AN INTRODUCTION TO SINGULAR HOMOLOGY AND THE HAIRY BALL THEOREM

ILAN HAAS

Abstract. The goal of this paper is to provide a readable introduction to singular homology, an essential topic in Algebraic Topology. We provide intuition and motivation for the algebraic structures involved, cover several basic properties, and finally show an application of homology, using properties of degree to prove the Hairy Ball Theorem.

Contents

1. Thinking about Holes 1
2. Formalizing the Idea of Filling in Loops 2
3. Introduction to Exact Sequences and Homology 5
4. Singular Homology 6
5. Properties of Homology 8
6. The Homology of Spheres 11
7. An Introduction to Degree 12
8. Proof of the Hairy Ball Theorem 13
Acknowledgments 15
References 16

1. Thinking about Holes

As stated by Chad Berkich in a previous paper for the University of Chicago REU program [2], "Holes are interesting to measure because the mechanisms used to measure them, their homology and cohomology groups, are algebraic topological invariants." We note that holes themselves are the topological invariants. There are multiple algebraic methods for detecting and measuring them, including fundamental groups, homotopy groups, and homology groups. While fundamental groups and homotopy groups are elegant approaches to measuring holes (as introduced in textbooks like Hatcher [1]), they become difficult to compute for higher-order spaces. Homology provides a more tractable alternative.

Consider a hole in a coin, depicted in Figure 1, considered as subset of $\mathbb{R}^2$. While we can clearly see the hole, specifying rigorous criteria for the existence of a hole is surprisingly challenging. We present several informal approaches that lead to different mathematical frameworks.

Approach 1: Connectivity after cutting. Draw a line from the boundary of the shape, through the interior, to another point on the boundary, then cut along

Date: August 2025.

this line. In a region without holes, this cut always produces two disconnected pieces. However, if a hole is present, certain cuts (such as the red line in Figure 1(a)) leave the shape connected. While intuitive, this approach does not correspond to homological hole detection.

Approach 2: Loop contractibility. Choose a basepoint x_0 in the space and consider loops that start and end at this point. In a space without holes, any such loop can be continuously contracted to the basepoint. When holes are present, some loops become "trapped" around the holes and cannot be contracted, as illustrated by the large green loop in Figure 1(a). This idea gives rise to the fundamental group and, when generalized, to homotopy groups.

Approach 3: Filling in a Loop. Draw a closed loop within the space. In a space without holes, the region enclosed by any loop can always be "filled in" and lies entirely within the space. When holes are present, some loops cannot be filled in because their interior intersects the hole. This geometric intuition leads directly to homology theory and is illustrated by the triangle in Figure 1(b).

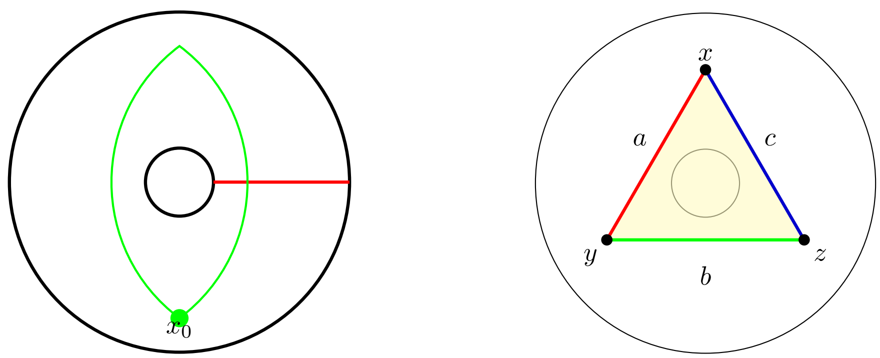

(A) Connectivity and loop approaches

(B) Boundary filling approach

FIGURE 1. Approaches to detecting holes in a topological space

2. Formalizing the Idea of Filling in Loops

To formalize the notion that "every loop can be filled in," we must first rigorously define what we mean by "all loops" in a topological space. Intuitively, a loop is any path through a space that begins and ends at the same point, regardless of whether it is curved, angular, or straight. More generally, multiple paths can combine to form a loop, as illustrated by the three edges of the triangle in Figure 1(b).

Since we are fundamentally concerned with paths, we begin by formalizing this concept.

Definition 2.1. Given a topological space X, any continuous map $\sigma : [0,1] \to X$ is called a *singular 1-simplex* or *path.*

To create loops by "adding" paths together, we need a formal algebraic structure for combining these maps.

Definition 2.2. Given a topological space X, the *singular 1-chain group* $C_1(X)$ is the free abelian group generated by all singular 1-simplices $\sigma : [0,1] \to X$.

Each element of $C_1(X)$ is a finite formal sum

$$\sum_i a_i \sigma_i,$$

where $a_i \in \mathbb{Z}$ and each σ_i is a singular 1-simplex. Such an element is called a *singular 1-chain.*

This algebraic structure allows us to formally "add" paths. The identity element corresponds to having zero copies of every path, and the inverse of a singular 1-simplex σ is just $-\sigma$, or "-1 copies of σ". It does *not* represent the same path traversed in the opposite direction; such a path would be considered an entirely unrelated object.

To determine when a singular 1-chain forms a loop, we observe that what matters are the endpoints: after traversing all paths in the chain, we should return to our starting point. Equivalently, all endpoints should "cancel out" in the formal sum. This requires us to track the 0-dimensional endpoints formally.

Definition 2.3. The *singular 0-chain group* $C_0(X)$ is the free abelian group generated by all points in X. A singular 0-simplex is a map $\sigma : \{0\} \to X$, which simply picks out a point in X.

