition 11.1.1 (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 11.1.2. 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 11.1.3. Suppose (X, P_X) and (Y, P_Y) are decomposed spaces, then a **decompositionpreserving** map is a continuous map $f : X \to Y$ that sends pieces to pieces, i.e. we have a commutative square



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 11.1.4 (Blow-Ups). Consider the map

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

This map is not triangulable, see [Shi] 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{RP}^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 S¹. 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{RP}^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 as in Figure 51 whose



FIGURE 51: Blowing up at a Point

boundary traces out a torus knot. See [AK10] for a treatment of different constructions of real blow-ups and their functorial properties.

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 52 can be decomposed into the two pieces

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



FIGURE 52: Topologist's Sine Curve

that satisfy the axiom of the frontier $X_{\sigma} \subset \bar{X}_{\tau}$, but it does not have the intuitively desired property that [Lu₇₆, p.1₃₁].

 $\dim X_{\sigma} < \dim X_{\tau}$

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 [Scho3, Rmk. 4.1.9]. 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.

11.1.1 Whitney 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.¹ In 1965, Whitney, whose approach relies on the geometry of tangent planes and secant lines, defined two properties that a stratified space should possess [Whi65a, Whi65b]. Thom, who proposed in a 1962 paper [Tho62] 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 [Tho69]. Thom's definition was first articulated carefully by John Mather in his famous 1970 Harvard "Notes on Topological

¹ For more historical context of these approaches, written by experts, we recommend the recent article [Gor12] and Part I, Section 1 of [GM88].



FIGURE 53: Diagram for Whitney Condition (b)

Stability" [Mat12], 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 proof² 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 11.2.17.

Definition 11.1.5 (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_σ ≤ X_τ 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 ∈ X_σ and the tangent spaces T_{y_i}X_τ 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 ℓ ⊆ T.

Remark 11.1.6. We have omitted **condition (a)** because it is implied by condition (b) [Mat12, 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.

² A proof appears in Mark Goresky's thesis [Gor76] that was never published and which he graciously provided to the author. We have since modified that proof to suit our purposes.

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:³

- **Dimension is Well-Behaved:** If $X_{\sigma} \subseteq fr(X_{\tau}) := \bar{X}_{\tau} X_{\tau}$, then dim $X_{\sigma} < \dim X_{\tau}$. See Proposition 2.7 of [Mat12] for a proof. This rules out the topologist's sine curve in Figure 52 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 ([Mat12] 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 $d_{\sigma} : T_{\sigma} \to \mathbb{R}_{\geq 0}$. 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 $\pi : S_{\sigma}(\epsilon) \to X_{\sigma}$ ([Gor78] 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 [Gor78].

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. This condition will be taken as primary when considering Thom-Mather stratifications.

11.1.2 Stratified Maps and a Counterexample

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 11.1.4 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 11.1.3.

³ Here we follow part of MacPherson's summary in the appendix of his 1991 Colloquium notes [Mac91].

Definition 11.1.7 (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 11.1.8. 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 11.1.9. 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 $X_{\sigma} f|_{X_{\sigma}}$ 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 [GM88, p. 41].

Lemma 11.1.10 (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 11.1.11. 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.

Counterexamples Creep In

One would like to say that given a general (not necessarily Thom) stratified map $f : X \rightarrow Y$, one could take a path $\gamma : [0, 1] \rightarrow Y$ so that the pullback $\gamma^* f : f^{-1}(\gamma) \rightarrow I$ 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.⁴

⁴ We are indebted to Mark Goresky for suggesting the key ideas of this example.



FIGURE 54: Preimage of the Spiral is Not Stratifiable

Example 11.1.12. 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"

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

is also Whitney stratified despite wrapping around the origin infinitely many times ([Pflo1] Example 1.4.8). However, the inverse image $\pi^{-1}(S)$ cannot be stratified because the inverse image of (0,0) is S¹, 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; see Figure 54. Since being Whitney stratified implies a drop in dimension of the frontier, contraposition shows that the inverse image cannot be Whitney stratified.

In Theorem 11.2.17 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 11.1.3.

11.1.3 *O-minimal Structures*

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 11.1.12 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 [Gro97]. 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 [Mac91]. Namely, a subset S is good if there is a Whitney stratification of M such that S is a 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" [vdDM96]. Taking requests from sheaf theorists [SV96] 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 $\forall, \exists, \lor, \land$ 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" [vdD98] as it is an excellent introduction that requires virtually no pre-requisites.

Definition 11.1.13 ([vdD98], p. 2). An o-minimal structure on \mathbb{R} is a sequence $\mathcal{O} = \{\mathcal{O}_n\}_{n \ge 0}$ 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 $i \leq j$ 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 O_1 are the finite unions of open intervals and points.

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

Remark 11.1.14. 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 **11.1.12** is not definable.

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

Lemma 11.1.15. 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 [vdDM96] to conclude that the restriction to the pullback of the projection maps to X and Y is definable as well. \Box

There are surprising facts that follow from the axioms of an o-minimal structure. For example, if $A \in O$, then the closure \overline{A} is in O ([vdD98] Ch. 1, 3.4). Another surprising fact is that definable sets can be Whitney stratified [Loi98]. 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 ([vdD98] Ch. 3).

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

Definition 11.1.16. 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 [Coso2]. 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 [Shi97]. This allows us to consider the following example of a semi-algebraic family of sets:

Example 11.1.17 (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

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


FIGURE 55: Point Cloud Data

By the previously stated facts we know that the sets

$$\mathbf{B}_z := \{ \mathbf{x} \in \mathbb{R}^{n+1} \mid \mathbf{f}_z(\mathbf{x}_1, \dots, \mathbf{x}_n) \leqslant \mathbf{x}_{n+1} \}$$

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

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

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

We conclude with the definition Miller and van den Dries proposed in section 1 of [vdDM96]. This definition allows us to verify definability locally, and allows us to work inside manifolds other than \mathbb{R}^n .

Definition 11.1.18 (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 11.1.19. 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 9-sets and 9-maps, although we will prefer to use the term "definable," has all the properties one could desire, including being closed under inverse images [vdDM96, D.7] (as long as the domain is closed) and Whitney stratifiability [vdDM96, D.16].⁵

11.1.4 Thom-Mather Stratifications

Definition 11.1.20 (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 $X_{\sigma} \leq X_{\tau}$, then

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

whenever both sides of the equations are defined.

Remark 11.1.21. In Figure 56 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.

⁵ The authors of [vdDM96] acknowledge that there is a gap in the proof of Whitney stratifiability of 9-maps, but Ta Lê Loi [Loi10] and others [NTT14] have since filled in this gap.



FIGURE 56: A System of Control Data

Definition 11.1.22 (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 $X_{\sigma}, \pi_{\sigma} : T_{\sigma} \to X_{\sigma}$ 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

$$(\pi_{\sigma,\tau}, \mathbf{d}_{\sigma,\tau}): \mathsf{T}_{\sigma,\tau} \to X_{\sigma} \times (0,\infty)$$

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

$$egin{array}{rll} \pi_{\sigma, au}\circ\pi_{ au,\lambda}&=&\pi_{\sigma,\lambda}\ d_{\sigma, au}\circ\pi_{ au,\lambda}&=&d_{\sigma,\lambda} \end{array}$$

whenever both sides of the equation are defined.

Remark 11.1.23. 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 [Gor78], among others. His proof views the lines of Figure 56 not as fibers of the retraction map π_{σ} , but rather⁶ as fibers of a radial projection map to the boundary of a tubular neighborhood.

Definition 11.1.24 (Family of Lines). A **family of lines** on a Thom-Mather stratified space is a system of radial projections

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

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

- 1. $r_{\sigma}(\varepsilon) \circ r_{\tau}(\varepsilon') = r_{\tau}(\varepsilon') \circ r_{\sigma}(\varepsilon) \in S_{\sigma}(\varepsilon) \cap S_{\tau}(\varepsilon')$
- 2. $d_{\sigma} \circ r_{\tau}(\varepsilon) = d_{\sigma}$
- 3. $d_{\tau} \circ r_{\sigma}(\varepsilon) = d_{\tau}$
- 4. $\pi_{\sigma} \circ r_{\tau}(\varepsilon) = \pi_{\sigma}$
- 5. If $0 < \varepsilon < \varepsilon' < \delta$, then $r_{\sigma}(\varepsilon') \circ r_{\sigma}(\varepsilon) = r_{\sigma}(\varepsilon')$

6.
$$\pi_{\sigma} \circ r_{\sigma}(\epsilon) = \pi_{\sigma}$$

7. $r_{\sigma}(\varepsilon)|_{T_{\sigma}(\varepsilon)\cap X_{\tau}}: T_{\sigma}(\varepsilon)\cap X_{\tau} \to S_{\sigma}(\varepsilon)\cap X_{\tau}$ is smooth

Remark 11.1.25. Every Thom-Mather stratified space admits a family of lines. This is the first proposition of [Gor₇₈].

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

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

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.

⁶ This description does not hold in higher dimensions.

Proposition 11.1.26 (Regular Neighborhoods of Closed Unions of Strata). Let X be a Thom-Mather stratified space and $W = \bigcup_{\sigma \in P_W} X_{\sigma}$ be a closed union of strata of X. The inclusion

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

is a homotopy equivalence.

Proof. Given a Thom-Mather stratified space, we can equip it with a family of lines [Gor78]. 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 $U_W(\epsilon) := \bigcup_{\sigma \in P_W} T_{\sigma}(\epsilon)$. 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 57 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:

$$f(x) = 0 \quad \text{if} \quad x \leq \frac{1}{2}$$
$$f(x) = 1 \quad \text{if} \quad x \geq \frac{3}{4}$$
$$f'(x) > 0 \quad \text{if} \quad x \in (\frac{1}{2}, \frac{3}{4})$$

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

$$H_{\sigma}(p,t) := \left\{ \begin{array}{ll} p & \text{if } p \notin T_{\sigma}(\varepsilon) \\ h_{\sigma}^{-1}(r_{\sigma}(\varepsilon)(p), d_{\sigma}(p)[(1-t)f(d_{\sigma}(p)/\varepsilon) + t]) & \text{if } p \in T_{\sigma}(\varepsilon) \end{array} \right.$$

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 11.1.24 imply that if $X_{\sigma} \leq X_{\tau}$ then

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

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



FIGURE 57: Two Sets of Tubular Neighborhoods

defines $H_i(p,t) = H_{\sigma}(p,t)$ if $p \in T_{\sigma}(\varepsilon)$. The commutation relation now extends to the statement that for any i and j

$$H_{i}(H_{i}(p,t),s) = H_{i}(H_{i}(p,s),t).$$

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 $r_1 = id$. It is easy to show that this implies that $W \hookrightarrow U_W(\epsilon/2)$ is a homotopy equivalence. \Box

Remark 11.1.27. 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 [Gor76]. This makes his homotopy

$$H_i^G(p,t) := \left\{ \begin{array}{ll} p & \text{if } p \notin T_\sigma(\varepsilon) \\ h_\sigma^{-1}(r_\sigma(\varepsilon)(p), d_\sigma(p)[(1-s_i(t))f(d_\sigma(p)/\varepsilon) + s_i(t)]) & \text{if } p \in T_\sigma(\varepsilon) \end{array} \right.$$

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

Remark 11.1.28. Of course, for a given stratum X_{σ} , away from its frontier the retraction map r_0 coincides with the tubular projection π_{σ} .

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 11.1.37 and Theorem 11.2.17. There is another technical tool that we need that can only be developed in the presence of control data.

Definition 11.1.29. 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 11.1.30. A stratified vector field η on X is said to be **controlled** by {T_{σ}, π_{σ} , d_{σ} } if the following compatibility conditions are satisfied for any pair of strata $X_{\sigma} \leq X_{\tau}$:

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

where ever both sides of the equation are defined.

11.1.5 Thom Mappings

Definition 11.1.31. A **Thom mapping** is a stratified map $f : (X, M) \to (Y, N)$ that satisfies **condition** a_f for every pair of strata $X_{\tau} \ge X_{\sigma}$: 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 $d(f|_{X_{\sigma}})_p \subseteq K$.

In Figure 58, we have drawn an example of a mapping that is not a Thom mapping.⁷ Other non-examples include the blow-up map discussed in Example 11.1.4. 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 [Shioo], but we have chosen not to rely on this conjecture. Instead, we only need the following proposition of Mather's (Proposition 11.3 of [Mat12]).

Proposition 11.1.32. 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 $X_{\sigma} \leq X_{\tau}$, 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.

⁷ This example is borrowed from [Lu76].



FIGURE 58: Not a Thom Mapping

(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 11.1.33. 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 11.1.34. 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 $X_{\sigma} \leq X_{\tau}$, 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

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

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 11.1.35. 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 $f : (X, M) \rightarrow (Y, \mathbb{R}^{k+1})$ is a Whitney stratified map where Y is the upper half plane in \mathbb{R}^{k+1} , i.e. $Y := \{(y_1, \cdots, y_{k+1}) | y_{k+1} \ge 0\}$. 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 $K_{\infty} \subseteq T_p M$ 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 $T_{\infty} \subseteq T_p M$. By Whitney's condition (a), $T_pX_{\sigma} \subset T_{\infty}$.

Denote by $\rho_Y(y) := \pi_{k+1}(y_1, \dots, y_{k+1}) = y_{k+1}$ 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

 $\nu \in T_i$ with $d\rho_X(x_i)(\nu) \neq 0$ must also have $df|_{X_\tau}(x_i)(\nu) \neq 0$ 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 $\nu \notin K_i$.

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 C¹ 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.

This lemma is instrumental for our proof of Theorem 11.2.17. On it's own, it has a useful corollary.

Corollary 11.1.36. Any stratified map $f : (X, M) \to (Y, \mathbb{R})$ that is C^1 on the ambient manifold is a Thom map.

11.1.6 Stratified Maps to the Real Line

Lemma 11.1.37. 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:

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

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 **11.1.26** 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 $U_{t_i}(\epsilon) = \bigcup_{\sigma \in X_{t_i}} T_{\sigma}(\epsilon/2)$ 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 $f^{-1}([s_i^-, t_i]) \subseteq U_{t_i}(\varepsilon)$. The composition of the inclusion followed by the retraction



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} .

$$H_{i}(X_{s_{i}^{-}};k) \rightarrow H_{i}(X_{t_{i}};k)$$

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 the persistence module introduced in Section 8.1.

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

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.

Corollary 11.1.38. The semialgebraic function $f : X \to \mathbb{R}$ in example 11.1.17 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 $n \ge 2$ 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 11.2, 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 11.1.37 contains the essential ideas of this more general picture. By considering certain definable paths in the parameter space, and analyzing their inverse



FIGURE 59: Two Entrance Paths and a Homotopy Between Them

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 11.2.17 associates a constructible cosheaf to a general definable map.

11.2 REPRESENTATIONS OF THE 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.

Definition 11.2.1. Let $(X, \{X_{\sigma}\}_{\sigma \in P_X})$ be a Whitney (or Thom-Mather) stratified space. We define the **entrance path category** $Entr(X, \{X_{\sigma}\})$ to be the category whose objects are points of X and whose morphisms are homotopy classes of entrance paths. An **entrance path** is a path $\gamma(t)$ whose ambient dimension (the pure dimension of the containing stratum) is non-increasing with t. Moreover we require the homotopies h(s, t) to be entrance paths for every fixed s. We write Entr(X) when a given stratification is understood.

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

Remark 11.2.2 ("Tame" Homotopies). David Treumann's thesis [Treo9], which was written under MacPherson's direction, contains one of the first published accounts of the exit path category. However, he added an additional hypothesis that the homotopies should

be "tame," which he defines by saying that $h : [0, 1]^2 \rightarrow X$ 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 [Wooo9] 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 11.2.3 (Definable Entrance Path Category). For a fixed analytic-geometric category 9 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 $Entr_9(X, \{X_{\sigma}\})$. Dually, we have a definable exit path category $Exit_9(X)$.

Remark 11.2.4. We will not need to use the definable entrance path category until Theorem **11.2.17**, 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 subcategories by choosing a single point from each connected component in the stratification and passing to a skeletal subcategory.

Example 11.2.5 (Entrance Path Category for S¹). Now consider the circle S¹ stratified as a single pure stratum. The argument above shows that we can view the entrance path category of S¹ 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. Hom(\star, \star) $\cong \mathbb{Z}$.

Example 11.2.6 (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 11.2.7** (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

 $\operatorname{Entr}(X) \cong \operatorname{Cell}(X)^{\operatorname{op}} = \operatorname{P}_X^{\operatorname{op}}$ and $\operatorname{Exit}(X) \cong \operatorname{Cell}(X) = \operatorname{P}_X$.

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.

11.2.1 Homotopy Links

Definition 11.2.8 (Homotopy Link). Suppose X is a decomposed space and $X_{\sigma} \leq X_{\tau}$ 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 ([Woo09], Lemma 3.2) to our situation.

Lemma 11.2.9. 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 $i \leq j$, we have that $X^i \leq X^j$.

We claim that one can show that every entrance path γ starting in a stratum X^k and ending in a stratum Xⁱ 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 $x_j := \gamma(t_n - \varepsilon) \in X^j$ 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 $x_j = \gamma(t_n - \varepsilon)$ 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



FIGURE 60: Modifying a Homotopy to Stay Inside an Open Set

 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 Xⁱ 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 [Milo6]. At a high level, this provides a proof of Proposition 11.2.7, which can also be seen using easier methods.

