Sperner's Lemma and Homeomorphisms

Rich Schwartz

March 19, 2024

1 Introduction

The purpose of these notes is to give a self-contained proof that $\boldsymbol{R}^m$ and $\boldsymbol{R}^n$ are not homeomorphic when $m > n$. The proof bypasses homology entirely, and uses Sperner's Lemma instead. The main thing you have to know about is higher dimensional triangulations both of a simplex and of a sphere. In the spherical case, the simplices in the triangulation are made of pieces of lower-dimensional spheres.

2 Sperner's Lemma

2.1 Statement of Result

Sperner's Lemma works in any dimension but I'm going to concentrate on the 2-dimensional case. I'll indicate how to generalize as I go along.

An example of a *cleanly triangulated triangle* is shown on the left hand side of Figure 1.

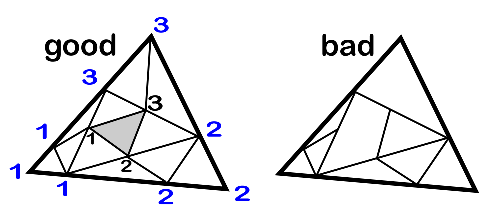

Figure 1: Two good triangulations and a bad one in action.

The drawing on the right shows the kind of junk we want to avoid. The main things you want from the triangulation are that the small triangles meet edge to edge and pairwise have disjoint interiors. In general, a cleanly divided n-simplex would be a subdivision of an n-simplex into smaller n-simplices which meet face to face.

Let T be the triangle, and suppose it is cleanly triangulated into n smaller triangles $T_1, ..., T_n$. Suppose also that the vertices of the triangulation are labeled by integers $\{1, 2, 3\}$ so that the kth side of T has no k label. Sperner's Lemma says that some T_j gets all 3 labels. In general you would have an n-dimensional simplex cleanly subdivided into smaller n-dimensional simplices, and the labeling on the boundary would be such that the kth face gets no k labels. Sperner's Lemma would then say that some sub-simple gets every label.

2.2 The Genius Proof

Here is the reference for the amazing proof to follow.

A. McLennan and R. Tourkey, *Using volume to prove Sperner's lemma* Econ. Theory **35** (2008) pp 593-597

We're going to assume that we have a labeling in which no triangle sees all three labels and derive a contradiction. We normalize so that T has area 1. Let's call the vertices of the triangulation $V_1, ..., V_m$. Let $L_1, ..., L_m$ be the labels of these vertices. Let $W(1)$, $W(2)$, $W(3)$ be the vertices of the big triangle T. For each $t \in [0, 1]$, define the new point

$$V_k(t) = (1 - t)V_k + tW(L_k). \tag{1}$$

In other words, we think of t as time and we think of a the curve $t \to V_k(t)$ as a path which starts at V_k and ends at $W(L_k)$ and moves in a straight line at constant speed. Supposing that the triangle T_k has vertices A, B, C, let $T_k(t)$ be the triangle with vertices $A(t), B(t), C(t)$. As the points move with t these triangles change shape.

Consider the function

$$f(t) = \sum_{j=1}^{n} \text{area}(T_j(t)). \tag{2}$$

First of all, this function is a polynomial, thanks to the simple formulas – e.g. determinants – one can use to compute the areas of the triangles. Second of all, $f(t) = 1$ for all t sufficiently close to 0. The point is that, for small t, we still have a triangulation even though the points have moved a little. Since f is a polynomial, we must have $f(t) = 1$ for all $t \in [0, 1]$. However, suppose that no triangle T_k sees all three labels. Then $T_k(t)$ converges either a single vertex or to an edge of T as $t \to 1$. Hence

$$\lim_{t \to 1} \text{area}((T_k(t)) = 0. \tag{3}$$

Since this is true for all triangles, and there are only finitely many of them in the triangulation, we see that $f(t) \to 0$ as $t \to 1$. This is a contradiction. That's the end of the proof. In the general case you'd use n-volume rather than area.

2.3 The Traditional Proof

Say that a *flag* is a pair (T_k, e), where T_k is one of the triangles of the triangulation and e is an edge of T_k. Each triangle participates in 3 flags and each edge either participates in 1 or 2 flags, depending on whether the edge is in the boundary of T. Say that $(1, 2)$-flag is a flag which has the labels 1 and 2 on its edge. We're going to count the $(1, 2)$-flags in two ways.