We can now define a map that tracks the endpoints of our 1-chains.

Definition 2.4. The *boundary map* $\partial_1 : C_1(X) \to C_0(X)$ is defined on generators by

$$\partial_1(\sigma) = \sigma(1) - \sigma(0)$$

for a singular 1-simplex $\sigma : [0,1] \to X$, and extended linearly to all of $C_1(X)$:

$$\partial_1\left(\sum_i a_i \sigma_i\right) = \sum_i a_i \partial_1(\sigma_i).$$

This map sends each path to the formal difference between its endpoint and starting point. The beauty of using subtraction is that endpoints naturally cancel in closed loops. For the triangle in Figure 1(b) with edges a, b, and c, we have:

$$\partial_1(a+b+c) = \partial_1(a) + \partial_1(b) + \partial_1(c) = (z-y) + (x-z) + (y-x) = 0.$$

Definition 2.5. A singular 1-chain $\gamma \in C_1(X)$ such that $\partial_1(\gamma) = 0$ is called a *1-cycle.*

We have now partially formalized "every cycle can be filled in." Next, we must formalize what it means to "fill in" a cycle.

Just as we needed to consider all possible loops, we now need to consider all possible 2-dimensional shapes that can fill them. These shapes will be continuous images of triangular regions.

Definition 2.6. The *singular 2-chain group* $C_2(X)$ is the free abelian group generated by all singular 2-simplices. A *singular 2-simplex* is a continuous map $\sigma : \Delta^2 \to X$, where Δ^2 a 2-simplex:

$$\Delta^2 = \left\{(t_0, t_1, t_2) \in \mathbb{R}^3 \mid t_0 + t_1 + t_2 = 1,\ t_i \geq 0\right\}.$$

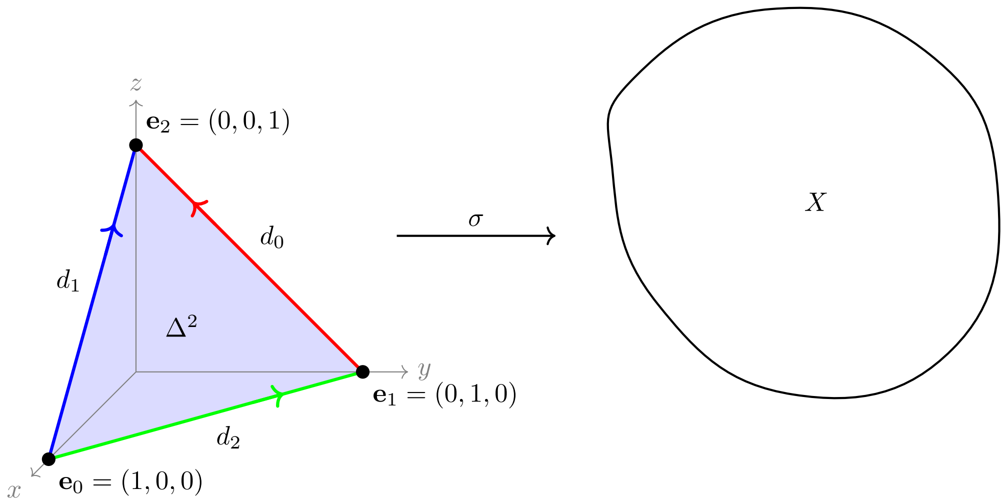

FIGURE 2. Singular 2-Simplex

Any 2-dimensional shape can be represented as an element of $C_2(X)$ by decomposing it into triangular pieces. Note that different triangulations of the same geometric shape yield different elements of $C_2(X)$, and even the same triangulation can be represented by different elements, since the same shapes with different orientations are considered as unrelated objects. The term "singular" reflects the fact that these maps may have singularities—the image need not resemble a triangle and could be collapsed to a line or point.

To formalize "filling in a cycle," we define a boundary map that relates 2-chains to their 1-dimensional boundaries.

Definition 2.7. The *boundary map* $\partial_2 : C_2(X) \to C_1(X)$ is defined on a singular 2-simplex $\sigma : \Delta^2 \to X$ by

$$\partial_2(\sigma) = \sigma|_{d_0} - \sigma|_{d_1} + \sigma|_{d_2},$$

where d_0, d_1, and d_2 are the three edges of Δ^2 with appropriate orientations, and $\sigma|_{d_i}$ denotes the restriction of σ to edge d_i.

The boundary map sends each 2-simplex to the alternating sum of its edges, respecting orientation. Crucially, orientation determines the signs in this alternating sum.

Definition 2.8. A *1-boundary* is an element of $\operatorname{im}(\partial_2) \subseteq C_1(X)$—that is, a 1-chain that is the boundary of some 2-chain.

We can now pose our central question in precise terms: is every 1-cycle a 1-boundary? A method to answer this question is not immediately obvious; the groups $C_1(X)$ and $C_2(X)$ are of monstrous size, beyond direct computation in all but the most trivial cases. Instead, we need to use algebraic tools. We have established a sequence of groups and homomorphisms:

$$C_2(X) \xrightarrow{\partial_2} C_1(X) \xrightarrow{\partial_1} C_0(X).$$

Since every 2-simplex maps to a 1-cycle under ∂_2, we have that every 1-boundary is a 1-cycle; in other words, $\operatorname{im}(\partial_2) \subseteq \ker(\partial_1)$. The question of whether cycles can always be filled reduces to whether this inclusion is an equality. The gap between $\ker(\partial_1)$ and $\operatorname{im}(\partial_2)$ is precisely what homology measures, leading us to later define homology groups as

$$H_1(X) = \ker(\partial_1)/\operatorname{im}(\partial_2).$$

3. Introduction to Exact Sequences and Homology

We begin with a short sequence of abelian groups and homomorphisms between them:

$$0 \xrightarrow{0} A \xrightarrow{\alpha} B \xrightarrow{\beta} C \xrightarrow{0} 0,$$

where 0 represents the trivial group, and the maps labeled 0 represent the zero homomorphisms.

