IAS/Park City Mathematics Series
Volume 00, Pages 000–000
S 1079-5634(XX)0000-0

Homological Algebra and Data

Robert Ghrist

Abstract. These lectures are a quick primer on the basics of applied algebraic topology with emphasis on applications to data. In particular, the perspectives of (elementary) homological algebra, in the form of complexes and co/homological invariants are sketched. Beginning with simplicial and cell complexes as a means of enriching graphs to higher-order structures, we define simple algebraic topological invariants, such as Euler characteristic. By lifting from complexes of simplices to algebraic complexes of vector spaces, we pass to homology as a topological compression scheme. Iterating this process of expanding to sequences and compressing via homological algebra, we define persistent homology and related theories, ending with a simple approach to cellular sheaves and their cohomology. Throughout, an emphasis is placed on expressing homological-algebraic tools as the natural evolution of linear algebra. Category-theoretic language (though more natural and expressive) is deemphasized, for the sake of access. Along the way, sample applications of these techniques are sketched, in domains ranging from neuroscience to sensing, image analysis, robotics, and computation.

Contents

2010 *Mathematics Subject Classification.* Primary 55-01; Secondary 18G35, 55N30.
Key words and phrases. cohomology, complexes, homology, persistence, sheaves.
RG supported by the Office of the Assistant Secretary of Defense Research & Engineering through ONR N00014-16-1-2010.

Introduction & Motivation

These lectures are meant as an introduction to the methods and perspectives of Applied Topology for students and researchers in areas including but not limited to data science, neuroscience, complex systems, and statistics. Though the tools are mathematical in nature, this article will treat the formalities with a light touch and heavy references, in order to make the subject more accessible to practitioners. See the concluding section for a roadmap for finding more details. The material is pitched to a level that a beginning graduate student in any of the applied mathematical sciences will have no difficulty.

What is Homology?

Homology is an algebraic compression scheme that excises all but the essential topological features from a particular class of data structures arising naturally from topological spaces. Homology therefore pairs with topology. Topology is the mathematics of abstract space and transformations between them. The notion of a space, X, requires only a set together with a notion of nearness, expressed as a system of subsets comprising the "open" neighborhoods satisfying certain consistency conditions. Metrics are permissible but not required. So many familiar notions in applied mathematics – networks, graphs, data sets, signals, imagery, and more – are interpretable as topological spaces, often with useful auxiliary structures. Furthermore, manipulations of such objects, whether as comparison, inference, or metadata, are expressible in the language of mappings, or continuous relationships between spaces. Topology concerns the fundamental notions of equivalence up to the loose nearness of what makes a space. Thus, connectivity and holes are significant; bends and corners less so. Topological invariants of spaces and mappings between them record the essential qualitative features, insensitive to coordinate changes and deformations.

Homology is the simplest, general, computable invariant of topological data. In its most primal manifestation, the homology of a space X returns a sequence of vector spaces $H_\bullet(X)$, the dimensions of which count various types of *linearly independent* holes in X. Homology is inherently linear-algebraic, but transcends linear algebra, serving as the inspiration for *homological algebra*. It is this algebraic engine that powers the subject.

When is Homology Useful?

Homological methods are, almost by definition, robust, relying on neither precise coordinates nor careful estimates for efficacy. As such, they are most useful in settings where geometric precision fails. With great robustness comes both great flexibility and great weakness. Topological data analysis is more fundamental than revolutionary: such methods are not intended to supplant analytic, probabilistic, or spectral techniques. They can however reveal a deeper basis for why some data sets and systems behave the way they do. It is unwise to wield topological techniques in isolation, assuming that the weapons of unfamiliar "higher" mathematics are clad in incorruptible silver.

Scheme

There is far too much material in the subject of algebraic topology to be surveyed here. Existing applications alone span an enormous range of principles and techniques, and the subject of applications of homology and homological algebra is in its infancy still. As such, these notes are selective to a degree that suggests caprice. For deeper coverage of the areas touched on here, complete with illustrations, see [51]. For alternate ranges and perspectives, there are now a number of excellent sources, including [40,62,76]. These notes will deemphasize formalities and ultimate formulations, focusing instead on principles, with examples and exercises. The reader should not infer that the theorems or theoretic minutiae are anything less than critical in practice.

These notes err or the side of simplicity. The many included exercises are not of the typical lemma-lemma-theorem form appropriate for a mathematics course; rather, they are meant to ground the student in examples. There is an additional layer of unstated problems for the interested reader: these notes are devoid of figures. The student apt with a pen should endeavor to create cartoons to accompany the various definitions and examples presented here, with the aim of minimality and clarity of encapsulation. The author's attempt at such can be found in [51].

Lecture 1: Complexes & Homology

This lecture will introduce the initial objects and themes of applied algebraic topology. There is little novel here: all definitions are standard and found in standard texts. The quick skip over formalities, combined with a linear-algebraic sensibility, allows for a rapid ascent to the interesting relationships to be found in homology and homological algebra.

Spaces

A *space* is a set X together with a compendium of all subsets in X deemed "open," which subcollection must of necessity satisfy a list of intuitively obvious properties. The interested reader should consult any point-set topology book (such as [70]) briefly or until interest wanes. All the familiar spaces of elementary calculus – surfaces, level sets of functions, Euclidean spaces – are indeed topological spaces and just the beginning of the interesting spaces studied in manifold theory, algebraic geometry, and differential geometry. These tend to be frustratingly indiscrete. Applications involving computation prompt an emphasis on those spaces that are easily digitized. Such are usually called *complexes*, often with an adjectival prefix. Several are outlined below.

Simplicial Complexes Consider a set X of discrete objects. A k-*simplex* in X is an unordered collection of $k+1$ distinct elements of X. Though the definition is combinatorial, for X a set of points in a Euclidean space [*viz.* point-cloud data set] one visualizes a simplex as the geometric convex hull of the $k+1$ points, a "filled-in" clique: thus, 0-simplices are points, 1-simplices are edges, 2-simplices are filled-in triangles, etc. A *complex* is a collection of multiple *simplices*.[1] In particular, a *simplical complex* on X is a collection of simplices in X that is *downward closed*, in the sense that every subset of a simplex is also a simplex. One says that X *contains all its faces*.

Exercise 1.1. Recall that a collection of random variables $\mathcal{X} = \{X_i\}_1^k$ on a fixed domain are statistically independent if their probability densities f_{X_i} are jointly multiplicative (that is, the probability density $f_{\mathcal{X}}$ of the combined random variable $(X_1, \ldots, X_k)$ satisfies $f_{\mathcal{X}} = \prod_i f_{X_i}$). Given a set of n random variables on a fixed domain, explain how one can build a simplicial complex using statistical independence to define simplices. What is the maximal dimension of this *independence complex*? What does the number of connected components of the independence complex tell you? Is it possible to have all edges present and no higher-dimensional faces?

Exercise 1.2. Not all interesting simplicial complexes are simple to visualize. Consider a finite-dimensional real vector space V and consider V to be the vertex set of a simplicial complex defined as follows: a k-simplex consists of $k+1$ linearly independent members of V. Is the resulting independence complex finite? Finite-dimensional? What does the dimension of this complex tell you?

Simplicial complexes as described are purely combinatorial objects, like the graphs they subsume. As a graph, one topologizes a simplicial complex as a quotient space built from topological simplices. The *standard* k-*simplex* is the

---

[1] The etymology of both words is salient.

following incarnation of its Platonic ideal:

$$\Delta^k = \left\{ x \in [0,1]^{k+1} \colon \sum_{i=0}^{k} x_i = 1 \right\}. \tag{1.3}$$

One topologizes an abstract simplicial complex into a space X by taking one formal copy of Δ^k for each k-simplex of X, then identifying these together along faces inductively. Define the k-*skeleton* of X, $k \in \mathbb{N}$, to be the quotient space:

$$X^{(k)} = \left(X^{(k-1)} \bigcup \coprod_{\sigma:\mathsf{dim}\,\sigma = k} \Delta^k \right) \Big/ \sim \ , \tag{1.4}$$

where $\sim$ is the equivalence relation that identifies faces of Δ^k with the corresponding combinatorial faces of σ in $X^{(j)}$ for $j < k$.

Exercise 1.5. How many total k-simplices are there in the closed n-simplex for $k < n$?

Vietoris-Rips Complexes A data set in the form of a finite metric space (X, d) gives rise to a family of simplicial complexes in the following manner. The *Vietoris-Rips complex* (or VR-complex) of (X, d) at scale $\epsilon > 0$ is the simplicial complex $\mathsf{VR}_\epsilon(X)$ whose simplices are precisely those collections of points with pairwise distance $\leqslant \epsilon$. Otherwise said, one connects points that are sufficiently close, filling in sufficiently small holes, with sufficiency specified by ϵ.

These VR complexes have been used as a way of associating a simplicial complex to point cloud data sets. One obvious difficulty, however, lies in the choice of ϵ: too small, and nothing is connected; too large, and everything is connected. The question of which ϵ to use has no easy answer. However, the perspectives of algebraic topology offer a modified question. *How to integrate structures across all* ϵ *values?* This will be considered in *Lecture Two* of this series.

Flag/clique complexes The VR complex is a particular instance of the following construct. Given a graph (network) X, the *flag complex* or *clique complex* of X is the maximal simplicial complex X that has the graph as its 1-skeleton: $X^{(1)} = X$. What this means in practice is that whenever you "see" the skeletal frame of a simplex in X, you fill it and all its faces in with simplices. Flag complexes are advantageous as data structures for spaces, in that you do not need to input/store all of the simplices in a simplicial complex: the 1-skeleton consisting of vertices and edges suffices to define the rest of the complex.

Exercise 1.6. Consider a combinatorial simplicial complex X on a vertex set of size n. As a function of this n, how difficult is it to store in memory enough information about X to reconstruct the list of its simplices? (There are several ways to approach this: see Exercise 1.5 for one approach.) Does this worst-case complexity improve if you know that X is a flag complex?

Nerve Complexes This is a particular example of a *nerve complex* associated to a collection of subsets.

Let $\mathcal{U} = \{U_\alpha\}$ be a collection of open subsets of a topological space X. The *nerve* of $\mathcal{U}$, $\mathcal{N}(\mathcal{U})$, is the simplicial complex defined by the intersection lattice of $\mathcal{U}$. The k-simplices of $\mathcal{N}(\mathcal{U})$ correspond to nonempty intersections of $k+1$ distinct elements of $\mathcal{U}$. Thus, vertices of the nerve correspond to elements of $\mathcal{U}$; edges correspond to pairs in $\mathcal{U}$ which intersect nontrivially. This definition respects faces: the faces of a k-simplex are obtained by removing corresponding elements of $\mathcal{U}$, leaving the resulting intersection still nonempty.

Exercise 1.7. Compute all possible nerves of four bounded convex subsets in the Euclidean plane. What is and is not possible? Now, repeat, but with two *nonconvex* subsets of Euclidean $\mathbb{R}^3$.

Dowker Complexes There is a matrix version of the nerve construction that is particularly relevant to applications, going back (at least) to the 1952 paper of Dowker [39]. For simplicity, let X and Y be finite sets with $\mathcal{R} \subset X \times Y$ representing the ones in a binary matrix (also denoted $\mathcal{R}$) whose columns are indexed by X and whose rows are indexed by Y. The *Dowker complex* of $\mathcal{R}$ on X is the simplicial complex on the vertex set X defined by the rows of the matrix $\mathcal{R}$. That is, each row of $\mathcal{R}$ determines a subset of X: use these to generate a simplex and all its faces. Doing so for all the rows gives the Dowker complex on X. There is a *dual Dowker complex* on Y whose simplices on the vertex set Y are determined by the ones in *columns* of $\mathcal{R}$.

Exercise 1.8. Compute the Dowker complex and the dual Dowker complex of the following relation $\mathcal{R}$:

$$\mathcal{R} = \begin{bmatrix} 1 & 0 & 0 & 0 & 1 & 1 & 0 & 0 \\ 0 & 1 & 1 & 0 & 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 & 1 & 1 & 0 & 1 \\ 1 & 0 & 1 & 0 & 1 & 0 & 0 & 1 \\ 1 & 0 & 1 & 0 & 0 & 1 & 1 & 0 \end{bmatrix}. \tag{1.9}$$

Dowker complexes have been used in a variety of social science contexts (where X and Y represent *agents* and *attributes* respectively) [7]. More recent applications of these complexes have arisen in settings ranging from social networks [89] to sensor networks [54]. The various flavors of *witness complexes* in the literature on topological data analysis [37,57] are special cases of Dowker complexes.

Cell Complexes There are other ways to build spaces out of simple pieces. These, too, are called complexes, though not simplicial, as they are not necessarily built from simplicies. They are best described as *cell complexes*, being built from cells of various dimensions sporting a variety of possible auxiliary structures.

A *cubical complex* is a cell complex built from cubes of various dimensions, the formal definition mimicking Equation (1.4): see [51,62]. These often arise as the natural model for pixel or voxel data in imagery and time series. Cubical complexes have found other uses in modelling spaces of phylogenetic trees [17, 77] and robot configuration spaces [1,53,55].

There are much more general cellular complexes built from simple pieces with far less rigidity in the gluing maps. Perhaps the most general useful model of a cell complex is the *CW complex* used frequently in algebraic topology. The idea of a CW complex is this: one begins with a disjoint union of points $X^{(0)}$ as the 0-skeleton. One then inductively defines the n-skeleton of X, $X^{(n)}$ as the $(n-1)$-skeleton along with a collection of closed n-dimensional balls, $\mathbb{D}^n$, each glued to $X^{(n-1)}$ via attaching maps on the boundary spheres $\partial\mathbb{D}^n \to X^{(n-1)}$. In dimension one, [finite] CW complexes, simplicial complexes, and cubical complexes are identical and equivalent to [finite] graphs.[2] In higher dimensions, these types of cell complexes diverge in expressivity and ease of use.

Spaces and Equivalence

Many of the spaces of interest in topological data analysis are finite metric spaces [point clouds] and simplicial approximations and generalizations of these. However, certain spaces familiar from basic calculus are relevant. We have already referenced $\mathbb{D}^n$, the closed unit n-dimensional ball in Euclidean $\mathbb{R}^n$. Its boundary defines the standard sphere $\mathbb{S}^{n-1}$ of dimension $n-1$. The 1-sphere $\mathbb{S}^1$ is also the 1-torus, where, by n-torus is meant the [Cartesian] product $\mathbb{T}^n = (\mathbb{S}^1)^n$ of n circles. The 2-sphere $\mathbb{S}^2$ and 2-torus $\mathbb{T}^2$ are compact, orientable surfaces of *genus* 0 and 1 respectively. For any genus $g \in \mathbb{N}$, there is a compact orientable surface Σ_g with that genus: for $g > 1$ these look like g 2-tori merged together so as to have the appearance of having g holes. All orientable genus g surfaces are *"topologically equivalent"*, though this is as yet imprecise.

One quickly runs into difficulty with descriptive language for spaces and equivalences, whether via coordinates or visual features. Another language is needed. Many of the core results of topology concern equivalence, detection, and resolution: are two spaces or maps between spaces qualitatively the same? This presumes a notion of equivalence, of which there are many. In what follows, *map* always implies a *continuous* function between spaces.

Homeomorphism & Homotopy A *homeomorphism* is a map $f\colon X \to Y$ with continuous inverse. This is the strongest form of topological equivalence, distinguishing spaces of different (finite) dimensions or different essential features (e.g., genus of surfaces) and also distinguishing an open from a closed interval. The more loose and useful equivalence is that generated by homotopy. A *homotopy* between maps, $f_0 \simeq f_1\colon X \to Y$ is a continuous 1-parameter family of maps

---

[2] With the exception of loop edges, which are generally not under the aegis of a graph, but are permissible in CW complexes.

$f_t \colon X \to Y$. A *homotopy equivalence* is a map $f \colon X \to Y$ with a homotopy inverse, $g \colon Y \to X$ satisfying $f \circ g \simeq \mathrm{Id}_Y$ and $g \circ f \simeq \mathrm{Id}_X$. One says that such an X and Y are *homotopic*. This is the core equivalence relation among spaces in topology.

Exercise 1.10. A space is *contractible* if it is homotopic to a point. (1) Show explicitly that $\mathbb{D}^n$ is contractible. (2) Show that $\mathbb{D}^3$ with a point in the interior removed is homotopic to $\mathbb{S}^2$. (3) Argue that the twice-punctured plane is homotopic to a "figure-eight." It's not so easy to do this with explicit maps and coordinates, is it?

Many of the core results in topology are stated in the language of homotopy (and are not true when *homotopy* is replaced with the more restrictive *homeomorphism*). For example:

Theorem 1.11. *If $\mathcal{U}$ is a finite collection of open contractible subsets of X with all non-empty intersections of subcollections of $\mathcal{U}$ contractible, then $\mathcal{N}(\mathcal{U})$ is homotopic to the union $\cup_\alpha U_\alpha$.*

Theorem 1.12. *Given any binary relation $\mathcal{R} \subset X \times Y$, the Dowker and dual Dowker complexes are homotopic.*

Homotopy invariants are central both to topology and its applications to data (noise perturbs spaces often in a non-homeomorphic but homotopic manner). Invariants of finite simplicial and cell complexes invite a computational perspective, since one has the hope of finite inputs and felicitous data structures.

Euler Characteristic The simplest nontrivial topological invariant of finite cell complexes dates back to Euler. It is elementary, combinatorial, and sublime. The *Euler characteristic* of a finite cell complex X is:

$$\chi(X) = \sum_{\sigma} (-1)^{\dim \sigma}, \tag{1.13}$$

where the sum is over all cells σ of X.

Exercise 1.14. Compute explicitly the Euler characteristics of the following cell complexes: (1) the decompositions of the 2-sphere, $\mathbb{S}^2$, defined by the boundaries of the five regular Platonic solids; (2) the CW complex having one 0-cell and one 2-cell disc whose boundary is attached directly to the 0-cell ("collapse the boundary circle to a point"); and (3) the thickened 2-sphere $\mathbb{S}^2 \times [0,1]$. How did you put a cell structure on this last 3-dimensional space?

Completion of this exercise suggests the following result:

Theorem 1.15. *Euler characteristic is a homotopy invariant among finite cell complexes.*

That this is so would seem to require a great deal of combinatorics to prove. The modern proof transcends combinatorics, making the problem hopelessly uncomputable before pulling back to the finite world, as will be seen in *Lecture Three*.

Exercise 1.16. Prove that the Euler characteristic distinguishes [connected] trees from [connected] graphs with cycles. What happens if the connectivity requirement is dropped?

Euler characteristic is a wonderfully useful invariant, with modern applications ranging from robotics [42, 47] and AI [68] to sensor networks [9–11] to Gaussian random fields [4, 6]. In the end, however, it is a numerical invariant, and has a limited resolution. The path to improving the resolution of this invariant is to enrich the underlying algebra that the Eulerian ±1 obscures.

