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 detecting 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 are isomorphic, 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 homotopy 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 same thing, it counts n-dimensional holes. In the end, it is extremely computable; but we will need a fair bit of machinery before it will become transparent to calculate. The short version is that 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.

We will focus on two approached to homology: simplicial and singular. 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, at the start, hard to compute! But for spaces where they are both defined, they are isomorphic. So between the two we get an invariant that is quick to compute.

Simplicial homology: This is a sequence of groups defined for spaces called Δ-complexes. They are a particular kind of CW-complex, defined by gluing simplices together using "nice enough" maps.

More precisely, the *standard n-simplex* Δ^n is the set of points

$$\{(x_1, \dots x_{n+1}) \in \mathbb{R}^{n+1} : \sum x_i = 1, x_i \geq 0 \text{ for all } i\}.$$

This is the set of *convex* linear combinations of the points $e_i = (0, \dots, 0, 1, 0, \dots, 0)$, the *vertices* of the standard simplex. An n-simplex is the set $[v_0, \dots v_n]$ of convex linear combinations of points $v_0, \dots, v_n \in \mathbb{R}^k$ for which $v_1 - v_0, \dots, v_n - v_0$ are linearly independent. Any bijection $\{\text{vertices of } \Delta^n\} \to \{v_0, \dots, v_N\}$ extends (linearly) to a homeo b/w Δ^n and $[v_0, \dots v_n]$. The *faces* of a simplex, each 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 $[v_0, \dots, v_{i-1}, v_{i+1}, \dots, v_n] = [v_0, \dots, \widehat{v_i}, \dots, v_n]$.

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) simplex. As before, we use the weak topology on the space; a set is open iff it's inverse image under the induced map of each simplex into the complex is open. Each n-cell comes equipped with a characteristic 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 characteristic map for that $(n-1)$-simplex. We will typically blur the distinction between the characteristic map σ 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 as a cell structure for the 2-torus T^2, and cutting the rectangle into two triangles (= 2-simplices) along a diagonal, we obtain a Δ-structure for T^2 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.

As with CW-cplxes, we typically think of building a Δ-complex X inductively. The *0-simplices* or *vertices* form the 0-skeleton $X^{(0)}$. n-simplices $\sigma^n = [v_0, \dots v_n]$ 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 an $(n-1)$-simplex in X. This attaching map is really determined by a map $\{v_0, \dots, v_n\} \to X^{(0)}$, 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.

The reverse is not true; the vertices of two different n-simplices in X can be the same. For example, build 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) sub-simplex 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 $f : \{v_0, \dots, v_n\} \to X^{(0)}$; an orientation on σ^n is an (equivalence class of) the ordered $(n+1)$-tuple $(f(v_0), \dots f(v_n)) = (V_0, \dots, V_n)$. 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.