Homological and Combinatorial Proofs of the Brouwer Fixed-Point Theorem

Everett Cheng

June 6, 2016

Three major, fundamental, and related theorems regarding the topology of Euclidean space are the Borsuk-Ulam theorem [6], the Hairy Ball theorem [3], and Brouwer's fixed-point theorem [6]. Their standard proofs involve relatively sophisticated methods in algebraic topology; however, each of the three also admits one or more *combinatorial* proofs, which rely on different ways of counting finite sets—for presentations of these proofs see [11], [8], [7], respectively. These combinatorial proofs tend to be somewhat elementary, and yet they manage to extract the same information as their algebraic-topology counterparts, with equal generality. The motivation for this paper, then, is to present the two types of proof in parallel, and to study their common qualities. We will restrict our discussion to Brouwer's fixed-point theorem, which in its most basic form states that a continuous self-map of the closed unit ball must have a fixed point.

We begin by presenting some basic formulations of Brouwer's theorem. We then present a proof using a combinatorial result known as Sperner's lemma, before proceeding to lay out a proof using the concept of homology from algebraic topology. While we will not introduce a theory of homology in its full form, we will come close enough to understand the essence of the proof at hand. The final portion of the paper is left for a comparison of the two methods, and a discussion of the insights that each method represents. Much of this discussion in inspired by Nikolai Ivanov's work in [7].

1 Definitions and background

First, let us denote by B^n and S^n the n-dimensional Euclidean closed unit ball and sphere, where $n \geq 0$:

$$B^n = \{ x \in \mathbb{R}^n : \|x\| \leq 1 \}, \quad S^n = \{ x \in \mathbb{R}^{n+1} : \|x\| = 1 \}.$$

Let Δ^n and $\mathring{\Delta}^n$ be the closed and open *standard n-simplex*;

$$\Delta^n = \{ x \in \mathbb{R}^{n+1} : x_i \geq 0, \sum x_i = 1 \},$$

$$\mathring{\Delta}^n = \{ x \in \mathbb{R}^{n+1} : x_i > 0, \sum x_i = 1 \}.$$

The following are some of the basic definitions of abstract topology, starting with the main object of study:

Definition 1.1. A *topological space* is a set S equipped with a set τ of subsets of S, such that

(i) The empty set $\emptyset$ and the entire set S are both members of τ, i.e., $\emptyset, S \in \tau$.

(ii) If $U_1, U_2, \ldots, U_n$ are members of τ, then their intersection $U_1 \cap U_2 \cap \cdots \cap U_n$ is a member of τ.

(iii) If U_i is a member of τ for every $i \in I$, where I is possibly infinite, then the union of all the U_i is a member of τ.

The collection of subsets τ is called the *topology* of the topological space. A subset $U \subseteq S$ is defined to be an *open set* of the topological space if $U \in \tau$. Likewise, a subset $V \subseteq S$ is defined to be a *closed set* of the topological space if its complement $S \setminus V$ is in τ.

If we take $S = \mathbb{R}^n$ and let τ contain the open sets defined in the normal Euclidean way, we can verify that n-dimensional Euclidean space is a topological space. In fact, the topological space is a way of generalizing the familiar properties of open and closed sets in Euclidean space. It turns out that these properties are even enough to allow us to define continuity:

Definition 1.2. For a function f from a topological space X to a topological space Y to be *continuous* means that for any open subset U in Y, the pullback $f^{-1}(U)$ is an open set in X.

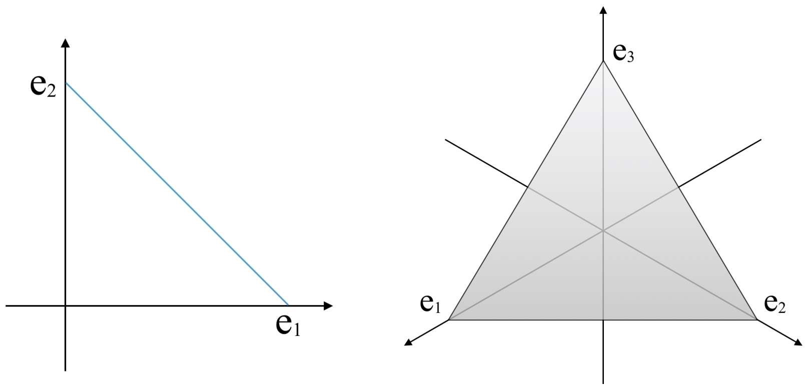

Figure 1: The standard 1-simplex as a set in $\mathbb{R}^2$, and the standard 2-simplex as a set in $\mathbb{R}^3$. Here e_i represents the i-th standard basis vector.

If X and Y are both Euclidean spaces, this definition is equivalent to the definition in terms of distance. However, for a general topological space the idea of distance is not defined, and so we are forced to define continuity only in terms of open sets. It is the case that many of the usual properties of continuous functions hold for the above definition, but we will omit such derivations here.

Now we will define two special kinds of continuous function:

Definition 1.3. A *retraction* r is a continuous function from a topological space X to a subspace $S \subseteq X$ such that $r(x) = x$ for all $x \in S$.

Here the term subspace indicates that the open sets in S are taken to be the intersections with S of the open sets in X. In the future, when we treat Δ^n, B^n or S^n as a topological space, we mean to consider it as a subspace of Euclidean space.

Definition 1.4. A *homeomorphism* h is a continuous function between topological spaces X and Y such that h is one-to-one and onto, and such that the inverse function $f^{-1}\colon Y \to X$ is continuous.

Definition 1.5. Topological spaces X and Y are *homeomorphic* or *topologically equivalent* if there exists a homeomorphism between them.

Homeomorphisms are important because they preserve many important properties of topological spaces, which we will observe more throughout the paper. As a first example, consider the property of compactness:

Definition 1.6. A topological space X is *compact* when every open covering of X has a finite subcovering. Here "open covering" refers to a covering of X whose each member is open in X.

Compactness is not a property that will change if we modify the size, shape, or orientation of a space. In fact, it is preserved by all homeomorphisms:

Proposition 1.1. *If X and Y are topological spaces, $h\colon X \to Y$ is a homeomorphism, and X is compact, then Y is compact.*

Proof. Let $C = \{ U_i : i \in I \}$ be any open covering of Y. Since h^{-1} is onto, $C' = \{ f^{-1}(U_i) : i \in I \}$ covers X, and each $f^{-1}(U_i)$ is open in X by the definition of continuity. Therefore C' has a finite subcovering $f^{-1}(U_1), f^{-1}(U_2), \ldots, f^{-1}(U_n)$. Since h is onto, $U_1, U_2, \ldots, U_n$ covers Y, and Y is compact. □

1.1 Simplices

The standard n-simplex is a topological object which is easy to reason about. It is so well-behaved, in fact, that we may want to understand other topological spaces in terms of the different continuous maps from the standard simplex.

Definition 1.7. The j-th *face* $\Delta^n[j]$ of the standard n-simplex is the intersection $\Delta^n \cap \{ x \in \mathbb{R}^n : x_j = 0 \}$, where $1 \leq j \leq n+1$. The *boundary* of the standard n-simplex $\partial\Delta^n$ is the union of its faces.

A face of a simplex has all of the same vertices as the original simplex, save one. In particular, the j-th face excludes the j-th vertex, so that excluding a certain vertex of the simplex uniquely determines one of its faces.

Definition 1.8. An n-dimensional *simplex* σ in a topological space X is a continuous map $\sigma\colon \Delta^n \to X$. If $n \geq 1$, the j-th *face* of σ is the $(n-1)$-simplex $\sigma[j] : \Delta^{n-1} \to X$ which maps

$$(x_1, x_2, \ldots, x_n) \mapsto \sigma(x_1, x_2, \ldots, x_{j-1}, 0, x_j, \ldots, x_n),$$

where $1 \leq j \leq n+1$. In general, a *k-dimensional face* of σ is a face of a $k+1$-dimensional face of σ. A *vertex* of σ is a 0-dimensional face, whose image is a single point. The *interior* $\operatorname{int}(\sigma)$ is the *image* of $\mathring{\Delta}^n$ under σ.

