The Power of Sperner's Lemma

Richard Evan Schwartz*

December 1, 2025

1 Introduction

Sperner's Lemma is a powerful combinatorial tool that can be used to give easy proofs of key topological results that otherwise would rely on algebraic topology or real analysis. In this note I'll use Sperner's Lemma to prove a number of topological results, including the Brouwer Fixed Point Theorem, the Inverse Function Theorem and the result that $\boldsymbol{R}^m$ and $\boldsymbol{R}^n$ are not homeomorphic when $m > n$. These notes assume that you are familiar with cubes, simplices, and Euclidean space of all dimensions.

2 Triangulations

The *cone* of a simplex τ to a point p is the new simplex $\langle \tau, p\rangle$ having all the original vertices of τ, and additionally p, as vertices. We will only use this construction when the dimension $\langle \tau, p\rangle$ is one more than the dimension of τ.

Let P be either a cube or a simplex. The *barycenter* of P is the average of its vertices. We can subdivide P into smaller simplices in an inductive way. Assuming we have already subdivided each boundary face of P into simplices we can cone all of these simplices to the barycenter of P. This construction starts by cutting all the edges of P in half, and then moves up through the dimensions. The process is called *barycentric subdivision*.

*Supported by N.S.F. grant DMS-2505381.
Department of Mathematics, Brown University, 151 Thayer St., Providence, RI, 02912.
Email1: Richard.Evan.Schwartz@gmail.com
Email2: res@math.brown.edu

Given any $\epsilon > 0$ we can iterate the barycentric subdivision to produce a decomposition of P into simplices which all have diameter less than ϵ. Figure 1 shows the second iterate of barycentric subdivision applied to an equilateral triangle. These are all examples of *triangulations* of P. The simplices meet "face to face" in a triangulation.

3 Sperner's Lemma

Suppose we have a big triangle T that is triangulated into small 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. Figure 1 illustrates this.

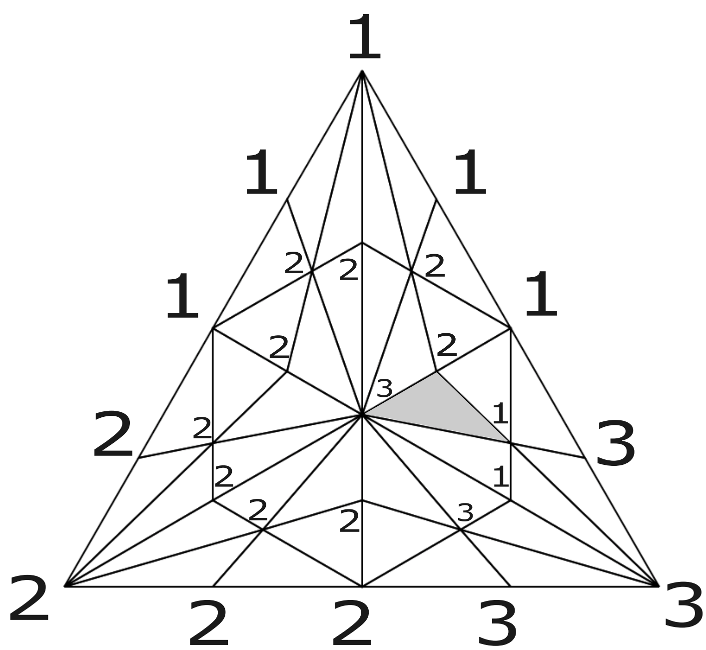

Figure 1: Sperner's Lemma in action

Sperner's Lemma works in any dimension. In general, we have a big n-dimensional simplex triangulated by smaller simplices. The labeling on the boundary is such that that the kth face gets no k labels. Sperner's Lemma says that some small simplex gets every label.

If you don't want to think about general triangulations of a simplex, you can always take the triangulations you get from iterated barycentric subdivision, as shown in Figure 2. These are sufficient for all the applications.

3.1 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 will give the argument in the 2-dimensional case and at the end explain how to generalize it. 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 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 to 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 general you'd use n-volume rather than area.

3.2 A Traditional Proof

Sperner's Lemma has a number of other proofs. Here I will explain a more traditional proof of the result. Let's start with the 2-dimensional case.

Let $T_1, ..., T_n$ be the triangulation of T, as above. 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.

