07-10-2015

Contents

We introduce the homology of posets, taking great pains to explain the geometric content of the definitions. The reader is advised that many of the basic calculations of homological algebra are best learned by doing. Trying to follow someone else's work can often be more confusing and difficult than simply working it out for oneself. Consequently, we have left some standard verifications as opportunities for the reader to get his/her feet wet.[1]

1. Introduction

In this course, we have been interested in the possibility of modeling (infinite) spaces – like spheres – with (weakly homotopy equivalent) finite ones. Finite objects, in priniciple, should be easier to understand; if worse comes to worse, one can always try to write down every possibility and check if something holds. Thus, having finite models could help us resolve difficult problems.

For example, we've already talked about the homotopy groups π_n, and how computing even $\pi_q(S^n) = \mathrm{Map}(S^q, S^n)/\text{homotopy}$ is an extremely difficult problem. Imagine if we could find finite models $\widetilde{S^n}$ for the n-spheres, say with $2n+2$ points. . . then we could just write down all of the (finitely many) maps $\widetilde{S^q} \to \widetilde{S^n}$ and check what happens! Wouldn't that be nice. There's just one problem: **this can't possibly work**.

There are, in fact, finite models of spheres. For instance, there's one for S^1 with just 4 points – let's call it $\widetilde{S^1}$. Then $\mathrm{Map}(\widetilde{S^1}, \widetilde{S^1}) \leqq 4^4$, so $\pi_0\mathrm{Map}(\widetilde{S^1}, \widetilde{S^1}) < \infty$. But $\pi_1(S^1) = \pi_0\mathrm{Map}(S^1, S^1) \cong \mathbb{Z}$. This is the "conundrum": even though we have a space $\widetilde{S^1}$ that is in many respects like the ordinary sphere S^1, the homotopy groups

Date: July 16, 2015.

[1]For those who get stuck, some of them should be in Hatcher or *Concise*.

of our model are not seeing everything. They are seeing something, but something strange is happening and it's subtle.

We're going to develop a new invariant called **homology**. Unlike homotopy groups, homology groups are going to be something that we can compute (relatively) easily, and that computers can be programmed to compute. Of course, there is a tradeoff: rich invariants that "see" a lot about our spaces are hard to compute (because they contain so much information), and therefore computable invariants cannot be as rich. Homology is a coarser invariant than homotopy.

2. Simplices

The idea is that spaces are built up out of simplices (just a fancy word for "triangle") of varying dimension[2]. One hopes that by looking at the (combinatorial) data of how simplices are attached to one another, we will be able to say something about the structure of our space. The homology groups we are going to define will, in some sense, measure how many holes (of varying dimension) are present.

2.1. Topological Simplices.

To make things precise: a 0-simplex is a point, a 1-simplex is an edge, a 2-simplex is a triangle, a 3-simplex is a tetrahedron, etc. Geometrically, there is the "standard topological n-simplex"

$$\Delta_n^{top} := \left\{ (x_1, \dots, x_{n+1}) \in \mathbb{R}^{n+1} \,\middle|\, 0 \leqq x_1, \dots, x_n \leqq 1 \text{ and } \sum_{i=1}^{n+1} x_i = 1 \right\}$$

which can be thought of as the convex hull of (= smallest convex set containing) the vectors

$$e_i := (0, \dots, \underbrace{1}_{i\text{th spot}}, \dots, 0)$$

for $i = 1, \dots, n+1$.

Exercise: A subset $X \subseteqq \mathbb{R}^n$ (or more generally, a real vector space) is called **convex** if for each pair of points $x, y \in X$, the line segment

$$l := \{x + t(y - x) \mid t \in [0, 1]\}$$

is also contained in X. Prove that the convex hull of the vectors $v_1, \dots, v_k \in \mathbb{R}^n$ is

$$\left\{ t_1 v_1 + \cdots + t_k v_k \,\middle|\, 0 \leqq t_1, \dots, t_k \leqq 1 \text{ and } \sum_{i=1}^{k} t_i = 1 \right\}.$$

Hint: induction.

Exercise: Draw pictures of the standard topological n-simplices for $n = 0, 1, 2$ to see that these really are just a point, edge, and triangle.

For a simple example of how a space might be built out of simplices, consider the unit circle S^1. We can obtain this space by gluing the endpoints of an edge ($\cong [0, 1]$) to a single point. In terms of diagrams, we have a pushout square

[2] or at the very least, can be approximated by spaces constructed in this way

$$\begin{array}{ccc} \{0,1\} & \longrightarrow & * \\ \partial \downarrow & & \downarrow \\ [0,1] & \longrightarrow & S^1 \end{array}$$

Exercise: Generalize the above to arbitrary $n > 0$, i.e. construct an n-sphere from a point and a single n-simplex.

It is worth noting that for *topological* purposes, an n-simplex is just a piece of $\mathbb{R}^n$, and is interchangeable with a closed ball or a closed cube. The precise geometry – i.e. the fact that simplices have sides – only becomes relevant when we introduce algebra/combinatorics into the picture.

Now, suppose X is a topological space. A **singular n-simplex** in X is defined to be a continuous map $\Delta_n^{top} \to X$. You can think of it as a figure in X parametrized by the standard topological n-simplex. It is called "singular" because we do *not* require the map to be an embedding or smooth in any way – only continuous.

