Homotopy Groups and the Fundamental Group Action

Algebraic Topology I | Fall 2026 | October 2026

9.1 Based Homotopy Groups

The homology groups studied in Lectures 5–8 measure cycles modulo boundaries. Homotopy groups start with a different geometric object: a map from a sphere. Two such maps represent the same element when one can be continuously deformed into the other while a chosen point remains fixed. Keeping track of that point is essential. It makes multiplication possible, and the ambiguity in moving it is itself measured by the fundamental group.

This lecture begins Part III of the syllabus. We construct homotopy groups, prove their basic properties, and explain their dependence on basepoints. We then state the Whitehead and Hurewicz theorems, whose proofs belong to Math 616. The sphere calculations at the end use the statement of Hurewicz and the integral homology of spheres already computed in Lecture 7. Throughout, $I = [0, 1]$; all maps and homotopies are continuous. Every occurrence of space homology includes its coefficient group explicitly.

9.1.1 Based Maps and Homotopy Classes

A basepoint is part of the data of a space in this discussion. It must be preserved at every stage of a based homotopy, not just at its two endpoints.

Definition 9.1.1. A *based space* (X, x) is a space with a distinguished point $x \in X$. A based map $f : (X, x) \to (Y, y)$ is a map satisfying $f(x) = y$. A based homotopy from f to g is a map $H : X \times I \to Y$ satisfying

$$H(z, 0) = f(z), \qquad H(z, 1) = g(z), \qquad H(x, s) = y.$$

We write $[X, Y]_*$ for the set of based homotopy classes when the basepoints have been specified. An ordinary homotopy, without the last condition, is also called a *free homotopy* in this context.

Based homotopy is an equivalence relation. The constant homotopy proves reflexivity; replacing s by $1 - s$ proves symmetry; concatenating two homotopies in the parameter s proves transitivity. The pasting lemma proves continuity of concatenation because the two formulas agree at the joining time. These same elementary operations will occur repeatedly below.

Definition 9.1.2. The set $\pi_0(X, x)$ is the set of path components of X, with the component containing x distinguished. A based map induces a function of pointed sets by sending the component of z to the component of $f(z)$. In general, $\pi_0(X, x)$ has no natural group operation.

Equivalently, give the two-point space S^0 one distinguished point. A based map $\mathrm{S}^0 \to X$ is specified by the image of its other point, and a based homotopy is a path between those images. Thus $\pi_0(X, x) = [\mathrm{S}^0, X]_*$ as pointed sets. Path components, rather than connected components, are the appropriate invariant: a homotopy supplies actual paths.

For $n \geq 1$, write I^n for the n-cube and ∂I^n for the points having at least one coordinate equal to 0 or 1. A map of pairs $f : (I^n, \partial I^n) \to (X, x)$ means simply a map $f : I^n \to X$ with $f(\partial I^n) = \{x\}$.

Definition 9.1.3. For $n \geq 1$, define $\pi_n(X, x)$ to be the set of maps $f : (I^n, \partial I^n) \to (X, x)$ modulo homotopies relative to ∂I^n. Such a homotopy is a map $H : I^n \times I \to X$ with the prescribed endpoint maps and $H(u, s) = x$ whenever $u \in \partial I^n$. The constant map c_x gives the distinguished class.

The quotient $I^n/\partial I^n$ is a sphere S^n, with the collapsed boundary as basepoint. One way to see this is to identify the cube with a disk, radially from their centers, and then identify the disk with a closed hemisphere whose boundary is collapsed to the opposite pole. Equivalently, the quotient is the one-point compactification of the cube interior, which is homeomorphic to $\mathbb{R}^n$. The quotient universal property identifies cube maps with based sphere maps. It also identifies their homotopies: the quotient map times I is a quotient map because it is a continuous surjection between compact Hausdorff spaces. Consequently

$$\pi_n(X, x) = [\mathrm{S}^n, X]_*.$$

A representative can therefore be pictured either as a sphere in X or as a cube whose entire boundary maps to x.

For $n = 1$ this is a loop at x. The group we are about to construct is the *fundamental group*. The interval model has two endpoints, but both are sent to the same point; this condition makes one loop follow another.

9.1.2 Concatenation and the Group Axioms

Our convention is that paths and loops are multiplied in *travel order*: $a * b$ traverses a first and b second. Composition of maps retains the usual convention, so $g \circ f$ applies f first. Distinguishing these conventions prevents a reversal later in the basepoint action.

Construction 9.1.4. Write $u = (t, v)$ with $t \in I$ and $v \in I^{n-1}$. For based cube maps $f, g : I^n \to X$, set

$$(f * g)(t, v) = \begin{cases} f(2t, v), & 0 \leq t \leq \tfrac{1}{2}, \\ g(2t - 1, v), & \tfrac{1}{2} \leq t \leq 1. \end{cases}$$

The two formulas agree on $t = \frac{1}{2}$ because each has value x. They define a continuous based cube map. Define $[f][g] = [f * g]$.