Lifting to Linear Algebra One of the core themes of this lecture series is the lifting of cell complexes to algebraic complexes on which the tools of homological algebra can be brought to bear. This is not a novel idea: most applied mathematicians learn, e.g., to use the adjacency matrix of a graph as a means of harnessing linear-algebraic ideas to understand networks. What is novel is the use of higher-dimensional structure and the richer algebra this entails.

Homological algebra is often done with modules over a commutative ring. For clarity of exposition, let us restrict to the nearly trivial setting of finite-dimensional vector spaces over a field $\mathbb{F}$, typically either $\mathbb{R}$ or, when orientations are bothersome, $\mathbb{F}_2$, the binary field.

Given a cell complex, one lifts the topological cells to algebraic objects by using them as bases for vector spaces. One remembers the dimensions of the cells by using a sequence of vector spaces, with dimension as a *grading* that indexes the vector spaces. Consider the following sequence $\mathcal{C} = (C_\bullet)$ of vector spaces, where the grading is over $\mathbb{N}$.

$$\cdots \qquad C_k \qquad C_{k-1} \qquad \cdots \qquad C_1 \qquad C_0\ . \tag{1.17}$$

For a finite (and thus finite-dimensional) cell complex, the sequence becomes all zeros eventually. Such a sequence does not obviously offer an algebraic advantage over the original space; indeed, much of the information on how cells are glued together has been lost. However, it is easy to "lift" the Euler characteristic to this class of algebraic objects. For $\mathcal{C}$ a sequence of finite-dimensional vector spaces with finitely many nonzero terms, define:

$$\chi(\mathcal{C}) = \sum_k (-1)^k \mathsf{dim} C_k. \tag{1.18}$$

Chain Complexes Recall that basic linear algebra does not focus overmuch on vector spaces and bases; it is in linear transformations that power resides. Augmenting a sequence of vector spaces with a matching sequence of linear transformations adds in the assembly instructions and permits a fuller algebraic representation of a topological complex. Given a simplicial[3] complex X, fix a field $\mathbb{F}$

---

[3] Cell complexes in full generality can be used with more work put into the definitions of the linear transformations: see [58].

and let $\mathcal{C} = (C_k, \partial_k)$ denote the following sequence of $\mathbb{F}$-vector spaces and linear transformations.

$$\cdots \longrightarrow C_k \xrightarrow{\partial_k} C_{k-1} \xrightarrow{\partial_{k-1}} \cdots \xrightarrow{\partial_2} C_1 \xrightarrow{\partial_1} C_0 \xrightarrow{\partial_0} 0\ . \tag{1.19}$$

Each C_k has as basis the k-simplices of X. Each ∂_k restricted to a k-simplex basis element sends it to a linear combination of those basis elements in C_{k-1} determined by the $k+1$ faces of the k-simplex. This is simplest in the case $\mathbb{F} = \mathbb{F}_2$, in which case orientations can be ignored; otherwise, one must affix an orientation to each simplex and proceed accordingly: see [51,58] for details on how this is performed.

The chain complex is the primal algebraic object in homological algebra. It is rightly seen as the higher-dimensional analogue of a graph together with its adjacency matrix.

Homology Homological algebra begins with the following suspiciously simple statement about simplicial complexes.

Lemma 1.20. *The boundary of a boundary is null:*

$$\partial^2 = \partial_{k-1} \circ \partial_k = 0, \tag{1.21}$$

for all k.

Proof. For simplicity, consider the case of an abstract simplicial complex on a vertex set $V = \{v_i\}$ with chain complex having $\mathbb{F}_2$ coefficients. The *face map* D_i acts on a simplex by removing the i^{th} vertex v_i from the simplex's list, if present; else, do nothing. The graded boundary operator $\partial\colon C_\bullet \to C_\bullet$ is thus a formal sum of face maps $\partial = \bigoplus_i D_i$. It suffices to show that $\partial^2 = 0$ on each basis simplex σ. Computing the composition in terms of face maps, one obtains:

$$\partial^2 \sigma = \sum_{i \neq j} D_j D_i \sigma. \tag{1.22}$$

Each $(k-2)$-face of the k-simplex σ is represented exactly twice in the image of $D_j D_i$ over all $i \neq j$. Thanks to $\mathbb{F}_2$ coefficients, the sum over this pair is zero. □

Inspired by what happens with simplicial complexes, one defines an algebraic *complex* to be any sequence $\mathcal{C} = (C_\bullet, \partial)$ of vector spaces and linear transformations with the property that $\partial^2 = 0$. Going two steps along the sequence is the zero-map.

Exercise 1.23. Show that for any algebraic complex $\mathcal{C}$, $\mathsf{im}\ \partial_{k+1}$ is a subspace of $\mathsf{ker}\ \partial_k$ for all k.

The homology of an algebraic complex $\mathcal{C}$, $H_\bullet(\mathcal{C})$, is a complex of vector spaces defined as follows. The k-*cycles* of $\mathcal{C}$ are elements of C_k with "zero boundary", denoted $Z_k = \mathsf{ker}\ \partial_k$. The k-*boundaries* of $\mathcal{C}$ are elements of C_k that are the boundary of something in C_{k+1}, denoted $B_k = \mathsf{im}\ \partial_{k+1}$. The *homology* of $\mathcal{C}$ is the complex $H_\bullet(\mathcal{C})$ of quotient vector spaces $H_k(\mathcal{C})$, for $k \in \mathbb{N}$, given by:

$$\begin{aligned} H_k(\mathcal{C}) &= Z_k/B_k \\ &= \mathsf{ker}\ \partial_k / \mathsf{im}\ \partial_{k+1} \\ &= \text{cycles}/\,\text{boundaries}. \end{aligned} \tag{1.24}$$

Homology inherits the grading of the complex $\mathcal{C}$ and has trivial (zero) linear transformations connecting the individual vector spaces. Elements of $H_\bullet(\mathcal{C})$ are *homology classes* and are denoted $[\alpha] \in H_k$, where $\alpha \in Z_k$ is a k-cycle and $[\cdot]$ denotes the equivalence class modulo elements of B_k.

Exercise 1.25. If $\mathcal{C}$ has boundary maps that are all zero, what can you say about $H_\bullet(\mathcal{C})$? What if all the boundary maps (except at the ends) are isomorphisms (injective and surjective)?

Homology of Simplicial Complexes For X a simplicial complex, the chain complex of $\mathbb{F}_2$-vector spaces generated by simplices of X and the boundary attaching maps is a particularly simple algebraic complex associated to X. The homology of this complex is usually denoted $H_\bullet(X)$ or perhaps $H_\bullet(X; \mathbb{F}_2)$ when the binary coefficients are to be emphasized.

Exercise 1.26. Show that if X is a connected simplicial complex, then $H_0(X) = \mathbb{F}_2$. Argue that $\mathsf{dim}\ H_0(X)$ equals the number of connected components of X.

Exercise 1.27. Let X be a one-dimensional simplicial complex with five vertices and eight edges that looks like $\bigotimes$. Show that $\mathsf{dim}\ H_1(X) = 4$.

One of the important aspects of homology is that it allows one to speak of cycles that are *linearly independent*. It is true that for a graph, $\mathsf{dim}\ H_1$ is the number of *independent* cycles in the graph. One might have guessed that the number of cycles in the previous exercise is five, not four; however, the fifth can be always be expressed as a linear combination of the other four basis cycles.

Graphs have nearly trivial homology, since there are no simplices of higher dimension. Still, one gets from graphs the [correct] intuition that H_0 counts connected components and H_1 counts loops. Higher-dimensional homology measures higher-dimensional "holes" as detectable by cycles.

Exercise 1.28. Compute explicitly the $\mathbb{F}_2$-homology of the cell decompositions of the 2-sphere, $\mathbb{S}^2$, defined by the boundaries of the five regular Platonic solids. When doing so, recall that ∂_2 takes the various types of 2-cells (triangles, squares, pentagons) to the formal sum of their boundary edges, using addition with $\mathbb{F}_2$ coefficients.

These examples testify as to one of the most important features of homology.

Theorem 1.29. *Homology is a homotopy invariant.*

As stated, the above theorem would seem to apply only to cell complexes. However, as we will detail in *Lecture Three*, we can define homology for any topological space independent of cell structure; to this, as well, the above theorem applies. Thus, we can talk of the homology of a space independent of any cell structure or concrete representation: homotopy type is all that matters. It therefore makes sense to explore some basic examples. The following are the homologies of the n-dimensional sphere, $\mathbb{S}^n$; the n-dimensional torus, $\mathbb{T}^n$; and the oriented surface Σ_g of genus g.

$$\dim H_k(\mathbb{S}^n) = \begin{cases} 1 & k = n, 0 \\ 0 & k \neq n, 0 \end{cases}, \tag{1.30}$$

$$\dim H_k(\mathbb{T}^n) = \binom{n}{k}, \tag{1.31}$$

$$\dim H_k(\Sigma_g) = \begin{cases} 0 & k > 2 \\ 1 & k = 2 \\ 2g & k = 1 \\ 1 & k = 0 \end{cases}. \tag{1.32}$$

Betti Numbers We see in the above examples that the *dimensions* of the homology are the most notable features. In the history of algebraic topology, these dimensions of the homology groups — called *Betti numbers* $\beta_k = \dim H_k$ — were the first invariants investigated. They are just the beginning of the many connections to other topological invariants. For example, we will explain the following in *Lecture Three*:

Theorem 1.33. *The Euler characteristic of a finite cell complex is the alternating sum of its Betti numbers:* $\chi(X) = \sum_k (-1)^k \beta_k$.

For the moment, we will focus on applications of Betti numbers as a topological statistic. In *Lecture Two* and following, however, we will go beyond Betti numbers and consider the richer internal structure of homologies.

Application: Neuroscience

Each of these lectures ends with a sketch of some application(s): this first sketch will focus on the use of Betti numbers. Perhaps the best-to-date example of the use of homology in data analysis is the following recent work of Giusti, Pastalkova, Curto, & Itskov [56] on network inference in neuroscience using parametrized Betti numbers as a statistic.

Consider the challenge of inferring how a collection of neurons is wired together. Because of the structure of a neuron (in particular the length of the axon), mere physical proximity does not characterize the wiring structure: neurons which are far apart may in fact be “wired” together. Experimentalists can measure the responses of individual neurons and their firing sequences as a response to stimuli. By comparing time-series data from neural probes, the correlations of neuron activity can be estimated, resulting in a correlation matrix with entries, say, between zero and one, referencing the estimated correlation between neurons, with the diagonal, of course, consisting of ones. By thresholding the correlation matrix at some value, one can estimate the "wiring network" of how neurons are connected.

Unfortunately, things are more complicated than this simple scenario suggests. First, again, the problem of which threshold to choose is present. Worse, the correlation matrix is not the truth, but an experimentally measured estimation that relies on how the experiment was performed (Where were the probes inserted? How was the spike train data handled?). Repeating an experiment may lead to a very different correlation matrix – a difference not accountable by a linear transformation. This means, in particular, that methods based on spectral properties such as PCA are misleading [56].

What content does the experimentally-measured correlation matrix hold? The entries satisfy an *order principle*: if neurons A and B seem more correlated than C and D, then, in truth, they are. In other words, repeated experiments lead to a nonlinear, but order-preserving, homeomorphism of the correlation axis. It is precisely this nonlinear coordinate-free nature of the problem that prompts a topological approach.

The approach is this. Given a correlation matrix $\mathcal{R}$, let $1 \geqslant \epsilon \geqslant 0$ be a decreasing threshold parameter, and, for each ϵ, let $\mathcal{R}_\epsilon$ be the binary matrix generated from $\mathcal{R}$ with ones wherever the correlation exceeds ϵ. Let X_ϵ be the Dowker complex of $\mathcal{R}_\epsilon$ (or dual; the same, by symmetry). Then consider the k^{th} Betti number distribution $\beta_k : [1, 0] \to \mathbb{N}$. These distributions are unique under change of correlation axis coordinates up to order-preserving homeomorphisms of the domain.

What do these distributions look like? For $\epsilon \to 1$, the complex is an isolated set of points, and for $\epsilon \to 0$ it is one large connected simplex: all the interesting homology lies in the middle. It is known that for such an increasing sequence of simplicial complexes, the Betti distributions β_k for $k > 0$ are unimodal. Furthermore, it is known that homological peaks are ordered by dimension [5]: the peak ϵ value for β_1 precedes that of β_2, etc. Thus, what is readily available as a signature for the network is the ordering of the heights of the peaks of the β_k distributions.

The surprise is that one can distinguish between networks that are wired *randomly* versus those that are wired *geometrically*. This is motivated by the neuroscience applications, since it has been known since the Nobel prize-winning work of O'Keefe et al. that certain neurons in the visual cortex of rats act as *place cells*, encoding the geometry of a learned domain (e.g., a maze) by how the neurons are wired [74], in manner not unlike that of a nerve complex [35]. Other neural networks are known to be wired together randomly, such as the olfactory system of a fly [29]. Giusti et al., relying on theorems about Betti number distributions for random geometric complexes by Kahle [63], show that one can differentiate between geometrically-wired and randomly wired networks by looking at the peak signatures of β_1, β_2, and β_3 and whether the peaks increase [random] or decrease [geometric]. Follow-on work gives novel signature types [87]. The use of these methods is revolutionary, since actual physical experiments to rigorously determine neuron wiring are prohibitively difficult and expensive, whereas computing homology is, in principle, simple. *Lecture Three* will explore this issue of computation more.

Lecture 2: Persistence

We have covered the basic definitions of simplicial and algebraic complexes and their homological invariants. Our goal is to pass from the mechanics of invariants to the principles that animate the subject, culminating in a deeper understanding of how data can be qualitatively compressed and analyzed. In this lecture, we will begin that process, using the following principles as a guide:

1. A simplicial [or cell] complex is the right type of discrete data structure for capturing the significant features of a space.
2. A chain complex is a linear-algebraic representation of this data structure – an algebraic set of assembly instructions.
3. To prove theorems about how cell complexes behave under deformation, study instead deformations of chain complexes.
4. Homology is the optimal compression of a chain complex down to its qualitative features.

Towards Functoriality

Our chief end is this: homology is *functorial*. This means that one can talk not only about homology of a complex, but also of the homology of a map between complexes. To study continuous maps between spaces algebraically, one translates the concept to chain complexes. Assume that X and Y are simplicial complexes and $f\colon X \to Y$ is a *simplicial map* – a continuous map taking simplices to simplices.[4] This does not imply that the simplices map homeomorphically to simplices of the same dimension.

In the same manner that X and Y lift to algebraic chain complexes $C_\bullet(X)$ and $C_\bullet(Y)$, one lifts f to a graded sequence of linear transformations $f_\bullet\colon C_\bullet(X) \to C_\bullet(Y)$, generated by basis n-simplices of X being sent to basis n-simplices of Y, where, if an n-simplex of X is sent by f to a simplex of dimension less than n, then the algebraic effect is to send the basis chain in $C_n(X)$ to $0 \in C_n(Y)$. The continuity of the map f induces a *chain map* $f_\bullet$ that fits together with the boundary maps of $C_\bullet(X)$ and $C_\bullet(Y)$ to form the following diagram of vector spaces and

---

[4] For cell complexes, one makes the obvious adjustments.

linear transformations:

$$\begin{tikzcd}
\cdots \arrow[r] & C_{n+1}(X) \arrow[r, "\partial"] \arrow[d, "f_\bullet"] & C_n(X) \arrow[r, "\partial"] \arrow[d, "f_\bullet"] & C_{n-1}(X) \arrow[r, "\partial"] \arrow[d, "f_\bullet"] & \cdots . \\
\cdots \arrow[r] & C_{n+1}(Y) \arrow[r, "\partial"] & C_n(Y) \arrow[r, "\partial"] & C_{n-1}(Y) \arrow[r, "\partial"] & \cdots
\end{tikzcd} \tag{2.1}$$

Commutative Diagrams Equation (2.1) is important – it is our first example of what is known as a *commutative diagram*. These are the gears for algebraic engines of inference. In this example, commutativity means precisely that the chain maps respect the boundary operation, $f_\bullet\partial = \partial f_\bullet$. *This is what continuity means for linear transformations of complexes.* There is no need for simplices or cells to be explicit. One defines a *chain map* to be any sequence of linear transformations $f_\bullet \colon \mathcal{C} \to \mathcal{C}'$ on algebraic complexes making the diagram commutative.

Exercise 2.2. Show that any chain map $f_\bullet \colon \mathcal{C} \to \mathcal{C}'$ takes cycles to cycles and boundaries to boundaries.

Induced Homomorphisms Because of this commutativity, a chain map $f_\bullet$ acts not only on chains but on cycles and boundaries as well. This makes well-defined the *induced homomorphism* $H(f) \colon H_\bullet(\mathcal{C}) \to H_\bullet(\mathcal{C}')$ on homology. For α a cycle in $\mathcal{C}$, recall that its homology class is denoted $[\alpha]$. One may thus define $H(f)[\alpha] = [f_\bullet\alpha] = [f \circ \alpha]$. This is well-defined: if $[\alpha] = [\alpha']$, then, as chains, $\alpha' = \alpha + \partial\beta$ for some β, and,

$$f_\bullet\alpha' = f \circ \alpha' = f \circ (\alpha + \partial\beta) = f \circ \alpha + f \circ \partial\beta = f_\bullet\alpha + \partial(f_\bullet\beta), \tag{2.3}$$

so that $H(f)[\alpha'] = [f_\bullet\alpha'] = [f_\bullet\alpha] = H(f)[\alpha]$ in $H_\bullet(\mathcal{C}')$.

The term *homomorphism* is used to accustom the reader to standard terminology. Of course, in the present context, an induced homomorphism is simply a graded linear transformation on homology induced by a chain map.

Exercise 2.4. Consider the disc in $\mathbb{R}^2$ of radius π punctured at the integer points along the x and y axes. Although this space is not a cell complex, let us assume that its homology is well-defined and is "the obvious thing" for H_1, defined by the number of punctures. What are the induced homomorphisms on H_1 of the continuous maps given by (1) rotation by $\pi/2$ counterclockwise; (2) the folding map $x \mapsto |x|$; (3) flipping along the y axis?

Functoriality Homology is functorial, meaning that the induced homomorphisms on homology are an algebraic reflection of the properties of continuous maps between spaces. The following are simple properties of induced homomorphisms, easily shown from the definitions above:

- Given a chain map $f_\bullet \colon \mathcal{C} \to \mathcal{C}'$, $H(f) \colon H_\bullet(\mathcal{C}) \to H_\bullet(\mathcal{C}')$ is a (graded) sequence of linear transformations.
- The identity transformation $\mathsf{Id} \colon \mathcal{C} \to \mathcal{C}$ induces the identity isomorphism $\mathsf{Id} \colon H_\bullet(\mathcal{C}) \to H_\bullet(\mathcal{C})$.

- Given $f_\bullet \colon \mathcal{C} \to \mathcal{C}'$ and $g_\bullet \colon \mathcal{C}' \to \mathcal{C}''$, $H(g \circ f) = H(g) \circ H(f)$.