2.2. **Simplices in a Poset.** Now that we've discussed the geometric/topological situation, we're going to abstract to posets ($\simeq$ Alexandroff T_0 spaces).

Start with a poset C, and regard it as a category $\mathscr{C}$ by letting the objects of $\mathscr{C}$ be the elements of C, and placing a unique morphism $x \to y$ if $x \leqq y$. One can think of a morphism $x \to y$ as "witnessing" the relation $x \leqq y$. Accordingly, we will use $x \to y$ and $x \leqq y$ interchangeably. Just as we can define singular simplices in a topological space, we can define simplices in the poset $\mathscr{C}$. This is done as follows:

- The 0-simplices of $\mathscr{C}$ are points (objects):

  $$x$$

  Let [0] be the poset $\{0\}$ in the usual order. Then a 0-simplex is equivalent to an order preserving map $[0] \to \mathscr{C}$.
- The 1-simplices of $\mathscr{C}$ are pairs $x \leqq y$:

  $$x \longrightarrow y$$

  Let [1] be the poset $\{0, 1\}$ in the usual order. Then a 1-simplex is equivalent to an order-preserving map $[1] \to \mathscr{C}$.
- The 2-simplices in $\mathscr{C}$ are triples $x \leqq y \leqq z$:

  $$\begin{array}{ccccc} & & z & & \\ & \nearrow & & \nwarrow & \\ x & & \longrightarrow & & y \end{array}$$

  Let [2] be the poset $\{0, 1, 2\}$ in the usual order. Then a 2-simplex is equivalent to an order-preserving map $[2] \to \mathscr{C}$.
- The 3-simplices in $\mathscr{C}$ are quadruples $x \leqq y \leqq z \leqq w$:

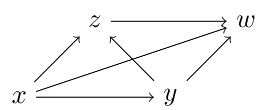

- ...
- The n-simplices in $\mathscr{C}$ are sequences $x_0 \leqq \cdots \leqq x_n$. Let $[n]$ be the poset $\{0, \ldots, n\}$ in the usual order. Then an n-simplex is equivalent to an order-preserving map $[n] \to \mathscr{C}$.

Thus, we can think of the posets $[n]$ as being the "standard categorical n-simplices" (with vertices given by the integers $0, \ldots, n$), and we can think of n-simplices in a poset $\mathscr{C}$ as figures parametrized by the standard n-simplex.

Exercise: A simplex $x_0 \leqq \cdots \leqq x_n$ is said to be **nondegenerate** if $x_i < x_{i+1}$ (strictly) for $i = 0, \ldots, n-1$. How many 3-simplices does $\{0, 1, 2, 3, 4\}$ have? How many are nondegenerate?

3. The Definition of Homology

3.1. **A Motivating Example.** Before diving into the formal definitions and accompanying algebra, consider the following geometric example. Let $X = \mathbb{R}^2 \setminus \{0\}$ be the punctured plane. From before, we know $X \simeq S^1$ via the deformation shrinking each vector v to $v/||v||$ over time, so that $\pi_1(X) \cong \pi_1(S^1) \cong \mathbb{Z}$. This is a reflection of the fact that there is a hole in X, and hence homotopically nontrivial loops exist, like $\gamma(t) = (\cos(2\pi t), \sin(2\pi t))$.

Let's describe another way of detecting the hole at 0. Note that γ above can also be thought of as the inclusion map $S^1 \to X$. Let

$$D^2 := \{(x, y) \in \mathbb{R}^2 \mid x^2 + y^2 \leqq 1\}$$

be the two-dimensional closed unit disc. The fact that γ is homotopically nontrivial implies that there is no continuous map $D^2 \to X$ whose restriction to $\partial D = S^1$ is γ; indeed, if there were such a $\Gamma : D^2 \to X$, then the map $h(x, t) = \Gamma((1-t)x)$ would be a homotopy of γ to a point, taking place in X. Geometrically, all we are saying is that you cannot map a disc continuously into $\mathbb{R}^2 \setminus \{0\}$ while requiring its bounday to be included into $\mathbb{R}^2 \setminus \{0\}$ as S^1.

Thus, the existence of a hole in $X = \mathbb{R}^2 \setminus \{0\}$ implies there is a closed figure $S^1 \to X$ (a **cycle**) that is not the **boundary** map of any figure $D^2 \to X$. It is this geometric situation that is formalized in the definition of homology.

Exercise: Generalize the above to $S^n \subset R^{n+1} \setminus \{0\}$. You may assume $\pi_n(S^n) \cong \mathbb{Z}$.

3.2. **The Formal Definition.**

"Algebra is what we do when we can't see what's going on geometrically."

Or at the very least, algebra is something we can fall back on when our vision (of, say, 15-dimensional shapes in a 32-dimensional space) fails us. We're going to formalize the situation described above for posets and generalize it to all dimensions. Accordingly, the definitions are somewhat abstract. The newcomer is urged to keep the geometric picture in mind going forward.

