THE BORSUK-ULAM AND HAM SANDWICH THEOREMS

BRIAN LIBGOBER

Abstract. In this paper I describe the way one might begin proving the Borsuk-Ulam theorem using measure theory and what remains to be done for such a proof. I then provide a proof of Borsuk-Ulam using graph theory and use the Borsuk-Ulam theorem to prove the Ham Sandwich theorem.

Contents

1. Introduction

The Ham Sandwich Theorem states something like the following ordinary language proposition: Take a sandwich made of a slice of ham and two slices of bread. No matter where one places the pieces of the sandwich in the kitchen, or house, or universe, so long as one's knife is long enough one can cut all three pieces in half in only one pass. The precise mathematical statement of the theorem, generalized to n dimensions, is that given n compact sets in $\mathbb{R}^n$ there is a hyperplane which bisects each compact set so that the two halves of both sets have equal measure. Somewhat surprisingly, since the statement at first glance appears to result from a wedding of plane geometry and measure theory, the proof is actually an easy consequence of the Borsuk-Ulam Theorem, a theorem which imposes certain requirements upon continuous maps from the n-sphere in $\mathbb{R}^{n+1}$ to $\mathbb{R}^n$. As it turns out, the Borsuk Ulam theorem is a far deeper theorem in that it can be recapitulated in equivalent form with rules governing various other types of continuous maps. In my paper I will discuss two proofs of the theorem which draw on very different ideas, one from combinatorics and the other from measure theory. Having established the Borsuk-Ulam Theorem, I will prove the Ham Sandwich Theorem.

2. A Theorem of Many Monikers

One of the most common variations on the Borsuk-Ulam theorem is the claim that:

Theorem 2.1. *Borsuk-Ulam. If $n \geq 0$ then for any continuous mapping $f : S^n \to \mathbb{R}^n$ there is a point $x \in S^n$ for which $f(x) = f(-x)$*

Date: DEADLINE AUGUST 22, 2008.

But another common take on the theorem is as follows.

Theorem 2.2. *Borsuk-Ulam. For every continuous mapping $f : S^n \to \mathbb{R}^n$ that is antipodal there is a point $x \in \mathbb{S}^n$ for which $f(x) = 0$, where an antipodal map is understood to be a map such that for all $x \in S^n$, $f(-x) = -f(x)$.*

To show that these two are equivalent we present the following proof:

Proposition 2.3. *Theorems 1.1 and 1.2 are equivalent.*

Proof. *Theorem 1.1* $\implies$ *Theorem 1.2.* Pick a continuous antipodal function $f : S^n \to \mathbb{R}^n$. By the antipodality of f we have that for all $x \in S^n$

$$f(-x) = -f(x).$$

Because f is continuous from the $n+1$ sphere to $\mathbb{R}^n$ we can apply Theorem 1.1 and find that there is some x for which

$$f(x) = f(-x)$$

Combining these two expressions we get $f(x) = -f(x) \implies 2f(x) = 0 \implies f(x) = 0$.

Theorem 1.2 $\implies$ *Theorem 1.1.* Take a continuous mapping $f : S^n \to \mathbb{R}^n$. Then let $g(x) = f(x) - f(-x)$. This means that $g(-x) = f(-x) - f(x) = -g(x)$. Therefore $g(x)$ is an antipodal mapping, and it is clearly continuous as it is the sum of two continuous functions. We apply Theorem 1.2 and get that there is an x for which $g(x) = 0$. But for that x, $f(x) - f(-x) = 0 \implies f(x) = f(-x)$. □

Both these previous variants are statements regarding existence of certain kinds of points for certain maps of the sphere. One can also show that the theorem is equivalent to the non-existence of a certain kind of map between a sphere and a sphere in a different dimension.

Theorem 2.4. *Both versions of Borsuk-Ulam already mentioned are equivalent to the statement that there is no continuous antipodal mapping $f : S^n \to S^{n-1}$*

Proof. Assume that there exists a continuous antipodal mapping $f : S^n \to S^{n-1}$. By definition $S^{n-1} = \{x \in \mathbb{R}^n | x_1^2 + x_2^2 + \ldots x_n^2 = 1\}$. Thus we can extend the codomain of f to the whole of $\mathbb{R}^n$ and consider it as a function from $S^n \to \mathbb{R}^n$. As a function from $S^n \to \mathbb{R}^n$ we know that if $f(x)$ is in the image of f then $f(x) \in S^{n-1}$. Since 0 is not in S^{n-1} we have that f cannot equal zero anywhere. But by the Borsuk-Ulam theorem, all antipodal continuous maps from $S^n \to \mathbb{R}^n$ must equal zero somewhere. Therefore, we have a contradiction.

For the other direction, let's say we know that there is no continuous antipodal mapping from $S^n \to S^{n-1}$. Take any antipodal mapping from $f : S^n \to \mathbb{R}^n$. If $f(x) \neq 0$ for any x, then we can define $g(x) = \frac{f(x)}{\|f(x)\|}$. Then this maps to S^{n-1} and is antipodal so we have a contradiction. □

The following variant on Borsuk Ulam, fairly similar to Theorem 2.4, will be useful later.

Theorem 2.5. *The versions of Borsuk-Ulam given before are equivalent to the statement that there is no continuous mapping $B^n \to S^{n-1}$ that is antipodal on the boundary, i.e. for all $x \in \partial B^n = S^{n-1}$ we have $f(-x) = -f(x)$.*

Proof. First we note that there is a homeomorphism $\pi : (x_1, \dots, x_{n+1}) \to (x_1, x_2, \dots, x_n)$ between the upper hemisphere of the n sphere and the ball B^n. For an idea of what this function looks like, take the crust of the earth on the northern hemisphere to be the upper hemisphere of a realization of S^2. Then the B^2 ball could be considered as that portion of the earth (including inside the crust) that is on the longitude of the equator. Each point on the crust of the Northern-Hemisphere sits above a point on the disc of the equator, and so our function associates these points. In any event, if there existed an antipodal mapping $f : S^n \to S^{n-1}$ we could define $f \circ \pi^{-1}$ a function from $B^n \to S^{n-1}$ that would be antipodal on the boundary of B^n. Therefore we have proved by contrapositive that if our version of Borsuk Ulam in Theorem 2.5 is true then the other variants of Borsuk-Ulam introduced so far also hold.

