Math 872 Algebraic Topology

Running lecture notes

Homology theory: Fundamental groups are a remarkably powerful tool for studying spaces; they capture a great deal of the global structure of a space, and so they are very good a distinguishing between homotopy-inequivalent spaces. In theory! But in practice, they suffer from the fact that deciding whether two groups are isomorphic or not is, in general, undecideable! Homology theory is designed to get around this deficiency; the theory, by design, builds (a sequence of) *abelian* groups $H_i(X)$ from a topological space. And deciding whether or not two abelian groups, at least if you're given a presentation for them, is, in the end, a matter of fairly routine linear algebra. Mostly because of the Fundamental Theorem of Finitely-generated Abelian groups; each such has a unique representation as $\mathbb{Z}^m \oplus \mathbb{Z}_{m_1} \oplus \cdots \oplus \mathbb{Z}_{m_n}$ with $m_{i+1}|m_i$ for every i .

There are also "higher" homotopy groups beyond the fundamental group π_1, (hence the name pi-*one*); elements are homtopy classes, rel boundary, of based maps $(I^n, \partial I^n) \to (X, x_0)$. Multiplication is again by concatenation. But unlike π_1, where we have a chance to compute it via Seifert-van Kampen, nobody, for example knows what all of the homotopy groups $\pi_n(S^2$ are (except that nearly all of them are non-trivial!). Like π_1, it describes, essentially, maps of S^n into X which don't extend to maps of D^{n+1}, i.e., it turns the "n-dimensional holes" of X into a group.

Homology theory does the exact same thing, counting n-dimensional holes. In the end we will find it to be extremely computable; but it will require building a fair bit of machinery before it will become so transparent to calculate. But the short version is that the homology groups compute "cycles mod boundaries", that is, n-dimesional objects/subsets that have no boundary (in the appropriate sense) modulo objects that are the boundary of $(n+1)$-dimensional ones. There are, in fact, probably as many ways to *define* homology groups as there are people actively working in the field; we will focus on two, simplicial homology and singular homology. The first is quick to define and compute, but hard to show is an invariant! The second is quick to see is an invariant, but, on the face of it, hard to compute! Luckily, for spaces where they are both defined, they are isomorphic. So, in the end, we get an invariant that is quick to compute. Of course, so is the invariant "4"; but this one will be a bit more informative....

First, simplicial homology. This is a sequence of groups defined for spaces for which they are easiest to define, which Hatcher calls Δ-complexes. Basically, they are spaces defined by gluing simplices together using nice enough maps. More precisely, the *standard n-simplex* Δ^n is the set of points $\{(x_1, \ldots x_{n+1}) \in \mathbb{R}^{n+1} : \sum x_i = 1, x_i \geq 0 \text{ for all } i\}$. This can also be expressed as convex linear combinations (literally, that's the conditions on the x_i's) of the points $e_i = (0, \ldots, 0, 1, 0, \ldots, 0)$, the *vertices* of the standard simplex. More generally, an n-simplex is the set $[v_0, \ldots v_n]$ of convex linear combinations of points $v_0, \ldots, v_n \in \mathbb{R}^k$ for which $v_1 - v_0, \ldots, v_n - v_0$ are linearly independent. Any bijection from the vertices of the standard simplex to the points $v_0, \ldots, v_n$ extends (linearly) to a homeomorphism of the simplices. The $n+1$ *faces* of a simplex, each sitting opposite a vertex v_i, are obtained by setting the corresponding coefficient x_i to 0. Each forms an $(n-1)$-simplex, which we

denote $[v_0, \ldots, v_{i-1}, v_{i+1}, \ldots, v_n]$ or $[v_0, \ldots, \widehat{v_i}, \ldots, v_n]$. A Δ-*complex* X is a cell complex obtained by gluing simplices together, but we insist on an extra condition: the restriction of the attaching map to any face is equal to a (lower-dimensional) cell. As before, we use the weak topology on the space; a set is open iff it's inverse image under the induced map of a cell into the complex is open. Each n-cell comes equipped with a (continuous) map $\sigma : \Delta^n \to X$, which is one-to-one on its interior, whose restriction to the boundary is the attaching map, and whose restriction to each face is the associated map for that $(n-1)$-simplex. We will typically blur the distinction between the map σ (called the *characteristic map* of the simplex) and its image, and denote the image by σ (or σ^n), when this will cause no confusion, and call σ an n-simplex *in* X. When we feel the need for the distinction, we will use e^n for the image and σ^n for the map.