Proof of **11.2.7**. 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

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

which was wanted.

11.2.2 Van Kampen Theorem for Entrance Paths

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 [Treo9] proves the van Kampen theorem for the exit path 2-category, but the elegant inductive argument in proposition 5.9 appears to have an error.⁸ Jacob Lurie has a proof for the ∞ -category case [Luro9b]. Jon Woolf has outlined another argument [Woo14] based on his classification of **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 60. 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 I² such that $h : I^2 \rightarrow X$ sends open cells of T to strata of X. We then, after sufficient refinement, define a homeomorphism of I² 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 I². Let us state our desired version of the van Kampen theorem and give the first step of the proof.

Theorem 11.2.10 (Van Kampen Theorem for Entrance Paths). If X is a Whitney stratified space and $\mathcal{U} = \{U_i\}$ is a cover, then

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

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.

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

⁸ The argument inducts on the number of triangles in a triangulation of I². The statement that the closure of the complement of a single triangle is homeomorphic to I² is not true if, for example, the triangle has two vertices on one side of the square and the third on another side.

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 11.2.13 together with Proposition 11.2.14.

Definition 11.2.11. 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 60 for an illustrative cartoon.

Remark 11.2.12. 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 11.2.13. 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 I² 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 $X_f \ge X_e \ge X_v$. 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 I². 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 $\gamma_i = w_a w_b$ labeled 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 I² 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 $\gamma_i = \nu f$ or $f\nu$, then replace γ_i with $\gamma'_i := \nu e f$ or $fe\nu$, 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 $s(f) \leq s(v) \leq s(e)$ or if $\gamma_i = ef$ and $s(e) \leq s(v) \leq s(f)$ 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.



FIGURE 61: Forcing Move (1c) to Apply

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.

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 $s(f) \leq s(v) \leq s(e)$ 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 I², 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 $s(f') \leq s(e_{34})$, which would imply that $2s_2 \leq s_3 + s_4$, but this would

imply that $s_2 = s(v_2) \leq s(f')$ 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 61. 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 11.2.14. Suppose that a (definable) triangulation $\varphi : |K| \rightarrow I^2$ 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 I² 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 $l_i(s)$ 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 $(s, l_i(s))$. 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.

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

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.

$$g_{\mathfrak{i}}(s,\mathfrak{t}) := (s,2 \cdot |\varphi_{\mathfrak{i}}(s) - \ell_{\mathfrak{i}}(s)| \cdot \mu_{\pm}(\frac{\mathfrak{t} - \ell_{\mathfrak{i}}(s)}{2|\varphi_{\mathfrak{i}}(s) - \ell_{\mathfrak{i}}(s)|}) + \ell_{\mathfrak{i}}(s))$$

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.

11.2.3 The Equivalence

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

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

defines a constructible cosheaf.

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

$$\widehat{F}(U) := \varinjlim_{\mathsf{Entr}(U)} \widehat{F}|_{U}$$

This is clearly a pre-cosheaf since if $U \hookrightarrow V$, the colimit over $\widehat{F}|_V$ defines by restriction a cocone over $\widehat{F}|_U$ and thus a unique map $\widehat{F}(U) \to \widehat{F}(V)$. 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 $Entr(U_x)$, thus the colimit returns the value of $\widehat{F}(x)$. 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 11.2.10:

$$\varinjlim_{I \in \mathsf{N}(\mathfrak{U})} \mathsf{Entr}(U_I) \cong \mathsf{Entr}(X)$$

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

$$\underbrace{\lim_{I \in \mathcal{N}(\mathcal{U})} \widehat{F}(\mathcal{U}_{I}) := \lim_{I \in \mathcal{N}(\mathcal{U})} \underbrace{\lim_{I \in \mathcal{N}(\mathcal{U})} \widehat{F}|_{\mathcal{U}_{I}} \widehat{F}|_{\mathcal{U}_{I}} \cong \underbrace{\lim_{I \in \mathcal{N}(\mathcal{U})} \underbrace{\lim_{I \in \mathcal{N}(\mathcal{U})} \widehat{F}|_{\mathcal{U}_{I}} \cong \underbrace{\lim_{I \in \mathcal{N}(\mathcal{U})} \widehat{F}|_{\mathcal{U}_{I}}}_{\mathsf{Entr}(\mathcal{U})} = \underbrace{\lim_{I \in \mathcal{N}(\mathcal{U})} \widehat{F}|_{\mathcal{U}_{I}} = \underbrace{\lim_{I \in \mathcal{N}(\mathcal{U})} \widehat{F}|_{\mathcal{U}_{I}}}_{\mathsf{Entr}(\mathcal{U})} = \underbrace{\lim_{I \in \mathcal{N}(\mathcal{U})} \widehat{F}|_{\mathcal{U}_{I}} = \underbrace{\lim_{I \in \mathcal{N}(\mathcal{U})} \widehat{F}|_{\mathcal{U}_{I}}}_{\mathsf{Entr}(\mathcal{U})} = \underbrace{\lim_{I \in \mathcal{N}(\mathcal{U})} \widehat{F}|_{\mathcal{U}_{I}} = \underbrace{\lim_{I \in \mathcal{N}(\mathcal{U})} \widehat{F}|_{\mathcal{U}_{I}}}_{\mathsf{Entr}(\mathcal{U})} = \underbrace{\lim_{I \in \mathcal{N}(\mathcal{U})} \widehat{F}|_{\mathcal{U}}}_{\mathsf{Entr}(\mathcal{U})} = \underbrace{\lim_{I \in \mathcal{N}(\mathcal{U})} \widehat{F}|_{\mathcal{U}}}_{\mathsf{Entr}(\mathcal{U})}_{\mathsf{Entr}(\mathcal{U})} = \underbrace{\lim_{I \in \mathcal{N}(\mathcal{U})} \widehat{F}|_{\mathcal{U}}}_{\mathsf{Entr}(\mathcal{U})}_{\mathsf{Entr}(\mathcal{U})}_{\mathsf{Entr}(\mathcal{U})}_{\mathsf{Entr}(\mathcal{U})}_{\mathsf{Entr}(\mathcal{U})}_{\mathsf{Entr}(\mathcal{U})}_{\mathsf{Entr}(\mathcal{U})}_{\mathsf{Entr}(\mathcal{U})}_{\mathsf{Entr}(\mathcal{U})}_{\mathsf{Entr}(\mathcal{U})}_{\mathsf{Entr}(\mathcal{U})}_{\mathsf{Entr}(\mathcal{U})}_{\mathsf{Entr}(\mathcal{U})}_{\mathsf{Entr}(\mathcal{U})}_{\mathsf{Entr}(\mathcal{U})}_{\mathsf{Entr}(\mathcal{U})}_{\mathsf{Entr}(\mathcal{U})}_{\mathsf{Entr}(\mathcalU)}$$

This establishes the cosheaf axiom.

Theorem 11.2.16 (Representations of the Entrance Path Category). Every cosheaf \widehat{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.

$$Entr(X, \{X_{\alpha}\}) \rightarrow vect$$

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 \widehat{F} is a constructible cosheaf, which we assume has finite-dimensional costalks. We claim that for U_x suitably small we can show that

$$\widehat{F}_{\mathbf{x}} \cong \widehat{F}(\mathbf{U}_{\mathbf{x}}).$$

This is not so easy to see and a proof would require substantial more development of cosheaf theory. Heuristically, if the value of \widehat{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 $\operatorname{Hom}_{vect}(-, k)$, which is an equivalence, and we can apply the proof for constructible sheaves found on p. 84 of [GM83].

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:

$$\widehat{F}_{x} \xrightarrow{\cong} \widehat{F}(U_{x}) \longleftrightarrow \widehat{F}(U_{y}) \xleftarrow{\cong} \widehat{F}_{y}$$

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 3.2.6 to define a map from $\widehat{F}_{y'}$ to \widehat{F}_{y} . Postcomposing with the above map defines the map $\widehat{F}_{y'} \rightarrow \widehat{F}_x$. 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 11.2.10 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 3.2.6 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 $\beta(t)$.



11.2.4 Representations from Stratified Maps

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 11.2.17 (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 C¹ 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 11.1.15, 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 11.1.37 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 [vdDM96]), there exists a finite partition of [0, 1] such that over each interval the inverse image is homeomorphic to the product:



By properness we can, for any fixed $\varepsilon > 0$, find an s_i^+ such that $f^{-1}([t_i, s_i^+])$ is contained in $U_i(\varepsilon) := \cup T_{\sigma}(\varepsilon/2)$ for $Y_{\sigma} \subseteq f^{-1}(t_i) := Y_i$. The retraction we constructed in Proposition 11.1.26 gives a retraction map $r_i^+ := H(p, 0)$ from $U_i(\varepsilon) \to Y_i$. This allows us to define a map on fibers

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

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.

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

Similarly, one can find a point s_{i+1}^- so that it's fiber is contained in $U_{i+1}(\varepsilon)$ 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 [Mat12]. 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 **11.2.10**, it suffices to check homotopy invariance on an elementary homotopy, such as the one depicted in Figure 59. 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



FIGURE 62: Argument for Homotopy Invariance

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 $\{t = 1\} \cup \{(1,t) | t \ge 1/2\}$ to the stratum X_{λ} . This guarantees that the fibers over x, x', y, y', a, a', b, b', c' and d' in Figure 62 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 $I^2 \times_X Y \subset \mathbb{R}^2 \times \mathbb{N}$. 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 $\mathbb{R} \times \mathbb{N}$. The manner in which Mather constructs control data in Proposition 7.1 of [Mat12] 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 $\mathbb{R}^2 \times \mathbb{N}$ 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

$$(r_{y})_{*} \circ (\phi_{y,x})_{*} \circ (r_{x})_{*}^{-1} : H_{n}(Y_{(0,0)}) \to H_{n}(Y_{(0,1)})$$

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

For $\beta(t)$, the action is similar:

$$(\mathbf{r}_{d})_{*} \circ (\varphi_{d,c})_{*} \circ (\mathbf{r}_{c})_{*}^{-1} \circ (\mathbf{r}_{b})_{*} \circ (\varphi_{b,a})_{*} \circ (\mathbf{r}_{a})_{*}^{-1} : H_{n}(Y_{(1,0)}) \to H_{n}(Y_{(1,1/2)}) \to H_{n$$

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 α', b', c' and d' in Figure 62. The representation associated to $\gamma(t)$ is

$$(\mathbf{r}_1)_* \circ (\varphi_{d',c'})_* \circ (\mathfrak{i}_{c'})_*^{-1} \circ (\mathfrak{i}_{b'})_* \circ (\varphi_{b',a'})_* \circ (\mathbf{r}_0)_*^{-1}.$$

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_{α} can be taken to be the restriction of a retraction map $r_0 : U_{Y_{(1,0)}}(\varepsilon) \rightarrow Y_{(1,0)}$ constructed in Proposition 11.1.26. This in turn can be taken to be the restriction of the tubular projections used to define a retraction map $r_{s=1} : U_{s=1}(\varepsilon) \rightarrow Y_{s=1}$. The commutation relations for control data allow us to imagine first taking the fiber $Y_{\alpha'}$ over α' 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

$$\mathbf{r}_0 = \mathbf{r}_a \circ \mathbf{r}_{e_0},$$

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 11.1.35 and property (b) of Proposition 11.1.32.

By Lemma 11.1.35, 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 $e_0 := \{(1,t) \mid 0 < t < 1/2\}$, there exists a system of control data $\{T'\}$ over T_{e_0} and the interior of I² by Proposition 11.1.32. If we restrict to those tubular neighborhoods coming from strata in Y_{e_0} , then property (a) of Proposition 11.1.32 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,⁹ there is a homeomorphism ψ_{e_0} of Y_{e_0} that takes $\{T'\}_{Y_{e_0}}$ to $\{T\}_{Y_{e_0}}$. This implies that if Y_{σ} is a stratum that is mapped to (1, 0) and Y_{τ} is mapped to e_0 , then

$$\pi_{\sigma} = \pi_{\sigma} \circ \psi_{e_0} \circ \pi_{\tau'}$$
 since $\pi_{\tau} = \psi_{e_0} \circ \pi_{\sigma'}$

By repeating the construction of a retraction outlined in Proposition 11.1.26, 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:

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 I² minus the

⁹ Mather mentions at the bottom of page 492 of [Mat12] 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 Corollary 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 Corollary 10.3 to the set where just $Z \times \mathbb{R}$ is given the extension of just the control data {T}.

vertices using the control data {T'} over {T_f, T_{e0}, T_{e1}}. 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

$$\varphi_{\mathfrak{b},\mathfrak{a}}\circ \mathfrak{r}_{\mathfrak{e}_0'}=\mathfrak{r}_{\mathfrak{e}_0'}\circ \varphi_{\mathfrak{b}',\mathfrak{a}'}$$

Again, the commutation relations in Proposition 11.1.26 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 $r_{e_0} = \psi_{e_0} \circ r'_{e_0}$ 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.

Remark 11.2.18 (Alternative Idea for a Proof). An alternative approach makes use of the properties of o-minimal structures. The generic triviality theorem 4.11 of [vdDM96] guarantees that we have a definable trivialization of the map over (0, 1).



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 $\alpha_x : (0,1) \rightarrow f^{-1}(\gamma([0,1]))$.

Mário Edmundo and Luca Prelli, in their recent note [EP12] 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 [PS99] that shows that being definably compact (equivalently, closed and bounded) is equivalent to being able to to complete curves. A map $f : Y \rightarrow X$ is definably proper if for every definable curve $\alpha : (0, 1) \rightarrow Y$ and every definable map $[0, 1] \rightarrow X$ there is at least one way to complete the diagram:



If one assumes all the maps are continuous as well as definable then the completion in the diagram above is unique.¹⁰

¹⁰ 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.

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 is showing that this map g is continuous and is invariant under homotopy.

12

DUALITY: EXCHANGE OF SHEAVES AND COSHEAVES

"Quod superius sicut quod inferius, et quod inferius sicut quod superius, ad perpetranda miracula rei unius."

— Hermes Trismegistus [VW88, p. xxix]

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 12.1.3 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 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



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 $\tau \leq \sigma$ 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:



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:



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 $\widehat{\mathcal{P}}(F)$

$$\widehat{\mathcal{P}}(F): \quad \sigma \quad \rightsquigarrow \quad F(\sigma) \to \bigoplus_{\sigma \leqslant_1 \tau} F(\tau) \to \bigoplus_{\sigma \leqslant_2 \gamma} F(\gamma) \to \cdots.$$

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 $\sigma \leq_2 \gamma$ 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 $d^{i+1} := \oplus [\tau: \gamma] \rho_{\gamma,\tau}^F$. Now let's consider a cell λ that is a codimension one face of σ , then the extension map $r_{\lambda,\sigma}^{\bullet}$ is defined to be the chain map

$$\begin{array}{c|c} 0 & \longrightarrow & F(\sigma) & \stackrel{d^{i}}{\longrightarrow} \bigoplus_{\sigma \leqslant_{1}\tau} F(\tau) & \stackrel{d^{i+1}}{\longrightarrow} \bigoplus_{\sigma \leqslant_{2}\gamma} F(\gamma) \\ r^{i-1}_{\lambda,\tau} & & r^{i}_{\lambda,\tau} & r^{i+1}_{\lambda,\tau} & r^{i+2}_{\lambda,\tau} \\ F(\lambda) & \xrightarrow[d^{i}-]{\longrightarrow} h_{\lambda \leqslant_{1}\sigma} F(\sigma) & \xrightarrow[d^{i}]{\longrightarrow} h_{\lambda \leqslant_{2}\tau} F(\tau) & \stackrel{d^{i+1}}{\longrightarrow} h_{\lambda \leqslant_{3}\gamma} F(\gamma) \end{array}$$

The reason it is a chain map is clear from the fact that if $\lambda \leq \sigma$ 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 12.1.1 (Poincaré-Verdier Equivalence Functor). Let X be a cell complex and let Shv(X) and CoShv(X) denote the categories of cellular sheaves and cosheaves respectively. We define the **Poincaré-Verdier Equivalence Functor** $\widehat{\mathcal{P}}$: $D^b(Shv(X)) \rightarrow D^b(CoShv(X))$ by the following formula: to a sheaf $F \in Shv(X)$ 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

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 12.1.2. If F is a cell sheaf on a cell complex X, then

$$H_{c}^{i}(X;F) \cong H_{-i}(X;\widehat{\mathcal{P}}(F)).$$

Proof. First we apply the equivalence functor $\widehat{\mathcal{P}}$ to F

$$0 \longrightarrow \bigoplus_{\nu \in X} [\hat{\nu}]^{F(\nu)} \xrightarrow{[\nu:e]\rho e, \nu} \bigoplus_{e \in X} [\hat{e}]^{F(e)} \xrightarrow{[e:\sigma]\rho_{\sigma,e}} \bigoplus_{\sigma \in X} [\hat{\sigma}]^{F(\sigma)} \longrightarrow \cdots$$

Taking colimits (pushing forward to a point) term by term produces the complex of vector spaces

$$0 \longrightarrow \bigoplus_{\nu \in X} F(\nu)_{[\nu:e]\rho e,\nu} \bigoplus_{e \in X} F(e)_{[e:\sigma]\rho_{\sigma,e}} \bigoplus_{\sigma \in X} F(\sigma) \longrightarrow \cdots$$

which the reader should recognize as being the computational formula for computing compactly supported sheaf cohomology. $\hfill \Box$

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 12.1.3. 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

$$H^{\iota}(X; F) \cong H_{n-\iota}(X; F).$$

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.