Counting by Edges: 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.

Counting By Triangles: 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 the 3-dimensional case, the (123)-face of the big tetrahedron gives a 2-dimensional instance of Sperner's Lemma. Call this the *boundary instance*. Counting (123)-flags in the tetrahedron by faces shows that the number of these flags has the same parity as the number of (123)-triangles in the boundary instance. So, in the 3-dimensional case, the first count gives an odd number. The second way of counting then shows that there is an odd number of (1234)-tetrahedra.

In this way, the general case is proved by induction on the dimension.

4 The No Retraction Theorem

The main result in this section says, intuitively, that you cannot map a ball into its boundary if you keep the boundary fixed. Usually you prove this result using homology theory. We begin with a technical lemma and then prove the main result.

Lemma 4.1 (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 two sequences $\{a_n\}$ 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. ♠

Theorem 4.2 (No Retraction) *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 $f(v)$. 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 an exact generalization of what you see in Figure 1.) So, this labeling contradicts Sperner's Lemma. ♠

Before moving on, let me give several of the main consequences of the No Retraction Theorem.

Corollary 4.3 (Surjective Image) *Let B be the unit ball in $\boldsymbol{R}^n$ and suppose $I : \partial B \to \boldsymbol{R}^n$ is the identity map. Suppose $p \in \boldsymbol{R}^n$ lies in the interior of B. There is no way to extend I to a continuous map $\Psi : B \to \boldsymbol{R}^n - \{p\}$.*

Proof: There is a continuous map ρ_p from $\boldsymbol{R}^n$ to ∂B, which is the identity on ∂B. We pick any $q \in \boldsymbol{R}^n$ and define $\rho_p(q) = \overrightarrow{pq} \cap \partial B$. Here $\overrightarrow{pq}$ is the ray starting at p and going through q. The map ρ_p is continuous on $\boldsymbol{R}^n - \{p\}$. If our map Ψ exists then $\rho_p \circ \Psi$ violates the No Retraction Theorem. ♠

Here is the famous Brouwer Fixed Point Theorem. This is the classic reduction of the Fixed Point Theorem to the No Retraction Theorem.

Corollary 4.4 (Brouwer Fixed Point) *Let B be the unit ball in $\boldsymbol{R}^n$. Any continuous map $g : B \to B$ has a fixed point.*

Proof: Suppose there is no $p \in B$ such that $g(p) = p$. Let ρ_p be the map defined in the previous proof. Define

$$f(p) = \rho_{g(p)}(p).$$

That is, $f(p)$ is the point where the ray from $g(p)$ through p intersects ∂B. By construction $f : B \to \partial B$ is continuous and restricts to the identity on ∂B. This contradicts the No Retraction Theorem. ♠

5 The No Homeomorphism Theorem

In this section we prove that there is no homeomorphism $h : \boldsymbol{R}^n \to \boldsymbol{R}^m$ if $m > n$. We first introduce some notation and terminology.

- Let O denote the origin in $\boldsymbol{R}^n$ and in $\boldsymbol{R}^m$.
- Let Q denote the unit cube in $\boldsymbol{R}^n$ centered at the origin O.
- Let B_r denote the ball of radius r in $\boldsymbol{R}^m$ centered at O.
- An *affine map* on a simplex is the composition of a linear map and a translation. A *piecewise affine map* on a cube (or on its boundary) is a continuous map which is affine when restricted to each simplex of a triangulation.

Lemma 5.1 *There is no map $G : Q \to \boldsymbol{R}^n - O$ which is the identity on ∂Q.*

Proof: We can find a homeomorphism $q : \boldsymbol{R}^n \to \boldsymbol{R}^n$ which maps Q to B and maps O to O. This map preserves rays through the origin, is linear on each ray, and scales the distances so as to map Q to B. Using q, we convert this lemma to a special case of the Surjective Image Theorem. ♠

Lemma 5.2 (Extension) *If $m > n$ and $f : \partial Q \to \boldsymbol{R}^m - O$ is continuous, then there is a continuous map $F : Q \to \boldsymbol{R}^m - O$ such that $F = f$ on ∂Q.*

Before proving the Extension Lemma, let us use it to finish the proof of the main result. Suppose $h : \boldsymbol{R}^n \to \boldsymbol{R}^m$ is a homeomorphism and $m > n$. We can normalize by translating to arrange that $h(O) = O$. Let $f : \partial Q \to \boldsymbol{R}^n$ be the identity map. By the Extension Lemma, the map $h \circ f$ can be extended to a map $F : Q \to \boldsymbol{R}^m - O$. But then $G = h^{-1} \circ F : Q \to \boldsymbol{R}^n - O$ is a continuous map which is the identity on ∂Q. This contradicts Lemma 5.1.