There is hardly a more important feature of homology than this functoriality.

Exercise 2.5. Show using functoriality that homeomorphisms between spaces induce isomorphisms on homologies.

Exercise 2.6. Can you find explicit counterexamples to the following statements about maps f between simplicial complexes and their induced homomorphisms $H(f)$ (on some grading for homology)?

(1) If f is surjective then $H(f)$ is surjective.
(2) If f is injective then $H(f)$ is injective.
(3) If f is not surjective then $H(f)$ is not surjective.
(4) f is not injective then $H(f)$ is not injective.
(5) f is not bijective then $H(f)$ is not bijective.

Functorial Inference It is sometimes the case that what is desired is knowledge of the qualitative features [homology] of an important but unobservable space X; what is observed is an approximation Y to X, of uncertain homological fidelity. One such observation is unhelpful. Two or more homological samplings may lead to increased confidence; however, functoriality can relate observations to truth. Suppose the observed data comprises the homology of a *pair* of spaces Y_1, Y_2, which are related by a map $f \colon Y_1 \to Y_2$ that factors through a map to X, so that $f = f_2 \circ f_1$ with $f_1 \colon Y_1 \to X$ and $f_2 \colon X \to Y_2$. If the induced homomorphism $H(f)$ is known, then, although $H_\bullet(X)$ is hidden from view, inferences can be made.

$$\begin{tikzcd}
H_\bullet Y_1 \arrow[rr, "H(f)"] \arrow[dr, dotted, "H(f_1)"'] & & H_\bullet Y_2 \, . \\
& H_\bullet X \arrow[ur, dotted, "H(f_2)"'] &
\end{tikzcd} \tag{2.7}$$

Exercise 2.8. In the above scenario, what can you conclude about $H_\bullet(X)$ if $H(f)$ is an isomorphism? If it is merely injective? Surjective?

The problem of measuring topological features of experimental data by means of sensing is particularly vulnerable to threshold effects. Consider, e.g., an open tank of fluid whose surface waves are experimentally measured and imaged. Perhaps the region of interest is the portion of the fluid surface above the ambient height $h = 0$; the topology of the set $A = \{h \geqslant 0\}$ must be discerned, but can only be approximated by imprecise pixellated images of $\{h \gtrsim 0\}$. One can choose a measurable threshold above and below the zero value to get just such a situation as outlined above. Similar scenarios arise in MRI data, where the structure of a tissue of interest can be imaged as a pair of pixellated approximations, known to over- and under-approximate the truth.

Sequences

Induced homomorphisms in homology are key, as central to homology as the role of linear transformations are in linear algebra. In these lectures, we have seen how, though linear transformations between vector spaces are important, what a great advantage there is in the chaining of linear transformations into sequences and complexes. The advent of induced homomorphisms should prompt the same desire, to chain into sequences, analyze, classify, and infer. This is the plan for the remainder of this lecture as we outline the general notions of persistence, persistent homology, and topological data analysis.

Consider a sequence of inclusions of subcomplexes $\iota\colon X_k \subset X_{k+1}$ of a simplicial complex X for $1 \leqslant k \leqslant N$. These can be arranged into a sequence of spaces with inclusion maps connecting them like so:

$$\varnothing = X_0 \xrightarrow{\iota} X_1 \xrightarrow{\iota} \cdots \xrightarrow{\iota} X_{N-1} \xrightarrow{\iota} X_N \xrightarrow{\iota} X. \tag{2.9}$$

Sequences of spaces are very natural. One motivation comes from a sequence of Vietoris-Rips complexes of a set of data points with an increasing sequence of radii $(\epsilon_i)_{i=1}^N$.

Exercise 2.10. At the end of *Lecture 1*, we considered a correlation matrix $\mathcal{R}$ on a set V of variables, where correlations are measured from 0 to 1 and used this matrix to look at a sequence of Betti numbers. Explain how to rephrase this as a sequence of homologies with maps (assuming some discretization along the correlation axis). What maps induce the homomorphisms on homologies? What do the homologies look like for very large or very small values of the correlation parameter?

A topological sequence of spaces is converted to an algebraic sequence by passing to homology and using induced homomorphisms:

$$H_\bullet(X_0) \xrightarrow{H(\iota)} H_\bullet(X_1) \xrightarrow{H(\iota)} \cdots \xrightarrow{H(\iota)} H_\bullet(X_{N-1}) \xrightarrow{H(\iota)} H_\bullet(X_N) \xrightarrow{H(\iota)} H_\bullet(X). \tag{2.11}$$

The individual induced homomorphisms on homology encode local topological changes in the X_i; thanks to functoriality, the *sequence* encodes the *global* changes.

Exercise 2.12. Consider a collection of 12 equally-spaced points on a circle — think of tick-marks on a clock. Remove from this all the points corresponding to the prime numbers (2, 3, 5, 7, 11). Use the remaining points on the circle as the basis of a sequence of Vietoris-Rips [VR] complexes based on an increasing sequence $\{\epsilon_i\}$ of distances starting with $\epsilon_0 = 0$. Without worrying about the actual values of the ϵ_i, describe what happens to the sequence of VR complexes. What do you observe? Does H_0 ever increase? Decrease? What about H_1?

What one observes from this example is the evolution of homological features over a sequence: homology classes are born, can merge, split, die, or persist. This evolutionary process as written in the language of sequences is the algebraic means of encoding notions of geometry, significance, and noise.

Persistence Let us formalize some of what we have observed. Consider a sequence of spaces (X_i) and continuous transformations $f_i\colon X_i \to X_{i+1}$, without requiring subcomplexes and inclusions. We again have a sequence of homologies with induced homomorphisms. A homology class in $H_\bullet(X_i)$ is said to *persist* if its image in $H_\bullet(X_{i+1})$ is also nonzero; otherwise it is said to *die*. A homology class in $H_\bullet(X_j)$ is said to be *born* when it is not in the image of $H_\bullet(X_{j-1})$.

One may proceed with this line of argument, at the expense of some sloppiness of language. Does every homology class have an unambiguous birth and death? Can we describe cycles this way, or do we need to work with classes of cycles modulo boundaries? For the sake of precision and clarity, it is best to follow the pattern of these lectures and pass to the context of linear algebra and sequences.

Consider a sequence $V_\bullet$ of finite-dimensional vector spaces, graded over the integers $\mathbb{Z}$, and stitched together with linear transformations like so:

$$V_\bullet = \quad \cdots \longrightarrow V_{i-1} \longrightarrow V_i \longrightarrow V_{i+1} \longrightarrow \cdots . \tag{2.13}$$

These *sequences* are more general than algebraic *complexes*, which must satisfy the restriction of composing two incident linear transformations yielding zero. Two such sequences $V_\bullet$ and $V'_\bullet$ are said to be *isomorphic* if there are isomorphisms $V_k \cong V'_k$ which commute with the linear transformations in $V_\bullet$ and $V'_\bullet$ as in Equation (2.1). The simplest such sequence is an *interval indecomposable* of the form

$$I_\bullet = \quad \cdots \longrightarrow 0 \longrightarrow 0 \longrightarrow \mathbb{F} \xrightarrow{\mathsf{Id}} \mathbb{F} \xrightarrow{\mathsf{Id}} \cdots \xrightarrow{\mathsf{Id}} \mathbb{F} \longrightarrow 0 \longrightarrow 0 \longrightarrow \cdots , \tag{2.14}$$

where the *length* of the interval equals the number of Id maps, so that an interval of length zero consists of $0 \to \mathbb{F} \to 0$ alone. Infinite or bi-infinite intervals are also included as indecomposables.

Representation Theory A very slight amount of representation theory is all that is required to convert a sequence of homologies into a useful data structure for measuring persistence. Consider the following operation: sequences can be formally added by taking the direct sum, $\oplus$, term-by-term and map-by-map. The interval indecomposables are precisely *indecomposable* with respect to $\oplus$ and cannot be expressed as a sum of simpler sequences, even up to isomorphism. The following theorem, though simple, is suitable for our needs.

Theorem 2.15 (Structure Theorem for Sequences)**.** *Any sequence of finite-dimensional vector spaces and linear transformations decomposes as a direct sum of interval indecomposables, unique up to reordering.*

What does this mean? It's best to begin with the basics of linear algebra, and then see how that extends to homology.

Exercise 2.16. Any linear transformation $\mathbb{R}^n \xrightarrow{A} \mathbb{R}^m$ extends to a biinfinite sequence with all but two terms zero. How many different isomorphism classes of decompositions into interval indecomposables are there? What types of intervals are present? Can you interpret the numbers of the various types of intervals? What well-known theorem from elementary linear algebra have you recovered?

Barcodes. When we use field coefficients, applying the Structure Theorem to a sequence of homologies gives an immediate clarification of how homology classes evolve. Homology classes correspond to interval indecomposables, and are born, persist, then die at particular (if perhaps infinite) parameter values. This decomposition also impacts how we illustrate evolving homology classes. By drawing pictures of the interval indecomposables over the [discretized] parameter line as horizontal *bars*, we obtain a pictograph that is called a *homology barcode*.

Exercise 2.17. Consider a simple sequence of four vector spaces, each of dimension three. Describe and/or draw pictures of all possible barcodes arising from such a sequence. Up to isomorphism, how many such barcodes are there?

The phenomena of homology class *birth*, *persistence*, and *death* corresponds precisely to the *beginning*, *middle*, and *end* of an interval indecomposable. The barcode is usually presented with horizontal intervals over the parameter line corresponding to interval indecomposables. Note that barcodes, like the homology they illustrate, are graded. There is an H_k-barcode for each $k \geqslant 0$. Since, from the Structure Theorem, the order does not matter, one typically orders the bars in terms of birth time (other orderings are possible).

The barcode provides a simple descriptor for topological significance: the shorter an interval, the more ephemeral the hole; long bars indicate robust topological features with respect to the parameter. This is salient in the context of point clouds $\mathcal{Q}$ and Vietoris-Rips complexes $\mathsf{VR}_\epsilon(\mathcal{Q})$ using an increasing sequence $\{\epsilon_i\}$ as parameter. For ϵ too small or too large, the homology of $\mathsf{VR}_\epsilon(\mathcal{Q})$ is unhelpful. Instead of trying to choose an *optimal* ϵ, choose them *all*: the barcode reveals significant features.

Exercise 2.18. Persistent homology is useful and powerful in topological data analysis, but sometimes one can get lost in the equivalence relation that comprises homology classes. Often, in applications, one cares less about the homology and more about a particular *cycle* (whose homology class may be too loose to have meaning within one's data). Given a sequence of chain complexes and chain maps, what can be said about *persistent cycles* and *persistent boundaries*? Are these well-defined? Do they have barcodes? How would such structures relate to persistent homology barcodes?

Persistent Homology Let us summarize what we have covered with slightly more formal terminology. A *persistence complex* is a sequence of chain complexes $\mathcal{P} = (\mathcal{C}_i)$, together with chain maps $x\colon \mathcal{C}_i \longrightarrow \mathcal{C}_{i+1}$. For notational simplicity, the index subscripts on the chain maps x are suppressed. Note that each $\mathcal{C}_i = (C_{\bullet,i}, \partial)$ is itself a complex: we have a sequence of sequences. The *persistent homology* of a persistence complex $\mathcal{P}$ is not a simple homology theory, but rather a homology associated to closed intervals in the "parameter domain". Over the interval $[i, j]$, its persistent homology, denoted $H_\bullet(\mathcal{P}[i, j])$, is defined to be the image of the induced homomorphism $H(x^{j-i})\colon H_\bullet(\mathcal{C}_i) \to H_\bullet(\mathcal{C}_j)$ induced by x^{j-i}. That is, one looks at the composition of the chain maps from $\mathcal{C}_i \to \mathcal{C}_j$ and takes the image of the induced homomorphism on homologies. This persistent homology consists of homology classes that *persist*: $\dim H_k(\mathcal{P}[i,j])$ equals the number of intervals in the barcode of $H_k(\mathcal{P})$ containing the parameter interval $[i,j]$.

Exercise 2.19. If, in the indexing for a persistence complex, you have $i < j < k < \ell$, what is the relationship between the various subintervals of $[i,\ell]$ using $\{i,j,k,\ell\}$ as endpoints? Draw the lattice of such intervals under inclusion. What is the relationship between the persistent homologies on these subintervals?

Persistence Diagrams. Barcodes are not the only possible graphical presentation for persistent homology. Since there is a decomposition into homology classes with well-defined initial and terminal parameter values, one can plot each homology class as a point in the plane with axes the parameter line. To each interval indecomposable (homology class) one assigns a single point with coordinates *(birth, death)*. This scatter plot is called the *persistence diagram* and is more practical to plot and interpret than a barcode for very large numbers of homology classes.

Exercise 2.20. In the case of a homology barcode coming from a data set, the "noisy" homology classes are those with the smallest length, with the largest bars holding claim as the "significant" topological features in a data set. What do these noisy and significant bars translate to in the context of a persistence diagram? For a specific example, return to the "clockface" data set of Exercise 2.12, but now consider the set of all even points: 2, 4, 6, 8, 10, and 12. Show that the persistent H_2 contains a "short" bar. Are you surprised at this *artificial bubble* in the VR complex? Does a similar bubble form in the homology when all 12 points are used? In which dimension homology?

One aspect worth calling out is the notion of persistent homology as a homological data structure *over* the parameter space, in that one associates to each interval $[i,j]$ its persistent homology. This perspective is echoed in the early literature on the subject [32,36,41,91], in which a continuous parameter space was used, with a continuous family of (excursion sets) of spaces X_t, $t \in \mathbb{R}$, was used: in this setting, persistent homology is assigned to an interval $[s,t]$. The discretized parameter interval offers little in the way of restrictions (unless you are working with fractal-like or otherwise degenerate objects) and opens up the simple setting of the Structure Theorem on Sequences as used in this lecture.

Stability

The idea behind the use of barcodes and persistence diagrams in data is grounded in the intuition that essential topological features of a domain are robust to noise, whether arising from sensing, sampling, or approximation. In a barcode, noisy features appear as short bars; in a persistence diagram, as near-diagonal points. To solidify this intuition of robustness, one wants a more specific statement on the *stability* of persistent homology. Can a small change in the input — whether a sampling of points or a Dowker relation or a perturbation of the metric — have a large impact on how the barcode appears?

There have of late been a plethora of stability theorems in persistent homology, starting with the initial result of Cohen-Steiner et al. [31] and progressing to more general and categorical forms [12,18,21,22,30]. In every one of these settings, the stability result is given in the context of persistence over a continuous parameter, ϵ, such as one might use in the case of a Vietoris-Rips filtration on a point-cloud. The original stability theorem is further framed in the setting of sublevel sets of a function $h\colon X \to \mathbb{R}$ and the filtration is by sublevel sets $X_t = \{h \leqslant t\}$. Both the statements of the stability theorems and their proofs are technical; yet the technicalities lie in the difficulties of the continuum parameter. For the sake of clarity and simplicity, we will assume that one has imposed a uniform discretization of the real line with step size a fixed $\epsilon > 0$ (as would often be the case in practice).

Interleaving. At present, the best language for describing the stability of persistent homology and barcode descriptors is the recent notion of *interleaving*. In keeping with the spirit of these lectures, we will present the theory in the context of sequences of vector spaces and linear transformations. Assume that one has a pair of sequences, $V_\bullet$ and $W_\bullet$, of vector spaces and linear transformations. We say that a T-*interleaving* is a pair of degree-T mappings:

$$f_\bullet\colon V_\bullet \to W_{\bullet+T} \qquad g_\bullet\colon W_\bullet \to V_{\bullet+T}, \tag{2.21}$$

such that the diagram commutes:

$$\begin{tikzcd}
\cdots \arrow[r] & V_n \arrow[r] \arrow[drr, "f"' near end] & \cdots \arrow[r] & V_{n+T} \arrow[r] \arrow[drr, "f" near start] & \cdots \arrow[r] & V_{n+2T} \arrow[r] & \cdots . \\
\cdots \arrow[r] & W_n \arrow[r] \arrow[urr, "g" near end] & \cdots \arrow[r] & W_{n+T} \arrow[r] \arrow[urr, "g"' near start] & \cdots \arrow[r] & W_{n+2T} \arrow[r] & \cdots
\end{tikzcd} \tag{2.22}$$

In particular, at each n, the composition $g_{n+T} \circ f_n$ equals the composition of the 2T horizontal maps in $V_\bullet$ starting at n; and, likewise with $f \circ g$ on $W_\bullet$. One defines the *interleaving distance* between two sequences of vector spaces to be the minimal $T \in \mathbb{N}$ such that there exists a T-interleaving.

Exercise 2.23. Verify that the interleaving distance of two sequences is zero if and only if the two sequences are isomorphic.

Exercise 2.24. Assume that $V_\bullet$ is an interval indecomposable of length 5. Describe the set of all $W_\bullet$ that are within interleaving distance one of $V_\bullet$.

Note that the conclusions of Exercises 2.23-2.24 are absolutely dependent on the discrete nature of the problem. In the case of a continuous parameter, the interleaving distance is *not* a metric on persistence complexes, but is rather a *pseudometric*, as one can have the infimum of discretized interleaving distances become zero without an isomorphism in the limit.

Application: TDA

Topological Data Analysis, or TDA, is the currently popular nomenclature for the set of techniques surrounding persistence, persistent homology, and the extraction of significant topological features from data. The typical input to such a problem is a point cloud $\mathcal{Q}$ in a Euclidean space, though any finite metric space will work the same. Given such a data set, assumed to be a noisy sampling of some domain of interest, one wants to characterize that domain. Such questions are not new: linear regression assumes an affine space and returns a best fit; a variety of locally-linear or nonlinear methods look for nonlinear embeddings of Euclidean spaces.

Topological data analysis looks for global structure — homology classes — in a manner that is to some degree decoupled from rigid geometric considerations. This, too, is not entirely novel. Witness clustering algorithms, which take a point cloud and return a partition that is meant to approximate *connected components*. Of course, this reminds one of H_0, and the use of a Vietoris-Rips complex makes this precise: *single linkage clustering* is precisely the computation of $H_0(\mathsf{VR}_\epsilon(\mathcal{Q}))$ for a choice of $\epsilon > 0$. Which choice is best? The lesson of persistence is to take all ϵ and build the homology barcode. Notice however, that the barcode returns only the dimension of H_0 — the number of clusters — and to more carefully specify the clusters, one needs an appropriate basis. There are many other clustering schemes with interesting functorial interpretations [26,27].

The ubiquity and utility of clustering is clear. What is less clear is the prevalence and practicality of higher-dimensional persistent homology classes in "organic" data sets. Using again a Vietoris-Rips filtration of simplicial complexes on a point cloud $\mathcal{Q}$ allows the computation of homology barcodes in gradings larger than zero. To what extent are they prevalent? The grading of homology is reminiscent of the grading of polynomials in Taylor expansions. Though Taylor expansions are undoubtedly useful, it is acknowledged that the lowest-order terms (zeroth and first especially) are most easily seen and used. Something like this holds in TDA, where one most readily sees clusters (H_0) and simple loops (H_1) in data. The following is a brief list of applications known to the author. The literature on TDA has blown-up of late to a degree that makes it impossible to give an exhaustive account of applications. The following are chosen as illustrative of the basic principles of persistent homology.