For example, taking our standard, identifications of the sides of a rectangle, cell structure for the 2-torus, and cutting the rectangle into two triangles (= 2-simplices) along a diagonal, we obtain a Δ-structure with 2 2-simplices, 3 1-simplices, and 1 0-simplex. A genus g surface can be built, by cutting the $2g$-gon into triangles, with $g+1$ 2-simplices, $3g$ 1-simplices, and 1 0-simplex.

We typically think of building a Δ-complex X inductively. The *0-simplices* (i.e., points), or *vertices*, form the 0-skeleton $X^{(0)}$. n-simplices $\sigma^n = [v_0, \ldots v_n]$ attach to the $(n-1)$-skeleton to form the n-skeleton $X^{(n)}$; the restriction of the attaching map to each face of σ^n is, by definition, an $(n-1)$-simplex in X. The attaching map is (by induction) really determined by a map $\{v_0, \ldots, v_n\} \to X^{(0)}$, since this determines the attaching maps for the 1-simplices in the boundary of the n-simplex, which gives 1-simplices in X, which then give the attaching maps for the 2-simplices in the boundary, etc. Note that the reverse is not true; the vertices of two different n-simplices in X can be the same. For example, think of the 2-sphere as a pair of 2-simplices whose boundaries are glued by the identity. Δ-complexes generalize *simplicial complexes* where the simplices are required to attach by homeomorphisms to the skeleton, and the intersection of two simplices are a (single) subsimplex of each. This has the advantage over Δ-complexes that an n-simplex is determined uniquely by the set of vertices in $X^{(0)}$ that it attaches to. This means that, in principle, a simplicial complex (and everything associated with it, e.g., its homology groups!) can be treated purely combinatorially; the complex is "really" a certain collection of subsets of the vertices (since these determine the simplices), with the property that any subset B of a subset A that has been declared to be a simplex is also a simplex. But they have the disadvantage that it typically takes far more simplices to build a simplicial structure on a space X that it does to build a Δ-structure. This makes the computations we are about to do take far longer.

The final detail that we need before defining (simplicial) homology groups is the notion of an *orientation* on a simplex of X. Each simplex σ^n is determined by a map $f : \{v_0, \ldots, v_n\} \to X^{(0)}$; an orientation on σ^n is an (equivalence class of) the ordered $(n+1)$-tuple $(f(v_0), \ldots f(v_n)) = (V_0, \ldots, V_n)$. Another ordering of the same vertices represents the same orientation if there is an *even* permutation taking the entries of the first $(n+1)$-tuple to the second. This should be thought of as a generalization of the right-hand rule for $\mathbb{R}^3$, interpreted as orienting the vertices of a 3-simplex. Note that there are

precisely two orientations on a simplex.

Now to define homology! We start by defining *n-chains*; these are (finite) formal linear combinations of the (oriented!) n-simplices of X, where $-\sigma$ is interpreted as σ with the opposite (i.e., other) orientation. Adding formal linear combinations formally, we get the n-th *chain group* $C_n(X) = \{\sum n_\alpha \sigma_\alpha : \sigma_\alpha \text{ an oriented } n\text{-simplex in } X\}$. We next define a *boundary operator* $\partial : C_n(X) \to C_{n-1}(X)$, whose image will be the $(n-1)$-chains that are the "boundaries" of n-chains. The idea is that the boundary of a 2-simplex, for example, should be a "sum" of its three faces (since they do make up the boundary of the simplex), but we need to take into account their orientations, in order to be getting the correct sum. Thinking of the orientation on a 1-simplex $[v, w]$ as an arrow pointing from v to w, we are lead to believe that the boundary of a 2-simplex $[u, v, w]$ should be $[u, v] + [v, w] + [w, u]$. Similarly, the boundary of $[u, v]$, on reflection, should be $[v] - [u]$, to distinguish the head of the arrow (the + side) from the tail (the − side). On the basis of these examples, trying to find a consistent formula, one might eventually be led to the following formulation. We define the boundary on the basis elements $\sigma_\alpha = \sigma$ of $C_n(X)$ as $\partial\sigma = \sum(-1)^i \sigma|_{[v_0,\dots,\widehat{v_i},\dots,v_n]}$, where $\sigma : [v_0, \dots, v_n] \to X$ is the characteristic map of σ_α . $\partial\sigma$ is therefore an alternating sum of the faces of σ. We then extend the definition by linearity to all of $C_n(X)$. When a notation indicating dimension is needed, we write $\partial = \partial_n$. We define $\partial_0 = 0$.

