Exercises in Algebraic Topology

Zack M. Davis

August–December 2025

Abstract

Exercises mostly from crashing Prof. Emily Clader's "Advanced Topics in Mathematics: Algebraic Topology" at San Francisco State University. Other books on hand as exercise sources: Theodore W. Gamelin and Robert Everist Greene's *Introduction to Topology* (Dover Books paperbook), James Munkres's *Topology* (borrowed from the SFSU library), and James Munkres's *Elements of Algebraic Topology* and Kammeyer's *Introduction to Algebraic Topology* (PDFs downloaded using SFSU library privileges).

Theorem. Let X be a convex subset of $\mathbb{R}^n$, let $a, b \in X$, and let g_0 and g_1 be paths in X between a and b. Then g_0 and g_1 are homotopic with endpoints fixed via the "straight-line homotopy" $F(s,t) := (1-t)g_0(s) + tg_1(s)$.

Proof. First note that the codomain of the given definition of F can indeed be taken to be X, because the fact that X is convex, means that $(1-t)g_0(s) + tg_1(s)$ will be in X if $g_0(s)$ and $g_1(s)$ are.

We then check the boundary conditions.

$F(s,0) = (1-0)g_0(s) + \cancel{0g_1(s)} = g_0(s)$ ✓

$F(s,1) = \cancel{(1-1)g_0(s)} + 1g_1(s) = g_1(s)$ ✓

$F(0,t) = (1-t)g_0(0) + tg_1(0) = (1-t)a + ta = a$ ✓

$F(1,t) = (1-t)g_0(1) + tg_1(1) = (1-t)b + tb = b$ ✓

We must then confirm that F is continuous. Fix ε.

Because g_0 and g_1 are continuous, there exists δ_0 and δ_1 such that $|s - s'| < \delta_0$ implies $|g_0(s) - g_0(s')| < \frac{\varepsilon}{4}$ and $|s - s'| < \delta_1$ implies $|g_1(s) - g_1(s')| < \frac{\varepsilon}{4}$.

Let $M := \max_{s \in [0,1]} (||g_0(s)|| + ||g_1(s)||)$.

Let $\delta := \min(\delta_0, \delta_1, \frac{\varepsilon}{2M})$.

Suppose that $||(s,t) - (s',t')|| < \delta$. Then

$$\|F(s,t) - F(s',t')\| = \|(1-t)g_0(s) + tg_1(s) - (1-t')g_0(s') - t'g_1(s')\|$$

$$\leq \|(1-t)g_0(s) - (1-t')g_0(s')\| + \|tg_1(s) - t'g_1(s')\|$$

$$= \left\| (1-t)g_0(s) - (1-t')g_0(s') + \underbrace{(1-t)g_0(s') - (1-t)g_0(s')}_{\text{fancy zero}} \right\| + \left\| tg_1(s) - t'g_1(s') + \underbrace{tg_1('s) - tg_1(s')}_{\text{fancy zero}} \right\|$$

$$\leq \|(1-t)g_0(s) - (1-t)g_0(s')\| + \|(1-t)g_0(s') - (1-t')g_0(s')\| + \|tg_1(s) - tg_1(s')\| + \|tg_1('s) - t'g_1(s')\|$$

$$= \|1-t\| \, \|g_0(s) - g_0(s')\| + \|g_0(s')\| \, \left\|\cancel{1} - t \cancel{-1} + t'\right\| + \|t\| \, \|g_1(s) - g_1(s')\| + \|g_1(s')\| \, \|t - t'\|$$

$$< \|1-t\| \frac{\varepsilon}{4} + \|t\| \frac{\varepsilon}{4} + \|t' - t\| \left(\|g_0(s')\| + \|g_1(s')\|\right) \leq \frac{\varepsilon}{2} + ||(s,t) - (s',t')||M \leq \frac{\varepsilon}{2} + \frac{\varepsilon}{2\cancel{M}}\cancel{M} = \varepsilon$$

so F is continuous, which is *quod erat demonstrandum.*

HW1#1. The gluing lemma. We have $A, B \subseteq X$ such that $X = A \cup B$, Y another topological space, $f : A \to Y$ and $g : B \to Y$ such that $f(x) = g(x)$ where $x \in A \cap B$, and continuous $h = \begin{cases} f(x) & x \in A \\ g(x) & x \in B \end{cases}$.

a. *Proposition.* $h^{-1}(C) = f^{-1}(C) \cup g^{-1}(C)$ for $C \subseteq Y$.
Commentary. Seems pretty trivial?
Proof. $h^{-1}(C) = \{x \in X : h(x) \in C\} = \{x \in A : f(x) \in C\} \cup \{x \in B : g(x) \in C\} = f^{-1}(C) \cup g^{-1}(C)$.

b. If C is closed, then $f^{-1}(C)$ is closed, because the inverse image under a continuous function of an open set is open, which by taking complements implies that the inverse image of a closed set is closed.

c. *Proposition.* Preimages of closed sets are closed under h.
Proof. $h^{-1}(C) = f^{-1}(C) \cup g^{-1}(C)$ by part a. Both disjuncts are closed by part b. The finite union of closed sets is closed.

HW1#2. *Proposition.* Path homotopy is an equivalence relation.
Proof. (Reflexive.) A path $f(s) : [0,1] \to X$ is path-homotopic to itself with the trivial homotopy $H(s,t) : [0,1] \times [0,1] \to X = f(s)$. Just ignore the second argument. ✓
(Symmetric.) If f is path-homotopic to g under $H(s,t)$, just do the homotopy in the other direction $H(s, 1-t)$. ✓
(Transitive.) If f is path-homotopic to g and g to h, then you can just stitch the homotopies together: go from f to g at double speed, then g to h at double speed. ✓

HW1#3. We are asked: why use the equivalence relation of path homotopy rather than (full) homotopy when defining the path product. Uh, kind of a weird question? If you want to compose paths, you want the second path to start where the first ends. Right? How could it be otherwise?

HW1#4. Let f be a path from x to y, $\overline{f}$ be a path from y to x, and e_x be the constant path at x.
a. We want a path homotopy between f and $e_x * f$. This is annoying because we have to figure out how to normalize the parameter. Or maybe it's just piecewise?
(An earlier attempt here was insane clown garbage and has been deleted.)
When $t = 0$, we want f. When $t = 1$, we want to do the constant path e_x for half the time, then do f at double speed. For intermediate t, we want to do the constant path e_x for $\frac{t}{2}$ of the time, and then do f in time $1 - \frac{t}{2}$. From Gemini Pro 2.5 tutoring, the key formula is that when we map $[0,1]$ to $[a,b]$, the new parametrization is $\frac{s-a}{b-a}$ (the numerator shifting to the new start time, the denominator normalizing to the new length). Here, we're mapping $[0,1]$ to $[\frac{t}{2}, 1]$, so $a := \frac{t}{2}$ and $b := 1$. (Gemini Pro 2.5 caught an error in an earlier attempt: I wrote $b := 1 - \frac{t}{2}$, mixing up the duration and the endpoint.) So the new parametrization should be $\frac{s-\frac{t}{2}}{1-\frac{t}{2}}$. So we should have

$$H(s,t) := \begin{cases} x & s < \frac{t}{2} \\ f\left(\frac{s-\frac{t}{2}}{1-\frac{t}{2}}\right) & s \geq \frac{t}{2} \end{cases}$$

b. We want a path homotopy between $f * \overline{f}$ and e_x. At $t = 0$, we have $\begin{cases} f(s) & s < \frac{1}{2} \\ f(1-s) & s \geq \frac{1}{2} \end{cases}$. (Uh, modulo speed normalization.) At $t = 1$, we want the constant path at x. The obvious strategy is to continuously "pull back" the "turnaround point" at $f(1)$. If we think of that as mapping $[0,1]$ onto $[0, 1-t]$, that suggests a parametrization of $\frac{s}{1-t}$, yielding $H(s,t) := \begin{cases} f(\frac{s}{1-t}) & s < \frac{1}{2} \\ f(1 - \frac{s}{1-t}) & s \geq \frac{1}{2} \end{cases}$. We check the endpoints: $H(s,0)$ confirms. But $H(s,1)$ is a division by zero. We were wrong: mapping $[0,1]$ onto $[0, 1-t]$ just does the whole route successively faster (and goes discontinuous at "lightspeed"); it doesn't pull back the turnaround point. For example, you could think that at $t := \frac{3}{4}$, we want to only go out to $f(\frac{3}{4})$—except t goes in the other direction ($t = 0$ is the full route, $t = 1$ is the constant path). So the next guess is that we actually want

$$H(s,t) := \begin{cases} f(s(1-t)) & s < \frac{1}{2} \\ f(1 - s(1-t)) & s \geq \frac{1}{2} \end{cases}$$

But this still has a discontinuity: $t = 1$ implies the second branch goes to $f(1)$, and we still haven't done speed normalization.
At this point we turn to Gemini 2.5 Pro again. It explains that the outward leg is $2s(1-t)$ (going out to $1 - t$ at double speed) and the inward leg is going to come from the linear function that maps $\frac{1}{2}$ (second leg start

time) to $1-t$ (second leg start displacement) and 1 (second leg end time) to zero (finish line). That's a slope of $-\frac{1-t}{\frac{1}{2}} = 2t-2$. We need to find the intercept; we know $(1,0)$ is on the line (we've returned to the starting point at time 0). So we have $0 = (2t-2)(1) + b$ implies $b = 2-2t$. So the inbound leg is $(2t-2)s + 2 - 2t$. So our homotopy is

$$H(s,t) := \begin{cases} f(2s(1-t)) & s < \frac{1}{2} \\ f((2t-2)s + 2 - 2t) & s \geq \frac{1}{2} \end{cases}$$

HW2#1. Let $\varphi : X \to Y$ be a continuous map between topological spaces, and $x \in X$. Let $\varphi_*([f]) : \pi_1(X,x) \to \pi_1(Y, \varphi(x)) = [\varphi \circ f]$.

a. *Proposition.* φ_* is well-defined: if $[f] = [g]$, then $[\varphi \circ f] = [\varphi \circ g]$.

Commentary. I'm not quite sure how to begin. Presumably it's not OK to say "just apply φ inside the brackets on both sides of $[f] = [g]$": we want to understand what how equivalence classes interact with maps such that "compose inside the brackets" is allowed. If f can be deformed into g, then why can $\varphi \circ f$ be deformed into $\varphi \circ g$? Gemini 2.5 Pro hints that we should work with a homotopy H between f and g—I shouldn't have stooped to ask.

Proof. Let $H(s,t)$ be a homotopy between f and g: that is, $H(0,t) = x$, $H(1,t) = x$, $H(s,0) = f(s)$, and $H(s,1) = g(s)$.

Now consider $\varphi \circ H$. That's continuous as the composition of continuous functions. We have $\varphi \circ H(0,t) = \varphi(x)$ and $\varphi \circ H(1,t) = \varphi(x)$. But we also have $\varphi \circ H(s,0) = \varphi \circ f(s)$ and $\varphi \circ H(s,1) = \varphi \circ g(s)$. But those are exactly the conditions we need for $\varphi \circ H$ to be a homotopy betwen $\varphi \circ f$ and $\varphi \circ g$, so $[\varphi \circ f] = [\varphi \circ g]$.

b. *Proposition.* φ_* is a group homomorphism.

Commentary. We need to show that $\varphi_*([f] * [g]) = \varphi_*([f]) * \varphi_*([g])$. What does that entail? I tried writing out what I thought was a proof (supposing we have a homotopy F between two members of $[f]$, G between members of $[g]$, and getting the same result from doing φ and the each-at-double-speed operation on the inside/outside), but Claude Opus 4.1 and Gemini 2.5 Pro aren't buying it. What am I misunderstanding?

The each-at-double-speed path product is defined on specific paths, and the product of equivalence classes is the equivalence class of product of representatives: $[f] * [g] = [f * g]$.

Proof. We have $\varphi_*([f] * [g]) = \varphi_*([f * g]) = [\varphi \circ (f * g)] = \varphi \circ \begin{cases} f(2s) & s \leq \frac{1}{2} \\ g(2s-1) & s > \frac{1}{2} \end{cases} = \begin{cases} \varphi \circ f(2s) & s \leq \frac{1}{2} \\ \varphi \circ g(2s-1) & s > \frac{1}{2} \end{cases}$.

But $\varphi_*([f]) * \varphi_*([g]) = [\varphi \circ f] * [\varphi \circ g] = [(\varphi \circ f) * (\varphi \circ g)] = \begin{cases} \varphi \circ f(2s) & s \leq \frac{1}{2} \\ \varphi \circ g(2s-1) & s > \frac{1}{2} \end{cases}$.

HW2#2. *Theorem.* Fundamental groups of the same space but different basepoint are isomorphic.

a. Let $x_0, x_1 \in X$, and $\hat{\alpha}([f]) : \pi_1(X, x_0) \to \pi_1(X, x_1) = [\overline{\alpha} * f * \alpha]$. We want to show that $\hat{\alpha}$ is well-defined: if $[f] = [g]$, then $\hat{\alpha}([f]) = \hat{\alpha}([g])$.

Proof. We have $\hat{\alpha}([f]) = [\overline{\alpha} * f * \alpha]$ and $\hat{\alpha}([g]) = [\overline{\alpha} * g * \alpha]$. Not being sure what operations we're formally "allowed to do inside the brackets", should not be a blocker.

$\overline{\alpha} * (f * \alpha)(s) = \begin{cases} \overline{\alpha}(2s) & s \in [0, \frac{1}{2}] \\ (f * \alpha)(2s-1) & s \in [\frac{1}{2}, 1] \end{cases}$ but $f * \alpha(s) = \begin{cases} f(2s) & s \in [0, \frac{1}{2}] \\ \alpha(2s-1) & s \in [\frac{1}{2}, 1] \end{cases}$, so $\overline{\alpha} * (f * \alpha)(s) = \begin{cases} \overline{\alpha}(2s) & s \in [0, \frac{1}{2}] \\ f(4s-2) & s \in [\frac{1}{2}, \frac{3}{4}] \\ \alpha(4s-3) & s \in [\frac{3}{4}, 1] \end{cases}$

(where we've simplified $2(2s-1) = 4s-2$ and $2(2s-1) - 1 = 4s-3$).

I notice that the path-concat operation itself (as opposed to the equivalence classes) is not associative (becuase in $a * (b * c)$, a takes half the time, but in $(a * b) * c$, a takes a quarter of the time), even though the equivalence classes are.

The same calculation without loss of generality also gives us $\overline{\alpha} * (g * \alpha)(s) = \begin{cases} \overline{\alpha}(2s) & s \in [0, \frac{1}{2}] \\ g(4s-2) & s \in [\frac{1}{2}, \frac{3}{4}] \\ \alpha(4s-3) & s \in [\frac{3}{4}, 1] \end{cases}$.