Medical imaging data: Some of the earliest and most natural applications of TDA were to image analysis [2,13,25,79]. One recent study by Benditch et al. looks at the structure of arteries in human brains [14]. These are highly convoluted pathways, with lots of branching and features at multiple scales, but which vary in dramatic and unpredictable ways from patient to patient. The topology of arterial structures are globally trivial — sampling the arterial structure though standard imaging techniques yields a family of trees (acyclic graphs). Nevertheless, since the geometry is measurable, one can filter these trees by sweeping a plane across the three-dimensional ambient domain, and look at the persistent H_0. Results show statistically significant correlations between the vector of lengths of the top 100 bars in the persistent H_0 barcode and features such as patient age and sex. For example, older brains tend to have shorter *longest bars* in the H_0 barcode. The significance of the correlation is very strong and outperforms methods derived from graph theory and phylogenetic-based tree-space geometry methods. It is interesting that it is not the "longest bar" that matters so much as the ensemble of longest bars in this barcode. Work in progress includes using H_0 barcode statistics to characterize global structure of graph-like geometries include examinations of insect wing patterns, tree leaf vein networks, and more. Other exciting examples of persistent H_0 to medical settings feature an analysis of breast cancer by Nicolau et al. [73].

Distinguishing illness from health and recovey: Where does persistent homology beyond H_0 come into applications? A recent excellent paper of Torres et al. uses genetic data of individuals with illnesses to plot a time series of points in a *disease space* of traits [88]. Several examples are given of studies on human and mouse subjects tracking the advancement and recovery from disease (including, in the mice, malaria). Genetic data as a function of time and for many patients gives a point cloud in an abstract space for which geometry is not very relevant (for example, axes are of differing and incomparable units). What is interesting about this study is the incorporation of data from subjects that extends from the onset of illness, through its progress, and including a full recovery phase back to health. Of interest is the question of recovery — does recovery from illness follow the path of illness in reverse? Does one recover to the same state of health, or is there a monodromy? The study of Torres et al. shows a single clear long-bar in the H_1-barcode in disease space, indicating that all instances of the illness are *homologous*, as are all instances of recovery, but that *illness and recovery are homologically distinct events*" [88]. In contrast to prior studies that performed a linear regression on a pair of variables and concluded a linear relationship between these variables (with a suggestion of a causal relationship), the full-recovery data set with its loopy phenomenon of recovery suggests skepticism: indeed, a careful projection of a generator for the H_1 barcode into this plane recovers the loop.

Robot path planning: A very different set of applications arises in robot motion planning, in which an autonomous agent needs to navigate in a domain $X \subset \mathbb{R}^n$ (either physical or perhaps a configuration space of the robot) from an initial state to a goal state in X. In the now-familiar case of self-driving vehicles, autonomous drones, or other agents with sensing capabilities, the navigable ("obstacle-free") subset of X is relatively uncertain, and known only as a probabilistic model. Bhattacharya et al. consider such a probability density $\rho\colon X \to [0,\infty)$ and use a combination of persistent homology and graph search-based algorithms in order to compute a set of best likely paths-to-goal [15]. The interesting aspects of this application are the following. (1) The persistence parameter is the probability, used to filter the density function on X. This is one natural instance of a continuous as opposed to a discrete persistence parameter. (2) The H_1 homology barcode is used, but on a slight modification of X obtained by abstractly identifying the initial and goal states (by an outside edge if one wants to be explicit), so that a path from initial to goal corresponds precisely to a 1-cycle that intersects this formal edge. (3) Once again, the problem of "where to threshhold" the density is avoided by the use of a barcode; the largest bars in the H_1 barcode correspond to the path classes most likely to be available and robust to perturbations in the sensing. For path classes with short bars, a slight update to the system might invalidate this path.

Localization and mapping: Computing persistent homology is also useful as a means of building topological maps of an unknown environment, also of relevance to problems in robotics, sensing, and localization. Imagine a lost traveler wandering through winding, convoluted streets in an unfamiliar city with unreadable street signs. Such a traveler might use various landmarks to build up an internal map: *"From the jewelry store, walk toward the cafe with the red sign, then look for the tall church steeple."* This can be accomplished without reference to coordinates or odometry. For the general setting, assume a domain $\mathcal{D}$ filled with landmarks identifiable by observers that register landmarks via local sensing/visibility. A collection of observations are taken, with each observation recording only those landmarks "visible" from the observation point. Both the landmarks and observations are each a discrete set, with no geometric or coordinate data appended. The sensing data is given in the form of an unordered sequence of pairs of observation-landmark identities encoding who-sees-what. From this abstract data, one has a Dowker relation from which one can build a pair of (dual, homotopic) Dowker complexes that serve as approximations to the domain topology. In [54], two means of inferring a topological map from persistent homology are given. (1) If observers record in the sensing relation a visibility strength (as in, say, the strengths of signals to all nearby wireless SSIDs), then filtering the Dowker complexes on this (as we did in the neuroscience applications of the previous lecture) gives meaning to long bars as significant map features. (2) In the binary (hidden/seen) sensing case, the existence of non-unique landmarks (*"ah, look! a Starbucks! I know exactly where I am now..."*) confounds the topology, but filtering according to witness weight can eliminate spurious simplices: see [54] for details and [38] for an experimental implementation.

Protein compressibility: One of the first uses of persistent H_2 barcodes has appeared recently in the work of Gameiro et al. [48] in the context of characterizing compressibility in certain families of protein chains. Compressibility is a particular characterization of a protein's *softness* and is key to the interface of structure and function in proteins, the determination of which is a core problem in biology. The experimental measurement of protein compressibility is highly nontrivial, and involves fine measurement of ultrasonic wave velocities from pressure waves in solutions/sovlents with the protein. Hollow cavities within the protein's geometric conformation are contributing factors, both size and number. Gameiro et al. propose a *topological compressibility* that is argued to measure the relative contributions of these features, but with minimal experimental measurement, using nothing more as input than the standard molecular datasets that record atom locations as a point cloud, together with a van der Waals radius about each atom. What is interesting in this case is that one does not have the standard Vietoris-Rips filtered complex, but rather a filtered complex obtained by starting with the van der Waals radii (which vary from atom to atom) and then adding to these radii the filtration parameter $\epsilon > 0$. The proposed topological compressibility is a ratio of the number of persistent H_2 intervals divided by the number of persistent H_1 intervals (where the intervals are restricted to certain parameter ranges). This ratio is meant to serve as proxy to the experimental measurement of cavities and tunnels in the protein's structure. Comparisons with experimental data suggest, with some exceptions, a tight linear correlation between the expensive experimentally-measured compressibility and the (relatively inexpensive) topological compressibility.

These varied examples are merely summaries: see the cited references for more details. The applications of persistent homology to data are still quite recent, and by the time of publication of these notes, there will have been a string of novel applications, ranging from materials science to social networks and more.

Lecture 3: Compression & Computation

We now have in hand the basic tools for topological data analysis: complexes, homology, and persistence. We are beginning to develop the theories and perspectives into which these tools fit, as a higher guide to how to approach qualitative phenomena in data. We have not yet dealt with issues of computation and effective implementation. Our path to doing so will take us deeper into sequences, alternate homology theories, cohomology, and Morse theory.

Sequential Manipulation

Not surprisingly, these lectures take the perspective that the desiderata for homological data analysis include a calculus for complexes and sequences of complexes. We have seen hints of this in *Lecture Two*; now, we proceed to introduce a bit more of the rich structure that characterizes (the near-trivial linear-algebraic version of) homological algebra. Instead of focusing on spaces or simplicial complexes per se, we focus on algebraic complexes; this motivates our examination of certain types of algebraic sequences.

Exact Complexes Our first new tool is inspired by the following question: among all complexes, which are simplest? Simplicial complexes might suggest that the simplest sort of complex is that of a single simplex, which has homology vanishing in all gradings except zero. However, there are simpler complexes still. We say that an algebraic complex $\mathcal{C} = (C_\bullet, \partial)$ is *exact* if its homology completely vanishes, $H_\bullet(\mathcal{C}) = 0$. This often written termwise as:

$$\ker \partial_k = \operatorname{im} \partial_{k+1} \quad \forall k. \tag{3.1}$$

Exercise 3.2. What can you say about the barcode of an exact complex? (This means the barcode of the complex, not its [null] homology).

The following simple examples of exact complexes help build intuition:

- Two vector spaces are isomorphic, $V \cong W$, iff there is an exact complex of the form:
$$0 \longrightarrow V \longrightarrow W \longrightarrow 0\ .$$
- The 1st Isomorphism Theorem says that for a linear transformation φ of V, the following sequence is exact:
$$0 \longrightarrow \ker \varphi \longrightarrow V \xrightarrow{\varphi} \operatorname{im} \varphi \longrightarrow 0\ .$$
Such a 5-term complex framed by zeroes is called a *short exact complex*. In any such short exact complex, the second map is injective; the penultimate, surjective.
- More generally, the kernel and cokernel of a linear transformation $\varphi\colon V \to W$ fit into an exact complex:
$$0 \longrightarrow \ker \varphi \longrightarrow V \xrightarrow{\varphi} W \longrightarrow \operatorname{coker} \varphi \longrightarrow 0\ .$$

Exercise 3.3. Consider $C = C^\infty(\mathbb{R}^3)$, the vector space of differentiable functions and $\mathcal{X} = \mathcal{X}(\mathbb{R}^3)$, the vector space of C^∞ vector fields on $\mathbb{R}^3$. Show that these fit together into an exact complex,

$$0 \longrightarrow \mathbb{R} \longrightarrow C \xrightarrow{\nabla} \mathcal{X} \xrightarrow{\nabla\times} \mathcal{X} \xrightarrow{\nabla\cdot} C \longrightarrow 0\ , \tag{3.4}$$

where ∇ is the gradient differential operator from vector calculus, and the initial $\mathbb{R}$ term in the complex represents the constant functions on $\mathbb{R}^3$. This one exact complex compactly encodes many of the relations of vector calculus.

Mayer-Vietoris Complex There are a number of exact complexes that are used in homology, the full exposition of which would take us far afield. Let us focus on one particular example as a means of seeing how exact complexes assist in computational issues. The following is presented in the context of simplicial complexes. Let $X = A \cup B$ be a union of two simplicial complexes with intersection $A \cap B$. The following exact complex is the *Mayer-Vietoris* complex:

$$\longrightarrow H_n(A \cap B) \xrightarrow{H(\phi)} H_n(A) \oplus H_n(B) \xrightarrow{H(\psi)} H_n(X) \xrightarrow{\delta} H_{n-1}(A \cap B) \longrightarrow\ .$$

The linear transformations between the homologies are where all the interesting details lie. These consist of: (1) $H(\psi)$, which adds the homology classes from A and B to give a homology class in X; (2) $H(\phi)$, which reinterprets a homology class in $A \cap B$ to be a homology class of A and a (orientation reversed!) homology class in B respectively; and (3) δ, which decomposes a cycle in X into a sum of chains in A and B, then takes the boundary of one of these chains in $A \cap B$. This unmotivated construction has a clean explanation, to be addressed soon.

For the moment, focus on what the Mayer-Vietoris complex means. This complex captures the additivity of homology. When, for example, $A \cap B$ is empty, then every third term of the complex vanishes — the $H_\bullet(A \cap B)$ terms. Because every-third-term-zero implies that the complementary pairs of incident terms are isomorphisms, this quickly yields that homology of a disjoint union is additive, using $\oplus$. When the intersection is nonempty, the Mayer-Vietoris complex details exactly how the homology of the intersection impacts the homology of the union: it is, precisely, an inclusion-exclusion principle.

Exercise 3.5. Assume the following: for any $k \geqslant 0$, (1) $H_n(\mathbb{D}^k) = 0$ for all $n > 0$; (2) the 1-sphere $\mathbb{S}^1$ has $H_1(\mathbb{S}^1) \cong \mathbb{F}$ and $H_n(\mathbb{S}^1) = 0$ for all $n > 1$. The computation of $H_\bullet(\mathbb{S}^k)$ can be carried out via Mayer-Vietoris as follows. Let A and B be upper and lower hemispheres of $\mathbb{S}^k$, each homeomorphic to $\mathbb{D}^k$ and intersecting at an equatorial $\mathbb{S}^{k-1}$. Write out the Mayer-Vietoris complex in this case: what can you observe? As $H_\bullet(\mathbb{D}^k) \cong 0$ for $k > 0$, one obtains by exactness that $\delta\colon H_n(\mathbb{S}^k) \cong H_{n-1}(\mathbb{S}^{k-1})$ for all n and all k. Thus, starting from a knowledge of $H_\bullet\mathbb{S}^1$, show that $H_n(\mathbb{S}^k) \cong 0$ for $k > 0$ unless $n = k$, where it has dimension equal to one.

Sequences of Sequences One of the themes of these lectures is the utility of composed abstraction: if spaces are useful, so should be spaces of spaces. Later, we will argue that an idea of "homologies of homologies" is sensible and useful (in the guise of the *sheaf theory* of *Lecture 4*). In this lecture, we argue that sequences of sequences and complexes of complexes are useful. We begin with an elucidation of what is behind the Mayer-Vietoris complex: what are the maps, and how does it arise?

Consider the following complex of algebraic complexes,

$$0 \longrightarrow C_\bullet(A \cap B) \xrightarrow{\phi_\bullet} C_\bullet(A) \oplus C_\bullet(B) \xrightarrow{\psi_\bullet} C_\bullet(A+B) \longrightarrow 0, \tag{3.6}$$

with chain maps $\phi_\bullet\colon c \mapsto (c, -c)$, and $\psi_\bullet\colon (a, b) \mapsto a + b$. The term on the right, $C_\bullet(A + B)$, consists of those chains which can be expressed as a sum of chains on A and chains on B. In cellular homology with A, B subcomplexes, $C_\bullet(A + B) \cong C_\bullet(X)$.

Exercise 3.7. Show that this complex-of-complexes is exact by construction.

In general, any such *short* exact complex of complexes can be converted to a *long exact complex* on homologies using a method from homological algebra called the *Snake Lemma* [50,58]. Specifically, given:

$$0 \longrightarrow A_\bullet \xrightarrow{i_\bullet} B_\bullet \xrightarrow{j_\bullet} C_\bullet \longrightarrow 0, \tag{3.8}$$

there is an induced exact complex of homologies

$$\longrightarrow H_n(\mathcal{A}) \xrightarrow{H(i)} H_n(\mathcal{B}) \xrightarrow{H(j)} H_n(\mathcal{C}) \xrightarrow{\delta} H_{n-1}(\mathcal{A}) \xrightarrow{H(i)} . \tag{3.9}$$

Moreover, the long exact complex is *natural*: a commutative diagram of short exact complexes and chain maps

$$\begin{tikzcd}
0 \arrow[r] & A_\bullet \arrow[r] \arrow[d, "f_\bullet"] & B_\bullet \arrow[r] \arrow[d, "g_\bullet"] & C_\bullet \arrow[r] \arrow[d, "h_\bullet"] & 0 . \\
0 \arrow[r] & \tilde{A}_\bullet \arrow[r] & \tilde{B}_\bullet \arrow[r] & \tilde{C}_\bullet \arrow[r] & 0
\end{tikzcd} \tag{3.10}$$

induces a commutative diagram of long exact complexes

$$\begin{tikzcd}
{} \arrow[r] & H_n(\mathcal{A}) \arrow[r] \arrow[d, "H(f)"] & H_n(\mathcal{B}) \arrow[r] \arrow[d, "H(g)"] & H_n(\mathcal{C}) \arrow[r, "\delta"] \arrow[d, "H(h)"] & H_{n-1}(\mathcal{A}) \arrow[r] \arrow[d, "H(f)"] & {} . \\
{} \arrow[r] & H_n(\tilde{\mathcal{A}}) \arrow[r] & H_n(\tilde{\mathcal{B}}) \arrow[r] & H_n(\tilde{\mathcal{C}}) \arrow[r, "\delta"] & H_{n-1}(\tilde{\mathcal{A}}) \arrow[r] & {}
\end{tikzcd} \tag{3.11}$$

The induced *connecting homomorphism* $\delta\colon H_n(\mathcal{C}) \to H_{n-1}(\mathcal{A})$ comes from the boundary map in $\mathcal{C}$ as follows:

1. Fix $[\gamma] \in H_n(\mathcal{C})$; thus, $\gamma \in C_n$.
2. By exactness, $\gamma = j(\beta)$ for some $\beta \in B_n$.
3. By commutativity, $j(\partial\beta) = \partial(j\beta) = \partial\gamma = 0$.
4. By exactness, $\partial\beta = i\alpha$ for some $\alpha \in A_{n-1}$.
5. Set $\delta[\gamma] = [\alpha] \in H_{n-1}(\mathcal{A})$.

Exercise 3.12. This is a tedious but necessary exercise for anyone interested in homological algebra: (1) show that $\delta[\gamma]$ is well-defined and independent of all choices; (2) show that the resulting long complex is exact. Work at *ad tedium*: for help, see any textbook on algebraic topology, [58] recommended.

Euler Characteristic, Redux Complexes solve the mystery of the topological invariance of the Euler characteristic. Recall that we can define the Euler characteristic of a (finite, finite-dimensional) complex $\mathcal{C}$ as in Equation (1.13). The alternating sum is a binary exactness. A short exact complex of vector spaces $0 \to A \to B \to C \to 0$ has $\chi = 0$, since $C \cong B/A$. By applying this to individual rows of a short exact complex of (finite, finite-dimensional) chain complexes, we can lift once again to talk about the Euler characteristic of a (finite-enough) complex of complexes:

$$0 \longrightarrow A_\bullet \longrightarrow B_\bullet \longrightarrow C_\bullet \longrightarrow 0 . \tag{3.13}$$

One sees that χ of this complex also vanishes: $\chi(A_\bullet) - \chi(B_\bullet) + \chi(C_\bullet) = 0$.