Now we prove the Extension Lemma. We begin with a preliminary lemma which is similar in spirit to the Uniform Continuity Lemma.

Lemma 5.3 *Suppose $f : \partial Q \to \boldsymbol{R}^m - O$ is continuous. Then there is some $\epsilon > 0$ so that $B_\epsilon \cap f(\partial Q) = \emptyset$.*

Proof: If this is false then we can find a sequence $\{p_n\}$ in ∂Q such that $f(p_n) \to O$. But then we can find a subsequence which converges to $p \in \partial Q$ and we would have $f(p) = O$. This is is a contradiction. ♠

By Lemma 5.3, there is some $\epsilon > 0$ so that $f(\partial Q) \cap B_\epsilon = \emptyset$. By scaling, we can assume that $\epsilon = 10$. Strategically, we are scaling so that $f(\partial Q)$ misses a big ball around the origin. This gives us some wiggle room.

Strategy: Our goal is to extend f to all of Q. We make our extension in two stages. Let $Q' \subset Q$ be the open cube centered at the origin in $\boldsymbol{R}^n$ which has half the diameter of Q. We first extend f from ∂Q to a map $F : Q - Q' \to \boldsymbol{R}^m - O$. The extension has the property that F is piecewise affine on $\partial Q'$. We then use a coning trick to extend F to Q'. At some moment in the second stage we will see where we use the hypothesis that $m > n$.

Stage 1: By the same argument as in the Uniform Continuity Lemma, there is a triangulation of ∂Q so that the diameter of $f(\tau)$ is less than 1 for each simplex τ in the triangulation. Let τ' be the scaled copy of τ on Q'. We define $f' : \tau' \to \boldsymbol{R}^m$ to be the affine map so that $f'(v') = f(v)$ for each vertex v of τ. By construction f' agrees on any intersection of the form $\tau_1' \cap \tau_2'$ for any pair of simplices τ_1, τ_2 which have a nontrivial intersection. For this reason, $f' : Q' \to \boldsymbol{R}^m$ is continuous and piecewise affine. Also $f'(Q') \cap B_9 = \emptyset$.

Let ρ be any ray through the origin. The intersection $\rho \cap (Q - Q')$ is a line segment, and we have already defined F on its endpoints. We extend F to be affine on the whole segment. Doing this for every ray we get $F : Q - Q' \to \boldsymbol{R}^m - O$ which is piecewise affine on $\partial Q'$.

Stage 2: The image $F(\partial Q')$ is contained in a finite union of n-dimensional subspaces. Since $m > n$ we can choose $q \in \boldsymbol{R}^m$ outside this union. But then, for every simplex τ in the triangulation of Q, the coned n-dimensional simplex $\langle f'(\tau'), q\rangle$ is disjoint from O.

Note that Q' is partitioned into n-dimensional simplices of the form $\langle \tau', O\rangle$. Define F to be affine on each of these n-simplices so that

$$F(\langle \tau', O\rangle) = \langle f'(\tau'), q\rangle.$$

This extends F to Q' so that $F(Q') \subset \boldsymbol{R}^m - O$. But then $F(Q) \subset \boldsymbol{R}^m - O$. ♠

6 The Inverse Function Theorem

Suppose now that $F : \boldsymbol{R}^n \to \boldsymbol{R}^n$ is a smooth map. This means that partial derivatives of F exist of all orders and also, at each point $p \in \boldsymbol{R}^n$,

$$\lim_{t \to 0} \frac{F(p + th) - F(p)}{t} = dF|_p(h). \tag{4}$$

Here $dF|_p$ is the usual matrix of first partial derivatives of the component functions of F, evaluated at p.

The Inverse Function Theorem says that F is a local diffeomorphism in an open neighborhood of each point where dF is an invertible matrix. The part that Sperner's Lemma helps out with – *via* the No Retraction Theorem – is the open mapping condition. The rest of the argument I give is pretty standard, but I will try to do it cleanly.