Suppose we have a homotopy $H(s,t)$ such that $H(s,0) = f$ and $H(s,1) = g$ (and $H(0,t) = x_0$ and $H(1,t) = x_1$). Then it would stand to reason that $\begin{cases} \overline{\alpha}(2s) & s \in [0, \frac{1}{2}] \\ H(4s-2, t) & s \in [\frac{1}{2}, \frac{3}{4}] \\ \alpha(4s-3) & s \in [\frac{3}{4}, 1] \end{cases}$ would be a homotopy between $\overline{\alpha} * (f * \alpha)$ and $\overline{\alpha} * (g * \alpha)$, so we indeed have $\hat{\alpha}([f]) = \hat{\alpha}([g])$.

Commentary. Gemini and Claude are saying we didn't need to do all that work. There's a shorter

Proof. $[f] = [g]$ implies $[\overline{\alpha}] * [f] = [\overline{\alpha}] * [g]$ and $[\overline{\alpha}] * [f] * [\alpha] = [\overline{\alpha}] * [g] * [\alpha]$ by group operations. But by the definition of the operation, $[\overline{\alpha}] * [f] * [\alpha] = [\overline{\alpha} * f] * [\alpha] = [\overline{\alpha} * f * \alpha]$ (and the same applies without loss of generality for $[\overline{\alpha}] * [g] * [\alpha]$. So $[\overline{\alpha} * f * \alpha] = [\overline{\alpha} * g * \alpha]$, which is $\hat{\alpha}([f]) = \hat{\alpha}([g])$.

b. *Proposition.* $\hat{\alpha}$ is a group homomorphism.

Proof. We need to show that $\hat{\alpha}([f * g]) = \hat{\alpha}([f]) * \hat{\alpha}([g])$.

We have $\hat{\alpha}([f * g]) = [\overline{\alpha} * f * g * \alpha]$ and $\hat{\alpha}([f]) * \hat{\alpha}([g]) = [\overline{\alpha} * f * \alpha] * [\overline{\alpha} * g * \alpha] = [\overline{\alpha} * f * \cancel{\alpha * \overline{\alpha}} * g * \alpha] = [\overline{\alpha} * f * g * \alpha]$ ✓

c. *Proposition.* $\hat{\alpha}$ is a bijection.

Proof. (Injectivity.) We need to show that if $\hat{\alpha}([f]) = \hat{\alpha}([g])$, then $[f] = [g]$. Suppose $[\overline{\alpha} * f * \alpha] = [\overline{\alpha} * g * \alpha]$. Then $[\cancel{\overline{\alpha}}] * [f] * [\cancel{\alpha}] = [\cancel{\overline{\alpha}}] * [f] * [\cancel{\alpha}]$ implies $[f] = [g]$. ✓

(Surjectivity.) We need to show that for all $f' \in \pi_1(X, x_1)$, there exists $f \in \pi_1(X, x_0)$ such that $\hat{\alpha}([f]) = [f']$. Consider arbitrary $[f']$. Then let $f := \alpha * f' * \overline{\alpha}$. Then $\hat{\alpha}([f]) = [\overline{\alpha} * f * \alpha] = [\cancel{\overline{\alpha} * \alpha} * f' * \cancel{\overline{\alpha} * \alpha}] = [f']$. ✓

HW2#3. *Theorem.* Let X be simply-connected, $x, y \in X$, and let $f : [0, 1] \to X$, and $g : [0, 1] \to X$ be paths from x to y. Then f is path-homotopic to g.

Comment. This is Lemma 52.3 in Munkres.

Proof. "Simply connected" means that the fundamental group is trivial. We notice that $f * \overline{g}$ is a loop, and path-homotopic to the constant loop at x. $[g] = [f * \overline{g}] * [g] = [f * \cancel{\overline{g} * g}] = [f]$.

HW2#4. *Proposition.* The "complement-of-wedge" shape $r \leq 1$, $\theta \in [0, \frac{3}{2}\pi]$ is simply connected but not convex.

Comment. It's obviously simply connected, but to prove that, it would seem we have to compute the fundamental group. Gemini 2.5 Pro hints that we can invoke star-shapedness. Munkres calls this property "star-convexity", which is a clue that our proof for convex sets might generalize. Since every point on in a ray segment from the center to a point on the loop, is itself inside the loop (that's what star-shapedness means), we can "tighten the loop" that way.

Proof. (Non-convexity.) It doesn't contain the straight line between $(1, 0)$ and $(0, -1)$, which suffices to show it's not convex. ✓

(Simply connected.) Let $f(s)$ be a loop based at the origin inside the complement-of-wedge shape with $f(0) = f(1) = (0, 0)$. Then $H(s, t) = (1 - t)f(s)$ is a homotopy between f and the constant path at 0: we see that $H(s, 0) = (1 - 0)f(s) = f(s)$ ✓, $H(s, 1) = (1 - 1)f(s) = 0$ ✓, $H(0, t) = (1 - t)f(0) = 0$ ✓, and $H(1, t) = (1 - t)f(1) = 0$ ✓, and $H(s, t)$ belongs to the complement-of-wedge by star-shapedness ✓.

HW3#1. *Proposition.* $p(x) : \mathbb{R} \to S^1 = (\cos 2\pi x, \sin 2\pi x)$ is a covering map.

Commentary. The prompt says to find a collection of open sets that cover S^1 and show that each is evenly covered. At first, that didn't seem right: why some specific open cover, don't we need *every* point to be evenly covered? But a specific open cover suffices: every point has one of the pieces of the cover as a neighborhood.

I'm using four because that was the argument suggested in Munkres §53, but Gemini 2.5 Pro confirms we can use two (each the circle with a point removed, the preimages being translates of $\mathbb{R}\backslash\mathbb{Z}$).

Proof. Let $U_1 := \{p(x) : x \in (-\frac{1}{4}, \frac{1}{4})\}$, $U_2 := \{p(x) : x \in (0, \frac{1}{2})\}$, $U_3 := \{p(x) : x \in (\frac{1}{4}, \frac{3}{4})\}$,$U_4 := \{p(x) : x \in (\frac{1}{2}, 1)\}$.

Then we have $p^{-1}(U_1) = \bigcup_{n \in \mathbb{Z}}\{(n - \frac{1}{4}, n + \frac{1}{4})\}$, $p^{-1}(U_2) = \bigcup_{n \in \mathbb{Z}}\{(n, n + \frac{1}{2})\}$, $p^{-1}(U_3) = \bigcup_{n \in \mathbb{Z}}\{(n + \frac{1}{4}, n + \frac{3}{4})\}$, $p^{-1}(U_4) = \bigcup_{n \in \mathbb{Z}}\{(n + \frac{1}{2}, n + 1)\}$.

Every point in S^1 has at least one of the U_i as an open neighbhorhood, and the inverse image of every U_i is the union of disjoint open sets such that the restriction of p to that set is a homeomorphism onto the neighborhood.

HW3#2. **a**. *Proposition.* For covering map $p : E \to B$, let $B_k := \{b \in B : p^{-1}(b)$ has exactly k elements$\}$. For each k, B_k is open in B.

Proof. If B_k is the empty set, then it's open, so suppose it's not empty.

Suppose for a contradiction that B_k is not open. Then there exists a point $b \in B_k$ such that there does not exist a neighborhood of b contained in B_k: i.e., there are points arbitrarily close to b that don't have exactly k preimage points.

By even covering, there exists a neighborhood U of b such that $p^{-1}(U)$ is a disjoint union of sets that are locally homeomorphic to U.

Suppose for a contradiction that the number of slices in $p^{-1}(U)$ is more or less than k, not exactly k.

If more, then we have more than k slices that are homeomorphic to U, which contains b, so b (having a preimage in each slice) would have more than k preimages, despite our assumption that b has exactly k preimages. Contradiction!

If less, then b would need to have other preimages not in $p^{-1}(U)$, but that doesn't make sense: a preimage of U contains the preimages of all the point in U, including b. Contradiction!

Thus, the neighborhood U of b is contained in B_k, despite our assumption that no such neighborhood exists. Contradiction!

b. *Proposition.* B_k is closed.

Proof. B_k is closed iff its complement is open, but the complement $B_k^{\complement} = \bigcup_{j \in \mathbb{N}, j \neq k} B_j$ is a union of open sets, which is open.

c. *Theorem.* If B is connected and some fiber of p has size k, then every fiber of p has size k.

Proof. The only open and closed nontrivial subset of a connected set is the set itself, but we've shown above that B_k is open and closed.

d. Cooking up an arbitrary concrete counterexample for the disconnected-B case is easy: let $B_1 := [0, 1]$ and $B_2 := [2, 3]$ and $B := B_1 \cup B_2$ and then say p maps $[4, 5]$ onto $[0, 1]$ and maps both $[6, 7]$ and $[8, 9]$ onto $[2, 3]$.

HW3#3. *Proposition.* Let $p : E \to B$ be a covering map, and let $b \in B$. The subspace topology on $p^{-1}(b)$ is the discrete topology.

Commentary. Munkres §53 mentions this argument inline. Gemini 2.5 Pro points out the key property is that each slice is locally homeomorphic to U, so for a given preimage slice, only one point maps to b (because it's a bijection).

Proof. By even covering, the point b has a neighborhood U whose inverse image comes in open slices, each of which intersects $p^{-1}(b)$ in a single point.

HW4#1. **b**. Let f be the path in $S^1 \times S^1$ given by $f(s) := ((\cos 2\pi s, \sin 2\pi s), (\cos 4\pi s, \sin 4\pi s))$.

The lift of f starting at $(0, 0)$ would be the line from $(0, 0)$ to $(1, 2)$ in $\mathbb{R}^2$. (We're going around the second circle "twice as fast"—twice during the $[0, 1]$ parameter domain—so the preimage in $\mathbb{R}^2$ goes twice as far.)

Asked to propose another lift, I nominated the line from $(\frac{1}{2}, \frac{1}{2})$ to $(\frac{3}{2}, \frac{5}{2})$, thinking that I could start mapping the circle from π radians, $(-1, 0)$: but that's not the same path as f; that would be a different path on the image (B) torus. Gemini Pro 2.5 explains that lifts of f take the form $\{(m + s, n + 2s) : m, n \in \mathbb{Z}\}$.

HW4#2. *Theorem.* Let $p : E \to B$ be a covering map. If E is path-connected and B is simply-connected, then p is a homeomorphism.

Commentary. Somehow these two conditions must imply that p has only one preimage sheet. I thought that this was basically Munkres's Theorem 54.4: consider the lifting correspondence $\phi : \pi_1(B, b_0) \to p^{-1}(b_0)$ that maps elements of the fundamental group onto the endpoint of its lift based at $e_0 \in E$. If E is path connected, then for an arbitrary $e_1 \in p^{-1}(b_0)$, there's a path $\tilde{f}$ in E from e_0 to e_1. Then $f := p \circ \tilde{f}$ is a loop in B. ... no, this isn't actually Theorem 54.4, which is talking about the consequences of E being both path- and simply-connected; here, it's B that's simply connected.

Proof. Because p is a covering map, every point $b \in B$ belongs to an open neighborhood U whose preimage $p^{-1}(U)$ can be written as a disjoint collection of open sets V_α that are each locally homeomorphic to U. If p is itself a homeomorphism, then the "disjoint collection of open sets" just has the one set, and conversely, if the disjoint collection is just one set, then it's a homeomorphism.

Suppose for a contradiction that p is a k-fold covering for $k > 1$. Then $b \in B$ would have (at least) two preimages in E, call them e_0 and e_1. Because E is path-connected, we could create a loop L containing e_0 and e_1. Because B is simply connected, the image of the loop, $p(L)$, could be tightened to a point, either at $p(e_0)$ or at $p(e_1)$. But without loss of generality, the preimage of a tiny loop at $p(e_1)$ couldn't contain e_0. Contradiction!?

Commentary. I'm sure this is correct in some intuitive sense, but I don't think I've articulated it very rigorously.

Criticism. Gemini 2.5 Pro explains the standard argument. A *path* between e_0 and e_1 (it doesn't have to be a loop) projects down to a loop in B. Because B is simply connected, the loop is homotopic to a constant path (like I said, "could be tightened to a point"). The formal version of "the preimage of a tiny loop at $p(e_1)$ couldn't contain e_0" has to do with the uniqueness of homotopy lifts.

HW4#3. A *retraction* of X onto A is a continuous map $r : X \to A$ such that for all $a \in A$, $r(a) = a$.

a. *Example.* A retraction of $\mathbb{R}^2$ onto the unit ball normalizes all vectors of magnitude greater than 1.

b. *Theorem.* If r is a retraction of X onto A, then for any $a \in A$, the homomorphism $r_* : \pi_1(X, a) \to \pi_1(A, a)$ is surjective.

Commentary. I started out reasoning by contradiction: suppose not. Then there would exist a homotopy equivalence class in $\pi_1(A, a)$, call it $[g]$, for which there does not exist $f \in \pi_1(X, a)$ such that $r_*([f]) = [r \circ f] = [g]$. But then I remember that $A \subset X$ suggests a more direct proof. A has to get mapped to itself, and if $[g] \neq [r \circ g]$, then it's not.

Proof. Let $[f] \in \pi_1(A, a)$. We know that $r|_A$ is the identity on A, so $r_*([f]) = [r \circ f] = [f]$. Every member of $\pi_1(A, a)$ is mapped to itself, therefore every member of $\pi_1(A, a)$ is mapped to by something.

Criticism. This prompt was used on the second assessment. Prof. Clader contends that $[f]$ isn't a valid input to r_*, because f is a "loop in A", not a "loop in X": we need $i \circ f : [0, 1] \to X$ (for the inclusion $i : A \to X$), not $f : [0, 1] \to A$.

c. *Corollary.* There is no retraction of $\mathbb{R}^2$ onto S^1.

Proof. $\pi_1(\mathbb{R}^2) = 0$, which has one element, but $\pi_1(S^1) = \mathbb{Z}$, which has countably infinitely many. The domain is too small for the range to possibly be surjective.

HW4#4. Let $p : S^1 \to S^1 = z^2$ (construing the circle as the unit circle in the complex plane). Then $p_* : \pi_1(S^1) \to \pi_1(S^1)$ "can be viewed as" a homomorphism $h(z) : \mathbb{Z} \to \mathbb{Z} : 2z$, because multiplying complex numbers adds their arguments, so squaring doubles the argument.