Theorem 9.1.5. For every $n \geq 1$, concatenation makes $\pi_n(X, x)$ a group. Its identity is $[c_x]$, and the inverse of $[f]$ is represented by

$$\bar{f}(t, v) = f(1 - t, v).$$

Proof. First suppose F and G are based homotopies of f and g. Applying the concatenation formula to $F(-, s)$ and $G(-, s)$ for every s gives a based homotopy of the products. Thus the operation is well-defined on classes.

For associativity, allocate intervals of lengths $a, b, c > 0$, with $a + b + c = 1$, to the three maps f, g, h. On these intervals use respectively

$$f(t/a, v), \qquad g((t - a)/b, v), \qquad h((t - a - b)/c, v).$$

These formulas agree at the joining faces. The lengths for $(f * g) * h$ are $(1/4, 1/4, 1/2)$, while those for $f * (g * h)$ are $(1/2, 1/4, 1/4)$. Interpolate linearly between these triples. All lengths stay positive, so the piecewise formulas vary continuously and define a based homotopy. This proves associativity of classes; it does not assert equality of the original parametrized maps.

For the right identity, let $a_s = (1 + s)/2$ and put

$$R_s(t, v) = \begin{cases} f(t/a_s, v), & t \leq a_s, \\ x, & t \geq a_s. \end{cases}$$

At $s = 0$ this is $f * c_x$, and at $s = 1$ it is f. The closed-piece pasting lemma gives continuity also at $t = a_s$. The analogous left-identity homotopy keeps a constant interval $0 \le t \le (1-s)/2$ and rescales the remaining interval to run through f.

Finally, $f * \bar{f}$ runs the first coordinate from 0 to 1 and back to 0. A contraction of that excursion is

$$K(t, v, s) = \begin{cases} f(2t(1-s), v), & t \le \frac{1}{2}, \\ f(2(1-t)(1-s), v), & t \ge \frac{1}{2}. \end{cases}$$

At $s = 0$ it is $f * \bar{f}$; at $s = 1$ it is x, since $f(0, v) = x$. All boundary points stay at x. Replacing f by $\bar{f}$ proves that $\bar{f} * f$ also contracts. This proves both inverse identities and completes all group axioms. □

The same argument applies to paths with different endpoints whenever the concatenations are defined. In particular, a path followed by its reverse is homotopic relative to its endpoints to the constant path. The identity and associativity statements concern endpoint-preserving homotopy classes, exactly as for loops.

9.1.3 The Eckmann–Hilton Argument

When $n \ge 2$, there is a second coordinate in which to concatenate. The resulting interchange law is the geometric reason higher homotopy groups are abelian.

Lemma 9.1.6. Suppose a set has two operations $\star$ and $\circ$ with the same two-sided identity e, and suppose

$$(a \star b) \circ (c \star d) = (a \circ c) \star (b \circ d)$$

for every four elements. Then the operations coincide and are commutative.

Proof. Inserting identity elements in the interchange law gives

$$a \circ b = (a \star e) \circ (e \star b) = (a \circ e) \star (e \circ b) = a \star b.$$

Inserting them in the other two positions gives

$$a \circ b = (e \star a) \circ (b \star e) = (e \circ b) \star (a \circ e) = b \star a.$$

Together these equalities prove the assertions. □

Theorem 9.1.7. For $n \ge 2$, the group $\pi_n(X, x)$ is abelian, and concatenation in any coordinate induces the same operation.

Proof. Let $\star$ concatenate in the first coordinate and let $\circ$ concatenate in the second. The group-law proof applies to either coordinate, so both operations have identity $[c_x]$. Their interchange law holds already for representative maps. Subdivide the first two coordinates at $1/2$. Both sides send the four resulting subcubes through f, g, h, k, respectively, in the same positions and with the same linear rescalings. All joining faces map to x, so no extra homotopy is required. The preceding lemma proves that the two operations agree and commute. Repeating the argument with the first and the jth coordinates proves independence of the coordinate. □

We therefore use additive notation for $\pi_n(X, x)$ when $n \ge 2$: the constant class is 0, reflection represents the negative, and concatenation represents addition. We retain multiplicative notation for $\pi_1(X, x)$. A one-dimensional interval has no second independent coordinate, so the preceding argument says nothing about commutativity of the fundamental group.

9.2 Functoriality and Basepoint Transport

9.2.1 Functoriality and Elementary Examples

Postcomposition carries a sphere in one space to a sphere in another. The definitions make this operation compatible with both homotopies and concatenation.

Proposition 9.2.1. A based map $f : (X, x) \to (Y, y)$ induces

$$f_* : \pi_n(X, x) \longrightarrow \pi_n(Y, y), \qquad [a] \longmapsto [f \circ a].$$

For $n = 0$ this is a function of pointed sets; for $n \geq 1$ it is a homomorphism, and for $n \geq 2$ a homomorphism of abelian groups. Identity maps induce identities, and $(g \circ f)_* = g_* \circ f_*$.