Note that the simplex is defined to be the continuous map itself, not the image. Note also that the simplex need not be a homeomorphism; in fact, an n-simplex could have an image consisting of only one point. For this reason, we would like to impose stronger conditions on the simplices we consider. It is better to define these conditions for groups of simplices rather than single ones; this leads us to simplicial complexes, which intuitively are ways of covering a region using non-overlapping simplices:

Definition 1.9. A *simplicial n-complex* $\mathcal{K}$ over a topological space X is a finite set S of simplices in X such that

(i) If $\sigma \in S$, then each face $\sigma[j]$ is in S.

(ii) The maximum dimension of any simplex in S is n.

(iii) A simplex in S is one-to-one over its interior, and each point $x \in X$ lies in the interior of exactly one simplex.

(iv) A subset $A \subseteq X$ is open in X if and only if the pullback $\sigma^{-1}(A)$ is open for every $\sigma \in S$.

(v) No two simplices in S share the exact same set of vertices.

In a simplicial complex, we identify each simplex with its vertices, so that $\sigma := [v_1, v_2, \ldots, v_{k+1}]$. It is easy to see that the face $\sigma[j]$ is given by $[v_1, \ldots, v_{j-1}, v_{j+1}, \ldots, v_{k+1}]$.

2 The Brouwer fixed-point theorem

The version of Brouwer's theorem stated here was proved in different ways by L. E. J. Brouwer [1] and by Jacques Hadamard [5], both in 1910. As one of the first results of algebraic topology, it remains important throughout many areas of mathematics. Among other results, Brouwer's theorem can be used to prove the Fundamental Theorem of Algebra [4], the Perron-Frobenius theorem in linear algebra [2], the existence of Nash equilibria in game theory [10], and the Jordan curve theorem [9], which is another fundamental and celebrated topological result. Due to its importance, Brouwer's fixed-point theorem has been generalized in numerous ways. However, for the purposes of this paper the following formulations will be sufficient:

Theorem 2.1 (Brouwer's fixed-point theorem)**.** *Any continuous function f mapping $B^n \to B^n$ has a fixed point, i.e., $f(x) = x$ for some $x \in B^n$. Alternatively, any continuous function $g\colon \Delta^n \to \Delta^n$ has a fixed point.*

The theorem holds for a large class of topological spaces besides just B^n and Δ^n. The following more general form is not proved directly, but as a consequence of the first form:

Theorem 2.2. *If X is a topological space homeomorphic to B^n, then any continuous function $f\colon X \to X$ has a fixed point.*

Proof. Take a homeomorphism $h\colon X \to B^n$ and consider the continuous composition $h \circ f \circ h^{-1}\colon B^n \to B^n$. By Theorem 2.1 it must have a fixed point $h(f(h^{-1}(x))) = x$, and applying h^{-1} to both sides we have $f(h^{-1}(x)) = h^{-1}(x)$. □

A common version of Brouwer's theorem is as follows. Although we state it without proof, this statement is implied by Theorem 2.2, since sets satisfying the hypothesis are homeomorphic to B^n.

Theorem 2.3. *If S is homeomorphic to a compact, convex subset of $\mathbb{R}^n$ with an interior point, then any continuous function $f\colon S \to S$ has a fixed point.*

A fourth statement turns out to be equivalent to Brouwer's theorem, and is often proved instead of the formulations above:

Theorem 2.4 (No Retraction theorem)**.** *For $n \geq 1$, there is no continuous retraction from B^n to S^{n-1}. Alternatively, there is no continuous retraction from Δ^n to $\partial\Delta^n$. In general, if $h\colon B^n \to X$ is a homeomorphism, there is no continuous retraction from X to $h(S^{n-1})$.*

The proof of this equivalence is given in many sources, one of which is [6].

Proof. Given Theorem 2.1, there cannot exist a continuous retraction r mapping $X \to h(S^{n-1})$, since the map $x \mapsto h^{-1}(-r(h(x)))$ would have no fixed point in B^n. Conversely, supposing that the fixed-point theorem is false for B^n, we can construct a retraction. Let $f\colon B^n \to B^n$ have no fixed points, and consider for any $\mathbf{x} \in B^n$ the ray $\rho_{\mathbf{x}}(t) = (1-t)f(\mathbf{x}) + t\mathbf{x}, t \geq 0$, which is nontrivial since $\mathbf{x} \neq f(\mathbf{x})$. This ray must make its closest approach to the origin at a point inside B^n, since it starts inside B^n. Further, after its closest approach, its distance from the origin only increases, so after it intersects S^{n-1} for the first time it cannot intersect it again. Therefore, let $r(\mathbf{x})$ be the point $\rho_{\mathbf{x}}(t)$ where $\|\rho_{\mathbf{x}}(t)\| = 1$, and notice that r is a retraction. Finally, $\rho_{\mathbf{x}}(t)$ is jointly continuous in $\mathbf{x}$ and t, by the continuity of f. The value of t at which $\|\rho_{\mathbf{x}}\| = 1$ also varies continuously in $\mathbf{x}$, since B^n is convex. So g is a continuous retraction from B^n to S^{n-1}, and $h \circ r \circ h^{-1}$ is the desired retraction. □

Having understood some of the forms of Brouwer's Theorem, we can proceed to give proofs.

3 Combinatorial proof of Brouwer's theorem

The following proof uses Sperner's lemma, which is a statement about the possible ways to label vertices in a grid of simplices (a simplicial complex). By doing so, it restricts its consideration to only finite amounts of information regarding the continuous function in question. This is part of what makes the method "combinatorial"; since we concern ourselves only with finite concepts—such as the vertices of simplices and their labellings—the infinite nature of the general continuous function is irrelevant, and we can focus on counting finitely many combinations of behaviors.

We reproduce the general proof of Sperner's lemma given in [7]; for a particularly elegant proof of the two-dimensional case, see [8].

Theorem 3.1 (Sperner's lemma)**.** *If Δ^n has a simplicial n-complex $\mathcal{K}$ with a set of vertices V, and a a labelling function $\ell\colon V \to \{1, 2, \ldots, n+1\}$ is such that $\ell(v) \neq j$ for $v \in \Delta^n[j]$, then there are an odd number of n-simplices σ for which ℓ attains every one of its values on the vertices of σ. We say that such σ are fully-labelled.*

Proof. We use induction on n. If $n = 1$, V is an ordered list of vertices $v_1, v_2, \ldots, v_r$, and the number of fully-labelled line segments between them is the number of times $\ell(v_i)$ changes between 1 and 2 when stepping through. But since $\ell(v_1) = 1$ and $\ell(v_r) = 2$, this number must be odd.

Next, assume Sperner's lemma holds for some $n-1$, and take a simplicial n-complex $\mathcal{K}$ over Δ^n with a set of vertices V and a labelling function ℓ satisfying the hypothesis. Let $\alpha_1, \ldots, \alpha_e$ be the fully-labelled n-simplices of $\mathcal{K}$. Each of these has exactly one face whose labels include all $1 \leq j \leq n$, which is found by excluding the vertex labelled $n+1$. Let $\beta_1, \ldots, \beta_f$ be the n-simplices on whose vertices ℓ attains every value except $n+1$. Each of these β_i has two distinct vertices which have the same label; by excluding exactly one of these vertices, we can form exactly two faces of β_i whose labels include all $1 \leq j \leq n$.

Let $\gamma_1, \ldots, \gamma_g$ be the $(n-1)$-simplices of $\mathcal{K}$ contained in $\Delta^n[n+1]$ on whose vertices ℓ attains every value $1 \leq j \leq n$. Note that these are these are the only simplices in $\partial\Delta^n$ with this property, since any simplex on the boundary which is not contained in $\Delta^n[n+1]$ is contained within $\Delta^n[j]$, meaning none of its vertices have j as a label.

