Homology

11 Nov, 2010

1 Overview

Today we have two main goals:

- Define the homology groups of a simplicial complex.
- Give the matrix reduction algorithm for computing the homology groups of a simplicial complex.

2 Review: Vector Space Terminology

Here we recall some standard definitions and terminology from linear algebra. First a set V along with a notion of addition $+$ forms an *Abelian group* if the following properties are satisfied

- $v, v' \in \mathsf{V}$ imply $v + v' \in \mathsf{V}$
- $v + v' = v' + v, (v + v') + v'' = v + (v' + v'')$.
- There exists a neutral element $0 \in \mathsf{V}$ which satisfies $v + 0 = v$ for every $v \in \mathsf{V}$.
- For every $v \in \mathsf{V}$, there is an inverse element $w \in \mathsf{V}$ such that $v + w = 0$.

Suppose we also have some field of scalars $\mathbb{F}$ and a notion of scalar multiplication ($r \in \mathbb{F}, v \in V group$ leads to $rv \in \mathsf{V}$. Then V is a *vector space* over $\mathbb{F}$ if vector addition and scalar multiplication interact in the expected way; we omit further details.

A set of vectors $v_1, \ldots, v_n \in \mathsf{V}$ forms a *basis* for V if every element $v \in \mathsf{V}$ can be written as $v = r_1 v_1 + \ldots + r_n v_n$ for a unique choice of scalars $r_1, \ldots, r_n \in \mathbb{F}$. Although V can have many choices of basis, the number of elements in such a basis can be shown to be fixed; this number, $\mathrm{rank}(\mathsf{V})$, is called the *dimension* or *rank* of V.

A subset $\mathsf{W} \subseteq \mathsf{V}$ forms a *vector subspace* of V if it is closed under addition and scalar multiplication. Given a fixed $v \in \mathsf{V}$, we define the *coset* $v + \mathsf{W} = \{v + w \mid w \in \mathsf{W}\}$. The *quotient space* V/W is then defined to be the set of all such cosets; it is itself a vector space over the same field, with vector addition defined by $(v + W) + (v' + W) = (v + v') + W$ and scalar multiplication given by $r(v + W) = (rv) + W$. The neutral element in V/W is of course just $\mathsf{W} = 0 + \mathsf{W}$. It can be shown that

$\mathrm{rank}(\mathsf{V}/\mathsf{W}) = \mathrm{rank}(\mathsf{V}) - \mathrm{rank}(\mathsf{W})$. If V and U are two vector spaces over $\mathbb{F}$, we can form their *direct sum* $\mathsf{V} \oplus \mathsf{U} = \{(v, u) \mid v \in \mathsf{V}, u \in \mathsf{U}\}$. This is also a vector space over $\mathbb{F}$, with vector addition and scalar multiplication defined componentwise.

A mapping $f : \mathsf{V} \to \mathsf{U}$ between vector spaces is called a *homomorphism* (or linear transformation) if f respects all vector space structure: namely, if $f(rv + sv') = rf(v) + sf(v')$ holds for all $v, v' \in V$ and $r, s \in \mathbb{F}$. We define the *kernel* of such a mapping by $\ker f = \{v \in \mathsf{V} \mid f(v) = 0\}$; note that $\ker f$ is a subspace of V. We also define the *image* of this mapping to be $\mathrm{im}\, f = \{f(v) \mid v \in \mathsf{V}\}$; note that $\mathrm{im}\, f$ is a subspace of U. Finally, it can be shown that $\mathrm{rank}(\mathsf{V}) = \mathrm{rank}(\ker f) + \mathrm{rank}(\mathrm{im}\, f)$.

3 Simplicial Homology

We now begin the definition of the simplicial homology groups for a given simplicial complex K. Our working example is shown in Figure 1.