Proof. Postcomposing a based homotopy of cube maps with f gives a based homotopy in Y, so the formula is independent of representatives. For positive n, the defining formulas give the equality of maps

$$f \circ (a * b) = (f \circ a) * (f \circ b),$$

which proves preservation of multiplication. Constant maps remain constant. Identity and composition laws follow from the corresponding equalities for ordinary maps. For $n = 0$, a path in X maps to a path in Y, and the distinguished component maps to the distinguished component. □

Proposition 9.2.2. If $f, g : (X, x) \to (Y, y)$ are based homotopic, then $f_* = g_*$ on $\pi_n(X, x)$ for all $n \geq 0$. A based homotopy equivalence induces isomorphisms in positive degrees and a bijection in degree zero.

Proof. Let H be a based homotopy from f to g. For a based cube map a, the formula $(u, s) \mapsto H(a(u), s)$ is a homotopy relative to the boundary between $f \circ a$ and $g \circ a$. For $n = 0$, the path $s \mapsto H(z, s)$ connects $f(z)$ to $g(z)$. If g is a based homotopy inverse of f, functoriality now gives $g_* f_* = (gf)_* = \mathrm{id}$ and $f_* g_* = \mathrm{id}$. □

Thus π_n is a functor on the based homotopy category, with target pointed sets, groups, or abelian groups according to the degree. This is the based counterpart of the homotopy-invariant functors considered in Lecture 5. An ordinary homotopy need not be based; its precise effect is the subject of the next section.

Before making difficult calculations, it is useful to record what the construction does to elementary geometric operations. These arguments work directly with representatives.

Proposition 9.2.3. For based spaces (X, x) and (Y, y), the projections induce a natural isomorphism

$$\pi_n(X \times Y, (x, y)) \cong \pi_n(X, x) \times \pi_n(Y, y)$$

for $n \geq 1$, and a bijection of pointed sets for $n = 0$.

Proof. A map into $X \times Y$ is exactly a pair of maps into the factors. The same is true of homotopies, and the boundary condition holds exactly when it holds in both factors. The correspondence on classes is consequently bijective. Concatenation is performed coordinate-wise in the target product, so it is a group isomorphism in positive degrees. For $n = 0$, two pairs of points are joined by a path precisely when their coordinates can be joined by paths; the two coordinate paths use the same parameter interval. □

Proposition 9.2.4. If $i : (A, a) \to (X, a)$ and $r : (X, a) \to (A, a)$ satisfy $r \circ i = \mathrm{id}_A$, then i_* is injective and r_* is surjective in every degree. If A is a deformation retract of X by a homotopy fixed on A, then these maps are inverse isomorphisms in positive degrees and inverse bijections in degree zero.

Proof. Functoriality gives $r_*i_* = \mathrm{id}$, proving the first assertions, including for pointed sets. In the deformation-retract case, ir is based homotopic to id_X, since $a \in A$ stays fixed. Hence $i_*r_* = \mathrm{id}$ as well. □

Example 9.2.5. A point has trivial homotopy groups in positive degrees and a one-element π_0. If X is discrete, every map from a connected cube is constant, so $\pi_n(X,x) = 0$ for $n \geq 1$, whereas $\pi_0(X,x)$ is the underlying pointed set of X. A convex subset of a real vector space contracts to any of its points by straight lines, fixing that point; its positive homotopy groups therefore vanish.

For arbitrary contractible spaces, a given contraction need not fix a prescribed basepoint. We will prove the same vanishing conclusion after accounting for moving basepoints. That small distinction is a useful test of whether a homotopy-invariance argument has actually handled its basepoints.

9.2.2 Transport Along a Path

Let $\alpha : I \to X$ be a path with $\alpha(0) = x$ and $\alpha(1) = y$. We want to carry a based sphere at x to one at y. Geometrically, shrink its cube and surround it by a collar on which the path runs from x on the inner edge to y on the outer edge.

For explicit formulas, let $c = (1/2, \dots, 1/2)$ and

$$r(u) = \max_{1 \leq i \leq n} |2u_i - 1|.$$

The equation $r(u) = 1$ describes the boundary; $r(u) \leq 1/2$ describes the concentric cube with half the side length. Scalar multiplication of $u - c$ below is coordinatewise.

Construction 9.2.6. For $f : (I^n, \partial I^n) \to (X, x)$, define $T_\alpha f : I^n \to X$ by

$$(T_\alpha f)(u) = \begin{cases} f(c + 2(u-c)), & r(u) \leq \frac{1}{2}, \\ \alpha(2r(u) - 1), & r(u) \geq \frac{1}{2}. \end{cases}$$

At the common boundary both expressions equal x, and on ∂I^n the value is y. Thus $T_\alpha f$ is a based cube at y. Set

$$\tau_\alpha : \pi_n(X,x) \longrightarrow \pi_n(X,y), \qquad \tau_\alpha([f]) = [T_\alpha f].$$

Applying the same formula to a based homotopy of f, or to an endpoint-preserving homotopy of α, proves that τ_α depends only on the two relevant homotopy classes. It remains to prove the algebraic properties. The following lemma is a convenient way to do all of them without relying on a picture of a moving collar.