Suppose $\mathscr{C}$ is a poset. We want to interpret the situation above in terms of the poset structure of $\mathscr{C}$. Since the disc D^2 is topologically the same as a 2-simplex, the figure $D^2 \to X$ will be analgous to $[2] \to \mathscr{C}$, i.e. a 2-simplex in $\mathscr{C}$. But how do we make sense of $S^1 \to X$?

The boundary of Δ_2^{top} is a triangle without interior, which is definitely not a simplex. It is, however, a union or "sum" of the simplices corresponding to the sides of Δ_2^{top}. How do we introduce this into the context above? Easy: we allow ourselves to *formally* add simplices together. It is critical that simplices have only finitely many sides – that's what allows us to stay in the realm of algebra, rather than moving into analysis.

In general, suppose A is a set. The **free abelian group on** A, denoted $F(A)$ is the abelian group whose elements are formal sums

$$\sum_{i=1}^{n} c_i a_i$$

where n ranges over the natural numbers, c_i over the integers, and a_i over elements of A. This is the group that one gets if he/she just starts formally adding the elements of A together without introducing any relations other than those forced by the axioms of an abelian group (e.g. associativity, commutativity, etc.) – hence the name "free".[3]

Now, so that we can talk about the boundary of a simplex $[n+1] \to \mathscr{C}$, define:

$$C_n(\mathscr{C}) := \text{ free abelian group generated by all } n\text{-simplices } [x_0 \leqq \cdots \leqq x_n] \text{ in } \mathscr{C}$$

For $n < 0$, define $C_n(\mathscr{C}) := 0$. The elements of $C_n(\mathscr{C})$ for $n > 0$ are formal sums

$$\sum_{j=1}^{m} a_j [x_0^j \leqq \cdots \leqq x_n^j]$$

(where $a_j \in \mathbb{Z}$) and are called n**-chains**. Geometrically, an n-chain can be thought of as a figure built out of finitely many n-simplices, together with an integer weight on each simplex. Note, just keeping track of which simplices occur is not enough to determine an n-chain; we need to keep track of the number of times a simplex is "repeated". For example, what is the geometric difference between the simplex σ and the chain $\sigma + \sigma$?

Exercise: Draw some pictures of chains in a space or poset. How would you write down something like the 2-sphere?

As observed above, the boundary of a simplex should be some kind of sum of its sides. Since we now understand the sum of simplices in $\mathscr{C}$, it remains to understand 'side'. By way of example, suppose $x \leqq y \leqq z$ is 2-simplex. The three sides of this simplex are obtained by omitting each of x, y, and z in succession:

2-simplex	x-side	y-side	z-side
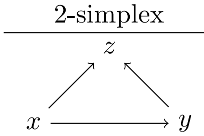	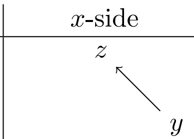	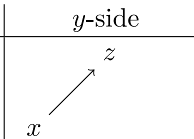	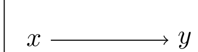
$x \leqq y \leqq z$	$y \leqq z$	$x \leqq z$	$x \leqq y$

[3]If one wants to be formal about it, this is the group whose elements are functions $c : A \to \mathbb{Z}$ such that for all but finitely many $a \in A$, $c_a := c(a) = 0$. The function c is thought of as specifying the coefficient attached to each element $a \in A$. Addition is componentwise, the 0 function is the identity, and $(-c)(a) := -[c(a)]$. Equivalently, $F(A)$ is the A-fold direct sum of copies of $\mathbb{Z}$, $\bigoplus_A \mathbb{Z}$.

Thus, for a general n-simplex $\sigma = [x_0 \leqq \cdots \leqq x_n]$, we define the ith side of σ to be the $(n-1)$ simplex obtained by omitting the ith vertex x_i. This is denoted

$$[x_0 \leqq \cdots \leqq \widehat{x_i} \leqq \cdots \leqq x_n].$$

The final piece to consider is orientation. Typically, the boundary of a reasonable subspace (e.g. a manifold) of $\mathbb{R}^n$ is oriented cyclically or with some condition on the normal vector. In the above table for the 2-simplex $\sigma = [x \leqq y \leqq z]$, the y-side is oriented opposite to the standard counterclockwise orientation. Thus, to get a cyclic orientation on the boundary of σ, we formally turn the y-side around by taking its additive inverse (we're in a group!). Define

$$\partial[x \leqq y \leqq z] := [y \leqq z] - [x \leqq z] + [x \leqq y].$$

In general, we define the **boundary of an n-simplex** by taking the alternating sum of sides

$$\partial[x_0 \leqq \cdots \leqq x_n] := \sum_{i=0}^{n} (-1)^i [x_0 \leqq \cdots \leqq \widehat{x_i} \leqq \cdots \leqq x_n].$$

The signs correspond to giving the boundary of the simplex a consistent orientation.

Exercise: Check that for the topological 3-simplex $\Delta \subset \mathbb{R}^3$ defined as the convex hull of $(0,0,0)$, $(1,0,0)$, $(0,1,0)$, and $(0,0,1)$, these signs orient $\partial\Delta$ with the outward normal vector.