By the inductive hypothesis, g is odd. Also, since the γ_i are on the boundary of Δ^n, each one is a face of exactly one n-simplex of $\mathcal{K}$.[1] Let $\delta_1, \ldots, \delta_h$ be the $(n-1)$-simplices not on the boundary of Δ^n, but on whose vertices ℓ attains every value $1 \leq j \leq n$ as above. Since these δ_i are not on the boundary of Δ^n, each is a face of exactly two n-simplices of $\mathcal{K}$.[2]

Let us count the number of pairs σ, σ' in $\mathcal{K}$ such that σ is an n-simplex, σ' is a face of σ, and the labels of the vertices of σ' include all $1 \leq j \leq n$. On the one hand, the result should be $e + 2f$, since letting $\sigma = \gamma_i$ yields one such pair, setting $\sigma = \delta_i$ yields two such pairs, and if σ is neither then no faces of σ have the desired labellings. On the other hand, we should get $g + 2h$, since setting $\sigma' = \alpha_i$ yields one pair, setting $\sigma' = \beta_i$ yields two pairs, and no other σ' have the desired labellings. We conclude that

$$e + 2f = g + 2h.$$

Since the right-hand side is odd, the left-hand side is also odd, making e odd. □

We may now derive Theorem 2.4 from Theorem 3.1 following [7]; in particular we will prove the second form of Theorem 2.4, which is stated for the standard n-simplex rather than the ball. Naturally, the proof relies on taking simplicial complexes over the standard simplex, and at this point we have not proven the existence of any such simplicial complex. As such, let the simplicial n-complex $\mathcal{K}_0$ consist of a simplex which is the identity map on Δ^n, along with all of the k-dimensional faces of this simplex for $k \geq 0$. One can verify that the axioms of a simplicial complex are satisfied for $\mathcal{K}_0$.

We can divide the standard n-simplex into smaller n-simplices as follows:[3] for $1 \leq i \leq n+1$, take $\tau_i \colon \Delta^n \to \Delta^n$ to be the simplex which maps $\mathbf{x} = (x_1, x_2, \ldots, x_{n+1})$ to

$$\frac{1}{2}\mathbf{x} + \frac{1}{2}\mathrm{e}_i, \tag{1}$$

[1]Take x_0 in the interior of γ_i. Each point y within $\epsilon > 0$ of x_0 which is inside the interior of Δ^n must lie in the interior of some simplex σ_y of $\mathcal{K}$. Since simplicial complexes are finite in this paper, all other simplices of dimension smaller than n are separated from x_0 by a nonzero distance, meaning σ_y will be an n-simplex at each y as long as ϵ is small. Let σ^* be one of the finitely-many σ_y which is the σ_y for y values with ϵ arbitrarily small. Then $x_0 \in \partial\sigma^*$, and the face of σ^* containing x_0 must equal γ_i, since x_0 can only be in the interior of one simplex. There are no other n-simplices with γ_i as a face, since the points y do not lie in the interiors of multiple simplices.

[2]The argument is similar to the above, except the points within ϵ of x_0 now lie in the interiors of two different n-simplices, one for each side of δ_i.

[3]There are many ways of subdividing simplices, and this one was chosen because it is intuitive and easy to visualize.

where e_i is the i-th standard basis vector in $\mathbb{R}^{n+1}$. Take $\tau_{n+2}\colon \Delta^n \to \Delta^n$ to be the simplex which maps $\mathbf{x}$ to

$$\mathbf{c} - \frac{1}{2}(\mathbf{x} - \mathbf{c}) = \frac{3}{2}\mathbf{c} - \frac{1}{2}\mathbf{x}, \tag{2}$$

where $\mathbf{c} = (1/(n+1), 1/(n+1), \ldots, 1/(n+1))$ is the center point of Δ^n. The images of these maps cover Δ^n without overlapping interiors, and if $1 \leq i \leq n+1$, then the i-th face of τ_i coincides with the i-th face of τ_{n+2}. The rest of the faces of the first $n+1$ simplices lie on the boundary of Δ^n.

Using these maps, we may subdivide a simplex σ into $n+2$ smaller simplices by taking the compositions $\sigma \circ \tau_i$ for $1 \leq i \leq n+2$. Indeed, the n-simplex of $\mathcal{K}_0$ can be divided into $n+2$ smaller simplices using this method, and those $n+2$ simplices along with their faces form a simplicial complex $\mathcal{K}_1$. Repeating the process, we take the n-simplices of $\mathcal{K}_h$ and apply each of the $n+2$ subdivision maps to each one. The resulting simplices, along with their faces of all dimensions, constitute the simplicial complex $\mathcal{K}_{h+1}$. Some properties to note are that $\mathcal{K}_h$ has $(n+2)^h$ different n-simplices, and that the diameter of one of these n-simplices is half that of one of the n-simplices of $\mathcal{K}_{h-1}$ (the diameter of a simplex is the maximum Euclidean distance between any two points in its image).

The most important part of this construction is that for any positive ϵ, beyond a certain k all of the n-simplices of $\mathcal{K}_h$ have diameter less than ϵ. This makes our proof of the No Retraction Theorem possible.

First, though, we prove the Lebesgue Number Lemma, which is used in the proof.

Lemma 3.1 (Lebesgue's Number Lemma)**.** *If X is a compact subset of $\mathbb{R}^n$ and $\mathcal{U}$ is an open cover of X, then there is some number $\delta > 0$ such that every closed ball of radius δ in X is contained within a some element of $\mathcal{U}$.*

Proof. Let $B_r(x)$ denote the portion of the closed ball of radius r centered at a point x which lies in X. For a given δ, let $T_\delta \subseteq X$ consist of the points $x \in X$ such that $B_\delta(x)$ is contained within one element of $\mathcal{U}$. This T_δ is open, since if $x_0 \in T_\delta$ has $B_\delta(x_0) \subseteq U_0 \in \mathcal{U}$, then $B_\delta(x_0)$ is a finite distance a from $X \setminus U_0$ and $B_{a/2}(x_0)$ is contained in T_δ.

The sets T_δ for $\delta > 0$ cover X, since each $x \in X$ is contained in some element of $\mathcal{U}$ which must also contain a ball around x. By compactness, a finite number of T_δ cover X, which means that $X = T_\delta$ for some δ. $\square$

As a final definition, let the *open star* $\operatorname{st} v$ of a vertex v in one of the $\mathcal{K}_h$ be the union of the interiors of the simplices in $\mathcal{K}_h$ for which v is a vertex. Note that the diameter of $\operatorname{st} v$ cannot be larger than twice the maximum diameter of simplices in $\mathcal{K}_h$, by the Triangle Inequality.

Proof of Theorem 2.4. Assume for the sake of contradiction that $r\colon \Delta^n \to \partial\Delta^n$ is a continuous retraction. The set differences $\partial\Delta^n \setminus \Delta^n[j]$ are open as subsets of $\partial\Delta^n$, since the sets $\mathbb{R}^{n+1} \setminus \Delta^n[j]$ are open in $\mathbb{R}^{n+1}$ (recall that $\partial\Delta^n$ can be treated as a topological subspace of $\mathbb{R}^{n+1}$). By continuity of r, the preimages $U_j = r^{-1}(\partial\Delta^n \setminus \Delta^n[j])$ are open in Δ^n, and they must cover Δ^n since the $\partial\Delta^n \setminus \Delta^n[j]$ cover $\partial\Delta^n$. By the Lebesgue number lemma, there is some δ such that balls of radius δ inside Δ^n lie inside single sets U_j. If we choose $\mathcal{K}_h$ with h large enough that the diameters of simplices are smaller than $\delta/2$, then we can guarantee that the open stars of the vertices of $\mathcal{K}_h$ lie inside single sets U_j.