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

$$\mathsf{H}_{-\mathfrak{i}}(\mathsf{X};\widehat{\mathcal{P}}(\mathsf{F}))\cong\mathsf{H}^{\mathsf{B}\mathsf{M}}_{\mathfrak{n}-\mathfrak{i}}(\widetilde{\mathsf{X}};\widehat{\mathsf{F}}).$$

However, we already observed in Theorem 7.4.3 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 F. Thus using compactness to drop the Borel-Moore label and chaining together the previous theorem we get

$$\mathsf{H}^{\mathfrak{i}}(X; \mathsf{F}) \cong \mathsf{H}_{-\mathfrak{i}}(X; \widehat{\mathcal{P}}(\mathsf{F})) \cong \mathsf{H}_{\mathfrak{n}-\mathfrak{i}}(\widetilde{X}; \widehat{\mathsf{F}}) \cong \mathsf{H}_{\mathfrak{n}-\mathfrak{i}}(X; \mathsf{F}).$$

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" [Sch98], which is a follow-up of a longer paper connecting sheaves, buildings and representation theory [Sch97]. 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 [Vyb99], but it took up until 2005 for Kohji Yanagawa to explicitly state that Vybornov's work implied the derived equivalence of sheaves and cosheaves [Yano5].

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 12.2.1 (Equivalence). $\widehat{\mathcal{P}}$: $D^b(Shv(X)) \to D^b(CoShv(X))$ 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 ([Wei94] 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

$$\widehat{\mathcal{P}}: [\sigma]^V \longrightarrow \bigoplus_{\tau^0 \subset \sigma} [\widehat{\tau}^0]^V \longrightarrow \cdots \longrightarrow \bigoplus_{\tau^i \subset \sigma} [\widehat{\tau}^i]^V \longrightarrow \cdots \longrightarrow [\widehat{\sigma}]^V$$

which is nothing other than the projective cosheaf resolution of the skyscraper (or stalk) cosheaf \hat{S}_{σ}^{V} supported on σ , i.e.

$$\hat{S}_{\sigma}^{V}(\tau) = \begin{cases} V & \sigma = \tau \\ 0 & o.w. \end{cases}$$

Consequently, there is a quasi-isomorphism $q : \widehat{\mathcal{P}}([\sigma]^V) \to \widehat{S}^V_{\sigma}[-\dim \sigma]$ where \widehat{S}^V_{σ} is placed in degree equal to the dimension of σ assuming that $[\sigma]^V$ is initially in degree o. By abusing notation slightly and letting \mathcal{P} send cosheaves to sheaves, we see that

$$\mathfrak{P}(\mathbf{q}):\mathfrak{P}\widehat{\mathfrak{P}}([\sigma]^{\mathsf{V}})\to\mathfrak{P}(\mathsf{S}_{\sigma}^{\mathsf{V}})=[\sigma]^{\mathsf{V}}$$

and thus we can define a natural transformation from $\widehat{\mathcal{PP}}$ to $id_{D^b(Shv)}$ when restricted to elementary injectives. However, by Lemma 7.1.6 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 $\widehat{\mathcal{PP}} \cong id_{D^b(Shv)}$. Repeating the argument starting from co-sheaves shows that

$$\widehat{\mathcal{P}}: \mathsf{D}^{\mathfrak{b}}(\mathsf{Shv}(\mathsf{X})) \leftrightarrow \mathsf{D}^{\mathfrak{b}}(\mathsf{CoShv}(\mathsf{X})): \mathcal{P}$$

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 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 \rightarrow 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 12.2.2 (Linear Duals). Define $\hat{V} : D^b(Shv(X))^{op} \to D^b(CoShv(X))$ as follows

- $\hat{V}(F^{\bullet}) = (F^*)^{-\bullet}$, i.e. take a sheaf in slot i, dualize its internal restriction maps $\rho_{\sigma,\tau}^{F^i}$ to extension maps $r_{\tau,\sigma}^{F^{i*}}$ 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 $\partial^{-i-1}:=(d^i)^*$

$$\star(\cdots \longrightarrow F^{i} \xrightarrow{d^{i}} F^{i+1} \longrightarrow \cdots) = \cdots \longrightarrow [(F^{i+1})^{*}]^{-i-1} \xrightarrow{a^{-i-1}} [(F^{i})^{*}]^{-i} \longrightarrow \cdots$$

We'll adopt the convention that lowering the index increases the degree $\partial^{-i-1} \rightarrow \partial_i$.

We will reserve the right to abuse notation and let V map from co-sheaves to sheaves in the obvious manner, i.e. $V : D^b(CoShv(X))^{op} \to D^b(Shv(X))$ or formally equivalent $V : D^b(CoShv(X)) \to D^b(Shv(X))^{op}$.

Lemma 12.2.3. $\hat{V}_f : D^b(Shv_f(X)) \to D^b(CoShv_f(X))^{op}$ 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 $K^{b}(Inj - S_{f})^{op}$ to $K^{b}(Proj - C_{f})$ and composed twice $V\hat{V} : K^{b}(Inj - S_{f}) \rightarrow K^{b}(Inj - S_{f})$ 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.

12.2.2 Verdier Dual Anti-Involution

Definition 12.2.4 (Verdier Dual). The Verdier dual functor $D : D(Shv_f(X)) \rightarrow D(Shv_f(X))^{op}$ 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 $Hom(F|_{st(\sigma)}, G|_{st(\sigma)})$, i.e. natural transformations between the restrictions to the star of σ .

The complex of injective sheaves ω_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 $D^{\bullet}F := \mathcal{H}om(F, \omega_{\chi}^{\bullet})$. Written out explicitly it is

$$\cdots \longrightarrow \oplus_{|\tau|=i+1} [\tau]^{F(\tau)^*} \underset{\oplus}{\longrightarrow} \oplus_{|\gamma|=i} [\gamma]^{F(\gamma)^*} \underset{\oplus}{\longrightarrow} \oplus_{|\sigma|=i-1} [\sigma]^{F(\sigma)^*} \longrightarrow \cdots$$

Proposition 12.2.5. The functor $\widehat{\mathcal{P}}$: $D^{b}(Shv_{f}(X)) \to D^{b}(CoShv_{f}(X))$ composed with linear duality $V : D^{b}(CoShv_{f}(X)) \to D^{b}(Shv_{f}(X))^{op}$ gives the Verdier dual anti-equivalence, i.e. $D \cong V\widehat{\mathcal{P}}$.

Proof. Just check by hand.

Remark 12.2.6. 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

"Speech is the twin of my vision....it is unequal to measure itself."

- Walt Whitman's Song of Myself [25]

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

$$\operatorname{Hom}_{\mathsf{R}}(\star,\star)\cong\mathsf{R}$$

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 Ab — the category of abelian groups.

An additive functor $B : R \to Ab$ 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 $A : R^{op} \to Ab$ 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

$$A \otimes B : \mathbb{R}^{op} \times \mathbb{R} \to Ab$$
 $(\star, \star) \mapsto A \otimes_{\mathbb{Z}} B.$

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.



In other words there is a coequalizer

$$A \otimes_{\mathbb{Z}} R \otimes_{\mathbb{Z}} B \xrightarrow{(a,r,b) \mapsto (a,rb)} A \otimes_{\mathbb{Z}} B \xrightarrow{(a,r,b) \mapsto (ar,b)} A \otimes_{\mathbb{Z}} B \xrightarrow{} A \otimes_{\mathbb{R}} B$$

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 13.1.1 (Tensoring Sheaves with Cosheaves). Let X be a topological space and let \widehat{G} and F be a pre-cosheaf and a pre-sheaf respectively, both valued in Vect. Note that for every pair of objects $U \rightarrow V$ in **Open**(X) we have a diagram



which is the building block in defining the **coend** or **tensor product over** X

$$\bigoplus_{U \to V} \widehat{G}(U) \otimes F(V) \rightrightarrows \bigoplus_{W} \widehat{G}(W) \otimes F(W) \to \int^{Open(X)} \widehat{G}(W) \otimes F(W) \eqqcolon \widehat{G} \otimes_X F.$$

We illustrate this definition with an immediate example.

Example 13.1.2 (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 G and let F vary then we get a functor

$$\widehat{\mathsf{G}} \otimes_{\mathsf{X}} - : \mathbf{Shv}(\mathsf{X}) \to \mathbf{Vect}$$

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 CoShv(X) and Shv(X), i.e.

$$-\otimes_X -: \mathbf{CoShv}(X) \times \mathbf{Shv}(X) \to \mathbf{Vect} \qquad (\widehat{\mathsf{G}},\mathsf{F}) \mapsto \widehat{\mathsf{G}} \otimes_X \mathsf{F} := \int^{\mathbf{Open}(X)} \widehat{\mathsf{G}}(U) \otimes \mathsf{F}(U),$$

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 [Sch87], 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 [Pit85].

Theorem 13.2.1. Let X be any topological space. Every colimit-preserving functor on sheaves arises by tensoring with a cosheaf, i.e.

$$CoShv(X; Set) \cong Fun^{co-cts}(Shv(X; Set), Set).$$

This theorem is also stated in Marta Bunge and Jonathan Funk's 2006 book "Singular Coverings of Toposes" [BFo6] 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,¹ 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 $p_! : Shv(X) \to Shv(\star) \cong$ 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 Rp₁ admits a right adjoint on the level of the derived category. This adjunction is sometimes called **global Verdier duality**:

$$\operatorname{Hom}(\operatorname{Rp}_{!}\mathsf{F},\mathsf{G})\cong\operatorname{Hom}(\mathsf{F},\mathsf{p}^{!}\mathsf{G}).$$

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.

¹ We delay this for another paper.

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}_{σ} when using the Alexandrov topology

$$[\hat{\sigma}](au) = \left\{ egin{array}{cc} k & ext{if } au \leqslant \sigma \ 0 & ext{o.w.} \end{array}
ight.$$

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 13.2.2. Let X be a cell complex, then the operation $Rp_! : Shv(X) \rightarrow Vect$ on cellular sheaves is equivalent to tensoring with the image of the constant sheaf through the derived equivalence, i.e.

$$\widehat{\mathcal{P}}(k_X) = \bigoplus_{\nu \in X} [\hat{v}] \to \bigoplus_{e \in X} [\hat{e}] \to \bigoplus_{\sigma \in X} [\hat{\sigma}] \to \cdots.$$

Proof. The proof is immediate given the previous description of taking stalks, i.e. one can check directly the formula

$$\widehat{\mathcal{P}}(k_X) \otimes_X F \cong \bigoplus_{\nu \in X} F(\nu) \to \bigoplus_{e \in X} F(e) \to \bigoplus_{\sigma \in X} F(\sigma) \to \cdots$$

whose cohomology is by definition the compactly supported cohomology of a cellular sheaf F. $\hfill \Box$

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{k}_X should act on sheaves by returning its colimit, i.e. zeroth sheaf homology

$$\hat{k}_X \otimes_X - : \mathbf{Shv}(X) \to \mathbf{Vect} \qquad F \rightsquigarrow H_0(X; F) = p_{\dagger}F.$$

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.

14

GRADED DESCRIPTIONS OF THE DERIVED CATEGORY

"Mathematics is the art of giving the same names to different things."

— Henri Poincaré [Poio8]

In Chapter 7 we introduced the derived category of cellular sheaves. The fundamental objects there are chain complexes parametrized by a cell complex. However, the derived category takes a further step by identifying objects that are "essentially the same" when viewed through the lens of cohomology sheaves.

In Section 14.1 we understand this principle better by demonstrating the well known fact that the derived category of chain complexes of vector spaces (sheaves over a point) is equivalent to the graded category of vector spaces. Our proof follows the standard proof in [Wei94] except we use the barcode method for chain complexes introduced in Example 6.3.7 to visualize explicitly what is happening. Roughly speaking, the derived category of chain complexes allows us to remove the green bars in Figure 63 as they "graph" the chain homotopy between the identity and the map that projects onto and then includes the red dots.

This sets us up for Section 14.2, which culminates in a proof that the derived category of cellular sheaves over a one-dimensional base space is equivalent to a graded category of sheaves. This should seem plausible because over each cell, a chain complex is equivalent to a graded vector space. Indeed, one could repeat the proof of Lemma 14.1.1 verbatim if it weren't for the pesky fact that projecting onto the cohomology cell-by-cell fails to define a sheaf map. However, the proof of the equivalence follows by constructing an explicit replacement of every object in the derived category with an object where such a naïve map will exist.



FIGURE 63: A Chain Complex for the 2-Sphere as Barcodes

14.1 THE DERIVED CATEGORY FOR COMPLEXES OF VECTOR SPACES

Here we warm-up with an alternative approach to the derived category of complexes of vector spaces. Recall that a chain complex simply consists of a collection of vector spaces and maps satisfying $d^2 = 0$.

$$\cdots \to V^{i-1} \to V^i \to V^{i+1} \to \cdots$$

Cohomology defines a functor from chain complexes to graded vector spaces simply by placing the ith cohomology in degree i.

$$H^*: C^{\mathfrak{b}}(\operatorname{Vect}) \to \operatorname{qrVect} \qquad (V^{\bullet}, d_{\mathcal{V}}) \rightsquigarrow \{H^*(V^{\bullet}, d_{\mathcal{V}})\}$$

A map of chain complexes $f^{\bullet} : V^{\bullet} \to W^{\bullet}$ is a quasi-isomorphism if the maps $H^{i}(f) : H^{i}(V^{\bullet}) \to H^{i}(W^{\bullet})$ is an isomorphism for every non-negative integer i.

The derived category of chain complexes is defined to be the category of chain complexes localized at the collection of quasi-isomorphisms. This is often simplified by saying that

In the derived category quasi-isomorphisms are formally inverted.

To illustrate this slogan we will prove that every chain complex is quasi-isomorphic to a graded vector space. This is the simplest instance of a more general theorem that we prove in this chapter.

Lemma 14.1.1. A chain complex (V, d_V) is quasi-isomorphic to its cohomology $(H^{\bullet}(V, d_V), 0)$, viewed as complex with zero differentials.

Proof. It suffices to define a chain map

$$\pi^{\bullet}: (V, d_V) \to (\mathsf{H}^{\bullet}(V, d_V), 0) \qquad \text{or} \qquad \iota^{\bullet}: (\mathsf{H}^{\bullet}(V, d_V), 0) \to (V, d_V)$$

that induces the obvious isomorphism $H^{\iota}(\pi) : H^{\iota}(V) \to H^{\iota}(V)$. We will use the decomposition for persistence modules from Theorem 6.3.3 to define this map.

In Example 6.3.7 we observed that any chain complex (V, d_V) is isomorphic to a direct sum

$$V \cong \bigoplus_{i \in \mathbb{Z}} S_i^{n_i} \oplus P_i^{m_i}$$

where

$$S_i: \cdots \to 0 \to k \to 0 \to \cdots$$

is a length zero interval module and

$$P_i: \cdots \to 0 \to k \to k \to 0 \to \cdots$$

is a length one interval module, with the first non-zero term in degree i. If we group terms so that

$$B^{i-1} := P_i^{\mathfrak{m}_{i-1}} \qquad H^i(V) := S_i^{\mathfrak{n}_i} \qquad B^i := P_i^{\mathfrak{m}_i}$$

then we get an obvious map, which precomposed with the above isomorphism is our desired quasi-isomorphism.



One can in fact show more. Let $\pi : B^i \oplus H^i(V) \oplus B^{i-1} \to H^i(V)$ be the obvious projection and $\iota : H^i(V) \to B^{i-1} \oplus H^i(V) \oplus B^i$ be the obvious inclusion, then one can construct an explicit chain homotopy s joining $\iota \circ \pi$ to the identity.

First it is useful to observe that, in this basis, the differentials have the form

$$d^{i} = \left[\begin{array}{ccc} 0 & 0 & 0 \\ 0 & 0 & 0 \\ id & 0 & 0 \end{array} \right]$$

A clear candidate for the map s^i is the projection onto B^{i-1} that then identifies it with its isomorphic copy as the third summand in the decomposition of V^{i-1} , i.e. the matrix

$$s^{i} = \left[\begin{array}{ccc} 0 & 0 & id \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{array} \right]$$

One can then check directly the equation

$$\mathrm{id} - \iota \circ \pi = \mathrm{d}^{\mathrm{i} - 1} \circ \mathrm{s}^{\mathrm{i}} + \mathrm{s}^{\mathrm{i} + 1} \circ \mathrm{d}^{\mathrm{i}}.$$

Remark 14.1.2. Using the barcode description, the map sⁱ simply follows length one barcodes to the left. Thus the length one barcodes can be interpreted as the "graph" of a chain homotopy.