Just as we may define a linear map between vector spaces by specifying values on a basis, so too many we define an additive map out of $F(A)$ by specifying a value on each generator $a \in A$. Explicitly, if B is an abelian group and A is a set, then a set map $\varphi : A \to B$ extends uniquely to a group homomorphism $F(A) \to B$ via the formula

$$\varphi(\sum c_j a_j) := \sum c_j \varphi(a_j).$$

Exercise: Check this claim.

Thus, there are induced **boundary maps** $\partial : C_n(\mathscr{C}) \to C_{n-1}(\mathscr{C})$ for each positive integer n. For $n \leqq 0$, we define $\partial := 0$, the zero homomorphism. For $n > 0$, the additivity of ∂ allows it to capture the "cancellations" that can occur when two simplices Δ_1 and Δ_2 in a chain share a boundary simplex σ. As a result, the ∂ map on chains often gives the "geometric boundary" of the figure determined by a chain. See the following exercise.

Exercise: Suppose $x_0 \leqq \cdots \leqq x_n$. Show that the boundary of the 1-chain $[x_0 \leqq x_1] + \cdots + [x_{n-1} \leqq x_n]$ is the boundary of $[x_0 \leqq x_n]$:

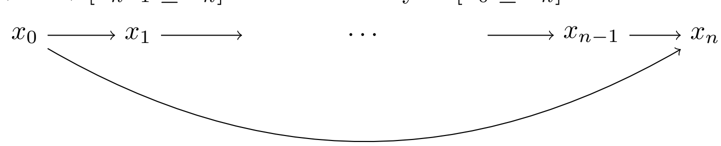

Thus, the "interior" 0-simplices $[x_1], \ldots, [x_{n-1}]$ cancel each other out.

We are poised to reinterpret the geometry in this algebraic setting. For each n, an n-chain is called a **boundary** if it is the boundary of an $(n+1)$-chain. An n-chain is called a **cycle** if it has no boundary, and therefore "closes up". Said differently, the n-boundaries are the elements of $\mathrm{im}(\partial : C_{n+1}(\mathscr{C}) \to C_n(\mathscr{C}))$, and the n-cycles are the elements of $\ker(\partial : C_n(\mathscr{C}) \to C_{n-1}(\mathscr{C}))$.

Suppose c is an n-chain that is a cycle, but not a boundary. In analogy to the geometric situation of $S^1 \subset \mathbb{R}^2 \setminus \{0\}$, we think of c as detecting a "hole": c is tracing out a closed figure that cannot be filled in by a figure one dimension higher. Since we're dealing with abelian groups, we could try to form the quotient (n-cycles)/(n-boundaries) to measure the extent to which there are cycles that are not boundaries. There is the following important result, which tells us that this quotient makes sense.

Proposition. *For each $n \geqq 2$, the composite $C_n(\mathscr{C}) \overset{\partial}{\to} C_{n-1}(\mathscr{C}) \overset{\partial}{\to} C_{n-2}(\mathscr{C})$ is the zero homomorphism, i.e. $C_*(\mathscr{C})$ is a chain complex. It follows that there is an inclusion $im(\partial : C_n(\mathscr{C}) \to C_{n-1}) \subseteqq ker(\partial : C_{n-1}(\mathscr{C}) \to C_{n-2})$ and the quotient $ker\partial/im\partial$ is well-defined.*[4]

Proof. Exercise. It is enough to check on simplices, since they generate the entire group of chains. □

We (finally) define the **homology groups** $H_n(\mathscr{C})$ of a poset $\mathscr{C}$ by the formula

$$H_n(\mathscr{C}) := \frac{\ker\Big(\partial : C_n(\mathscr{C}) \to C_{n-1}(\mathscr{C})\Big)}{\mathrm{im}\Big(\partial : C_{n+1}(\mathscr{C}) \to C_n(\mathscr{C})\Big)} = \frac{n\text{-cycles}}{n\text{-boundaries}}.$$

In particular, for $n = 0$ we have $(\partial : C_0(\mathscr{C}) \to C_{-1}(\mathscr{C})) = 0$, so that every 0-chain is a 0-cycle and

$$H_0(\mathscr{C}) = \frac{C_0(\mathscr{C})}{\mathrm{im}\Big(\partial : C_1(\mathscr{C}) \to C_0(\mathscr{C})\Big)} \overset{\text{Def.}}{=} \mathrm{coker}\Big(\partial : C_1(\mathscr{C}) \to C_0(\mathscr{C})\Big).$$

For $n < 0$, the groups of n-cycles and n-boundaries are both just 0, and $H_n(\mathscr{C}) = 0$.

Retracing our steps, we see that the definition of $H_n(\mathscr{C})$ is *algebraic*, so that homology groups can be (and often are) manipulated without any geometric considerations whatsoever. Despite this fact, the previous discussion should make it clear that homology groups are – in some sense – detecting holes in our space of various (co)dimension. In broad strokes, we have defined combinatorial notions of cycles (closed figures) and boundaries, and the existence of a cycle that is not the (algebraic) boundary of a chain one dimension higher, i.e. a nontrivial element in homology, is (algebraically) capturing a geometric condition for a hole in the space. (see section 3.1!)