Lemma 9.2.7. There is a homotopy from f to $T_\alpha f$ whose boundary has value $\alpha(s)$ at time s. More generally, if $H : I^n \times I \to X$ satisfies

$$H(u,0) = f(u), \quad H(u,1) = g(u), \quad H(u,s) = \alpha(s) \text{ for } u \in \partial I^n,$$

then $[g] = \tau_\alpha([f])$ in $\pi_n(X,y)$.

Proof. For the first assertion, put $a_s = 1 - s/2$ and define

$$M(u,s) = \begin{cases} f(c + (u-c)/a_s), & r(u) \leq a_s, \\ \alpha(2r(u) - 2 + s), & r(u) \geq a_s. \end{cases}$$

On the joining surface, the first argument lies on ∂I^n and the second argument is 0, so both values are x. The formula is continuous on its two closed regions, including at $s = 0$, since $a_s \geq 1/2$. At time 0 it gives f, at time 1 it gives $T_\alpha f$, and on the boundary it gives $\alpha(s)$.

For the general assertion, convert H into a homotopy with boundary fixed at y. Define

$$K(u,s) = \begin{cases} H(c + 2(u-c), s), & r(u) \le \frac{1}{2}, \\ \alpha(s + (1-s)(2r(u)-1)), & r(u) \ge \frac{1}{2}. \end{cases}$$

The joining values are $\alpha(s)$, and the outer boundary value is $\alpha(1) = y$. Hence K is a based homotopy at y. It starts at $T_\alpha f$ and ends at the map obtained by shrinking g into the central cube and filling the collar constantly with y. That last map is based homotopic to g: use the first formula of this proof for the constant path at y, and run the resulting homotopy backward. Therefore $[T_\alpha f] = [g]$. □

9.2.3 Composition, Inverses, and Naturality

Transport should compose in the same order that paths are traversed. The moving-boundary lemma proves this and also shows that transport respects the group operation.

Theorem 9.2.8. For every $n \ge 1$, transport is a group isomorphism. If α runs from x to y, β runs from y to z, and c_x is the constant path, then

$$\tau_{c_x} = \mathrm{id}, \qquad \tau_{\alpha * \beta} = \tau_\beta \circ \tau_\alpha, \qquad \tau_{\bar\alpha} = \tau_\alpha^{-1},$$

where $\bar\alpha(t) = \alpha(1-t)$. For any map $q : X \to Y$, transport is natural:

$$q_* \circ \tau_\alpha = \tau_{q\circ\alpha} \circ q_*.$$

Here the induced maps have the basepoints supplied by the displayed path endpoints.

Proof. For the constant path, the first homotopy in the moving-boundary lemma is based, proving $\tau_{c_x} = \mathrm{id}$. Apply that lemma first along α to f and then along β to $T_\alpha f$. Concatenating the two homotopies gives a moving-boundary homotopy from f to $T_\beta T_\alpha f$, with boundary path $\alpha * \beta$. The second part of the lemma gives

$$[T_\beta T_\alpha f] = \tau_{\alpha*\beta}([f]),$$

which is the composition formula. The paths $\alpha * \bar\alpha$ and $\bar\alpha * \alpha$ contract relative to their endpoints by the excursion contraction used in Theorem 9.1.5. Since transport depends only on the endpoint-preserving path class, the composition and identity formulas prove the inverse formula.

To prove preservation of multiplication, choose the moving-boundary homotopies M_f and M_g for two cubes at x. At each time, concatenate them in the first cube coordinate. Along the common face their values are both $\alpha(s)$, so the formulas paste continuously. The resulting homotopy runs from $f * g$ to $(T_\alpha f) * (T_\alpha g)$ and has boundary path α. The moving-boundary lemma gives

$$\tau_\alpha([f][g]) = \tau_\alpha([f])\tau_\alpha([g]).$$

Together with the inverse formula, this proves that transport is a group isomorphism. Finally, the collar formula gives the equality of representative maps $q \circ T_\alpha f = T_{q\circ\alpha}(q \circ f)$, proving naturality. □

For $n = 1$, reading the collar from the left endpoint to the right gives the especially useful formula

$$\tau_\alpha([f]) = [\bar\alpha * f * \alpha].$$

The three pieces have lengths $1/4, 1/2, 1/4$ in the collar model. Changing these positive lengths does not change their homotopy class, by the reparametrization argument in the group-law

proof. The order has a geometric meaning: a loop at y must first travel backward along α to reach the old basepoint x.

If X is path connected, all its groups $\pi_n(X,x)$ with fixed $n \geq 1$ are therefore isomorphic. A particular isomorphism still requires a path. The next section identifies exactly the ambiguity in that choice.

9.3 The Fundamental Group Action and Homotopy Invariance

9.3.1 A Left Action and Free Homotopy Classes

Returning a basepoint to its starting point produces an automorphism rather than a new group. Reversing the transporting loop converts this construction into a left action with the usual conjugation formula in degree one.

Definition 9.3.1. For $n \geq 1$, define a left action of $\pi_1(X,x)$ on $\pi_n(X,x)$ by

$$[a] \cdot u = \tau_{\bar{a}}(u).$$

For $n \geq 2$ these are automorphisms of an abelian group. In degree 1 the same formula is the action by inner automorphisms.