This definition, it utrns out, is cooked up to make the maxim "boundaries have no boundary" true; that is, $\delta_{n-1} \circ \delta_n = 0$, the 0 map. This is because, for any simplex $\sigma = [v_0, \dots v_n]$,

$$\delta \circ \delta(\sigma) = \delta(\sum_{i=0}^{n}(-1)^i \sigma|_{[v_0,\dots,\widehat{v_i},\dots,v_n]})$$
$$= (\sum_{j<i}(-1)^j(-1)^i \sigma|_{[v_0,\dots,\widehat{v_j},\dots,\widehat{v_i},\dots,v_n]}) + (\sum_{j>i}(-1)^{j-1}(-1)^i \sigma|_{[v_0,\dots,\widehat{v_i},\dots,\widehat{v_j},\dots,v_n]})$$

The distinction between the two pieces is that in the second part, v_j is actually the $(j-1)$-st vertex of the face. Switching the roles of i and j in the second sum, we find that the two are negatives of one another, so they sum to 0, as desired.

And this little calculation is all that it takes to define homology groups! What this tells us is that $\text{im}(\delta_{n+1}) \subseteq \ker(\delta_n$ for every n. $\ker(\delta_n = Z_n(X)$ are called the *n-cycles* of X; they are the n-chains with 0 (i.e., empty) boundary. They form a (free) abelian subgroup of $C_n(X)$. $\text{im}(\delta_{n+1} = B_n(X)$ are the *n-boundaries* of X; they are, of course, the boundaries of $(n+1)$-chains in X. The n-th homology group of X, $H_n(X)$ is the quotient $Z_n(X)/B_n(X)$; it is an abelian group.

A key observation is that the boundary maps δ_n are linear, that is, they are homomorphisms between the free abelian groups $\delta_n : C_n(X) \to C_{n-1}(X)$. Consequently, they can be expressed as (integer-valued) matrices Δ_n. Row reducing Δ_n (over the integers!) allows us to find a basis $v_1, \dots, v_k$ for $Z_n(X)$ (clearing denomenators to get vectors over $\mathbb{Z}$). Then since $\Delta_n\Delta_{n+1} = 0$, the columns of Δ_{n+1} are in the kernel of Δ_n, so can be expressed as linear combinations of the v_i . These combinations can be determined by row reducing the augmented matrix $(v_1 \cdots v_k|\Delta_{n+1})$. This will row reduce to $\left(\begin{array}{c|c} I & C \\ 0 & 0 \end{array}\right)$, and C basically

describes the boundaries $B_n(X)$ in terms of the basis $v_1, \ldots, v_k$. The homology group $H_n(X)$ is then the *cokernel* of C, i.e., $\mathbb{Z}^k/\mathrm{im}C$. Note that C will have integer entries, since we know that the columns of Δ_{n+1} can be expressed as integer linear combinations of the v_i, and, being a basis, there is only one such expression.