Definition 3.1. A sequence of groups and homomorphisms is called an *exact sequence* if for every pair of adjacent maps,

$$\ker(\textbf{second map}) = \operatorname{im}(\textbf{first map}).$$

Let us analyze what exactness means for our short sequence. The image of the map $0 \to A$ is $\{e_A\}$, where e_A is the identity element of A. Since α is a homomorphism, we have $\alpha(e_A) = e_B$. The exactness condition

$$\operatorname{im}(0) = \{e_A\} = \ker(\alpha)$$

implies that α maps only the identity element to e_B, and therefore α is injective.

Similarly, for the latter part of the sequence $B \xrightarrow{\beta} C \to 0$, exactness requires

$$\operatorname{im}(\beta) = \ker(C \to 0) = C.$$

Since the map $C \to 0$ has kernel equal to all of C, we conclude that β is surjective.

Finally, the middle exactness condition gives us

$$\operatorname{im}(\alpha) = \ker(\beta).$$

By the first isomorphism theorem, this implies

$$C \approx B/\operatorname{im}(\alpha) \approx B/A,$$

where we identify A with its image $\operatorname{im}(\alpha)$ in B.

The importance of short exact sequences lies in their universality: any long exact sequence can be decomposed into short exact sequences. A long exact sequence has the form

$$\cdots \to A_{n+1} \xrightarrow{d_{n+1}} A_n \xrightarrow{d_n} A_{n-1} \to \cdots \to A_2 \xrightarrow{d_2} A_1 \xrightarrow{d_1} 0,$$

where $\operatorname{im}(d_{n+1}) = \ker(d_n)$ for all n.

Homology arises when we relax the exactness condition. Instead of requiring that $\operatorname{im}(d_{n+1}) = \ker(d_n)$, we consider sequences where $\operatorname{im}(d_{n+1}) \subset \ker(d_n)$.

Definition 3.2. A *chain complex* is a sequence of abelian groups and homomorphisms

$$\cdots \to A_{n+1} \xrightarrow{d_{n+1}} A_n \xrightarrow{d_n} A_{n-1} \to \cdots$$

such that $d_n \circ d_{n+1} = 0$ for all n. Equivalently, $\operatorname{im}(d_{n+1}) \subseteq \ker(d_n)$.

Definition 3.3. The ***nth homology group*** **of a chain complex is**

$$H_n = \ker(d_n)/\operatorname{im}(d_{n+1}).$$

When $\operatorname{im}(d_{n+1}) = \ker(d_n)$, the nth homology group is trivial, indicating exactness at position n. When $\operatorname{im}(d_{n+1})$ is smaller than $\ker(d_n)$, the homology group H_n is nontrivial, and its structure and size measure precisely how the sequence fails to be exact at the nth term of the chain complex.

4. Singular Homology

We now construct singular homology rigorously. This requires generalizing our earlier definitions to arbitrary dimensions, beginning with the notion of a simplex.

A simplex is the higher-dimensional generalization of a triangle: a point of dimension 0, a line segment of dimension 1, a triangle of dimension 2 (as shown in Figure 2), a tetrahedron of dimension 3, and so forth.

Definition 4.1. An ***n-simplex*** **is the convex hull of $n+1$ ordered, affinely independent points $\mathbf{v_0}, \mathbf{v_1}, \dots, \mathbf{v_n}$:**

$$\left\{ \sum_{i=0}^{n} t_i \mathbf{v_i} \mid t_i \geq 0, \sum_{i=0}^{n} t_i = 1 \right\},$$

denoted as $[\mathbf{v_0}, \mathbf{v_1}, \dots, \mathbf{v_n}]$. The ***standard n-simplex*** **Δ^n is**

$$\Delta^n = \left\{ (t_0, t_1, \dots, t_n) \in \mathbb{R}^{n+1} \mid t_i \geq 0, \sum_{i=0}^{n} t_i = 1 \right\} = [\mathbf{e_0}, \mathbf{e_1}, \dots, \mathbf{e_n}],$$

where $\mathbf{e_i}$ is the ith standard basis vector in $\mathbb{R}^{n+1}$.

For our purposes, space homeomorphic to Δ^n Note that an n-simplex has dimension n but requires $n+1$ vertices. The ordering of vertices is crucial in homology, as different orderings yield different orientations.

Definition 4.2. Given an n-simplex $[\mathbf{v_0}, \mathbf{v_1}, \dots, \mathbf{v_n}]$, a ***face*** **is a simplex defined by a non-empty subset of these vertices, with order preserved.**

We often refer specifically to the $(n-1)$-simplexes as faces, usually denoted as $[\mathbf{v_0}, \dots, \widehat{\mathbf{v_i}}, \dots, \mathbf{v_n}]$, where the pointy hat $\widehat{\mathbf{v_i}}$ indicates that $\mathbf{v_i}$ is omitted. A consequence of this is the natural ordering of these faces, the ith $(n-1)$-dimensional face being the face with the ith vertex removed. The vertices of each face inherit their ordering from the original simplex, as shown by the arrows in Figure 2.