Proposition 9.3.2. The preceding formula is a well-defined left action by group automorphisms. In degree 1 it is

$$[a] \cdot [f] = [a][f][a]^{-1}.$$

For a based map $q : (X,x) \to (Y,y)$, induced maps respect the actions:

$$q_*([a] \cdot u) = q_*([a]) \cdot q_*(u).$$

Proof. Path-homotopy independence and the automorphism property follow from Theorem 9.2.8. The constant loop acts as the identity. Reversal of a concatenation gives $\overline{a * b} = \bar{b} * \bar{a}$, so

$$\tau_{\overline{a*b}} = \tau_{\bar{b}*\bar{a}} = \tau_{\bar{a}} \circ \tau_{\bar{b}}.$$

This says exactly $([a][b]) \cdot u = [a] \cdot ([b] \cdot u)$, the left-action law. In degree 1, apply the explicit transport formula to the path $\bar{a}$ to obtain $[a * f * \bar{a}]$. Naturality of transport, together with $q \circ \bar{a} = \overline{q \circ a}$, proves equivariance. □

Using τ_a itself would instead define a right action. Both conventions occur in the literature. Our convention is fixed by two requirements: paths multiply in travel order, and the left action on π_1 is conjugation $u \mapsto aua^{-1}$.

Proposition 9.3.3. Suppose α and β both run from x to y. Then

$$\tau_{\beta}^{-1} \circ \tau_{\alpha} = \tau_{\alpha * \bar{\beta}} \quad \text{on } \pi_n(X,x).$$

Consequently, if X is simply connected, transport between any two points is independent of the chosen path.

Proof. The inverse and composition formulas give $\tau_{\beta}^{-1}\tau_{\alpha} = \tau_{\bar{\beta}}\tau_{\alpha} = \tau_{\alpha*\bar{\beta}}$. By definition, simply connected means path connected with trivial fundamental group. In that case every loop at x, including $\alpha * \bar{\beta}$, represents the constant class, so the displayed automorphism is the identity.□

A path-connected space is called *simple* if the fundamental-group action on every positive homotopy group is trivial. In degree 1 this requires the fundamental group to be abelian.

Abelianity alone does not assert that the actions in higher degrees vanish. A simply connected space is automatically simple. The point of this terminology is that even abelian higher homotopy groups can carry additional information through their fundamental-group actions.

Write $[\mathrm{S}^n, X]$ for free homotopy classes of maps, with no basepoint condition on maps or homotopies. Forgetting the basepoint gives a function from $\pi_n(X, x)$ to this set.

Proposition 9.3.4. If X is path connected and $n \geq 1$, forgetting the basepoint induces a bijection

$$\pi_n(X, x)/\pi_1(X, x) \longrightarrow [\mathrm{S}^n, X],$$

where the left side is the set of orbits for the action just defined. In particular, free homotopy classes of loops correspond to conjugacy classes in $\pi_1(X, x)$.

Proof. First consider an arbitrary sphere map f. Its chosen sphere basepoint has some image $y \in X$. Choose a path from y to x. In the cube model, the moving-boundary homotopy transports f to a map based at x and is an ordinary sphere homotopy throughout. Thus forgetting the basepoint is surjective.

Now let f and g both be based at x, and suppose H is a free homotopy between them. The trace $a(s) = H(s_0, s)$ of the sphere basepoint s_0 is a loop at x. Pulling back to the cube gives a moving-boundary homotopy with boundary path a. Hence $[g] = \tau_a([f]) = [\bar{a}] \cdot [f]$, so the two based classes belong to the same orbit.

Conversely, if $[g] = [a] \cdot [f] = \tau_{\bar{a}}([f])$, the collar construction supplies a free homotopy from f to a representative of $[g]$. Follow it by a based homotopy to g. Thus the fibers of the forgetful map are exactly the orbits. In degree 1, Proposition 9.3.2 identifies these with conjugacy classes. □

For a simply connected target, the action is trivial, so every free homotopy class of sphere maps has a unique based homotopy class after a basepoint is specified. For a general target, forgetting the basepoint can genuinely identify different elements; the orbit set need not itself carry a useful group structure.

9.3.2 Homotopies and Homotopy Equivalences

A free homotopy carries its own basepoint comparison: follow the chosen point through the homotopy. This path supplies precisely the correction needed for the induced homomorphisms.

Theorem 9.3.5. Let $H : X \times I \to Y$ be a homotopy from f to g, and choose $x \in X$. The path $\alpha(s) = H(x, s)$ runs from $f(x)$ to $g(x)$. For $n \geq 1$,

$$g_* = \tau_\alpha \circ f_* : \pi_n(X, x) \longrightarrow \pi_n(Y, g(x)).$$

For $n = 0$, the two maps induce the same function on the underlying sets of path components.