Some examples: the Klein bottle K has a Δ-complex structure with 2 2-simplices, 3 1-simplices, and 1 0-simplex; we will call them $f_1 = [0,1,2], f_2 = [1,2,3], e_1 = [0,2] = [1,3], e_2 = [1,0] = [2,3], e_3 = [1,2]$, and $v_1 = [0] = [1] = [2] = [3]$. Computing, we find $\partial_2 f_1 = \partial[0,1,2] = [1,2] - [0,2] + [0,1] = e_3 - e_1 - e_2$, $\partial_2 f_2 = e_2 - e_1 + e_3$, $\partial_1 e_1 = \partial_1 e_2 = \partial_1 e_3 = 0$ and $\partial_i = 0$ for all other i (as well). So we have the chain complex $\cdots \to 0 \to \mathbb{Z}^2 \to \mathbb{Z}^3 \to \mathbb{Z} \to 0$
and all of the boundary maps are 0, except for ∂_2, which has the matrix $\begin{pmatrix} -1 & -1 \\ -1 & 1 \\ 1 & 1 \end{pmatrix}$. This matrix is injective, so $\ker \partial_2 = 0$, so $H_2(K) = 0$, on the other hand, $H_1(K) = \mathrm{coker}(\partial_2)$, and applying column operations we can transform the matrix for ∂_2 to $\begin{pmatrix} 1 & 0 \\ 1 & 2 \\ -1 & 0 \end{pmatrix}$, which implies that the cokernel is $\mathbb{Z} \oplus \mathbb{Z}_2$, since $\begin{pmatrix} 1 \\ 1 \\ -1 \end{pmatrix}, \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix}$ is a basis for $\mathbb{Z}^3$. Finally, $H_0(K) = \mathbb{Z}$, since $\partial_1, \partial_0 = 0$, and all higher homology groups are also 0, for the same reason.

As another example, the topologist's dunce hat has a Δ-structure with 1 2-simplex, 1 1-simplex, and 1 0-simplex. The boundary maps, we can work out (starting from $C_2(X)$), are $(1), (0)$, and (0), so $H_2(X) = H_1(X) = 0$, and $H_0(X) = \mathbb{Z}$. all higher groups are also 0.

These homology groups are, in the end, fairly routine to calculate from a Δ-complex structure. But there is one very large problem; the calculations depend on the Δ structure! This is not a group defined from the space X; it is defined from the space and a Δ structure on it. A priori, we don't know that if we chose a different structure on the same space, that we would get isomorphic groups! We should really denote our groups by $H_i^\Delta(X)$, to acknowledge this dependence on the structure.

But we don't *want* a group that depends on this structure. We want groups that just depend on the topological space X, i.e., which are topological invariants. In really turns out that these groups $H_i^\Delta(X)$ are topological invariants, but we will need to take a very roundabout route to show this. What we will do now is to define another sequence $H_i(X)$ of groups, the *singular homology groups*, which their definition makes apparent from the outset that they are topological invariants. But this definition will also make it very unclear how to really compute them! Then we will show that for Δ-complexes these two sequences of groups are really the same. In so doing, we will have built a sequence of topological invariants that for a large class of spaces are fairly routine to compute. Then all we will

need to show is that they also capture useful information about a space (i.e., we can prove useful theorems with them!).

And the basic idea behind defining them is that, with simplicial homology, we have already done all of the hard work. What we do is, as before, build a sequence of (free) abelian groups, the chain groups $C_n(X)$, and boundary maps between them, with consecutive maps composing to 0. Then, as before, the homology groups are kernels mod images, i.e., cycles mod boundaries. And, as before, the basis elements for each of our chain groups $C_n(X)$ will be the n-simplices in X. But now X is any topological space. So how do we get n-simplices in such a space? We do the only thing we can; we *map* them in.

More precisely, we work with *singular n-chains*, that is, formal (finite) linear combinations $\sum a_i\sigma_i$, where $a_i \in BbbZ$ and the σ_i are *singular simplices*, that is, (continuous) maps $\sigma_i : \Delta^n \to X$ from the (standard) n-simplex into X. The boundary maps are really exactly as before; they are the alternating sum of the restrictions of σ_i to the $n+1$ faces of Δ^n . (Formally, we must precompose these face maps with the (orientation-preserving) linear isomorphism from the standard $(n-1)$-simplex to each of the faces, preserving the ordering of their vertices.) The same proof as before (except that we interpret the faces as restrictions of the map σ_i, instead of as physical faces) shows that the composition of two successive boundaries are 0, and so all of the machinery is in place to define the *singular homology groups* $H_i(X)$ as the kernel of ∂_i modulo the image of $\partial_{i+1} = Z_i(X)/B_i(X)$. They are, by their definition, groups defined using the topological space X as input, and so are topological invariants of X. The elements are equivalence classes of i-cycles, where $z_1 \simeq z_2$ if $z_1 - z_2 = \partial w$ for some $(i+1)$-chain w . We say that z_1 and z_2 are *homologous*.