The other direction is equally simple. Assume that there is a function g that goes from $B^n \to S^{n-1}$ that is antipodal on the boundary. Then $f(x) = g(\pi(x))$ is defined on the upper hemisphere of S^n. Further we can extend f to the whole sphere by letting $f(-x) = -f(x)$. The fact that B^n is antipodal on the boundary assures us that these associations are consistent and that our function is well-defined. This function is continuous and is an antipodal mapping of $S^n \to S^{n-1}$. □

Another variant of Borsuk Ulam, which goes in a somewhat different direction, is a statement congenial to elementary point-set topology.

Theorem 2.6. *Lyusternik-Shnirel'man. If $U_1, U_2, \dots, U_{n+1}$ is a cover of S^n with with each U_i open then for some i, U_i contains a pair of antipodal points.*

Remark 2.7. Lyusternik-Shnirel'man is also equivalent to the same statement with all the U_i taken to be closed sets. Formalizing the argument with full rigor would be a substantial digression, but seeing why it should be true is not difficult.

Assume that the theorem holds for closed covers. Take any $\{U_i\}_{i=1}^{n+1}$ that is an open cover of S^n. One can shave off a tiny bit from each element in this set and get a collection of closed sets which covers S^n. Then the antipodal points are guaranteed to be in one of these sets by Lyusternik Shnirel'man for closed sets. But whichever set has antipodal points is a subset of one of the open sets so Lyusternik-Shnirel'man is true for open sets.

Assume that the theorem holds for open covers. Take a closed cover $\{F_i\}_{i=1}^{n+1}$. Let $U_i^\epsilon = \{x \in S^n | \text{dist}(x, F_i) < \frac{1}{n}$. So there is for each ϵ an open cover of S^n and so for each ϵ we have an x_ϵ and $-x_\epsilon$ contained in the same set. It is possible to find a sequence $x_1, x_2, \dots$ such that $\lim_{j\to\infty} \text{dist}(x, F_i) = \lim_{j\to\infty} \text{dist}(-x, F_i) = 0$ for a fixed i. Because S^n is compact we know that there is a convergent subsequence and we take the limit of this sequence to be x. Then this x must have $-x$ in F_i as well.

Proposition 2.8. *The Lyusternik-Shnirel'man theorem is equivalent to the Borsuk-Ulam theorem.*

Proof. Lyusternik-Shnirel'man $\implies$ Borsuk-Ulam. It is a useful fact for this theorem that S^{n-1} can be covered using $n+1$ open sets in such a way as to have none of the open sets contain antipodal points. Let us denote this cover $\{U_i\}_{i=1}^{n+1}$.

If we assume Borsuk-Ulam is false and that there is a continuous antipodal map $f : S^n \to S^{n-1}$ then by the definition of continuity $f^{-1}(U_i)$ is open for all U_i. Further, it covers S^n. Therefore by Lyusternik-Shnirel'man we have $x, -x \in f^{-1}(U_i)$

for some i. But this is impossible. Since f is an antipodal map $f(-x) = -f(x)$ so $f(x)$ and $f(-x)$ are antipodes on S^{n-1}. But they are both in U_i by definition of $f^{-1}(U_i)$. But by our construction we have that no U_i contains antipodes and so we have reached a contradiction.

Borsuk-Ulam $\Longrightarrow$ Lyusternik-Shnirel'man. It is convenient for the proof to use closed covers instead of open covers. Take $\{U_i\}_{i=1}^{n+1}$ to be a closed cover of S^n.

Now we define a function $f : S^n \to \mathbb{R}^n$ coordinate-wise. For all $x \in S^n$ let $f_i(x) = \inf_{y \in F_i} |y - x|$. Then by Borsuk-Ulam we have that $f(x) = f(-x)$ for some $x \in S^n$. If the i^{th} coordinate is 0 at the point $f(x)$ then $x, -x \in F_i$. If none of the coordinates are zero, then it means that neither x nor $-x$ are in any of the sets F_i. But because all the F_i are a cover we have that these two must both be in F_{n+1}. □

3. A Measure Theoretic Approach to Lyustenik-Shnirel'man

Although the last section established that there was a large family of mutually equivalent statements it did little to show how one might get started actually proving any of them. As a first example, when things are very concrete and easily visualizable, we now set out to prove the case of the Borsuk-Ulam theorem when $n = 1$ which concerns maps of the circle into the real line. The Lyusternik-Shnirel'man variant of Borsuk Ulam can be proved quite easily.

Theorem 3.1. *If S^1 is covered by two open sets U_1 and U_2 then one of the two sets contains a pair of antipodal points.*

Proof. First, it is a basic fact that S^1 is not a separable space, so $U_1 \cap U_2 \neq \emptyset$. The proof is immediate. Pick $x \in U_1 \cap U_2$. Then the antipodal point $-x \in S^1 \subset U_1 \cup U_2$ and therefore $-x \in U_1$ or $-x \in U_2$. If $-x \in U_1$, U_1 contains x and $-x$. If $-x \in U_2$ then U_2 is the set which contains antipodal points. □

This proof is startlingly simple, but it is apparent that this strategy would not work for higher dimensions. Consider, for example, the Earth as an example of S^2. Then when we take the Northern Hemisphere without the equator as one open set, the Southern Hemisphere without the equator as another, and an open set that looks like a slim belt around the equator, we get an open cover of S^2 but the mutual intersection of all three sets is empty.

So let's return to the drawing board and see if we can come up with another way to prove Lyusternik-Shnirel'man in S^1.