The following lemma is the homological version of the Rank-Nullity Theorem from linear algebra:

Lemma 3.14. *The Euler characteristic of a chain complex* $C_\bullet$ *and its homology* $H_\bullet$ *are identical, when both are defined.*

Proof. From the definitions of homology and chain complexes, one has two short exact complexes of chain complexes:

$$\begin{array}{c} 0 \longrightarrow B_\bullet \longrightarrow Z_\bullet \longrightarrow H_\bullet \longrightarrow 0\ . \\ 0 \longrightarrow Z_\bullet \longrightarrow C_\bullet \longrightarrow B_{\bullet-1} \longrightarrow 0 \end{array} \tag{3.15}$$

Here, $B_{\bullet-1}$ is the shifted boundary complex whose k^{th} term is B_{k-1}. By exactness, the Euler characteristic of each of these two complexes is zero; thus, so is the Euler characteristic of their concatenation.

$$0 \longrightarrow B_\bullet \longrightarrow Z_\bullet \longrightarrow H_\bullet \longrightarrow 0 \longrightarrow Z_\bullet \longrightarrow C_\bullet \longrightarrow B_{\bullet-1} \longrightarrow 0\ .$$

Count the $+/-$ signs: the Z terms cancel, and the B terms cancel, since $\chi(B_{\bullet-1}) = -\chi(B_\bullet)$. This leaves two terms left with the conclusion $\chi(H_\bullet) - \chi(C_\bullet) = 0$. ☐

Euler characteristic thus inherits its topological invariance from that of homology. Where does the invariance of homology come from? Something more complicated still?

Homology Theories

Invariance of homology is best discerned from a singularly uncomputable variant that requires a quick deep dive into the plethora of homologies available. We begin with a reminder: homology is an algebraic compression scheme — a way of collapsing a complex to the simplest form that respects its global features. The notion of homology makes sense for any chain complex. Thus far, our only means of generating a complex from a space X has been via some finite auxiliary structure on X, such as a simplicial, cubical, or cellular decomposition. There are other types of structures a space may carry, and, with them, other complexes. In the same way that the homology of a simplicial complex is independent of the simplicial decomposition, the various homologies associated to a space under different auspices tend to be isomorphic.

Reduced homology Our first alternate theory is not really a different type of homology at all; merely a slight change in the chain complex meant to make contractible (or rather acyclic) spaces fit more exactly into homology theory. Recall that a contractible cell complex — such as a single simplex — has homology $H_k = 0$ for all $k > 0$, with H_0 being one-dimensional, recording the fact that the cell complex is connected. For certain results in algebraic topology, it would be convenient to have the homology of a contractible space vanish completely. This can be engineered by an *augmented complex* in a manner that is applicable to any $\mathbb{N}$-graded complex. Assume for simplicity that $\mathcal{C} = (C_k, \partial)$ is a $\mathbb{N}$-graded complex of vector spaces over a field $\mathbb{F}$. The *reduced complex* is the following augmentation:

$$\cdots \xrightarrow{\partial} C_3 \xrightarrow{\partial} C_2 \xrightarrow{\partial} C_1 \xrightarrow{\partial} C_0 \xrightarrow{\epsilon} \mathbb{F} \longrightarrow 0\ , \tag{3.16}$$

where the aumentation map $\epsilon\colon C_0 \to \mathbb{F}$ sends a vector in C_0 to the sum of its components (have fixed a basis for C_0). The resulting homology of this complex is called *reduced homology* and is denoted $\tilde{H}_\bullet$.

Exercise 3.17. Show that the reduced complex is in fact a complex (i.e., that $\epsilon\partial = 0$). How does the reduced homology of a complex differ from the "ordinary" homology? What is the dependence on the choice of augmentation map ϵ? Show that the augmented complex of a contractible simplicial complex is exact.

Čech Homology One simple structure associated to a topological space is an open cover — a collection $\mathcal{U}$ of open sets $\{U_\alpha\}$ in X. The *Čech complex* of $\mathcal{U}$ is the complex $\mathcal{C}(\mathcal{U})$ with basis for $C_k(\mathcal{U})$ being all (unordered) sets of $k+1$ distinct elements of $\mathcal{U}$ with nonempty intersection. (The usual complexities arise for coefficients not in $\mathbb{F}_2$, as one needs to order the elements of $\mathcal{U}$ up to even permutations.) The boundary maps $\partial\colon C_\bullet(\mathcal{U}) \to C_{\bullet-1}(\mathcal{U})$ act on a basis element by forgetting one of the terms in the set, yielding *face maps*.

For a finite collection $\mathcal{U}$ of sets, the Čech complex is identical to the simplicial chain complex of the nerve $\mathcal{N}(\mathcal{U})$. With the further assumption of contractible sets and intersections, the resulting homology is, by the Nerve Lemma, identical to $X = \cup_\alpha U_\alpha$. However, even in the non-finite case, the result still holds if the analogous contractibility assumptions hold. In short, if the all basis elements for the Čech complex are nullhomologous, then the *Čech homology*, $H_\bullet(\mathcal{C}(\mathcal{U}))$ is isomorphic to $H_\bullet(X)$. Of course, the Čech complex and its homology are still well-defined even if the local simplicity assumptions are violated. This Čech homology has in the past been used in the context of a sequence of covers, with limiting phenomena of most interest for complex fractal-like spaces [60]. This is perhaps one of the earliest incarnations of persistent homology.

Singular Homology If you find it risky to think of the Čech homology of a cover of non-finite size, then the next homology theory will seem obscenely prodigal. Given a topological space X, the *singular chain complex* is the complex $\mathcal{C}^{\text{sing}}$ whose k-chains have as basis elements all maps $\sigma\colon \Delta^k \to X$, where Δ^k is the Platonic k-simplex. Note: there are no restrictions on the maps σ other than continuity: images in X may appear crushed or crumpled. The boundary maps are the obvious restrictions of σ to the $(k-1)$-dimensional faces of Δ^k, taking a linear combination with orientations if the field $\mathbb{F}$ demands. The resulting singular chain complex is wildly uncountable, unless X should happen to be a trivial space.

Exercise 3.18. Do the one explicit computation possible in singular homology: show that for X a finite disjoint union of points, the singular homology $H_n^{\text{sing}}(X)$ vanishes except when $n = 0$, in which case it has β_0 equal to the number of points in X. Note: you cannot assume, as in cellular homology, that the higher-dimensional chains C_n vanish for $n > 0$.

There is little hope in computing the resulting *singular homology*, save for the fact that this homology is, blessedly, an efficient compression.

Theorem 3.19. *For a cell complex, singular and cellular homology are isomorphic.*

The proof of this is an induction argument based on the n-skeleton of X. The previous exercise establishes the isomorphism on the level of H_0. To induct to higher-dimensional skeleta requires a few steps just outside the bounds of these lectures: see [51,58] for details.

Homotopy Invariance Why pass to the uncomputable singular theory? There is so much *room* in $\mathcal{C}^{\text{sing}}$ that it is easy to deform continuously and prove the core result.

Theorem 3.20. *Homology [singular] is a homotopy invariant of spaces.*

When combined with Theorem 3.19, we obtain a truly useful, computable result. The proof of Theorem 3.20 does not focus on spaces at all, but rather, in the spirit of these lectures, pulls back the notion of homotopy to complexes. Recall that $f, g\colon X \to Y$ are homotopic if there is a map $F\colon X \times [0,1] \to Y$ which restricts to f on $X \times \{0\}$ and to g on $X \times \{1\}$. A *chain homotopy* between chain maps $\varphi_\bullet, \psi_\bullet \colon \mathcal{C} \to \mathcal{C}'$ is a graded linear transformation $F\colon \mathcal{C} \to \mathcal{C}'$ sending n-chains to $(n+1)$-chains so that $\partial F - F\partial = \varphi_\bullet - \psi_\bullet$:

(3.21)
$$\begin{tikzcd}
\cdots \arrow[r] & C_{n+1} \arrow[r, "\partial"] \arrow[dl, "F"'] \arrow[d, "\psi_\bullet"', "\varphi_\bullet"] & C_n \arrow[r, "\partial"] \arrow[dl, "F"'] \arrow[d, "\psi_\bullet"', "\varphi_\bullet"] & C_{n-1} \arrow[r, "\partial"] \arrow[dl, "F"'] \arrow[d, "\psi_\bullet"', "\varphi_\bullet"] & \cdots . \arrow[dl, "F"] \\
\cdots \arrow[r] & C'_{n+1} \arrow[r, "\partial"'] & C'_n \arrow[r, "\partial"'] & C'_{n-1} \arrow[r, "\partial"'] & \cdots
\end{tikzcd}$$

One calls F a map of *degree* $+1$, indicating the upshift in the grading.[5] Note the morphological resemblance to homotopy of maps: a chain homotopy maps each n-chain to a $n+1$-chain, the algebraic analogue of a 1-parameter family. The difference between the ends of the homotopy, $\partial F - F\partial$, gives the difference between the chain maps.

Exercise 3.22. Show that chain homotopic maps induce the same homomorphisms on homology. Start by considering $[\alpha] \in H_\bullet(\mathcal{C})$, assuming $\varphi_\bullet$ and $\psi_\bullet$ are chain homotopic maps from $\mathcal{C}$ to $\mathcal{C}'$.

The proof of Theorem 3.20 follows from constructing an explicit chain homotopy [58].

Morse Homology All the homology theories we have looked at so far have used simplices or cells as basis elements of chains and dimension as the grading. There is a wonderful homology theory that breaks this pattern in a creative and, eventually, useful manner. Let M be a smooth, finite-dimensional Riemannian

---

[5] The overuse of the term *degree* in graphs, maps of spheres, and chain complexes is unfortunate.

manifold. There is a homology theory based on a dynamical system on M. One chooses a function $h\colon M \to \mathbb{R}$ and considers the (negative) gradient flow of h on M — the smooth dynamical system given by $dx/dt = -\nabla h$.

The dynamics of this vector field are simple: solutions either are fixed points (critical points of h) or flow *downhill* from one fixed point to another. Let $\mathsf{Cr}(h)$ denote the set of critical points, and assume for the sake of simplicity that all such critical points are *nondegenerate* – the second derivative (or *Hessian*) is nondegenerate (has nonzero determinant) at these points. These nondegenerate critical points are the basis elements of a *Morse complex*. What is the grading?

Nondegenerate critical points have a natural grading – the number of negative eigenvalues of the Hessian of second derivatives of h at p. This is called the *Morse index*, $\mu(p)$, of $p \in \mathsf{Cr}(h)$ and has the more topological interpretation as the dimension of the set of points that converge to p in negative time. The Morse index measures how unstable a critical point is: minima have the lowest Morse index; maxima the highest. Balancing a three-legged stool on k legs leads to an index $\mu = 3 - k$ equilibrium.

One obtains the *Morse complex*, $\mathcal{C}^h = (MC_\bullet, \partial)$, with MC_k the vector space with basis $\{p \in \mathsf{Cr}(h)\,;\, \mu(p) = k\}$. The boundary maps encode the global flow of the gradient field: ∂_k counts (modulo 2 in the case of $\mathbb{F}_2$ coefficients) the number of *connecting orbits* – flowlines from a critical point with $\mu = k$ to a critical point with $\mu = k - 1$. One hopes (or assumes) that this number is well-defined. The difficult business is to demonstrate that $\partial^2 = 0$: this involves careful analysis of the connecting orbits, as in, e.g., [8, 84]. The use of $\mathbb{F}_2$ coefficients is highly recommended. The ensuing *Morse homology*, $MH_\bullet(h)$, captures information about M.

Theorem 3.23 (Morse Homology Theorem)**.** *For M a compact manifold and $h\colon M \to \mathbb{R}$ Morse, $MH_\bullet(h; \mathbb{F}_2) \cong H_\bullet(M; \mathbb{F}_2)$, independent of h.*

Exercise 3.24. Compute the Morse homology of a 2-sphere, $\mathbb{S}^2$, outfitted with a Morse function having two maxima and two minima. How many saddle points must it have?

Our perspective is that Morse homology is a *precompression* of the complex onto its *critical* elements, as measured by h.

Discrete Morse Theory As given, Morse homology would seem to be greatly disconnected from the data-centric applications and perspectives of this chapter — it uses smooth manifolds, smooth functions, smooth flows, and nondegenerate critical points, all within the delimited purview of smooth analysis. However, as with most of algebraic topology, the smooth theory is the continuous limit of a correlative discrete theory. In ordinary homology, the discrete [that is, simplicial] theory came first, followed by the limiting case of the singular homology theory. In the case of Morse theory, the smooth version came long before the following discretization, which first seems to have appeared in a mature form in the work of Forman [45, 46]; see also the recent book of Kozlov [65].

Consider for concreteness a simplicial or cell complex X. The critical ingredient for Morse theory is *not* the Morse function but rather its gradient flow. A *discrete vector field* is a pairing V which partitions the cells of X (graded by dimension) into pairs $V_\alpha = (\sigma_\alpha \lhd \tau_\alpha)$ where σ_α is a codimension-1 face of τ_α. All leftover cells of X not paired by V are the *critical cells* of V, $\mathsf{Cr}(V)$. A *discrete flowline* is a sequence (V_i) of distinct paired cells with codimension-1 faces, arranged so that

$$\overbrace{\sigma_1 \lhd \tau_1}^{V_1} \unrhd \overbrace{\sigma_2 \lhd \tau_2}^{V_2} \unrhd \cdots \unrhd \overbrace{\sigma_N \lhd \tau_N}^{V_N} . \tag{3.25}$$

A flowline is *periodic* if $\tau_N \unrhd \sigma_1$ for $N > 1$. A *discrete gradient field* is a discrete vector field devoid of periodic flowlines.

It is best to lift everything to algebraic actions on the chain complex $\mathcal{C} = (C_\bullet^{\mathrm{cell}}, \partial)$ associated to the cell complex X. By linearity, the vector field V induces a chain map $V\colon C_k \to C_{k+1}$ induced by the pairs $\sigma \lhd \tau$ – one visualizes an arrow from the face σ to the cell τ. As with classical Morse homology, $\mathbb{F}_2$ coefficients are simplest.

To every discrete gradient field is associated a discrete Morse complex, $\mathcal{C}^V = (MC_\bullet, \tilde{\partial})$ with MC_k the vector space with basis the critical cells $\{\sigma \in \mathsf{Cr}(V)\,;\, \mathsf{dim}(\sigma) = k\}$. Note that dimension plays the role of Morse index.

Exercise 3.26. Place several discrete gradient fields on a discretization of a circle and examine the critical cells. What do you notice about the number and dimension of critical cells? Does this make sense in light of the Euler characteristic of a circle?

The boundary maps $\tilde{\partial}_k$ count (modulo 2 in the case of $\mathbb{F}_2$ coefficients; with a complicated induced orientation else) the number of discrete flowlines from a critical simplex of dimension k to a critical simplex of dimension $k-1$. Specifically, given τ a critical k-simplex and σ a critical $(k-1)$-simplex, the contribution of $\tilde{\partial}_k(\tau)$ to σ is the number of gradient paths from a face of τ to a coface of σ. In the case that $\sigma \lhd \tau$, then this number is 1, ensuring that the trivial V for which all cells are critical yields $\mathcal{C}^V$ the usual cellular chain complex. It is not too hard to show that $\tilde{\partial}^2 = 0$ and that, therefore, the homology $MH_\bullet(V) = H_\bullet(\mathcal{C}^V)$ is well-defined. As usual, the difficulty lies in getting orientations right for $\mathbb{Z}$ coefficients.

Theorem 3.27 ([45])**.** *For any discrete gradient field* V, $MH_\bullet(V) \cong H_\bullet^{\mathrm{cell}}(X)$.

Discrete Morse theory shows that the classical constraints – manifolds, smooth dynamics, nondegenerate critical points – are not necessary. This point is worthy of emphasis: the classical notion of a critical point (maximum, minimum, saddle) is distilled away from its analytic and dynamical origins until only the algebraic spirit remains.

Applications of discrete Morse theory are numerous and expansive, including to combinatorics [65], mesh simplification [67], image processing [79], configuration spaces of graphs [43, 44], and, most strikingly, efficient computation of homology of cell complexes [69]. This will be our focus for applications to computation at the end of this lecture.

Application: Algorithms

Advances in applications of homological invariants have been and will remain inextricably linked to advances in computational methods for such invariants. Recent history has shown that potential applications are impotent when divorced from computational advances, and computation of unmotivated quantities is futile. The reader who is interested in applying these methods to data is no doubt interested in knowing the best and easiest available software. Though this is not the right venue for a discussion of cutting-edge software, there are a number of existing software libraries/packages for computing homology and persistent homology of simplicial or cubical complexes, some of which are exciting and deep. As of the time of this writing, the most extensive and current benchmarking comparing available software packages can be found in the preprint of Otter et al. [75]. We remark on and summarize a few of the issues involved with computing [persistent] homology, in order to segue into how the theoretical content of this lecture impacts how software can be written.

Time complexity: Homology is known to be output-sensitive, meaning that the complexity of computing homology is a function of how large the homology is, as opposed to how large the complex is. What this means in practice is that the homology of a simple complex is simple to compute. The time-complexity of computing homology is, by output-sensitivity, difficult to specify tightly. The standard algorithm to compute $H_\bullet(X)$ for X a simplicial complex is to compute the *Smith normal form* of the graded boundary map $\partial : C_\bullet \to C_\bullet$, where we concatenate the various gradings into one large vector space. This graded boundary map has a block structure with zero blocks on the block-diagonal (since $\partial_k \colon C_k \to C_{k-1}$ and is nonzero on the superdiagonal blocks). The algorithm for computing Smith normal form is really a slight variant of the ubiquitous Gaussian elimination, with reduction to the normal form via elementary row and column operations. For field coefficients in $\mathbb{F}_2$ this reduction is easily seen to be of time-complexity $O(n^3)$ in the size of the matrix, with an expected run time of $O(n^2)$. This is not encouraging, given the typical sizes seen in applications. Fortunately, compression preprocessing methods exist, as we will detail.

Memory and inputs: Time-complexity is not the only obstruction; holding a complex in memory in nontrivial, as is the problem of *inputting* a complex. A typical simplicial complex is specified by fixing the simplices as basis and then specifying the boundary matrices. For very large complexes, this is prohibitive and unnecessary, as the boundary matrices are typically sparse. There are a number of ways to reduce the input cost, including inputting (1) a distance matrix spelling out an explicit metric between points in a point cloud, using a persistent Dowker complex (see Exercise 2.10) to build the filtered complex; (2) using voxels in a lattice as means of coordinatizing top-dimensional cubes in a cubical complex, and specifying the complex as a list of voxels; and (3) using the Vietoris-Rips complex of a network of nodes and edges, the specification of which requires only a quadratic number of bits of data as a function of nodes.

Exercise 3.28. To get an idea of how the size of a complex leads to an inefficient complexity bound, consider a single simplex, Δ^n, and cube, I^n, each of dimension n. How many total simplices/cubes are in each? Include all faces of all dimensions. Computing the [nearly trivial] homology of such a simple object requires, in principle, computing the Smith normal form of a graded boundary matrix of what net size?

Morse theory & Compression: One fruitful approach for addressing the computation of homology is to consider alternate intermediate compression schemes. If instead of applying Smith Normal Form directly to a graded boundary operator, one modifies the complex first to obtain a smaller chain-homotopic complex, then the resulting complexity bounds may collapse with a dramatic decrease in size of the input. There have been many proposals for reduction and coreduction of chain complexes that preserve homology: see [62] for examples. One clear and successful compression scheme comes from discrete Morse theory. If one puts a discrete gradient field on a cell complex, then the resulting Morse complex is smaller and potentially much smaller, being generated only by critical cells. The process of defining and constructing an associated discrete Morse complex is roughly linear in the size of the cell complex [69] and thus gives an efficient approach to homology computation. This has been implemented in the popular software package `Perseus` (see [75]).