Proof. For a based cube $a : I^n \to X$, the composite $(u, s) \mapsto H(a(u), s)$ has endpoint maps $f \circ a$ and $g \circ a$. On the cube boundary it equals $H(x, s) = \alpha(s)$. The moving-boundary lemma gives the formula. In degree zero, $s \mapsto H(z, s)$ is a path from $f(z)$ to $g(z)$ for every z, so both images lie in the same component. □

If f and g preserve a basepoint but the homotopy does not, α is a loop. Their induced homomorphisms need not be equal: the formula relates them by the action of its reversed class. This is why an unqualified assertion that “homotopic maps induce the same map on homotopy groups” is incomplete. Equality follows when the homotopy is based, or when the relevant transport action is trivial.

Even if neither a homotopy inverse nor its inverse homotopies preserve the selected basepoints, the trace formula still gives isomorphisms. The following argument keeps the intermediate basepoints explicit.

Corollary 9.3.6. If $f : X \to Y$ is an ordinary homotopy equivalence, then for every $x \in X$ the map

$$f_* : \pi_n(X, x) \longrightarrow \pi_n(Y, f(x))$$

is an isomorphism for $n \geq 1$, and f induces a bijection on path components.

Proof. Choose a homotopy inverse $g : Y \to X$. For each x, a homotopy $gf \simeq \mathrm{id}_X$ and the trace formula show that

$$g_* f_* : \pi_n(X, x) \longrightarrow \pi_n(X, gf(x))$$

is an isomorphism: composing it with the transport along the trace gives the identity. Thus f_* is injective at every basepoint. Applying the same argument to $fg \simeq \mathrm{id}_Y$ shows that g_* is injective at every basepoint of Y.

To see that f_* is surjective at the chosen x, let $v \in \pi_n(Y, f(x))$. Since $g_* f_*$ is surjective, choose $u \in \pi_n(X, x)$ with $g_* f_*(u) = g_*(v)$. Injectivity of g_* at $f(x)$ then gives $f_*(u) = v$. This proves the isomorphism without requiring either homotopy to fix a basepoint. The homotopies induce identities on component sets, so f and g give inverse bijections on π_0 as well. □

Example 9.3.7. If X is nonempty and contractible, the map $X \to *$ is a homotopy equivalence. Therefore $\pi_0(X, x)$ has one element and $\pi_n(X, x) = 0$ for every $n \geq 1$ and every choice of x. No assumption that a chosen contraction fixes x is needed.

9.4 The Whitehead and Hurewicz Theorems

9.4.1 Weak Equivalences and the Whitehead Theorem

A homotopy equivalence preserves all the invariants constructed so far. The converse needs geometric hypotheses on the spaces and a single map realizing the isomorphisms.

Definition 9.4.1. A map $f : X \to Y$ is a *weak homotopy equivalence* if it induces a bijection on path components and, for every $x \in X$ and every $n \geq 1$, an isomorphism

$$f_* : \pi_n(X, x) \longrightarrow \pi_n(Y, f(x)).$$

By naturality of transport, it is enough to test one basepoint in each path component of X.

The last assertion follows by choosing a path from the tested point to another point in its component. The two transport maps give a commutative square for f_* with isomorphisms on the vertical sides. Thus the isomorphism property is unchanged by the choice of basepoint within a component. The condition on π_0 must be stated separately; positive homotopy groups at one point cannot see other components.

Theorem 9.4.2. A weak homotopy equivalence $f : X \to Y$ between CW complexes is a homotopy equivalence. In particular, for connected CW complexes it is enough that f_* be an isomorphism on every positive homotopy group at one basepoint.

The proof of the Whitehead theorem is deferred to Math 616. The CW hypotheses are part of the theorem: for arbitrary spaces, a weak homotopy equivalence need not be a homotopy equivalence. The conclusion also requires an actual map inducing the isomorphisms; abstract isomorphisms between lists of homotopy groups are not the stated hypothesis. Recall from

Lecture 8 that CW complexes are built by attaching disks in successive dimensions, with their weak topology. That geometry is what makes the theorem possible.

As an immediate application of the statement, a nonempty CW complex with one path component and all positive homotopy groups zero is contractible. Indeed, its map to a point is a weak homotopy equivalence, and Whitehead makes it a homotopy equivalence. This is stronger than the already-proved fact that a contractible space has vanishing homotopy groups.

9.4.2 The Hurewicz Map and Theorem

Fix an orientation class $[\mathrm{S}^n] \in \mathrm{H}_n(\mathrm{S}^n;\mathbb{Z})$ for each $n \geq 1$. We use the cube orientation and the associated quotient sphere orientation. The natural way to turn a sphere map into a homology class is to push forward this class.

Construction 9.4.3. For $n \geq 1$, the *Hurewicz map* is

$$h_n : \pi_n(X,x) \longrightarrow \mathrm{H}_n(X;\mathbb{Z}), \qquad h_n([f]) = f_*([\mathrm{S}^n]).$$

It is well-defined because singular homology is homotopy invariant. Its target is ordinary integral singular homology, without a basepoint choice.

Proposition 9.4.4. The Hurewicz map is a natural homomorphism. It is unchanged by basepoint transport: for every path α from x to y,

$$h_n \circ \tau_\alpha = h_n.$$

In particular, it is constant on the orbits of the fundamental-group action.