Proof. Let's say that we are given two open sets U_1 and U_2 which cover S^1. The naive guess we might have is that it will that the "'larger"' of the two covers which will contain antipodal points. But what does "large" mean for us in this situation? One good measure of the size of these sets is the arc length of that portion of the circle covered by each i.e. the Lebesgue measure on S^1. As it turns out, it is possible to show that the larger of the two covers contains

We let μ denote the arc-length of a set and say that $\mu(S^1) = 2\pi$ which fits with our standard notion of arc-length. Because U_1, U_2 are open we know they are measurable and since they cover S^1 we have that $\mu(U_1) + \mu(U_2) > 2\pi$. This means that the average value of the two is greater than π, which implies that one or the other has value greater than π. Therefore without loss of generality we take $\mu(U_1) > \pi$.

Because μ is invariant under rotation we know that if $-U_1 = \{x \in S^1 \mid -x \in U_1\}$, then $\mu(-U_1) = \mu(U_1) > \pi$.

Since $\mu(U_1) + \mu(S^1 \setminus U_1) = 2\pi$ we have that $\mu(S^1 \setminus U_1) < \pi$. If $U_1 \cap -U_1 = \emptyset$ then $-U_1 \subset S^1 \setminus U_1$. But this implies that $\mu(-U_1) \leq \mu(S^1 \setminus U_1)$. Therefore we have that $\pi < \mu(-U) \leq \mu(S^1 \setminus U) < \pi$. This is a contradiction.

Actually, we can make this argument a little better without the assumption that $U_1 \cap -U_1 = \emptyset$. Note that

$$\mu(-U_1) = \mu(U_1 \cap -U_1) + \mu(-U_1 \setminus U_1)$$

Further, we observe that

$$-U_1 \setminus U_1 \subset S^1 \setminus U_1$$

Combining these observations with what was already stated we get following inequality.

$$\pi < \mu(-U) = \mu(-U_1 \cap U_1) + \mu(-U_1 \setminus U_1) < \mu(-U_1 \cap U_1) + \mu(S^1 \setminus U_1) < \mu(-U_1 \cap U_1) + \pi$$

Clearly whenever $\mu(-U_1 \cap U_1) = 0$ we get a contradiction, meaning that our set of intersection has positive measure. $\square$

Comparing the two proofs of Lyusternik-Shnirel'man in S^1 that have been given so far, it appears that the measure theoretic proof has at least two comparative advantages over the set theoretic proof. One potential comparative advantage is that the first proof generated only a single antipodal point, while this second one gives us an uncountable number of antipodal points. But one could easily reach the same conclusion from the first proof with only a little more work. Here's how. The map which sends each point of the sphere to its antipode is a bijective homeomorphism. Thus if $\nu(x) = -x$ denotes this bijective homeomorphism then when we take $\nu(U_1 \cap U_2)$ we know that we get an open set. $\nu(U_1 \cap U_2) \cap U_i$ is open for $i = 1$ and $i = 2$ and for one of these numbers is not equal to the empty set. The only open set containing a countable number of elements is the empty set so $\nu(U_1 \cap U_2) \cap U_i$ is uncountable for some i.

The real comparative advantage of the second proof is how it uses the fact that $U_1 \cap U_2$ is not empty. The first proof uses it as the entire proof, while the second proof only uses the fact to show that $U_1 \cap U_2$ has positive measure and therefore that $\mu(U_1) + \mu(U_2) > 2\pi$. The reason this is a real comparative advantage is that what we need in the first proof is that all the sets in the two element cover intersect. As mentioned before in the earth with belt example, this does not hold up in higher dimensions. In the second proof all we need is that the arc length covered by the elements in the cover is greater than the arc-length of the entire circle. Analogues of this fact do hold up in higher dimensions, each element in an open cover of the sphere must overlap with its adjacent neighbors, meaning that the sum of all the surface areas must be greater than the surface area total. In fact, it yields immediately, the following Borsuk-Ulam-type theorem.

Theorem 3.2. *If we take an open cover of S^n and any one element U in that cover has surface area greater than half the surface area of the entire sphere then that element U contains a non-trivial subset that is invariant under the antipodal mapping $v(x) = -x$. Further, that subset has positive Lebesgue measure.*

Proof. Let μ denote the Lebesgue measure on a sphere S^n. Further let U_i be an element of an open cover of S^n with $\mu(U_i) > \frac{\mu(S^n)}{2}$. Then because surface area of U_i is equal to the surface area of $\mu(-U_i)$ we have that $\mu(-U_i) > \frac{\mu(S^n)}{2}$. Further, similar to the argument in the proof above we have the following inequality.

$$\frac{\mu(S^n)}{2} < \mu(-U_i) = \mu(-U_i \backslash U_i) + \mu(-U_i \cap U_i) \leq \mu(S^n \backslash U_i) + \mu(-U_i \cap U_i) < \frac{\mu(S^n)}{2} + \mu(-U_i \cap U_i)$$

Whenever $\mu(-U_i \cap U_i) = 0$ we have a contradiction, yielding the theorem. □

The simplicity of this application of measure theory may belie the non-triviality of its result. This theorem says that so long as an open set on the sphere is sufficiently big it will have a large number of points with antipodes still in the set. Additionally, it actually reduces what needs to be proved for Lyusternik Shnirel'man in S^2 to the following.

Conjecture 3.3. *If U_1, U_2, U_3 is an open cover of S^2 and for each U_i the surface area measure $\mu(U_i) \leq 2\pi$ then one of the elements contains a pair of antipodal points.*

While this initially does not seem like we have moved ahead that far, after some thought it is apparent that we can use the fact that Lyusternik-Shnirel'man holds for S^1. That is to say, if the following conjecture were true:

Conjecture 3.4. *If U_1, U_2, U_3 is an open cover of S^2 and for each U_i the surface area measure $\mu(U_i) \leq 2\pi$ then for some i and j the union $U_i \cup U_j$ must contain an equator.*

Then one could say two open sets contain an equator and so must contain a set of positive measure invariant under the map which exchanges antipodes. Thus one might get the somewhat nice result that if the cover contains an element sufficiently big then that element has antipodes, and if none of the elements in the cover are sufficiently big then none is big enough to stop the other two from containing an equator and thus proving our theorem. After substantial thought about lots of messy covers, the conjecture still seems true. Sadly, this problem is a difficult one to solve and the author was unable to come up with any techniques to do it. Indeed, this last piece were it found might really yield a complete solution and it might generalize to higher dimensions. Nevertheless, there are proofs of Borsuk-Ulam already well-known from a variety of different angles, one of which we approach here.