We conclude this section with a useful technical fact. As always, assume $\mathscr{C}$ is a poset. Define the group $C_n^<(\mathscr{C})$ of **nondegenerate n-chains on $\mathscr{C}$** to be the free abelian group generated by strictly increasing n-chains $[x_0 < \cdots < x_n]$. It

[4]If the reader was unsure of the sign conventions on boundaries, this should be taken as justification. The alternating sum makes the algebra work out.

is a subgroup of $C_n(\mathscr{C})$, and since each face of a nondegenerate simplex is nondegenerate, it follows that the boundary map $\partial : C_n(\mathscr{C}) \to C_{n-1}(\mathscr{C})$ restricts to a homomorphism $\partial : C_n^<(\mathscr{C}) \to C_{n-1}^<(\mathscr{C})$. We therefore obtain a sub-chain complex $C_*^<(\mathscr{C}) \subset C_*(\mathscr{C})$, and can define the homology groups

$$H_n^<(\mathscr{C}) := \frac{\ker\Big(\partial : C_n^<(\mathscr{C}) \to C_{n-1}^<(\mathscr{C})\Big)}{\operatorname{im}\Big(\partial : C_{n+1}^<(\mathscr{C}) \to C_n^<(\mathscr{C})\Big)}.$$

Fact. *The homology groups $H_n(\mathscr{C})$ and $H_n^<(\mathscr{C})$ are isomorphic.*

We might prove this result at a later time. For the time being, we will not concern ourselves with this precise details; the intuition is that degenerate simplices are geometrically trivial, and that homology shouldn't care about them.

Here is the importance of the isomorphism $H_n \cong H_n^<$: it turns out H_n is easier to prove things about because we have a more inclusive notion of chain, while $H_n^<$ is easier to compute because there are fewer chains to worry about. This is typical in algebraic topology: often we have two descriptions of a single object or construction, one well-suited for theory, but unwieldy for computation, the other streamlined for computation, but hard to prove things about. As a result, there is sometimes a sharp divide between the theory and the practice (calculations).

A final remark: if you repeat the construction of homology just outlined for the singular simplices of a topological space X, you get something called the **singular homology of X**. See *Concise* or Hatcher for more details.

4. Applications and First Properties of Homology

Now that we've defined homology and given some indication of what it is measuring, it's time to look at some examples and to analyze its properties.

4.1. **Sample Calculations.** First, some examples. We will always compute the homology of a poset $\mathscr{C}$ in concrete cases using $H_n^<(\mathscr{C})$ (i.e. using nondegenerate simplices $[x_0 < \cdots < x_n]$), since there is less data to worry about.

Example: Suppose $\mathscr{C} = *$ is a poset with exactly one element. Then $C_n(\mathscr{C}) = 0$ whenever $n > 0$ because there are no relations $x < y$ in $\mathscr{C}$. Said differently, we cannot fit in any chains of positive length. On the other hand, there is exactly one 0-simplex, namely $[*]$. Therefore $C_0(*) \cong \mathbb{Z}$ and

- $H_0(*) = \ker(C_0(*) \to 0)/\operatorname{im}(0 \to C_0(*)) = C_0(*)/0 \cong \mathbb{Z}$.
- For $n > 0$, $C_n(*) = 0$ implies that $\ker(C_n(*) \to C_{n-1}(*)) = 0$, and hence $H_n(*) = 0/\operatorname{im}(C_{n+1}(*) \to C_n(*)) \cong 0$.

Exercise: Prove that in general, $H_0(\mathscr{C}) \cong \mathbb{Z}^{\#\text{ of conn. components}}$. Hint: write down ∂ of a generic 1-chain $\sum_i (\pm 1)[x_1^i < x_2^i]$ and set it equal to $[y] - [x]$. By considering what cancellations must occur, show that two 0-cycles $[x]$ and $[y]$ differ by a 0-boundary if and only if they are in the same (path) component of the poset.

Example: Suppose $\mathscr{C}$ is the poset below.

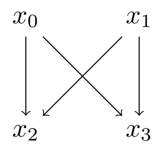

This is a model for S^1. We have drawn the poset in this configuration (in layers with all arrows pointing down) to emphasize that there are no strictly increasing sequences $a_0 < \cdots < a_n$ for $n > 1$. Now:

- We only have nondegenerate n-simplices for $n = 0, 1$, hence for $n \geqq 2$, $C_n^<(\mathscr{C}) \cong 0$ and $H_n(\mathscr{C}) \cong 0$.
- $H_0(\mathscr{C})$ is given using generators and relations by

$$\Big\langle [x_0], [x_1], [x_2], [x_3] \;\Big|\; [x_2] - [x_0], [x_2] - [x_1], [x_3] - [x_0], [x_3] - [x_1] \Big\rangle$$

  i.e. we have four free generators $[x_0], \ldots, [x_3]$ that are identified in pairs. Therefore $H_0(\mathscr{C}) \cong \mathbb{Z}$. This can be seen as a special case of the preceding exercise $H_0(\mathscr{C}) \cong \mathbb{Z}^{\#\text{ of conn. components}}$.
- $H_1(\mathscr{C}) = \ker(\partial)/\mathrm{im}(\partial) \cong \ker(\partial : C_1^<(\mathscr{C}) \to C_0^<(\mathscr{C}))$ because there are no nondegenerate 2-simplices and hence $\mathrm{im}\,\partial = 0$. The rest of the computation boils down to "linear algebra over $\mathbb{Z}$": we need to find the kernel of the $(\mathbb{Z})$-linear map

$$\partial : \Big\langle [x_0 < x_2], [x_1 < x_2], [x_0 < x_3], [x_1 < x_3] \Big\rangle \to \Big\langle [x_0], [x_1], [x_2], [x_3] \Big\rangle.$$

  It has a corresponding 4×4 matrix with respect to these bases. We can do elementary row and column operations **using only integer coefficients** to show that the kernel is $\cong \mathbb{Z}$.