We now label the vertices of $\mathcal{K}_h$ as follows: given a vertex v, $\ell(v)$ will be some j such that $r(\operatorname{st}(v)) \subseteq \partial\Delta^n \setminus \Delta^n[j]$. In other words, v will be labelled by some j such that the retraction r never maps points "near" v to the j-th face of Δ^n. This labelling satisfies the hypothesis of Sperner's lemma, since if v is on $\Delta^n[j]$ then $\operatorname{st}(v)$ has a point in the interior of $\Delta^n[j]$ as a limit point, and by continuity of r, $r(\operatorname{st}(v))$ also has this interior point as a limit point. Therefore $r(\operatorname{st}(v))$ must intersect $\Delta^n[j]$, and v is not labelled by j. Hence, by Sperner's lemma there exists an n-simplex σ of $\mathcal{K}_h$ which is fully-labelled. But the interior $\operatorname{int}(\sigma)$ is in the open star of each of the vertices of σ, and so $r(\operatorname{int}(\sigma))$ is contained in $\partial\Delta^n \setminus \Delta^n[j]$ for all j. Since $\operatorname{int}(\sigma)$ is non-empty, this is a contradiction. □

4 Homological proof of Brouwer's theorem

The idea of homology is to understand a topological space by examining the number of "holes" it has of a particular dimension. For example, the Euclidean space $\mathbb{R}^n$ is free of holes, while the sphere S^n possesses one n-dimensional hole. There are many useful theories of homology; here we will introduce the theory of *singular homology*, since it leads to a relatively short proof of Brouwer's theorem. The adjective *singular* refers to the fact that we are working with arbitrary continuous simplices, as opposed to members of simplicial complexes, and that the simplices can have "singularities" where they are not one-to-one.

4.1 Singular homology

To develop singular homology, we first must define the *singular chain*, which is an abstract structure built from simplices. Unlike simplices, though, they carry a notion of *orientation*; a chain can be thought of as a collection of simplices in a topological space, along with information about the "direction" in which each simplex should be traversed. For example, the two possible directions for traversing the faces of a 2-simplex correspond to clockwise and counterclockwise. In general, there are only two possible orientations, which we represent as a sign, "+" or "−". We can formally represent this collection of simplices as a sum:

Definition 4.1. A *singular n-chain* C of a topological space X is a finite formal sum of n-simplices of X, or of their negations. Symbolically,

$$C := \sum_{i=1}^{k} \pm\sigma_i \ ,$$

where $\sigma_i \colon \Delta^n \to X$ are n-simplices of X. These σ_i are not required to be distinct, so $\sigma_0 + \sigma_0$ is a valid chain. However, any repeated occurrences of a simplex should have the same sign. The negation $-C$ is formed by reversing the sign of each term. The sum $C_1 + C_2$ of two n-chains is formed by joining their terms and cancelling any pairs of opposite terms such as $\sigma_0 + (-\sigma_0)$. When working with chains, we use 0 to denote the empty chain.

Finally, we use the notation $C_n(X)$ to represent the set of all singular n-chains of X.

It is important to realize that the "+" and "−" appearing here do not represent any real operations related to the simplices. They are simply formal symbols which allow us to define singular chains. The sum of two simplices is not defined, and does not yield a new simplex. The only objects we may add are singular chains of the same dimension, and doing so produces another chain.

Definition 4.2. If $C \in C_n(X)$ is a singular n-chain on a topological space X which consists of only one simplex σ_0, then its *boundary* $\partial C \in C_{n-1}(X)$ is the singular $(n-1)$-chain

$$\partial C := \sum_{i=1}^{n+1} (-1)^{i+1} \sigma_0 \, [i] \, .$$

In order to define boundaries for chains C of multiple simplices, we let $\partial(-C) = -\partial C$ and $\partial(C_1 + C_2) = \partial C_1 + \partial C_2$. The boundary of a 0-chain is defined to be the empty chain 0.

Homology involves the study of special types of chains, in particular the chains without boundary. Intuitively, these are chains whose simplices combine to form

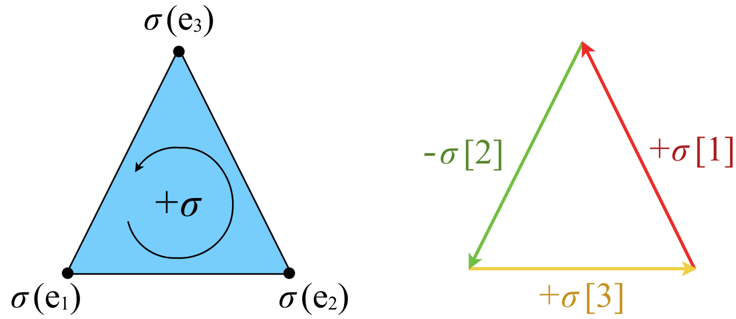

Figure 2: A representation of a chain consisting of one 2-simplex, and a representation of its boundary; in truth, we are only depicting the images of these simplices. We treat the page as part of the topological space $\mathbb{R}^2$. The arrows are shorthands for the orientation of each simplex: when the sign of a simplex is positive, the arrow points towards the vertices in ascending order, and when the sign is negative it points to the vertices in descending order.

---

closed loops or closed surfaces, so that any piece of the boundary of one of its simplices is also a piece of the boundary of another simplex. Chains without boundary must also have a consistent orientation, so that each piece of boundary of one of its simplices appears with each of its two possible orientations an equal number of times.

Definition 4.3. A singular chain $C \in C_n(X)$ is a *cycle* when $\partial C = 0$.

We are also interested in chains which are themselves boundaries; if $C = \partial B$, then one can roughly imagine that C encloses some region of the topological space. In other words, C does not enclose a hole in the space, since every point enclosed by C is contained in the images of the simplices of B. Note, however, that this intuition is not rigorous in the general case.

Definition 4.4. A singular n-chain C of a topological space X is a *boundary* when there is some $B \in C_{n+1}(X)$ such that $C = \partial B$.

Proposition 4.1. *A boundary $C = \partial B$ in $C_n(X)$ is a cycle.*

Proof. We will assume B consists of a single $(n+1)$-simplex σ_0. The general case will then follow from the fact that if C_1 and C_2 are cycles, so are $-C_1$ and $C_1 + C_2$. By the definition of a chain's boundary,

$$C = \sum_{i=1}^{n+2} (-1)^{i+1} \sigma_0 [i] \, ,$$

and

$$\partial C = \sum_{i=1}^{n+2} (-1)^{i+1} \partial (\sigma_0 [i]) = \sum_{i=1}^{n+2} \sum_{j=1}^{n+1} (-1)^{i+j} \sigma_0 [i] [j] . \tag{3}$$

Recall that the i-th face of σ_0 is defined to be the map

$$\Delta^n \to X \colon (x_1, \dots , x_{n+1}) \mapsto \sigma_0(x_1, \dots , x_{i-1}, 0, x_i, \dots , x_{n+1}).$$

Assuming $i \leq j$, the j-th face of the i-th face of σ_0 is the map from Δ^{n-1} to X given by

$$\begin{aligned} \sigma_0 [i] [j] : (x_1, \dots , x_n) &\mapsto \sigma_0 [i] (x_1, \dots , x_{j-1}, 0, x_j, \dots , x_n) \\ &= \sigma_0(x_1, \dots , x_{i-1}, 0, x_i, \dots , x_{j-1}, 0, x_j, \dots , x_n). \end{aligned}$$

Similarly, the i-th face of the $(j + 1)$-st face of σ_0 is the map from Δ^{n-1} to X given by

$$\begin{aligned} \sigma_0 [j+1] [i] : (x_1, \dots , x_n) &\mapsto \sigma_0 [j+1] (x_1, \dots , x_{i-1}, 0, x_i, \dots , x_n) \\ &= \sigma_0(x_1, \dots , x_{i-1}, 0, x_i, \dots , x_{j-1}, 0, x_j, \dots , x_n), \end{aligned}$$

and so it turns out that $\sigma_0 [i] [j] = \sigma_0 [j + 1] [i]$ whenever $i \leq j$. Now we split (3) into two sums, the first with $i > j$ and the second with $i \leq j$:

$$\partial C = \sum_{i=2}^{n+2} \sum_{j=1}^{i-1} (-1)^{i+j} \sigma_0 [i] [j] + \sum_{i=1}^{n+1} \sum_{j=i}^{n+1} (-1)^{i+j} \sigma_0 [j+1] [i] .$$