4. Tucker's Lemma: Detour De Force

Many of the most commonly known proofs of Borsuk-Ulam theorem are relatively simple consequences of quite sophisticated machinery. The following sketch should give a reader un-initiated in this machinery a feel for what's missing and the advanced reader an idea of how it might work.

Proposition 4.1. *The Borsuk-Ulam theorem is true.*

Proof. First it is a fact, in itself something to prove, that antipodal maps of the sphere must have something called an odd degree, where by degree we mean a specialized notion attached to Homology theory. Now let's say there were an existent continuous antipodal map f of S^n to S^{n-1}. Then the S^{n-1} we are mapping to is the boundary of a disc in $\mathbb{R}^n$. Consider the restriction of f to an equator of $S^{n-1} \subset S^n$. We know that it must have odd degree. But if we look at f restricted to the upper hemisphere of S^n we know that f must be nullhomotopic when restricted to S^{n-1}. This means that our restriction of f to S^{n-1} must have degree zero. □

In contrast with the above homological proof, which require a lot of prerequisite knowledge in homology theory, there is a proof using graph theory which can be found in Jiri Matousek's book on the Borsuk-Ulam theorem. As it turns out, one can prove a weak version of Tucker's Lemma, derive from it the Borsuk-Ulam theorem, and then double back and use Borsuk-Ulam to derive Tucker's Lemma in its full glory. We intend to recount this proof in this paper.

But how does graph theory get introduced to the problem at all? The Borsuk Ulam theorem typically lays down rules governing continuous maps of the sphere. Continuous maps have a certain liquidity about them, however, and it does not take much effort to see that the Borsuk-Ulam theorem is saying something about a larger class of shapes than might be obvious at first. After all, if one takes a shape different from the sphere but that nevertheless has a continuous map onto the sphere, one can compose the map from the shape to the sphere with a map from the sphere to say another sphere and see that the composed function certainly is governed in some way by the Borsuk-Ulam theorem. The hope is that if one can prove a statement about a shape that is similar enough to the sphere than what goes for that shape will hold for the sphere as well.

This provides some motivation for the notion of *triangulations*, but to understand what exactly triangulations are we need a little background in simplices.

Definition 4.2. If $v_0, v_1, \ldots, v_k$ are points in $\mathbb{R}^n$ then we say that these points are affinely independent if $v_1 - v_0, v_2 - v_0, \ldots, v_k - v_0$ are linearly independent.

Which is why affinely independent sets look like this (the lines are added to make the linear independence more clear, the affinely independent set is just the vertices):

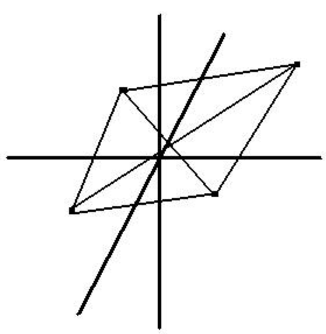

The convex hull of a set of points is the intersection of all convex sets which contain those points. In other words the convex hull of any of the points in the figure above is the same sets but instead of containing nothing in the interior we have that the interior is filled in. This suggests the following definition.

Definition 4.3. A **simplex** σ is the convex hull of a finite collection of affinely independent points in $\mathbb{R}^n$. The elements in the finite collection are called the vertices of σ and we denote them $V(\sigma)$. The dimension of a simplex is 1 less than the number of vertices it has.

Definition 4.4. If A is a subset of $V(\sigma)$ then we call the convex hull of A a face of σ. From this it readily follows that the face of a simplex is also a simplex.

The requirements for something to be a simplex are rather rigid and sometimes we want to use the more loose concept of simplicial complexes.

Definition 4.5. A nonempty family Δ of simplices is called a simplicial complex so long as the following two conditions are met.

(1) Each face of a simplex $\sigma \in \Delta$ is also a simplex in Δ
(2) The intersection of two simplices in Δ is a face of each simplex.

Further if one takes the union of all simplices in a simplicial complex Δ then this union is called the **polyhedron** of Δ and is denoted by $\|\Delta\|$

Roughly speaking a major difference between the two is that simplexes are convex, while simplicial complexes are only cobbled together from things which are convex and are not necessarily convex themselves. Further, we can easily go back and forth between simplicial complexes geometrically realized in the plane and the so called "abstract" simplicial complexes.

Definition 4.6. An abstract simplicial complex is a pair (V, K) where V is a set of "vertices" and $K \subset P(V)$ the power set of V. We require that if $F \in K$ and $G \subset F$ then $G \in K$. Dimension is defined similarly as one less than the dimension of the vertex set.

The idea of triangulation is to analyze these "triangular" looking sets instead of analyzing sets with any other shape. From the standpoint of topology, looking at these shapes comes cost free so long as the two spaces are homeomorphic, that is to say if there is a continuous bijection with continuous inverse between two spaces.

Definition 4.7. If X is a topological space and Δ is a simplicial complex such that X is homeomorphic to $\|\Delta\|$ then we call Δ a triangulation of X.

To really be able to switch over from talking about topological spaces to simplicial complexes we need one more piece of the puzzle. The functions which go from the topology of one space to the topology of another are called continuous. Likewise we can talk about functions which map simplices from one complex to another.

Definition 4.8. For two abstract simplicial complexes K and L if f is a function from $V(K)$ to $V(L)$ and has the property that $f(F) \in L$ for all $F \in K$ then we say that f is a *simplicial mapping*.

Such a mapping on abstract simplicial complex induces a mapping between their geometric realizations.