Corollary 14.1.3. For any collection of integers m_i and n_i we have the following isomorphisms in the derived category.

$$\bigoplus_{i\in\mathbb{Z}}S_i^{n_i}\simeq\bigoplus_{i\in\mathbb{Z}}S_i^{n_i}\oplus P_i^{m_i}$$

The upshot of the above corollary is that

"Indecomposables do not survive the derived category!" [Mac14a]

14.2 DERIVED COMPLEXES OF CELLULAR SHEAVES

Recall from Chapter 7 that a complex of cellular sheaves F^{\bullet} assigns to every cell σ a chain complex and to every pair of incident cells $\sigma \leq \tau$ a chain map $\rho^{\bullet}_{\tau,\sigma} : F^{\bullet}(\sigma) \to F^{\bullet}(\tau)$. For each i we can define the ith cohomology sheaf as the assignment

$$\mathcal{H}^{i}(\mathsf{F}^{\bullet}): \qquad \sigma \rightsquigarrow \mathsf{H}^{i}(\mathsf{F}^{\bullet}(\sigma)).$$

The restriction maps being defined as the map induced on cohomology by $\rho_{\tau,\sigma}^{\bullet}$. A quasiisomorphism is a map of complexes $f^{\bullet} : F^{\bullet} \to G^{\bullet}$ that induces isomorphisms on each cohomology sheaf. We can also view all of the cohomology sheaves as a single graded cohomology sheaf \mathcal{H}^*F^{\bullet} .

14.2.1 Counterexample to the Naïve Approach

To begin, let us consider the simplest possible space where a complex of sheaves has interesting behavior. Let X = [0, 1) with cell structure $v = \{0\}$ and e = (0, 1). A complex of sheaves over X is completely described by two chain complexes and a chain map between them:

$$\rho_{e,\nu}^{\bullet}: F^{\bullet}(\nu) \to F^{\bullet}(e)$$



FIGURE 64: The Counterexample

At first glance there should be a simple adaptation of Lemma 14.1.1 to a pair of chain complexes. Indeed, we can use Theorem 6.3.3 over each cell to obtain a candidate inclusion map from the cohomology sheaf into the complex:



However such a map *does not commute* in general and as such fails to define a sheaf map. For example, if F[•] was the following pair of complexes and chain map



then the map ρ takes the one and only generator of H_{ν}^{i} to an element of the boundaries B_{e}^{i-1} , which is non-zero, but zero in cohomology, since $H^{i}(\rho) = 0$. We have provided a barcode version of this counter example in Figure 64.

14.2.2 Using the Calculus of Fractions Formulation

In order to address the counterexample to the naïve approach, we will employ a zig-zag of morphisms. In this section we briefly review why such a zigzag is natural, when viewing the derived category as the **left calculus of fractions** [GZ67]. An alternative description of the derived category goes as follows [She85, p.52-3].

Definition 14.2.1 (Derived Category via Fractions). The **bounded derived category of cellular sheaves**, $D^{b}(Shv(X))$, has the same collection of objects as Fun(X, Ch(Vect)), but with a modified class of morphisms. A morphism from F^{\bullet} to G^{\bullet} is a diagram of the following form



where the arrow decorated with a \simeq is a quasi-isomorphism. We declare two such morphisms (diagrams) to be the same if there is a larger, commutative diagram that fits in between them:



14.2.3 *The Equivalence*

We now proceed with a general method for addressing the earlier counterexample. The basic method of argument is that we will define a quasi-isomorphic complex of sheaves J• where the naïve approach does work. This is visualized in Figure 65.



Let's first illustrate our method over the simple base space X = [0, 1).



FIGURE 65: Replacing the Counterexample with a Quasi-Isomorphic Sheaf

Lemma 14.2.2. Over X = [0, 1) with cell structure $v = \{0\}$ and e = (0, 1), every bounded complex of cellular sheaves F^{\bullet} is quasi-isomorphic to its graded cohomology sheaf \mathcal{H}^*F^{\bullet} , i.e. F^{\bullet} is *isomorphic* to \mathcal{H}^*F^{\bullet} in the derived category.

Proof. As before, we decompose the chain complex over the vertex v and the edge e so that F[•] has the following form

$$\begin{array}{c} B_{\nu}^{i+1} \oplus H_{\nu}^{i+1} \oplus B_{\nu}^{i} \xrightarrow{\rho_{F}^{i+1}} B_{e}^{i+1} \oplus H_{e}^{i+1} \oplus B_{e}^{i} \\ & d_{\nu}^{i} \uparrow & \uparrow \\ B_{\nu}^{i} \oplus H_{\nu}^{i} \oplus B_{\nu}^{i-1} \xrightarrow{\rho_{F}^{i}} B_{e}^{i} \oplus H_{e}^{i} \oplus B_{e}^{i-1} \end{array}$$

The map ρ_F^i has an easily described form in this basis.

$$\rho_F^i = \left[\begin{array}{ccc} \alpha^i & 0 & 0 \\ \gamma^i & H^i \rho & 0 \\ \delta^i & \beta^{i-1} & \alpha^{i-1} \end{array} \right]$$

The naïve inclusion map's failure to commute is precisely described by non-zero terms in the submatrix β^{i-1} , as can be seen by inspecting the diagram below.



However, the complex of sheaves F^{\bullet} is quasi-isomorphic to the complex J^{\bullet} defined as follows.

The matrix representation for ρ_I^i has the desired form:

$$\rho_{J}^{i} = \begin{bmatrix} 0 & id & 0 & 0 & 0 \\ \gamma^{i} & 0 & H^{i}\rho & 0 & 0 \\ \delta^{i} & 0 & 0 & id \end{bmatrix}$$

The quasi-isomorphism $q^{\bullet} : F^{\bullet} \to J^{\bullet}$ is defined over the vertex ν in any given degree i as

$$q_{\nu}^{i} = \begin{bmatrix} id & 0 & 0 \\ \alpha^{i} & 0 & 0 \\ 0 & id & 0 \\ 0 & 0 & id \\ 0 & \beta^{i-1} & \alpha^{i-1} \end{bmatrix}$$

and over the edge *e* as $q_e^i = id$. By construction, the complex J[•] has a well defined sheaf map

which extends to a quasi-isomorphism $\iota : \mathcal{H}^*F^\bullet \hookrightarrow J^\bullet$ since each of the differentials $\mathcal{H}^iF^\bullet \to \mathcal{H}^{i+1}F^\bullet$ are zero. This completes the proof.

We can now prove the general theorem of interest.

Theorem 14.2.3. Let X be an arbitrary one dimensional cell complex and F^{\bullet} a complex of sheaves over X. Then F^{\bullet} is quasi-isomorphic to its graded cohomology sheaf. In particular, we have the equivalence of categories

$$D^{\mathfrak{b}}(\mathsf{Shv}(X)) \simeq \mathsf{Shv}(X; \mathsf{grVect}).$$

Proof. The bulk of the argument is contained in Lemma 14.2.2, which we show extends to the desired generality. The definition of $J^{i}(v)$ where v is a vertex with more than one incident edge is easily modified as follows:

$$J^{i}(\nu) := B^{i}_{\nu} \bigoplus_{e \geqslant \nu} B^{i}_{e} \oplus H^{i}_{\nu} \oplus B^{i-1}_{\nu} \bigoplus_{e \geqslant \nu} B^{i-1}_{e}$$

The restriction map to a single edge e' is defined exactly as in Lemma 14.2.2 with the stipulation that factors B_e^i and B_e^{i-1} where $e \neq e'$ are mapped to zero.

This argument implies that we can define the desired morphism



in an open neighborhood of a vertex. However, since over each edge $q_e^i : F^i(e) \to J^i(e)$ is defined to be the identity map, these locally defined maps agree on the edges and hence give a globally defined map of sheaves. This shows the equivalence with the graded cohomology sheaf.

One can check that every map of sheaves $f : F^{\bullet} \to G^{\bullet}$ extends to a map of the associated sheaves $J_f : J_F^{\bullet} \to J_G^{\bullet}$ so the construction is functorial and hence defines an equivalence of categories.

15

A METRIC ON THE CATEGORY OF SHEAVES

ΑΓΕΩΜΕΤΡΗΤΟΣ ΜΗΔΕΙΣ ΕΙΣΙΤΩ¹

— Purported Inscription at The Academy

Science depends on knowing that measurements and observations should only be trusted up to an interval of uncertainty. Saying that one is traveling about 30 miles per hour depends on their being a continuum of speeds. It makes no sense to say that one is traveling at a speed that happens to be an algebraic number. ² Persistent homology has addressed this issue in a rather elegant way. Given two functions $f, g : Y \to \mathbb{R}$ such that

$$\|f-g\|_{\infty} := \sup_{y} |f(y) - g(y)| < \varepsilon$$

one can say that the sublevel sets obey the inclusions

$$f^{-1}(-\infty,t] \subset g^{-1}(-\infty,t+\varepsilon]$$
 and $g^{-1}(-\infty,t] \subset t^{-1}(-\infty,t+\varepsilon]$

for every value of t. In the language of [CCSG⁺09], one has an **interleaving** of sublevel sets:

$$f^{-1}(-\infty, t] \longrightarrow f^{-1}(-\infty, t+\varepsilon] \longrightarrow f^{-1}(-\infty, t+2\varepsilon]$$
$$g^{-1}(-\infty, t] \longrightarrow g^{-1}(-\infty, t+\varepsilon] \longrightarrow g^{-1}(-\infty, t+2\varepsilon]$$

By functoriality of homology, one obtains a notion of interleaving of functors in $Fun(\mathbb{R}, Vect)$, where \mathbb{R} with its partial order is viewed as a category. Defining the interleaving distance to be the infimum over all ϵ such that there is an ϵ -interleaving gives an extended pseudo-metric on the category of such functors.

^{1 &}quot;Let no one who cannot think geometrically enter." [OR]

² This metaphor was used by Bob MacPherson in his opening remarks on "Continuity and the philosophy of science." [Mac14b]



FIGURE 66: Do Close Maps Give Rise to "Close" Sheaves?

In this chapter, we will consider a generalization of this approach, first for pre-sheaves and then for sheaves. We will try to answer analogous questions such as, "Suppose we have two maps f, $g: Y \to X$ to a metric space that are close in the supremum norm, i.e.

$$d_{\infty}(f,g) = \sup_{y \in Y} d_X(f(y),g(y)) < \varepsilon.$$

Is there any reasonable sense where the pushforward sheaves f_*k_Y and g_*k_Y are close?" It turns out that the answer is "yes," but studying higher invariants of the fiber, such as H^i for $i \ge 1$, is unstable to large perturbations. By studying interleavings of complexes of sheaves, one obtains a derived stability result. To conclude the chapter, we provide preliminary results towards equating the interleaving distance with a modified version of the bottleneck distance described in [CSEH07] for definable sheaves with finite support over the real line. The most important takeaway from this chapter is that interleavings for sheaves and pre-sheaves are obstructed by global sections.

15.1 INTERLEAVINGS FOR PRE-SHEAVES

Definition 15.1.1. Let (X, d) be a metric space. The ϵ -thickening of open sets is the map of posets

$$\epsilon$$
 : **Open**(X) \rightarrow **Open**(X)

given by

$$U \rightsquigarrow U^{\epsilon} := \bigcup_{x \in U} B(x, \epsilon).$$

Of course, any map of posets dualizes to a map of posets ϵ : **Open**(X)^{op} \rightarrow **Open**(X)^{op}

Remark 15.1.2. For metric spaces like \mathbb{R}^n with the Euclidean metric, $(U^{\epsilon})^{\epsilon} = U^{2\epsilon}$. In general, the triangle inequality implies that $(U^{\epsilon})^{\epsilon} \subseteq U^{2\epsilon}$. The reverse containment is also true if (X, d) is convex, for example. Convexity guarantees that intuitive results are true, but it is not strictly necessary for any of the following arguments.

Definition 15.1.3 (Thickened Pre-Sheaf). Using the previous two definitions, we can define the ϵ -thickening of a pre-sheaf F via the formula

$$F^{\varepsilon} := F \circ \varepsilon$$
 i.e. $F^{\varepsilon}(U) := F(U^{\varepsilon}).$

Moreover, the thickening operation is functorial. If $\varphi : F \to G$ is a natural transformation, then we get for free a natural transformation between the thickened pre-sheaves $\varphi_{\varepsilon} : F^{\varepsilon} \to G^{\varepsilon}$. Consequently, we can define the ε -thickening functor to be

$$\epsilon^* : \operatorname{PreShv}(X) \to \operatorname{PreShv}(X) \qquad F \rightsquigarrow F^{\epsilon}.$$

One of the most important observations for working with interleavings is that since F is a pre-sheaf we have a canonical natural transformation

$$\eta_{\epsilon}^{F}: F^{\epsilon} \to F$$

coming from $\rho_{U,U^\varepsilon}: F^\varepsilon(U) = F(U^\varepsilon) \to F(U)$. This follows by showing that for every pair $V \subseteq U$ the square

commutes by virtue of F being a pre-sheaf:

$$\rho_{V,U} \circ \rho_{U,U^{\epsilon}} = \rho_{V,U^{\epsilon}} = \rho_{V,V^{\epsilon}} \circ \rho_{V^{\epsilon},U^{\epsilon}}$$

Of course, the map

$$\eta_{2\varepsilon}^{\mathsf{F}}:\mathsf{F}^{2\varepsilon}\to\mathsf{F}$$

always exists as well and for metric spaces where $U^{2\epsilon} = (U^{\epsilon})^{\epsilon}$ it is equal to the composition $\eta_{\epsilon}^{F} \circ \epsilon^{*} \eta_{\epsilon}^{F}$.

Remark 15.1.4 (Notation 2ϵ vs. $\epsilon\epsilon$). To avoid cumbersome notation, we may sometimes substitute $U^{2\epsilon}$ for $(U^{\epsilon})^{\epsilon}$, $\eta_{2\epsilon}^{F}$ for $\epsilon^*\eta_{\epsilon}^{F}$, $F^{2\epsilon}$ for ϵ^*F^{ϵ} , and so on. For metric spaces such as \mathbb{R}^n with the Euclidean metric these differences do not exist and can be safely ignored.

A version of the following definition was communicated to the author by Amit Patel [Pat14].

Definition 15.1.5 (Interleaving of Pre-Sheaves). Let $F, G : Open(X)^{op} \to D$ be two presheaves on a metric space X. We define an ϵ -interleaving of F and G to be a pair of natural transformations

$$\varphi_{\varepsilon}: F^{\varepsilon} \to G \qquad \psi_{\varepsilon}: G^{\varepsilon} \to F$$

that satisfy the compatibility relations

$$\eta^F_{2\varepsilon}=\psi_\varepsilon\circ\varepsilon^*\phi_\varepsilon\qquad \eta^G_{2\varepsilon}=\phi_\varepsilon\circ\varepsilon^*\psi_\varepsilon.$$

An interleaving is better summarized via the following commutative diagram:



Observe that we have abused notation by writing $F^{2\epsilon}$ for $\epsilon^* F^{\epsilon}$. Also observe that there is a logical dualization for two pre-cosheaves.

Lemma 15.1.6. If two presheaves F and G are ϵ -interleaved for some $\epsilon \ge 0$, then they are ϵ' -interleaved for every $\epsilon' \ge \epsilon$.

Proof. Suppose that we have an ϵ interleaving, i.e. maps $\varphi_{\epsilon} : F^{\epsilon} \to G$ and $\psi_{\epsilon} : G^{\epsilon}$ that induce the correct commutative diagram. If $\epsilon' \ge \epsilon$ then the natural map $F^{\epsilon'} \to F$ factors through $F^{\epsilon} \to F$, allowing us to define $\varphi_{\epsilon'} = \varphi_{\epsilon} \eta^{F}_{\epsilon,\epsilon'}$ and symmetrically for $\psi_{\epsilon'}$. The map $(\epsilon')^* \varphi_{\epsilon'}$ is defined by

$$F((u^{\varepsilon'})^{\varepsilon'}) = F^{\varepsilon'}(u^{\varepsilon'}) \to F^{\varepsilon}(u^{\varepsilon'}) \to G(u^{\varepsilon'}) = G^{\varepsilon'}(u).$$

Observing that the map φ_{ε} is natural for a pair of open sets $U^{\varepsilon} \subset U^{\varepsilon'}$ proves that we have the commutative diagram in Figure 67.

Remark 15.1.7. One should note that if F and G are ϵ -interleaved and G and H are ϵ' -interleaved, then F and H are $\epsilon + \epsilon'$ -interleaved.

One should also note that if F and G are 0-interleaved, then they are isomorphic.



FIGURE 67: Diagram for the Proof of Lemma 15.1.6

Definition 15.1.8. We can define the interleaving distance on pre-sheaves by declaring

 $d(F, G) := \inf\{\epsilon \ge 0 \mid \exists \epsilon - interleaving\}.$

If no interleaving exists, we define $d(F,G) = \infty$. This is what we mean by an extended metric.

For pre-sheaves the interleaving distance is an extended pseudo-metric. "Extended" means that the distance ∞ is allowed and "pseudo" means that if d(F, G) = 0, then it does not follow that F = G. For sheaves, it is true that if a map induces isomorphisms on stalks, then the map is an isomorphism of sheaves. This suggests that if for every $\varepsilon > 0$ there is an ε -interleaving, then perhaps the sheaves are isomorphic. We will present an argument in this direction later on.

15.1.1 Easy Stability

The following results were obtained independently of [BdS13].

Lemma 15.1.9. Suppose $f, g : Y \to X$ are continuous maps to a metric space that are less than ϵ distance apart in the supremum norm, then the functors

 $\widehat{F}: U \rightsquigarrow f^{-1}(U) \qquad \text{and} \qquad \widehat{G}: U \rightsquigarrow g^{-1}(U)$

are ϵ -interleaved, when viewed as cosheaves.