Definition 4.3. Given a topological space X, a ***singular n-simplex*** **is a continuous map**

$$\sigma : \Delta^n \to X.$$

Definition 4.4. Given a topological space X, the ***singular n-chain group*** **$C_n(X)$ is the free abelian group generated by all singular n-simplices $\sigma : \Delta^n \to X$. Each element of $C_n(X)$ is a finite formal sum**

$$\sum_i a_i \sigma_i$$

where $a_i \in \mathbb{Z}$ and each σ_i is a singular n-simplex. Such an element is called a ***singular n-chain.***

As before, the boundary map connects chains of different dimensions by relating each simplex to its boundary.

Definition 4.5. The *boundary map* $\partial_n : C_n(X) \to C_{n-1}(X)$ is defined on a singular n-simplex $\sigma : \Delta^n \to X$ by

$$\partial_n(\sigma) = \sum_{i=0}^{n} (-1)^i \sigma \circ d_i,$$

where $d_i : \Delta^{n-1} \to \Delta$ is the canonical inclusion of the ith face $[v_0, ...\widehat{v_i}...v_n]$ onto Δ^n. We extend ∂_n linearly to all n-chains:

$$\partial_n \left(\sum_i a_i \sigma_i \right) = \sum_i a_i \partial_n(\sigma_i).$$

The term $\sigma \circ d_i$ could also be considered as $\sigma|_{[v_0,...\widehat{v_i}...v_n]}$, the restriction of σ to the ith face of Δ^n. The alternating signs take orientation into account.

$$\begin{aligned} \partial_2([v_0, v_1, v_2]) &= (-1)^2[\widehat{v_0}, v_1, v_2] + (-1)^1[v_0, \widehat{v_1}, v_2] + (-1)^0[v_0, v_1, \widehat{v_2}] \\ &= [v_1, v_2] - [v_0, v_2] + [v_0, v_1] \end{aligned}$$

The orientation of Δ^2 is counterclockwise, but the orientation of the sides are determined by their inherited order of vertices, which leads to the edge $[v_0, v_2]$ inheriting opposite orientation, as indicated by Figure 2. We recommend the explanation found in Hatcher [1] on page 105.

We can now define the fundamental structure of singular homology:

Definition 4.6. The *singular chain complex* of a topological space X is the sequence

$$\cdots \longrightarrow C_{n+1}(X) \xrightarrow{\partial_{n+1}} C_n(X) \xrightarrow{\partial_n} C_{n-1}(X) \longrightarrow \cdots \longrightarrow C_1(X) \xrightarrow{\partial_1} C_0(X) \xrightarrow{\partial_0} 0.$$

Proposition 4.7. *For all n, we have $\partial_n \circ \partial_{n+1} = 0$, or equivalently,*

$$\operatorname{im}(\partial_{n+1}) \subseteq \ker(\partial_n).$$

Proof sketch. This follows from the fact that $\partial \circ \partial = 0$ for simplices, which can be verified by direct computation using the alternating sum formula, because each face appears exactly twice with opposite signs when computing $\partial_{n-1} \circ \partial_n$. □

This fundamental property ensures that this sequence of groups and homomorphisms is indeed a chain complex, allowing us to define homology groups that measure the degree to which a singular chain complex fails to be exact.

Definition 4.8. The n*th singular homology group* of X is

$$H_n(X) = \ker(\partial_n) / \operatorname{im}(\partial_{n+1}).$$

Elements of $\ker(\partial_n)$ are called n*-cycles*, and elements of $\operatorname{im}(\partial_{n+1})$ are called n*-boundaries*.

The nth homology group $H_n(X)$ captures the n-dimensional holes in X: cycles that cannot be filled in by $(n+1)$-chains—cycles that aren't boundaries.

5. Properties of Homology

There are quite a few nice properties of homology that make it a powerful topological invariant. I will give an overview and give proofs for some, but not all, of these statements.

Proposition 5.1. *Let X be a non-empty, path-connected topological space. Then $H_0(X) \approx \mathbb{Z}$.*

Proof. By definition,

$$H_0(X) = \ker(\partial_0)/\operatorname{im}(\partial_1) = C_0/\operatorname{im}(\partial_1),$$

since ∂_0 sends all of C_0 to 0. We define a map $I : C_0 \to \mathbb{Z}$ to be

$$I(\sum_i a_i\sigma_i) = \sum_i a_i,$$

called the *index* of the chain (this map is also called the *augmentation map*). In order to show that $H_0(X) \approx \mathbb{Z}$, we only need to show that $\ker(I) = \operatorname{im}(\partial_1)$ and that I is surjective. Then we can apply the Fundamental Theorem on Homomorphisms, which says that $\mathbb{Z} \approx C_0/\ker(I)$. Then if $\ker(I) = \operatorname{im}(\partial_1)$, we have

$$\mathbb{Z} \approx C_0/\ker(\partial_1) \approx H_0(X),$$

as desired. Now let us show that, indeed, $\operatorname{im}(\partial_1) = \ker(I)$.
($\subset$) Let $\sum a_i\sigma_i \in C_1$ be arbitrary. Then

$$\partial_1(\sum a_i\sigma_i) = \sum a_i\sigma_i \mid_{[v_0]} - \sum a_i\sigma_i \mid_{[v_1]} .$$

Thus

$$I(\partial_1(\sum a_i\sigma_i)) = \sum a_i - \sum a_i = 0,$$

so $\sum a_i\sigma_i \in \ker(I)$ and therefore $\operatorname{im}(\partial_1) \subseteq \ker(I)$.
($\supset$) Let $\alpha = \sum a_i\sigma_i \in \ker(I) \subset C_0(X)$ be arbitrarily chosen. Then $\sum a_i = 0$, by the definition of I. We will now construct an element v such that $\partial_i(v) = \alpha$, showing that $\alpha \in \operatorname{im}(\partial_1)$, and hence $\ker(I) \subset \operatorname{im}(\partial_1)$.

