A characteristic feeling: exploring characteristic classes - Part 1, cohomology

rapha • 13 Feb 2025

I'm planning on writing a series of posts that explore the theory and applications of characteristic classes in algebraic topology, following the book from Milnor and Stasheff. Here I start with the appendix A, where homology is discussed and basic theorems are laid out. This post will also serve me well as a quick reminder, since I keep forgetting small details and ideas regarding (singular) (co)homology. Many results and explanations can be found in Hatcher's book.

Singular homology

The **standard n-simplex** is the set $\Delta^n \subseteq \mathbb{R}^{n+1}$ consisting of all $(n+1)$-tuples $(t_0, t_1, \ldots, t_n)$ with the following two properties:

- for each i, we have $t_i \geq 0$;
- and $\sum t_i = 1$.

The second property says that points in Δ^n are such that the dot product of the vector $(t_0 - 1, t_1, \ldots, t_n)$ with the vector $n = (1, 1, \ldots, 1)$ is zero, i.e. Δ^n lies in the n-hyperplane with normal n, translated by one unit in any direction. Hence each Δ^n is an affine space of dimension n. For instance:

- Δ^0 is a single point $1 \in \mathbb{R}$;
- Δ^1 is the line segment in $\mathbb{R}^2$ going up from $(1, 0)$ to $(0, 1)$;
- Δ^2 is the (filled) triangle in $\mathbb{R}^3$ having as vertices the three standard basis vectors;
- etc.

As a special case, we also define Δ^{-1} to be the empty set.

In general each Δ^n has $(n+1)$ **vertices**, which are the points in Δ^n corresponding to the standard basis vectors. Another way to talk about the standard n-simplex is to say it is the *convex hull* of the standard basis vectors in $\mathbb{R}^{n+1}$. We can label each vertex with an integer i, where i is the (zero-based) position of the unique 1 in the standard basis vector corresponding to that vertex.

For each $0 \leq i \leq n$, we have a way to talk about the i**-th side** of a standard n-simplex via the function $\phi_i : \Delta^{n-1} \to \Delta^n$, which is defined as

$$\phi_i(t_0, \ldots, \widehat{t_i}, \ldots, t_n) = (t_0, \ldots, t_{i-1}, 0, t_{i+1}, \ldots, t_n).$$

(As usual the hat over a variable in an enumeration means that variable is actually omitted from the enumeration.) Thus the i-th side is the convex hull of all vertices of Δ^n that are not labelled with i. Notice that ϕ_i is actually an affine embedding. Moreover, it gives an orientation to each side of a standard simplex: the orientation is positive i is even, and negative otherwise.

Again, as a special case, we define the 0-th side of Δ^0 to be the unique function $\phi_0 : \Delta^{-1} \to \Delta^0$ from the empty set (i.e. Δ^{-1}) to the singleton $\{1\}$ (i.e. Δ^0).

Let X be any topological space. A **singular** n**-simplex** in X is a continuous map from Δ^n to X. The idea here is to identify such a map with its image; since there are no restrictions on what this image may look like except that nearby points stay nearby (continuity), there could be collapsing or other weird things happening to Δ^n when viewed through the map. That's why it's called a *singular* simplex.

For $0 \leq i \leq n$, the i**-th face** of a singular n-simplex $\sigma : \Delta^n \to X$ is the singular $(n-1)$-simplex given by

$$\sigma \circ \phi_i : \Delta^{n-1} \to X.$$

For each $n \geq 0$, the **singular chain group** $C_n(X; \Lambda)$ with coefficients in a commutative ring Λ is the free Λ-module having one generator for each singular n-simplex σ in X. In other words, $C_n(X; \Lambda)$ consists of the formal Λ-linear combinations of n-simplices in X:

$$C_n(X; \Lambda) = \bigoplus_{\sigma : \Delta^n \to X} \Lambda.$$

Notice that this is a *module*, even though we use the term *group*. Just another fun little opportunity to be confused down the line. When $n < 0$, the singular chain group is defined to be the zero module. In the special case $\Lambda = \mathbb{Z}$, the singular chain group at n is the free abelian group on the singular n-simplices.