Now we reindex the second sum with $r = j + 1$, $s = i$, to get

$$\partial C = \sum_{i=2}^{n+2} \sum_{j=1}^{i-1} (-1)^{i+j} \sigma_0 [i] [j] - \sum_{s=1}^{n+1} \sum_{r=s+1}^{n+2} (-1)^{r+s} \sigma_0 [r] [s] .$$

If we were to switch the order of the summations in the second term, r should range from 2 to $n + 2$, and for any fixed r, s should range from 1 to $r - 1$. Therefore, the two sums are equal, and their difference is zero. □

The most important types of chains for understanding homology are those which are cycles but not boundaries. In contrast with the boundaries, these chains can be thought of as enclosing one or more holes in the underlying space. Of course, there will be many such cycles corresponding to any one hole in the space; in fact, two cycles might be said to enclose the same hole exactly when they differ by a boundary. This insight motivates the following definitions:

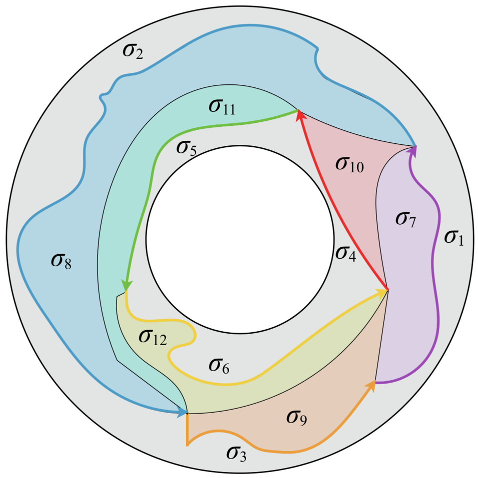

Figure 3: Simplices of a topological space which is shaped as an annulus. Each 2-simplex has the counterclockwise orientation. The 1-cycle $C = \sigma_1 + \sigma_2 + \sigma_3$ is homologous to the other 1-cycle $D = \sigma_4 + \sigma_5 + \sigma_6$, since $C - D$ is equal to the boundary of the 2-chain $\sigma_7 + \cdots + \sigma_{12}$. Neither is a boundary; intuitively, C and D enclose the same hole in the underlying space.

Definition 4.5. Two n-cycles C_1 and C_2 of a space X are *homologous* when their difference $C_1 - C_2$ is a boundary.

Definition 4.6. If C is a cycle of a space X, the *homology class* of C is denoted $[C]$ and is defined as the set of cycles which are homologous to C. Every space has a trivial homology class $[0]$, consisting of all of its boundaries. Given a space X, we denote by $H_n(X)$ the set of all homology classes of singular n-chains of X.

A topological space X has n-cycles which are not boundaries if and only if it has a homology class in $H_n(X)$ which is not $[0]$. This occurs whenever X has a "hole." As one might imagine, the space consisting of a single point, $\{0\}$, has no n-dimensional holes for $n \geq 1$:

Proposition 4.2. *A cycle in $C_n(\{0\})$ is a boundary for $n \geq 1$.*

Proof. There is only one n-simplex of $\{0\}$ for each n, which maps all of Δ^n to 0. We will call this simplex σ_n, and take C_n to be the chain consisting of just σ_n. The boundary of C_n is just an alternating sum of $n+1$ copies of C_{n-1}, which is equal to 0 for odd n and C_{n-1} for even $n \geq 2$. The cycles of $\{0\}$ with $n \geq 1$ are therefore either zero chains, or chains with odd dimension. For odd n, C_n is the boundary of C_{n+1}, and so since each n-cycle is a sum of some number of C_n, each n-cycle is a boundary. Combined with the fact that 0 is a boundary, this completes the proof. □

4.1.1 Maps between chains

An important property of homology is related to the idea that we can apply continuous functions to singular chains. If $f\colon X \to Y$ is a continuous map and $\sigma\colon \Delta^n \to X$

is an n-simplex of X, then f naturally associates to σ the n-simplex $f \circ \sigma$ of Y. Moreover, if C is a singular n-chain of X given by

$$\sum_{i=1}^{k} \pm\sigma_i,$$

then f associates to C the singular n-chain of Y given by

$$f(C) := \sum_{i=1}^{k} \pm(f \circ \sigma_i),$$

where the sign of each simplex is preserved. As usual, we cancel terms which have the same simplex but opposite signs. This is necessary since the $f \circ \sigma_i$ are not guaranteed to be distinct for distinct σ_i. Finally, if 0 is the chain consisting of no simplices, then $f(0) = 0$.

This application of f to chains is linear, in that $f(-C) = -f(C)$ and $f(C_1+C_2) = f(C_1) + f(C_2)$. The most important property of this operation, however, is that

$$f(\partial C) = \partial f(C) \tag{4}$$

for any chain C of X. By appealing to linearity as before, we only need to show the equality when C consists of a single n-simplex σ_0. In this case, $f(C) = f \circ \sigma_0$, so

$$\partial f(C) = \sum_{j=1}^{n+1} (-1)^{j+1} (f \circ \sigma_0)\,[j]\,.$$

The j-th face of $f \circ \sigma_0$ is the simplex

$$(x_1, \ldots, x_n) \mapsto f(\sigma_0(x_1, \ldots, x_{j-1}, 0, x_j, \ldots, x_n)),$$

and so it is equal to $f(\sigma_0\,[j])$. We now have

$$\partial f(C) = \sum_{j=1}^{n+1} (-1)^{j+1} f(\sigma_0\,[j]) = f(\partial C)$$

as desired.

Among other results, this shows that applying f to a boundary yields a boundary, and that applying f to a cycle yields another cycle. The identity is also useful because of the following fact, and because of the definition it makes possible:

Proposition 4.3. *If two n-cycles C_1 and C_2 of a space X are homologous, then $f(C_1)$ and $f(C_2)$ are homologous n-cycles of Y.*

Proof. Let $C_1 - C_2 = \partial B$ for some $B \in C_{n+1}(X)$. Then $f(C_2) - f(C_1) = f(\partial B) = \partial f(B)$. □

Definition 4.7. If $f\colon X \to Y$ is continuous, then for each n there is a function $f_*\colon H_n(X) \to H_n(Y)$ which maps a homology class $[C]$ of X to the homology class $[f(C)]$ of Y. By Proposition 4.2, f_* is well-defined in that it only assigns a single result to a given homology class.

A simple fact about the map f_* is that $f_*([0]) = [f(0)] = [0]$. This will be useful in our proof of Brouwer's theorem, which now follows.

4.2 Proof of No Retraction theorem

Rather than proving Brouwer's theorem directly, we will prove the No Retraction theorem which we know to be equivalent. The following version is taken from [3], and given the tools we have developed so far, the proof is for the most part very simple. However, it relies on the following facts:

Proposition 4.4. *Every cycle in $C_k(B^n)$ is a boundary as long as $k \geq 1$. There exists a cycle in $C_n(S^n)$ which is not a boundary.*

We will discuss Proposition 4.4 afterwards, since it is more difficult to prove.

Proof of Theorem 2.4. Assume for the sake of contradiction that $r\colon B^n \to S^{n-1}$ is a continuous retraction, and let $i\colon S^{n-1} \to B^n$ be the "inclusion" or identity map (recall that S^{n-1} is a subset of B^n in Euclidean space). By the definition of a retraction, $r \circ i$ is the identity map on S^{n-1}. Let $[C]$ be any homology class in $H_{n-1}(S^{n-1})$, and note that $(r \circ i)_*([C]) = [C]$. On the other hand, $(r \circ i)_*([C])$ equals $r_*([i(C)])$, and since $i(C)$ is an n-cycle in B^n, Proposition 4.4 tells us that $[i(C)] = [0]$. This shows that $[C] = r_*([0]) = [0]$, or equivalently that all $(n-1)$-cycles of S^{n-1} are boundaries. By Proposition 4.4, we have a contradiction. □

We proceed to discuss Proposition 4.4.