Commentary. What's the smart way to say "can be viewed as"? "Isomorphic to a homomorphism" feels clunky (stacking morphisms on top of morphisms), but I'd expect it to be correct: why shouldn't a homomorphism be isomorphic to something like any other mathematical object can?

HW5#1. Let X be the cylinder in $\{(x, y, z)) : x^2 + y^2 = 1\} \subset \mathbb{R}^3$, A be the unit circle, $i : A \to X$ be the inclusion (identity with expanded codomain) map, and $r : X \to A$ be the retraction of the cylinder onto the circle.

a. We're looking for a homotopy relative to $(1, 0, 0)$ between $\mathrm{id} : X \to X$ and $i \circ r : X \to X$.

Commentary. I feel like I don't understand the motivation here. $i \circ r$ is retracting, then fixing up the codomain. If I want a homotopy between the cylinder identity and the retraction-to-the-circle-with-full-codomain, I can presumably use the t parameter to shrink the height of the cylinder. That would be relative to all of S^1, not just $(1, 0, 0)$. Is it relative to $(1, 0, 0)$ just because we need a basepoint to work with fundamental groups?

To shrink the height of the cylinder, I need to map the parameter $[0, 1]$ onto max-height $[\infty, 0]$. There's probably an elegant way to do this, but piecewise would work.

Proof. Let $f(t) : (0, 1] \to \mathbb{R} := \begin{cases} \frac{1}{t} & t < \frac{1}{2} \\ -4t + 4 & t \geq \frac{1}{2} \end{cases}$ and let $g(z, t) := \begin{cases} f(t) & t > 0 \ \& \ z \geq f(t) \\ z & \text{else} \end{cases}$. Consider $H((x, y, z), t) := (x, y, g(z, t))$.

H is a homotopy from the identity on X to the inclusion of the retraction onto the circle: $H((x, y, z), 0) = (x, y, z)$ and $H((x, y, z), 1) = (x, y, 0)$.

Gemini Pro 2.5 points out that we could just linearly scale every point on the cylinder. I guess I found this counterintuitive because it doesn't "converge uniformly" (there are arbitrarily high points on the cylinder than require arbitrarily tiny multipliers to get within ε of 0), but pointwise convergence is fine here.

b. *Proposition.* $i_* : \pi_1(A, a) \to \pi_1(X, a)$ and $r_* : \pi_1(X, a) \to \pi_1(A, a)$ are inverses.

Commentary. Um, i and r are *homotopy* inverses because we have a homotopy between the identity and $i \circ r$. What does that imply about their induced homomorphisms? Munkres Theorem 58.7 says that if f is a homotopy equivalence, then f_* is an isomorphism, which sounds stronger than what we're being asked here, but actually isn't. (If bijective $f : A \to B$ is an isomorphism, it doesn't follow that bijective $g : B \to A$ equals f^{-1}, because there could be more than one isomorphism between A and B.)

Socratic tutoring from Claude Sonnet 4.5 indicates that we need to understand why maps that are homotopic relative to a basepoint yield the same induced homomorphisms. Then we can apply the functoriality of "$_*$". That is, if we know that $\mathrm{id}_X \simeq i \circ r$ rel $(1, 0, 0)$, we can infer that $(\mathrm{id}_X)_* = (i \circ r)_* = i_* \circ r_*$, and if $i_* \circ r_* = \mathrm{id}_{\pi_1(X,(1,0,0))}$, it follows that $i_* = r_*^{-1}$ ... but now Sonnet 4.5 is complaining that that reasoning assumes r_*^{-1} exists (circular!).

Proof. We know that $i \circ r \simeq \mathrm{id}_X$ rel A (from the homotopy in part a). We also know that $r \circ i = \mathrm{id}_A$ (because the inclusion $i : A \to X$ just widens the codomain without changing anything, and the retraction $r : X \to A$ collapses it down). Pick a basepoint $a \in A$.

Maps $f, g : X \to Y$ that are homotopic relative to a basepoint a induce the same homomorphism $\pi_1(X, a) \to \pi_1(Y, a)$: if α is a loop in X based at a (and thus a member of an equivalence class in $\pi_1(X, a)$), then $f \circ \alpha$ and $g \circ \alpha$ are loops based at $f(a) = g(a)$ in Y with a homotopy between them, and thus belong to the same equivalence class in $\pi_1(Y, a)$.

Thus $(i \circ r)_* = (\mathrm{id}_X)_*$ and thus $i_* \circ r_* = \mathrm{id}_{\pi_1(X,a)}$. Meanwhile, $(r \circ i)_* = (\mathrm{id}_A)_*$ implies $r_* \circ i_* = \mathrm{id}_{\pi_1(A,a)}$.

c. *Theorem.* $\pi_1(X, a) \simeq \mathbb{Z}$.

Proof. We know that A is a circle, so we know that $\pi_1(A, a) \simeq \mathbb{Z}$, but i_* (from part b) is a bijection from $\pi_1(A, a)$ to $\pi_1(X, a)$.

HW5#2. Let $X := D \times S^1 = \{(x, y, z, w) : x^2 + y^2 \leq 1 \;\&\; w^2 + z^2 = 1\}$, the "filled-in torus", and let $A := \{(0, 0, z, w) : w^2 + z^2 = 1\}$. Then $\pi_1(X) \simeq \mathbb{Z}$.

Proof. Consider $r(x, y, z, w) : X \to A = (0, 0, z, w)$, and the inclusion $i : A \to X$. Let $H(x, y, z, w, t) : X \to X := (tx, ty, z, w)$. Then $H(x, y, z, w, 1) = \mathrm{id}_X$ and $H(x, y, z, w, 0) = r$, and $H(0, 0, z, w, t) = (0, 0, z, w)$ so H is a homotopy relative to A between id_X and $i \circ r$, and A is a deformation retract of X, so they have the same fundamental group. But A is a circle.

HW5#3. **a**. $\pi_1(\{(x, y) : x^2 + y^2 > 1\})$ is $\mathbb{Z}$ because it can be deformation-retracted onto the circle with $p/\|p\|$.

b. $\pi_1(S^1 \cup \{(x, 0) : x > 0\}$ is $\mathbb{Z}$ because it can be deformation-retracted onto S^1 by "pushing in" the ray to $(1, 0)$.

c. $\pi_1(\mathbb{R}^2 \backslash (\mathbb{R}^+ \times \{0\}))$ is $\{1\}$ because—I was sure there was a map between the plane and the slitted plane, but I didn't remember it. Gemini 2.5 Pro points out that there's a map from the slitted plane to the half plane (take $\sqrt{z}$, which halves the argument), which is simply connected.

HW5#4. *Theorem.* If A is a deformation retract of X, then there exists a homotopy equivalence between A and X.

Proof. If A is a deformation retract of X, that means that there's a homotopy between the identity on X and a retraction onto A: some $H(x, t)$ such that for all x, $H(x, 0) = x$, $H(x, 1) \in A$, and for all $a \in A$ and all $t \in [0, 1]$, $H(a, t) = a$.

A homotopy equivalence between A and X would mean there would exist functions $f : X \to A$ and $g : A \to X$ such that $f \circ g \simeq \mathrm{id}_A$ and $g \circ f \simeq \mathrm{id}_X$. Let f be the retraction of X onto A (call it r instead of f) and let g be the inclusion map of A in X (call it i). Then $r \circ i = \mathrm{id}_A$ (they're actually equal) and $i \circ r \simeq \mathrm{id}_X$ (they're homotopic by virtue of the deformation-retraction homotopy).

§58#1. *Theorem.* If A is a deformation retract of X and B is a deformation retract of A, then B is a deformation retract of X.

Proof. Let $H_1 : X \times [0, 1] \to X$ be such that for all x, $H_1(x, 0) = x$ and $H_1(x, 1) \in A$, and for all $a \in A$ and for all t, $H_1(a, t) = a$.

H_1 is a homotopy between the identity on X, and the composition of the inclusion map of A in X with a retraction of X onto A. It exists because we know that A is a deformation retract of X.

Similarly, let $H_2 : A \times [0, 1] \to A$ be such that for all $a \in A$, $H_2(a, 0) = a$ and $H_2(a, 1) \in B$, and for all $b \in B$ and for all t, $H_2(b, t) = b$. It exists because we know that B is a deformation retract of A.

The elementary "one at double-speed, then the other" technique should give us a homotopy between the identity on X, and the composition of the inclusion map of B in X with a retraction of X onto B. We now confirm this.

Let $H(x, t) := \begin{cases} H_1(x, 2t) & t \leq \frac{1}{2} \\ H_2(H_1(x, 1), 2t - 1) & t > \frac{1}{2} \end{cases}$.

($\forall x\; H(x, 0) = x$) $H(x, 0) = H_1(x, 0) = x$ ✓

($\forall x\; H(x, 1) \in B$) $H(x, 1) = H_2(H_1(x, 1), 1)$. We know that $H_1(x, 1) \in A$. But then the fact that for all $a \in A$, $H_2(a, 1) \in B$ implies that $H_2(H_1(x, 1), 1) \in B$. ✓

($\forall b \in B\, \forall t\; H(b, t) = b$) Suppose $t \leq \frac{1}{2}$. Then $H(b, t) = H_1(b, 2t)$. We know that $B \subseteq A$, because B is a deformation retract of A, so $b \in A$ and we have $H_1(b, t) = b$. On the other hand, suppose $t > \frac{1}{2}$. Then $H(b, t) = H_2(H_1(b, 1), 2t - 1)$. We know that $H_1(b, 1) = b$ and we know that $H_2(b, 2t - 1) = b$. ✓

Commentary. Writing out the proof demands attention to detail which is not demanded by "yeah, the double-speed trick, I get it": the nested call to H_1 in the $t > \frac{1}{2}$ case was not intuitively obvious.

§58#3. *Theorem.* Homotopy equivalence is an equivalence relation.

Reminder. If $f : X \to Y$ and $g : Y \to X$ are continuous maps such that $g \circ f : X \to X$ is homotopic to id_X and $f \circ g : Y \to Y$ is homotopic to id_Y, then f and g are homotopy equivalences between X and Y.

Proof. (Reflexive.) We want to show that X is homotopy equivalent to X. Let $f = g = \mathrm{id}_X$. The identity is continuous, and $g \circ f = \mathrm{id}_X \simeq \mathrm{id}_X$, and $f \circ g = \mathrm{id}_X \simeq \mathrm{id}_X$. ✓

(Symmetric.) We want to show that if X is homotopy equivalent to Y, then Y is homotopy equivalent to X. But this follows from the definition itself being symmetric: you can interchange the roles of X and Y and just use the same maps. ✓

(Transitive.) We want to show that if X is homotopy equivalent to Y, and Y is homotopy equivalent to Z, then X is homotopy equivalent to Z.

Suppose we have $f : X \to Y$ and $f' : Y \to X$ such that $f' \circ f : X \to X$ is homotopic to id_X and $f \circ f' : Y \to Y$ is homotopic to id_Y, and $g : Y \to Z$ and $g' : Z \to Y$ such that $g' \circ g : Y \to Y$ is homotopic to id_Y and $g \circ g' : Z \to Z$ is homotopic to id_Z.

Consider $g \circ f : X$ ~~$\to Y$~~ $\to Z$ *vis-à-vis* $f' \circ g' : Z$ ~~$\to Y$~~ $\to X$. We have $(g \circ f) \circ (f' \circ g') = g \circ f \circ f' \circ g' \simeq g \circ \mathrm{id}_Y \circ g' = g \circ g' \simeq \mathrm{id}_Z$ and similarly $(f' \circ g') \circ (g \circ f) = f' \circ g' \circ g \circ f \simeq f' \circ \mathrm{id}_Y \circ f = f' \circ f \simeq \mathrm{id}_X$. ✓

§58#6. *Proposition.* The retract of a contractible space is contractible.

Proof. Let X be contractible. That means there exists a c such that there exists a homotopy $G(x,t) : X \times [0,1] \to X$ such that for all x, $G(x,0) = x$ and $G(x,1) = c$.

Let A be a deformation retract of X. That means there exists a homotopy $H(x,t) : X \times [0,1] \to X$ such that for all x, $H(x,0) = x$ and $H(x,1) \in A$, and for all $a \in A$ and $t \in [0,1]$, $H(a,t) = a$.

We want to show that there exists a c such that there is a homotopy $K(a,t) : A \times [0,1] \to A$ such that for all a, $K(a,0) = a$ and $K(a,1) = c$. But $K := G|_A$ suffices: for all $a \in A \subseteq X$, we have $G(a,0) = a$ and $G(a,1) = c$.

§58#9. Defining the *degree* of a map $S^1 \to S^1$. Let $b_0 := (1,0)$. We choose a generator γ for the infinite cyclic group $\pi_1(S^1, b_0)$. We recall that the elements of $\pi_1(S^1, b_0)$ are homotopy equivalence classes of loops in S^1 based at b_0. For arbitrary $x_0 \in S^1$, we choose a path α from b_0 to x_0. We define $\gamma(x_0) := \hat{\alpha}(\gamma)$. (It looks like this is overloading γ? It's a loop-equivalence class, and a function that can take an argument?) Recall the meaning of the hat: being that α is a path from b_0 to x_0, $\hat{\alpha}$ is a function from $\pi_1(S^1, b_0)$ to $\pi_1(S^1, x_0)$ given by $\hat{\alpha}(f) := [\overline{\alpha}] * [f] * [\alpha]$. We had $\gamma \in \pi_1(S^1, b_0)$, so $\gamma(x_0) := \hat{\alpha}(\gamma) \in \pi_1(S^1, x_0)$.

For $h : S^1 \to S^1$, we choose $x_0 \in S^1$ and let $x_1 := h(x_0)$. Then we have the induced homomorphism $h_* : \pi_1(S^1, x_0) \to \pi_1(S^1, x_1)$. Both groups are infinite cyclic, so we have $h_*(\gamma(x_0)) = d \cdot \gamma(x_1)$ for some d. Because ... that's the only form a homomorphism $\mathbb{Z} \to \mathbb{Z}$ can take? Or rather, because if $\gamma(x_1)$ generates $\pi_1(S^1, x_1)$, then all elements of $\pi_1(S^1, x_1)$ have the form $d \cdot \gamma(x_1)$.

Then d is the *degree* of h.

That is: h is a map from the circle to itself. It induces a homomorphism between fundamental groups based at x_0 and at $x_1 = h(x_0)$.

a. *Proposition.* d is indpendent of the choice of x_0.

Proof. Consider some other point y_0. In addition to the path α from b_0 to x_0, we also have the path β from x_0 to y_0. Then we presumably define $\gamma(y_0) := (\widehat{\alpha\beta})(\gamma)$.