Exercise: Write down the matrix and do the calculation. Why can we use a matrix to represent this homomorphism? (Hint: we have free generators $\simeq$ a basis.) Why is it valid to use row and column operations in this way? (Hint: they correspond to composition/precomposition with certain invertible transformations, and therefore define isomorphisms between the image/kernel of the map and something more recognizable.) What is a generator for $H_1(\mathscr{C})$?

It is worth emphasizing that homology is computable; we can even program computers to do the calculations. This makes it a very useful invariant, even if it is not as sensitive as homotopy.

Going forward, we will be interested in computing the homology of poset models for the n-sphere, $n > 0$. It turns out that for a certain $\mathscr{C}$ modeling S^n, we have:

$$H_k(\mathscr{C}) \cong \begin{cases} \mathbb{Z} & k = 0, n \\ 0 & \text{else} \end{cases}.$$

We'll prove this using a piece of machinery called the *Mayer-Vietoris* sequence, but it is also possible to show this by direct inspection of the chains.

Exercise: Suppose $\mathscr{C}$ is a poset. We define the **poset suspension $\Sigma\mathscr{C}$ of $\mathscr{C}$** to be the poset obtained by adjoining two new points $\{N_1, N_2\}$ and declaring that $N_i > x$ for all $x \in \mathscr{C}$ and $i = 1, 2$. In other words, we add two new (incomparable) points that are bigger than everything in $\mathscr{C}$. By a direct inspection of chains in $\mathscr{C}$ and $\Sigma\mathscr{C}$, prove directly from the definitions that $H_n(\Sigma\mathscr{C}) \cong H_{n-1}(\mathscr{C})$

for $n > 1$. Can this be true for $n = 1$? (Hint: what happens for $\mathscr{C} = *$?)[5]

If we accept that $\overbrace{\Sigma \cdots \Sigma}^{n \text{ times}}(* \coprod *)$ is a model for S^n (for all $n > 0$), then the desired calculation follows by induction.

4.2. **Some Basic Theoretical Properties of Homology.** We now develop some of the machinery we'll need to work effectively with homology.

Lemma. *For each $n \geqq 0$, H_n extends to a functor* ***Poset*** $\to$ ***Ab****, going from the category of posets and order-preserving maps to the category of abelian groups and group homomorphisms.*

Proof. We've already defined $H_n(\mathscr{C})$, so what we need to do now is explain how an order-preserving map $\mathscr{C} \to \mathscr{D}$ between posets induces a group homomorphism $H_n(\mathscr{C}) \to H_n(\mathscr{D})$. If $f : \mathscr{C} \to \mathscr{D}$ is order-preserving, then it preserves chains defined using $\leqq$ (but *not* those defined with $<$). For each $n \geqq 0$, the assignment

$$C_n(f)[x_0 \leqq \cdots \leqq x_n] = [f(x_0) \leqq \cdots \leqq f(x_n)]$$

between n-simplices extends by additivity to a homomorphism

$$C_n(f) : C_n(\mathscr{C}) \to C_n(\mathscr{D})$$

between the groups of n-chains. For $n < 0$, we define $C_n(f) := 0 : 0 \to 0$. The sequence of homomorphisms $(C_n(f))_{n \in \mathbb{Z}}$ has the property that, for every $n \in \mathbb{Z}$, $\partial \circ C_n(f) = C_{n-1}(f) \circ \partial$, i.e. the square

$$\begin{array}{ccc} C_n(\mathscr{C}) & \xrightarrow{C_n(f)} & C_n(\mathscr{D}) \\ \Big\downarrow \partial & & \Big\downarrow \partial \\ C_{n-1}(\mathscr{D}) & \xrightarrow{C_{n-1}(f)} & C_{n-1}(\mathscr{D}) \end{array}$$

commutes (covering up the ith factor and then applying f is the same as applying f and then covering up the ith factor). It follows that for each n, $C_n(f)$ maps n-cycles to n-cycles, and n-boundaries to n-boundaries. Therefore $C_n(f)$ descends to a homomorphism

$$H_n(f) : H_n(\mathscr{C}) \to H_n(\mathscr{D})$$

between homology groups. Explicitly, if $\alpha \in H_n(\mathscr{C})$ is represented by the n-cycle $\sum_i a_i[x_0^i \leqq \cdots \leqq x_n^i]$, then $H_n(f)(\alpha) \in H_n(\mathscr{D})$ is the homology class represented by $\sum_i a_i[f(x_0^i) \leqq \cdots \leqq f(x_n^i)]$. It follows that $H_n(g \circ f) = H_n(g) \circ H_n(f)$ and $H_n(\mathrm{id}) = \mathrm{id}$, so H_n is indeed a functor. $\square$

Exercise: Check the details.