4.2.1 Homotopy equivalence and the homology of B^n

Intuition tells us that all cycles in B^n should be boundaries, since B^n has no holes. In order to formalize this notion, it is best to introduce a new type of equivalence between topological spaces.

Definition 4.8. Two continuous functions $f, g\colon X \to Y$ are *homotopic* when there exists a continuous function $H\colon X \times [0,1] \to Y$ such that $H(x,0) = f(x)$ and $H(x,1) = g(x)$ for $x \in X$. This function H is called a *homotopy* between f and g.

If $f = g$, then f and g are homotopic, since $H(x,t) = f(x)$ is a homotopy.

Definition 4.9. A topological space X is *homotopy equivalent* to another space Y when there exist continuous maps $f\colon X \to y$ and $g\colon Y \to X$ such that $g \circ f$ is homotopic to the identity map in X and $f \circ g$ is homotopic to the identity map in Y. In this case, f and g are called *homotopy equivalences*.

The use of the variable t in $H(x,t)$ is suggestive of time, and indeed these definitions capture the intuitive idea of *continuously deforming* one map or space into another. Two functions are homotopic when we can continuously change one into the other over some time interval, where the intermediary maps are $f_t\colon X \to Y : x \mapsto H(x,t)$.

Viewed from another angle, homotopy equivalence captures the same fundamental idea as homeomorphism, namely that two spaces have similar topological properties as long as we can change between them in a continuous fashion. Homotopy equivalence, though, is a weaker condition. For example, an important property of a homeomorphism is that it preserves dimension, i.e., an open subset of $\mathbb{R}^n$ and an open subset of $\mathbb{R}^m$ can only be homeomorphic if $m = n$; on the other hand, all of the balls B^n are mutually homotopy equivalent since B^n can be continuously expanded to fill B^{n+1}.

Proposition 4.5. *If X and Y have a homeomorphism $h\colon X \to Y$, then X and Y are homotopy equivalent.*

Proof. Let the homotopy equivalences be given by $h\colon X \to Y$ and $h^{-1}\colon Y \to X$. Both are continuous, and $h \circ h^{-1}$ and $h^{-1} \circ h$ are the identity functions of Y and X. Since the identity functions are homotopic to themselves, the conditions for homotopy equivalence are satisfied. □

For our purposes, the most important feature of homotopy-equivalent spaces is that they have the same number of homology classes. In order to show this, we need the following fact:

Proposition 4.6. *If the continuous maps $f, g\colon X \to Y$ are homotopic, then the maps $f_*, g_*\colon H_n(X) \to H_n(Y)$ coincide.*

The proof of Proposition 4.6 is beyond the scope of this paper; see [6], page 112.

Proposition 4.7. *If X and Y are homotopy equivalent, there exists a one-to-one and onto correspondence between $H_n(X)$ and $H_n(Y)$.*

Proof. Let $f\colon X \to Y$ and $g\colon Y \to X$ make up a homotopy equivalence. Then $(g\circ f)_*$ is the identity map on $H_n(X)$, so it is one-to-one. Since $(g \circ f)_* = g_* \circ f_*$, f_* must be one-to-one. On the other hand, $(f \circ g)_*$ is the identity map on $H_n(Y)$, so it is onto. Since $(f \circ g)_* = f_* \circ g_*$, f_* must be onto. So f_* gives the desired correspondence. □

The fact that will be most useful to us is as follows: if X and Y are homotopy equivalent, X has n-cycles which are not boundaries if and only if Y has n-cycles which are not boundaries.

In particular, it happens that B^n is homotopy equivalent to the topological space $\{0\}$ consisting of a single point. Indeed, if we let f map B^n to $\{0\}$ in the only way possible, and g map $\{0\}$ to the center point $\mathbf{0} \in B^n$, then $f \circ g$ is the identity on $\{0\}$ and $g \circ f$ maps B^n to its center point. By letting $H\colon B^n \times [0, 1] \to B^n$ map $\mathbf{x}$ to $t\mathbf{x}$, we see that $g \circ f$ is homotopic to the identity on B^n.

By Proposition 4.2, every k-cycle in $\{0\}$ is a boundary when $k \geq 1$. We have shown that the same is true for B^n.

4.2.2 The homology of S^n

We now turn our attention to the sphere S^n. The simplest possible case is that of S^0, which only consists of the two isolated points $-1, 1 \in \mathbb{R}$.

Lemma 4.1. *There exists a 0-cycle of S^0 which is not a boundary.*

Proof. There are only two k-simplices of S^0, each of which maps its entire domains to a single one of the points -1 or 1; this is because continuous functions map connected sets to connected sets. If σ is one of the two 1-simplices, the faces of σ map Δ^0 onto the same point, and so $\sigma\,[1] = \sigma\,[2]$. The chain consisting of σ therefore has an empty boundary, and by the linearity of the boundary operator, so does any other 1-chain of S^0.

Let σ_1 be the 0-simplex mapping Δ^0 to the point 1, and let σ_2 be the other 0-simplex which maps to -1. Let C_1 be the chain consisting of σ_1 and likewise for C_2. These chains are nonzero and so cannot be boundaries. Since the boundary of a 0-chain is defined to be zero, C_1 and C_2 are also cycles. □

By the linearity of the boundary operator, any nonzero combination of C_1 and C_2 will be a cycle which is not a boundary. One such cycle which will be useful later is

$$C_1 - C_2. \tag{5}$$

In what follows, we will sketch a proof of the second part of Proposition 4.4 using induction on the dimension n. In particular, we want to show that the existence of an $(n-1)$-cycle of S^{n-1} which is not a boundary implies the existence of an n-cycle of S^n which is not a boundary. The proof is a special case of the *Mayer-Vietoris sequence*, which is useful for computing the homology classes of many topological spaces. Unfortunately, the proof of the Mayer-Vietoris sequence relies on a property of singular homology about *excision*, the proof of which is beyond the scope of this paper. We will therefore leave out the derivation of Fact 4.1. Before stating it, it will be helpful to begin the proof.

In the remainder of this section, it is assumed that $n \geq 1$. The first step is to divide S^n into two overlapping pieces,

$$S_1^n = \left\{ (x_1, \ldots, x_{n+1}) \in S^n : x_{n+1} \leq \frac{1}{2} \right\} ,$$
$$S_2^n = \left\{ (x_1, \ldots, x_{n+1}) \in S^n : x_{n+1} \geq -\frac{1}{2} \right\} .$$

The subset of $S_1^n \cap S_2^n$ with $x_{n+1} = 0$ is homeomorphic to S^{n-1}, and the entire intersection $S_1^n \cap S_2^n$ is homotopy equivalent to this subset. For example, a homotopy equivalence $f\colon (S_1^n \cap S_2^n) \times [0,1] \to (S_1^n \cap S_2^n)$ can be constructed by taking $(x_1, \ldots, x_{n+1}, t)$ to $(x_1, \ldots, x_n, tx_{n+1})$, and normalizing the result so that it lies on the unit sphere. Therefore the homology classes of $S_1^n \cap S_2^n$ are the same as those of S^{n-1}.

The chain in (5) naturally extends to a chain in $C_0(S_1^1 \cap S_2^1)$, where we map the point $1 \in S^0$ to $(1,0) \in S_1^1 \cap S_2^1$ and the point $-1 \in S^0$ to $(-1,0)$. It also extends to a chain in $C_0(S_1^1)$ defined in the same way, as well as one in $C_0(S_2^1)$. The latter two 0-chains are boundaries in their respective spaces; for example, the one in S_1^1 is the boundary of the 1-chain consisting of the simplex $\sigma\colon \Delta^1 \to S_1^1\colon (x, 1-x) \mapsto (\cos \pi x, -\sin \pi x)$.