Singular homology groups are very quick to define, but what do they measure? The basic idea is that we are trying to mimic simplicial homology, but because a general topological space X cannot be thought of as being built out of simplices, we do the next best thing; we study the space by mapping simplices in. Formally, this is what we did with simplicial homology anyway, except that we restricted ourselves to a very few special singular simplices (the characteristic maps of the building blocks for X). In the end an n-cycle $\sum a_i\sigma_i^n$, since the faces of the σ_i must match up precisely, in order to cancel in the sum, can be thought of as a map of an n-complex into X, made by gluing the n-simplices σ_i together before mapping in. The fact that faces cancel really means that these simplices restrict to the same maps on their faces. The integer coefficients can really be interpreted as taking multiple copies of Δ^n and gluing them together along their boundaries (the signs tell us the underlying orientations). The idea is that this n-complex is being mapped "around a hole", unless it extends to a map of an $(n+1)$-complex into X (having our n-complex as boundary). So singular homology really is trying to detect holes, it is just doing it with maps.....

The "fun" with singular homology groups, though, comes when you try to compute them. $C_n(X) = \{\sum a_i\sigma_i : a_i \in \mathbb{Z} \text{ and } \sigma_i : \Delta^n \to X \text{ is continuous}\}$ is typically a huge group, since there will be immense numbers of maps $\Delta^n \to X$. About the only space for which this is not true is the one-point space $*$; then there is, for each n, exactly one (distinct) map $\sigma_n : \Delta^n \to *$; the constant map. Therefore each face of Δ^n gives the same restriction

map σ^{n-1}, and so the boundary maps can be dirctly computed (the depend on the parity of the number $n+1$ of faces an n-simplex has). We find that $\partial_{2n} = Id$ and $\partial_{2n-1} = 0$. so in computing homology groups, we either have kernel everything ($\partial_i = 0$) and image everything ($\partial_{i+1} = Id$) or kernel nothing ($\partial_i = Id$) and image nothing ($\partial_{i+1} = 0$), so in both cases $H_i(*) = 0$. Except for $i = 0$; then $\partial_0 = 0$ (by definition) and $\partial_1 = 0$, so $H_0(*) = \mathbb{Z}$. But other than this example (and, well, OK, spaces with the discrete topology; it's the same calculation as above for every point!), computing singular homology from the definition is quite a chore! so we need to build some labor-saving devices, namely, some theorems to help us break the problem of computing these groups into smaller, more managable pieces.

First set of managable pieces: if we decompose X into its path components, $X = \bigcup X_\alpha$, then $H_i(X) \cong \bigoplus H_i(X_\alpha)$ for every i. This is because every singular simplex, since Δ^i is path-connected, maps into some X_α . So $C_i(X) \cong \bigoplus C_i(X_\alpha)$. Since the boundary of a simplex mapping into X_α consists of simplices in X_α, the boundary maps respect the decomposistions of the chain groups, so $B_i(X) \cong \bigoplus B_i(X_\alpha)$ and $Z_i(X) \cong \bigoplus Z_i(X_\alpha)$, and so the quotients are $H_i(X) \cong \bigoplus H_i(X_\alpha)$.

So, if we wish to, we can focus on computing homology groups for path-connected spaces X. For such a space, $H_0(X) \cong \mathbb{Z}$, generated by any map of a 0-simplex (= a point) into X. This is because any pair of 0-simplices are homologous; given any two points $x, y \in X$, there is a path $\gamma : I \to X$ from x to y, This path can be interpreted as a singular 1-simplex, and $\partial\gamma = y - x$. So $H_0(X)$ is generated by a single point $[x]$. No multiple of this point is null-homologous, because for any 1-chain $\sum n_i\sigma_i$, the sum of the coefficients of its boundary is 0 (since this is true for each singular 1-simplex), and any 0-chain $\sum n_i[x_i]$ is homologous to $(\sum n_i)[x]$ by the above argument.