Observe that this construction really breaks down into two steps. First, we defined the maps $C_*(f)$, and then we passed to homology. Formalizing what we see above, we say that a sequence $(\varphi_n : A_n \to B_n)$ of homomorphisms between the component groups of two chain complexes (A_n) and (B_n) is a **chain map** if and only if $\partial \circ \varphi_n = \varphi_{n-1} \circ \partial$ for all n. Chain complexes (of abelian groups) and chain

[5] This is something reflecting the difference between *reduced* and *unreduced* homology theories.

maps form a category, which we shall denote **Ch**. Homology groups make sense for any chain complex (by the condition $\partial \circ \partial = 0$), and the arguments above show that for each n, there is an nth homology functor $H_n : \mathbf{Ch} \to \mathbf{Ab}$ sending a chain complex A_* to its nth homology, and a chain map (f_n) to the map on homology induced by the homomorphism f_n. The homology functors for posets are really a composite of two functors $\mathbf{Poset} \to \mathbf{Ch} \to \mathbf{Ab}$, the first sending a poset to its chains, and the second extracting homology.

Next, we describe a criterion for when two order-preserving maps $f, g : \mathscr{C} \rightrightarrows \mathscr{D}$ induce the same (equal) maps on homology. Indeed, the assumption $f \neq g$ does *not* imply $H_n(f) \neq H_n(g)$.

Lemma. *If $f, g : \mathscr{C} \rightrightarrows \mathscr{D}$ are order-preserving maps between posets $\mathscr{C}$ and $\mathscr{D}$, and $f(x) \leqq g(x)$ for every $x \in \mathscr{C}$, then $H_n(f) = H_n(g) : H_n(\mathscr{C}) \to H_n(\mathscr{D})$ for every n.*

Proof. This is actually quite geometric. Suppose we have parallel order preserving maps $f, g : \mathscr{C} \rightrightarrows \mathscr{D}$ such that $f(x) \leqq g(x)$ for every $x \in \mathscr{C}$. Recall this implies that if we regard $\mathscr{C}$ and $\mathscr{D}$ as Alexandroff spaces, then there is a homotopy $h(x, t) : \mathscr{C} \times I \to \mathscr{D}$ from f to g defined by

$$h(x,t) := \begin{cases} f(x) & t < 1 \\ g(x) & t = 1 \end{cases} .$$

Regard this homotopy h as fixed data. Given any n-simplex $\sigma : \Delta \to \mathscr{C}$ in $\mathscr{C}$, we now obtain a figure (depending only on σ)

$$P(\sigma) : \Delta \times I \overset{\sigma \times I}{\to} \mathscr{C} \times I \overset{h}{\to} \mathscr{D}$$

$P(\sigma)$ parametrizes an "$(n+1)$-prism" in $\mathscr{D}$. The face corresponding to $t = 0$ is $C_n(f)\sigma = f \circ \sigma$, while the face at $t = 1$ is $C_n(g)\sigma = g \circ \sigma$. Thus, $P(\sigma)$ can be thought of as "witnessing" the deformation of $C_n(f)\sigma$ to $C_n(g)\sigma$. The idea now is to see what this geometric situation implies about the algebraic relationship between $C_n(f)\sigma$ and $C_n(g)\sigma$. Observe that the boundary of a prism $\Delta \times I$ can be decomposed into its faces as below

$$\partial(\Delta \times I) = \Big(\Delta \times \{0\}\Big) \cup \Big(\Delta \times \{1\}\Big) \cup \Big(\partial\Delta \times I\Big).$$

If we hit this with the prism map $P(\sigma) = h \circ (\sigma \times I) : \Delta \times I \to \mathscr{D}$, and remember that $P(\tau)$ is the prism "witnessing the deformation of $C_n(f)\tau$ to $C_n(g)\tau$", we expect the following geometric situation: (draw the picture)

$$\partial P(\sigma) = C_n(f)\sigma \cup C_n(g)\sigma \cup P\partial(\sigma).$$

Thus, up to signs related to orientation, we expect a formula of the form

$$\pm C_n(f)\sigma \pm C_n(g)\sigma = \pm \partial P(\sigma) \pm P\partial(\sigma)$$

for some suitably defined "prism operator" on chains.

The technical stumbling block for all this is that chains are built out of simplices. Thus, we need to find a way of formally subdividing prisms so that we get a useful formula as above. We define $P_n : C_n(\mathscr{C}) \to C_{n+1}(\mathscr{D})$ by taking

$$P_n[x_0 \leqq \cdots \leqq x_n] := \sum_{i=0}^{n} (-1)^i [f(x_0) \leqq \cdots \leqq \overbrace{f(x_i) \leqq g(x_i)}^{\text{repeat } i} \leqq \cdots \leqq g(x_n)]$$

on n-simplices, and extending by additivity.[6] As with the boundary maps ∂, the correctness of this definition can be justified by the fact that *it makes the algebra work*. Indeed, with this construction, we have the following equality of group homomorphisms $C_n(\mathscr{C}) \to C_n(\mathscr{D})$:

$$C_n(f) - C_n(g) = \partial \circ P_n + P_{n-1} \circ \partial.$$

This is great news. Suppose that $\alpha \in C_n(\mathscr{C})$ is a cycle. Then $\partial(\alpha) = 0$, hence $C_n(f)(\alpha) - C_n(g)(\alpha) = \partial(P_n(\alpha))$ is a boundary. Thus, when we pass to homology (setting all boundaries to 0), the cycles $C_n(f)(\alpha)$ and $C_n(g)(\alpha)$ will represent the same homology class, i.e. are identified in $H_n(\mathscr{D})$. Since every element of $H_n(\mathscr{C})$ is represented by a cycle, it follows $H_n(f) = H_n(g)$. □