Proof. By hypothesis, for every open set we have the following inclusions:

$$f^{-1}(U) \subseteq g^{-1}(U^{\varepsilon})$$
 and $f^{-1}(U^{\varepsilon}) \supseteq g^{-1}(U)$

This implies that we have an interleaving of pre-images:



Which is equivalent to saying that the functors \hat{F} and \hat{G} are ϵ -interleaved.

Corollary 15.1.10. If f, g : Y \rightarrow X are continuous maps such that $d_{\infty}(f,g) < \varepsilon$, then the cohomology presheaves

$$H^{i}F: U \rightsquigarrow H^{i}(f^{-1}(U); k)$$
 and $H^{i}G: U \rightsquigarrow H^{i}(g^{-1}(U); k)$

are ϵ -interleaved for every $i \ge 0$.

Proof. This follows from the lemma since cohomology is a functor.

Corollary 15.1.11. If f, g : $Y \to X$ are continuous maps such that $d_{\infty}(f,g) < \varepsilon$, then the pushforwards of the constant sheaf on Y,

$$f_*k_Y$$
 and g_*k_Y ,

are ϵ -interleaved.

Proof. The result follows by recalling the definition of the pushforward and that the constant sheaf records H^0 of an open set.

15.1.2 Global Sections Obstruct Interleavings

In this section, we introduce the first major difference between the interleaving distance for functors $S : (\mathbb{R}, \leq) \rightarrow$ Vect and interleavings for pre-sheaves.

Lemma 15.1.12 (Obstruction to Interleavings). If F and G are pre-sheaves on a metric space X and $F(X) \ncong G(X)$, then there is no ϵ -interleaving, for any value of ϵ .

Proof. Clearly $X^{\epsilon} = X$. Recalling the compatibility condition for interleavings



implies that $G(X) \rightarrow F(X)$ is a surjection. Considering the analogous triangle



implies that $F(X) \rightarrow G(X)$ is a surjection, which together proves that $F(X) \cong G(X)$. Contraposition proves the result.

Example 15.1.13 (Skyscraper Sheaf vs. Ephemeral Module). Suppose $\delta_x : \mathbb{R} \to \text{Vect}$ is the functor that assigns k to the points $x \in \mathbb{R}$ and 0 everywhere else. Such a persistence module is sometimes referred to as an **ephemeral module**. Without too much trouble, one can see that this functor is interleaving distance 0 from the zero functor.

In contrast, the skyscraper sheaf S_x , which assigns to every open set $U \ni x$ the vector space k and 0 to any open set not containing x, is infinite distance away from the zero sheaf.

15.2 INTERLEAVINGS FOR SHEAVES

Now we want to emulate the above constructions for functors $F : Open(X)^{op} \to \mathcal{D}$ that satisfy the sheaf axiom. Of course, we can regard any sheaf F as a pre-sheaf and apply the thickening construction to produce a pre-sheaf F^{ϵ} . Is the resulting pre-sheaf also a sheaf? No, because thickening can create intersections where there shouldn't be any, as in the following example.

Example 15.2.1. Let $F = k_X$ be the locally constant sheaf on $X = \mathbb{R}$. If $U_1 = (0, 1)$ and $U_2 = (1, 2)$, then for any $\epsilon > 0$ we have that



is not a limit diagram. This is due to the obvious defect

$$(\mathbf{U}_1 \cap \mathbf{U}_2)^{\epsilon} \neq \mathbf{U}_1^{\epsilon} \cap \mathbf{U}_2^{\epsilon}.$$

In light of the above example, thickening the underlying pre-sheaf must be followed by sheafification. Thus one approach to defining the thickening of a sheaf is to use the the sheafification of the thickened pre-sheaf. However, there is a nicer definition of the thickening of a sheaf that has the advantage that it can be defined in the derived setting, as well as being easier to prove theorems with. We state this definition now:

Definition 15.2.2 (Thickening via Convolution). Let X be a metric space and let $\bar{\Delta}_X^{\epsilon}$ be the subset of the product $\{(x, y) | d(x, y) \leq \epsilon\}$. Denote by ϵ_i the projection onto the *i*th coordinate of $\bar{\Delta}_X^{\epsilon}$. Define the ϵ -thickened sheaf \tilde{F}^{ϵ} to be $\epsilon_{2*}\epsilon_1^*F$. When the context is clear we will omit the tilde.

We will not show that the sheafification of the thickened pre-sheaf is the same as the above definition, as they very well could be different in certain cases.

Definition 15.2.3 (Interleavings for Sheaves). We say two sheaves F and G are ϵ -interleaved if there are maps $\varphi_{\epsilon}: \widetilde{F}^{\epsilon} \to G$ and $\psi_{\epsilon}: \widetilde{G}^{\epsilon} \to F$ such that the following diagram commutes



Derived Stability

Theorem 15.2.4. If f, g : $Y \rightarrow X$ are ϵ -close in the sup norm, then the derived pushforwards of the constant sheaf k_Y are ϵ -interleaved, i.e.

$$d(\mathsf{Rf}_*\mathsf{k}_{\mathsf{Y}},\mathsf{Rg}_*\mathsf{k}_{\mathsf{Y}}) \leqslant d_{\infty}(\mathsf{f},\mathsf{g})$$

Proof.

$$f^{-1}(U) \subseteq g^{-1}(U^{\epsilon})$$
 and $f^{-1}(U^{\epsilon}) \supseteq g^{-1}(U)$

The functors \widehat{F} , \widehat{G} : **Open**(X) \rightarrow **Top** are ε -interleaved, as already noted. This implies that the complexes of singular cochains are interleaved, which, after subdivision, defines a flabby resolution of the constant sheaf.

15.2.1 The Effect of Sheafification

In this section we investigate the impact of sheafification on interleavings of pre-sheaves. After all, every sheaf is a presheaf, so we can apply the notion of interleaving to both structures. As we will show by example, two pre-sheaves can be finite distance apart, but then be infinite distance apart (not ϵ -interleaved for any ϵ) after sheafification. Conversely, we can produce two pre-sheaves that are infinite distance apart, whose sheafifications are interleaved.

Let the map drawn on the left of Figure 68 be called f and let the map on the right be called g. As already observed in Corollary 15.1.10, the presheaves

$$H^{1}F: U \rightsquigarrow H^{1}(f^{-1}(U); k)$$
 and $H^{1}G: U \rightsquigarrow H^{1}(g^{-1}(U); k)$

are ϵ -interleaved for ϵ larger than the radius of the circle. However, the sheafification of both of these presheaves produce radically different sheaves.



FIGURE 68: Two Maps and their associated Leray Sheaves

Recall from Definition 2.5.6 that the sheafification can be viewed as the sheaf of sections of the product over all the stalks mapping down to the base space.

$$\prod_{\mathbf{x}\in\mathbb{R}}\mathsf{F}_{\mathbf{x}}$$

$$\downarrow^{\pi}$$

$$\mathbb{R}$$

For H^1F , every stalk is the zero vector space. For H^1G , the stalk at p is non-zero and all other ones are zero. Consequently the sheafifications are

$$\widetilde{F}^1 \cong 0$$
 and $\widetilde{G}^1 \cong S_p$

where S_p is the skyscraper sheaf at p. By lemma 15.1.12, these two sheaves are not interleaved.

One might conjecture in light of the above example that sheafification is a distance increasing operation. This is not the case. Consider the presheaf H¹F from the example above. One can easily see that it is *not* interleaved with the zero sheaf. However, the sheafification of H¹F is the zero sheaf. So sheafification took two presheaves that were infinite distance apart and returned isomorphic (distance zero) sheaves.

15.2.2 Thickening Global Sections

We will now prove that global sections remain unaffected by this thickening procedure. This provides us with a sheafified version of Lemma 15.1.12.

Lemma 15.2.5. If X is a metric space such that the closed ϵ -ball is connected for every point of x and the projection map $\epsilon_1 : \overline{\Delta}_X^{\epsilon} \to X$ is a closed map, then thickening preserves global sections, i.e.

$$F^{\epsilon}(X) \cong F(X).$$

Proof. This follows from a simpler version of the Vietoris mapping theorem, stated as Theorem 11.7 of [Bre97]. Essentially, the Vietoris mapping theorem guarantees that

$$H^{0}(X; F) \cong H^{0}(\overline{\Delta}_{X}^{\epsilon}; \epsilon_{1}^{*}F)$$

and since the pushforward $\epsilon_{2*}\epsilon_1^*F$ has the same global sections of ϵ_1^*F by functoriality (global sections are gotten by pushing forward to a point), then the result follows.

15.2.3 *Metric on Sheaves*

Claim 15.2.6. If F, G \in **Shv**(X) are interleaving-distance zero apart, i.e. there is a sequence of { ϵ_n } converging to zero where F and G are ϵ_n interleaved for each n, then F \cong G.

Sketch. We are going to define maps between the étalé spaces

$$\phi: \prod_{x \in X} F_x \to \prod_{x \in X} G_x \qquad \text{and} \qquad \psi: \prod_{x \in X} G_x \to \prod_{x \in X} F_x$$

with the property that they are inverses of one another.

Given any element $s_x \in F_x$ there exists a $U \ni x$ and a section s_U such that $(s_U)_x = s_x$. However, since U is open, there exists an r > 0 and $\varepsilon_n > 0$ such that $B(x, r + 2\varepsilon_n) \subset U$. This implies in turn that $B(x, r)^{2\varepsilon_n} \subset U$. Consequently, there is a lift $s_x^{2\varepsilon_n} \in \tilde{F}_x^{2\varepsilon_n}$ of s_x , gotten by taking the image of s_u under the following composition:

$$F(U) \to F(B(x, r+2\varepsilon_n)) \to F(B(x, r)^{2\varepsilon_n}) \to \widetilde{F}_{\chi}^{2\varepsilon_n}$$

We then define

$$\phi(s_{x}) := \phi_{x}^{\varepsilon_{n}} \circ \widetilde{\eta_{\varepsilon_{n},2\varepsilon_{n}x}}^{F}(s_{x}^{2\varepsilon_{n}})$$

Of course, by the definition of interleaving, we have a natural choice of a lift of $\varphi(s_x)$ to $\widetilde{G}_x^{\varepsilon_n}$ given by

$$\widetilde{\eta_{\varepsilon_n}}_{x}^{G}\phi_{x}^{2\varepsilon_n}(s_{x}^{2\varepsilon_n})\text{,}$$

which by construction has the property that applying ψ_x^{ϵ} yields s_x .

Of course it needs to be checked that such a lift is picked out by the symmetric construction of the map ψ . It could have happened that a lift to $\widetilde{G}_{x}^{\epsilon'_{n}}$ for $\epsilon'_{n} \neq \epsilon_{n}$. However, one only needs to observe that for any sheaf F we have the following commutative diagram for any pair $0 < \epsilon'_{n} < \epsilon_{n}$:



Arguing symmetrically and checking continuity of the resulting maps φ and ψ completes the proof.

Theorem 15.2.7 (Skeletal Metric Space). Let X be an arbitrary metric space. The interleaving distance induces an extended metric on the skeleton of the category of sheaves Shv(X). Here extended means that the value $d(F,G) = \infty$ is allowed, i.e. there is no interleaving between F and G whatsoever.

Proof. Recall that the skeleton of a category C is a full, isomorphism-dense subcategory S in which no two distinct objects are isomorphic.

If we take C = Shv(X), on which the interleaving distance already defines an extended pseudo-metric, then the above result implies that if d(F, G) = 0, then $F \cong G$ and hence in a skeletal subcategory F = G. This implies that the interleaving distance is an extended metric when restricted to any skeletal subcategory of Shv(X).

As one can imagine, the space of sheaves viewed as a metric space can be enormously complicated. Every map $f : X \to \mathbb{R}^n$ has an associated sheaf on \mathbb{R}^n , simply by considering the pushforward of the constant sheaf. This includes every possible subspace $Y \subset \mathbb{R}^n$ with the pushforward of the constant sheaf along this inclusion serving as a sort of "indicator function" on it. The interleaving distance would give us one notion of distance between all these possible subspaces. In this case, there is a ready comparison to be made with the Gromov-Hausdorff distance between metric spaces, which is more refined than the interleaving distance. However, the interleaving distance also gives a distance between *information* on top of a metric space, where information is encoded via a sheaf.

15.3 The space of constructible sheaves over ${\mathbb R}$

In this section we will give an explicit description of the space of constructible/cellular sheaves on \mathbb{R} with the interleaving distance as a metric. It turns out that one can use the

indecomposable sheaves to give a set of "coordinates" on this space. It will turn out that the space resembles a disjoint union of configuration spaces, where there the components are divvied up by global sections, i.e. H⁰. A comparison to McDuff's construction of the tangent space use a configuration space of points, where points can disappear, is made.

First we recall the definition of constructible sheaf pertinent to this section. Because there are competing, more general notions of a constructible sheaf, we will use slightly different terminology.

Definition 15.3.1. Let $Shv_d(\mathbb{R})$ denote the **category of definable sheaves** over the real line \mathbb{R} , equipped with the usual Euclidean topology. Specifically, a sheaf F will be regarded as a contravariant functor from the open set category with the necessary gluing properties. Such a sheaf F is **definable** if \mathbb{R} can be written as the finite union of open intervals and points

$$\mathbb{R} = (-\infty, a_0) \cup \{a_0\} \cup \cdots (a_i, a_{i+1}) \cdots \cup \{a_n\} \cup (a_n, \infty)$$

such that when restricted to each interval the sheaf is locally constant. We do **not** assume that every sheaf is constructible with respect to the same set of intervals.

We will find it convenient to work with a subcategory of this category given by the definable sheaves with **finite support**, $Shv_{d,f}(X)$, where the sheaf must restrict to zero on the two half-open intervals including $\pm \infty$.

As already established in this thesis, such a sheaf is completely described via a zig-zag of vector spaces and linear maps

$$F(a_0) \leftarrow F(x_0) \rightarrow F(a_1) \leftarrow F(x_1) \rightarrow \cdots \leftarrow F(x_n) \rightarrow F(a_n)$$

where we have abbreviated the intervals and points by using a's and x's along with subscripts appropriately.

By Gabriel's theorem, we know that such a diagram amounts to a representation of an A_n -type quiver and can be decomposed into finitely many indecomposable representations. These indecomposables, in view of the cell structure, can be regarded as one of the following sheaves:

- $k_{[x_i,x_i]}$ the constant sheaf on the closed interval $[x_i, x_j]$
- $k_{(x_i,x_i)}$ the constant sheaf on the open interval (x_i, x_j)
- $k_{[x_i,x_i]}$ the constant sheaf on the half-open interval $[x_i,x_j]$
- $k_{(x_i,x_i]}$ the constant sheaf on the half-open interval $(x_i, x_j]$
- $k_{\mathbb{R}}$ the constant sheaf supported on the whole real line \mathbb{R}

• S_x — the skyscraper sheaf concentrated on the vertex x

The skyscraper sheaf, S_x , is just a special instance of the constant sheaf on a closed interval $[x_i, x_j]$ where $x_i = x_j$. As such, we will not always distinguish the skyscraper sheaf from the constant sheaf on the closed interval. Similarly, one can view the constant sheaf on \mathbb{R} , $k_{\mathbb{R}}$, as a degenerate version of the open interval.

To keep the notation clean and free ourselves from a particular declaration of cell structure on the real line, we will speak of the **four indecomposable sheaves** on the real line:

$$k_{[b,d]}$$
 $k_{(b,d)}$ $k_{[b,d)}$ $k_{(b,d)}$

Remark 15.3.2. The constant sheaf, $k_{\mathbb{R}}$, and any of the others where b or d is $\pm \infty$ are excluded from the category of definable sheaves with finite support.

In the following sections it will be paramount to understand when there is and isn't a non-zero map of sheaves between these four types.

Proposition 15.3.3 (Dévissage for 1D Indecomposable Sheaves). We have the following explicit characterizations for the space of sheaf morphisms:

• For closed intervals $I_1 = [b_1, d_1]$ and $I_2 = [b_2, d_2]$ we have

$$\operatorname{Hom}_{\operatorname{Shv}}(k_{I_1},k_{I_2}) = \begin{cases} k & \text{if } I_2 \subseteq I_1, \\ 0 & \text{o.w.} \end{cases}$$

• For open intervals $I_1 = (b_1, d_1)$ and $I_2 = (b_2, d_2)$ we have

$$\operatorname{Hom}_{\operatorname{Shv}}(k_{I_1},k_{I_2}) = \begin{cases} k & \text{if } I_1 \subseteq I_2, \\ 0 & \text{o.w.} \end{cases}$$

• For half open intervals of the form $I_n = [b_n, d_n)$ we have

$$\operatorname{Hom}_{\mathsf{Shv}}(k_{I_1},k_{I_2}) = \begin{cases} k & \text{if } b_1 \leqslant b_2 \text{ and } b_2 < d_1 \leqslant d_2, \\ 0 & \text{o.w.} \end{cases}$$

• For half open intervals of the form $I_n = (b_n, d_n]$ we have

$$\operatorname{Hom}_{\mathsf{Shv}}(k_{I_1}, k_{I_2}) = \begin{cases} k & \text{if } b_2 \leqslant b_1 \text{ and } b_1 < d_2 \leqslant d_1, \\ 0 & \text{o.w.} \end{cases}$$

• For a closed interval $I_1 = [b_1, d_1]$ and any non-compact interval I_2

$$Hom_{Shv}(k_{I_1}, k_{I_2}) = 0$$

Proof. The calculations all follow from considering the behavior of a cellular sheaf map near the endpoints of the above indecomposables. Specifically, in a small enough neighborhood of any point in the real line, an indecomposable cellular sheaf has one of the following three forms, possibly after reflection.

$$M = 0 \leftarrow k \rightarrow k$$
$$U = k \leftarrow k \rightarrow k$$
$$H = k \leftarrow 0 \rightarrow 0$$

Clearly there are non-zero natural transformations $H \Rightarrow U \Rightarrow M$, but every natural transformation $M \Rightarrow U \Rightarrow H$ must be zero.