To begin, let us fix a point $x_0 \in X$, which exists because X is nonempty. Since X is path connected, we can find a path v_1 from x_0 to σ_1, where σ_1 is the first element of $\alpha = \sum a_i\sigma_i$ (each σ_i is just a point). We can similarly do this for each other point σ_i in α. Let $v \in C_1$ be the sum $v = \sum_i a_i v_i$, where each a_i is taken from α. Essentially, v is the collection of paths from x_0 to the points σ_i, with exactly one path for each instance of a point.

Consider that

$$\partial_1(v) = \partial_1(\sum_i a_i v_i) = \sum_i a_i x_0 - \sum_i a_i\sigma_i.$$

Since $\sum_i a_i = 0$, we have that

$$\partial_1(v) = \sum_i a_i\sigma_i = \sigma.$$

Hence $\alpha \in \operatorname{im}(\partial_1)$. □

Here we are considering $H_0(X)$ for a connected space. In $C_0(X)$, points are considered as distinct. However in $H_0(X)$, two points are equivalent if their difference is a boundary. This boundary can be constructed, because the space is connected, as the path between two points. Hence, under homology, all points are the same; so every arrangement of n points is the same under homology, and two arrangement only differ if they include different numbers of points.

Proposition 5.2. *If X can be decomposed into path-connected components X_α, then there is an isomorphism between $H_n(X)$ and the direct sum of its path connected components $\bigoplus_\alpha X_\alpha$.*

Proof. For each map σ, the image $\sigma(\Delta^n)$ is connected and hence is a subset of some X_α. Therefore $C_n(X)$ also splits into a direct sum of the subgroups $C_n(X_\alpha)$. In addition, the boundary maps preserve this split, since ∂_n takes an element from $C_n(X_\alpha)$ to $C_{n-1}(X_\alpha)$. Hence $\ker(\partial_n)$ and $\mathrm{im}(\partial_n)$ also split as direct sums, so

$$H_n(X) \approx \bigoplus_\alpha X_\alpha.$$

□

Proposition 5.3. *If X is a single point, then $H_n(X) = 0$ for $n > 0$, and $H_0(0) \approx \mathbb{Z}$.*

Apropros to the above statement, we generally want $H_n(X)$ to measure the n-dimensional holes, of which X in the above case X clearly has none. That $H_0(X) \approx \mathbb{Z}$ instead of 0 is a quirk that throws of the otherwise neat description of $H_n(X)$ above; so for the sake of convenience, we introduce another definition.

Definition 5.4. We define $\widetilde{H}_n(X)$, called the *reduced homology of* X, to be the homology of the augmented chain complex

$$\ldots \longrightarrow C_{n+1}(X) \xrightarrow{\partial_{n+1}} C_n(X) \xrightarrow{\partial_n} C_{n-1}(X) \longrightarrow \ldots \xrightarrow{\partial_1} C_0 \xrightarrow{I} \mathbb{Z} \to 0,$$

where I is the index map introduced in Proposition 5.1.

The result of this is that $H_0(X)$, which would previously be a copy of $\mathbb{Z}$ for each connected component, now has one less copy of $\mathbb{Z}$. On a side note, here we assume that X is non-empty, since the index map I is otherwise ill-defined.

Proposition 5.5. *Given a continuous map $f : X \to Y$, then f induces a homomorphism $f_\# : C_n(X) \to C_n(Y)$ and a homomorphism $f_* : H_n(X) \to H_n(Y)$ for all n.*

Theorem 5.6. *Let $f, g : X \to Y$ be two continuous maps. If f and g are homotopic, then the induced maps f_*, g_* induce the same homomorphism $f_* = g_* : H_n(X) \to H_n(Y)$ for all n.*

Proof. To show that $f_* = g_*$, iit sufficed to show that given an $n-$cycle $\sigma \in C_n(X)$, the chains $f_\#(\sigma)$ and $g_\#(\sigma)$ represent the same homology class in $H_n(Y)$. More specifically, we want to show that

$$f_\#(\sigma) - g_\#(\sigma) \in \mathrm{im}(\partial_{n+1}).$$

So, there must exist some $(n+1)$-chain $c \in C_{n+1}(Y)$ such that:

$$f_\#(\sigma) - g_\#(\sigma) = \partial c.$$

Why is this sufficient? Because we are concerned with equivalence classes. The map $f_*(z)$ of a cycle z is the equivalence class $[f_\#(z)]$, $f_\#(z)$ being a cycle.

The relationship between the homology groups $H_n(X)$ and $H_n(Y)$ are determined by the homotopy

$$H : X \times [0,1] \to Y,$$

which includes all the information of f and g and moreover how to continuously transition between the two. So, given this homotopy, we can construct an element $c \in C_{n+1}(Y)$ such that

$$\partial_{n+1}(c) = f_\#(\sigma) - g_\#(\sigma).$$

First, we have our homotopy $H : X \times I \to Y$ and a singular simplex $\sigma : \Delta^n \to X$. So, we can use composition to combine our map σ with our homotopy H:

$$H \circ (\sigma \times [0, 1]) : \Delta^n \times [0, 1] \to X \times [0, 1] \to Y.$$

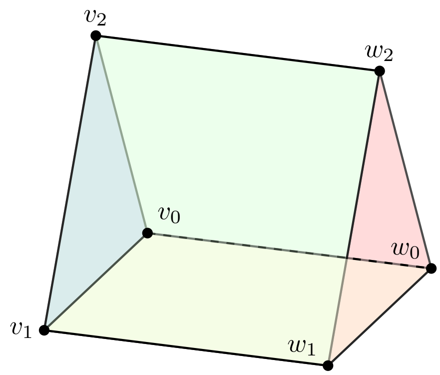

In order for this to be an element of $C_{n+1}(Y)$, we have to divide the prism $\sigma \times [0, 1]$ into simplices. The prism can be represented as two copies of our original simplex, one for the left: $\Delta^n \times \{0\} = [v_0, v_1, ...v_n]$ and the other for the right: $\Delta^n \times \{1\} = [w_0, w_1, ...w_n]$.