Definition 4.9. Let K_1 and K_2 be abstract simplicial complexes. Let f be a simplicial mapping from the first complex to the second. If Δ_1 and Δ_2 are the aforementioned geometric realization then define $\|f\| : \|\Delta_1\| \to \|\Delta_2\|$ as follows. If x is in the relative interior of a face $\sigma \in \Delta_1$ then by convexity of σ we have that $x = \sum_{i=1}^{d} a_i v_i$ where $v_i \in V(\sigma)$ and $\sum_{i=1}^{d} a_i = 1$. Therefore we say $\|f\|(x) = \sum_{i=1}^{d} a_i f(v_i)$. We call $\|f\|$ the *affine extension* of f.

Showing that $\|f\|$ is well-defined is not challenging, if the reader is interested in doing so on their own it is helpful to first show that x must be in the relative interior of only one face of Δ_1. Further, the reader may try to verify the also true statement that $\|f\|$ is continuous.

Now we can understand the statement of Tucker's Lemma.

Theorem 4.10. *Let T be a triangulation of B^n with a finite number of vertices and which is antipodally symmetric on the boundary. By antipodally symmetric on the boundary, we mean that the set of all simplices of T contained in $\partial B^n = S^{n-1}$ is a triangulation of S^{n-1} and further that if $\sigma \subset \partial B^n$ then $-\sigma \in T$. Given such a triangulation if there is a map*

$$\lambda : V(T) \to \{+1, -1, +2, -2, \ldots, +n, -n\}$$

for which $\lambda(-v) = -\lambda(v)$ for all vertices $v \in \partial B^n$ then there is a 1-dimensional simplices, or an edge, contained in T such that if v_1 and v_2 are its vertices $\lambda(v_1) = -\lambda(v_2)$.

One can think of λ as labeling the points and Tucker's lemma as stating that somewhere there is an edge connecting two vertices which have been given opposite labels.

Before we go forward, however, it is helpful to also recast Tucker's lemma as saying something not just about labellings but about simplicial maps. Recall that analogous to the conceptualization of continuity as a description of the way topologies go to other topologies, we have that simplicial maps are the way that we describe simplices as going into simplices. Indeed, if the point at the end is to say something about continuous maps by using simplicial maps as a way of getting a handle on the problem then we had better cast Tucker's Lemma in terms of simplicial maps.

But how to do that? Thankfully the machinery we developed before in terms of abstract simplicial complices allows us to quickly make a link between the set $\{\pm 1, \pm 2, \ldots, \pm n\}$ and a simplicial complex. Let us consider the abstract simplicial complex Λ formed by this set of labels. Tucker's Lemma gives us the fact that if there is a simplicial mapping of the vertex set of our triangulation T into Λ then there must be a complementary edge. In other words, there must be a simplex that looks like an edge connecting the vertices $+i$ and $-i$ for some i. Further, we could form a new simplicial complex from Λ by cutting out of Λ all edges connecting opposite labels (we will also have to cut out additional simplices from Λ that had these edges as their faces or else our result after our little simplicial surgery would have ceased to be a simplicial complex at all). For this new simplex, we would find that there could not be a simplicial map from $V(T)$ to this complex.

This discussion suggests we define $\diamond^{n-1}$ to be the abstract simplicial complex formed by the vertex set $V(\diamond^{n-1}) = \{\pm 1, \pm 2, \ldots, \pm n\}$ and under the requirement that a set $F \subset V(\diamond^{n-1})$ is a simplex if and only if there is no i such that $+i$ and $-i$ are both in F. If one thinks about the geometric realization of $\diamond^{n-1}$ it is most natural to consider all the vertices as a unit direction along the corresponding coordinate axes because this restriction agrees with our requirement for simplices to be affinely independent. That is to say $V(\diamond^{n-1}) = \pm e_1, \pm e_2, \ldots, \pm e_n$. At this point it is helpful to look at the picture.

The only requirement we put on $\diamond^{n-1}$ is that it cannot have a face which contains two opposite vertices, i and $-i$,which means that there is no simplex in $\diamond^{n-1}$ that is on the interior of the convex hull of $V(\diamond^{n-1})$. Thus we get from Tucker's Lemma the following theorem.

Theorem 4.11. *Let T be a triangulation of B^n that is antipodally symmetric on the boundary. Then there is no map $\lambda : V(T) \to V(\diamond^{n-1})$ that is a simplicial map of T into $\diamond^{n-1}$ and is antipodal on the boundary.*

There is a real impact to this theorem which needs to be teased out. $\diamond^{n-1}$ is the set of faces of the boundary complex formed by a set of coordinate vectors, and so is a triangulation of S^{n-1}. Tucker's Lemma thus appears to have as a consequence a statement stating that it is impossible to have certain kinds of maps from triangulations of the $n-ball$ to triangulations of the $n-1$ sphere. Tucker's Lemma is starting to look a lot more like Borsuk-Ulam.

Further, it is apparent that if one were to have a labeling of the triangulation and that there were no complementary edge, or edge with opposite labels, then one could reason similarly as I have done above and find that there *was* a simplicial map from the triangulation into $\diamond^{n-1}$. The non-existence of such a simplicial map thus implies that there is no complementary edge. Thus we have that these theorems are both equivalent and henceforth we refer to both as Tucker's Lemma.

The statement of the theorem should now be somewhat understood and we can now talk about proving the statement. As mentioned at the outset of this section, the theorem is difficult to approach directly. Instead of doing it for all triangulations we do it for a special class of triangulations. What exactly is this special class? We define it below.

Definition 4.12. The ball B^n is homeomorphic to the set $\hat{B}^n = \{x = x_1, x_2, \ldots, x_n \mid \sum_{i=1}^n x_i = 1\}$ which is the unit ball with the ℓ_1 norm. Let $\triangle$ be the "natural" triangulation of $\hat{B}^n$ induced by the coordinate hyperplanes. Explicitly that means that $\sigma \in \triangle$ means that $\sigma \in \diamond^{n-1}$ or $\sigma = \tau \cup \{0\}$ where $\tau \in \diamond^{n-1}$. Then we say that T is a *special triangulation* of $\hat{B}^n$ if

1. T is antipodally symmetric on the boundary, i.e. is contained in S^{n-1} and if $\sigma \in T$ and is a subset of S^{n-1} then $-\sigma$ is also a simplex of T.
2. For each $\sigma \in T$ we have that there is a $\tau \in \triangle$ such that $\sigma \subset \tau$.