Remark 15.3.4 (Dévissage for Constructible Sheaves). In David Nadler's beautiful application of constructible sheaves to the study of the Fukaya category [Nad11], he refers to the diagram



as the "dévissage pattern for constructible sheaves" — an ode to Grothendieck's method for studying *coherent* sheaves. Here $j : V \rightarrow X$ is the inclusion of an open set and $i : Y := X - V \rightarrow X$ is the inclusion of the closed complement. Here the categories **Shv**(–) refer to the full differential graded category of constructible complexes of sheaves, H⁰ of which is the usual derived category.

15.3.1 Interleavings and Dynamics on Indecomposable Sheaves

Since indecomposable sheaves are the underlying elements that build up a sheaf, we investigate the behavior of each of these sheaves under epsilon-thickening. We summarize the result of each of these calculations below:

• If
$$F = k_{[b,d]}$$
 and $\epsilon \ge 0$, then $F^{\epsilon} = k_{[b-\epsilon,d+\epsilon]}$.

• If $F = k_{(b,d)}$ and $0 \leq \varepsilon < d-b$, then $\widetilde{F}^{\varepsilon} = k_{(b+\varepsilon,d-\varepsilon)}$. If $d-b \leq \varepsilon$, then $\widetilde{F}^{\varepsilon} = 0$.

- If $F = k_{[b,d]}$ and $\epsilon \ge 0$, then $\widetilde{F}^{\epsilon} = k_{[b-\epsilon,d-\epsilon]}$.
- If $F = k_{(b,d)}$ and $\epsilon \ge 0$, then $\widetilde{F}^{\epsilon} = k_{[b+\epsilon,d+\epsilon]}$.

With these calculations in hand, we can then discuss distance between these sheaves.

- The sheaf $k_{[b,d]}$ is interleaving distance r = (d b)/2 from the skyscraper sheaf S_m where m = (d + b)/2 the midpoint and infinite distance from any of the other three types of indecomposables, as well as the zero sheaf.
- The sheaf $k_{(b,d)}$ is interleaving distance r = (d b)/2 from the zero sheaf 0.
- The sheaves $k_{[b,d]}$ and $k_{(b,d]}$ are interleaving distance r = (d-b)/2 from the zero sheaf 0.

We now give the supporting arguments for these calculations.

Proposition 15.3.5. If $F = k_{[b,d]}$ and $\epsilon \ge 0$, then $\widetilde{F}^{\epsilon} = k_{[b-\epsilon,d+\epsilon]}$.

Proof. To construct \tilde{F}^{ε} it suffices to consider the stalks of the thickened pre-sheaf. It suffices to consider the extreme points. Consider the point $x = b - \varepsilon$, then any ball B(x,r) has the property that $[b,d] \cap B(x,r)^{\varepsilon} = [b,b+r) \neq \emptyset$. Since F is defined as the pushforward of the constant sheaf along $j : X = [b,d] \hookrightarrow \mathbb{R}$, then $F^{\varepsilon}(B(x,r)) = k_X([b,b+r)) = k$. This proves that $F_x \cong k$.

Proposition 15.3.6. If $F = k_{(b,d)}$ and $0 \le \varepsilon < d-b$, then $\widetilde{F}^{\varepsilon} = k_{(b+\varepsilon,d-\varepsilon)}$. If $d-b \le \varepsilon$, then $\widetilde{F}^{\varepsilon} = 0$.

Proof. If $j : W = (b, d) \hookrightarrow \mathbb{R}$ denotes the inclusion of the open interval, then we can identify $F = j_! k_W$. Here $j_!$ denotes the pushforward with compact supports functor. For the inclusion of a locally closed subspace *W* into a general topological space X [Ive86] provides a precise description. For a sheaf F on *W* the sheaf $j_!F$ has sections on an open set U given by

$$\Gamma(\mathbf{U}, \mathbf{j}_{!}\mathbf{F}) := \{ \mathbf{s} \in \Gamma(\mathbf{W} \cap \mathbf{U}, \mathbf{F}) | \operatorname{supp}(\mathbf{s}) \operatorname{closed} \operatorname{rel}. \mathbf{U} \}$$

The support of a section is the set of points where a section has non-vanishing stalks.

For an open set $U \subset \mathbb{R}$, we can describe the sheaf pertinent to us even more explicitly: $j_!k_W(U)$ is non-zero if and only if there is a closed set Y such that $U \subset Y \subset W$ — we say that U is **completely contained** in W. Note that W = (b, d) is *not* completely contained in itself, so in particular (b, d) is assigned the zero vector space by $j_!k_W$.

Consequently, any point x within distance ϵ of the boundary of (b, d) will fail to possess a ball B(x, r) such that $B(x, r)^{\epsilon}$ is completely contained in (b, d). Thus we say that (b, d) is "eroded" by distance ϵ .

Proposition 15.3.7. If $F = k_{[b,d]}$ and $\epsilon \ge 0$, then $\widetilde{F}^{\epsilon} = k_{[b-\epsilon,d-\epsilon]}$.

Proof. This follows by considering each of the above arguments separately about each endpoint. \Box

Because the other case is obviously symmetric, we omit a separate argument.

An extremely important observation that distinguishes interleavings of sheaves from the usual context of an interleavings of persistence modules [BL13] is that the sheaf $k_{[b,d]}$ is not interleaved with the zero sheaf. This follows from the fact that the space of global sections is preserved by thickening.

15.3.2 Coordinates for the Category of Sheaves

We can now use the decomposition theorem for constructible sheaves over the real line to give explicit coordinates for each isomorphism class of a definable sheaf with finite support.

Let $\mathbb{H} = \{(x, y) \in \mathbb{R}^2 | y \ge 0\}$ denote the closed upper half plane. To each of the four indecomposable sheaves with finite support — $k_{[b,d]}$, $k_{(b,d)}$, $k_{[b,d)}$, and $k_{(b,d]}$ — we can associate a point in \mathbb{H} as follows:

$$I \rightsquigarrow (x,y) = (\mathfrak{m}(I), \mathfrak{r}(I)) := (\frac{b+d}{2}, \frac{d-b}{2})$$

Here I is a stand-in for any of the four types of indecomposable sheaves. The variable names m and r are meant to connote the midpoint and the radius, respectively, of the underlying bar in the "barcode". One can then associate to any definable sheaf F with finite support the following coordinates

$$F \cong \bigoplus_{k=1}^{n} k_{I_k} \rightsquigarrow \{(\mathfrak{m}(I_1), \mathfrak{r}(I_1)), \cdots, (\mathfrak{m}(I_n), \mathfrak{r}(I_n))\} \in \mathbb{H}^{n_1} \sqcup \mathbb{H}^{n_2} \sqcup \mathbb{H}^{n_3} \sqcup \mathbb{H}^{n_4}$$

where n_1 , n_2 , n_3 , n_4 refers to the number of closed, open, half-open on the right and halfopen on the left indecomposables occurring in the decomposition for F, respectively. Of course, $n = n_1 + n + 2 + n_3 + n_4$.

As presented, the space has too many points for the simple reason that the line r = 0 in \mathbb{H} must be identified with the zero sheaf for the non-closed indecomposable types. To capture the full category of sheaves we must append a distinguished basepoint \star to represent the zero sheaf, quotient the upper half plane so as to identify $\{r = 0\} \sim \{\star\}$ and then form a few infinite symmetric products. The first part is simple. We define

$$\mathsf{Z} := (\mathbb{H} \sqcup \{\star\}) / \sim \qquad \text{where} \qquad (\mathfrak{m}, \mathfrak{r}) \sim \star \text{ iff } \mathfrak{r} = \mathfrak{0}.$$

The next construction begins by observing that Z is naturally a pointed space, where \star serves as the distinguished basepoint. To every pointed space we can associate a new space called the infinite symmetric product.

Definition 15.3.8. Recall that the n-fold symmetric product of X, denoted $SP_n(X)$, is given by forming the n-fold Cartesian product and quotienting out by the action of the symmetric group, i.e. $SP_n(X) := X^n / \Sigma_n$.

Let Z be a pointed topological space, whose distinguished point is \star . There is a system of embeddings $Z^n \hookrightarrow Z^{n+1}$ given by sending any point

$$(z_1,\ldots,z_n)\mapsto(z_1,\ldots,z_{n+1},\star).$$

This embedding descends to an embedding $SP_n(Z) \hookrightarrow SP_{n+1}(Z)$, which forms a directed system of spaces. The **infinite symmetric product** SP(Z) is the direct limit of this system, i.e.

$$SP(Z) = \varinjlim_{n>0} SP_n(Z).$$

We can now state a theorem.

Theorem 15.3.9. Isomorphism classes of $\mathsf{Shv}_{d,f}(\mathbb{R})$ are in bijective correspondence with points in the following space:

$$\mathbb{B} := \bigsqcup_{n>0} \operatorname{SP}_{n}(\mathbb{H}) \times \left(\operatorname{SP}(Z) \bigvee_{\star} \operatorname{SP}(Z) \bigvee_{\star} \operatorname{SP}(Z) \right)$$

Remark 15.3.10. If one could show that the topology induced by the interleaving distance made \mathbb{H} and Z into cell complexes, then the Dold-Thom theorem would tell us the singular homology of the space of definable sheaves on the real line is isomorphic to a countably infinite number of copies of \mathbb{Z} in degree zero and zero in all higher degrees.

It remains to be seen what geometry is induced on \mathbb{B} by pulling back the interleaving distance. Conjecturally, this should be accomplished by the bottleneck distance [CCSG⁺09, BL13], with the stipulation that only points in SP(Z) can be matched with zero.

15.3.3 Towards a Bottleneck Distance for Sheaves

We now investigate simpler descriptions of the interleaving distance for definable sheaves with finite support on the real line. Our goal is to establish a connection between the interleaving distance for sheaves and the bottleneck distance, which we now define. **Definition 15.3.11** (The Bottleneck Distance). Let $\Delta_+ := \{(x, y) \in \mathbb{R}^2 | x \leq y\}$ be equipped with the sup norm

$$d_{\infty}(p,p') = \sup\{|x_1 - x_2|, |y_1 - y_2|\}$$

A multiset D in Δ_+ is a subset |D| of Δ_+ equipped with a multiplicity function $\mu : |D| \rightarrow \mathbb{N}$. Any multiset can be considered as a set via disjoint unions

$$D = \bigcup_{p \in |D|} \coprod_{i=1}^{\mu(p)} \{p\}.$$

A multi-bijection $m : D \to D'$ is a bijection between the underlying sets, where if $\mu(p) = k$, then the underlying set has k elements corresponding to p. The **bottleneck distance** is a distance between multisets defined by the formula

$$d_{B}(D,D') := \inf_{\mathfrak{m}} \sup_{p \in D} d_{\infty}(p,p').$$

In usual sub-level set persistence, the multisets D consist of the diagonal $\Delta = \{(x, x)\}$, equipped with infinite multiplicity, and the points (b, d) corresponding to the interval modules $k_{[b,d]}$ making up the interval decomposition of theorem 6.3.3. The stipulation that the diagonal has infinite multiplicity reflects the fact that for half-open intervals, the module can be interleaved with zero. For sheaves, the obstruction by global sections result implies that the number of points corresponding to indecomposables $k_{[b,d]}$ is an invariant of the multi-set — infinite multiplicity of the diagonal cannot be used there. Nevertheless, we can describe the geometry there in our choice of coordinates.

Lemma 15.3.12. Suppose $F = k_{[m_1-r_1,m_1+r_1]}$ and $G = k_{[m_2-r_2,m_2+r_2]}$ are two indecomposable sheaves supported over closed intervals, then the interleaving distance for F and G is their distance in a taxicab metric on \mathbb{H} , i.e.

$$d(k_{[m_1-r_1,m_1+r_1]},k_{[m_2-r_2,m_2+r_2]}) = |m_1-m_2| + |r_1-r_2|.$$

Remark 15.3.13. If we write $[m_1 - r_1, m_1 + r_1] =: [b_1, d_1]$ and $[m_2 - r_2, m_2 + r_2] =: [b_2, d_2]$ then the taxicab metric specializes to a sup-norm on the space $\{(b, d) \in \mathbb{R}^2 | b \leq d\}$, that is to say

$$|\mathbf{m}_1 - \mathbf{m}_2| + |\mathbf{r}_1 - \mathbf{r}_2| = \sup\{|\mathbf{b}_1 - \mathbf{b}_2|, |\mathbf{d}_1 - \mathbf{d}_2|\},\$$

which is a special instance of the bottleneck distance on persistence diagrams.

Proof. Without loss of generality we can assume that $m_1 \leq m_2$. We know that there can only be a non-zero map from $F^{\varepsilon} \to G$ if $[m_2 - r_2, m_2 + r_2] \subset [m_1 - r_1 - \varepsilon, m_1 + r_1 + \varepsilon]$, i.e. if $\varepsilon \geq (m_2 - m_1) + (r_2 - r_1)$. Similarly, there is a non-zero map $G^{\varepsilon} \to F$ only if
$\varepsilon \ge (m_2 - m_1) + (r_1 - r_2)$. Since in order for a non-zero interleaving to exist both maps must be non-zero, we conclude that

$$\epsilon = \sup\{(m_2 - m_1) + (r_2 - r_1), (m_2 - m_1) + (r_1 - r_2)\} = m_2 - m_1 + |r_1 - r_2|.$$

This proves that there is a non-zero interleaving for this value of ϵ . However, the interleaving distance is the infimum over all such ϵ . However, for any smaller ϵ one of the maps φ_{ϵ} or ψ_{ϵ} must be zero. However, global sections obstructs such a pair of maps from defining an interleaving, since for every ϵ the maps

$$\eta_{2\epsilon}^{\mathsf{F}}:\mathsf{F}^{2\epsilon}(\mathsf{X})\to\mathsf{F}(\mathsf{X})$$
 and $\eta_{2\epsilon}^{\mathsf{G}}:\mathsf{G}^{2\epsilon}(\mathsf{X})\to\mathsf{G}(\mathsf{X})$

are non-zero. They are, in fact, the identity map $id : k \rightarrow k$.

The situation for constant sheaves supported on half-open intervals or the open interval is more complicated since they have no global sections and can be interleaved with the zero sheaf.

Lemma 15.3.14. Suppose F and G are each interleaved with the zero sheaf, i.e. $d(F,0), d(G,0) < \infty$. Then the interleaving distance between F and G is bounded above by the greater of the two distances from zero,

$$d(F,G) \leq \sup\{d(F,0), d(G,0)\}.$$

Proof. Choose any $\epsilon \ge \sup\{d(F, 0), d(G, 0)\}$, then the zero maps define an ϵ interleaving, factoring through zero, between F ang G. Since the interleaving distance is the infimum, the inequality follows.

We can use the above lemma to establish what the interleaving distance between pairs of sheaves of the other three types looks like.

Lemma 15.3.15. Let $F = k_{[m_1-r_1,m_1+r_1]}$ and $G = k_{[m_2-r_2,m_2+r_2]}$ be indecomposable sheaves supported on half-open intervals, then

$$d(k_{[m_1-r_1,m_1+r_1]},k_{[m_2-r_2,m_2+r_2]}) = \inf\{\max(r_1,r_2), |m_1-m_2| + |r_1-r_2|\}$$

Proof. Without loss of generality, we assume that $F = k_{[-R,R]}$ and $G = k_{[m-r,m+r]}$ where $m \ge 0$. This can be done since the ϵ -thickening operation simply translates sheaves of this form to the left by ϵ . In view of lemma 15.3.14 it suffices to consider the possible values for m and r such that G^{ϵ} admits a non-zero map to F. By proposition 15.3.3 we can determine precisely what inequalities ϵ must satisfy.

First we assume that $r \leq R$. In this case, the inequalities are

$$-r - R \leq m - \epsilon \leq r - R$$
 and $R - r \leq m + \epsilon < R + r$.