3.1 Chain Complexes

Fix a dimension p and a field $\mathbb{F}$. A p-*chain* is a formal sum of p-simplices $\Sigma_i r_i \sigma_i$, where $r_i \in \mathbb{F}$ and the sum is taken over all possible p-simplices $\sigma_i \in K$. The set of all such p-chains is denoted $C_p(K)$. We can add two p-chains: given $c = \Sigma_i r_i \sigma_i$ and $d = \Sigma_i s_i \sigma_i$, we define $c + d = \Sigma_i (r_i + s_i)\sigma_i$. We can also multiply p-chains by scalars in the obvious way. Hence $C_p(K)$ forms a vector space over $\mathbb{F}$, and is called the group of p-chains in K; note that the set of p-simplices forms an obvious basis for $C_p(K)$ but that other bases are also possible. Hence the rank of $C_p(K)$ is simply n_p, the number of p-simplices in K. Note that the neutral element is $0 = \Sigma_i 0\sigma_i$. For $p < 0$ and $p > dim(K)$, we have $C_p(K) = 0$ since there are no simplices in those dimensions.

From now on, we will make the simplifying assumption that our field $\mathbb{F}$ is simply the binary field $\mathbb{Z}/2\mathbb{Z}$; in this case, a p-chain can be thought of as just a collection of p-simplices and adding two p-chains corresponds to taking the symmetric difference of the collections.

Boundary maps. For each p, there is a homomorphism $\partial_p : C_p \to C_{p-1}$. We define it first on p-simplices, Given a p-simplex $\sigma = [u_0, u_1, \ldots, u_p] \in C_p(K)$, we define $\partial_p(\sigma)$ to be the sum of all $(p-1)$-dimensional faces of σ. In other words,

$$\partial_p(\sigma) = \Sigma_{j=0}^{p} [u_0, \ldots, \hat{u_j}, \ldots, u_p],$$

where the hat indicates that the u_j is omitted. We then extend the definition of ∂_p by linearity, setting $\partial_p(\Sigma_i \sigma_i) = \Sigma_i \partial_p(\sigma_i)$.

For an example, consider Figure 1. Then $\partial_2(S) = A + D + E \in C_1(K)$, while $\partial_2(S + T) = \partial_2(S) = \partial_2(T) = (A + D + E) + (E + F + G) = A + D + F + G$. Note also that $\partial_1(\partial_2(S + T)) = \partial_1(A + D + F + G) = 0$.

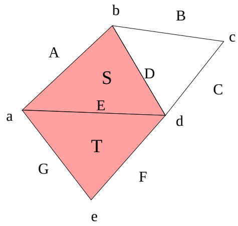

Figure 1:

We often visualize all the boundary maps together in the following sequence:

$$\cdots \xrightarrow{\partial_{p+2}} C_{p+1} \xrightarrow{\partial_{p+1}} C_p \xrightarrow{\partial_p} C_{p-1} \xrightarrow{\partial_{p-1}} \cdots$$

It is easy to see (and easy to believe from the above example) that $\partial_p \circ \partial_{p+1} = 0$ for every integer p. This implies that $\operatorname{im} \partial_{p+1}$ is a subspace of $\ker \partial_p$.

3.2 Cycles, Boundaries, Homology

Using the boundary maps, we now distinguish two special types of chains and use them to define homology groups. First we define a *p-cycle* to be a p-chain with empty boundary, that is, $c \in C_p(K)$ is a p-cycle iff $\partial_p(c) = 0$. The set of all such p-cycles forms a subspace of $C_p(K)$, which we denote as $\mathsf{Z}_p(K) = \ker \partial_p$; we set its rank to z_p. For the complex in Figure 1, we see that the group of 1-cycles $\mathsf{Z}_1(K)$ contains the chains $B + C + D, A + D + F + G$, as well as many more. On the other hand, $\mathsf{Z}_2(K) = 0$ as can be seen by direct computation. Finally, every single 0-chain forms a 0-cycle, since $C_{-1}(K) = 0$.