That is, $(\widehat{\alpha\beta})(f) := [\overline{\alpha\beta}] * [f] * [\alpha\beta]$. The path $\alpha\beta$ goes from b_0 to y_0 (via x_0). $(\widehat{\alpha\beta})(\gamma)$ is a generator of $\pi_1(S^1, y_0)$ because it's going from y_0 to b_0, then doing the path of a generator of $\pi_1(S^1, b_0)$, then going from b_0 back to y_0 (closing the loop at y_0).

Let $y_1 := h(y_0)$. We have the induced homomorphism $h_* : \pi_1(S^1, y_0) \to \pi_1(S^1, y_1)$.

Somehow we need to show that in the equations $h_*(\gamma(x_0)) = d \cdot \gamma(x_1)$ and $h_*(\gamma(y_0)) = e \cdot \gamma(y_1)$, we must have $d = e$.

But $\gamma(y_0) = \widehat{\alpha\beta}(\gamma)$, whereas $\gamma(x_0) = \hat{\alpha}(\gamma)$. So that's $h_*(\hat{\alpha}(\gamma)) = d \cdot \gamma(h(x_0))$ and $h_*(\widehat{\alpha\beta}(\gamma)) = e \cdot \gamma(h(y_0))$. ...?

§69#1. Consider the symmetric group with three elements S_3. A transposition generates a subgroup of order 2. The elements (1 2) and (2 3) generate the group, but the free product $(1\ 2) * (2\ 3) \neq S_3$. Both (1 2)(2 3)(1 2) and (2 3)(1 2)(2 3) are reduced words, but both collapse to (1 3).

§69#2. **a**. *Proposition.* The free product of nontrivial groups is not abelian.

Proof. The free product of groups G_α is defined as a group G such that for all $g \in G$, there is only one reduced word in the groups G_α that represents g. Suppose g were abelian. Then $g_1 g_2 = g_2 g_1$. But $g_1 g_2$ and $g_2 g_1$ are different words; we just said every element in the free product only has one reduced word. Contradiction!

b. *Proposition.* In the free product of two nontrivial groups, a reduced word of even length at least two does not have finite order.

Proof. $(ab)^n = \underbrace{abab...ab}_{n \text{ times}}$. For every value of n, we get a word of length $2n$ which does not reduce further: $(ab)^n \neq 1$ for $n > 0$.

Proposition. In the free product of two nontrivial groups, a reduced word of odd length at least three is conjugate to an element of shorter length.

Proof. Without loss of generality, consider aba. Then conjugate by a: $aab\cancel{aa^{-1}} = a^2b$.

c. *Proposition.* The only elements of a free product of two nontrivial groups that have finite order are the elements of the component groups that have finite order and their conjugates.

Commentary. What makes this subtle is "proving the negative" about non-conjugates.

Proof. Consider the free product $A * B$. Without loss of generality, if $a \in A$ has finite order and $b \in B$, the conjugate bab^{-1} has order $|a|$, because $(bab^{-1})^{|a|} = \underbrace{ba\cancel{b^{-1}b}ab^{-1}...ba\cancel{b^{-1}b}ab^{-1}}_{|a| \text{ times}} = b\cancel{a^{|a|}}b^{-1} = \cancel{bb^{-1}} = 1$.

Suppose for a contradiction that $c \in A * B$ is not a conjugate of an element of one of the component groups with finite order. In order for $c^n = 1$ for some $n > 0$, we need all the "letters" in the reduced word to cancel.

Let's break down to cases.

(If c is length 1.) If c has finite order, then it's an element of finite order of one of the component groups; if not, then not.

(If c is length 2.) Length-two words with both letters from the same group like a_1a_2 or b_1b_2 reduce to a single group element. Length-two words with a letter from each group don't have finite order: $(ab)^n = abab...ababab$: the end of one copy of ab doesn't cancel with the start of the next.

(If c is length ≥ 3.) In order to not be a conjugate, the first letter and last letter of c can't be inverses of each other. If $c := c_1xc_k$, then $c^2 = c_1xc_kc_1xc_k$. The c_kc_1 isn't going to cancel all the way down to 1, because c_1 and c_k aren't inverses. (It could cancel down to a single non-identity letter if c_1 and c_k are in the same component group.)

§69#3. *Proposition.* For $G = G_1 * G_2$ and $c \in G$, $cG_1c^{-1} \cap G_2 = \{1\}$.

Proof. If $c \in G_2$ and $x \in G_1$, either cxc^{-1} doesn't reduce (adjacent elements from different component groups), or is 1 if x is 1 (because $c1c^{-1} = cc^{-1} = 1$), or is in G_1 if c is 1 (because $1x1^{-1} = x$). If $c \in G_1$, then cxc^{-1} is an element of G_1.

HW6#1.b. *Theorem.* If A and B are deformation retracts of the same space X, then A is homotopy equivalent to B.

Proof. If A and B are deformation retracts of the same space X, then A is homotopy equivalent to X and B is homotopy equivalent to X (from HW5#4). But homotopy equivalence is an equivalence relation (§58#3). By transitivity, if $A \simeq X$ and $X \simeq B$, then $A \simeq B$.

HW6#3.a. *Theorem.* X is contractible (homotopy equivalent to a single point) iff the identity map $\mathrm{id} : X \to X$ is homotopic to a constant map.

Proof. ($\Rightarrow$) If X is homotopy equivalent to a single point c, then there exist maps $f : X \to c$ and $g : c \to X$ such that $f \circ g \simeq \mathrm{id}_c$ and $g \circ f \simeq \mathrm{id}_X$. But notice that $f : X \to c$ must be constant, because there's only one point in the codomain. If f is constant, then $g \circ f$ is constant, but we know that $g \circ f$ is homotopic to the identity, giving us our desired homotopy between the identity and a constant map.

($\Leftarrow$) Suppose we have a homotopy $H(x,t)$ between the identity and a constant map: $H(x,0) = \mathrm{id}_X$ and $H(x,1) = c$. Let $g : c \to X$ be the inclusion map and $f : X \to c = c$. Then $g \circ f : X \to X = c$ equals $H(x,1)$ and is homotopic to id_X via the homotopy H. Meanwhile, $f \circ g : c \to c$ is trivially equal and therefore homotopic to id_c.

b. *Theorem.* A topological space X is contractible iff for topological space Y, any two continuous maps $f, g : Y \to X$ are homotopic.

Proof. ($\Rightarrow$) Suppose X is contractible: homotopy equivalent to a single point c. Then if arbitrary $f, g : Y \to c$ are homotopic, it will follow that $f, g : Y \to X$ are homotopic, because homotopy equivalence is an equivalence relation. But $f, g : Y \to c$ are trivially homotopic, because there's only one map $Y \to c$. (Every element of Y gets mapped to c, because c is the only point in the codomain.)

($\Leftarrow$) *Attempt.* Suppose arbitrary $f, g : Y \to X$ are homotopic: we have a homotopy $H : Y \times [0,1] \to X$ such that $H(y,0) = f(y)$ and $H(y,1) = g(y)$. In particular, we can let f be surjective, and let g be constant. (This direction seemed promising to me, but I didn't see how to convert it: we would need a way to turn a map $Y \to X$ into a map $X \to X$ that's continuous with t; you can't "just use the image" of the map from Y.)

Hint. Claude Sonnet 4.5 points out that Y can be any topological space, even a convenient one. Kind of a spoiler.

Proof (continued). Let $Y := X$. If any two continuous maps $f, g : X \to X$ are homotopic, then in particular, that would include the identity and a constant map.

HW6#4. **a**. *Proposition.* Let $A \subseteq X$, and let $j : A \to X$ be the inclusion map. If $f : X \to A$ is a retraction and $j \circ f \simeq \mathrm{id}_X$, then j_* is an isomorphism.

Commentary. $j_* : \pi_1(A) \to \pi_1(X)$ is a homomorphism by definition, so to show that it's an isomorphism, we merely need show that it's injective and surjective.

Inclusion maps are always injective, right? For injectivity, we need to show that if $[a] \neq [b]$ in $\pi_1(A)$, then $j_*([a]) \neq j_*([b])$ in $\pi_1(X)$. Suppose $[a] \neq [b]$. Then $j_*([a]) = [j \circ a]$ and $j_*([b]) = [j \circ b]$. But the inclusion map j is just broadening the domain. If loops a and b aren't homotopic in A, then they're not going to be homotopic in X, either ... right? Actually, that's not obvious: if A is in two disconnected pieces, and $X \backslash A$ fills in the gap and simply-connects the pieces, then a and b would become homotopic in X—so it's not true that inclusion maps are always injective; we need to be using the other conditions.

For injectivity, a hint from Claude Sonnet 4.5 suggests starting with $j_*([\alpha]) = j_*([\beta])$. We have $j_*([\alpha]) = j_*([\beta])$ for $\alpha, \beta \in A$ implies $f_* \circ j_*([\alpha]) = f_* \circ j_*([\beta])$, which amounts to $[f \circ j \circ \alpha] = [f \circ j \circ \beta]$, but that doesn't immediately help us: we already established that "it just broadens the codomain, so it doesn't matter" is bogus reasoning, and while $j \circ f$ would cancel inside the brackets as homotopic to the identity, what we have is $f \circ j$. Sonnet 4.5 in its tutor aspect asks me what a retraction is.

Proof. (Injective.) $j_*([\alpha]) = j_*([\beta])$ implies $f_* \circ j_*([\alpha]) = f_* \circ j_*([\beta])$ implies $[f \circ j \circ \alpha] = [f \circ j \circ \beta]$, but $f \circ j = \mathrm{id}_A$, because for points in A, j acts like the identity, and the retraction f also acts like the identity on A, so we have $[\alpha] = [\beta]$.

(Surjective.) We need to show that for all $[\gamma] \in \pi_1(X)$, there exists $[\tau] \in \pi_1(A)$ such that $j_*([\tau]) = [\gamma]$. But as an inclusion map, j isn't even defined on $X \backslash A$. So what we're really saying here is that every loop in $X \backslash A$ is homotopic to a loop in A alone. Our $j \circ f \simeq \mathrm{id}_X$ assertion should be buying us that somehow. Applying the "star functor", we have $j_* \circ f_* = \mathrm{id}_{\pi_1(X)}$. That is, for all $[\gamma] \in X$, $j_*(f_*([\gamma])) = [\gamma]$, and therfore $f_*([\gamma])$ is the τ that fulfills our surjectivity criterion.

III.8#2. *Proposition.* If Y is contractible, then $X \times Y$ is homotopically equivalent to X.

Proof. Y being contractible means that it's homotopy-equivalent to a single point, which is to say that there's a homotopy between id_Y and $Y \to c$: a continuous function $H(y, t)$ such that $H(y, 0) = y$ and $H(y, 1) = c$.

We want to find a homotopy equivalence between $X \times Y$ and X, which is to say a continuous $f : X \times Y \to X$ such that there exists continuous $g : X \to X \times Y$ such that $f \circ g \simeq \mathrm{id}_X$ and $g \circ f \simeq \mathrm{id}_{X \times Y}$.

Consider $f(x, y) = x$ and $g(x) = (x, c)$. We see that $f \circ g = f(g(x)) = x$, so $f \circ g = \mathrm{id}_X$. Then $g \circ f = g(f(x, y)) = (x, c)$. But then $G(x, y, t) = (x, H(y, t))$ is a homotopy between $\mathrm{id}_{X \times Y}$ and $g \circ f$, because $G(x, y, 0) = (x, y)$ and $G(x, y, 1) = (x, c)$, so $\mathrm{id}_{X \times Y} \simeq g \circ f$.

III.8#5. *Proposition.* $\mathbb{R}^{n+1} \backslash \{0\}$ is homotopically equivalent to S^n.

Proof. Consider $f(v) : \mathbb{R}^{n+1} \backslash \{0\} \to S^n = \frac{v}{\|v\|}$, and the inclusion map $i : S^n \to \mathbb{R}^{n+1} \backslash \{0\}$. We see that $f \circ i = \mathrm{id}_{S^n}$. We also have $i \circ f \simeq \mathrm{id}_{\mathbb{R}^{n+1} \backslash \{0\}}$ via a homotopy that looks like f but which uses t to vary the "scaling factor" between 1 and $\frac{1}{\|v\|}$.

HW7#1. Let $F_2 := \langle a, b \rangle$ and $F_3 := \langle c, d, e \rangle$ be free groups. We consider the free product $G := F_2 *_{F_3} \{1\}$ amalgamated over the homomorphism $\varphi_1(c) = a^4$, $\varphi_1(d) = b^2$, $\varphi_1(e) = (ab)^2$, and $\varphi_2 = 1$.

a. $F_2 * \{1\}$ with no amalgamation is just the free group on two generators: $\langle a, b : \varnothing \rangle$. (The "$\varnothing$" is the absence of relations.) To amalgamate over the given homomorphisms from F_3 amounts to adding relations for all $f \in F_3$, $\varphi_1(f) = \varphi_2(f)$. So we have $\langle a, b : a^4 = 1,\ b^2 = 1,\ (ab)^2 = 1 \rangle$.

b. Dummit and Foote give a presentation of D_4 as $\langle s, r : s^2 = r^4 = 1,\ rs = sr^{-1} \rangle$, which is the same thing in different symbols. If $(rs)^2 = 1$, then $rsrs = 1$ and $rs = s^{-1}r^{-1}$, but $s^2 = 1$ implies that $s = s^{-1}$, so we have $rs = sr^{-1}$.

HW7#2. Now we consider the free product $F_2 *_{F_3} \{1\}$ amalgamated over the homomorphism $\varphi_1(c) = a^2$, $\varphi_1(d) = b^2$, and $\varphi_1(e) = aba^{-1}b^{-1}$, and $\varphi_2 = 1$. That's $\langle a, b : a^2 = b^2 = 1,\ aba^{-1}b^{-1} = 1 \rangle$. But $aba^{-1}b^{-1} = 1$ just amounts to a commutativity assertion $ab = ba$. So the elements of the group is $\{1, a, b, ab\}$. It's isomorphic to the

Klein 4-group.

HW7#4. **a**. *Theorem.* $\pi_1(\mathbb{RP}^2) \simeq \mathbb{Z}/2\mathbb{Z}$.

Commentary. This is Munkres Corollary 60.4. I'm finding this hard to visualize. Loops that stay in the northern hemisphere can be collapsed to a point no matter how many times they go around: that's $n0 = 0$ and is fine. But why is it the case that a north–south loop is different, but going north–south twice is the same? Actually, it's easy: go north-to-south twice along different arcs, which you can imagine forming a lens shape. But the lens is just a loop which can be tightened.

