Functoriality, Singular Homology, Relative Homology

18 Nov, 2010

1 Overview

Last week, we defined the simplicial homology groups of a triangulated space and gave a matrix reduction algorithm for their computation. This week, we will:

- Discuss *functoriality*: how continuous maps between topological spaces induce algebraic maps between their respective homology groups.
- Give a slightly different definition of homology, called *singular homology*, which is not so intimately tied to a triangulation.
- Define and compute the *relative homology* groups for pairs of spaces.

2 Functoriality

Now we come to probably the most important property of homology: the fact that maps between spaces induce maps between homology groups. We'll demonstrate this first for simplicial maps between simplicial complexes. The whole discussion becomes a lot cleaner after we define singular homology groups.

Induced maps. Suppose we have a simplicial map $f : K \to L$ between two simplicial complexes; recall that this means f maps each simplex σ in K linearly onto some simplex, of the same or lower dimension, $f(\sigma)$ of L. For each integer p, we thus see that f maps p-chains in K to p-chains in L. Formally, given a p-chain $c = \Sigma_i a_i \sigma_i \in \mathsf{C}_p(K)$, we define $f_\#(\sigma) = \Sigma_i a_i \tau_i \in \mathsf{C}_p(L)$, where $t_i = f(\sigma_i)$ if $f(\sigma_i)$ also has dimension p, and $t_i = 0$ if $f(\sigma_i)$ has dimension less than p.

Note that this defines homomorphisms $f_\# : \mathsf{C}_p(K) \to \mathsf{C}_p(L)$ in each dimension p. We also have, for each p, the two boundary maps $\partial_p^K : \mathsf{C}_p(K) \to \mathsf{C}_{p-1}(K)$ and $\partial_p^L : \mathsf{C}_p(L) \to \mathsf{C}_{p-1}(L)$. And it is easy to see that $f_\#$ commutes with these two boundary maps: $\partial_p^L \circ f_\# = f_\# \circ \partial_p^K$. In other words, $f_\#$ has the defining characteristic of a *chain map* between two chain complexes.

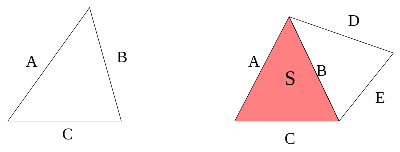

Figure 1: K is on the left, L is on the right, and the map is indicated by the labelling.

Hence $f_\#$ maps cycles to cycles and boundaries to boundaries: $f_\#(\mathsf{Z}_p(K)) \subseteq \mathsf{Z}_p(L)$ and $f_\#(\mathsf{B}_p(K)) \subseteq \mathsf{B}_p(L)$. In other words, $f_\#$ induces a map on homology groups, which we denote by $f_* : \mathsf{H}_p(K) \to \mathsf{H}_p(L)$; note that although we have dropped all mention of p for notational clarity, there is really one map f_* for each dimension p.

Example. Probably the easiest situation to understood is a homology map induced by an injective map, as illustrated in Figure 1. On the left we have a simplicial complex K and on the right we have L. We define $f : K \to L$ in such a way the labelled edges in Figure 1 correspond. Note that in this case, our map f is merely a way to mentally include K as a subcomplex of L, and the chain map $f_\#$ performs a similar service for chains. The induced homology map f_* is more interesting, and we'll discuss it in each dimension.

- $f_* : \mathsf{H}_0(K) \to \mathsf{H}_0(L)$ is an isomorphism: each complex has only one component, and $f(K)$ contains a representative of this component.
- Notice that $\mathsf{H}_1(K)$ and $\mathsf{H}_1(L)$ are both rank one, but that $f_* : \mathsf{H}_1(K) \to \mathsf{H}_1(L)$ is not an isomorphism, because it has a non-zero kernel: namely the homology class in $\mathsf{H}_1(K)$ represented by the 1-chain $\gamma = A + B + C$. This class does not bound in K, but it does bound in L, since $\gamma = \partial_2^L(S)$, and hence $f_*(\gamma) = 0$. Note also that the 1-dimensional homology class in L, represented by $C+D+E$, is not in the image of f_*.
- There are no 2-dimensional classes in either complex, so $f_* : \mathsf{H}_2(K) \to \mathsf{H}_2(L)$ is an isomorphism for trivial reasons.