Lemma 6.1 *Suppose $dF|_p$ is invertible. Then F is injective on a ball centered at p. Also, for all sufficiently small balls B centered at p, the image $F(B)$ contains a ball about $F(p)$.*

Proof: We normalize so that $p = O$ and $F(O) = O$. Composing by an invertible linear transformation we can assume that $dF|_O$ is the identity map. Given a ball B centered at the origin let rB be the r-scaled copy.

Since dF is continuous, all sufficiently small balls B centered at the origin have the following property: For any unit vector V and any $p \in B$, we have $\angle(V, dF_p(V)) < 10^{-100}$ and $\|dF_p(V)\| > 1 - 10^{-100}$. So, F maps line segments of length λ in B to almost straight curves which have speed almost λ. So, F is injective on B, and $F(\partial B)$ lies outside $\frac{3}{4}B$.

Define $G : 2B - B \to \boldsymbol{R}^n$ as follows: For each ray ρ through the origin, let $\zeta = \rho \cap \partial B$ and let the restriction of G to $\rho \cap (2B - B)$ be the linear interpolation between $F(\zeta)$ and 2ζ. Since $F(\zeta)$ and ζ are very close, the image $G(\rho \cap (2B - B))$ avoids $\frac{1}{2}B$. Hence

$$G(2B - B) \cap \frac{1}{2}B = \emptyset. \tag{5}$$

The union map $F \cup G : 2B \to \boldsymbol{R}^n$ is continuous, and the restriction of $F \cup G$ to $\partial(2B)$ is the identity map. By the Surjectivity Lemma (and scaling), $\frac{1}{2}B \subset F(B) \cup G(2B - B)$. By Equation 5, we get $\frac{1}{2}B \subset F(B)$. ♠

Choose some point p where dF_p is invertible. By the previous lemma we can find a small open ball U centered at p such that $F : U \to \boldsymbol{R}^n$ is injective. By continuity we can assume that dF is invertible at every point of U. Let $V = F(U)$.

Lemma 6.2 *F is a homeomorphism from U to V.*

Proof: We first show that F maps open subsets $U' \subset U$ open subsets of V. Choose some $q = F(p) \in V' = F(U')$. Since U' is open, U' contains all sufficiently small balls centered at p. But then, by the previous lemma, V' contains a ball centered at q. Hence V' is open.

The map $F : U \to V$ is a continuous bijection which maps open sets to open sets. But then F^{-1} is also continuous. Hence F is a homeomorphism from U to $F(U)$. ♠

Now we deal with differentiability. We work entirely inside the set U. We first show that F^{-1} is differentiable at each point $p \in V$. We translate so that $p = O$ and $F(O) = O$ and so that $dF|_O$ is the identity.

Here is a geometric interpretation of Equation 4. Let $E_r : \boldsymbol{R}^n \to \boldsymbol{R}^n$ be the map which dilates distances by r. To say that $dF|_O$ is the identity is to say that the composition $E_r \circ F \circ E_{1/r}$ converges uniformly on the unit ball to the identity map as $r \to \infty$. But then

$$E_r \circ F^{-1} \circ E_{1/r} = (E_r \circ F \circ E_{1/r})^{-1}$$

also converges uniformly on the unit ball to the identity. Interpreting this convergence analytically, this means that $d(F^{-1})|_O$ exists and is also the identity.

Unravelling our normalization, we have really proved that

$$d(F^{-1}) = (dF)^{-1} \circ F^{-1} \tag{6}$$

on V. You may recognize this as the Chain Rule applied to the composition Identity $= F^{-1} \circ F$ and then suitably rearranged. Now we give the usual bootstrap argument. We know that $(dF)^{-1}$ is a smooth function and F^{-1} is once differentiable. So, the composition is once differentiable. This means that $d(F^{-1})$ is once differentiable. Hence F^{-1} is twice differentiable. Running the argument again, we see that $d(F^{-1})$ is twice differentiable and so F^{-1} is three times differentiable. And so on. Hence F^{-1} is also smooth. This completes the proof.

Generative AI Statement: I used ChatGPT 5.1 to help spot typos in this note and also to check for correctness. I did not accept any writing suggestions, however.