Proof. Let $p : S^2 \to \mathbb{RP}^2$ map the sphere to the northern hemisphere antipodally. By Munkres 54.4, the path-connected codomain and simply connected domain means that the lifting correspondence is a bijection. So the two preimages of a neighborhood (in the north and south hemispheres of S^2) imply that $|\pi_1(\mathbb{RP}^2)| = 2$, and there's only the one group of order 2.

b. *Commentary.* The argument that the fundamental group of the punctured torus is $\mathbb{Z} * \mathbb{Z}$ (identify the corners; then the edges a and b are the two loops of a figure-eight) seems like it shold generalize to the projective plane; the "mirrored" edges having opposite orientations doesn't seem like it should change the argument? But if that were true, then SvK would give the projective plane the same fundamental group as the torus (the rest of the argument carrying as well), $\mathbb{Z} \times \mathbb{Z}$, which we know isn't true.

This was discussed at Prof. Clader's office hours on 20 October.

Proof (in the style of Seifert–van Kampen). Let U be an open "cap" of the northern hemisphere with intersects V (the bottom part of the hemisphere) in an annulus. We have $\pi_1(U) \simeq \{1\}$ and $\pi_1(U \cap V) \simeq \mathbb{Z}$. Then this is the tricky part: $\pi_1(V) \simeq \mathbb{Z}$ (it has the one hole in it left by U's absence), but what matters for the SvK computation is the homomorphism from $\pi_1(U \cap V)$ to $\pi_1(V)$. We can deformation-retract the punctured hemisphere to a circle with antipodal points identified. Going once around the annulus $U \cap V$ amounts to going *twice* around the half-circle at the edge of V.

§71#4. *Proposition.* If X is an infinite wedge of circles, then X is not first-countable.

Commentary. Recall that X is first-countable if for every $x \in X$, there exists a sequence of open neighborhoods of x, $\{U_k\}_{k=1}^{\infty}$, such that each neighborhood of x includes one of the U_k. Because the wedge isn't embedded anywhere in particular, I had gotten the idea that an open arc of one of the circles could be an open neighborhood of the wedgepoint, even though this wouldn't make sense in metric space (analogously to how $[-1, 1]$ isn't open in $\mathbb{R}^2$). Then we could do a Cantor diagonalization on whether the neighborhoods in the sequence intersect $C_i \backslash \{p\}$ or not. Gemini 2.5 Pro says my conception of an open set is wrong, but (despite being told not to spoil the exercise) says that the diagonalization idea can still work. (And then I got the inclusion direction wrong when writing the proof the first time, embarrassing.)

Proof. Suppose for a contradiction that X is first-countable, with a corresponding sequence of open neighborhoods $\{U_k\}_{k=1}^{\infty}$ of the wedgepoint p. Every U_k includes an open arc from each C_i. Then we can do a Cantor diagonalization: if we define V such that for all k, $(V \cap C_k) \subset (U_k \cap C_k)$, then $U_k \not\subseteq V$.

After the second assessment, Prof. Clader's class is moving on from homotopy to homology (even as the books make it clear that there's much more homotopy to do). Gamelin & Greene and big Munkres (*Topology*) only cover homotopy theory. I had a bad feeling looking at Hatcher. My SFSU library privileges (still not revoked) let me download PDFs of Munkres's specialized algebraic topology book (*Elements of* ...) and a 2022 *Introduction to* by Holger Kammeyer which takes a category-theoretic viewpoint.

... I'm liking the looks of little Munkres over Kammeyer.

Proposition. A simplex is a compact, convex set in $\mathbb{R}^n$.

Proof. (Compact.) Compactness follows from being closed and bounded in $\mathbb{R}^n$. ✓

(Convex.) Let $x := \sum_{j=0}^{n} x_j a_j$ and let $y := \sum_j y_j a_j$ where $\sum_j x_j = 1$ and $\sum_j y_j = 1$; that is, x and y are in the simplex. The line segment between them is given by $tx + (1-t)y = \sum_{j=0}^{n} (tx_j + (1-t)y_j)\, a_j$. If $\sum_{j=0}^{n} (tx_j + (1-t)y_j) = 1$, then $tx + (1-t)y$ will also be in the simplex. But $\sum_{j=0}^{n} (tx_j + (1-t)y_j) = \sum_{j=0}^{n} tx_j + \sum_{j=0}^{n} (1-t)y_j = t\sum_{j=0}^{n} x_j + (1-t)\sum_{j=0}^{n} y_j = t(1) + (1-t)(1) = 1$. ✓

HW9#1. *Theorem.* $\{p_k - p_0\}_{k=1}^{n}$ is linearly independent iff for all $\{\lambda_k\}_{k=0}^{n}$, $\sum_{k=0}^{n} \lambda_k p_k = 0$ and $\sum_{k=0}^{n} \lambda_k = 0$ imply that for all k, $\lambda_k = 0$.

Commentary. This doesn't look like it should be hard at all, but it gave me a ton of trouble. In order to reduce the cognitive load of thinking about implications-of-implications, I had tried chunking it as $A \to B \Rightarrow C \,\&\, D \to B$, where B represented "all the coëfficients are zero", but Gemini 2.5 Pro pointed out that that choice was concealing the essence of the problem: "all the coëfficients" *means something different* for the 1-indexed linear-independence-of-vector-differences criterion and the 0-indexed zero-linear-combination-and-zero-sum-of-coëfficients criterion.

Proof. ($\Rightarrow$) Suppose $\{p_k - p_0\}_{k=1}^n$ is linearly independent, which is to say that if $\sum_{k=1}^n \lambda_k(p_k - p_0) = 0$, then for all $k \in \{1...n\}$, $\lambda_k = 0$. Furthermore suppose that $\sum_{k=0}^n \lambda_k p_k = 0$ and $\sum_{k=0}^n \lambda_k = 0$. Crucially, $\sum_{k=0}^n \lambda_k = 0$ implies that $\lambda_0 = -\sum_{k=1}^n \lambda_k$. Thus we can write $0 = \sum_{k=0}^n \lambda_k p_k = \lambda_0 p_0 + \sum_{k=1}^n \lambda_k p_k = (-\sum_{k=1}^n \lambda_k)p_0 + \sum_{k=1}^n \lambda_k p_k = \sum_{k=1}^n \lambda_k p_k - \sum_{k=1}^n \lambda_k p_0 = \sum_{k=1}^n \lambda_k(p_k - p_0)$. Then the linear independence of $\{p_k - p_0\}_{k=1}^n$ guarantees that $\{\lambda_k\}_{k=1}^n = \{0\}$.

($\Leftarrow$) Suppose $\sum_{k=0}^n \lambda_k p_k = 0$ and $\sum_{k=0}^n \lambda_k = 0$ imply that for all $k \in \{0...n\}$ $\lambda_k = 0$, and furthermore suppose that $\sum_{k=1}^n \lambda_k(p_k - p_0) = 0$. [Incomplete, with regrets.]

HW9#3. **a**. The boundary of a 3-simplex is homeomorphic to the 2-sphere.

b. The boundary to two 2-simplices that meet at one vertex is homeomorphic to the figure-eight.

c. A triangular prism (without the "end caps", whose three sides can be formed from two 2-simplices) is homeomorphic to the cylinder $S^1 \times [0, 1]$.

HW9#4. *Proposition.* $\partial_1 \circ \partial_2 : C_2(K) \to C_0(K) = 0$.

Proof. Suppose we have an oriented 2-simplex $[v_0, v_1, v_2]$. The boundary $\partial_2[v_0, v_1, v_2]$ is $[v_1, v_2] - [v_0, v_2] + [v_0, v_1]$.

For an oriented 1-simplex $[v_0, v_1]$, the boundary $\partial_1[v_0, v_1]$ is (by the same procedure, somewhat degeneratelly) $v_1 - v_0$.

Then $\partial_1 \circ \partial_2[v_0, v_1, v_2] = \partial_1([v_1, v_2] - [v_0, v_2] + [v_0, v_1]) = v_2 - v_1 - v_2 + v_0 + v_1 - v_0 = \cancel{v_2 - v_2} \cancel{- v_1 + v_1} + \cancel{v_0 - v_0} = 0$.

Proposition. $\partial_2 \circ \partial_3 : C_3(K) \to C_1(K) = 0$.

Proof. $\partial_3[v_0, v_1, v_2, v_3] = [v_1, v_2, v_3] - [v_0, v_2, v_3] + [v_0, v_1, v_3] - [v_0, v_1, v_2]$.

$\partial_2 \circ \partial_3[v_0, v_1, v_2, v_3] = \partial_2\left([v_1, v_2, v_3] - [v_0, v_2, v_3] + [v_0, v_1, v_3] - [v_0, v_1, v_2]\right)$.

Even just the $3 \cdot 4 = 12$ terms here are annoying enough that I want to reach for Python. (I don't expect it to be faster in absolute terms, just less demoralizing.) After some debugging, we get—

```
class SignedSimplex:
    def __init__(self, vertices, positive=True):
        self.vertices = vertices
        self.positive = positive
    def n(self):
        return len(self.vertices) - 1
    def boundary(self):
        boundary_components = []
        positive = True
        for i in range(len(self.vertices)):
            boundary_components.append(
                SignedSimplex(self.vertices[:i] + self.vertices[i+1:], positive=positive)
            )
            positive = not positive
        return boundary_components
def one_simplex_latex(one_simplices):
    out = ""
    for one_simplex in one_simplices:
        out += "+{}-{}".format(
            one_simplex.vertices[1 if one_simplex.positive else 0],
            one_simplex.vertices[0 if one_simplex.positive else 1]
        )
    return out
if __name__ == "__main__":
    σ = SignedSimplex(["v_{0}", "v_{1}", "v_{2}", "v_{3}"])
    out = ""
    for τ in σ.boundary():
        out += one_simplex_latex(τ.boundary())
```

```
print(out)
```

which yields $+v_3-v_2+v_1-v_3+v_2-v_1+v_3-v_2+v_0-v_3+v_2-v_0+v_3-v_1+v_0-v_3+v_1-v_0+v_2-v_1+v_0-v_2+v_1-v_0$, which is

$$\cancel{+v_3-v_3}\cancel{+v_3+v_3-v_3-v_3}\cancel{-v_2+v_2}\cancel{-v_2+v_2}\cancel{+v_2-v_2}\cancel{+v_1-v_1}\cancel{-v_1+v_1}\cancel{-v_1+v_1}\cancel{+v_0-v_0}\cancel{+v_0-v_0}\cancel{+v_0-v_0}=0$$

✓.

Theorem. $\partial_{n-1}\circ\partial_n : C_n(K)\to C_{n-2}(K)=0$.

Proof. By induction.

(Base.) The $n:=2$ case $\partial_1\circ\partial_2 : C_2(K)\to C_0(K)=0$ is proven above.

(Induction.) Suppose inductively that $\partial_{n-1}\circ\partial_n : C_n(K)\to C_{n-2}(K)=0$. We need to show that $\partial_n\circ\partial_{n+1} : C_{n+1}(K)\to C_{n-1}(K)=0$.

Let c be a member of $C_{n+1}(K)$. Then $\partial_{n+1}c$ is an alternating sum of members of C_n. So $\partial_n\circ\partial_{n+1}c$ [Incomplete, with regrets.]

§2#2. *Proposition.* The star and closed star of a vertex are path-connected.

Proof. Let p be a point in a simplex σ adjacent to vertex v. We have a path from p to v: start with the barycentric coordinates of p, and continuously increase the "weight" corresponding to v and decrease all the others. All of the points on the path except the endpoint are in $\operatorname{Int}\sigma$, and the endpoint v is in the star of itself. By the same construction in reverse, we have a path from v to any point q in another simplex τ adjacent to v. The same strategy clearly works for the closed star.

§2#3. **a**. *Proposition.* The collection of 1-simplices in $\mathbb{R}^2$ having vertices $(0,0)$ and $(1,\frac{1}{k})$ is not locally compact.

Attempt. Consider an neighborhood U of the origin. It contains portions of the countably infinitely many line segments σ_k from $(1,\frac{1}{k})$. Then $\bigcup_{k\in\mathbb{N}}\{U\cap\sigma_k\}$ is an open cover without a finite subcover.

Criticism. Gemini 2.5 Pro isn't having it. [Incomplete, with regrets.]

b. *Theorem.* If K is not locally finite, then $|K|$ is not locally compact.

Proof. Suppose K is not locally finite: there's a vertex that belongs to infinitely many simplices σ_k. An open neighborhood U of the vertex contains part of each σ_k.

Recall that a subset $A\subseteq|K|$ is closed iff $A\cap\sigma_k$ is closed for each k.

Our result will follow immediately if we can show that $U\cap\sigma_k$ is open (because those will be the pieces of our open cover with no finite subcover).

We know that the topology of $|K|$ is coherent with respect to the σ_k, which the text notes is equivalent to U being open iff $U\cap\sigma_k$ is open for each k. But we know that U is open, because we defined it as an open neighborhood! So we're done.

§2#4. The collection of 1-simplices in $\mathbb{R}^2$ having vertices $(0,0)$ and $(1,\frac{1}{k})$ is not metrizable.

Commentary. We are given a hint to show that it's not first-countable. We should be able to use the same diagonalization strategy that we used above for §71#4 from the other Munkres book?

Proof. Suppose for a contradiction that it is too metrizable, thus first-countable, thus there exists a sequence of open neighborhoods $\{U_k\}_{k=1}^{\infty}$ of the origin such that each neighborhood of the origin contains one of the U_k. We can construct an open neighborhood V such that $(V\cap\sigma_k)\subset(U_k\cap\sigma_k)$. Then for all k, $U_k\not\subseteq V$, because V is strictly smaller than U_1 in σ_1, smaller than U_2 in σ_2, *&c.*

§2#5. *Theorem.* Let $g:|K|\to|L|$ map the vertices of σ onto the vertices of τ. Then g maps some face of σ homeomorphically onto τ.

Commentary. Recall that a face of σ is just a simplex spanned by a subset of the vertices of σ. The preimage of the vertices of τ is the vertices of σ. Presumably the reason this isn't trivial is the "homeomorphically" part. The map g can non-injectively "collapse" multiple vertices of σ onto the same vertex of τ.

Proof. Let the vertices of σ be $\{v_i\}_{i=0}^m$ and the vertices of τ be $\{w_j\}_{j=0}^n$. We know that for all j, $w_j=f(v_i)$ for some i.

Pick some subset of $\{v_i\}_{i=0}^m$ such that the map from the subset of the vertices of τ is injective. Then g maps $x=\sum_{j=1}^n a_j v_j\mapsto g(x)=\sum_{j=1}^n b_j w_j$. There exists an invertible affine (thus continuous) map from an affine combination of n vectors to an affine combination of n vectors.