Modernizing Morse Theory: Morse homology, especially the discrete version, has not yet been fully unfolded. There are several revolutionary approaches to Morse theory that incorporate tools outside the bounds of these lectures. Nevertheless, it is the opinion of this author that we are just realizing the full picture of the centrality of Morse theory in Mathematics and in homological algebra in particular. Two recent developments are worth pointing out as breakthroughs in conceptual frameworks with potentially large impact. The first, in the papers by Nanda et al. [71,72], gives a categorical reworking of discrete Morse that relaxes the notion of a discrete vector field to allow for any acyclic pairing of cells and faces without restriction on the dimension of the face. It furthermore shows how to reconstruct the topology of the original complex (up to homotopy type, not homology type) using only data about critical cells and the critical discrete flowlines. Though the tools used are formidable (2-categories and localization), the results are equally strong.

Matroid Morse Theory: The second contribution on the cusp of impact comes in the thesis of Henselman [59] which proposes *matroid theory* as the missing link between Morse theory and homological algebra. *Matroids* are classical structures in the intersection of combinatorial topology and linear algebra and have no end of interesting applications in optimization theory. Henselman recasts discrete Morse theory and persistent homology both in terms of matroids, then exploits matroid-theoretic principles (rank, modularity, minimal bases) in order to generate efficient algorithms for computing persistent homology and barcodes. This work has already led to an initial software package `Eirene`[6] that, as of the time of this writing, has computed persistent H_k for $0 \leqslant k \leqslant 7$ of a filtered 8-dimensional simplicial complex obtained as the Vietoris-Rips complex of a random sampling of 50 points in dimension 20 with a total of $3.3E+11$ simplices on a 2012 Apple laptop in about an hour. This computation is more than two orders of magnitude faster than the computation using the fastest-available software, `Ripser`[7], on a cluster of machines, as recorded in the survey of Otter et al. [75]. This portends much more to come, both at the level of conceptual understanding and computational capability.

This prompts the theme of our next lecture, that in order to prepare for increased applicability, one must ascend and enfold tools and perspectives of increasing generality and power.

Lecture 4: Higher Order

Having developed the basics of topological data analysis, we focus now on the theories and principles to which these point.

Cohomology & Duality

One of the first broad generalizations of all we have described in these lectures is the theory of *cohomology*, an algebraic dual of homology. There are many ways to approach cohomology — dual spaces, Morse theory, differential forms, and configuration spaces all provide useful perspectives in this subject. These lectures will take the low-tech approach most suitable for a first-pass. We have previously considered a general chain complex $\mathcal{C}$ to be a graded sequence of vector spaces $C_\bullet$ with linear transformations $\partial_k \colon C_k \to C_{k-1}$ satisfying $\partial^2 = \partial_{k-1}\partial_k = 0$. These chain complexes are typically graded over the naturals $\mathbb{N}$, and any such complex compresses to its homology, $H_\bullet(\mathcal{C})$, preserving homological features and forgetting all extraneous data.

In a chain complex, the boundary maps descend in the grading. If that grading is tied to dimension or local complexity of an assembly substructure, then the boundary maps encode how more-complex objects are related, attached, or projected to their less-complex components. Though this is a natural data structure

---

[6] Available at gregoryhenselman.org/eirene.

[7] Available at https://github.com/Ripser/ripser.

in many contexts, there are instances in which one knows instead how objects are related to larger superstructures rather than smaller substructures. This prompts the investigation of *cochain complexes*. For purposes of these lectures, a cochain complex is a sequence $\mathcal{C} = (C^\bullet, d^\bullet)$ of vector spaces and linear transformation which increment the grading ($d^k\colon C^k \to C^{k+1}$) and satisfy the complex condition ($d^{k+1}d^k = 0$). One uses subscripts for chain complexes and superscripts for cochain complexes. The *cohomology* of the cochain complex is the complex $H^\bullet(\mathcal{C}) = \ker d/\operatorname{im} d$ consisting of *cocycles* equivalent up to *coboundaries*.

The simplest example of a cochain complex comes from dualizing a chain complex. Given a chain complex (C_k, ∂, k) of $\mathbb{F}$-vector spaces, define $C^k = C_k^\vee$, the vector space of linear functionals $C_k \to \mathbb{F}$. The coboundary d^k is then the adjoint (*de facto*, transpose) of the boundary ∂_{k+1}, so that

$$d \circ d = \partial^\vee \circ \partial^\vee = (\partial \circ \partial)^\vee = 0^\vee = 0. \tag{4.1}$$

In the case of a simplicial complex, the standard simplicial cochain complex is precisely such a dual to the simplicial chain complex. The coboundary operator d is explicit: for $f \in C^k$ a functional on k-simplices, its coboundary acts as $(df)(\tau) = f(\partial\tau)$. For σ a k-simplex, d implicates the *cofaces* – those $(k+1)$-simplices τ having σ as a face.

Dualizing chain complexes in this manner leads to a variety of cohomology theories mirroring the many homology theories of the previous section: simplicial, cellular, singular, Morse, Čech, and other cohomology theories follow.

Exercise 4.2. Fix a triangulated disc $\mathbb{D}^2$ and consider cochains using $\mathbb{F}_2$ coefficients. What do 1-cocycles look like? Show that any such 1-cocycle is the coboundary of a 0-cochain which labels vertices with 0 and 1 *on the left* and *on the right* of the 1-cocycle, so to speak: this is what a trivial class in $H^1(\mathbb{D}^2)$ looks like. Now fix a circle $\mathbb{S}^1$ discretized as a finite graph and construct examples of 1-cocycles that are (1) coboundaries; and (2) nonvanishing in H^1. What is the difference between the trivial and nontrivial cocycles on a circle?

The previous exercise foreshadows the initially depressing truth: nothing new is gained by computing cohomology, in the sense that $H^n(X)$ and $H_n(X)$ have the same dimension for each n. Recall, however, that there is more to co/homology than just the Betti numbers. Functoriality is key, and there is a fundamental difference in how homology and cohomology transform.

Exercise 4.3. Fix $f\colon X \to Y$ a simplicial map of simplicial complexes, and consider the simplicial cochain complexes $\mathcal{C}(X)$ and $\mathcal{C}(Y)$. We recall that the induced chain map $f_\bullet$ yields a well-defined induced homomorphism on homology $H(f)\colon H_\bullet(\mathcal{C}(X)) \to H_\bullet(\mathcal{C}(Y))$. Using what you know about adjoints, show that the induced homomorphism on cohomology is also well-defined but *reverses direction:* $H^\bullet(f)\colon H^\bullet(Y) \to H^\bullet(X)$. This allows one to lift cohomology cocycles from the codomain to the domain.

Alexander Duality There are numerous means by which duality expresses itself in the form of cohomology. One of the most useful and ubiquitous of these is known as *Alexander duality*, which relates the homology and cohomology of a subset of a sphere $\mathbb{S}^n$ (or, with a puncture, $\mathbb{R}^n$) and its complement. The following is a particularly simple form of that duality theorem.

Theorem 4.4 (Alexander Duality)**.** *Let* $A \subset \mathbb{S}^n$ *be compact, nonempty, proper, and locally-contractible. There is an isomorphism*

$$\mathsf{AD}\colon \tilde{H}_k(\mathbb{S}^n - A) \xrightarrow{\cong} \tilde{H}^{n-k-1}(A). \tag{4.5}$$

Note that the reduced theory is used for both homology and cohomology.

Cohomology & Calculus Most students initially view cohomology as more obtuse than homology; however, there are certain instances in which cohomology is the most natural operation. Perhaps the most familiar such setting comes from calculus. As seen in Exercise 3.3 from *Lecture* 2, the familiar constructs of vector calculus on $\mathbb{R}^3$ fit into an exact complex. This exactness reflects the fact that $\mathbb{R}^3$ is topologically trivial [contractible]. Later, in Exercise 4.2, you looked at simplicial 1-cocycles and hopefully noticed that whether or not they are null in H^1 depends on whether or not these cochains are simplicial *gradients* of 0-chains on the vertex set. These exercises together hint at the strong relationship between cohomology and calculus.

The use of gradient, curl, and divergence for vector calculus is, however, an unfortunate vestige of the philosophy of calculus-for-physics as opposed to a more modern calculus-for-data sensibility. A slight modern update sets the stage better for cohomology. For $U \subset \mathbb{R}^n$ an open set, let $\Omega^k(U)$ denote the differentiable k-form fields on U (a smooth choice of multilinear antisymmetric functionals on ordered k-tuples of tangent vectors at each point). For example, Ω^0 consists of smooth functionals, Ω^1 consists of 1-form fields, viewable (in a Euclidean setting) as duals to vector fields, Ω^n consists of signed densities on U times the volume form, and $\Omega^{k>n}(U) = 0$. There is a natural extension of differentiation (familiar from implicit differentiation in calculus class) that gives a coboundary map $d : \Omega^k \to \Omega^{k+1}$, yielding the *deRham complex*,

$$0 \longrightarrow \Omega^0(U) \xrightarrow{d} \Omega^1(U) \xrightarrow{d} \Omega^2(U) \xrightarrow{d} \cdots \xrightarrow{d} \Omega^n(U) \longrightarrow 0\ . \tag{4.6}$$

As one would hope, $d^2 = 0$, in this case due to the fact that mixed partial derivatives commute: you worked this out explicitly in Exercise 3.3. The resulting cohomology of this complex, the *deRham cohomology* $H^\bullet(U)$, is isomorphic to the singular cohomology of U using $\mathbb{R}$ coefficients.

This overlap between calculus and cohomology is neither coincidental nor concluded with this brief example. A slightly deeper foray leads to an examination of the Laplacian operator (on a manifold with some geometric structure). The well-known *Hodge decomposition theorem* then gives, among other things, an isomorphism between the cohomology of the manifold and the *harmonic* differential forms (those in the kernel of the Laplacian). For more information on these connections, see [19]. What is especially satisfying is that the calculus approach to cohomology and the deRham theory feeds back to the simplicial: one can export the Laplacian and the Hodge decomposition theorem to the cellular world (see [51, Ch. 6]). This, then, impacts data-centric problems of ranking and more over networks.

Cohomology & Ranking Cohomology arises in a surprising number of different contexts. One natural example that follows easily from the aforementioned calculus-based perspective on cohomology lives in certain Escherian optical illusions, such as impossible tribars, eternally cyclic waterfalls, or neverending stairs. When one looks at an Escher staircase, the drawn perspective is locally realizable – one can construct a local perspective function. However, a global extension cannot be defined. Thus, an Escherlike loop is really a non-zero class in H^1 (as first pointed out by Penrose [78]).

This is not disconnected from issues of data. Consider the problem of ranking. One simple example that evokes nontrivial 1-cocycles is the popular game of *Rock, Paper, Scissors*, for which there are local but not global ranking functions. A local gradient of *rock-beats-scissors* does not extend to a global gradient. Perhaps this is why customers are asked to conduct rankings (*e.g.*, *Netflix* movie rankings or *Amazon* book rankings) as a 0-cochain (*"how many stars?"*), and not as a 1-cochain (*"which-of-these-two-is-better?"*): nontrivial H^1 is, in this setting, undesirable. The *Condorcet paradox* – that locally consistent comparative rankings can lead to global inconsistencies – is an appearance of H^1 in ranking theory.

There are less frivolous examples of precisely this type of application, leveraging the language of gradients and curls to realize cocycle obstructions to perfect rankings in systems. The paper of Jiang et al. [61] interprets the simplicial cochain complex of the clique/flag complex of a network in terms of rankings. For example, the ($\mathbb{R}$-valued) 0-cochains are interpreted as numerical *score* functions on the nodes of the network; the 1-cochains (supported on edges) are interpreted as pairwise preference rankings (with oriented edges and positive/negative values determining which is preferred over the other); and the higher-dimensional cochains represent more sophisticated local orderings of nodes in a clique [simplex]. They then resort to the calculus-based language of grad, curl, and div to build up the cochain complex and infer from its cohomology information about existence and nonexistence of compatible ranking schemes over the network. Their use of the Laplacian and the Hodge decomposition theorem permits projection of noisy or inconsistent ranking schemes onto the nearest consistent ranking.

There are more sophisticated variants of these ideas, with applications passing beyond finding consistent rankings or orderings. Recent work of Gao et al. [49] gives a cohomological and Hodge-theoretic approach to synchronization problems over networks based on pairwise nodal data in the presence of noise. Singer and collaborators [85, 86] have published several works on cryo electron miscroscopy that is, in essence, a cohomological approach to finding consistent solutions to pairwise-compared data over a network. The larger lesson to be inferred from these types of results is that networks often support data above and beyond what is captured by the network topology alone (nodes, edges). This data blends with the algebra and topology of the system using the language of cohomology. It is this perspective of data that lives above a network that propels our next set of tools.

Cellular Sheaves

One of the most natural uses for cohomology comes in the form of a yet-more-abstract theory that is the stated end of these lectures: sheaf cohomology. Our perspective is that a sheaf is an algebraic data structure tethered to a space (generally) or simplicial complex (in particular). In keeping with the computational and linear-algebraic focus of this series, we will couch everything in the language of linear algebra. The more general approach [20,64,83] is much more general.

Fix X a simplicial (or regular cell) complex with $\trianglelefteq$ denoting the face relation: $\sigma \trianglelefteq \tau$ if and only if $\sigma \subset \bar{\tau}$. A *cellular sheaf* over X, $\mathcal{F}$, is generated by an assignment to each simplex σ of X a vector space $\mathcal{F}(\sigma)$ and to each face σ of τ a *restriction map* – a linear transformation $\mathcal{F}(\sigma \trianglelefteq \tau) \colon \mathcal{F}(\sigma) \to \mathcal{F}(\tau)$ such that faces of faces satisfy the composition rule:

$$\rho \trianglelefteq \sigma \trianglelefteq \tau \;\Rightarrow\; \mathcal{F}(\rho \trianglelefteq \tau) = \mathcal{F}(\sigma \trianglelefteq \tau) \circ \mathcal{F}(\rho \trianglelefteq \sigma). \tag{4.7}$$

The *trivial* face $\tau \trianglelefteq \tau$ by default induces the identity isomorphism $\mathcal{F}(\tau \trianglelefteq \tau) = \mathsf{Id}$. One thinks of the vector spaces as the data over the individual simplices. In the same manner that the simplicial complex is glued up by face maps, the sheaf is assembled by the system of linear transformations.

One simple example of a sheaf on a cell complex X is that of the *constant sheaf*, $\mathbb{F}_X$, taking values in vector spaces over a field $\mathbb{F}$. This sheaf assigns $\mathbb{F}$ to every cell and the identity map $\mathsf{Id} \colon \mathbb{F} \to \mathbb{F}$ to every face $\sigma \trianglelefteq \tau$. In constrast, the *skyscraper sheaf* over a single cell σ of X is the sheaf $\mathbb{F}_\sigma$ that assigns $\mathbb{F}$ to σ and 0 to all other cells and face maps.

Exercise 4.8. Consider the following version of a random rank-1 sheaf over a simplicial complex X. Assign the field $\mathbb{F}$ to every simplex. To each face map $\sigma \trianglelefteq \tau$ assign either Id or 0 according to some (your favorite) random process. Does this always give you a sheaf? How does this depend on X? What is the minimal set of assumptions you would need to make on either X or the random assignment in order to guarantee that what you get is in fact a sheaf?

One thinks of the values of the sheaf over cells as being data and the restriction maps as something like local constraints or relationships between data. It's very worthwhile to think of a sheaf as programmable – one has a great deal of freedom in encoding local relationships. For example, consider the simple linear recurrence $u_{n+1} = A_n u_n$, where $u_n \in \mathbb{R}^k$ is a vector of states and A_n is a k-by-k real matrix. Such a discrete-time dynamical system can be represented as a sheaf $\mathcal{F}$ of states over the time-line $\mathbb{R}$ with the cell structure on $\mathbb{R}$ having $\mathbb{Z}$ as vertices, where $\mathcal{F}$ has constant stalks $\mathbb{R}^k$. One programs the dynamics of the recurrence relation as follows: $\mathcal{F}(\{n\} \trianglelefteq (n, n+1))$ is the map $u \mapsto A_n u$ and $\mathcal{F}(\{n+1\} \trianglelefteq (n, n+1))$ is the identity. Compatibility of local solutions over the sheaf is, precisely, the condition for being a global solution to the dynamics.

Local & Global Sections One says that the sheaf is *generated* by its values on individual simplices of X: this data $\mathcal{F}(\tau)$ over a cell τ is also called the *stalk* or *local sections* of $\mathcal{F}$ on τ: one writes $s_\tau \in \mathcal{F}(\tau)$ for a local section over τ. Though the sheaf is generated by local sections, there is more to a sheaf than its generating data, just as there is more to a vector space than its basis. The restriction maps of a sheaf encode how local sections can be continued into larger sections. One glues together local sections by means of the restriction maps. The value of the sheaf $\mathcal{F}$ on all of X is defined to be those collections of local sections that *continue* according to the restriction maps on faces. The *global sections* of $\mathcal{F}$ on X are defined as:

$$\mathcal{F}(X) = \{(s_\tau)_{\tau \in X} : s_\sigma = \mathcal{F}(\rho \trianglelefteq \sigma)(s_\rho) \ \forall \rho \trianglelefteq \sigma\} \subset \prod_\tau \mathcal{F}(\tau) \,. \tag{4.9}$$

Exercise 4.10. Show that in the previous example of a sheaf for the recurrence relation$u_{n+1} = A_n u_n$, the global solutions to this dynamical system are classified by the global sections of the sheaf.

The observed fact that the value of the sheaf over all of X retains the same sort of structure as the type of data over the vertices — say, a vector space over a field $\mathbb{F}$ — is a hint that this space of global solutions is really a type of homological data. In fact, it is cohomological in nature, and, like zero-dimensional cohomology, it is measure of *connected components* of the sheaf.

Cellular Sheaf Cohomology

In the simple setting of a compact cell complex X, it is easy to define a cochain complex based on a sheaf $\mathcal{F}$ on X. Let $C^n(X; \mathcal{F})$ be the product of $\mathcal{F}(\sigma)$ over all n-cells σ of X. These cochains are connected by coboundary maps as follows:

$$0 \longrightarrow \prod_{\dim\ \sigma=0} \mathcal{F}(\sigma) \xrightarrow{d} \prod_{\dim\ \sigma=1} \mathcal{F}(\sigma) \xrightarrow{d} \prod_{\dim\ \sigma=2} \mathcal{F}(\sigma) \xrightarrow{d} \cdots \,, \tag{4.11}$$

where the coboundary map d is defined on sections over cells using the sheaf restriction maps

$$d(s_\sigma) = \sum_{\sigma \trianglelefteq \tau} [\sigma\colon \tau]\, \mathcal{F}(\sigma \trianglelefteq \tau) s_\sigma, \tag{4.12}$$

where, for a regular cell complex, $[\sigma\colon \tau]$ is either zero or ± 1 depending on the orientation of the simplices involved (beginners may start with all vector spaces using binary coefficients so that $-1 = 1$). Note that $d\colon C^n(X; \mathcal{F}) \to C^{n+1}(X; \mathcal{F})$, since $[\sigma\colon \tau] = 0$ unless σ is a codimension-1 face of τ. This gives a cochain complex: in the computation of d^2, the incidence numbers factor from the restriction maps, and the computation from cellular co/homology suffices to yield 0. The resulting *cellular sheaf cohomology* is denoted $H^\bullet(X;\mathcal{F})$.