Proof. The sphere model for concatenation is the pinch map

$$p : \mathrm{S}^n \longrightarrow \mathrm{S}^n \vee \mathrm{S}^n,$$

which collapses the middle face of the cube, together with its boundary. Orient the two quotient cubes by their inherited coordinates. The induced map on integral homology sends the sphere class to the sum of the two oriented sphere classes. To check the signs and coefficients, triangulate the cube so that its middle face is a subcomplex. Its oriented relative fundamental chain is the sum of the oriented chains of its two halves; the shared face cancels with opposite boundary orientations. After pinching, the two summands are exactly the respective sphere fundamental classes. This also works when $n = 1$, with the two positively oriented half-intervals.

The map representing the product is $(f \vee g) \circ p$, so pushing the last identity into X gives

$$h_n([f][g]) = h_n([f]) + h_n([g])$$

in multiplicative notation for the source. This proves the homomorphism assertion in all positive degrees. Naturality follows from $(qf)_* = q_* f_*$ on homology. Finally, f and $T_\alpha f$ are freely homotopic by Lemma 9.2.7; homotopy invariance of singular homology gives the transport identity. Replacing α by a loop or its reverse gives the orbit assertion. □

The range in which homotopy agrees with homology is controlled by the vanishing of lower homotopy groups. We first fix terminology so that the theorem's hypotheses are unambiguous.

Definition 9.4.5. A space is 0-*connected* if it is nonempty and path connected. For $r \geq 1$, it is r-*connected* if it is path connected and $\pi_i(X,x) = 0$ for $1 \leq i \leq r$. Transport makes this condition independent of the chosen basepoint. Thus 1-connected means simply connected.

Theorem 9.4.6. If $n \geq 2$ and X is $(n-1)$-connected, then $\widetilde{\mathrm{H}}_i(X;\mathbb{Z}) = 0$ for $0 \leq i < n$, and

$$h_n : \pi_n(X,x) \longrightarrow \mathrm{H}_n(X;\mathbb{Z})$$

is an isomorphism. If X is path connected, the degree-one map is surjective with kernel the commutator subgroup, and hence induces an isomorphism

$$\pi_1(X,x)^{\mathrm{ab}} \cong \mathrm{H}_1(X;\mathbb{Z}).$$

Both assertions of the Hurewicz theorem are stated here without proof; their proofs are deferred to Math 616. The degree-one assertion is also commonly called the first Hurewicz theorem. More explicitly, if $G = \pi_1(X,x)$ and $[G,G]$ is its commutator subgroup, it gives the exact sequence of groups

$$1 \longrightarrow [G,G] \longrightarrow G \xrightarrow{h_1} \mathrm{H}_1(X;\mathbb{Z}) \longrightarrow 0.$$

The final group is written additively. The kernel is exactly $[G,G]$, not merely a subgroup containing all commutators.

For $n = 2$, the hypothesis is simple connectivity and the conclusion is $\pi_2(X,x) \cong \mathrm{H}_2(X;\mathbb{Z})$. For larger n, all lower positive homotopy groups must vanish. The theorem compares the first possible nonzero positive homotopy group of a simply connected space with integral homology; it does not identify every homotopy group with homology in the same degree. The Hopf fibration in Lecture 10 will exhibit a nonzero third homotopy group of S^2, although $\mathrm{H}_3(\mathrm{S}^2;\mathbb{Z}) = 0$.

9.5 Sphere Calculations and the Next Step

9.5.1 Connectivity of Spheres Without Cellular Approximation

To apply the Hurewicz statement to spheres, we need a separate connectivity argument. The following proof uses only compactness, fine triangulations of a cube, and elementary Euclidean geometry. In particular, it does not presume the cellular approximation theorem.

Lemma 9.5.1. For $m \geq 1$, the sphere S^m is path connected, and

$$\pi_k(\mathrm{S}^m, x) = 0 \qquad (1 \leq k < m).$$

Thus S^m is $(m-1)$-connected.

Proof. The unit sphere in $\mathbb{R}^{m+1}$ is path connected for $m \geq 1$: nonantipodal points can be joined by normalizing their straight segment; antipodal points can be joined using any third point not on their common line. Fix $1 \leq k < m$ and a based cube map $f : I^k \to \mathrm{S}^m$ with boundary value x.