§2#7. *Theorem.* $|K|$ is metrizable iff K is locally finite.

Commentary. A hint points out that $d(x,y) := \sup |t_v(x) - t_v(y)|$ is a metric for each finite subcomplex of K.

Proof. ($\Rightarrow$) §2#4 above (and §71#4 above that from big Munkres) generalizes to prove the contrapositive, that if K is not locally finite, then $|K|$ is not metrizable.

($\Leftarrow$) If we have a metric for each subcomplex (per the hint), then we just need some canonical way to "stitch them together." This was discussed in office hours 27 October: if we can show that between any two points, we only have a finite subcomplex, that should do it.

Recall that a Δ-complex is a collection of maps from standard n-simplices to the space of interest: $\{\sigma_\alpha : \Delta^{\dim(\mathrm{im}(\sigma_\alpha))} \to X\}$. For a simplical complex, each simplex is uniquely determined by its vertices. For a Δ-complex, that's not true: the simplices live in "standard" space, and are getting mapped onto our topological space of interest (which lets us triangulate a torus or whatever with fewer triangles, since they can share vertices in the image space).

Our standard 1-simplex has vertices at $(1,0)$ and $(0,1)$; the 2-simplex at $(1,0,0)$, $(0,1,0)$, and $(0,0,1)$, *&c.*

HW10#2. We define a Δ-complex on the torus, conceptualized as $[0,1] \times [0,1]$ with the edges glued. We need to take identification of faces seriously. (An earlier attempt at defining the relevant Δ-complex defined σ_1^1 through σ_5^1 and led to horrible, soul-corrupting confusion. Thanks to Gemini 2.5 Pro and Claude Sonnet 4.5 for helping to right the ship.)

Let $\sigma^0 : \Delta^0 \to T = (0,0) \sim (1,0) \sim (0,1) \sim (1,1)$.

Let $\sigma_1^1(1-t,t) : \Delta^1 \to T = (t,0) = (t,1)$.

Let $\sigma_2^1(1-t,t) : \Delta^1 \to T = (0,t) = (1,t)$.

Let $\sigma_3^1(1-t,t) : \Delta^1 \to T = (t,t)$.

Let $\sigma_1^2\left(\begin{bmatrix} x \\ y \\ z \end{bmatrix}\right) : \Delta^2 \to T = \begin{bmatrix} 0 & 1 & 1 \\ 0 & 0 & 1 \end{bmatrix}\begin{bmatrix} x \\ y \\ z \end{bmatrix}$.

Let $\sigma_2^2\left(\begin{bmatrix} x \\ y \\ z \end{bmatrix}\right) : \Delta^2 \to T = \begin{bmatrix} 0 & 0 & 1 \\ 0 & 1 & 1 \end{bmatrix}\begin{bmatrix} x \\ y \\ z \end{bmatrix}$.

Proposition. $H_0^\Delta(T) \simeq \mathbb{Z}$

Attempt. We need to compute $\frac{\ker \partial_0}{\mathrm{im}\partial_1}$. The boundary map ∂_0 maps 0-chains to 0, so its kernel is the whole space, and the space is spanned by one generator for the one point, so $\ker \partial_0 \simeq \mathbb{Z}$.

The boundary map ∂_1 maps 1-chains to 0-chains. Given a formal linear combination $\sum_{j=1}^3 a_j \sigma_j^1$, then $\partial_1(\sum_{j=1}^3 a_j \sigma_j^1) = \sum_{j=1}^3 a_j \partial_1 \sigma_j^1$. So then we need to compute the boundary of our σ_j^1. But every edge is actually a loop based at the one vertex, so every boundary is $\sigma^0 - \sigma^0 = 0$.

Denouement. $\frac{\mathbb{Z}}{\{0\}} \simeq \mathbb{Z}$.

Proposition. $H_1^\Delta(T) \simeq \mathbb{Z} \oplus \mathbb{Z}$

Proof. We need to compute $\frac{\ker \partial_1}{\mathrm{im}\partial_2}$. We've seen that $\mathrm{im}\partial_1 \simeq \{0\}$, so the rank–nullity theorem implies that $\ker \partial_1 \simeq \mathbb{Z}^3$.

The generic 2-chain is $a_1\sigma_1^2 + a_2\sigma_2^2$, and its boundary is $a_1\partial_2(\sigma_1^2) + a_2\partial_2(\sigma_2^2)$.

Recall that $\partial_2(\sigma_j^2) = \sigma_j^2 \circ F_0 - \sigma_j^2 \circ F_1 + \sigma_j^2 \circ F_2$, where F_i is the face map $\Delta^{i-1} \to \Delta^i$ (in this case $\Delta^1 \to \Delta^2$). Note that the domains match up: the composition $\Delta^{i-1} \to \Delta^i \to T$ becomes a map $\Delta^{i-1} \to T$, which is the type of an $i-1$ chain. The face maps are indexed from 0 because the ith map is the one that "omits the ith vertex" when, *e.g.*, mapping $[s,t]$ to $[0,s,t]$.

Concretely, the face maps are $F_0\left(\begin{bmatrix} 1-t \\ t \end{bmatrix}\right) : \Delta^1 \to \Delta^2 = \begin{bmatrix} 0 \\ 1-t \\ t \end{bmatrix}$, $F_1\left(\begin{bmatrix} 1-t \\ t \end{bmatrix}\right) : \Delta^1 \to \Delta^2 = \begin{bmatrix} 1-t \\ 0 \\ t \end{bmatrix}$, and $F_2\left(\begin{bmatrix} 1-t \\ t \end{bmatrix}\right) : \Delta^1 \to \Delta^2 = \begin{bmatrix} 0 \\ 1-t \\ t \end{bmatrix}$.

We can compute $\partial_2(\sigma_1^2) = \sigma_1^2 \circ F_0 - \sigma_1^2 \circ F_1 + \sigma_1^2 \circ F_2 = \begin{bmatrix} 0 & 1 & 1 \\ 0 & 0 & 1 \end{bmatrix}\begin{bmatrix} 0 \\ 1-t \\ t \end{bmatrix} - \begin{bmatrix} 0 & 1 & 1 \\ 0 & 0 & 1 \end{bmatrix}\begin{bmatrix} 1-t \\ 0 \\ t \end{bmatrix} + \begin{bmatrix} 0 & 1 & 1 \\ 0 & 0 & 1 \end{bmatrix}\begin{bmatrix} 1-t \\ t \\ 0 \end{bmatrix} = \begin{bmatrix} \cancel{1-t+t} \\ t \end{bmatrix} - \begin{bmatrix} t \\ t \end{bmatrix} + \begin{bmatrix} t \\ 0 \end{bmatrix} = \sigma_2^1 - \sigma_3^1 + \sigma_1^1$.

Likewise, $\partial_2(\sigma_2^2) = \sigma_2^2 \circ F_0 - \sigma_2^2 \circ F_1 + \sigma_2^2 \circ F_2 = \begin{bmatrix} 0 & 0 & 1 \\ 0 & 1 & 1 \end{bmatrix} \begin{bmatrix} 0 \\ 1-t \\ t \end{bmatrix} - \begin{bmatrix} 0 & 0 & 1 \\ 0 & 1 & 1 \end{bmatrix} \begin{bmatrix} 1-t \\ 0 \\ t \end{bmatrix} + \begin{bmatrix} 0 & 0 & 1 \\ 0 & 1 & 1 \end{bmatrix} \begin{bmatrix} 1-t \\ t \\ 0 \end{bmatrix} =$ $\begin{bmatrix} t \\ 1\cancel{-t+t} \end{bmatrix} - \begin{bmatrix} t \\ t \end{bmatrix} + \begin{bmatrix} 0 \\ t \end{bmatrix} = \sigma_1^1 - \sigma_3^1 + \sigma_2^1$.

We see that $\text{im}\partial_2$ has a single generator, $\sigma_1^1 - \sigma_3^1 + \sigma_2^1$, and conclude that $\frac{\ker \partial_1}{\text{im}\partial_2} \simeq \frac{\mathbb{Z} \oplus \mathbb{Z} \oplus \mathbb{Z}}{\langle \sigma_1^1 - \sigma_3^1 + \sigma_2^1 \rangle} \simeq \mathbb{Z} \oplus \mathbb{Z}$.

Proposition. $H_2^\Delta(T) \simeq \mathbb{Z}$

Proof. We need to compute $\frac{\ker \partial_2}{\text{im}\partial_3}$.

We've seen that $\text{im}\partial_2 \simeq \mathbb{Z}$, and there are two 2-simplexes, so the rank–nullity theorem implies that $\ker \partial_2 \simeq \mathbb{Z}$.

But there aren't any 3-simplexes in our space, so the image of ∂_3 (which maps 3-chains to their boundaries) is zero. $\frac{\ker \partial_2}{\text{im}\partial_3} \simeq \frac{\mathbb{Z}}{\{0\}} \simeq \mathbb{Z}$.

HW10#3. We define a Δ-complex on the Klein bottle, conceptualized as $[0, 1] \times [0, 1]$ with the edges glued with (crucially) one pair of edges in reversed orientations. (Both pairs reversed would be $\mathbb{RP}^2$; I misremembered this and had to be corrected by Gemini 2.5 Pro.)

Let $\sigma^0 : \Delta^0 \to T = (0, 0) \sim (1, 0) \sim (0, 1) \sim (1, 1)$.

Let $\sigma_1^1 \left(\begin{bmatrix} 1-t \\ t \end{bmatrix} \right) : \Delta^1 \to T = \begin{bmatrix} t \\ 0 \end{bmatrix} = \begin{bmatrix} 1-t \\ 1 \end{bmatrix}$.

Let $\sigma_2^1 \left(\begin{bmatrix} 1-t \\ t \end{bmatrix} \right) : \Delta^1 \to T = \begin{bmatrix} 0 \\ t \end{bmatrix} = \begin{bmatrix} 1 \\ t \end{bmatrix}$.

Let $\sigma_3^1 \left(\begin{bmatrix} 1-t \\ t \end{bmatrix} \right) : \Delta^1 \to T = \begin{bmatrix} t \\ t \end{bmatrix}$.

Let $\sigma_1^2 \left(\begin{bmatrix} 1-s-t \\ s \\ t \end{bmatrix} \right) : \Delta^2 \to T = \begin{bmatrix} 0 & 1 & 1 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1-s-t \\ s \\ t \end{bmatrix}$.

Let $\sigma_2^2 \left(\begin{bmatrix} 1-s-t \\ s \\ t \end{bmatrix} \right) : \Delta^2 \to T = \begin{bmatrix} 0 & 1 & 0 \\ 0 & 1 & 1 \end{bmatrix} \begin{bmatrix} 1-s-t \\ s \\ t \end{bmatrix}$.

Proposition. $H_0^\Delta(K) \simeq \mathbb{Z}$

Proof. We need to compute $\frac{\ker \partial_0}{\text{im}\partial_1}$. We know that $\ker \partial_0 \simeq C_0^\Delta \simeq \mathbb{Z}$ (one generator for the one vertex). We know that $\text{im}\partial_1 \simeq \{0\}$ (every edge is a loop based at the one vertext). Thus $\frac{\ker \partial_0}{\text{im}\partial_1} \simeq \frac{\mathbb{Z}}{\{0\}} \simeq \mathbb{Z}$.

Proposition. $H_1^\Delta(K) \simeq \mathbb{Z} \oplus \mathbb{Z}/2\mathbb{Z}$

Proof. We need to compute $\frac{\ker \partial_1}{\text{im}\partial_2}$. We've seen that $\text{im}\partial_1 \simeq \{0\}$, so the rank–nullity theorem implies that $\ker \partial_1 \simeq \mathbb{Z}^3$.

Inspecting our diagram, we can guess that the generators of $\text{im}\partial_2$ are going to be $\sigma_1^1 + \sigma_2^1 - \sigma_3^1$ and $\sigma_1^1 - \sigma_2^1 + \sigma_3^1$.

More formally, $\partial_2(\sigma_1^2) = \sigma_1^2 \circ F_0 - \sigma_1^2 \circ F_1 + \sigma_1^2 \circ F_2 = \begin{bmatrix} 0 & 1 & 1 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 0 \\ 1-t \\ t \end{bmatrix} - \begin{bmatrix} 0 & 1 & 1 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1-t \\ 0 \\ t \end{bmatrix} +$ $\begin{bmatrix} 0 & 1 & 1 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1-t \\ t \\ 0 \end{bmatrix} = \sigma_2^1 - \sigma_3^1 + \sigma_1^1$

But $\partial_2(\sigma_2^2) = \sigma_2^2 \circ F_0 - \sigma_2^2 \circ F_1 + \sigma_2^2 \circ F_2 = \begin{bmatrix} 0 & 1 & 0 \\ 0 & 1 & 1 \end{bmatrix} \begin{bmatrix} 0 \\ 1-t \\ t \end{bmatrix} - \begin{bmatrix} 0 & 1 & 0 \\ 0 & 1 & 1 \end{bmatrix} \begin{bmatrix} 1-t \\ 0 \\ t \end{bmatrix} + \begin{bmatrix} 0 & 1 & 0 \\ 0 & 1 & 1 \end{bmatrix} \begin{bmatrix} 1-t \\ t \\ 0 \end{bmatrix} =$ $\begin{bmatrix} 1-t \\ 1\cancel{-t+t} \end{bmatrix} - \begin{bmatrix} 0 \\ t \end{bmatrix} + \begin{bmatrix} t \\ t \end{bmatrix} = \sigma_1^1 - \sigma_2^1 + \sigma_3^1$, as expected.

With two generators for the image, one might think that we have $\frac{\ker \partial_1}{\text{im}\partial_2} \simeq \frac{\mathbb{Z}^3}{\langle \sigma_2^1 - \sigma_3^1 + \sigma_1^1, \sigma_1^1 - \sigma_2^1 + \sigma_3^1 \rangle} \simeq \mathbb{Z}$, but Gemini 2.5 Pro is pointing out that that's wrong (and provides followup tutoring on the right way to do it). The system $\begin{cases} \sigma_1^1 - \sigma_2^1 + \sigma_3^1 = 0 \\ \sigma_1^1 + \sigma_2^1 - \sigma_3^1 = 0 \end{cases}$ implies $\sigma_3 = \sigma_1 + \sigma_2$ (from the second equation) and thus $\sigma_1 \cancel{-\sigma_2} + \sigma_1 \cancel{+\sigma_2} = 2\sigma_1 = 0$ (from substituting into the first). Since σ_3 is redundant, we can informally reason that we have two generators plus the relation $2\sigma_1 = 0$, which amounts to $\mathbb{Z} \oplus \mathbb{Z}/2\mathbb{Z}$.