There exists a standard decomposition of the prism: the first $n + 1$ simplex is the left face with the last vertex given by w_n, resulting in

$$[v_0, v_1, ..., v_n, w_n].$$

The next one is $[v_0, v_1, ..., v_{n-1}, w_{n-1}, w_n]$, and so on.

Now we can construct prism operator $P : C_n(X) \to C_{n+1}(Y)$, defined to be

$$P(\sigma) = \sum_{i=0}^{n+1} (-1)^i H \circ (\sigma \times \mathbf{1}_{[0,1]}) \mid_{[v_0,...,v_i,w_i,...,w_n]} .$$

Having done this, we now show that

$$\partial P = g_\# - f_\# - P\partial.$$

The term ∂P represents the boundary of our entire prism; this is broken up into the left $f_\#$, the right $g_\#$, and the other sides (front, back, and bottom sides) $P\partial$. However, since we have broken up the prism into simplices, the term ∂P includes the internal faces as well. But all of these internal faces cancel each other out, since the decomposition of the prism into simplices has all adjacent faces have opposite orientations.

Now, given a cycle z, we have

$$g_\#(z) - f_\#(z) = \partial P(z) + P\partial(z) = \partial P(z),$$

since $\partial z = 0$ and thus $P\partial(z) = 0$. Hence $g_\#(z) - f_\#(z)$ is a boundary, so $g_\#(z) \cong f_\#(z)$ in $H_n(Y)$. Therefore $f_* = g_*$, as desired. □

Corollary 5.7. *If $f : X \to Y$ is a homotopy equivalence, then $f_* : H_n(X) \to H_n(Y)$ is an isomorphism for all n, or equivalently $H_n(X) \approx H_n(Y)$ for all n.*

Proof. The definition of homotopy equivalence is that

$$g \circ f \cong \mathbf{id}_X \text{ and } f \circ g \cong \mathbf{id}_Y.$$

By applying the previous theorem, we get that

$$(g \circ f)_* = (\mathbf{id}_X)_* = \mathbf{id}_{H_n(X)}.$$

Since $(g \circ f)_* = g_* \circ f_*$, we see that

$$g_* \circ f_* = \mathbf{id}_{H_n(X)}.$$

By similar logic $f_* \circ g_* = \mathbf{id}_{H_n(Y)}$, so f_* and g_* are inverses of each other, and must both be isomorphisms. Naturally this is also the case for reduced homology. □

These are most of the basic properties of singular homology. Other essential properties (of much greater complexity) are Mayer-Vietoris, Excision, and the equivalence between singular, simplical, and cellular homology.

Definition 5.8. Let A be a closed subset of a topological space X. Then (X, A) is a *good pair* if there is an open set $B \supset A$ such that A is a deformation retract of B.

Example 5.9. An example of a good pair is (D^n, S^{n-1}). That S^{n-1} is closed will not be shown. We only need to find an open set in D^n which deformation retracts onto S^{n-1}. Let

$$B = \{\mathbf{x} \in D^n \mid 0 < \|\mathbf{x}\| < 1\},$$

which is an open set. We can find a deformation retract: $f_t(x) = t\frac{\mathbf{x}}{\|\mathbf{x}\|} + (1-t)\mathbf{x}$, and hence (D^n, S^{n-1}) is a good pair.

Theorem 5.10. *(Mayer-Vietoris) Let (X, A) be a good pair. Then there is an exact sequence*

$$\cdots \to \widetilde{H}_{n+1}(X/A) \xrightarrow{\partial} \widetilde{H}_n(A) \xrightarrow{i_*} \widetilde{H}_n(X) \xrightarrow{q_*} \widetilde{H}_n(X/A) \xrightarrow{\partial} \widetilde{H}_{n-1}(A) \to \cdots$$

where:

- $i_* : \widetilde{H}_n(A) \to \widetilde{H}_n(X)$ *is induced by the inclusion* $i : A \hookrightarrow X$
- $q_* : \widetilde{H}_n(X) \to \widetilde{H}_n(X/A)$ *is induced by the quotient map* $q : X \to X/A$
- $\partial : \widetilde{H}_n(X/A) \to \widetilde{H}_{n-1}(A)$ *is the boundary homomorphism, not to be confused with* $\partial_n : C_n(X) \to C_{n-1}(X)$.

Although I am unable to go into detail here, the exact derivation of the ∂ map above is essential to a solid understanding of the above theorem.

6. The Homology of Spheres

The homology of spheres is very neat, as one would hope; after all, spheres are the most basic examples of holes in $\mathbb{R}^n$. The disk D^n is a shape without any holes; by removing a single point we end up with $D^n \setminus \{\mathbf{0}\}$, which is homotopic to S^{n-1}, as in Example 5.9. The hole inside of the sphere S^{n-1} is the "hole of largest dimension" that can exist in a given dimension of $\mathbb{R}^n$.