Let's first count these flags going edge by edge. Each interior edge contributes an even number of $(1, 2)$-flags to the total, because it participates in two flags and these two flags are either simultaneously $(1, 2)$-flags or not $(1, 2)$-flags. The only boundary edges which contribute a $(1, 2)$-flag are the ones on the side which has the 1 and 2 labels. This side is divided into finitely many edges. One endpoint of the side is labeled 1 and the other one is labeled 2. So, as we go from one endpoint to the other, we have to switch labels an odd number of times. Hence, there are an odd number of $(1, 2)$-flags coming from the boundary edges. But that means there is an odd number of $(1, 2)$-flags overall.

On the other hand, let's count the $(1, 2)$-flags triangle by triangle. As I mentioned above, each triangle participates in 3 flags. Just list out the possibilities and you can see that a triangle contributes an odd number of $(1, 2)$-flags to the count if and only if it gets all 3 labels. The triangles labeled $(1, 2, 2)$ $(1, 1, 2)$ each contribute two $(1, 2)$ flags to the count and the rest contribute zero. Hence, there must be an odd number of triangles which are labeled $(1, 2, 3)$. That completes the proof.

In general, you would count $(1, 2, ..., n-1)$-flags in two ways and use induction on the dimension.

3 Applications

3.1 A Topological Interlude

Before going further, we need a foundational result from topology.

Lemma 3.1 (Uniform Continuity) *Let Δ be a simplex in $\boldsymbol{R}^n$ and suppose $f : \Delta \to \boldsymbol{R}^n$ is a continuous map. There is some $\delta > 0$ so that $\|f(a) - f(b)\| < 1$ as long as $\|a - b\| < \delta$.*

Proof: Suppose not. Then we can find a two sequence $\{a_n\}$ and $\{b_n\}$ such that $\|a_n - b_n\| \to 0$ and $\|f(a_n) - f(b_n)\| \geq 1$. But we can pass to a subsequence so that $a_n \to p$. Evidently $b_n \to p$ as well. But then our conditions violate the continuity of f at p. ♠

3.2 No Retractions

Here is how to use Sperner's Lemma to prove a classic result from algebraic topology. Usually you use homology for this.

Theorem 3.2 *Let B be the unit ball in $\boldsymbol{R}^n$. There is no map $f : B \to \partial B$ which is the identity on ∂B.*

Proof: I'll give the 2 dimensional proof. You can think about how to generalize the proof to higher dimensions. Suppose f exists. We choose a homeomorphism h from B to Δ, the equilateral triangle having side length 100. It suffices to prove result for the map $h \circ f \circ h^{-1}$. In other words, we can replace the disk B by the big triangle Δ.

We will suppose that f exists and derive a contradiction. By the Uniform Continuity Lemma, we can find some δ so that $\|f(a) - f(b)\| < 1$ for all a, b in the triangle with $\|a - b\| < \delta$. Choose a triangulation of Δ where each triangle in the triangulation has side length less than δ.

Label a vertex v of the triangulation by the name of the vertex of Δ closest to $f(k)$. In case of a tie, choose the lower label. f maps the vertices of each small triangle so that they are each within 1 of each other. But then no triangle can be labeled $(1, 2, 3)$ because the three vertices of Δ are spread far apart. On the boundary, the labeling is more or less as in Figure 1. So, this labeling contradicts Sperner's Lemma. ♠

Corollary 3.3 *Let B be the unit ball in $\boldsymbol{R}^n$ and suppose $I : \partial B \to \boldsymbol{R}^n$ is the identity map. There is no way to extend I to a map $\Psi : B \to \boldsymbol{R}^n - O$, where O is the origin.*

Proof: Let $\rho : \boldsymbol{R}^n - O \to \partial B$ be the radial retraction map. In polar coordinates we have $\rho(r,\theta) = (1,\theta)$. This map is continuous. If Ψ exists then the map $\rho \circ \Psi$ would give us a continuous map from $B \to \partial B$ which is the identity on ∂B. This contradicts the previous result. ♠

4 Continuous Extension

4.1 The End of the Proof

Again let B be the unit ball in $\boldsymbol{R}^n$. Suppose now that $f_0 : \partial B \to \boldsymbol{R}^m - O$ with $m > n$. Below, we construct a continuous extension $F : B \to \boldsymbol{R}^m - O$.