By uniform continuity, choose a finite triangulation of I^k, with its boundary a subcomplex, sufficiently fine that the values of f at any two points of one simplex are less than $1/2$ apart in Euclidean distance. Such triangulations are obtained by subdividing a cubical grid into simplices. Define $F : I^k \to \mathbb{R}^{m+1}$ by taking the value $f(v)$ at each vertex v and extending affinely over each simplex. The affine definitions agree on common faces. Since every point of a simplex is a convex combination of its vertices,

$$\|F(u) - f(u)\| < \tfrac{1}{2} \quad \text{for every } u \in I^k.$$

Also $F = x$ on the boundary, because all boundary vertices have value x. Every vector $(1-s)f(u)+sF(u)$ is within $1/2$ of the unit vector $f(u)$ and is therefore nonzero. Normalizing gives a based homotopy from f to

$$g(u) = \frac{F(u)}{\|F(u)\|}.$$

On any simplex, F takes values in the linear span of at most $k+1$ vertex vectors, and g takes values in the intersection of that span with the sphere. Since $k+1 \le m < m+1$, each such span is a proper linear subspace of $\mathbb{R}^{m+1}$. Finitely many proper linear subspaces cannot cover the sphere: each is annihilated by a nonzero linear functional, and the product of these functionals is a nonzero homogeneous polynomial. If it vanished on the unit sphere, homogeneity would make it vanish at every nonzero vector and hence everywhere, which is impossible. Choose $p \in \mathrm{S}^m$ outside all these spans, also different from x. The image of g avoids p.

Let $U_p = \{z \in \mathrm{S}^m : z \neq p\}$ be the sphere punctured at p. Stereographic projection gives a homeomorphism $\sigma : U_p \to \mathbb{R}^m$. The formula

$$(u,s) \longmapsto \sigma^{-1}\big((1-s)\sigma(g(u)) + s\sigma(x)\big)$$

is a homotopy from g to the constant map at x, fixed on the boundary. Combining the two homotopies proves $[f] = 0$. □

9.5.2 The First Nonzero Group and Degree

The connectivity lemma verifies the hypothesis of the stated Hurewicz theorem for each positive-dimensional sphere. In dimensions at least two, the previously computed integral homology now determines the first nonzero homotopy group.

Corollary 9.5.2. For $n \ge 2$, the Hurewicz map identifies

$$\pi_n(\mathrm{S}^n, x) \cong \mathrm{H}_n(\mathrm{S}^n; \mathbb{Z}) \cong \mathbb{Z},$$

and the identity sphere map represents the positive generator when source and target orientations agree. In particular, $\pi_2(\mathrm{S}^2) \cong \mathbb{Z}$ and $\pi_3(\mathrm{S}^3) \cong \mathbb{Z}$.

Proof. Lemma 9.5.1 verifies the $(n-1)$-connectedness hypothesis of the stated Hurewicz theorem. Lecture 7 computed the integral homology group in dimension n as $\mathbb{Z}$, generated by the chosen sphere orientation. By its definition, the Hurewicz map sends the identity class to that generator. □

This proof uses Hurewicz as a stated theorem, rather than supplying its deferred proof. The case $n = 1$ will be established in Lecture 11 by lifting loops along $\mathbb{R} \to \mathrm{S}^1$; the higher-dimensional argument above does not replace that lifting calculation.

Corollary 9.5.3. For $n \ge 2$, two based maps $f, g : \mathrm{S}^n \to \mathrm{S}^n$ are based homotopic if and only if they have the same degree. The same classification by an integer holds for free homotopy classes of sphere self-maps.

Proof. Degree was defined in Lecture 8, Definition 8.1.6 by $f_*([\mathrm{S}^n]) = \deg(f)[\mathrm{S}^n]$ in $\mathrm{H}_n(\mathrm{S}^n;\mathbb{Z})$. Thus the preceding Hurewicz isomorphism sends $[f]$ to $\deg(f)$, giving the based classification and realizing every integer. For $n \ge 2$, the target sphere is simply connected by Lemma 9.5.1. Proposition 9.3.4 therefore identifies based and free homotopy classes, proving the second assertion. □

These results prepare the exact-sequence methods of Lecture 10. A fibration will compare the homotopy groups of its fiber, total space, and base, with basepoints and fundamental-group actions controlling the low-degree end. In the Hopf fibration, the group $\pi_3(\mathrm{S}^3) \cong \mathbb{Z}$ just

established becomes the source of a computation of $\pi_3(S^2)$, a group that cannot be read from $H_3(S^2;\mathbb{Z})$.

9.5.3 References and Scope

The core constructions and proofs above are self-contained, using homotopy invariance of singular homology and sphere homology from earlier lectures where indicated. The two major comparison theorems are deliberately statements only. No proof of Whitehead or Hurewicz is included, and no relative Hurewicz theorem, cellular approximation theorem, or covering-space classification is assumed.

For further reading, consult the following author-hosted texts. The page numbers are the printed page numbers; the links open the corresponding PDF locations.

1. J. P. May, *A Concise Course in Algebraic Topology*: Chapter 9, Sections 1 and 5, pp. 65 and 68–69, for homotopy groups and basepoint change; Chapter 10, Section 3, pp. 75–76, for Whitehead; and Chapter 15, Section 1, pp. 117–118, for the Hurewicz map and both forms of the theorem. The collar proofs here do not require the fibration machinery used in May's basepoint discussion.
2. Allen Hatcher, *Algebraic Topology*, provides supplementary statements in Theorem 4.5, p. 346, for Whitehead, Theorem 4.32, pp. 366–367, for Hurewicz, and Section 2.A, pp. 166–168, for abelianization in degree one. The conventions and proofs in this lecture follow the course's own map-based development.