Another fact is that the sets S_1^n and S_2^n are each homeomorphic to the ball B^n. We can construct a homeomorphism for S_1^n as follows: given $(x_1, \ldots, x_{n+1})$, take $(x_1, \ldots, x_n)$, normalize it so that it lies on the boundary of B^n, and then multiply by the scalar $\frac{2}{3}(x_{n+1} + 1)$. The inverse simply takes the spherical subset of B^n with a distance r from the origin, and maps it back to the spherical cross section $x_{n+1} = \frac{3}{2}r - 1$ of S_1^n in the obvious way. Therefore every k-cycle of S_1^n is a boundary for $k \geq 1$, and because it is homeomorphic to S_1^n by mapping $x_{n+1} \mapsto -x_{n+1}$, the same holds for S_2^n.

We have reduced Proposition 4.4 to the following:

Proposition 4.8. *If there exists a cycle of $C_{n-1}(S_1^n \cap S_2^n)$ which is not a boundary in $S_1^n \cap S_2^n$, then there exists an cycle in $C_n(S^n)$ which is not a boundary in S^n.*

Let $C_n(S_1^n + S_2^n) \subseteq C_n(S^n)$ consist of the chains whose each simplex is either contained in S_1^n or S_2^n. Equivalently, the chains in $C_n(S_1^n + S_2^n)$ are those which can be expressed as the sum of a chain contained in S_1^n and a chain contained in S_2^n. To help us with the proof of Proposition 4.8, we require the following fact, which is proved in full generality in [6], page 199.

Fact 4.1. *If there exists a cycle in $C_n(S_1^n + S_2^n)$ which is not the boundary of a chain in $C_{n+1}(S_1^n + S_2^n)$, then there exists a cycle in $C_n(S^n)$ which is not the boundary of a chain in $C_{n+1}(S^n)$.*

Fact 4.1 is related to the following intuitions: if we are searching for an $(n+1)$-chain $B = \sum \pm\sigma_i$ such that $\partial B = C = \sum \pm\sigma_i'$ for a given cycle $C \in C_n(S_1^n + S_2^n)$, the behavior of the simplex maps of B on the interior of their domain Δ^{n+1} is unimportant. Since we only care about the boundary of B, we only care about how the simplices σ_i of B behave on the faces of Δ^{n+1}. Of course, by continuity, the σ_i will need to send points near $\Delta^{n+1}[j]$ to points in S^n near $\sigma_i(\Delta^{n+1}[j])$. The assumption that C is in $C_n(S_1^n + S_2^n)$, then, means that continuity will not require the values of σ_i near $\Delta^{n+1}[j]$ to intersect both $S_1^n - S_2^n$ and $S_2^n - S_1^n$ for any particular j for which $\sigma_i[j]$ is supposed to be one of the σ'. By using many different σ_i such that most of the faces of the σ_i cancel in the boundary, it makes intuitive sense that we should be able to build B such that $\partial\sigma_i$ are each contained within either S_1^n or S_2^n, as long as any $B \in C_{n+1}(S^n)$ exists to begin with.[4] Now, since the behavior of σ_i on the interior of Δ^{n+1} is unimportant, we should be able to deform or modify σ_i such that its values on the interior of Δ^{n+1} lie in one of the sets S_1^n or S_2^n which also contains $\partial\sigma_i$. More formally, since $\partial\sigma_i$ is a cycle and thus a boundary in either S_1^n or S_2^n, we may replace σ_i with some simplex contained in S_1^n or S_2^n which has the same boundary as σ_i. Hence, if C is the boundary of a chain in $C_{n+1}(S_n)$, it should be the boundary of a chain in $C_{n+1}(S_1^n + S_2^n)$, and by contrapositive, is C is not the boundary of any chain in $C_{n+1}(S_1^n + S_2^n)$, it should not be the boundary of any chain in $C_{n+1}(S^n)$.

We proceed with the proof of Proposition 4.8. Let $C \in C_{n-1}(S_1^n \cap S_2^n)$ be a cycle but not a boundary. In the special case $n = 1$, let C be the chain in (5). Let i and j be the inclusion or identity maps from $S_1^n \cap S_2^n$ into its supersets S_1^n and S_2^n, respectively, and let k and l be the inclusion maps from S_1^n and S_2^n into S^n. These maps are all continuous, meaning $i(C)$ and $j(C)$ are cycles by (4). Since we showed

[4] This is the only step in the paragraph which is difficult to make rigorous.

above that all $(n-1)$-cycles in S_1^n and S_2^n are boundaries for $n > 1$, and that applying i or j to the chain in (5) gives a boundary, we are justified in letting $i(C) = \partial A$ and $j(C) = \partial B$. Here A and B are singular n-chains of S_1^n and S_2^n, respectively. Now, we claim that the chain $D = k(A) - l(B)$ in $C_n(S_1^n + S_2^n)$ is a cycle, but not a boundary of a chain in $C_{n+1}(S_1^n + S_2^n)$. The construction of D is shown in Figure 4.

By (4) we have $\partial D = \partial k(A) - \partial l(B) = k(i(C)) - l(j(C))$, and since $k \circ i$ and $l \circ j$ both represent the inclusion map from $S_1^n \cap S_2^n$ to S^n, these two terms cancel. This makes D a cycle.

In order to show that D is not a boundary of a chain in $C_{n+1}(S_1^n + S_2^n)$, we show the contrapositive: if D is constructed as above and is a boundary of a chain in $C_{n+1}(S_1^n + S_2^n)$, then C is a boundary.

Suppose $D = \partial P$ where $P \in C_{n+1}(S_1^n + S_2^n)$. By the definition of $C_{n+1}(S_1^n + S_2^n)$, we can write P as $k(Q) - l(R)$ for some $Q \in C_{n+1}(S_1)$ and $R \in C_{n+1}(S_2)$, and we have that $D = \partial k(Q) - \partial l(R) = k(\partial Q) - l(\partial R)$. But D also equals $k(A) - l(B)$, so $k(A - \partial Q) = l(B - \partial R)$. This $k(A - \partial Q)$ is in the image of both k and l, so $A - \partial Q$ is contained in $S_1^n \cap S_2^n$ and we can let $A - \partial Q = i(Y)$ for some $Y \in C_n(S_1^n \cap S_2^n)$. Finally, $\partial Y = C$, since $i(\partial Y) = \partial(A - \partial Q) = \partial A = i(C)$ and i is one-to-one. (We have used the fact that $\partial(\partial Q) = 0$; see Proposition 4.1.)

We have now shown that D is a cycle in $C_n(S_1^n + S_2^n)$ which is not the boundary of any chain in $C_{n+1}(S_1^n + S_2^n)$. By Fact 4.1, there exists a cycle in $C_n(S^n)$ which is not a boundary, thus completing the proof of Proposition 4.8. Proposition 4.8 completes the proof of Proposition 4.4, which in turn completes our homological proof of Brouwer's theorem.

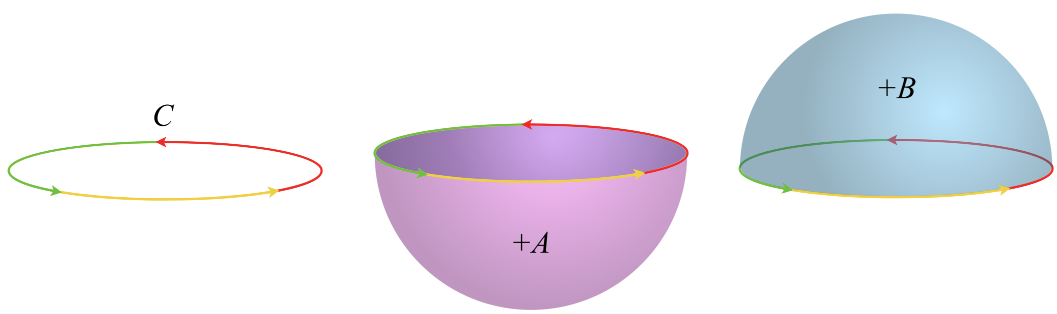

Figure 4: The chain C consists of three 1-simplices of the space $S_1^n \cap S_2^n$, which is not shown. The chains A of S_1^n and B of S_2^n each essentially have C as their boundary, and so their difference D in S^n has an empty boundary.