The singular chain group is the algebraic realization of how "sticking triangles on a space" works. If $\Lambda = \mathbb{Z}/2\mathbb{Z}$, then glueing two copies of the same simplex one on top of another means they "cancel out", and when only one of their sides are overlapping, these overlapping sides also cancel out; you obtain a square in the space instead of two triangles. Many geometric arguments use the fact that

sides on the boundaries cancel out. The precise meaning of this in $C_n(X;\Lambda)$ is given by the **boundary homomorphism**, which is a Λ-linear map $\partial : C_n(X;\Lambda) \to C_{n-1}(X;\Lambda)$ defined as

$$\partial\sigma = \sum_{i=0}^{n}(-1)^i(\sigma \circ \phi_i).$$

This definition is made for $n \geq 1$; if $n \leq 0$ we simply say $\partial = 0$. The sign in the formula represents the orientation of each side (recall: the i-th side has positive orientation when i is even, and negative orientation otherwise). This is done because we need "identical sides" that are "going in opposite directions" to cancel out in many geometrical arguments. For instance, the boundary of a tiling of some region of space by triangles (i.e. a sum in the singular chain group C_2) should be the boundary of the region, *not* the sum of the individual boundaries of each triangle (for instance, think about the proof of Stoke's formula).

An important property of the boundary homomorphism is this: $\partial^2 = 0$. Intuitively: the boundary of the boundary is empty. Think of Δ^2, which is a (filled) triangle, and think about it as a singular simplex in $\mathbb{R}^2$ (maybe via the projection on the plane in which Δ^2 lies). Its boundary is the formal sum of three line segments. These segments are all perfectly lined up so that the end of the one is the start of the next: they form a cycle. Moreover, the point corresponding to the end of one segment has the opposite orientation to the point corresponding to the start of the next segment, so they cancel out. Since these endpoints are the boundaries of each of the three line segment, and since they all cancel out, the boundary of the boundary is effectively zero.

In fact, the boundary of a singular chain is zero precisely when all of the summands in the chain are arranged so that their boundaries all cancel out, that is, when they form a cycle and "enclose" some region of space. We define the n**-cycles** to be the set of all such chains:

$$Z_n(X;\Lambda) = \ker(\partial : C_n(X;\Lambda) \to C_{n-1}(X;\Lambda)).$$

The n**-boundaries** is the set of all chains that can be expressed as the boundary of some $(n+1)$-dimensional singular chain:

$$B_n(X;\Lambda) = \operatorname{im}(\partial : C_{n+1}(X;\Lambda) \to C_n(X;\Lambda)).$$

The identity $\partial^2 = 0$ says we have a containment of Λ-submodules

$$B_n(X;\Lambda) \subseteq Z_n(X;\Lambda).$$

Hence we can consider the n**-th singular homology group**

$$H_n(X;\Lambda) = Z_n(X;\Lambda)/B_n(X;\Lambda).$$

The homology group captures in algebra an intuitive spatial fact. We have seen that any cycle "encloses" a region of space, by sticking together simplices of the same dimension along their boundaries until none are "left alone" (each boundary has a matching, opposite, boundary). Now suppose a given cycle is itself the boundary of some higher-dimensional simplex Δ. Now the cycle can "move inside" Δ without breaking appart, shrinking until it becomes a single point. If a cycle is not a boundary, it means that something about the space X obstructs the construction of a simplex Δ which would have the cycle as its boundary: there's a hole in X. Notice that the condition of continuity on singular simplices is essential here: the hole would basically force any Δ to be torn appart if it were to have the cycle as a boundary, breaking continuity. In this way, the homology construction detects holes in X and gives useful information about them. This technology could be used to make a precise definition of what a "hole" in a topological space is: a hole is a generator for the homology group.

Some abstract nonsense

Homology is a functor from the category of topological spaces up to homotopy, to the category of Λ-modules. The source category's objects are topological spaces, and the arrows are equivalence classes of continuous maps, where two maps are considered to be the same when they are homotopic. Since I always forget the details of what this means, here they are: consider two continuous maps $f, g : X \to Y$ between topological spaces. We say they are **homotopic** when there exists another continuous map $H : X \times [0, 1] \to Y$ (called an **homotopy**) such that $H(-, 0) = f$ and $H(-, 1) = g$. In the source category we're interested in, an arrow is a actually a set of continous maps, all homotopic to each other.