The first set of inequalities bounds when $G^{\epsilon} \to F$ can be non-zero and the second set of inequalities bounds when $F^{\epsilon} \to G$ can be non-zero. The smallest such ϵ that satisfies the first set is $\epsilon = m + R - r$, which only satisfies the second set if m < R. For these values of m and r, there is the smallest non-zero interleaving, so the distance is

$$\mathbf{d}(\mathbf{F},\mathbf{G}) = \mathbf{m} + \mathbf{R} - \mathbf{r}.$$

Similarly for $r \ge R$, $\varepsilon = m + r - R$ is the smallest possible value for $G^{\varepsilon} \to F$ to be non-zero. Requiring $F^{\varepsilon} \to G$ to be non-zero as well implies that $m \le R$ and for these values of m and r the distance is

$$\mathbf{d}(\mathsf{F},\mathsf{G})=\mathsf{m}+\mathsf{r}-\mathsf{R}.$$

Consequently, wherever a non-zero interleaving is possible, the smallest such value is given in the desired taxicab form m + |r - R|. Now we can apply lemma 15.3.14 to determine whether the non-zero interleavings are the smallest possible. We split into two cases. Suppose $r \leq R$, then $m + R - r \leq R$ — the largest distance from zero — when $m \leq r$, precisely where we determined a non-zero interleaving exists. Similarly if $R \leq r$, then we have that $m + r - R \leq r$ precisely if $m \leq R$.

With this evidence in hand, we believe a version of the isometry theorem [Les12, BS, BL13] should hold for definable sheaves with finite support over the real line. However, we must delay this for another time.

- [Acho7] 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.
- [AFT12] David Ayala, John Francis, and Hiro Lee Tanaka. Structured singular manifolds and factorization homology. arXiv:1206.5164 [math.AT], December 2012.
- [AGV72] 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.
- [AHS09a] Jiri Adamek, Horst Herrlich, and George E. Strecker. Abstract and Concrete Categories: The Joy of Cats. Dover Books on Mathematics. Dover Publications, 2009. Available online at http://katmat.math.uni-bremen.de/acc/ acc.pdf.
- [AHS09b] 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.
 - [AK10] 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.
 - [Ale28] Paul Alexandroff. AIJber den allgemeinen dimensionsbegriff und seine beziehungen zur elementaren geometrischen anschauung. *Mathematische Annalen*, 98:617–635, 1928.
 - [Ale37] P. Alexandroff. Diskrete räume. Mat. Sb, 2:501–518, 1937.
 - [Ale47] P. S. Alexandrov. *Combinatorial Topology*. Dover Publications, Inc., 31 East 2nd Street, Mineola, NY, 1947.
 - [Aluo9] Paolo Aluffi. *Algebra. Chapter o*, volume 104. Graduate Studies in Mathematics American Mathematical Society, 2009.

- [AM69] M.F. Atiyah and I.G. MacDonald. *Introduction To Commutative Algebra*. Addison-Wesley series in mathematics. Westview Press, 1969.
- [AoM62] M. Artin and Harvard University. Dept. of Mathematics. *Grothendieck topologies: notes on a seminar*. Harvard University, Dept. of Mathematics, 1962.
 - [AR94] J. Adámek and J. Rosicky. *Locally Presentable and Accessible Categories*. Cambridge Tracts in Mathematics. Cambridge University Press, 1994.
- [Arno4] V.I. Arnol'd. Lectures on Partial Differential Equations. Universitext (1979). Springer, 2004.
- [Art91] M. Artin. Algebra. Prentice Hall, 1991.
- [Ati56] Michael Atiyah. On the krull-schmidt theorem with application to sheaves. *Bull. Soc. Math. France*, 84, 1956.
- [Aw010] Steve Awodey. Category Theory. Oxford University Press, 2nd edition, 2010.
- [Bac75] Kenneth Bacławski. Whitney numbers of geometric lattices. *Advances in Mathematics*, 16(2):125–138, 1975.
- [BBR07] Morten Brun, Winfried Bruns, and Tim Römer. Cohomology of partially ordered sets and local cohomology of section rings. *Advances in Mathematics*, 208(1):210–235, 2007.
- [BD13] Dan Burghelea and Tamal K Dey. Topological persistence for circle-valued maps. *Discrete & Computational Geometry*, 50(1):69–98, 2013. arXiv:1104.5646 [math.AT].
- [BdS13] P. Bubenik, V. de Silva, and J. Scott. Metrics for generalized persistence modules. *ArXiv e-prints*, December 2013.
- [BFo6] Marta Bunge and Jonathan Funk. *Singular coverings of toposes*, volume 1890 of *Lecture Notes in Mathematics*. Springer-Verlag, 2006.
- [BL13] Ulrich Bauer and Michael Lesnick. Induced matchings of barcodes and the algebraic stability of persistence. *arXiv preprint arXiv:1311.3681*, 2013.
- [Bor48] Karol Borsuk. On the imbedding of systems of compacta in simplicial complexes. *Fund. Math.*, 35:217–234, 1948.
- [Bot72] Raoul Bott. Lectures on characteristic classes and foliations. In *Lectures on Algebraic and Differential Topology,* volume 279 of *Lecture Notes in Mathematics,* pages 1–94. Springer Berlin Heidelberg, 1972. Notes by Lawrence Conlon.
- [Bre68] Glen Bredon. Cosheaves and homology. Pacific J. Math., 25(1):1–32, 1968.

- [Bre97] Glen Bredon. *Sheaf Theory*, volume 170 of *Graduate Texts in Mathematics*. Springer-Verlag, 2nd edition, 1997.
- [Bro67] R Brown. Groupoids and van kampen's theorem. In *Proc. London Math. Soc.*(3), volume 17, pages 385–401, 1967.
 - [BS] Peter Bubenik and Jonathan A Scott. Categorification of persistent homology. *Discrete & Computational Geometry*, pages 1–28.
- [BT82] R. Bott and L.W. Tu. Differential Forms in Algebraic Topology. Springer, 1982.
- [Caro9] Gunnar Carlsson. Topology and data. *Bulletin of the American Mathematical Society*, 46(2):255–308, 2009.
- [CB12] William Crawley-Boevey. Decomposition of pointwise finite-dimensional persistence modules. *arXiv preprint arXiv:1210.0819*, 2012.
- [CCSG⁺09] Frédéric Chazal, David Cohen-Steiner, Marc Glisse, Leonidas J Guibas, and Steve Y Oudot. Proximity of persistence modules and their diagrams. In Proceedings of the 25th annual symposium on Computational geometry, pages 237–246. ACM, 2009.
 - [CdS10] Gunnar Carlsson and Vin de Silva. Zigzag persistence. Foundations of Computational Mathematics, 10(4):367–405, August 2010. http://arxiv.org/abs/ 0812.0197.
- [CdSM09] 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. http://www.mrzv.org/ publications/zigzags/.
 - [CG] Kevin Costello and Owen Gwilliam. Factorization algebras in quantum field theory. http://math.northwestern.edu/~costello/factorizationpdf.
- [CGN13] J. Curry, R. Ghrist, and V. Nanda. Discrete Morse theory for computing cellular sheaf cohomology. *ArXiv e-prints*, December 2013.
 - [Coso2] Michel Coste. An introduction to semialgebraic geometry. lecture notes, October 2002.
- [CSEH07] David Cohen-Steiner, Herbert Edelsbrunner, and John Harer. Stability of persistence diagrams. *Discrete & Computational Geometry*, 37(1):103–120, 2007.
 - [CZ09] Gunnar Carlsson and Afra Zomorodian. The theory of multidimensional persistence. *Discrete & Computational Geometry*, 42(1):71–93, 2009.

- [CZCG04] Gunnar Carlsson, Afra Zomorodian, Anne Collins, and Leonidas Guibas. Persistence barcodes for shapes. In *Proceedings of the 2004 Eurographics/ACM SIGGRAPH symposium on Geometry processing*, pages 124–135. ACM, 2004.
 - [DIo1] Daniel Dugger and Daniel Isaksen. Hypercovers in topology. K-theory Preprint Archives, http://www.math.uiuc.edu/K-theory/0528/, November 2001.
 - [DL12] Lee DeVille and Eugene Lerman. Dynamics on networks of manifolds. *arXiv* preprint arXiv:1208.1513, 2012.
 - [DM69] Pierre Deligne and David Mumford. The irreducibility of the space of curves of given genus. *Publications Mathématiques de l'IHES*, 36(1):75–109, 1969.
 - [Dou07] Christopher L. Douglas. Sheaves in homotopy theory. 2007.
- [dSGo6a] Vin de Silva and Robert Ghrist. Coordinate-free coverage in sensor networks with controlled boundaries via homology. *The International Journal of Robotics Research*, 25(12):1205–1222, 2006.
- [dSGo6b] 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.
 - [dSSo4] P.N. de Souza and J.N. Silva. *Berkeley Problems in Mathematics*. Problem Books in Mathematics. Springer, 2004.
 - [DW05] Harm Derkden and Jerzy Weyman. Quiver representations. *Notices of the AMS*, 52(2):200–206, Feb 2005.
 - [EFS56] 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.
- [EGH⁺11] 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.
 - [EM45] Samuel Eilenberg and Saunders MacLane. General theory of natural equivalences. *Transactions of the American Mathematical Society*, 58(2):pp. 231–294, 1945.
 - [Eme60] Ralph Waldo Emerson. *The Journals of Ralph Waldo Emerson*. Random House-/The Modern Library, 1960.

- [EP12] Mário Edmundo and Luca Prelli. The six grothendieck operations on ominimal sheaves. Available at http://www.ciul.ul.pt/~edmundo/6Go_03_ 10_12.pdf, 2012.
- [FF56] L. R. Ford and D. Fulkerson. Maximal flow through a network. *Canadian Journal of Mathematics*, 8:399–404, 1956.
- [Fri11] Joel Friedman. Sheaves on graphs, their homological invariants, and a proof of the hanna neumann conjecture. arXiv:1105.0129 [math.CO], 2011.
- [Fun95] J. Funk. The display locale of a cosheaf. *Cahiers de Topologie et Géométrie Différentielle Catégoriques*, 36(1):53–93, 1995.
 - [GG] Ryan Grady and Owen Gwilliam. l_{∞} spaces and derived loop spaces.
- [GH11] Robert Ghrist and Yasuaki Hiraoka. Applications of sheaf cohomology and exact sequences on network codings. http://www.math.upenn.edu/ ~ghrist/preprints/networkcodingshort.pdf, 2011.
- [Ghro8] Robert Ghrist. Barcodes: The persistent topology of data. *Bulletin Amer. Math. Soc.*, 45:61–75, 2008.
- [GM83] M. Goresky and R. MacPherson. Intersection homology ii. *Inventiones Mathematicae*, 72(1):77–129, 1983.
- [GM88] 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.
- [GMV96] Sergei Gelfand, Robert MacPherson, and Kari Vilonen. Perverse sheaves and quivers. *Duke Mathematical Journal*, 83(3):621–643, June 1996.
 - [Gor76] Robert Mark Goresky. Geometric Cohomology and Homology of Stratified Objects. PhD thesis, Brown University PhD Thesis, May 1976. Pages 15-21 are available, with the author's permission, at http://hans.math.upenn.edu/ ~jucurry/papers/regular_nbhd.pdf.
 - [Gor78] R.M. Goresky. Triangulation of stratified objects. *Proceedings of the American Mathematical Society*, pages 193–200, 1978.
 - [Gor12] 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.
 - [Gra79] John W. Gray. Fragments of the history of sheaf theory. *Lecture Notes in Mathematics*, 753:1–79, 1979.

- [Gro97] A. Grothendieck. Sketch of a programme. 1997. Orginally written in 1984. Translated and available online at http://matematicas.unex.es/~navarro/ res/esquisseeng.pdf.
- [GZ67] P. Gabriel and M. Zisman. *Calculus of Fractions and Homotopy Theory*. Ergebnisse der Mathematik und ihrer Grenzgebiete. Springer Berlin Heidelberg, 1967.
- [Hato2] A. Hatcher. *Algebraic Topology*. Cambridge University Press, 2002.
- [Isao2] Daniel C. Isaksen. Calculating limits and colimits in pro-categories. *Fund. Math.*, 175:175–194, 2002.
- [Ive86] B. Iversen. *Cohomology of Sheaves*. Universitext (Springer-Verlag). Springer Berlin Heidelberg, 1986.
- [Kas84] Masaki Kashiwara. The riemann-hilbert problem for holonomic systems. *Publ. RIMS, Kyoto Univ.*, 20:319–365, 1984.
- [KS02] 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.
- [KS06] Masaki Kashiwara and Pierre Schapira. *Categories and Sheaves*, volume 332 of *Grundlehren der Mathematischen Wissenschaften*. Springer-Verlag, 2006.
- [Lado8] Sefi Ladkani. On derived equivalences of categories of sheaves over finite posets. *Journal of Pure and Applied Algebra*, 212(2):435–451, 2008. arXiv:0705.0946 [math.RT].
- [Ler45] Jean Leray. Sur la forme des espaces topologiques et sur les points fixes des representations. *Journal de Math.*, 24:95–167, 1945.
- [Ler46] Jean Leray. L'anneau d'homologie d'une représentation. C.R. Acad. Sci. Paris, (222):1366–1368, 1946.
- [Les12] Michael Lesnick. Multidimensional interleavings and applications to topological inference. arXiv:1206.1365 [math.AT], 2012.
- [LHM⁺10] Peter D. Lax, Friedrich Hirzebruch, Barry Mazur, Lawrence Conlon, Edward B. Curtis, Harold M. Edwards, Johannes Huebschmann, and Herbert Shulman. Raoul bott as we knew him. Kotiuga, P. Robert (ed.), A celebration of the mathematical legacy of Raoul Bott. Based on the conference, CRM, Montreal, Canada, June 9–13, 2008. Providence, RI: American Mathematical Society (AMS). CRM Proceedings and Lecture Notes 50, 43-49 (2010)., 2010.

- [Lin65] F. E. J. Linton. Review: Peter freyd, abelian categories. an introduction to the theory of functors. Bulletin of the American Mathematical Society, 71(4):577–580, 07 1965.
- [Loi98] T.L. Loi. Verdier and strict thom stratifications in o-minimal structures. *Illinois Journal of Mathematics*, 42(2):347–356, 1998.
- [Loi10] 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.
- [LSL⁺13] PY Lum, G Singh, A Lehman, T Ishkanov, M Vejdemo-Johansson, M Alagappan, J Carlsson, and G Carlsson. Extracting insights from the shape of complex data using topology. *Scientific reports*, 3, 2013.
 - [Lu76] Y.C. Lu. *Singularity theory and an introduction to catastrophe theory*. Universitext (1979). Springer-Verlag, 1976.
- [Luro9a] J. Lurie. *Higher Topos Theory (AM-170)*. Annals of Mathematics Studies. Princeton University Press, 2009.
- [Luro9b] Jacob Lurie. Derived algebraic geometry vi: E_k algebras. *arXiv preprint arXiv:0911.0018*, 2009.
 - [LZ] Ryan H Lewis and Afra Zomorodian. Multicore homology.
- [Mac89] Saunders Mac Lane. Concepts and categories in perspective. A century of mathematics in America, Pt. I, Hist. Math. 1, 323-365 (1989)., 1989.
- [Mac91] Robert MacPherson. Intersection homology and perverse sheaves. Colloquium lecture notes distributed by the American Mathematical Society, hosted online by intrepid researchers., 1991.
- [Mac98] Saunders Mac Lane. *Categories for the Working Mathematician*, volume 5 of *Graduate Texts in Mathematics* 5. Springer-Verlag, 2nd edition, 1998.
- [Mac14a] Robert MacPherson. Personal communication, 2012–2014.
- [Mac14b] Robert MacPherson. Night seminars at the institute for advanced study, 2012-2014. multiple lectures.
 - [Mat12] John; Mather. Notes on topological stability. *Bulletin (New Series) of the American Mathematical Society*, 49(4):475–506, October 2012.
- [May99] J. P. May. *Concise Course in Algebraic Topology*. Chicago Lectures in Mathematics, University of Chicago Press, 1999.

- [McC] Jason McCarty. Limits and colimits in abelian categories. Notes from UVA Math 760 Course in Homological Algebra. To be found at http://www.math. virginia.edu/~ww9c/HomoAlg08/limits-Jason.pdf.
- [McC66] M.C. McCord. Singular homology groups and homotopy groups of finite topological spaces. *Duke Mathematical Journal*, 33(3):465–474, 1966.
- [McC01] J. McCleary. *A User's Guide to Spectral Sequences*. Cambridge Studies in Advanced Mathematics. Cambridge University Press, 2001.
- [Milo6] David Miller. Popaths and holinks. *Journal of Homotopy and Related Structures*, 1(1):1–9, 2006.
- [MM92] S. Mac Lane and I. Moerdijk. *Sheaves in Geometry and Logic: A First Introduction to Topos Theory*. Universitext (1979). Springer-Verlag, 1992.
- [MP08] J. Milton and P. Pullman. *Paradise Lost*. Oxford World's Classics. OUP Oxford, 2008.
- [MT12] Davesh Maulik and David Treumann. Constructible functions and lagrangian cycles on orbifolds. arXiv:1110.3866 [math.AG], 2012.
- [Munoo] J.R. Munkres. *Topology*. Prentice Hall, Incorporated, 2000.
- [Nad11] David Nadler. Fukaya categories as categorical morse homology. arXiv:1109.4848 [math.SG], 2011.
- [NTT14] N. Nguyen, S. Trivedi, and D. Trotman. A geometric proof of the existence of definable Whitney stratifications. *ArXiv e-prints*, April 2014.
 - [OR] John J. O'Connor and Edmund F. Robertson. The academy of plato.
 - [Pat14] Amit Patel. Personal communication, 2012–2014.
 - [Pflo1] M.J. Pflaum. Analytic and Geometric Study of Stratified Spaces: Contributions to Analytic and Geometric Aspects, volume 1768. Springer, 2001.
 - [Pit85] A. M. Pitts. On product and change of base for toposes. *Cahiers de Top et Geom. Diff. Categoriques*, 26(1):43–61, 1985.
 - [Poio8] Henri Poincaré. The future of mathematics. *Revue generale des Sciences pures et appliquees*, (23), December 1908.
 - [Pra11] Andrei V. Prasolov. Precoshieves of pro-sets and abelian pro-groups are smooth. arXiv:1105.3167 [math.AT], May 2011.