5 Analysis

First, it should be noted that our combinatorial proof of the No Retraction theorem in section 3 was nonstandard. Our proof was taken from [7], in which it is compared with the more traditional proof. The traditional proof also relies on Sperner's Lemma, and the difference is mostly one of notation. In any case, the most interesting comparisons are between the combinatorial and homological proofs.

One clear similarity between the two proofs we have presented is that both use the No Retraction theorem as an intermediary. More than Brouwer's theorem itself, the No Retraction theorem is exposed to attacks via topological methods, since it relates to maps between B^n and S^{n-1} which are fundamentally different topological objects. The particular connection I would like to discuss, however, is as follows: any continuous retraction r from B^n to S^{n-1} gives rise to a homotopy between the identity map on S^{n-1} and a map sending all of S^{n-1} to a single point. One such homotopy is given by

$$H(\mathbf{x}, t) = r\big((1-t)\mathbf{x}\big), \tag{6}$$

where $t\mathbf{x}$ is interpreted as a point in B^n. We see that $H(\mathbf{x}, 0)$ is constantly equal to $r(0)$, while $H(\mathbf{x}, 1) = r(\mathbf{x}) = \mathbf{x}$ on S^{n-1}. As such, a continuous retraction from B^n to S^{n-1} gives a way of continuously deforming $S^{n-1} = \partial B^n$ into a point, without it ever intersecting the interior of B^n. In some fashion or another, both the combinatorial and homological proofs of the No Retraction theorem can be related to this homotopy.

At the start of section 4.2, we gave a proof of the No Retraction theorem using the language of homology classes. This was done partly to emphasize the importance of homology classes. In particular, homology theory is more properly developed using the tools of *group theory*,[5] which makes the proof using homology classes more convenient than others. However, we do not even need homology classes for this proof:

Proof of Theorem 2.4. Let $i\colon S^{n-1} \to B^n$ the inclusion map, and let C be a cycle of S^{n-1} which is not a boundary. We know that $i(C)$ is a boundary in B^n, so let $i(C) = \partial B$. There cannot exist a continuous retraction $r\colon B^n \to S^{n-1}$, since C would equal $r(i(C)) = \partial r(B)$. □

[5]The sets $C_n(X)$ are thought of as groups, and the boundary operators are thought of as *homomorphisms* between them. The cycles form a *subgroup* of $C_n(X)$, and the boundaries form a subgroup of the cycle group. The homology classes $H_n(X)$ also form a group, and in fact $H_n(X)$ is constructed by taking the *quotient group* of the cycles with respect to the boundaries. Since it was not necessary for the proof of Brouwer's theorem, we ignored this group structure.

This follows the same line of reasoning as the version with homology classes. As a common thread, most proofs of the theorem by topological methods rely on some *invariant*, or property which is known to be preserved under suitable maps and transformations. In the proof above, we show that the property "every cycle is a boundary" is invariant under continuous retraction. There are many other important types of invariants in topology and throughout mathematics; for example, Proposition 1.1 states that the compactness of a topological space is invariant under homeomorphism. We now present a third homological proof of the No Retraction theorem which makes the use of such an invariant more explicit. In particular, Proposition 4.7 states that the homology classes of a space are invariant under homotopy equivalence.

Proof of Theorem 2.4. If the identity in S^{n-1} and the map to a single point $a = r(0) \in S^{n-1}$ are homotopic as in (6), then S^{n-1} is homotopy equivalent to the topological space $\{ a \}$. To see this, let f be the map from S^{n-1} to $\{ a \}$, and let g be the inclusion map which sends $a \in \{ a \}$ to $a \in S^{n-1}$; then $f \circ g$ is trivially the identity map in $\{ a \}$, and $g \circ f$ is homotopic to the identity in S^{n-1} by (6). Hence the homology classes of S^{n-1} and $\{ a \}$ should be in one-to-one correspondence, which is a contradiction by Propositions 4.2 and 4.4. □

This proof nicely represents the idea of both homological proofs we have given: show that two spaces have different properties, and show that those properties must be invariant under some map related to the hypothetical continuous retraction.

Let C be a cycle of S^{n-1} which is not a boundary, and consider the intermediary functions $r_t(\mathbf{x}) = H(\mathbf{x}, t)$ of of the homotopy (6). Clearly $r_0(C) = r(C) = C$ is a cycle and not a boundary, but the same is not true of $r_1(C)$, whose simplices map to single points. Our homological proof amounts to invoking the properties of homotopies: $r_0(C)$ and $r_1(C)$ must be homologous by Proposition 4.6. Ultimately, Proposition 4.6 is derived from nothing more than the properties of continuous functions, most importantly the fact that compositions of continuous functions are continuous.

By contrast, our combinatorial proof initially accepts the existence of such a homotopy. Rather than using an invariant property between r_0 and r_1 to rule out its existence, we explore the behavior of r_t between $t = 0$ and $t = 1$. Using Sperner's Lemma, we showed a contradiction as long as the retraction r_0 maps neighborhoods smaller than some δ to some $\partial\Delta^n \setminus \Delta^n [j]$. Equivalently, we showed that in order for the homotopy H to exist, it must map arbitrarily small sets of $(\mathbf{x}, t)$ to the entirety of $\partial\Delta^n$ (for now we substitute the sphere with the simplex and regard H as a homotopy from a point to $\partial\Delta^n$). In order to complete the proof, we use Lebesgue' number lemma and the properties of continuous functions, particularly that the preimages of the open sets $\partial\Delta^n \setminus \Delta^n[j]$ are open.

References

[1] Brouwer, L. E. J. *Über Abbildung von Mannigfaltigkeiten.* Mathematische Annalen, Vol. 71, pp. 97–115, 1911.

[2] Choe, Geon H. *A Topological Proof of the Perron-Frobenius Theorem.* Comm. Korean Mathematical Society, Vol. 9, pp. 565–570, 1994.

[3] Crossley, Martin D. *Essential Topology.* Springer, 2005.

[4] Fort, M. K., Jr. *Some Properties of Continuous Functions.* The American Mathematical Monthly, Vol. 59, pp. 372–375, 1952.

[5] Hadamard, Jacques. *Sur quelques applications de l'indice de Kronecker.* Introduction à la théorie des fonctions d'une variable, ed. J. Tannery. Hermann, pp. 875–915, 1910.

[6] Hatcher, Allen. *Algebraic Topology.* Cambridge University Press, 2002. http://www.math.cornell.edu/~hatcher/AT/AT.pdf.

[7] Ivanov, Nikolai V. *Sperner's Lemma, the Brouwer Fixed-Point Theorem, and Cohomology.* arXiv:0906.5193, 2009.

[8] Jarvis, Tyler and James Tanton. *The Hairy Ball Theorem via Sperner's Lemma.* The American Mathematical Monthly, Vol. 111, pp. 599–603, 2004.

[9] Maehara, Ryuji. *The Jordan Curve Theorem Via the Brouwer Fixed Point Theorem.* The American Mathematical Monthly, Vol. 91, pp. 641–643, 1984.

[10] Nash, John. *Non-Cooperative Games.* Annals of Mathematics, Vol. 54, pp. 286–295, 1951.

[11] Weiss, Benjamin. *A Combinatorial Proof of the Borsuk-Ulam Antipodal Point Theorem.* Israel Journal of Mathematics, Vol. 66, pp. 364–368, 1989.