Functoriality gives us the following for free: if two spaces are homotopic, then they have isomorphic homology groups. In particular, since $\mathbb{R}^n$ is contractible (i.e. homotopy equivalent to a point) for any $n \geq 0$, we have a concrete homology computation: for any $i \geq 1$,

$$H_i(\mathbb{R}^n; \Lambda) \cong H_i(\{p\}; \Lambda) = 0.$$

(It's easy to compute homology of a point: all singular simplices are the same!) When $i = 0$, things are a bit weird, and fixing weird things is the purpose of the next section.

Reduced homology

There's a technical point to address here. Consider what happens when we take the space X to be a single point p and we compute the 0-th homology group. Since singular 1-simplices are continuous maps $\Delta^1 \to \{p\}$, all 1-simplices are actually the same. Hence they are all cycles, so by definition their boundaries are zero: $B_0(\{p\};\Lambda) = 0$. Therefore $H_0(\{p\};\Lambda) \cong Z_0(\{p\};\Lambda)$. However, we set ∂ to be the zero map for all $n \leq 0$ earlier, so $H_0(\{p\};\Lambda)$ is $C_0(\{p\};\Lambda)$, the free module generated by all 0-simplices. Since they are all the same, there's actually only one generator, so $H_0(\{p\};\Lambda) \cong \Lambda$. For technical reasons, it's better for the zeroth homology of a point to be the zero module; also, it makes sense intuitively, since we expect the homology to measure holes in a space, and we feel a point doesn't have holes.

There's an easy fix to this. Instead of having $\partial = 0$ at degree zero, we set $C_{-1}(X;\Lambda)$ to be Λ (recall: we made a special case above, where the only singular (-1)-simplex is the unique function from the empty set Δ^{-1} to X; then the chain group of degree -1 has to be the free Λ-module generated by that single (-1)-simplex), and now we may define ∂ at degree zero with the same formula we used for positive degrees. For any singular 0-simplex $\sigma : \Delta^0 \to X$, we now have

$$\partial\sigma = \sigma \circ \phi_0,$$

which is the unique function from the empty set to X, and which is identified with $1 \in \Lambda$. Therefore ∂ corresponds to the identity map on Λ. This means $Z_0(\{p\};\Lambda)$ is trivial, and so is this modified homology at degree zero.

This modified homology is called **reduced singular homology** and its homology groups are denoted with a tilde:

$$\widetilde{H}_i(X;\Lambda).$$

Since the only modification happens at degree zero, we have

$$\widetilde{H}_i(X;\Lambda) = H_i(X;\Lambda)$$

for each $i \geq 1$. In general, we see that for any 0-chain $a_1\sigma_1 + \cdots + a_k\sigma_k$ (these are just formal linear combinations of points in X) we have

$$\partial_0(a_1\sigma_1 + \cdots + a_k\sigma_k) = \sum_{i=1}^{k} a_i \in \Lambda.$$

There is an easy way to get from unreduced homology to reduced homology: at all positive dimensions the groups are the same, and at dimension zero we have the equation

$$H_0(X;\Lambda) \cong \widetilde{H}_0(X;\Lambda) \oplus \Lambda.$$

(Hint: think about what happens if X is path-connected, and which 0-chains are boundaries.)

Mayer-Vietoris sequence

A great tool for computing with homology. It works for "unreduced" and reduced homology (just replace H with $\widetilde{H}$ everywhere). Let A and B be two subsets of X such that their interior cover X (and for reduced homology, we also want their intersection to be nonempty). Then there is a long exact sequence in homology

$$\ldots \longrightarrow H_{i+1}(X) \longrightarrow H_i(A \cap B) \longrightarrow H_i(A) \oplus H_i(B) \longrightarrow H_i(X) \longrightarrow \ldots$$

We can use this to compute the homology of the n-sphere S^n. Let A be the "open north cap" and B be the "open south cap", i.e. A and B are contractible open sets in S^n such that their intersection is homotopic to the "equator" S^{n-1}. Then the reduced Mayer-Vietoris sequence looks like

$$\ldots \longrightarrow \widetilde{H}_{i+1}(S^n) \longrightarrow \widetilde{H}_i(S^{n-1}) \longrightarrow \widetilde{H}_i(A) \oplus \widetilde{H}_i(B) \longrightarrow \widetilde{H}_i(S^n) \longrightarrow \ldots$$

Because A and B are contractible, the middle term above is zero for all $i \in \mathbb{Z}$, so we have a collection of isomorphisms $\widetilde{H}_i(S^n) \cong \widetilde{H}_{i-1}(S^{i-1})$. Since $\widetilde{H}_0(S^0) = \Lambda$ and zero otherwise, we find by induction on n the following calculation:

$$\widetilde{H}_i(S^n;\Lambda) = \begin{cases} \Lambda & \text{if } i = n \\ 0 & \text{otherwise.} \end{cases}$$

Relative homology

We now consider pairs (X, A) where A is any subspace of X (including the empty subspace, and the full subspace). We're going to look at homology "modulo A", in the sense that any singular simplex whose image lies completely in A is going to be considered as "completely collapsed", i.e. zero as a chain. Formally, we define the **relative n-th singular chain group** to be

$$C_n(X, A;\Lambda) = C_n(X;\Lambda)/C_n(A;\Lambda).$$

Because ∂ carries chains in A to chains in A, we obtain a chain complex "modulo A" and we can define relative homology as

$$H_n(X, A; \Lambda) = Z_n(X, A; \Lambda)/B_n(X, A; \Lambda).$$

Any pair (X, A) gives an exact sequence of Λ-modules

$$0 \longrightarrow C_n(A; \Lambda) \longrightarrow C_n(X; \Lambda) \longrightarrow C_n(X, A; \Lambda) \longrightarrow 0.$$

From the general theory of abelian categories, we obtain from it a **long exact sequence in homology**:

$$\begin{aligned} \ldots \longrightarrow H_{n-1}(A; \Lambda) \longrightarrow H_{n-1}(X; \Lambda) \longrightarrow H_{n-1}(X, A; \Lambda) \\ \hookrightarrow H_n(A; \Lambda) \longrightarrow H_n(X; \Lambda) \longrightarrow H_n(X, A; \Lambda) \\ \hookrightarrow H_{n+1}(A; \Lambda) \longrightarrow H_{n+1}(X; \Lambda) \longrightarrow H_{n+1}(X, A; \Lambda) \longrightarrow \ldots \end{aligned}$$

This long exact sequence also exists for reduced homology.

An important tool for working with homology is the **excision theorem**: let A and B be two subspaces of X such that their interior cover X; then the inclusion $(B, A \cap B) \hookrightarrow (X, A)$ induces isomorphisms

$$H_n(B, A \cap B) \cong H_n(X, A)$$

for all $n \in \mathbb{Z}$. Equivalently, for any subspaces $Z \subseteq A \subseteq X$ such that the closure of Z is contained in the interior of A, the obvious inclusion induces isomorphisms

$$H_n(X - Z, A - Z; \Lambda) \cong H_n(X, A; \Lambda).$$

This version of the statement is what justifies the name “excision”, since it gives us conditions under which we may excise Z from X without changing the homology groups. That is insanely powerful. For instance, here’s a proof of the so-called Brouwer’s invariance of domain: if $U \subseteq \mathbb{R}^m$ and $V \subseteq \mathbb{R}^n$ are two nonempty homeomorphic open sets, then $m = n$. The idea of the proof is to look at what happens locally around a point, so we define **local homology groups** around some point $p \in X$ by

$$H_{n,p}(X; \Lambda) = H_n(X, X - \{p\}; \Lambda).$$

Back to the proof. Let $f : U \to V$ be a homeomorphism and pick some point $x \in U$. Then, by excision, we have $H_{n,x}(U) \cong H_{n,x}(\mathbb{R}^m)$. (Hint: in the excision theorem, pick A to be the complement of $\{x\}$ and pick B to be U). The long exact sequence for the pair $(\mathbb{R}^m, \mathbb{R}^m - \{x\})$ looks like:

$$\ldots \longrightarrow H_{i-1,x}(\mathbb{R}^m) \longrightarrow H_i(\mathbb{R}^m - \{x\}) \longrightarrow H_i(\mathbb{R}^m) \longrightarrow H_{i,x}(\mathbb{R}^m) \longrightarrow \ldots$$

We saw that $H_i(\mathbb{R}^m) \cong 0$ for every $i \in \mathbb{Z}$. Hence we have a collection of isomorphisms

$$H_{i,x}(\mathbb{R}^m) \cong H_{i+1}(\mathbb{R}^m - \{x\}).$$

Moreover, because $\mathbb{R}^m - \{x\}$ is homotopic to S^{m-1}, the homology $H_{i+1}(\mathbb{R}^m - \{x\})$ is Λ for $i+1 = m$ and zero otherwise. This gives the following calculation: the homology $H_{i,x}(\mathbb{R}^m)$ is zero if and only if $i = m$. The same reasoning applied to V and $h(x)$ gives the same calculation, and since h induces an isomorphism of homology groups, we must have $m = n$.

Singular cohomology

Since all elements in a Λ-module M may be identified with the Λ-linear maps $\Lambda \to M$, we may dualize and consider linear maps of the same kind but having opposite polarity: this we do. The **n-th cochain group** is the dual module

$$C^n(X;\Lambda) = \mathrm{Hom}_\Lambda(C_n(X;\Lambda),\Lambda),$$

consisting of all Λ-linear maps going from the singular chain group into its ring of scalars. A cochain is just a way to (linearly) compute a scalar quantity from a singular chain. There's an analogy to be made with geometry: you have points in some affine space (singular chains), and you have coordinates (cochains), which in a way compute a number from each point. In geometry, there's a strong link between points and coordinates: studying algebraic varieties is essentially the same as studying rings of coordinates, which are basically rings of polynomial functions from the space to the underlying field. However, it is easier to work with coordinates than sets of points, because there's a natural ring structure. Hence, if this analogy is to hold, one would expect a link between homology and cohomology; moreover, it should be easier to work with cohomology than homology. And so it is.

The **value** of a cochain c on a chain γ will be denoted $\langle c, \gamma\rangle$ and is defined as

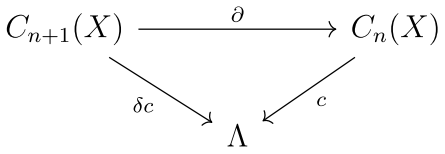
$$\langle c, \gamma\rangle = c(\gamma) \in \Lambda.$$

Obviously $\langle -, -\rangle : C^n(X;\Lambda) \times C_n(X;\Lambda) \to \Lambda$ is Λ-bilinear.

The **coboundary** of a cochain $c \in C^n(X;\Lambda)$ is defined to be the cochain $\delta c \in C^{n+1}(X;\Lambda)$ whose value on each $(n+1)$-chain α is determined by the identity

$$\langle \delta c, \alpha\rangle + (-1)^n \langle c, \partial\alpha\rangle = 0.$$

Hence δ is, up to sign, the dual of ∂, in the sense that for any cochain c, the cochain δc is, up to sign, the pullback of c along ∂:



This sign convention is used in Milnor and Stasheff's book, but not in Hatcher for instance, where he defines δ as precisely the dual of ∂. Since my goal is to understand characteristic classes, I'm going to keep the sign convention used in the M&S book.

Again, this definition has intuitive content: since c is able to "measure", or "give coordinates", any n-chain, then it should be possible to obtain a way to measure $(n+1)$-chains α by combining measures for the boundaries of α.

The coboundary homomorphism, just like its dual friend, verifies $\delta^2 = 0$. Therefore, if we define n**-cocycles** to be

$$Z^n(X;\Lambda) = \ker(\delta : C^n(X;\Lambda) \to C^{n+1}(X;\Lambda))$$

and n**-coboundaries** to be

$$B^n(X;\Lambda) = \mathrm{im}(\delta : C^{n-1}(X;\Lambda) \to C^n(X;\Lambda),$$

then we may also define the n**-th singular cohomology group** by

$$H^n(X;\Lambda) = Z^n(X;\Lambda)/B^n(X;\Lambda).$$

Universal coefficient theorem for cohomology

Instead of using the ring Λ as coefficients, we may also use any Λ-module M. If Λ is a principal ideal domain, then there is a natural split exact sequence

$$0 \longrightarrow \mathrm{Ext}^1_\Lambda(H_{n-1}(X;\Lambda), M) \longrightarrow H^n(X;\Lambda) \xrightarrow{h} \mathrm{Hom}_\Lambda(H_i(X;\Lambda), M) \longrightarrow 0.$$

The map h is the canonical map sending a cohomology class represented by a cochain c, to the map which sends any homology class represented by a chain α to the element $\langle c, \alpha \rangle$ of Λ.

This exact sequence measures how close the cohomology group is to be the dual of homology.