As we'll see in a few weeks, homology maps induced by inclusions will be the main theoretical tool needed for the definition of persistent homology.

Functorial properties. The fact that a continuous map between topological spaces (or a simplicial map between simplicial complexes, but it's really the same thing as we've seen!) induces an homomorphism between homology groups is incredibly important: basically, it means that we can associate algebraic data to each space, and then attempt to transfer this information from space to space in order to make comparisons. The following two properties show that this information transfer takes place in a consistent manner:

- Let K be a simplicial complex and $i : K \to K$ denote the identity map given by $i(v) = v$ for all vertices in K. Then the induced map i_* is also the identity homomorphism.
- Let $f : K \to L$ and $g : L \to M$ be two maps between simplicial complexes, and consider the composition $(g \circ f) : K \to M$. Then $(g \circ f)_* = g_* \circ f_*$; in other words, the induced map of the composition is the composition of the induced maps.

As an easy consequence of these two properties, we have the following essential fact: if $f : K \to L$ is a homeomorphism, then $f_* : \mathsf{H}_p(K) \to \mathsf{H}_p(L)$ is an isomorphism for every p.

proof:

This is most useful in the contrapositive. For example, we can now show that the sphere S^2 is not homeomorphic to the torus T^2, since there can never be an isomorphism between $\mathsf{H}_1(S^2) = 0$ and $\mathsf{H}_1(T^2) = \mathbb{Z}/2\mathbb{Z} \oplus \mathbb{Z}/2\mathbb{Z}$.

In fact, something much more powerful is true: if $f, h : K \to L$ are two homotopic maps, then the induced maps are equal: $f_* = h_*$. From this, we see that a homotopy equivalence will also induce a homology isomorphism. Hence S^2 and T^2 are not even homotopically equivalent!

3 Singular Homology

So far, we have defined homology in terms of various sets of simplices which arise from a given triangulation of a topological space. Although this is necessary for algorithmic purposes, it is also theoretically cumbersome; we do not actually want to mentally triangulate a space every time we wish to think about its homology groups! Fortunately, there's a way around this problem, although we are forced to deal with yet one more level of abstraction.

Let $\mathbb{X}$ be a topological space, and for each integer p, let Δ^p denote the standard p-simplex with the topology inherited from Euclidean space. A *singular p-simplex* in $\mathbb{X}$ is a continuous mapping $\sigma : \Delta^p \to \mathbb{X}$.

We can formally add singular p-simplices to form singular p-chains, and the set of such chains forms a vector space $S_p(\mathbb{X})$, called the p-th singular chain group in $\mathbb{X}$. As before, the chain groups are connected by boundary maps: $\partial_p : S_p(\mathbb{X}) \to S_{p-1}(\mathbb{X})$, defined on a basis as follows. Given a singular p-simplex $\sigma : \Delta^p \to \mathbb{X}$, we define $\partial_p(\sigma)$ to be the formal sum of the singular $(p-1)$-simplices obtained by restricting the map σ to each of the $(p-1)$-faces of Δ^p in turn.

Finally, we can define singular p-cycles and singular p-boundaries exactly as before, and then define the p-th singular homology group of $\mathbb{X}$, $\mathsf{H}_p(\mathbb{X})$, to be p-cycles modulo p-boundaries, again exactly as before. Here are several important facts about singular homology:

- Amazingly, the singular homology groups of a triangulable space $\mathbb{X}$ are isomorphic to the simplicial homology groups of any simplicial complex which triangulates $\mathbb{X}$.

- As defined, the singular homology groups are clearly not algorithmically computable (for one, the vector spaces $S_p(\mathbb{X})$ are uncountably infinite in dimension!); on the other hand, they're useful for mental arithmetic, and we will see later that there are some systematic ways of understanding them.