Given a p-chain $c \in C_p(K)$, we say that c is a *p-boundary* if there exists $d \in C_{p+1}(K)$ such that $c = \partial_{p+1}(d)$. The set of p-boundaries also forms a subspace of $C_p(K)$, and we denote it by $\mathsf{B}_p(K) = \operatorname{im} \partial_{p+1}$; we set its rank to b_p. Again referring to Figure 1, we note for example that $a+b \in \mathsf{B}_0(K)$, since $a+b = \partial_1(A)$. Furthermore, the 1-cycle $A+D+F+G$ is also a 1-boundary, while the 1-cycle $B+C+D$ is not a 1-boundary. From the remark above, it follows that every p-boundary is also a p-cycle; that is, $\mathsf{B}_p(K)$ is a subspace of $\mathsf{Z}_p(K)$.

Using this fact, we define the *p-th homology group* of K to be $\mathsf{H}_p(K) = \mathsf{Z}_p(K)/\mathsf{B}_p(K)$. The rank of this group is called the *p-th Betti number* of K, $\beta_p = \operatorname{rank}(\mathsf{H}_p(K))$. Note that $\beta_p = z_p - b_p$. (note also that $n_p = z_p + b_{p-1}$) The elements of H_p are called p-dimensional homology classes, and they correspond to cosets $c + \mathsf{B}_p$, where c is a p-cycle. Any two cycles in the same class are called homologous; note that $c, c' \in \mathsf{Z}_p(K)$ are homologous iff there is a $(p+1)$-chain d such that $c + c' = \partial(d)$.

For example, consider the three 1-cycles $\gamma_1 = B+C+D, \gamma_2 = B+C+F+G+A$, and $\gamma_3 = A + D + E$ in Figure 1. The first two of these are homologous, since $\gamma_1 + \gamma_2 = \partial(S + T)$, and so they represent the same homology class. On the other hand, γ_3 is homologous to zero, since $\gamma_3 = \partial(S)$; we say that γ_3 represents the trivial homology class. We will see soon that every 1-cycle in K is homologous either to zero or to γ_1; in other words, $\beta_1(K) = 1$. On the other hand, it is also easy to see that $\beta_2(K) = 0$ (there are no non-zero 2-cycles to begin with!) and that $\beta_0(K) = 1$.

Fun Facts about Homology Here are some amazing facts that we do not even begin to prove, although we will gesture at a proof of some of them next time:

- Despite their definition, the homology groups do not depend on choice of triangulation. In other words, no matter how we triangulate a given topological space, we will always get the same groups!
- Homology is a homeomorphism invariant.
- Even stronger: homology is a homotopy type invariant.

4 In-Class Exercises

A Let K consist of the boundary of a tetrahedron (so that K triangulates a 2-sphere). Prove formally that $H_2(K) = \mathbb{Z}/2\mathbb{Z} = H_0(K)$. Convince yourself that $H_1(K) = 0$ (a formal proof will come after the break):

B Let K be an arbitrary simplicial complex of dimension k and let K^i be the i-skeleton, with $0 \leq i < k$. Conjecture a relationship between $\mathsf{H}_p(K^i)$ and $\mathsf{H}_p(K)$, for all $p \neq i$. (hint: look very closely at the definitions!)

C Let K consist of the disjoint union of two simplicial complexes K' and K''. Conjecture a relationship between $\mathsf{H}_p(K), \mathsf{H}_p(K')$, and $\mathsf{H}_p(K'')$:

5 Matrix Reduction

We now give an algorithm which finds the Betti numbers, and indeed finds bases for each homology group, for any simplicial complex. First we produce a matrix to represent each boundary homomorphism ∂_p and then we reduce them to a simple form. From now on we assume a fixed simplicial complex K and drop all mention of it from our notation.

Boundary Matrices Recall that $\partial_p : C_p \to C_{p-1}$ is a linear transformation which takes each p-simplex to the sum of its $(p-1)$-dimensional faces. Every linear transformation can be represented by a matrix, once we fix a basis for the domain and the target spaces. We choose the set of p-simplices as a basis for C_p and the set of $(p-1)$-simplices as a basis for C_{p-1}, and we select some arbitrary but fixed ordering of these simplices. In these ordered bases, ∂_p is represented by the *boundary matrix* $D_p = [a_i^j]$ which has n_{p-1} rows and n_p columns; each row is indexed by a $(p-1)$-simplex and each column is indexed by a p-simplex. The j-th column contains a 1 for each row indexed by a $(p-1)$-face of the j-th p-simplex, and all other entries are zero. For example, the matrix D_1 corresponding to the complex in Figure 1 is drawn below:

We note that a collection of columns in D_p represents a p-chain and the sum of those columns represents the boundary of that chain.