Proposition. $H_2^{\Delta}(K) \simeq \{0\}$

Proof. We need to compute $\frac{\ker \partial_2}{\mathrm{im}\partial_3}$. We've seen that $\mathrm{im}\partial_2 \simeq \mathbb{Z}^2$, and there are two 2-simplexes, so by the rank–nullity theorem, $\ker \partial_2 \simeq \{0\}$. There are no 3-simplexes, so $\mathrm{im}\partial_3 \simeq \{0\}$, and we have $\frac{\{0\}}{\{0\}} \simeq \{0\}$.

HW10.5#1. *Theorem.* Induced homomorphisms in singular homology are functorial: $(g \circ f)_* = g_* \circ f_*$ and $(\mathrm{id}_X)_* = \mathrm{id}_{H_n(X)}$.

Preliminaries. Recall that for $f : X \to Y$ and $\sigma : \Delta^n \to X$, we define $f_\#(\sigma) : C_n(X) \to C_n(Y) = f \circ \sigma$, and $f_*([\alpha]) : H_n(X) \to H_n(Y) = [f_\#(\alpha)]$. The brackets represent the equivalence relation of being represented by the same element in the homology group: α is an n-chain in X, $f_\#(\alpha)$ is an n-chain in Y, and $[f_\#(\alpha)]$ is in the nth homology group of Y.

Question. This seems to be defined for arbitrary chains, but it's not clear to me how arbitrary chains fit into the homology group. The homology group is $\frac{\ker \partial_n}{\mathrm{im}\partial_{n+1}}$, that is, cycles quotiented by boundaries. What about chains that aren't boundaries?

Answer. Claude Sonnet 4.5 responds that the α in $f_*([\alpha])$ does have to be a cycle, even though $f_\#$ takes arbitrary chains.

Proof. Let $c := \sum_j a_j \sigma_j$ be a chain.

(Composition.) $(g \circ f)_*([c]) = [(g \circ f)_\#(c)] = [(g \circ f)(c)]$, but $g_* \circ f_*([c]) = g_* \circ [f_\#(c)] = [g_\# \circ f_\#(c)] = [g_\# \circ f(c)] = [(g \circ f)(c)]$. ✓

(Identity.) $(\mathrm{id}_X)_*([c]) = [(\mathrm{id}_X)_\#(c)] = [\mathrm{id}_X(c)] = [c]$ ✓

Question. Recall that a quotient group G/K is associated with a homomorphism $\varphi : G \to H$; K is the kernel of φ. The elements of G/K are translates of the kernel K, which functions as the identity. So the very existence of the homology group $H_n = \frac{\ker \partial_n}{\mathrm{im}\partial_{n+1}}$ is implicitly asserting that $\mathrm{im}\,\partial_{n+1}$ must be the kernel of a homomorphism with domain $\ker \partial_n$. What is that?

Answer. Claude Sonnet 4.5 responds that it's just the canonical projection $\ker \partial_n \to \frac{\ker \partial_n}{\mathrm{im}\partial_{n+1}}$. This feels incredibly underwhelming, but I guess that's how it goes sometimes.

HW10.5#2. **a**. The reduced singular homology groups append $\tilde{\partial}_0(\sum_j a_j \sigma_j) = \sum_j a_j$ (for $\sigma_j \in C_0(X)$) to the chain complex.

Theorem? $\tilde{H}_n(X) := \begin{cases} H_n(X) & n > 0 \\ \{0\} & n = 0 \end{cases}$

Attempt. For $n \geq 1$, $\tilde{H}_n = \frac{\ker \tilde{\partial}_n}{\mathrm{im}\tilde{\partial}_{n+1}} = \frac{\ker \partial_n}{\mathrm{im}\partial_{n+1}} = H_n$. For $n = 0$, $\tilde{H}_0 = \frac{\ker \tilde{\partial}_0}{\mathrm{im}\tilde{\partial}_1}$.

For $n = 0$, we reason casually: $\tilde{\partial}_0$ is summing up the coëfficients of a formal linear combination of 0-chains; speaking tautologically and pleonastically, its kernel is all formal linear combinations of 0-chains (points) whose coëfficients sum to zero.

$\tilde{\partial}_1$ takes a formal linear combination of 1-chains (line segments) to the sum of formal differences of their endpoints. But every formal difference $v_1 - v_0$ is a formal linear combination whose coëfficients sum to zero, so $\mathrm{im}\,\tilde{\partial}_1 = \ker \tilde{\partial}_0$ and $\frac{\ker \tilde{\partial}_0}{\mathrm{im}\tilde{\partial}_1} \simeq \{0\}$.

Criticism. Claude Sonnet 4.5 points out that we've only shown that $\mathrm{im}\,\tilde{\partial}_1 \subseteq \ker \tilde{\partial}_0$, and asks us to consider a space of two isolated points. The problem is apparent: $+v_1 - v_2$ has coëfficients that sum to zero, but isn't the image of a 1-chain if there's no 1-simplex connecting v_1 and v_2. We should actually expect $\mathbb{Z}^{r-1}$ when there are r connected components.

We atone for our sins by providing the correct

Theorem. Suppose a topological space X has $r \in \mathbb{N}^+$ path-connected components. Then $\tilde{H}_n(X) := \begin{cases} \mathbb{Z}^{r-1} & n = 0 \\ H_n(X) & \text{else} \end{cases}$.

Proof. For $n \geq 1$, $\tilde{H}_n = \frac{\ker \tilde{\partial}_n}{\mathrm{im}\,\widetilde{\partial_{n+1}}} = \frac{\ker \partial_n}{\mathrm{im}\,\partial_{n+1}} = H_n$.

For $n = 0$, $\tilde{H}_0 = \frac{\ker \tilde{\partial}_0}{\mathrm{im}\,\tilde{\partial}_1}$, which amounts to the set of formal linear combinations of 0-chains whose coëfficients sum to zero which are not the boundary of a 1-chain.

Suppose $\tau \in \ker \tilde{\partial}_0$ but $\tau \notin \mathrm{im}\,\tilde{\partial}_1$ where τ is an otherwise arbitrary 0-chain given by $\tau := \sum_{\alpha \in A_\tau} c_\alpha \sigma_\alpha$ for some index set A_τ. Every σ_α is wholly contained within exactly one path component, so we can write $\tau =$

$\sum_{k=1}^{r}\sum_{\alpha_k\in A_{\tau,k}} c_{\alpha_k}\sigma_{\alpha_k}$ for index sets $A_{\tau,k}$ such that σ_{α_k} belongs to the kth path component (and thus $\bigcup_{k=1}^{r} A_{\tau,k} = A_\tau$).

It is not the case that for all k, $\sum_{\alpha_k\in A_{\tau,k}} c_{\alpha_k} = 0$. Why? Suppose for a contradiction that it were. Then for all k, $\sum_{\alpha_k\in A_{\tau,k}} c_{\alpha_k} = 0$ could be written as $\sum_{p\in A^+_{\tau,k}} c_p + \sum_{n\in A^-_{\tau,k}} c_n$ where c_p are all positive and c_n are all negative, and $\sum_{p\in A^+_{\tau,k}} c_p = \sum_{n\in A^-_{\tau,k}} c_n$. Let $N := \sum_{p\in A^+_{\tau,k}} c_p = \sum_{n\in A^-_{\tau,k}} c_n$. Let $\gamma_{a,b} : \Delta^1 \to X$ be a 1-simplex whose image is a path from the image of σ_a to the image of σ_b. Consider the 1-chain β constructed as follows. Construct lists $C_+ := [(c_p, p)]_{p\in A^+_{\tau,k}}$ and $C_- := [(c_n, n)]_{n\in A^-_{\tau,k}}$. On each of N loop iterations, do the following: let (c_p, p) be the last element of C_+ and let (c_n, n) be the last element of C_-; add $\gamma_{n,p}$ to β; decrement the first tuple-element of the last element of C_+ and C_-; pop the last element from C_+ and C_- if its first tuple-element is 0; end loop iteration. Then we would have $\partial_1(\beta) = \tau$, but we specified that $\tau \notin \operatorname{im}\tilde{\partial}_1$. Contradiction!

Because $\tau \in \ker\tilde{\partial}_0$, we know that $\tau = \sum_{k=1}^{r}\sum_{\alpha_k\in A_{\tau,k}} c_{\alpha_k} = 0$. We can think of this equation as indicating the kernel of a surjective linear map $\mathbb{Z}^r \to \mathbb{Z}$ (where $\sum_{\alpha_k\in A_{\tau,k}} c_{\alpha_k}$ are the coordinates and $A_{\tau,k}$ are the "basis vectors"), which is $r-1$-dimensional; thus, $\tilde{H}_n(X) \simeq \mathbb{Z}^{r-1}$.

HW10.5#3. A loop $f : [0,1] \to X$ induces a singular 1-simplex $\hat{f} := f \circ p$ where $p(1-t, t) : \Delta^1 \to [0,1] = t$. (So $f \circ p$ goes $\Delta^1 \to [0,1] \to X$.) For any $x \in X$, let $h([f]) : \pi_1(X, x) \to H_1(X) = [\hat{f}]$.

a. *Proposition.* h is well-defined: $[f] = [g]$ in $\pi_1(X, x)$ implies $[\hat{f}] = [\hat{g}]$ in $H_1(X)$.

Proof. Because f and g are in the same homotopy class, we have a path homotopy $H : [0,1]\times[0,1] \to X$ between them. Let $p_2(1-s-t, s, t) : \Delta^2 \to [0,1]\times[0,1] = (s,t)$. Then $H \circ p_2$ goes $\Delta^2 \to [0,1]\times[0,1] \to X$ and is a 2-chain whose boundary is $\hat{f} - \hat{g}$.

Criticism. Gemini says this is wrong; we need two simplices to cover the homotopy square.

b. For $[f], [g] \in \pi_1(X, x)$, let $\sigma(1-s-t, s, t) : \Delta^2 \to X = (f * g)(\frac{1}{2}s + t)$.

Proposition. $\partial_1\sigma = \hat{f} + \hat{g} - \widehat{f * g}$.

Commentary. Um, I know how to take the boundary of a simplex-as-orientedvertex-list; I don't think I know how to take the boundary of $(f*g)(\frac{1}{2}s+t)$. But that is supposed to be a map *from* a simplex, the standard 2-simplex. The standard 2-simplex is $\left[\begin{bmatrix}1\\0\\0\end{bmatrix}, \begin{bmatrix}0\\1\\0\end{bmatrix}, \begin{bmatrix}0\\0\\1\end{bmatrix}\right]$ and its boundary is $\begin{bmatrix}0\\0\\1\end{bmatrix} - \begin{bmatrix}0\\1\\0\end{bmatrix} - \begin{bmatrix}0\\0\\1\end{bmatrix} + \begin{bmatrix}1\\0\\0\end{bmatrix} + \begin{bmatrix}0\\1\\0\end{bmatrix} - \begin{bmatrix}1\\0\\0\end{bmatrix}$ ... um, which is zero. [Incomplete, with regrets.]

I told Tristen Gann that I would help on his project about category theory, so I guess I will look at Kammeyer, who takes a category-theoretic approach.

1.1. Let G be a group, and consider the category G-Set, the category of left G-sets, whose objects are functors from G to Set and whose morphisms are natural transformations.

We want to find the left adjoint to the forgetful functor G-Set $\to$ Set. Obviously, the left adjoint will have the domain–codomain signature Set $\to$ G-Set, but what does it actually look like?

Let's make this concrete. Suppose $G := \mathbb{Z}_3$. An object in $\mathbb{Z}_3$-Set is a functor $\mathbb{Z}_3 \to$ Set, which means it takes objects in $\mathbb{Z}_3$ (the group construed as a category with one object) to objects in Set, and morphisms in $\mathbb{Z}_3$ (where the elements of $\mathbb{Z}_3$ are morphisms from the singleton object to itself which compose in a way that respects the group table) to morphisms in Set.

But which specific morphisms in Set? Presumably that's pinned down by the "natural transformation" criterion.

A natural transformation $n : \mathcal{F} \to \mathcal{G}$ between functors $\mathcal{F}, \mathcal{G} : \mathcal{C} \to \mathcal{D}$ is a morphism in the category of functors $\mathcal{C} \to \mathcal{D}$, such that for every A in $\operatorname{obj}(\mathcal{C})$, it assigns a morphism $n_A : \mathcal{F}(A) \to \mathcal{G}(A)$ in $\mathcal{D}$ such that for every morphism $f : A \to B \in \operatorname{Hom}_{\mathcal{C}}(A, B)$, we have $\mathcal{G}(f) \circ n_A = n_B \circ \mathcal{F}(f)$.

In this case, $\mathcal{C} := \mathbb{Z}_3$ and $\mathcal{D} :=$ Set, so for every $\bullet \in \operatorname{obj}(\mathbb{Z}_3)$ (that's the singleton in the group-construed-as-category), there's a morphism $n_\bullet : \mathcal{F}(\bullet) \to \mathcal{G}(\bullet)$ such that for every morphism $f : \bullet \to \bullet \in \operatorname{Hom}_{\mathbb{Z}_3}(\bullet, \bullet)$ (these are the three morphisms representing the elements of $\mathbb{Z}_3$), we have $\mathcal{G}(f) \circ n_\bullet = n_\bullet \circ \mathcal{F}(f)$.

It seems like we just have three different morphisms from $\bullet$ (the singleton member of the group-as-category) to an arbitrary set. Call the morphisms e,g, h, which are distinct because of the rules for how they compose ($g \circ g = h$, $h \circ g = e$, $e \circ e = e$), not because of where they point to or from. This is a disappointingly trivial answer considering how much work we did just to understand the question, but that's how it goes sometimes.

So what does that say about the left adjoint? It takes an arbitrary set, and makes it into a G-Set. Is there supposed to be a less trivial answer than that? At this point, we've earned the right to turn Gemini 3 Pro with the question itself, not just background definition tutoring.

Gemini straightens us out: an object in $\mathbb{Z}_3$-Set is a functor $\mathbb{Z}_3 \to \mathsf{Set}$ that picks out both some specific set X and an action of $\mathbb{Z}_3$ on X. Recall that an action of G on X is a function $G \times X \to X$ that respects the group operations ...

1.2. Let FinBij be the category whose objects are finite sets and whose morphisms are bijections. For a finite set X, let $\mathcal{B}(X)$ be the set of bijections $X \to X$ and let $\mathcal{O}(X)$ be the set of total orders on X.

Proposition. $\mathcal{B}$ is a functor $\mathsf{FinBij} \to \mathsf{Set}$.