A small techincal aside: the fact that $H_0(*) = \mathbb{Z}$ is annoying to some, and often requires treating 0-dimensional homology as a special case. But since the boundary of a singular 1-simplex is always of the form $v - w$, we find that the image of ∂_1 is always contained in the subgroup of $C_0(X)$ consisting of chains whose coefficients sum to 0. This means that we can, for free, *augment* the singular chain complex by a map $\cdots \to C_1(X) \xrightarrow{\partial_1} C_0)X) \xrightarrow{\alpha} \mathbb{Z} \to 0$ where α is the map $\alpha(\sum a_i\sigma_i^0) = \sum a_i$. This is still a chain complex (compositions of consecutive maps are 0); the resulting homology groups are called *reduced* homology $\widetilde{H}_i(X)$. The only affect this really has is to remove one copy of $\mathbb{Z}$ from H_0; $\widetilde{H}_0(X) \oplus \mathbb{Z} \cong H_0(X)$. All other homology groups are unchanged. There is a reduced relative homology as well, since we can augment with the same map (1-simplices always have 2 ends!), but in this case it has (essentially) no effect; $\widetilde{H}_i(X, A) \cong H_i(X, A)$ for all i <u>unless</u> $A = \emptyset$, in which case we lose the $\mathbb{Z}$ in dimension 0 that we expect to.

Perhaps the most important property of the fundamental group is that a continuouos map between spaces induces a homomorphism between groups. (This implied, for instance, that homeomorphic spaces have isomorphic π_1). The same is true for homology groups, for essentially the same reason. Given a map $f : X \to Y$, there is an induced map $f_\# : C_n(X) \to C_n(Y)$ defined by postcomposition; for a singular simplex σ, $f_\#(\sigma) = f \circ \sigma$, and we extend the map linearly. Since $f \circ (g|_A) = (f \circ g)|_A$ (postcomposition

commutes with restriction of the domain), $f_\#$ commutes with ∂ : $f_\#(\partial\sigma) = \partial(f_\#(\sigma))$. A homomorphism between chain complexes (i.e., a sequence of such maps, one for each chain group) which commutes with the boundaries maps in this way, is called a *chain map*. A chain map, such as $f_\#$, therefore, takes cycles to cycles, and boundaries to boundaries, and so $f_\# : Z_i(X) \to Z_i(Y)$ (which is linear, hence a homomorphism) induces a homomorphism $f_* : H_i(X) \to H_i(Y)$ by $f_*[z] = [f_\#(z)]$. Since it is defined by composition with singular simplices, it is immediate that, for a map $g : Y \to Z$, $(g \circ f)_* = g_* \circ f_*$. And since the identity map $I : X \to X$ satisfies $I_\# = Id$, so $I_* = Id$, homeomorphic spaces have isomorphic homology groups.