Proposition 4.13. *Tucker's Lemma is true for special triangulations.*

Proof. Let T be a special triangulation of $\hat{B}^n$ and $\lambda : V(T) \to \{\pm 1, \pm 2, \pm 3, \ldots, \pm n\}$. For a simplex $\sigma \in T$ we consider the set $\lambda(\sigma) = \{\lambda(v) : v \text{ is a vertex of } \sigma\}$. $\lambda(\sigma)$ can be thought of as the collection of all labels which have been proscribed by the mapping λ to that simplex.

What we do next is compare this list to a different list of labels that one could ascribe to the simplex based on its location in $\mathbb{R}^d$. Pick an element x in the relative interior of σ. Then we say that $S(\sigma) = \{+i \mid x_i > 0, i = 1, 2, \ldots n\} \cup \{-i : x_i < 0, i = 1, 2, ..., n\}$. It follows from the second requirement of special triangulations that no matter where we pick x so long as it is in the *relative* interior we will get the same value of $S(\sigma)$ (the reason we say the relative interior is that there are many simplices in a special triangulation which have one or more edge on the coordinate axes, which, were it picked instead, would certainly give a different set).

We will call σ a **happy simplex** if $S(\sigma) \subset \lambda(\sigma)$. In other words, happy simplices have been proscribed labels which fit with the ones we'd naturally imagine them being given.

Observation 4.14. Happy simplices have the following properties.

(1) The dimension of a happy simplex is equal to the number of elements in $S(\sigma)$ or it is one less than that number

The reason for this fact is that σ lies inside the linear subspace spanned by the coordinate axes x_i such that $i \in S(\sigma)$ or $-i \in S(\sigma)$. Therefore it must have dimension less than or equal to that subspace. At the same time, however, there must be at least as many vertices being considered in the set $\lambda(\sigma)$ as there are elements of $S(\sigma)$ for it to be happy. Therefore since the dimension of a simplex is defined to be one less than the number of vertices, σ must be no more than one less than the number of elements in $S(\sigma)$.

This suggests that we give the name **tight** to those happy simplices for which $\dim S(\sigma) = k - 1$ and **loose** for those happy simplices for which $\dim S(\sigma) = k$.

(2) A boundary simplex, if it is happy, is necessarily tight. Non-boundary simplices are either tight or loose.

The reason is that when a simplex is on the boundary it has as many vertices as there are prescribed labels.

(3) $\{0\}$ is a loose happy simplex

It must always be happy because $S(\{0\}) = \emptyset$, and so the number of elements in $S(\{0\})$ is zero. This also implies that it must always be loose because $\{0\}$ has one vertex and dimension zero.

For brevity of expression, let us agree to call τ a facet of a simplex σ if τ is a face of σ and τ has dimension only one less than that of σ. We will now define a graph G with vertices all happy simplices. We say that two vertices $\sigma, \tau \in T$ are connected by an edge if σ and τ are antipodal boundary simplices or σ is a facet of τ and the labels of σ alone already make τ happy, i.e. $S(\tau) \subset \lambda(\sigma)$.

$\{0\}$ has degree 1 since it is connected to the edge of the triangulation that is made happy by the label $\lambda(0)$. What we want to show is that if there is no complementary edge, then any other vertex σ of the graph G has degree 2. This will yield a contradiction since a finite graph cannot contain only one vertex of odd degree.

Observation 4.15. (1) Suppose σ is a tight happy simplex. Then if τ is a neighbor of σ then either τ is $-\sigma$ or has σ as a facet.

(a) If σ is on the boundary $\partial \hat{B}^n$ then $-\sigma$ is one of its neighbors. Any other neighbor τ has σ as a facet and is made happy by its labels. That is to say, it is a loose happy simplex. Thus it has to lie in the subspace L_σ mentioned above. Then $L_\sigma \cap \hat{B}^n$ is a k-dimensional crosspolytope and since σ is in the boundary of a $(k-1)$-dimensional simplex it is a facet of only one k-simplex.

(b) σ does not lie on the boundary means that σ is a facet of only two simplices made happy by its labels, and these are the two neighbors.

(2) Suppose σ is a loose happy simplex.

(a) We have that $S(\sigma) = \lambda(\sigma)$ and so one of the labels occurs twice on σ. Then σ is adjacent to exactly two of its facet.

(b) There is an extra label $i \in \lambda(\sigma) \setminus S(\sigma)$. It follows that $-i \notin S(\sigma)$ for otherwise we would have a complementary edge and Tucker's Lemma would be true. One of the neighbors of σ is a facet of σ not containing the vertex with the extra label i. Further σ is a facet of exactly one loose simplex σ' made happy by the labels of σ. Namely, one with $S(\sigma') = \lambda(\sigma) = S(\sigma) \cup \{i\}$ We enter that σ' if we go from an interior point of σ in the direction of the $x_{|i|} - axis$ in the positive direction for $i > 0$ and in the negative direction for $i < 0$.

Thus for each possibility we have two neighbors, and so $\{0\}$ is the only vertex where we can only go out one direction, a contradiction. □

Thus we have established Tucker's Lemma for special triangulations.

Proposition 4.16. *If Tucker's Lemma is true for special triangulations then Borsuk-Ulam is also true.*

Proof. Assume Borsuk-Ulam is false. Specifically, we assume that $f : B^n \to S^{n-1}$ is a continuous map and that it is antipodal on the boundary. We will show this implies that there is a map λ and a special triangulation T such that $\lambda : V(T) \to V(\diamond^{n-1})$ that is a simplicial map and is antipodal on the boundary.