Row and Column Operations We now want to calculate bases, and the rank, for the cycle group Z_p and the boundary group B_p. Recall that the former group is the kernel of ∂_p and hence the set of vectors which D_p sends to zero. We now perform some operations on D_p to make these vectors more apparent. We confine ourselves to two column operations which modify D_p without changing any ranks, and we do so by right-multiplying D_p by a matrix $V = [v_i^j]$:

- exchange column k with column l; here $v_k^l = v_l^k = 1$, $v_i^i = 1$, for all $i \neq k, l$ and all other entries zero.
- replace column l with the sum of column k and column l: here $v_k^l = 1, v_i^i = 1$ for all i, and all other entries zero.

The first column operation just swaps the name of basis elements, while the second replaces the l-th basis element with the sum of the k-th and the l-th, or by the sum of whatever the two columns represented before the operation.

We can also perform two row operations, each done by left-multiplying by $U = [u_i^j]$:

- exchange rows k and l: $u_k^l = u_l^k = 1$, $u_i^i = 1$, for all $i \neq k, l$ and all other entries zero.
- replace row l with the sum of row k and row l: here $u_l^k = 1$, $u_i^i = 1$ for all i, and all other entries zero.

The second row operation replaces the k-th basis element with the sum of the k-th and l-th, or by whatever these rows represented before the operation. Note that after every such operation, we still have valid bases for C_p and for C_{p-1}.

Smith Normal Form The end goal is to put our matrix into *Smith Normal Form* N_p, which means a form where some initial segment of the diagonal (or perhaps the entire diagonal) is 1 and the rest of the matrix is zero, as shown below:

We recall that the number of columns of this matrix is $\operatorname{rank}(C_p) = n_p = b_{p-1} + z_p$. We arrange it so that the leftmost b_{p-1} columns have ones on the diagonal and the rightmost z_p columns are all zero. Then the latter columns represent p-chains which have zero boundary; in other words, basis elements for Z_p. The former represent p-chains whose non-zero boundaries form a basis for B_{p-1} Hence, by reducing the matrices D_p for all p into Smith normal form, we can extract the numbers z_p and b_p, and thus the Betti numbers $\beta_p = z_p - b_p$ for every p.

To actually produce basses for Z_p and for B_{p-1}, we can keep track of the row and column operations. Writing the Smith Normal Form as $N_p = U_{p-1} D_p V_p$, we can show that the last z_p columns of V_p give a basis for Z_p and the first b_{p-1} columns of U_{p-1}^{-1} give a basis for B_{p-1}.

Reduction process and example How do we reduce D_p into Smith Normal Form? This is easy. First we perform exchanges to move a 1 to the top left corner. Using this 1, we destroy the first column and the first row. We then recurse on the smaller submatrix obtained by removing the first column and the first row. It's not hard to see that this reduction takes time at most cubic in the number of simplices. We now perform the algorithm for the three boundary matrices D_0, D_1, and D_2, corresponding the 2-dimensional simplicial complex in Figure 1

6 More Exercises

Turn your brain off and produce and then reduce the boundary matrices for the following simplicial complexes. Read the Betti numbers from your answer. If you're feeling adventurous, also read off bases for the cycle and boundary groups:

A your favorite triangulation of a three-dimensional closed ball (answer: $\beta_0 = 1, \beta_1 = \beta_2 = \beta_3 = 0$):

B your favorite triangulation of the 2-sphere ($\beta_0 = \beta_2 = 1, \beta_1 = 0$):

C your favorite triangulation of the torus ($\beta_0 = \beta_2 = 1, \beta_1 = 2$):