Exercise: Check that $C_n(f) - C_n(g) = \partial \circ P_n + P_{n-1} \circ \partial$ and the details in the argument that $H_n(f) = H_n(g)$.

Exercise: Draw pictures of $P(\sigma)$, where σ is a 1-simplex or 2-simplex. The point is that as we increase i in the alternating sum defining $P(\sigma)$, we "sweep out" simplices that fill up the prism.

Exercise: (If you are feeling brave.) Regard the topological n-simplex Δ_n^{top} as the convex hull $\mathrm{conv}(e_1, \dots, e_{n+1})$ of the standard basis vectors $e_1, \dots, e_{n+1} \in \mathbb{R}^n$, and the prism $P(\Delta_n^{top})$ as the subset $\Delta_n^{top} \times [0,1] \in \mathbb{R}^{n+2}$. Show that

$$P(\Delta_n^{top}) = \bigcup_{i=1}^{n+1} \mathrm{conv}\Big((e_1, 0), \dots, (e_i, 0), (e_i, 1), \dots, (e_{n+1}, 1)\Big).$$

Show, furthermore, that the points $(e_1, 0), \dots, (e_i, 0), (e_i, 1), \dots, (e_{n+1}, 1)$ are in general position (and hence determine an $(n+1)$-simplex) and that distinct simplices in this union can only intersect on their boundaries.

In the above, we saw how a homotopy h between the maps $f, g : \mathscr{C} \rightrightarrows \mathscr{D}$ gave rise to a "prism operator" $P : C_*(\mathscr{C}) \to C_{*+1}(\mathscr{D})$. P is reflecting $h : f \Rightarrow g$ on the level of chains, and accordingly can be thought of as a "homotopy" $C_*(f) \Rightarrow C_*(g)$ between the maps on chains. This leads to the following abstraction: suppose A_* and B_* are chain complexes, and $f_*, g_* : A_* \rightrightarrows B_*$ are a parallel pair of chain maps. A **chain homotopy** $P : f_* \Rightarrow g_*$ is a sequence of degree-increasing homomorphisms $(P_n : A_n \to B_{n+1})$ such that $f_n - g_n = \partial \circ P_n + P_{n-1} \circ \partial$ for all n.

Example: Here is an application of the previous lemma. Suppose $\mathscr{C}$ is a poset and assume $\mathscr{C}$ has a maximum (or analogously, a minimum) M. Let's take $f : \mathscr{C} \to \mathscr{C}$ to be the identity map $f(x) = x$, and $g : \mathscr{C} \to \mathscr{C}$ to be constant to M. Then $f(x) \leqq g(x)$ for all x, so the lemma applies and we conclude $H_*(f) = H_*(g)$.

Since $f = \mathrm{id}$, it follows $H_*(f) = \mathrm{id}$ by functoriality. What about $H_*(g)$? This is the map induced by the assignment $[x_0 \leqq \cdots \leqq x_n] \mapsto [M \leqq \cdots \leqq M]$ on simplices. Claim: for $n > 0$, $H_n(g)$ is the 0 homomorphism. Granting this, the only way the 0 homomorphism is the identity map is if $H_n(\mathscr{C}) = 0$ for $n > 0$. As for $n = 0$, the existence of a maximum M implies that $\mathscr{C}$ is connected, hence $H_0(\mathscr{C}) \cong \mathbb{Z}$.

[6] see the following exercises.

Exercise: Prove that $H_n(g) = 0$ for $n > 0$.

4.3. **Preview.** Here is a sneak preview of some (possible) things to come.

So far, we've introduced two invariants: homotopy and homology. The two are related. Geometrically, both measure the holes or twisting in a space, albeit in different ways. In terms of their relative strength, the preceding lemma involving the prism operator P showed that homology is a homotopy invariant. It is also true that homology is a *weak* homotopy invariant, but this is harder to prove. Thus, homology is a coarser invariant than homotopy, and is actually strictly so.

That said, homotopy and homology are not irreconcilable. Suppose X is a space or poset. In general, we cannot expect that $\pi_1(X) \cong H_1(X)$, because homology groups are always abelian, whereas fundamental groups need not be. However, the next best thing happens: $H_1(X)$ is (canonically) isomorphic to the **abelianization** of $\pi_1(X)$. Intuitively, the abelianization of a group G is the group obtained from G by formally requiring all elements to commute. It is obtained from G by introducing the relations $gh = hg$ for all pairs $g, h \in G$, and can be regarded as the "most efficient" way of turning a general group into an abelian group. Thus, while $\pi_1(X) \ncong H_1(X)$, they are as close as one could reasonably expect.