Functoriality. Induced maps are much easier to understand in the context of singular homology. Let $f : \mathbb{X} \to \mathbb{Y}$ be a continuous map between topological spaces and let $\sigma : \Delta^p \to \mathbb{X}$ be a singular p-simplex in $\mathbb{X}$. Then $f_\#(\sigma)$ is defined to be the singular p-simplex in $\mathbb{Y}$ given by the composition $f \circ \sigma : \Delta^p \to \mathbb{Y}$.

Application. As a simple demonstration of what can be shown with these ideas, we prove the following famous theorem in all dimensions. Let $\mathbb{B}^d$ denote the closed d-dimensional unit ball.

BROUWERS'S FIXED-POINT THEOREM: every continuous map $f : \mathbb{B}^d \to \mathbb{B}^d$ has a fixed point.

proof: for purposes of contradiction, suppose that we have a continuous map $f : \mathbb{B}^d \to \mathbb{B}^d$ which satisfies $f(x) \neq x$ for all $x \in \mathbb{B}^d$. Hence for every x, there is a unique ray starting at $f(x)$, passing through x, and ending at some point $r(x)$ on the boundary sphere $\mathbb{S}^{d-1}$. Mapping each x to its $r(x)$ in this manner, we define another continuous map $r : \mathbb{B}^d \to \mathbb{S}^{d-1}$.

Now let $I : \mathbb{S}^{d-1} \to \mathbb{S}^{d-1}$ be the identity map on the sphere and let $i : \mathbb{S}^{d-1} \hookrightarrow \mathbb{B}^d$ denote the inclusion of the sphere into the ball. Note that the composition $r \circ i = I$, and so, applying the second functoriality property above, we get $r_* \circ i_* = I_*$ for all dimensions p. Let's focus on dimension $p = d - 1$. By the first property above, we know that $I_* : \mathsf{H}_{d-1}(\mathbb{S}^{d-1}) \to \mathsf{H}_{d-1}(\mathbb{S}^{d-1})$ takes the non-zero (in fact, rank one) homology group of the sphere onto itself. On the other hand, $i_* : \mathsf{H}_{d-1}(\mathbb{S}^{d-1}) \to \mathsf{H}_{d-1}(\mathbb{B}^d) = 0$ must be the zero map, since the target is the zero group! Hence we have factored an isomorphism through a zero map, which must be a contradiction. $\boxminus$

4 In-Class Exercises

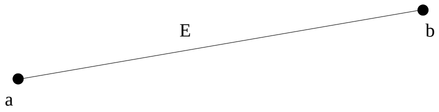

A Referring to the figure above, let K consist of the vertices a and b, and let L be the edge E along with its vertices. Let $i : K \to L$ denote the inclusion map.

(a) Is $i_{\#}(a) = i_{\#}(b)$?

(b) Let $\alpha, \beta \in \mathsf{H}_0(K)$ denote the homology classes represented by a and b, respectively. Is $i_*(\alpha) = i_*(\beta)$?. What is $i_*(\alpha + \beta)$?

B Look at Figure 1 again. Can you find a simplicial map $g : K \to L$ such that g_* is an isomorphism for every dimension p?

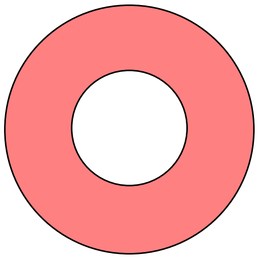

Figure 2: $\mathbb{X}$ is the annulus, and $\mathbb{X}_0$ is the union of the outer and inner circles

5 Relative Homology

The homology groups measure the higher-order connectivity of a space $\mathbb{X}$, counting, in some well-defined algebraic sense, the number of components, tunnels, voids, and so forth in $\mathbb{X}$. On the other hand, suppose we have a nested pair of spaces $(\mathbb{X}, \mathbb{X}_0)$ with $\mathbb{X}_0 \subseteq \mathbb{X}$, or pair of simplicial complexes (K, K_0) with K_0 a subcomplex of K. The *relative homology* groups of this pair will measure the connectivity of the large space relative to the small space, in a sense we now define. If a distinction is needed, we will use *absolute homology* to refer to the homology of a single space.