Attempt. (Objects to objects.) Consider $X \in \mathrm{obj}(\mathsf{FinBij})$. $\mathcal{B}(X)$ is a *set* of bijections, so $\mathcal{B}(X) \in \mathrm{obj}(\mathsf{Set})$. ✓

(Morphisms to morphisms.) Consider $m : X \to Y \in \mathrm{mor}(\mathsf{FinBij})$. It's an arrow between objects that are finite sets of the same size, representing a bijection of those sets. It can get mapped to the arrow $\mathcal{B}(X) \to \mathcal{B}(Y) \in \mathrm{mor}(\mathsf{Set})$.

(Identity.) For every $X \in \mathrm{obj}(\mathsf{FinBij})$, there's an identity morphism $\mathrm{id}_X \in \mathrm{Hom}_{\mathsf{FinBij}}(X, X)$.

Now $\mathrm{id}_{\mathcal{B}(X)}$ is such that for every $f : \mathcal{B}(X) \to K$ and $g : K \to \mathcal{B}(X)$, $f \circ \mathrm{id}_X = f$ and $\mathrm{id}_X \circ g = g$.

We need to check that the morphisms-to-morphisms rule sends id_X to $\mathrm{id}_{\mathcal{B}(X)}$. But it totally does: an arrow from a finite set X to itself is getting mapped by the functor to an arrow from the set of bijections $X \to X$ to itself.

(Composition.) We need to check that $\mathcal{B}(f \circ g) = \mathcal{B}(f) \circ \mathcal{B}(g)$. The composed bijections $X \xrightarrow{g} Y \xrightarrow{f} Z$ become a bijection $X \to Z$, and similarly $\mathcal{B}(X) \xrightarrow{\mathcal{B}(g)} \mathcal{B}(Y) \xrightarrow{\mathcal{B}(Z)} \mathcal{B}(Z)$ becomes $\mathcal{B}(X) \to \mathcal{B}(Z)$.

Criticism. Gemini 3 Pro thinks I'm being too vague with the morphisms-to-morphisms, as does Claude Opus 4.5 (released today). What we need to do is, given a bijection $m : X \to Y$ and $\sigma : X \to X \in \mathcal{B}(X)$, we need to produce $\mathcal{B}(m)(\sigma) = \tau : Y \to Y$, which is going to turn out to be conjugation. On followup questioning, Gemini points out that the "conjugation" insight is discoverable from the square commutative diagram: the down arrow on the right side of the square needs to be the same as going around "the long way" left–down–right.

I think I'm getting confused about whether to think of morphisms as being like a directed edge in a graph, or like a function with domain and codomain. Gemini clarifies that it's a directed multigraph, but an edge often represents a function. (*E.g.*, in Top,if X and Y are topological spaces, then the morphism $X \to Y$ is a continuous map with that domain and codomain.)

Atonement. (Morphisms to morphisms.) For $X, Y \in \mathrm{obj}(\mathsf{FinBij})$, consider $m : X \to Y \in \mathrm{mor}(\mathsf{FinBij})$. We want to define $\mathcal{B}(m) : \mathcal{B}(X) \to \mathcal{B}(Y)$, a morphism from the set of bijections $X \to X$ to the set of bijections $Y \to Y$, which in the category of Set would be a bijection, so we have to say how an arbitrary particular bijection $X \to X$ gets mapped to a bijection $Y \to Y$.

HW11#2. Let $A_\bullet$, $B_\bullet$, and $C_\bullet$ be chain complexes, and let $i_\bullet : A_\bullet \to B_\bullet$ and $j_\bullet : B_\bullet \to C_\bullet$ be chain maps such that $0 \to A_n \xrightarrow{i_n} B_n \xrightarrow{j_n} C_n \to 0$ is an exact sequence for all n.

Let $\delta([c]) : H_n(C_\bullet) \to H_{n-1}(A_\bullet) = [a]$ where $a \in A_{n-1}$ is such that $i_{n-1}(a) = \partial_n^B(b)$ for some $b \in B_n$ such that $c = j_n(b)$. We're going to show that this is well-defined, step by step.

Commentary. We're going to prove the snake lemma by infering the existence of the "snake map" from the properties of i_{n-1}, ∂^B, and j_n.

a. *Proposition.* There exists $b \in B_n$ such that $c = j_n(b)$.

Commentary. Presumably this is going to be a simple application of the rules for what exact sequences with 0s in them imply about injectivity and surjectivity, which I'm going to need to look up. I'm being told (by Ben Lynn's webpage at *https://crypto.stanford.edu/pbc/notes/commalg/exact.html*) that $0 \to A \xrightarrow{f} B$ is exact iff f is injective and $A \xrightarrow{g} B \to 0$ is exact iff g is surjective.

Proof. The claim that there exists $b \in B_n$ such that $c = j_n(b)$ is the same as saying that j_n is surjective, which must be true if $B_n \xrightarrow{j_n} C_n \to 0$ is a short exact sequence.

b. *Proposition.* For any such b as in the previous proposition, there exists $a \in A_{n-1}$ such that $i_{n-1}(a) = \partial_n^B(b)$.

Commentary. We need to show that there's something in A_{n-1} that hits the same target as $\partial_n^B : B_n \to B_{n-1}$. That would be easy (just like the previous proposition) if i_{n-1} were surjective, but I don't think we know that: $0 \to A_{n-1} \xrightarrow{i_{n-1}} B_{n-1}$ gives us *injectivity*, not surjectivity.

If it's not that, presumably we can work with ∂_n^B not being injective: there exist $n-1$-chains that aren't the boundary of an n-chain. (*E.g.*, the formal sum of three distinctpoints isn't the boundary of any 1-simplex, because the boundary is a formal sum of two distinct points.)

Glancing at *https://en.wikipedia.org/wiki/Snake_lemma#Construction_of_the_maps* gives me a hint that might be sufficient to proceed under my own power. By commutativity of the diagram, $j_{n-1}(\partial^B(b)) = \partial^C(j_n(b))$ ... okay, I'm actually not sure how that helps us; let's read the *Wikipedia* account in more detail.

Wikipedia seems to be claiming (in our notation) that $j_n(b)$ is in the kernel of ∂^C_n, and I don't see why that would have to be true.

Gemini 3 Pro points out that the fact that c is a homology group representative means that it has to be a cycle, and clarifies the last step.

Proof. We know that $j_{n-1}(\partial^B(b)) = \partial^C(j_n(b)) = 0$ (where the first equality follows from the commutativity of the diagram, and the second because $c = j_n(b)$ is a homology class representative, thus a cycle, thus its boundary is zero). But $j_{n-1}(\partial^B(b))$ getting mapped to zero means that $\partial^B(b) \in \ker(j_{n-1})$, so by exactness, $\partial^B(b) \in \mathrm{im}(i_{n-1})$.

c. *Proposition.* The previous propositions don't depend on the choice of a and b.

Attempt. Suppose b and b' are such that $j_n(b) = j_n(b') = c$, and suppose that a and a' are such that $i_{n-1}(a) = \partial^B_n(b)$ and $i_{n-1}(a') = \partial^B_n(b')$. We want to show that $[a] = [a']$ in $H_{n-1}(A_\bullet)$.

If $j_n(b), j_n(b') \in [c]$, then they "differ by a boundary", so $b' = b + d$ where d is a boundary. The boundary of a boundary is zero, so $\partial^B_n(b') = \partial^B_n(b + d) = \partial^B_n(b) + \cancel{\partial^B_n(d)} = \partial^B_n(b)$. So we know that $i_{n-1}(a) = i_{n-1}(a') = \partial^B_n(b)$. But the exactness of the short exact sequence $0 \to A_{n-1} \to B_{n-1}$ implies that i_{n-1} is injective, so $a = a'$.

Criticism and Atonement. Gemini 3 Pro disagrees. (I should have known $a = a'$ was too strong.) The homology equivalence is between $j_n(b)$ and $j_n(b')$ in C_n, not b and b' in B_n, dummy. If $j_n(b) = j_n(b') = c$, then $j_n(b) - j_n(b') = c - c = 0$. Gemini reassures me that j_n is "linear" (my word choice; it's a group homomorphism), such that $j_n(b) - j_n(b') = 0 \Rightarrow j_n(b - b') = 0$.

With more tutoring: so we know that $b - b' \in \ker j_n$, and thus by exactness that $b - b' \in \mathrm{im} i_n$, which is to say there exists $x \in A_n$ such that $b - b' = i_n(x)$. Taking the boundary, we get $\partial^B_n(b) - \partial^B_n(b') = \partial^B_n(i_n(x))$, which we know is $i_{n-1}(a) - i_{n-1}(a') = \partial^B_n(i_n(x))$. By commutativity of the square diagram, $\partial^B_n \circ i_n = i_{n-1} \circ \partial^A_n$, so we have $i_{n-1}(a) - i_{n-1}(a') = i_{n-1}(\partial^A_n(x))$, and thus $a - a' = \partial^A_n(x)$, *viz.*, that a and a' differ by a boundary and are therefore in the same homology class.

HW11#3. We dwell on how one would prove that the long sequence

$$... \to H_n(A_\bullet) \xrightarrow{i_*} H_n(B_\bullet) \xrightarrow{j_*} H_n(C_\bullet) \xrightarrow{\delta} H_{n-1}(A_\bullet) \xrightarrow{i_*} H_{n-1}(B_\bullet) \to ...$$

is exact.

Proposition. $\ker \delta = \mathrm{im} j_*$

Commentary. I was hoping to get through this one with less LLM tutoring, but right off the bat I needed a hint from Gemini that $\exists a' \in A_n \; a = \partial^A_n(a')$ isn't unconditionally true, and then another hint on how to establish $\partial^B_n(b) = \partial^B_n(i_n(a'))$, and then even a reminder that $j_n(b) = c$ and that exactness means double-composition is zero, and then even yet more guidance on part b. I'm going to die.

I also don't think I've been properly distinguishing between chain complexes and homology groups, which is another reason I'm going to die.

Proof. **a**. ($\ker \delta \subseteq \mathrm{im} j_*$) Suppose $[c] \in \ker \delta$. That means $\delta([c]) = [0]$, so if $\partial([c]) = [a]$, then $[a] = [0]$, which means that a is a boundary (because the homology group is cycles quotiented by boundaries).

We next want to show that $b - i_n(a') \in \ker(\partial^B_n)$, which is to say that $\partial^B_n(b - i_n(a')) = \partial^B_n(b) - \partial^B_n(i_n(a')) = 0$, so $\partial^B_n(b) = \partial^B_n(i_n(a'))$.

By diagram chasing, $\partial^B_n(i_n(a')) = i_{n-1}(\partial^A_n(a'))$, which is $i_{n-1}(a)$, but we defined a precisely such that $i_{n-1}(a) = \partial^B_n(b)$.

Then we want to show that $j_*([b - i_n(a')]) = [c]$. By the definition of the induced homomorphism, that amounts to $[j_n(b) - j_n(i_n(a'))] = [c]$. But b was defined such that $j_n(b) = c$, and the idea of "exactness" is a generalization of $\partial^2 = 0$; that is, $j_n \circ i_n = 0$ because the outputs (image) of j_n are the inputs to i_n that get mapped to zero (the kernel). All of which is to say $j_*([b - i_n(a')]) = [j_n(b) - j_n(i_n(a'))] = [c - 0] = c$. That's assuming we know that we have exactness at B_n, but I think we do know that. (The current endeavor is showing the exactness of the long exact sequence; we've long ago assumed for all n the exactness of the short exact sequence $0 \to A_n \to B_n \to C_n \to 0$.)

b. ($\mathrm{im} j_* \subseteq \ker \delta$) By definition, $\delta([c]) = [a]$ where $a \in A_{n-1}$ is such that $i_{n-1}(a) = \partial^B_n(b)$ for some $b \in B_n$ such that $c = j_n(b)$. If $[c] = j_*([b])$, then we have such a b that is, in particular, a cycle (because it's a homology class representative). So $\partial^B_n(b) = 0$, and we have $i_{n-1}(a) = 0$, which implies $a = 0$ because i_{n-1} is injective by virtue of its position in the short exact sequence $0 \to A_{n-1} \to B_{n-1}$.

HW11#4. *Theorem.* Let X be a path-connected topological space, and $x \in X$. Then $H_n(X, \{x\}) = \begin{cases} H_n(X) & n > 0 \\ \{0\} & \text{else} \end{cases}$.

Proof. We use the long exact sequence of the pair. Recall that

$$... \to H_n(X) \to H_n(X, A) \to H_{n-1}(A) \to H_{n-1}(X) \to ...$$

is an exact sequence. In the case $A := \{x\}$, this becomes

$$... \to \underbrace{H_n(\{x\})}_{\{0\} \text{ if } n>0} \xrightarrow{e} H_n(X) \xrightarrow{f} \underbrace{H_n(X, \{x\})}_{\text{we want this}} \xrightarrow{g} \underbrace{H_{n-1}(\{x\})}_{\{0\} \text{ if } n>1} \xrightarrow{h} H_{n-1}(X) \to ...$$

If $n > 1$, then f is both surjective (from $g = 0$ and exactness properties) and injective (from $e = 0$ and exactness properties), thus an isomorphism, so $H_n(X) \cong H_n(X, \{x\})$.

It remains to check the $n := 1$ and $n := 0$ cases.

For $n := 1$, we have

$$... \to \underbrace{H_1(\{x\})}_{\{0\}} \xrightarrow{e} H_1(X) \xrightarrow{f} H_1(X, \{x\}) \xrightarrow{g} \underbrace{H_0(\{x\})}_{\mathbb{Z}} \xrightarrow{h} \underbrace{H_0(X)}_{\mathbb{Z}}$$

(where we know that $H_0(\{x\}) \cong H_0(X) \cong \mathbb{Z}$ from path-connectedness and the path-connection-component-counting nature of H_0).

For $n := 0$, we have

$$... \to \underbrace{H_0(\{x\})}_{\mathbb{Z}} \xrightarrow{e} \underbrace{H_0(X)}_{\mathbb{Z}} \xrightarrow{f} H_0(X, \{x\})$$

Here we turn to Gemini 3 Pro for tutoring, where it teaches us an important fact about exact sequences. I knew that a "zero on the left" makes the next map injective, and a "zero on the right" makes the previous map surjective. But consider $A \xrightarrow{f} B \xrightarrow{g} C$ where g is an isomorphism. The isomorphism gives us $\ker g = \{0\}$, thus $\text{im} f = \{0\}$.

That gives us $H_1(X) \cong H_1(X, \{x\})$ (the isomorphism to the right becomes a zero on the right).

Then in the $n := 0$ sequence, e being an isomorphism means that its image is $\mathbb{Z}$, which means the kernel of f is $\mathbb{Z}$, which means that $H_0(X, \{x\}) \cong \{0\}$.