Another important property of π_1 is that homotopic maps give the same induced map (after correcting for basepoints). This is also true for homology; if $f \simeq g : X \to Y$, then $f_* = g_*$. The proof, however, is not quite as straightforward as for homotopy. And it requires some new technology; the chain homotopy. A chain homotopy H between the chain complexes $f_\#, g_\# : C_*(X) \to C_*(Y)$ is a sequence of homomorphisms $H_i : C_i(X) \to C_{i+1}(Y)$ satisfying $H_{i-1}\partial_i + \partial_{i+1}H_i = f_\# - g_\# : C_i(X) \to C_i(Y)$. The existence of H implies that $f_* = g_*$, since for an i-cycle z (with $\partial_i(z) = 0$) we have
$f_*[z] - g_*[z] = [f_\#(z) - g_\#(z)] = [H_{i-1}\partial_i(z) + \partial_{i+1}H_i(z)] = [H_{i-1}(0) + \partial_{i+1}(w)] = [\partial_{i+1}(w)] = 0$.
And the existence of a homotopy between f and g implies the existence of a chain homotopy between $f_\#$ and $g_\#$. This is because the homotopy gives a map $H : X \times I \to Y$, which induces a map $H_\# : C_{i+1}(X \times I) \to C_{i+1}(Y)$. Then we pull, from our back pocket, a *prism map* $P : C_i(X) \to C_{i+1}(X \times I)$; the composition $H_\# \circ P$ will be our chain homotopy. The prism map takes a (singular) i-simplex σ and sends it to a sum of singular $(i+1)$-simplices in $X \times I$. and the way we define it is to take the i-simplex Δ^i, and taking it to $\Delta^i \times I$ (i.e., a *prism*), and thinking of this as a sum of $(i+1)$-simplices. Using the map $\sigma' = \sigma \times Id : \Delta^i \times I \to X \times I$ restricted to each of these $(i+1)$-simplices yields the prism map. Now, there are many ways of decomposing a prism into simplices, but we need to be careful to choose one which restricts well to each of the <u>faces</u> of Δ^i, in order to get the chain homotopy property we require. In the end, what this requires is that the decomposition, when restricted to any face of Δ^i (which we think of as a copy of Δ^{i-1}), is the same as the decomposition we would have applied to a prism over an $(i-1)$-simplex. After some exploration, we are led to the following formulation.

If we write $\Delta^n \times \{0\} = [v_0, \ldots, v_n]$ and $\Delta^n \times \{1\} = [w_0, \ldots, w_n]$, then we can decompose $\Delta^n \times I$ as the (n+1)-simplices $[v_0, \ldots, v_i, w_i, \ldots, w_n]$. We then define $P(\sigma) = \sum(-1)^i \sigma'|_{[v_0,\ldots,v_i,w_i,\ldots,w_n]}$. A routine calculation verifies that $(\partial P + P\partial)(\sigma) = \sigma'|_{[w_0,\ldots,w_n]} - \sigma'|_{[v_0,\ldots v_n]}$; Composing with $H_\#$ yields our result.

Consequently, for example, homotopy equivalent spaces have isomporphic (reduced) homology groups; homotopy equivalences induce isomorphisms. So all contractible spaces have trivial reduced homology in all dimensions, since they are all homotopy to a point. If we think of a cell complex as a collection of disks glued together, this lends some hope that we can compute their homology groups, since we can compute the homology of the building blocks. Our next goal is to make turn this idea into action; but we need another tool, to frame our answer in the best way possible.

Exact sequences: Most of the fundamental properties of homology groups are described in terms of exact sequences. A sequence of homomorphisms $\cdots \xrightarrow{f_{n+1}} A_n \xrightarrow{f_n} A_{n-1} \xrightarrow{f_{n-1}} a_{n-2} \to \cdots$ of abelian groups is called *exact* if $\mathrm{im}(f_n) = \ker(f_{n-1})$ for every n. In most cases, we get the most mileage out of an exact sequence when some of the groups are trivial; $0 \to A \xrightarrow{f} B$ is exact $\Leftrightarrow$ f is injective, and $A \xrightarrow{f} B \to 0$ is exact $\Leftrightarrow$ f is surjective. An exact sequence $0 \to A \to B \to C \to 0$ is called a *short exact sequence.*

The main tool we will use turns a family of short exact sequences of chain maps between three chain complexes into a single *long exact homology sequence.* Given chain complexes $\mathcal{A} = (A_n, \partial)$, $\mathcal{B} = (B_n, \partial')$, and $\mathcal{C} = (C_n, \partial'')$ and short exact sequences of chain maps (i.e., $\partial' i_n = i_n \partial$, $\partial'' j_n = j_n \partial'$)
$0 \to A_n \xrightarrow{i_n} B_n \xrightarrow{j_n} C_n \to 0$ there is a general result which provides us with a long exact sequence

$$\cdots \xrightarrow{\partial} H_n(\mathcal{A}) \xrightarrow{i_*} H_n(\mathcal{B}) \xrightarrow{j_*} H_n(\mathcal{C}) \xrightarrow{\partial} H_{n-1}(\mathcal{A}) \xrightarrow{i_*} \cdots$$