This idea of *global compatibility* of sets of local data in a sheaf yield, through the language of cohomology, global qualitative features of the data structure. We have seen several examples of the utility of classifying various types of holes or large-scale qualitative features of a space or complex. Imagine what one can do with a measure of topological features of a data structure *over* a space.

Exercise 4.13. The cohomology of the constant sheaf $\mathbb{F}_X$ on a compact cell complex X is, clearly, $H^\bullet_{\text{cell}}(X;\mathbb{F})$, the usual cellular cohomology of X with coefficients in $\mathbb{F}$. Why the need for compactness? Consider the following cell complex: $X = \mathbb{R}$, decomposed into two vertices and three edges. What happens when you follow all the above steps for the cochain complex of $\mathbb{F}_X$? Show that this problem is solved if you include in the cochain complex only contributions from compact cells.

Exercise 4.14. For a closed subcomplex $A \subset X$, one can define the constant sheaf over A as, roughly speaking, the constant sheaf on A (as its own complex) with all other cells and face maps in X having data zero. Argue that $H^\bullet(X;\mathbb{F}_A) \cong H^\bullet(A;\mathbb{F})$. Conclude that it is possible to have a contractible base space X with nontrivial sheaf cohomology.

The elements of linear algebra recur throughout topology, including sheaf cohomology. Consider the following sheaf $\mathcal{F}$ over the closed interval with two vertices, a and b, and one edge e. The stalks are given as $\mathcal{F}(a) = \mathbb{R}^m$, $\mathcal{F}(b) = 0$, and $\mathcal{F}(e) = \mathbb{R}^n$. The restriction maps are $\mathcal{F}(b \trianglelefteq e) = 0$ and $\mathcal{F}(a \trianglelefteq e) = A$, where A is a linear transformation. Then, by definition, the sheaf cohomology is $H^0 \cong \ker A$ and $H^1 \cong \operatorname{coker} A$.

Cellular sheaf cohomology taking values in vector spaces is really a characterization of solutions to complex networks of linear equations. If one modifies $\mathcal{F}(b) = \mathbb{R}^p$ with $\mathcal{F}(b \trianglelefteq e) = B$ another linear transformation, then the cochain complex takes the form

$$0 \longrightarrow \mathbb{R}^m \times \mathbb{R}^p \xrightarrow{[A|-B]} \mathbb{R}^n \longrightarrow 0 \longrightarrow \cdots, \tag{4.15}$$

where $d = [A|-B] : \mathbb{R}^{m+p} \to \mathbb{R}^n$ is augmentation of A by $-B$. The zero$^{\text{th}}$ sheaf cohomology H^0 is precisely the set of solutions to the equation $Ax = By$, for $x \in \mathbb{R}^m$ and $y \in \mathbb{R}^p$. These are the global sections over the closed edge. The first sheaf cohomology measures the degree to which $Ax - By$ does not span $\mathbb{R}^n$. All higher sheaf cohomology groups vanish.

Exercise 4.16. Prove that sheaf cohomology of a cell complex in grading zero classifies global sections: $H^0(X;\mathcal{F}) = \mathcal{F}(X)$.

Cosheaves Sheaves are meant for cohomology: this direction of the restriction maps insures this. Is there a way to talk about sheaf homology? If one works in the cellular case, this is a simple process. As we have seen that the only real difference between the cohomology of a cochain complex and the homology of a chain complex is whether the grading ascends or descends, a simple matter of arrow reversal on a sheaf should take care of things. It does. A *cosheaf* $\hat{\mathfrak{F}}$ of vector spaces on a simplicial complex assigns to each simplex σ a vector space $\hat{\mathfrak{F}}(\sigma)$ and to each face $\sigma \trianglelefteq \tau$ of τ a *corestriction map* – a linear transformation $\hat{\mathfrak{F}}(\sigma \trianglelefteq \tau)\colon \hat{\mathfrak{F}}(\tau) \to \hat{\mathfrak{F}}(\sigma)$ that reverses the direction of the sheaf maps. Of course, the cosheaf must respect the composition rule:

$$\rho \trianglelefteq \sigma \trianglelefteq \tau \;\Rightarrow\; \hat{\mathfrak{F}}(\rho \trianglelefteq \tau) = \hat{\mathfrak{F}}(\rho \trianglelefteq \sigma) \circ \hat{\mathfrak{F}}(\sigma \trianglelefteq \tau), \tag{4.17}$$

and the identity rule that $\hat{\mathfrak{F}}(\tau \trianglelefteq \tau) = \mathsf{Id}$.

In the cellular context, there are very few differences between sheaves and cosheaves — the use of one over another is a matter of convenience, in terms of which direction makes the most sense. This is by no means true in the more subtle setting of sheaves and cosheaves over open sets in a continuous domain.

Splines & Béziers. Cosheaves and sheaves alike arise in the study of splines, Bézier curves, and other piecewise-assembled structures. For example, a single segment of a planar Bézier curve is specified by the locations of two endpoints, along with additional *control points*, each of which may be interpreted as a *handle* specifying tangency data of the resulting curve at each endpoint. The reader who has used any modern drawing software will understand the control that these handles give over the resulting smooth curve. Most programs use a cubic Bézier curve in the plane – the image of the unit closed interval by a cubic polynomial . In these programs, the specification of the endpoints and the endpoint handles (tangent vectors) completely determines the interior curve segment uniquely.

This can be viewed from the perspective of a cosheaf $\hat{\mathfrak{F}}$ over the closed interval $I = [0,1]$. The costalk over the interior $(0,1)$ is the space of all cubic polynomials from $[0,1] \to \mathbb{R}^2$, which is isomorphic to $\mathbb{R}^4 \oplus \mathbb{R}^4$ (one cubic polynomial for each of the x and y coordinates). If one sets the costalks at the endpoints of $[0,1]$ to be $\mathbb{R}^2$, the physical locations of the endpoints, then the obvious corestriction maps to the endpoint costalks are nothing more than evaluation at 0 and 1 respectively. The corresponding cosheaf chain complex is:

$$\cdots \longrightarrow 0 \longrightarrow \mathbb{R}^4 \oplus \mathbb{R}^4 \xrightarrow{\partial} \mathbb{R}^2 \oplus \mathbb{R}^2 \longrightarrow 0. \tag{4.18}$$

Here, the boundary operator ∂ computes how far the cubic polynomial (edge costalk) 'misses' the specified endpoints (vertex costalks).

Exercise 4.19. Show that for this simple cosheaf, $H_0 = 0$ and $H_1 \cong \mathbb{R}^2 \oplus \mathbb{R}^2$. Interpret this as demonstrating that there are four degrees of freedom available for a cubic planar Bézier curve with fixed endpoints: these degrees of freedom are captured precisely by the pair of handles, each of which is specified by a (planar) tangent vector. Repeat this exercise for a 2-segment cubic planar Bézier curve. How many control points are needed and with what degrees of freedom in order to match the H_1 of the cosheaf?

Note the interesting duality: the global solutions with boundary condition are characterized by the top-dimensional homology of the cosheaf, instead of the zero-dimensional cohomology of a sheaf. This simple example extends greatly, as shown originally by Billera (using cosheaves, without that terminology [16]) and Yuzvinsky (using sheaves [90]). Billera's work implies that the (vector) space of splines over a triangulated Euclidean domain is isomorphic to the top-dimensional homology of a particular cosheaf over the domain. This matches what you see in the simpler example of a Bézier curve over a line segment.

Splines and Béziers are a nice set of examples of cosheaves that have natural higher-dimensional generalizations — Bézier surfaces and surface splines are used in design and modelling of surfaces ranging from architectural structures to vehicle surfaces, ship hulls, and the like. Other examples of sheaves over higher-dimensional spaces arise in the broad generalization of the Euler characteristic to the *Euler calculus*, a topological integral calculus of recent interest in topological signal processing applications [9, 10, 80–82].

Towards Generalizing Barcodes One of the benefits of a more general, sophisticated language is the ability to reinterpret previous results in a new light with new avenues for exploration appearing naturally. Let's wrap up our brief survey of sheaves and cosheaves by revisiting the basics of persistent homology, following the thesis of Curry [33]. Recall the presentation of persistent homology and barcodes in *Lecture 2* that relied crucially on the Structure Theorem for linear sequences of finite-dimensional vector spaces (Theorem 2.15).

There are a few ways one might want to expand this story. We have hinted on a few occasions the desirability to work over a continuous line as a parameter: our story of sequences of vector spaces and linear transformations is bound to the discrete setting. Intuitively, one could take a limit of finer discretizations and hope to obtain a convergence with the appropriate assumptions on variability. Questions of stability and interleaving (recall Exercise 2.24) then arise: see [18, 21, 66]

Another natural question is: what about non-linear sequences? What if instead of a single parameter, there are two or more parameters that one wants to vary? Is it possible to classify higher-dimensional sequences and derive barcodes here? Unfortunately, the situation is much more complex than in the simple, linear setting. There are fundamental algebraic reasons for why such a classification is not directly possible. These obstructions originate from representation theory and *quivers*: see [24, 28, 76]. The good news is that quiver theory implies the existence of barcodes for linear sequences of vector spaces where the directions of the maps do not have to be uniform, as per the *zigzag persistence* of Carlsson and de Silva [24]. The bad news is that quiver theory implies that a well-defined barcode cannot exist as presently conceived for any 1-d sequence that is not a *Dynkin diagram* (meaning, in particular, that higher-dimensional persistence has no simple classification).

Nevertheless, one intuits that sheaves and cosheaves should have some bearing on persistence and barcodes. Consider the classical scenario, in which one has a sequence of finite-dimensional vector spaces V_i and linear transformations $\varphi_i\colon V_i \to V_{i+1}$. Consider the following sheaf $\mathcal{F}$ over a $\mathbb{Z}$-discretized $\mathbb{R}$. To each vertex $\{i\}$ is assigned V_i. To each edge $(i, i+1)$ is assigned V_{i+1} with an identity isomorphism from the right vertex stalk to the edge and φ_i as the left-vertex restriction map to the edge data. Note the similarity of this sheaf to that of the recurrence relation earlier in this lecture. As in the case of the recurrence relation, H^0 detects global solutions: something similar happens for intervals in the barcode.

Exercise 4.20. Recall that persistent homology of a persistence complex is really a homology that is attached to an interval in the parameter line. Given the sheaf $\mathcal{F}$ associated to $\mathbb{R}$ as above, and an interval I subcomplex of $\mathbb{R}$, let $\mathcal{F}_I$ be the restriction of the sheaf to the interval, following Exercise 4.14. Prove that the number of bars in the barcode over I is the dimension of $H^0(I; \mathcal{F}_I)$. Can you argue what changes in $\mathcal{F}$ could be made to preserve this result in the case of a sequence of linear transformations φ_i that do not all "go the same direction"? Can you adapt this construction to a collection of vector spaces and linear transformations over an arbitrary *poset* (partially-ordered set)? Be careful, there are some complications.

There are limits to what basic sheaves and cosheaves can do, as cohomology does not come with the descriptiveness plus uniqueness that the representation theoretic approach gives. Nevertheless, there are certain settings in which barcodes for persistent homology are completely captured by sheaves and cosheaves (see the thesis of Curry for the case of level-set persistence [33]), with more characterizations to come [34].

Homological Data, Redux We summarize by updating and expanding the principles that we outlined earlier in the lecture series into a more refined language:

(1) Algebraic co/chain complexes are a good model for converting a space built from local pieces into a linear-algebraic structure.
(2) Co/homology is the optimal compression scheme to collapse inessential structure and retain qualitative features.
(3) The classification of linear algebraic sequences yields barcodes as a decomposition of sequences of co/homologies, capturing the evolution of qualitative features.
(4) Exact sequences permit inference from partial sequential co/homological data to more global characterization.

(5) A variety of different co/homology theories exist, adapted to different types of structures on a space or complex, with functoriality being the tool for relating (and, usually, equating) these different theories.
(6) Sheaves and cosheaves are algebraic data structures over spaces that can be programmed to encode local constraints.
(7) The concomitant sheaf cohomologies [& cosheaf homologies] compress these data structures down to their qualitative core, integrating local data into global featuress.
(8) Classifications of sheaves and cosheaves recapitulates the classification of co/chain complexes into barcodes, but presage a much broader and more applicable theory in the making.

Application: Sensing & Evasion

Most of the examples of cellular sheaves given thus far have been simplistic, for pedagogical purposes. This lecture ends with an example of what one can do with a more interesting sheaf to solve a nontrivial inference problem. This example is chosen to put together as many pieces as possible of the things we have learned, including homology theories, persistence, cohomology, sheaves, computation, and more. It is, as a result, a bit complicated, and this survey will be highly abbreviated: see [52] for full details.

Consider the following type of *evasion game*, in which an evader tries to hide from a pursuer, and *capture* is determined by being "seen" or "sensed". For concreteness, the game takes place in a Euclidean space $\mathbb{R}^n$ and time progresses over the reals. At each time $t \in \mathbb{R}$, the observer sees or senses a *coverage region* $C_t \subset \mathbb{R}^n$ that is assumed to (1) be connected; and (2) include the region outside a fixed ball (to preclude the evader from running away off to infinity). The *evasion problem* is this: given the coverage regions over time, is it possible for there to be an *evasion path*: a continuous map $e\colon \{t\} \mapsto (\mathbb{R}^n - C_t)$ on all of the timeline $t \in \mathbb{R}$. Such a map is a *section* of the projection $p\colon (\mathbb{R}^n \times \mathbb{R}) - C \to \mathbb{R}$ from the complement of the full coverage region $C \subset \mathbb{R}^n \times \mathbb{R}$ to the timeline.

What makes this problem difficult (i.e., interesting) is that the geometry and topology of the complement of the coverage region, where the evader can hide, is not known: were this known, graph-theoretic methods would handily assist in finding a section or determining nonexistence. Furthermore, the coverage region C is not known geometrically, but rather topologically, with unknown embedding. The thesis of Adams [3] gives examples of two different time-dependent coverage regions, C and C', whose fibers are topologically the same (homotopic) for each t, but which differ in the existence of evasion path. The core difficulty is that, though C and C' each admit "tunnels" in their complements stretching over the entire timeline, one of them has a tunnel that snakes backwards along the time axis: topologically, legal; physically, illegal.

Work of Adams and Carlsson [3] gives a complete solution to the existence of an evasion path in the case of a planar ($n = 2$) system with additional genericity conditions and some geometric assumptions. Recently, a complete solution in all dimensions was given [52] using sheaves and sheaf cohomology. One begins with a closed coverage region $C \subset \mathbb{R}^n \times \mathbb{R}$ whose complement is uniformly bounded over $\mathbb{R}$. For the sake of exposition, assume that the time axis is given a discretization into (ordered) vertices v_i and edges $e_i = (v_i, v_{i+1})$ such that the coverage domains C_t are topologically equivalent over each edge (this is not strictly necessary). There are a few simple sheaves over the discretized timeline relevant to the evasion problem.

First, consider for each time $t \in \mathbb{R}$, the coverage domain $C_t \subset \mathbb{R}^n$. How many different ways are there for an evader to hide from C_t? This is regulated by the number of connected components of the complement, classified by $H_0(\mathbb{R}^n - C_t)$. Since we do not have access to C_t directly (remember – its embedding in $\mathbb{R}^n$ is unknown to the pursuer), we must try to compute this H_0 based only on the topology of C_t. That this can be done is an obvious but wonderful corollary of *Alexander duality*, which relates the homology and cohomology of complementary subsets of $\mathbb{R}^n$. In this setting, Alexander duality implies that $H_0(\mathbb{R}^n - C_t) \cong H^{n-1}(C_t)$: this, then, is something we can measure, and motivates using the Leray cellular sheaf $\mathcal{H}$ of $n-1$ dimensional cohomology of the coverage regions over the time axis. Specifically, for each edge e_i define $\mathcal{H}(e_i) = H^{n-1}(C_{(v_i, v_{i+1})})$ to be the cohomology of the region over the open edge. For the vertices, use the star: $\mathcal{H}(v_i) = H^{n-1}(C_{(v_{i-1}, v_{i+1})})$.

Exercise 4.21. Why are the stalks over the vertices defined in this way? Show that this gives a well-defined cellular sheaf using as restriction maps the induced homomorphisms on cohomology. *Hint:* which way do the induced maps in cohomology go?

The intuition is that the global sections of this sheaf over the time axis, $H^0(\mathbb{R}; \mathcal{H})$, would classify the different complementary "tunnels" through the coverage set that an evader could use to escape detection. Unfortunately, this is incorrect, for reasons pointed out by Adams and Carlsson [3] (using the language of zigzags). The culprit is the commutative nature of homology and cohomology — one cannot discern tunnels which illegally twirl backwards in time. To solve this problem, one *could* try to keep track of some sort of directedness or orientation. Thanks to the assumption that C_t is connected for all time, there is a global orientation class on $\mathbb{R}^n$ that can be used to assign a ± 1 to basis elements of $H^{n-1}(C_t)$ based on whether the complementary tunnel is participating in a time-orientation-preserving evasion path on the time interval $(-\infty, t)$.

However, to incorporate this orientation data into the sheaf requires breaking the bounds of working with vector spaces. As detailed in [52], one may use sheaves that take values in *semigroups*. In this particular case, the semigroups are *positive cones* within vector spaces, where a *cone* $K \subset V$ in a vector space V is a subset closed under vector addition and closed under multiplication by $\mathbb{R}^+ = (0, \infty)$, the [strictly] positive reals. A cone is *positive* if $K \cap -K = \varnothing$. With work, one can formulate co/homology theories and sheaves to take values in cones: for details, see [52]. The story proceeds: within the sheaf $\mathcal{H}$ of $n-1$ dimensional cohomology of C on $\mathbb{R}$, there is (via abuse of terminology) a "subsheaf" ${}^+\mathcal{H}$ of positive cones, meaning that the stalks of ${}^+\mathcal{H}$ are positive cones within the stalks of $\mathcal{H}$, encoding all the *positive* cohomology classes that can participate in a legal (time-orientation-respecting) evasion path. It is this sheaf of positive cones that classifies evasion paths.