Theorem 6.1. *For the n-sphere S^n (with $n \geq 0$), the reduced singular homology groups are:*

$$\widetilde{H}_k(S^n) = \begin{cases} \mathbb{Z} & \text{if } k = n \\ 0 & \text{otherwise} \end{cases}$$

Proof. We use the Mayer-Vietoris sequence from Theorem 5.10. Let

$$(X, A) = (D^n, S^{n-1})$$

be the good pair from Example 5.9. Here S^{n-1} is the boundary of D^n, and thus D^n/S^{n-1} becomes S^n, the sphere in $\mathbb{R}^{n+1}$.

The disk D^n is retractable, so $\widetilde{H}_k(D^n) = 0$ for all $k \geq 0$; this follows from Corollary 5.7. We retract the disk to a single point; and we know the homology of a point from Proposition 5.3.

Hence the long exact sequence

$$\cdots \to \widetilde{H}_{k+1}(X/A) \xrightarrow{\partial} \widetilde{H}_k(A) \xrightarrow{i} \widetilde{H}_k(X) \xrightarrow{q} \widetilde{H}_k(X/A) \xrightarrow{\partial} \widetilde{H}_{k-1}(A) \to \cdots$$

becomes, in our case,

$$\cdots \to \widetilde{H}_{k+1}(S^n) \xrightarrow{\partial} \widetilde{H}_k(S^{n-1}) \xrightarrow{i} \widetilde{H}_k(D^n) \xrightarrow{q} \widetilde{H}_k(S^n) \xrightarrow{\partial} \widetilde{H}_{k-1}(S^{n-1}) \to \cdots$$

which can further be simplified into

$$\cdots \to \widetilde{H}_{k+1}(S^n) \xrightarrow{\partial} \widetilde{H}_k(S^{n-1}) \to 0 \to \widetilde{H}_k(S^n) \xrightarrow{\partial} \widetilde{H}_{k-1}(S^{n-1}) \to 0 \to \cdots$$

Since this sequence is exact, we only need to consider each section

$$0 \to \widetilde{H}_k(S^n) \xrightarrow{\partial} \widetilde{H}_{k-1}(S^{n-1}) \to 0$$

to see that $\xrightarrow{\partial}$ must be an isomorphism for all k (since it must be both injective and surjective). Hence $\widetilde{H}_k(S^n) \approx \widetilde{H}_{k-1}(S^{n-1})$ for all $k \in \mathbb{N}$.

We use Proposition 5.3, which says that $H_0(X) \approx \mathbb{Z}$ when X is a single point. The zero-sphere S^0 is just two points; so by Proposition 5.2, $H_0(S^0)$ is the direct sum of its connected components, $\mathbb{Z} \oplus \mathbb{Z}$. The reduced homology is essentially removing one of the $\mathbb{Z}$'s, and therefore $\widetilde{H}_0(S^0) \approx \mathbb{Z}$.

Since $\widetilde{H}_0(S^0) \approx \mathbb{Z}$, we know, by our above step, that $\widetilde{H}_k(S^k) \approx \mathbb{Z}$ for all k.

Let $k > 0$. Again using Proposition 5.2, we see that $H_k(S^0) \approx H_k(\{x\}) \oplus H_k(\{y\})$, which by Proposition 5.3 becomes $H_k(S^0) \approx 0 \oplus 0 \approx 0$. Since $k > 0$, homology and reduced homology do not differ; so for all $k > n$, we have that

$$\widetilde{H}_k(S^n) \approx \widetilde{H}_{k-1}(S^{n-1}) \approx ... \approx \widetilde{H}_{k-n}(S^0) \approx 0.$$

Let $k < n$. By repeatedly applying our above result, we get that

$$\widetilde{H}_k(S^n) \approx \widetilde{H}_{k-1}(S^{n-1}) \approx ... \approx \widetilde{H}_0(S^{n-k}).$$

Since S^{n-k} is a sphere of dimension greater than zero, we know that it is connected, and hence by Proposition 5.1, we have that $\widetilde{H}_0(S^{n-k}) \approx 0$. □

7. An Introduction to Degree

With degree we are generally considering continuous maps between (orientable) manifolds of the same dimension. The object that has degree is the map itself. Simply put, degree measures the number of times that the map "wraps" the domain around its codomain. The degree will always be an integer, positive or negative depending on orientation.

We see in Theorem 6.1 that the homology group of a sphere will always be a free group isomorphic to $\mathbb{Z}$ or trivial. Any map $f : S^n \to S^n$ will induce a map on homology; so the induced map f_* is a homomorphism from $\mathbb{Z}$ to $\mathbb{Z}$.

It is a special property of homomorphisms from $\mathbb{Z} \to \mathbb{Z}$ that we only need to know what 1, the identity element, is mapped onto through f_*, since

$$f_*(n) = f_*(\sum_{i=1}^{n} 1) = \sum_{i=1}^{n} f_*(1) = nf_*(1).$$

Definition 7.1. Given a map $f : S^n \to S^n$, the *degree* of f, derived from the induced map on homology $f_* : \mathbb{Z} \to \mathbb{Z}$, is the integer $f_*(1)$.

Here are a few very nice properties of degree, which can be further explored in Hatcher [1], on page 134. If we let $f, g : S^n \to S^n$,

(1) $deg(\mathbf{id}) = 1$.
(2) If $f \cong g$ then $deg(f) = deg(g)$.
(3) $deg(f \circ g) = deg(f) \cdot deg(g)$.
(4) $deg(f) = -1$ if f is a reflection, i.e. a reflection across a plane of dimension $n-1$ that passes thought he origin.

This last property is especially important in our following proof.

8. Proof of the Hairy Ball Theorem

The statement of the hairy ball theorem is about vector fields on spheres S^n. I have not gone into great depth in my studies of vector fields on manifolds; undoubtedly it is a beautiful subject. But even without this background, the statement of the Hairy Ball theorem is still approachable.