- [PS99] 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.
- [Qui73] Daniel Quillen. Higher algebraic k-theory: I. In *Higher K-theories*, pages 85–147. Springer, 1973.
- [Ramo5] S. Ramanan. *Global Calculus*. Graduate studies in mathematics. American Mathematical Society, 2005.
- [Ree46] 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.
 - [Sch] Pierre Schapira. Algebra and topology. Lecture Notes from Paris VI University 2007/2008. Version 2 dated 1/9/11. To be found at http://www.math.jussieu.fr/~schapira/lectnotes/AlTo.pdf.
- [Sch87] Jean-Pierre Schneiders. Cosheaves homology. *Bull. Soc. Math. Belgique*, 39:1–31, 1987.
- [Sch97] Peter Schneider. Representation theory and sheaves on the bruhat-tits building. *Publications mathematiques de l' I.H.E.S.*, 85:97–191, 1997.
- [Sch98] Peter Schneider. Verdier duality on the building. J. reine angew Math., 494:205–218, 1998.
- [Scho3] Jörg Schürmann. Topology of Singular Spaces and Constructible Sheaves, volume 63 of Monografie Matematyczne Instytut Matematyczny New Series. Birkhäuser Verlag, 2003.
- [Seg68] Graeme Segal. Classifying spaces and spectral sequences. *Publications mathèmatiques de l 'I.H.È.S.*, 34, 1968.
- [Sey77] P.D. Seymour. The matroids with the max-flow min-cut property. *Journal of Combinatorial Theory, Series B*, 23(2):189–222, 1977.
- [She85] Allen Shepard. *A Cellular Description of the Derived Category of a Stratified Space*. PhD thesis, Brown University PhD Thesis, May 1985.
 - [Shi] Masahiro Shiota. *Geometry of Subanalytic and Semialgebraic Sets,* volume 150 of *Progress in Mathematics.* Birkhäuser, Boston, MA.
- [Shi97] M. Shiota. *Geometry of Subanalytic and Semialgebraic Sets*. Progress in Mathematics. Birkhäuser, 1997.

- [Shioo] Masahiro Shiota. ThomâĂŹs conjecture on triangulations of maps. *Topology*, 39:383–399, 2000.
- [Spa94] Edwin Spanier. *Algebraic Topology*. Springer-Verlag, 1994. originally published by McGraw-Hill in 1966.
- [Ste43] N. E. Steenrod. Homology with local coefficients. *Annals of Mathematics*, 44(4):pp. 610–627, 1943.
- [Str93] Gilbert Strang. The fundamental theorem of linear algebra. *American Mathematical Monthly*, pages 848–855, 1993.
- [Sugo1] Yuichi Sugiki. The category of cosheaves and laplace transforms. Preprint Series 33, The University of Tokyo Graduate School of Mathematical Sciences, December 2001.
- [SV96] W. Schmid and K. Vilonen. Characteristic cycles of constructible sheaves. *Inventiones Mathematicae*, 124(1):451–502, 1996.
- [Tho62] René Thom. La stabilité topologique des applications polynomiales. *Enseignement Math.* (2), 8:24–33, 1962.
- [Tho69] R. Thom. Ensembles et morphismes stratifiés. *Bull. Amer. Math. Soc*, 75(2):240–284, 1969.
- [Treo9] David Treumann. Exit paths and constructible stacks. *Compositio Mathematica*, 145:1504–1532, 2009. http://arxiv.org/abs/0708.0659v1.
- [vdD98] L.P.D. van den Dries. *Tame Topology and O-minimal Structures*. London Mathematical Society Lecture Note Series. Cambridge University Press, 1998.
- [vdDM96] 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.
 - [VW88] T. Vaughan and A.E. Waite. The Magical Writings of Thomas Vaughan (Eugenius Philatethes): A Verbatim Reprint of His First Four Treatises: Anthroposophia Theomagica, Anima Magica Abscondita, Magia Adamica, and the True Cœlum Terræ. G. Redway, 1888.
 - [Vyb97] Maxim Vybornov. Mixed algebras and quivers related to cell complexes. *Commun. Algebra*, 25(12):3985–3992, 1997.
 - [Vyb98] Maxim Vybornov. Constructible sheaves on simplicial complexes and koszul duality. *Mathematical Research Letters*, 5(675):675–683, 1998.

- [Vyb99] Maxim Vybornov. Sheaves on Triangulated Spaces and Koszul Duality. PhD thesis, Yale University, 1999. condensed version available at http://arxiv. org/abs/math/9910150.
- [Wei94] Charles A. Weibel. *An Introduction to Homological Algebra*, volume 38 of *Cambridge studies in advanced mathematics*. Cambridge University Press, 1994.
- [Whi65a] Hassler Whitney. Local properties of analytic varieties. In Differential and Combinatorial Topology (A Symposium in Honor of Marston Morse), pages 205– 244. Princeton Univ. Press, Princeton, N. J., 1965.
- [Whi65b] Hassler Whitney. Tangents to an analytic variety. *Annals of Mathematics*, 81(3):pp. 496–549, 1965.
- [Whi78] George Whitehead. *Elements of Homotopy Theory*, volume 61 of *Graduate Texts in Mathematics*. Springer-Verlag, 1978.
- [Woo09] Jonathan Woolf. The fundamental category of a stratified space. *Journal of Homotopy and Related Structures*, 4(1):359–387, 2009. arXiv:0811.2580.
- [Woo14] Jon Woolf, 2012-2014. Personal Communication.
- [Yano1] Kohji Yanagawa. Sheaves on finite posets and modules over normal semigroup rings. *Journal of Pure and Applied Algebra*, 161(3):341–366, 2001.
- [Yano5] Kohji Yanagawa. Dualizing complex of the incidence algebra of a finite regular cell complex. *Illinois Journal of Mathematics*, 49(4):1221–1243, 2005. arXiv:math/0407383 [math.RA].
- [Yuz81] Sergey Yuzvinsky. Linear representations of posets, their cohomology and a bilinear form. *European Journal of Combinatorics*, 2(4):385–397, 1981.
- [ZCo8] Afra Zomorodian and Gunnar Carlsson. Localized homology. *Computational Geometry*, 41(3):126–148, 2008.
- [Zee62a] Christopher Zeeman. Dihomology. i. relations between homology theories. *Proc. London Math. Soc.*, 3(12):609–638, 1962.
- [Zee62b] Christopher Zeeman. Dihomology. ii. *Proc. London Math. Soc.*, 3(12):639–689, 1962.
- [Zee63] Christopher Zeeman. Dihomology. iii. a generalization of the poincaré duality for manifolds. *Proc. London Math. Soc.*, 3(13):155–183, 1963.

INDEX

absolute homology cosheaf, 137 adjoint, 16 functor theorem, Freyd's, 17 Alexandrov topology, 73 unit interval example, 74 analytic-geometric category, 201 barcode, 100 associated to chain complex, 101, 246 filtration of Bott's torus, 101 fundamental theorem of linear algebra, 100 generalized, 157 length, 100 trick, 139 blow-up, 191 Borel-Moore cosheaf homology, 126 invariance under subdivision, 126 bottleneck distance, 272

capacity, 159 category, 3 of local systems, 53 abelian, 103 additive, 103 $C^{b}(A)$ of chain complexes, 121 Ch[•](Vect) of chain complexes, 93 (co)filtered, 40 data, 5 D^b(A) derived category, 121 calculus of fractions, 250 discrete, 4 exact, 102 finite, 3 fundamental groupoid $\pi_1(X)$, 5 groupoid, 4

 $K^{\mathfrak{b}}(\mathbf{A})$ homotopy category of chain complexes, 121 of a cell complex Cell(X), 68 of cellular sheaves Shv(X), cosheaves CoShv(X), 71 of cocones, 13 of cones, 10 of covers, 32, 40 of definable sheaves, 266 of functors, 8 of intervals, 135 of open sets, 4 of points and open sets, 4 of representations, 102 of sheaves coordinates, 270 opposite, 6Vect, **43** path/free, 107 skeletal subcategory, 8 small, 3 subcategory, 6 Čech homology, 28 $\partial^2 = 0, 29$ example of, 30 Čech sheaves, 145 cell complex, 68 subdivision, 83 cellular maps, 87 bijection with a circle, 90 inclusion of a point, 89 projection to a point, 88 cellular perverse sheaves, 109 cellular sheaves and cosheaves, 68 (co)chain complex, 93

homotopy, 120 cocones, 13 category of, 13 coend, 240 higher Euler calculus, 244 coequalizer, 15 cohomology, 94 of a cell complex, 95 of a cellular sheaf, 96 compactly-supported, 95, 243 coinvariants, 15 cokernel, 15 colimit, 13 as a coequalizer, 26 filtered, 40 commutes with finite limits, 41 is a union, 14 over a cover, 25 synonyms for, 13 cones, 10 category of, 10 control data, 202 coproduct, 14, 37 cosheaf absolute homology, 137 axiom as a coequalizer, 26 axiom as an exact sequence, 27 axiom for two opens, 24 behavior under refinement, 34 cellular, 68 Borel-Moore homology, 126 Borel-Moore homology of, 97 Borel-Moore homology under subdivision, 126 derived equivalence with sheaves, 235 height function on torus, 69 homology of, 98 homology via barcodes, 138 Klein bottle, 69 cellular homology

Lf_{*} left derived pushforward, 124 cohomology, cellular, 128 subdivision, 130 constructible, 189 from a definable map, 224 Entr(X), 223 cosheafification, 44 approaches, 43 Set, 45 universal property as a right adjoint, 45 definition of, 22 elementary injective on a poset, 114 elementary projective on a poset, 116 evasion, 182 given by H_0 , 32 locally constant cosheaf, 57 representation of $\pi_1(X)$, 57 Mayer-Vietoris, 31 of compactly supported functions, 24,63 of connected components, 61 of features, 60 of pre-images, 60 of singular chains, 62 on a poset, 76 on closed sets, 78, 86 push-pull adjunctions, 90 pushforward, 80, 83 pushforward with compact supports, 88 sections over a poset, 77 skyscraper, 60 three perspectives on, 18 costalk, 40 cover, 20 refinement of, 32 covering space, 51

data types, 5 decoding edge, 163 decomposition

generalized barcode, 157 interval/barcode, 100 into strata, 190 Remak, 104 sensing and evasion regions, 183 definable entrance path category, 213 $D^{b}(A)$ derived category, 121 calculus of fractions, 250 derived equivalence of cellular sheaves and cosheaves, 235 with graded sheaves, 250, 253 with graded vector spaces, 246 diagram, 9 direct image, 80, 83 discrete Morse theory for sheaves, 144 disjoint union, 37 duality derived equivalence, 235 for routing sheaves, 162 linear, 98, 237 network coding sheaf, 162 over a manifold, 129, 234 Poincaré-Verdier equivalence functor, 233 Verdier, 238 Entr(X) entrance path category, 212 definable, 213 representations are cosheaves, 222 specializes to Cell(X), 213 van Kampen theorem, 216 € thickening, 255 equalizer, 12 equivalence adjoint, 8 fully faithful and essentially surjective, 8 evasion cosheaf, 182 exact sequence of presheaves, 41

sequence of representations, 103

sequence of sheaves, 41 Exit(X) exit path category, 212 f[†] notation, 84 family of lines, 204 fiber bundle, 51 fiber byndle gives local systems, 54 functor, 6 fundamental groupoid $\pi_1(X)$, 5 and locally constant (co)sheaves, 57 determines locally constant (co)sheaves, 59 local systems, 53 Poincaré ∞ -groupoid, 5 Gabriel's theorem, 107 global sections pushforward to a point, 84 graded vector space, 93 graph, 159 Grothendieck Six Functor Formalism, 38 groupoid, 4 $K^{b}(A)$ homotopy category of chain complexes, 121 homology of a cellular cosheaf, 98 Borel-Moore, 97 via barcodes, 138 homotopy link, 214 incidence relation, 94 indecomposable representation, 104 constant (co)sheaves on connected, 105 initial object, 6 injective cellular sheaf, 114 cellular sheaves are sums, 115 cosheaf on a poset, 114

mapping identities for (co)sheaves, 116 objects, 113 properties, 113 resolutions, 117 sheaf on a poset, 114 interleaving of presheaves, 257 obstructed by global sections, 260 of sheaves, 261 over the real line, 268 interval module, 99 intruder problem, 170 approaches to, 172 invariants, 12 Kan extension, 76 motivation, 75 kernel, 12 Krull-Schmidt Theorem, 105 Lf_{*} left derived pushforward, 123 Leray sheaves, 146 limit, 10 as an equalizer, 26 cofiltered, 40 do not commute with finite colimits, 42 is an intersection, 11 over a cover, 25 synonyms for, 11 linear duality, 98, 237 local orientation, 94 local system, 52, 53 Klein bottle and torus examples, 55 from a fiber bundle, 54 representation of $\pi_1(X)$, 53 locale, 61 Möbius bundle, 51 Mayer-Vietoris, 31 as a cosheaf axiom, 62 blowup complex, 149

via cellular sheaf cohomology, 145 module interval, 99 left and right, 239 persistence, 99 natural transformation, 7 naturally isomorphic, 7 nerve, 21 generalization, 149 nerve theorem, 21, 144 generalization, 149 network coding sheaf, 159 duality, 162 o-minimal structure, 199 analytic-geometric category, 201 persistence module, 99 decomposition theorem, 100 persistent homology, 132 level set determines sublevel set, 150 level set/zigzag, 135 multi-dimensional sublevel set, 155 via derived pushforward, 156 sublevel set, 134 perversity, 109 Pitt's theorem, 241 point cloud, 132, 200 determines a stratified map, 211 pointless topology, 61 poset, 4, 72 maps, 81 pre-cosheaf, 22 behavior under refinement, 33 preordered set, 4, 72 presheaf, 22 behavior under refinement, 33 exact sequence, 41 inconsistent example, 36 non-local example, 35 stalk, 36

product, 11, 37 projective cosheaf on a poset, 116 mapping identities for (co)sheaves, 116 objects, 115 properties, 116 resolutions, 117 sheaf on a poset, 116 pullback, 12 of a sheaf over a poset, 82 of sheaves, 38 pushforward to a point, 84 pushforward with compact supports f!, 88 pushforward with open supports f_{\dagger} , 85 pushout, 14 quasi-isomorphism, 122 quiver, 106, 159 representation, 107 Reeb graph, 64 regular cell complex, 67 relation, 4 representation, 9 direct sum, 103 indecomposable, 104 of a quiver, 107 Entr(X), 222 Entr(X)from a stratified map, 224 $\pi_1(X)$, 53, 57 subrepresentation, 103 routing sheaf, 161 duality, 162 semialgebraic subset, 200 sensing sheaf, 181 sensors, 169 multi-modal, 179 sheaf axiom as an equalizer, 26

axiom as an exact sequence, 27 axiom for two opens, 24 behavior under refinement, 34 Čech, 145 cellular, 68 cohomology of, 96 compactly supported cohomology of, 95, 243 derived equivalence with cosheaves, 235 derived equivalence with graded sheaves, 250, 253 subdivision, 83 cellular perverse, 109 constant sheaf, 52 pullback from a point, 82 constructible, 189 definable, 266 definition of, 22 elementary injective cellular sheaf $[\sigma]^{W}$, 114 elementary injective on a poset, 114 elementary projective on a poset, 116 Hom, 39 homology, cellular, 128, 244 examples, 164 subdivision, 130 Leray, 146 locally constant sheaf, 52, 57 representation of $\pi_1(X)$, 57 network coding, 159 of sections, 49 on a poset, 76 on closed sets, 19, 78, 86 pullback, 38 push-pull adjunctions, 90 pushforward, 38, 80, 83 pushforward f_†, 85 pushforward with compact supports f₁, 88 restricted sheaf, 189

```
routing, 161
    sections over a poset, 77
    sensing, 181
   sheafification, 37, 43
     universal property as a left adjoint,
        38
    stalk, 36, 39
    \otimes, 39
    thickening, 261
    three perspectives on, 18
sheaves
    exact sequence of, 41
skeletal subcategory, 8
splitting lemma, 104
stability, 258
    for derived interleavings, 262
stacks, 68
stalk, 36, 39, 40
star, 73
strata, 190
   closed union, has regular
       neighborhood, 205
```

stratified map determines constructible cosheaves, 224 Thom mapping, 207, 209 Whitney, 196 stratified space general properties of, 195 Thom-Mather, 203 Whitney, 194 stratified submersion, 196 stratified submersion, 196 stratified vector field, 207 subdivision, 83 of a cellular sheaf, 83 tame topology, 198 tensoring sheaves with cosheaves, 240

tensoring sheaves with cosheaves, 240 terminal object, 6 Thom's first isotopy lemma, 196

van Kampen theorem, 58 Entr(X), 216 Verdier duality, 238 Vopenka's principle, 45