First of all it is apparent that the special triangulations can be made with arbitrarily small simplex diameter. We will use this fact. For the sake of convenience we will use the $||_{\ell^p}$ notation, where $|y|_{\ell^\infty}$ is equal to the largest absolute value of y_i. If $y \in S^{n-1}$ then because $\sum_{i=1}^n y_i^2 = 1$ we know that the absolute value of at least one of the y_i is greater than $\frac{1}{\sqrt{n}}$. That is, $|y|_{\ell^\infty} \geq \frac{1}{\sqrt{n}}$. Because f is a continuous function on a compact set we know that there is a δ such that $|x - x'| < \delta$ implies that $|f(x) - f(x')|_{\ell^\infty} < \frac{2}{\sqrt{n}}$. We then take T to be a triangulation with simplex diameter less than δ.

Now we define our labeling map which will lead to contradiction. First let $k(v) = \min\{i : |f(v)| \geq \frac{1}{\sqrt{n}}\}$. We define $\lambda : V(T) \to \{\pm 1, \pm 2, \ldots, \pm n\}$ as follows.

$$\lambda(v) = +k(v) \text{if} f(v)_{k(v)} > 0, -k(v) \text{if} f(v)_{k(v)} < 0$$

Since f is antipodal on ∂B^n we have that $\lambda(-v) = -\lambda(v)$. Therefore we know that Tucker's Lemma applies and that there is a pair v, v' that are connected by a complementary edge. But by definition of λ this can only happen if $f(v)_i \geq \frac{1}{\sqrt{n}}$ and $f(v')_i \leq \frac{-1}{\sqrt{n}}$ or vice-versa, but we assume that these are the inequalities without loss of generality. But then we have that $|f(v) - f(v)|_{\ell^\infty} \geq |f(v)_i - f(v')_i| \geq \frac{2}{\sqrt{n}}$. This contradicts the uniform continuity of f. Contradiction. □

Thus Borsuk-Ulam theorem has been established. The full-fledged Tucker's Lemma, for all triangulation not just special ones, is low hanging fruit.

Proposition 4.17. *Borsuk-Ulam implies Tucker's Lemma*

Proof. Assume there is an antipodal simplicial map λ of T into $\diamond^{n-1}$, where T is any Triangulation of B^n. But then $\|\lambda\| : \|T\| \to \left\|\diamond^{n-1}\right\|$ is a continuous map antipodal on the boundary. Then there is a continuous map $f : B^n \to T$ antipodal on the boundary and a continuous map $g : \diamond^{n-1} \to S^{n-1}$ that is antipodal on the boundary. Therefore $g \circ \|\lambda\| \circ f$ is a continuous map from $B^n \to S^{n-1}$ which is antipodal on the boundary. But this contradicts Borsuk-Ulam. □

5. The Ham Sandwich Theorem

At the beginning of the paper we stated that the Ham-Sandwich theorem is a result about hyperplanes bisecting a collection of compact sets. All the discussion of the Borsuk-Ulam theorem would be for naught if there was not some link between the two ideas. As it turns out, it is possible to think of hyperplanes in $\mathbb{R}^n$ being uniquely determined by points in S^n, and from there it is possible to design a function which measures how well the hyperplane defined by the point cuts the sets in $\mathbb{R}^n$. But first it is essential to see how one associates for each point in S^n a hyperplane in $\mathbb{R}^{\ltimes}$ and the related notion of *Half Space*

If one takes a point $u = (u_0, u_1, \ldots, u_n) \in S^n$ then each oriented hyperplane going through the origin is associated uniquely to a point on S^n by the dot product. That is to say, the set of all points $x = (x_0, x_1, \ldots, x_n)$ for which $u \cdot x = 0$ is the hyperplane going through the origin orthogonal to the vector going from the origin to u. When we take the set of all points with $u \cdot x \leq 0$ we get a half space, or the half of $\mathbb{R}^{n+1}$ below the hyperplane (although what exactly below means depends on how our axes are set up, so we should just ensure that we are consistent in our algebraic definition). But what we wanted was a hyperplane that divided $\mathbb{R}^n$. In order to get that hyperplane what we can do is we intersect that with the hyperplane defined by all the points $< x_0, x_1, \ldots, x_n >$ where $x_0 = -1$, that is we force the planes to have less freedom where they could be. The choice $x_0 = -1$ is somewhat arbitrary, in that we would get unique associations for every half-space in $\mathbb{R}^n$ if $x_0 = -2$, for example. Further we could have locked any single x_i at -1 and gotten rather different looking associations for all the hyperplanes in $\mathbb{R}^n$ but every oriented hyperplane would still be covered uniquely. Nevertheless picking $x_0 = -1$ will reveal the simplest form for our unique associations. That is to say, we have two requirements on x for it to be in our set.

- $x \cdot u \leq 0$.
- $x_0 = -1$

We can rewrite the first equation with the dot product as $\sum_{i=1}^{n} x_i u_i \leq u_0$. Hence, we have the motivation to make the following definition.

Definition 5.1. *Half Space.* Given a point $u \in S^n$ we define the *half space* at u, denoted as $h^+(u)$, as follows: $h^+(u) = \{(x_1, x_2, \ldots, x_n) \in \mathbb{R}^n \mid \sum_{i=1}^{n} x_i u_i < u_0\}$.

This definition is essential to the proof I will present of Ham Sandwich, and before jumping into the proof a few more words about this definition. As written in n-dimensional space the equation looks bewildering but really this function is not difficult to imagine even in three-dimensions. What the definition is saying, for the case of S^2 in three dimensions, is that when one picks a point on the sphere one draws a vector connecting that sphere to the origin. This vector in turn uniquely determines a plane going through the origin, and in fact every plane through the origin is determined by some vector. Additionally, antipodal points on the sphere generate planes with same slope but pointing in opposite directions, an important property true for all n-dimensions that we will use later. Once one has a given plane through the origin then one takes everything below it, where one has to be careful about what below means, and also takes the intersection of that with the plane which runs parallel to the yz-plane through the point $(-1, 0, 0)$. All that remains after this intersection is an infinite portion of the yz-plane through $(-1, 0, 0)$. This infinite-portion, deceptively called a half space because it divides the yz-plane into two equal but infinite chunks, is what we want. All the x coordinates in this half-space are -1, so we can safely collapse all our coordinates, which are in R^3, to merely (y, z) coordinates in $\mathbb{R}^2$. In n-dimensions, the process is similar although more difficult to visualize.

Theorem 5.2. *Ham Sandwich. Given n compact sets in $\mathbb{R}^n$ there is a hyperplane which bisects each set into two sets of equal measure.*

Proof. Let $K_1, \ldots, K_n$ be our collection of compact sets. Let $\mu_i(A) = \mu(A \cap K_i)$ for all sets K. It is a fact that all μ_i are still Radon measures. It is a property of Radon measures that the measure of compact sets is always finite, therefore $\mu_i(\mathbb{R}) = \mu(K_i) < \infty$.

The salient feature of the Half Space construction is that we get a series of planes defined by each point on S^n with points near each other leading to similar hyperplanes in $\mathbb{R}^n$. What we do next is define a function for each point on S^n based on the measure of the half space each point generates. For all $u \in S^n$ let f be a function from $S^n \to \mathbb{R}^n$ defined coordinate-wise as $f_i(u) = \mu_i(h^+(u))$. This function going from $S^n \to \mathbb{R}$ is well-defined for each u because the measure of a half space is finite.

It should be clear that this construction is a map from $S^n \to \mathbb{R}^n$ and that as a result the Borsuk-Ulam Theorem becomes a useful tool. In fact, its application proves the theorem. If one can apply Borsuk-Ulam to this function then one gets that there is a point u on the sphere for which $f(u) = f(-u)$. In other words, for all i we have $\mu_i(h^+(u)) = \mu_i(h^+(-u))$. But as mentioned before, antipodal points on the half plane produce opposite half-spaces. Therefore the space below the hyperplane generated by u has equal measure to the space above it. But the μ_i measure of the whole space is only the measure of K_i. Therefore we have that the measure of the portion of K_i above the plane is equal to the measure of that below it. Therefore if we can use Borsuk-Ulam we will have established the fact for all K_i and be done.

Nevertheless, to use Borsuk-Ulam we need to establish continuity. Pick $u \in S^n$ and a sequence $x^{(k)} \in S^n$ so that $\lim_{k\to\infty} x^{(k)} = u$. We want to show that $\lim_{k\to\infty} f(x^{(k)}) = f(u)$. Because we have only a finite number of coordinates, we can do that by showing that the limit holds termwise.

Therefore we only need to show for all i that:

$$\lim_{k\to\infty} \mu_i(h^+(x^(k))) = \mu_i(h^+(u))$$

Another way to state this fact is in terms of integrals. That is to say, if we set $g_k(y)$ as the characteristic function of $h^+(x^k)$ and g_u as the characteristic function of $h^+(u)$ then what we want to show is that:

$$\lim_{k\to\infty} \int g(y) d\mu_i = \int g_u(y) d\mu_i$$

An analyst will immediately recognize that the best way to show this fact is Lebesgue's Dominated Convergence Theorem. Lebesgue's Theorem will show that this equality is true if we can bound all the g_k and g_u by another function which has finite integral and show that $\lim_{k\to\infty} g_k(y) = g_u(y)$ almost everywhere with respect to the measure μ_i. That is to say, we need almost everywhere point-wise convergence. The bounding part is easy, the characteristic function is always less than equal to 1, and $\int 1\mu_i$ is finite for all i. Let $y \notin \delta h^+(u)$, the boundary of $h^+(u)$. Then we know $y =< y_1, y_2, \ldots, y_n >\in h^+(u)$ if and only if:

$$y_1u_1 + y_2u_2 + \ldots y_nu_n < u_0$$

Let $\delta = u_0 - y_1u_1 + y_2u_2 + \ldots y_nu_n > 0$. Then we say $x^{(k)} = x_0^{(k)}, x_1^{(k)}, \ldots, x_n^{(k)}$. $\lim_{k\to\infty} x^{(k)} = u$ means that there is for each coordinate an N_i such that for all $k > N_i$ we have $x_i^{(k)} < u_i + \frac{\delta}{2ny_i}$ and $x_i^{(k)} > u_i - \frac{\delta}{2ny_i}$, except in the case of y_0 which we have not defined yet but for convenience we simply let equal 1. Therefore there is a largest N so that we have for all i the same property when $k > N$. For $k > N$ we thus get that

$$\sum_{i=1}^{n} y_i x_i^{(k)} < y_1(u_1 + \frac{\delta}{2ny_1}) + y_2(u_i + \frac{\delta}{2ny_2}) + \ldots + y_n(u_n + \frac{\delta}{2ny_n}) =$$

$$= (\sum_{i=1}^{n} y_i u_i) + \frac{\delta}{2} = (\sum_{i=1}^{n} y_i u_i) + \delta - \frac{\delta}{2} < u_0 - \frac{\delta}{2} < x_0^{(k)}$$

.

Therefore we have that for all $y \in h^+(u)$ that are not in the boundary, that $\lim_{k\to\infty} g_k(y) = g_u(y)$. Now we need to check that for all $y \notin h^+(u)$ that $g_k(y)$ also converges on $g_y(u)$. That is to say, for large enough k, y will not be in $h^+(x^{(k)})$. But this is similar to the argument before. We know that $\sum_{i=1}^{n} y_i u_i > u_0$. If we let $\sum_{i=1}^{n} y_i u_i - u_0 = \delta > 0$ then we know that we can find an N such that for all $k > N$ we have $x_i^{(k)} < u_i + \frac{\delta}{2ny_i}$ and $x_i^{(k)} > u_i - \frac{\delta}{2ny_i}$, except in the case of y_0 which we again set to be 1. Then we have that

$$x_0^{(k)} < \frac{\delta}{2} + u_0 < (\sum_{i=1}^{n} y_i u_i) - \frac{\delta}{2} = (\sum_{i=1}^{n} y_i(u_i - \frac{\delta}{2ny_i}) < \sum_{i=1}^{n} y_i x_i^{(k)}$$

But this means that y not in $h^+(x^{(k)})$ for all sufficiently large k, so again g_k converges point wise. Thus we have shown that we can use Lebesgue Dominated Convergence and have established continuity of our function f, which by using Borsuk-Ulam finishes the proof of Ham Sandwich. □