Theorem 8.1. *Every continuous vector field on S^n will vanish somewhere if n is even.*

When n is odd, we can find such a continuous, non-vanishing field. One such field is shown here for the case of S^1.

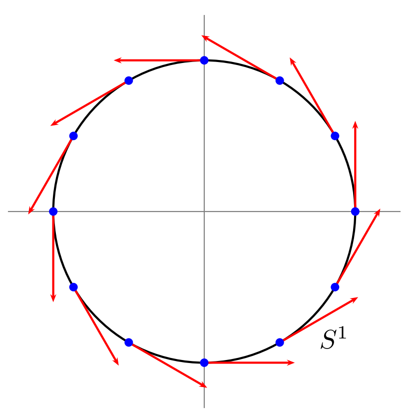

Figure 3. Illustration of continuous non-vanishing vector field

Similar vector fields can be found for larger odd values of n. While we will not prove the Hairy Ball theorem in full generality, we are able to prove it the case of $n = 2$ with our current tools.

Proof. We assume for the sake of contradiction that there exists a continuous, non-vanishing vector field F on the sphere S^2. We use F to construct a homotopy between

$$\mathbf{id}_{S^2} : S^2 \to S^2 \text{ and } -\mathbf{id}_{S^2} : S^2 \to S^2,$$

the identity map and the antipodal map. The antipodal map maps each point onto the opposite side of the sphere, and is explicitly given by

$$-\mathbf{id}_{S^2}(x_1, x_2, ..., x_n) = (-x_1, -x_2, ..., -x_n).$$

To construct this homotopy, we begin with an arbitrary vector $\mathbf{v}$ that lies on the unit circle. The field F gives us a tangent vector $F(\mathbf{v})$. If we normalize this vector $F(\mathbf{v})$ and consider it as originating from the origin, we see that $\frac{F(\mathbf{v})}{\|F(\mathbf{v})\|}$ also represents a vector in S^2 — necessarily a different one from $\mathbf{v}$ itself. We draw a great arc beginning at $\mathbf{v}$, passing through the vector $\frac{F(\mathbf{v})}{\|F(\mathbf{v})\|}$, and ending at $-\mathbf{v}$.

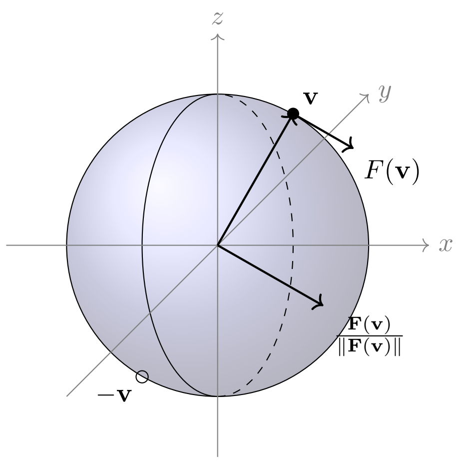

Thus our homotopy $H : S^n \times [0, 1] \to S^n$ is given by

$$H(\mathbf{v}, t) = (\cos \pi t)\mathbf{v} + (\sin \pi t)\frac{F(\mathbf{v})}{\|F(\mathbf{v})\|}.$$

Clearly sine and cosine are continuous; and by assumption, $F(\mathbf{v})$ is also continuous, so H is continuous and thus gives a homotopy between

$$H(\mathbf{v}, 0) = \mathbf{v} = \mathbf{id}_{S^n}(\mathbf{v}) \text{ and } H(\mathbf{v}, 1) = (\cos \pi)\mathbf{v} = -\mathbf{v} = -\mathbf{id}_{S^n}(\mathbf{v}).$$

Since degree respects homotopy, we can say that

$$deg(-\mathbf{id}_{S^n}) = deg(\mathbf{id}_{S^n}) = 1.$$

To arrive at a contradiction, we must use the fact that n is even. If we go back to the definition of the antipodal map,

$$-\mathbf{id}_{S^n}(x_0, x_1, x_2, ..., x_n) = (-x_0, -x_1, -x_2, ..., -x_n),$$

we can see that the antipodal map is the composition of $n + 1$ reflections maps r_i that each add a single negative sign to one component:

$$r_i(x_0, x_1, x_2, ..., x_n) = (x_0, ..., -x_i, ..., x_n).$$

If n is even, then the degree of the antipodal map is $(-1)^{n+1} = -1$, resulting in a contradiction. $\square$

Acknowledgments

I would like to thank my mentor Paweł Poczobut for his patience and thorough explanations. I would like to thank Peter May for running this program. Thank you to all the professors, grad students, and visiting speakers who gave wonderful lectures during this program. Thank you to Daniil Rudenko and other graduate students for the fantastic morning apprentice Lectures. Lastly I thank my fellow students in the apprentice and full programs who accompany me on the pursuit of mathematical knowledge, who have patience and are excited to explain and share. You are the reason that I am able to appreciate the beauty of mathematics.

References

[1] Allen Hatcher. Algebraic topology. Cambridge University Press, Cambridge, 2002. https://pi.math.cornell.edu/ hatcher/AT/AT.pdf

[2] Chad Berkich. homology, COhomology, AND THE DE RHAM THEOREM. August 28, 2021. https://math.uchicago.edu/ may/REU2021/REUPapers/Berkich.pdf

[3] John M. Lee. Introduction to Smooth Manifolds. Springer, 2012.

[4] Anthony Bosman. Recorded Lectures on Algebraic Topology. Sep 1, 2023. https://youtu.be/kCTpfqRJ2kk?si=s5V-30rOW7GXiXeh