Theorem 4.22. *For $n > 1$ and $C = \{C_t\} \subset \mathbb{R}^n \times \mathbb{R}$ closed and with bounded complement consisting of connected fibers C_t for all t, there is an evasion path over $\mathbb{R}$ if and only if $H^0(\mathbb{R}; {}^+\mathcal{H})$ is nonempty.*

Note that the theorem statement says *nonempty* instead of *nonzero*, since sheaf takes values in positive cones, which are $\mathbb{R}^+$-cones and thus do not contain zero. The most interesting part of the story is how one can compute this H^0. This is where the technicalities begin to weigh heavily, as one cannot use the classical definition of H^0 in terms of kernels and images. The congenial commutative world of vector spaces requires significant care when passing to the nonabelian setting. One defines H^0 for sheaves of cones using constructs from category theory (*limits*, specifically). Computation of such objects requires a great deal more thought and care than the simpler linear-algebraic notions of these lectures. That is by no means a defect; indeed, it is a harbinger. Increase in resolution requires an increase in algebra.

Conclusion: Beyond Linear Algebra

These lecture notes have approached topological data analysis from the perspective of homological algebra. If the reader takes from these notes the singular idea that linear algebra can be enriched to cover not merely linear transformations, but also sequences of linear transformations that form complexes, then these lectures will have served their intended purpose. The reader for whom this seems to open a new world will be delighted indeed to learn that the vector space version of homological algebra is almost too pedestrian to matter to mathematicians: as hinted at in the evader inference example above, the story begins in earnest when one works with rings, modules, and more interesting categories still. Nevertheless, for applications to data, vector spaces and linear transformations are a safe place to start.

For additional material that aligns with the author's view on topology and homological algebra, the books [50, 51] are recommended. It is noted that the perspective of these lecture notes is idiosyncratic. For a broader view, the interested reader is encouraged to consult the growing literature on topological data analysis. The book by Edelsbrunner and Harer [40] is a gentle introduction, with emphases on both theory and algorithms. The book on computational homology by Kaczynsky, Mischaikow, and Mrozek [62] has even more algorithmic material, mostly in the setting of cubical complexes. Both books suffer from the short shelf life of algorithms as compared to theories. Newer titles on the theory of persistent homology are in the process of appearing: that of Oudot [76] is one of perhaps several soon to come. For introductions to topology specifically geared towards the data sciences, there does not seem to be an ideal book; rather, a selection of survey articles such as that of Carlsson [23] is appropriate.

The open directions for inquiry are perhaps too many to identify properly – the subject of topological data analysis is in its infancy. One can say with certainty that there is a slow spread of these topological methods and perspectives to new application domains. This will continue, and it is not yet clear to this author whether neuroscience, genetics, signal processing, materials science, or some other domain will be the locus of inquiry that benefits most from homological methods, so rapid has been the advances in all these areas. Dual to this spread in applications is the antipodal expansion into Mathematics, as ideas from homological data engage with impact contemporary Mathematics. The unique demands of data have already prompted explorations into mathematical structures (e.g., interleaving distance in sheaf theory and persistence in matroid theory) which otherwise would seem unmotivated and be unexplored. It is to be hoped that the simple applications of representation theory to homological data analysis will enervate deeper explorations into the mathematical tools.

There is at the moment a frenzy of activity surrounding various notions of stability associated to persistent homology, sheaves, and related structures concerning representation theory. It is likely to take some time to sort things out into their clearest form. In general, one expects that applications of deeper ideas from algebraic topology exist and will percolate through applied mathematics into application domains. Perhaps the point of most optimism and uncertainty lies in the intersection with probabilistic and stochastic methods. The rest of this volume makes ample use of such tools; their absence in these notes is noticeable. Topology and probability are neither antithetical nor natural partners; expectation of progress is warranted from some excellent results on the topology of Gaussian random fields [6] and recent work on the homology of random complexes [63]. Much of the material from these lectures appears ripe for a merger with modern probabilistic methods. Courage and optimism — two of the cardinal mathematical virtues — are needed for this.

References

[1] A. Abrams and R. Ghrist. State complexes for metamorphic systems. *International Journal of Robotics Research*, 23(7,8):809–824, 2004. 7

[2] H. Adams and G. Carlsson. On the nonlinear statistics of range image patches. *SIAM J. Imaging Sci.*, 2(1):110–117, 2009. 22

[3] H. Adams and G. Carlsson. Evasion paths in mobile sensor networks. *International Journal of Robotics Research*, 34:90–104, 2014. 46, 47

[4] R. J. Adler. *The Geometry of Random Fields*. Society for Industrial and Applied Mathematics, 1981. 9

[5] R. J. Adler, O. Bobrowski, M. S. Borman, E. Subag, and S. Weinberger. Persistent homology for random fields and complexes. In *Borrowing Strength: Theory Powering Applications*, pages 124–143. IMS Collections, 2010. 13

[6] R. J. Adler and J. E. Taylor. *Random Fields and Geometry*. Springer Monographs in Mathematics. Springer, New York, 2007. 9, 49

[7] R. H. Atkin. *Combinatorial Connectivities in Social Systems*. Springer Basel AG, 1977. 6

[8] A. Banyaga and D. Hurtubise. *Morse Homology*. Springer, 2004. 32

[9] Y. Baryshnikov and R. Ghrist. Target enumeration via Euler characteristic integrals. *SIAM J. Appl. Math.*, 70(3):825–844, 2009. 9, 44

[10] Y. Baryshnikov and R. Ghrist. Euler integration over definable functions. *Proc. Natl. Acad. Sci. USA*, 107(21):9525–9530, 2010. 9, 44

[11] Y. Baryshnikov, R. Ghrist, and D. Lipsky. Inversion of Euler integral transforms with applications to sensor data. *Inverse Problems*, 27(12), 2011. 9

[12] U. Bauer and M. Lesnick. Induced matchingsand the algebraic stability of persistence barcodes. *Discrete & Computational Geometry*, 6(2):162–191, 2015. 21

[13] P. Bendich, H. Edelsbrunner, and M. Kerber. Computing roburobustand persistence for images. *IEEE Trans. Visual and Comput. Graphics*, pages 1251–1260, 2010. 22

[14] P. Bendich, J. Marron, E. Miller, A. Pielcoh, and S. Skwerer. Persistent homology analysis of brain artery trees. To appear in *Ann. Appl. Stat.*, 2016. 22

[15] S. Bhattacharya, R. Ghrist, and V. Kumar. Persistent homology for path planning in uncertain environments. *IEEE Trans. on Robotics*, 31(3):578–590, 2015. 23

[16] L. J. Billera. Homology of smooth splines: generic triangulations and a conjecture of Strang. *Trans. Amer. Math. Soc.*, 310(1):325–340, 1988. 44

[17] L. J. Billera, S. P. Holmes, and K. Vogtmann. Geometry of the space of phylogenetic trees. *Adv. in Appl. Math.*, 27(4):733–767, 2001. 7

[18] M. Botnan and M. Lesnick. Algebraic stability of zigzag persistence modules. ArXiv:160400655v2, Apr 2016. 21, 44

[19] R. Bott and L. Tu. *Differential Forms in Algebraic Topology*. Springer, 1982. 39

[20] G. Bredon. *Sheaf Theory*. Springer, 1997. 40

[21] P. Bubenik, V. de Silva, and J. Scott. Metrics for generalized persistence modules. *Found. Comput. Math.*, 15(6):1501–1531, 2015. 21, 44

[22] P. Bubenik and J. A. Scott. Categorification of persistent homology. *Discrete Comput. Geom.*, 51(3):600–627, 2014. 21

[23] G. Carlsson. The shape of data. In *Foundations of computational mathematics, Budapest 2011*, volume 403 of *London Math. Soc. Lecture Note Ser.*, pages 16–44. Cambridge Univ. Press, Cambridge, 2013. 49

[24] G. Carlsson and V. de Silva. Zigzag persistence. *Found. Comput. Math.*, 10(4):367–405, 2010. 44, 45

[25] G. Carlsson, T. Ishkhanov, V. de Silva, and A. Zomorodian. On the local behavior of spaces of natural images. *International Journal of Computer Vision*, 76(1):1–12, Jan. 2008. 22

[26] G. Carlsson and F. Mémoli. Characterization, stability and convergence of hierarchical clustering methods. *J. Mach. Learn. Res.*, 11:1425–1470, Aug. 2010. 22

[27] G. Carlsson and F. Mémoli. Classifying clustering schemes. *Found. Comput. Math.*, 13(2):221–252, 2013. 22

[28] G. Carlsson and A. Zomorodian. The theory of multidimensional persistence. *Discrete Comput. Geom.*, 42(1):71–93, 2009. 44

[29] S. Carson, V. Ruta, L. Abbott, and R. Axel. Random convergence of olfactory inputs in the *drosophila* mushroom body. *Nature*, 497(7447):113–117, 2013. 14

[30] F. Chazal, V. de Silva, M. Glisse, and S. Oudot. The structure and stability of persistence modules. *Arxiv preprint arXiv:1207.3674*, 2012. 21

[31] D. Cohen-Steiner, H. Edelsbrunner, and J. Harer. Stability of persistence diagrams. *Discrete Comput. Geom.*, 37(1):103–120, 2007. 21

[32] A. Collins, A. Zomorodian, G. Carlsson, and L. Guibas. A barcode shape descriptor for curve point cloud data. In M. Alexa and S. Rusinkiewicz, editors, *Eurographics Symposium on Point-Based Graphics*, ETH, Zürich, Switzerland, 2004. 20

[33] J. Curry. *Sheaves, Cosheaves and Applications*. PhD thesis, University of Pennsylvania, 2014. 44, 45

[34] J. Curry and A. Patel. Classification of constructible cosheaves. ArXiv 1603.01587, Mar 2016. 45

[35] C. Curto and V. Itskov. Cell groups reveal structure of stimulus space. *PLoS Comput. Biol.*, 4(10):e1000205, 13, 2008. 13

[36] M. d'Amico, P. Frosini, and C. Landi. Optimal matching between reduced size functions. Technical Report 35, DISMI, Univ. degli Studi di Modena e Reggio Emilia, Italy, 2003. 20

[37] V. de Silva and G. Carlsson. Topological estimation using witness complexes. In M. Alexa and S. Rusinkiewicz, editors, *Eurographics Symposium on Point-based Graphics*, 2004. 6

[38] J. Derenick, A. Speranzon, and R. Ghrist. Homological sensing for mobile robot localization. In *Proc. Intl. Conf. Robotics & Aut.*, 2012. 24

[39] C. Dowker. Homology groups of relations. *Annals of Mathematics*, pages 84–95, 1952. 6

[40] H. Edelsbrunner and J. Harer. *Computational Topology: an Introduction*. American Mathematical Society, Providence, RI, 2010. 3, 49

[41] H. Edelsbrunner, D. Letscher, and A. Zomorodian. Topological persistence and simplification. *Discrete and Computational Geometry*, 28:511–533, 2002. 20

[42] M. Farber. *Invitation to Topological Robotics*. Zurich Lectures in Advanced Mathematics. European Mathematical Society (EMS), Zürich, 2008. 9

[43] D. Farley and L. Sabalka. On the cohomology rings of tree braid groups. *J. Pure Appl. Algebra*, 212(1):53–71, 2008. 34

[44] D. Farley and L. Sabalka. Presentations of graph braid groups. *Forum Math.*, 24(4):827–859, 2012. 34

[45] R. Forman. Morse theory for cell complexes. *Adv. Math.*, 134(1):90–145, 1998. 33

[46] R. Forman. A user's guide to discrete Morse theory. *Sém. Lothar. Combin.*, 48, 2002. 33

[47] Ś. R. Gal. Euler characteristic of the configuration space of a complex. *Colloq. Math.*, 89(1):61–67, 2001. 9

[48] M. Gameiro, Y. Hiraoka, S. Izumi, M. Kramar, K. Mischaikow, and V. Nanda. Topological measurement of protein compressibility via persistent diagrams. *Japan J. Industrial & Applied Mathematics*, 32(1):1–17, Oct 2014. 24

[49] T. Gao, J. Brodzki, and S. Mukherjee. The geometry of synchronization problems and learning group actions. ArXiv:1610.09051, 2016. 39

[50] S. I. Gelfand and Y. I. Manin. *Methods of Homological Algebra*. Springer Monographs in Mathematics. Springer-Verlag, Berlin, second edition, 2003. 27, 48

[51] R. Ghrist. *Elementary Applied Topology*. Createspace, 1.0 edition, 2014. 3, 7, 10, 31, 39, 48

[52] R. Ghrist and S. Krishnan. Positive Alexander duality for pursuit and evasion. ArXiv:1507.04741, 2015. 46, 47, 48

[53] R. Ghrist and S. M. Lavalle. Nonpositive curvature and Pareto optimal coordination of robots. *SIAM J. Control Optim.*, 45(5):1697–1713, 2006. 7

[54] R. Ghrist, D. Lipsky, J. Derenick, and A. Speranzon. Topological landmark-based navigation and mapping. Preprint, 2012. 6, 24

[55] R. Ghrist and V. Peterson. The geometry and topology of reconfiguration. *Adv. in Appl. Math.*, 38(3):302–323, 2007. 7

[56] C. Giusti, E. Pastalkova, C. Curto, and V. Itskov. Clique topology reveal intrinsic structure in neural connections. *Proc. Nat. Acad. Sci.*, 112(44):13455–13460, 2015. 12, 13

[57] L. J. Guibas and S. Y. Oudot. Reconstruction using witness complexes. In *Proc. 18th ACM-SIAM Sympos. on Discrete Algorithms*, pages 1076–1085, 2007. 6

[58] A. Hatcher. *Algebraic Topology*. Cambridge University Press, 2002. 9, 10, 27, 28, 31

[59] G. Henselman and R. Ghrist. Matroid filtrations and computational persistent homology. ArXiv:1606.00199 [math.AT], Jun 2016. 36

[60] J. Hocking and G. Young. *Topology*. Dover Press, 1988. 30

[61] X. Jiang, L.-H. Lim, Y. Yao, and Y. Ye. Statistical ranking and combinatorial Hodge theory. *Math. Program.*, 127(1, Ser. B):203–244, 2011. 39

[62] T. Kaczynski, K. Mischaikow, and M. Mrozek. *Computational Homology*, volume 157 of *Applied Mathematical Sciences*. Springer-Verlag, New York, 2004. 3, 7, 35, 49

[63] M. Kahle. Topology of random clique complexes. *Discrete & Computational Geometry*, 45(3):553–573, 2011. 14, 49

[64] M. Kashiwara and P. Schapira. *Categories and Sheaves*, volume 332 of *Grundlehren der Mathematischen Wissenschaften*. Springer-Verlag, 2006. 40

[65] D. Kozlov. *Combinatorial Algebraic Topology*, volume 21 of *Algorithms and Computation in Mathematics*. Springer, 2008. 33, 34

[66] M. Lesnick. The theory of the interleaving distance on multidimensional persistence modules. arXiv:1106.5305, 2011. 44

[67] T. Lewiner, H. Lopes, and G. Tavares. Applications of Forman's discrete Morse theory to topology visualization and mesh compression. *IEEE Trans. Visualization & Comput. Graphics*, 10(5):499–508, 2004. 34

[68] M. Minsky and S. Papert. *Perceptrons: An Introduction to Computational Geometry*. MIT Press, 1987. 9

[69] K. Mischaikow and V. Nanda. Morse theory for filtrations and efficient computation of persistent homology. *Discrete Comput. Geom.*, 50(2):330–353, 2013. 34, 35

[70] J. Munkres. *Topology*. Prentice Hall, 2000. 4

[71] V. Nanda. Discrete Morse theory and localization. ArXiv:1510.01907 [math.AT], Oct 2015. 35

[72] V. Nanda, D. Tamaki, and K. Tanaka. Discrete Morse theory and classifying spaces. ArXiv:1612.08429v1, Dec 2016. 35

[73] M. Nicolau, A. J. Levine, and G. Carlsson. Topology based data analysis identifies a subgroup of breast cancers with a unique mutational profile and excellent survival. *Proceedings of the National Academy of Sciences*, 108(17):7265–7270, 2011. 23

[74] J. O'Keefe and J. Dostrovsky. The hippocampes as a spatial map. *Brain Research*, 34(1):171–175, 1971. 13

[75] N. Otter, M. Porter, U. Tillmann, P. Grindrod, and H. Harrington. A roadmap for the computation of persistent homology. ArXiv:1506.08903v5, Jan 2017. 34, 35, 36

[76] S. Oudot. *Persistence Theory: From Quiver Representations to Data Analysis*. American Mathematical Society, 2015. 3, 44, 49

[77] L. Pachter and B. Sturmfels. The mathematics of phylogenomics. *SIAM Rev.*, 49(1):3–31, 2007. 7

[78] R. Penrose. La cohomologie des figures impossibles. *Structural Topology*, 17:11–16, 1991. 39

[79] V. Robins, P. Wood, and A. Sheppard. Theory and algorithms for constructing discrete Morse complexes from grayscale digital images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 33(8):1646–1658, 2011. 22, 34

[80] M. Robinson. *Topological Signal Processing*. Springer, Heidelberg, 2014. 44

[81] P. Schapira. Operations on constructible functions. *J. Pure Appl. Algebra*, 72(1):83–93, 1991. 44

[82] P. Schapira. Tomography of constructible functions. In *Applied Algebra, Algebraic Algorithms and Error-Correcting Codes*, pages 427–435. Springer, 1995. 44

[83] J. Schürmann. *Topology of Singular Spaces and Constructible Sheaves*, volume 63 of *Mathematics Institute of the Polish Academy of Sciences. Mathematical Monographs (New Series)*. Birkhäuser Verlag, Basel, 2003. 40

[84] M. Schwarz. *Morse Homology*, volume 111 of *Progress in Mathematics*. Birkhäuser Verlag, Basel, 1993. 32

[85] Y. Shkolnisky and A. Singer. Viewing direction estimation in cryo-EM using synchronization. *SIAM J. Imaging Sci.*, 5(3):1088–1110, 2012. 39

[86] A. Singer. Angular synchronization by eigenvectors and semidefinite programming. *Appl. Comput. Harmonic Anal.*, 30(1):20–36, 2011. 39

[87] A. Sizemore, C. Giusti, and D. Bassett. Classification of weighted networks through mesoscale homological features. *J. Complex Networks*, 2016. 14

[88] B. Torres, J. Oliviera, A. Tate, P. Rath, K. Cumnock, and D. Schneider. Tracking resilience to infections by mapping disease space. *PLOS Biology*, 2016. 23

[89] A. Wilkerson, H. Chintakunta, H. Krim, T. Moore, and A. Swami. A distributed collapse of a network's dimensionality. In *Proceedings of Global Conference on Signal and Information Processing (GlobalSIP)*. IEEE, 2013. 6

[90] S. Yuzvinsky. Modules of splines on polyhedral complexes. *Math. Z.*, 210(2):245–254, 1992. 44

[91] A. Zomorodian. *Topology for Computing*. Cambridge Univ Press, 2005. 20

Departments of Mathematics and Electrical & Systems Engineering, University of Pennsylvania
E-mail address: `ghrist@math.upenn.edu`