Definition We give the simplicial definition and leave the obvious singular changes to you. Assume we have a nested pair of subcomplexes (K, K_0). The key concept is that of *relative chain groups*: for every p, we define $\mathsf{C}_p(K, K_0) = \mathsf{C}_p(K)/\mathsf{C}_p(K_0)$. In other words, a relative p-chain is a coset $c+\mathsf{C}_p(K_0)$. Thus, two p-chains $c, c' \in \mathsf{C}_p(K)$ define the same relative p-chain iff their difference consists entirely of p-simplices in K_0. For each p, we have a boundary map $\partial_p : \mathsf{C}_p(K, K_0) \to \mathsf{C}_{p-1}(K, K_0)$ which comes from the boundary map on K: on each coset, we define $\partial_p(c + \mathsf{C}_p(K_0)) = \partial_p(c) + \mathsf{C}_{p-1}(K_0)$. Notice that this map is indeed well-defined, since the boundary of a p-chain in K_0 is a $(p-1)$-chain in K_0.

Now we follow the same algebraic procedures as in the definition of absolute homology, setting the group of relative p-cycles $\mathsf{Z}_p(K, K_0) = \ker \partial_p$, the group of relative p-boundaries $\mathsf{B}_p(K, K_0) = \operatorname{im} \partial_{p-1}$, and finally the p-th relative homology group $\mathsf{H}_p(K, K_0) = \mathsf{Z}_p(K, K_0)/\mathsf{B}_p(K, K_0)$.

Example To help unpack the definitions above, we work an example in some detail. Consider the pair of spaces $(\mathbb{X}, \mathbb{X}_0)$ where $\mathbb{X}$ is the annulus $\mathbb{S}^1 \times [0, 1]$ as drawn in Figure 2, and $\mathbb{X}_0$ is the union of the outer and the inner circles. We will mentally compute the relative homology group $\mathsf{H}_p(\mathbb{X}, \mathbb{X}_0)$ in all dimensions. To make this easier, we imagine that we have triangulated $\mathbb{X}$ in such a way that $\mathbb{X}_0$ forms a subcomplex; if you don't like doing this, then you can also just think of singular chains.

We can show that $\mathsf{H}_0(\mathbb{X}, \mathbb{X}_0) = 0$:

and $\mathsf{H}_1(\mathbb{X}, \mathbb{X}_0) = \mathbb{Z}/2\mathbb{Z}$:

and $\mathsf{H}_2(\mathbb{X}, \mathbb{X}_0) = \mathbb{Z}/2\mathbb{Z}$

Matrix reduction algorithm. The matrix reduction algorithm for computing the relative homology groups of a pair of subcomplexes (K, K_0) is almost identical to the absolute homology algorithm; the only difference lies in the ordering of the columns and the interpretation of the reduced form. First we choose an ordered basis for each chain group $\mathsf{C}_p(K)$ in such a way that the simplices in K_0 come first. Then for each p, we again create a boundary matrix D_p to represent the boundary map ∂_p; note then the columns and the rows are first indexed by simplices from K_0 and then finally by simplices in $K - K_0$. We then reduce the matrix to Smith Normal Form exactly as before, as shown below. But now we are only interested in the lower right submatrix indexed by simplices in $K - K_0$:

The form of the algorithm makes an important fact apparent. Suppose have a further pair of subcomplexes (L, L_0) which satisfies the inclusions $L \subseteq K$, $L_0 \subseteq K_0$, and the equality $L - L_0 = K - K_0$.

In this case, the EXCISION THEOREM says that the pair have isomorphic relative homology groups; that is, $\mathsf{H}_p(K, K_0) \cong \mathsf{H}_p(L, L_0)$ for all dimensions p. For a proof, we need only look at the matrix reduction above. We can obtain the boundary matrix for L by removing all rows and columns which correspond to simplices in $K - L$. But it's easy to see that $K - L = K_0$, and thus our row and column removal only affects the irrelevant upper left submatrix.