Suppose $h : \boldsymbol{R}^n \to \boldsymbol{R}^m$ is a homeomorphism, with $m > n$. We normalize so that $h(O) = O$. Then h also gives a homeomorphism from $\boldsymbol{R}^n - O$ to $\boldsymbol{R}^m - O$. Let $I : \partial B \to \boldsymbol{R}^n$ be the identity map. Let $f_0 = h \circ I$. Let $F : B \to \boldsymbol{R}^m - O$ be the extension map. Then $\Psi = h^{-1} \circ F : B \to \boldsymbol{R}^n - O$ is a continuous extension of the identity map. This contradicts the corollary above. Hence h does not exist.

Now we construct the extension.

4.2 Linear Interpolation

Let B be the unit ball in $\boldsymbol{R}^n$ and suppose $f_0, f_1 : \partial B \to \boldsymbol{R}^m$ are 2 continuous maps. Here m might equal n or it might not. We define

$$f_t(p) = (1-t)f_0(p) + tf_1(p). \tag{4}$$

The map f_t is a continuous interpolation from f_0 to f_1. If $f_0(\partial B)$ and $f_1(\partial B)$ both miss the ball of radius 10 about O, and if the distance from $f_0(p)$ to $f_1(p)$ is at most 1, then $f_t(p) \in \boldsymbol{R}^m - O$ for all t. This is a criterion we apply to all $p \in \partial B$.

4.3 Simplicial Approximation

Now suppose we have $f_0 : \partial B \to \boldsymbol{R}^m - O$. We scale the picture so that the ball of radius 11 is disjoint from $f_0(\partial B)$. Next, we pick a very fine triangulation on ∂B. Here is a recipe for the map f_1:

1. We let $f_1 = f_0$ on the vertices of the triangulation.
2. Given a simplex τ of the triangulation, let $v(\tau)$ be the vertex set of τ.
3. Let Π_τ denote projection onto the convex hull of $f_0(v(\tau))$.
4. The restriction of f_1 to τ equals $\Pi_\tau \circ f_0$.

The map f_1 is continuous because the definitions agree on adjacent simplices in the triangulation. If the triangulation is fine enough, then $f_1(\partial B)$ misses the ball of radius 10 about the origin and $\|f_0(p) - f_1(p)\| < 1$ for all $p \in \partial B$. Hence the interpolation f_t maps ∂B into $\boldsymbol{R}^m - O$ for all $t \in [0, 1]$.

Here is the beauty of the map f_1. The image $f_1(\partial B)$ is contained in a finite union of $(n-1)$-dimensional simplices in $\boldsymbol{R}^m - O$. All these simpices are outside the unit ball.

4.4 Dimensionality

If we have an $(n-1)$-simplex $\Delta \subset \boldsymbol{R}^m$ and some point $q \in \boldsymbol{R}^m$ then we let (Δ, q) denote the n-simplex we get by just adding q as a vertex. Equivalently, (Δ, q) is the convex hull of $\Delta \cup q$. Here is the punchline: For almost every choice of $q \in \boldsymbol{R}^m$ the simplex (Δ, q) is disjoint from O. This uses $m > n$.

If we have a finite union of $(n-1)$-simplices $\Delta_1, ..., \Delta_k$ then we can find $q \in \boldsymbol{R}^m$ such that $(\Delta_i, q) \in \boldsymbol{R}^m - O$ for all $i = 1, ..., k$. In particular, we can apply this to the simplices comprising the image $f_1(\partial B)$. We choose such a q and we define $f_2 : \partial B \to \boldsymbol{R}^m - O$ to be the constant map $f_2(\partial B) = q$. By construction, the interpolating map f_t, for $t \in [1, 2]$, maps ∂B into $\boldsymbol{R}^m - O$.

4.5 The Final Extension

We put polar coordinates (r, θ) on our ball B. Here $r \in [0, 1]$ and θ is a parameter for ∂B. Define $F(r, \theta) = f_{2-2r}(\theta)$. We have $F(1, \theta) = f_0(\theta)$. So, our map F is a continuous extension of f_0 to B. Also, by construction $F(B) \subset \boldsymbol{R}^m - O$.

We're done.