Most of the work is in defining the "boundary" map ∂. Given an element $[z] \in H_n(\mathcal{C})$, a representative $z \in C_n$ satisfies $\partial''(z) = 0$. But j_n is onto, so there is a $b \in B_n$ with $j_n(b) = z$, Then $i_{n-1}\partial'(b) = \partial'' j_n(b) = 0$, so $\partial'(b) \in \ker(j_{n-1} = \mathrm{im}(a_{n-1})$. So there is an $a \in A_{n-1}$ with $i_{n-1}(a) = \partial'(b)$. But then $i_{n-2}\partial(a) = \partial' i_{n-1}(a) = \partial'\partial'(b) = 0$, so, since i_{n-2} is injective, $\partial a = 0$, so $a \in Z_{n-1}(\mathcal{A})$, and so represents a homology class $[a] \in H_n(\mathcal{A})$. We define $\partial([z]) = [a]$.
To show that this is well-defined, we need to show that the class $[a]$ we end up with is independent of the choices made along the way. The choice of a was not really a choice; i_{n-1} is, by assumption, injective. For b, if $j_n(b) = z = j_n(b')$, then $j_n(b - b') = 0$, so $b - b' = i_n(w)$ for some $w \in A_n$. Then $\partial' b' = \partial' b - \partial' i_n(w) = \partial' b - i_{n-1}\partial(w)$, so choosing $a' = a - \partial(w)$ we have $i_{n-1}(a') = \partial'(b')$. But then $[a'] = [a - \partial w] = [a] - [delw] = [a]$. Finally, there is actually a choice of z ; if $[z] = [z']$, then $z' = z + \partial'' w$ for some $w \in C_{n+1}$; but then choosing b', w' with $j_n(b') = z'$, $j_{n+1}(w') = w$, we have
$\partial'' w = \partial'' j_{n+1}(w') = j_n \partial'(w')$, so
$z' = z + \partial'' w = j_n(b + \partial' w')$, so we may choose $b' = b + \partial' w'$ (since the result is independent of this choice!), then since $\partial' b' = \partial' b$ everything continues the same.

Now to exactness! We need to show three (types of) equalities, which means six containments. Three (image contained in kernel) are shown basically by showing that compositions of two consecutive homomorphisms are trivial. $j_n i_n = 0$ immediately implies $j_* i_* = 0$. From the definition of ∂, $i_* \partial[z] = [i_n(a)] = [\partial'(b)] = 0$, and $\partial j_*[z] = \partial[j_n(z)] = [a]$, where $i_{n-1}(a) = \partial'(z) = 0$, so $a = 0$ (since i_{n-1} is injective), so $[a] = 0$.

For the opposite containments, if $j_*[z] = [j_n(z)] = 0$, then $j_n(z) = \partial'' w$ for some w. Since j_{n+1} is onto, $w = j_{n+1}(b)$ for some b. Then $j_n(z - \partial' b) = \partial'' w - \partial'' j_{n+1} b = 0$, so $z = \partial' b = i_n(a)$ for some a, so $i_*[a] = [z - \partial' b] = [z]$. So $\ker j_* \subseteq \mathrm{im} i_*$. If $i_*[z] = 0$, then $i_n(z) = \partial' w$ for some $w \in B_{n+1}$. Setting $c = j_{n+1}(w)$, then $\partial'' c = j_n \partial' w - i_n i_n(Z) = 0$, so $[c] \in h_{n+1}(\mathcal{C})$, and computing $\partial[c]$ we find that we can choose w for the first step and z for the second step, so $\partial[c] = [z]$. So $\ker j_n \subseteq \mathrm{im}\partial$. Finally, if $\partial[z] = 0$, then $z = j_n(b)$ for some b, and $\partial' b = i_{n-1}(a)$ with $[a] = 0$, i.e., $a = \partial w$ for some w. So $\partial' b = i_{n-1}\partial w = \partial' i_n w$ But then $\partial'(b - i_n w) = 0$, and $j_n(b - i_n w) = z - 0 = z$, so $z \in \mathrm{im}(j_n)$, so $[z] \in \mathrm{im}(j_*)$. So $\ker \partial \subseteq \mathrm{im}(j_n)$. Which finishes the proof!

Now all we need are some new chain complexes!