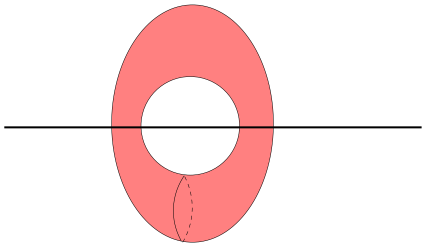

Figure 3: $\mathbb{X} = \mathbb{Y}$ is the torus, and $\mathbb{Y}_0$ is everything below the thick line.

Maps of pairs. Recall that a continuous (or simplicial) map between topological spaces (or simplicial complexes) induces a homomorphism between absolute homology groups. The same thing happens for relative homology. More precisely, suppose that we have a continuous *map of pairs* $f : (\mathbb{X}, \mathbb{X}_0) \to (\mathbb{Y}, \mathbb{Y}_0)$; this means a continuous map $f : \mathbb{X} \to \mathbb{Y}$ with the property that $f(\mathbb{X}_0) \subseteq \mathbb{Y}_0$. As before we can map every chain in $\mathbb{X}$ to a chain in $\mathbb{Y}$, but we also note that every chain in $\mathbb{X}_0$ maps to one in $\mathbb{Y}_0$. In other words, for each p we have a chain map $f_\# : \mathsf{C}_p(\mathbb{X}, \mathbb{X}_0) \to \mathsf{C}_p(\mathbb{Y}, \mathbb{Y}_0)$ which leads to a map between relative homology groups $f_* : \mathsf{H}_p(\mathbb{X}, \mathbb{X}_0) \to \mathsf{H}_p(\mathbb{Y}, \mathbb{Y}_0)$.

For a simple example, consider the map of pairs $i : (\mathbb{X}, \varnothing) \to (\mathbb{X}, \mathbb{X}_0)$, where $\mathbb{X}$ the 2-torus and $\mathbb{X}_0$ is everything sitting below the thick line in Figure 3. We describe the induced relative homology homomorphism $i_* \mathsf{H}_p(\mathbb{X}) \to \mathsf{H}_p(\mathbb{X}, \mathbb{X}_0)$ in every dimension p.

For $p = 0$, we note that $\mathsf{H}_0(\mathbb{X})$ is rank one, while $\mathsf{H}_0(\mathbb{X}, \mathbb{X}_0)$ is trivial, and so i_* must be the zero map in this dimension.

For $p = 1$, the group $\mathsf{H}_1(\mathbb{X})$ is rank two; let's take as a basis the homology class α represented by the loop drawn in Figure 3 and the homology class β represented by the meridian circle. It is easy to see that $i_*(\alpha) = 0$, since α has a representative in $\mathbb{X}_0$. On the other hand, $i_*(\beta)$ is in fact just the generator of the rank-one group $\mathsf{H}_1(\mathbb{X}, \mathbb{X}_0)$. Thus, in dimension one, i_* is surjective and has a rank-one kernel.

For $p = 2$, the group $\mathsf{H}_2(\mathbb{X})$ is rank one, with generator represented by the entire torus; this generator maps via i_* to the generator of the rank-one group $\mathsf{H}_2(\mathbb{X}, \mathbb{X}_0)$. Hence i_* is an isomorphism in dimension two.

6 Even More Exercises

A Compute (at least mentally) the relative homology groups $\mathsf{H}_p(\mathbb{B}^2, \mathbb{S}^1)$, where we think of the circle $\mathbb{S}^1$ as the boundary of the closed 2-dimensional unit ball $\mathbb{B}^2$. Generalize your answer to arbitrary dimension:

B Consider the annulus $\mathbb{X}$ as drawn in Figure 2 and recall that $\mathbb{X}_0$ is the union of the inner and the outer circles. Recall that we have already computed $\mathsf{H}_p(\mathbb{X}, \mathbb{X}_0)$ in every dimension.

- Compute the absolute homology groups $\mathsf{H}_p(\mathbb{X}_0)$ and $\mathsf{H}_p(\mathbb{X})$:

- Let i denote the inclusion of $\mathbb{X}_0$ into $\mathbb{X}$ and j denote the inclusion of pairs $(\mathbb{X}, \varnothing) \to (\mathbb{X}, \mathbb{X}_0)$. Prove that $\operatorname{im} i_* = \ker j_*$: