Ham Sandwich Theorem

Sophia Tatar

Ross Mathematics Program Lecture 2

Abstract

This is the second talk of a lecture series on groundbreaking theorems in geometry! Whether you're a first-year student or a seasoned peer mentor, this lecture should be equally accessible and compelling.

This session explores the Ham Sandwich Theorem, a powerful and versatile tool in incidence geometry and general geometric theory.

The folklore phrasing asks the following: Can you always slice a ham sandwich into two equal halves, no matter how the bread and ham are shaped? What if you add a slice of cheese? How do these cuts look, and how can we calculate where to make them? We will use these questions to start our conversation on how we can cut bounded open sets using hyperplanes and hypersurfaces, two extremely useful tools in geometry that everybody should know how to work with effectively. We will intuitively and rigorously define every notion along the way so do not fret if those words are unfamiliar. This theorem illustrates the beauty of hyperplanes and hypersurfaces in action. But that's not all! For our proof of the Ham Sandwich theorem, we need the Borsuk-Ulam Theorem, which illustrates how some theorems can have many many different rephrasings that seem different but are actually equivalent. Finally, we will venture out into the lesser-known (yet equally intriguing and useful) Polynomial Ham Sandwich theorem. Expect lots of stunning visuals that will bring textbook definitions to life.

Introduction to the Classical Ham Sandwich Theorem

Imagine you have a rectangular piece of cake and a circular drop of white frosting on top, as shown below.

Question: Can you draw a line going through the cake and frosting that bisects both the frosting and the cake no matter where the frosting is on the cake?

By bisect, I mean split into two sets of equal Lebesgue measures.[1]

Answer: The answer to the question away above is clearly yes since this can always be done by finding the center of each shape and drawing a line through the center points.

The question that follows is if we can do this for shapes without a symmetrical center. From now on in this lecture, we are curious about not just any shapes, but rather only bounded open sets.[2]

Pulling the idea of bisecting shapes without symmetrical centers and the

[1]See bottom of document for a formal definition of Lebesgue measure, but the intuitive idea is that Lebesgue measure is just a way to measure the size of a set and is especially useful for funky-looking sets. To find the size of a funky (or not funky) set, we first cover it with a minimum tilling (a covering set where the total measure–length, area, volume, etc.—of the tiles used is minimized). The sum of the measures of all the tiles in the minimum covering is the Lebesgue Measure. For a two-dimensional shape like a circle, the measure is the area. For a box, the measure is the volume, etc.

[2]See bottom of document for formal definitions of bounded and open. However, intuitively, a bounded set is a set such that if I were to take a big enough box, I could fix the whole set inside. Also intuitively, an open set is a closed set where we remove the boundary (can imagine taking a polygon and removing the edges).

idea of bounded open sets, we get the following question.

Question: Let's say I have two bounded open sets in $\mathbb{R}^2$. Can I draw a line to bisect both of these no matter what the two bounded open sets are?

The image below illustrates splitting two bounded open sets using a line. If the area of the red and pink are equal, then the line bisects that bounded open set. The same goes for the second bounded open set.

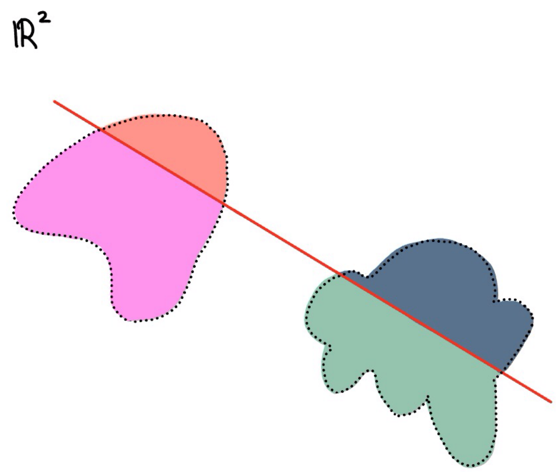

Answer: The picture seems to suggest that the answer to the question above is yes, meaning that we can always find a line to split two bounded open sets no matter how funky-looking they are. However, this statement is much more difficult to prove than the case for the rectangular cake and circular frosting. We will get back to proving this later.

What if we add a third open bounded set? Can we still find such a line that splits all sets into equal measures? More formally,

Question: Let's say I have three bounded open sets in $\mathbb{R}^2$. Can I draw a line to bisect both of these no matter what the three bounded open sets are?

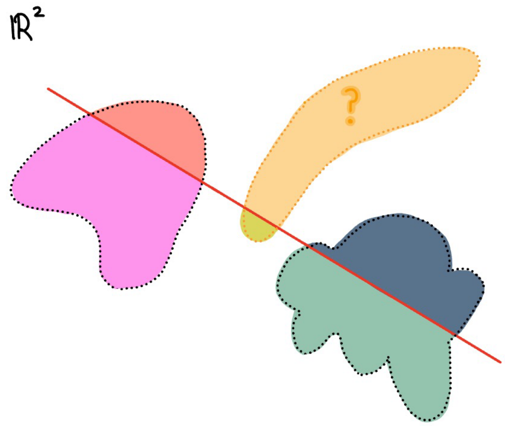

Answer: The picture suggests that something breaks down when we add more bounded open sets.

So when we are working in *two* dimensions, it seems like if we have *two* sets we can always draw a line to split them in half. But when we add more sets, it is obvious that most of the time there will not be such a line. How can we fix this issue? How can we split three bounded open sets?

One way to fix this would be to go into a higher dimension. If we have three open bounded sets in $\mathbb{R}^3$ rather than in $\mathbb{R}^3$, we can conjecture that there will always be such a way that each set is split into two sections of equal Lebesgue measure. The three bounded open sets can be thought of as two pieces of bread and a slice of ham, as shown below. So the question can be rephrased to ask:

Question: "Can we slice a ham sandwich into equal halves no matter how the bread and ham are shaped?"

The image below illustrates the situation.

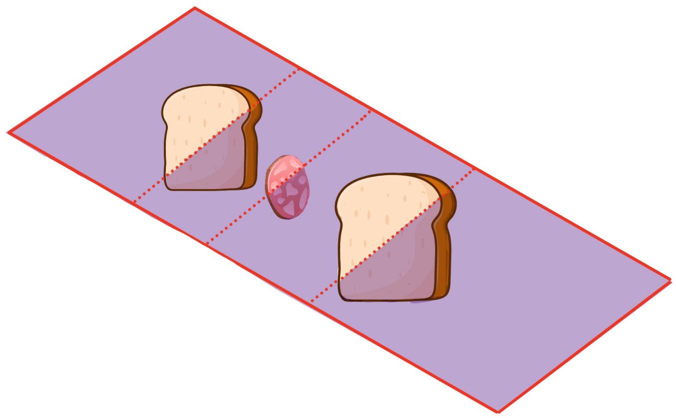

Answer: As the picture suggests, it is possible to cut 3 bounded open sets in R^3 using a plane.

It appears that there is a pattern...In $\mathbb{R}^2$ we can bisect up to 2 bounded open sets. In $\mathbb{R}^3$ we can bisect up to 3 bounded open sets. What happens in $\mathbb{R}^4$? What happens after that?

In order to generalize the number of bounded open sets we can bisect in higher dimensional spaces, we need to generalize what we can use to make our cuts.

To tackle this, we will formalize the cutting tools in the lower dimensions, and then try to find something equivalent for the higher ones.

In the first case, $\mathbb{R}^2$, we use a line, so the equation should be of the form $h(x_1, x_2) = h_0 + h_1x_1 + h_2x_2 = 0$. In R^3, we use a plane, which has the form $h_0 + h_1x_1 + h_2x_2 + h_3x_3 = 0$. Notice that for both, the function h splits our space into two sections: $h(x) > 0$ and $h(x) < 0$ (by intermediate value theorem). What we are describing here are specific cases of something called a hyperplane.

Hyperplane

In $\mathbb{R}^n$, a hyperplane is an (n-1)-dimentional subspace that divides $\mathbb{R}^n$ into two half-spaces. We define it as the set of points $x = (x_1, x_2, ...x_n)$ that satisfy a linear equation of the form $h_0 + h_1x_1 + h_2x_2...h_nx_n$.

We use the notation H^+ to denote the subspace where h has a positive value, and similarly use H^- to denote the subspace where h has a negative value.

Generalizing the idea of using a hyperplane to cut bounded open sets, we get the classical Ham Sandwich Theorem. Informally, the Ham Sandwich theorem says that any n (finite) bounded open sets in R^n can be simultaneously bisected by a hyperplane. Formally,

Classical Ham Sandwich Theorem

Let $U_1, ...U_n \subset \mathbb{R}^n$ be bounded open sets. Then there exists a hyperplane $H = \{x \in \mathbb{R}^n | h(x) = 0\}$ (with h(x) a degree one polynomial in n variables) such that for each $i \in [n]$ the two sets $U_i \cap H^+ = \{x \in U_i | h(x) > 0\}$ and $U_i \cap H^- | h(x) < 0$ have equal Lebesgue measure. In this case, we say that H bisects each of the $U_i's$.

Below is an example of what it would mean to bisect a bounded open set U_1 in $\mathbb{R}^2$ using a hyperplane $h(x_1, x_2) = h_0 + h_1 x_1 + h_2 x_2 = 0$ into the sections as stated in the theorem.

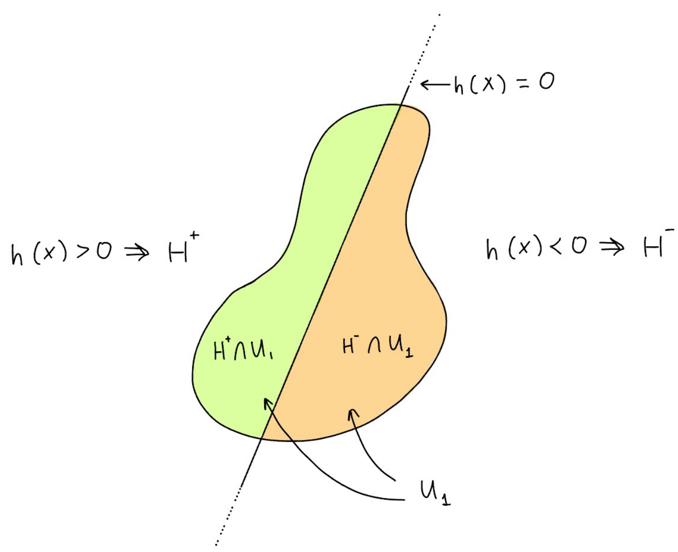

Proof of Ham Sandwich Theorem

Defining Tools for Proof of Ham Sandwich Theorem

A short proof to the classical Ham Sandwich Theorem uses a famous result called the Borsuk-Ulam Theorem. This theorem is really useful to this proof, but it is also interesting in and of itself because it is a theorem with several equivalent versions that seem completely different but actually imply the same thing. It is also a good theorem to remember because there are also numerous other applications of it beyond the Ham Sandwich Theorem.

In order to understand the Borsuk-Ulam theorem, we need to know one definition.

Imagine you had a one-dimensional unit sphere where points on opposite sides were assigned the same magnitude with different signs. This property is called **antipodal symmetry**. More generally,

Antipodal Symmetry

Let S^n be the n-dimensional unit sphere in $\mathbb{R}^{n+1}$. For any point x on S^n, the antipodal point of x is denoted by $-x$, which is the point directly opposite to x through the center of the sphere. A function f is said to have antipodal symmetry if it satisfies the following condition: $f(-x) = -f(x)$ for all x in the domain.

Many of the formulation of the Borsuk-Ulam theorem use antipodal symmetry.

One of the versions of the Borsuk–Ulam theorem, and in my opinion, the one that is easiest to remember, intuitively states that for any continuous function from the sphere to Euclidean space that respects antipodal symmetry, there is a point on the sphere that maps to the origin. Formally,

Borsuk-Ulam Theorem Formulation 1

For every continuous mapping $f : S^n \mapsto R^n$, there exists a point $x \in S^n$ such that $f(x) = f(-x)$.

To visualize this theorem in $\mathbb{R}^2$, we can imagine the following: Take a rubber ball, deflate and crumple it, and lay it on the ground. According to the Borsuk-Ulam theorem, there will be two points x and -x on the surface of the ball that were diametrically opposite (antipodal) and now are lying on top of one another!

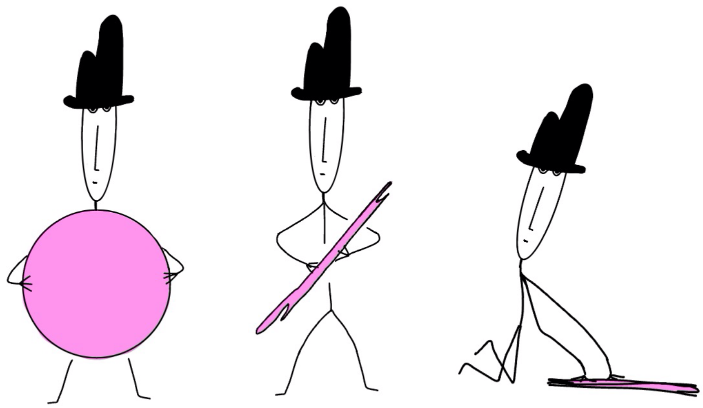

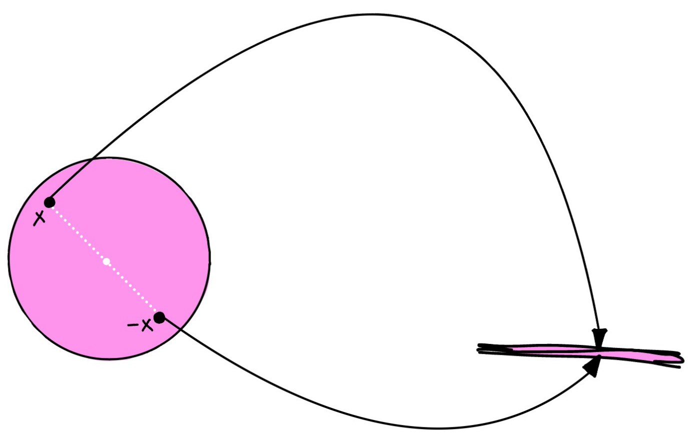

There are many equivalent statements that sound completely different. Some of the easiest-to-understand formulations are below.

Borsuk-Ulam Theorem Equivalent Statements

There is no continuous mapping $f : B^n \mapsto S^{n-1}$ that is antipodal on the boundary, i.e., satisfies $f(-x) = f(x)$ for all $x \in S^{n-1} = \delta B^n$

For any cover $F_1, ..., F_{n+1}$ of the sphere S^n by n+1 closed sets, there is at least one set containing a pair of antipodal points (that is, $F_i \cap (-F_i) \neq \emptyset$).

There is no antipodal mapping $f : S^n \mapsto S^{n-1}$

For any cover $U_1, ..., U_{n+1}$ of the sphere S_n by n+1 open sets, there is at least one set containing a pair of antipodal points.

One of the versions of the Borsuk–Ulam theorem, states the following. Let $S^n \subset \mathbb{R}^{n+1}$ be the n-dimensional unit sphere and let $f : S^n \mapsto \mathbb{R}^n$ be a continuous map such that $f(-x) = -f(x)$ for all $x \in S^n$ (such a map is called antipodal). Then there exists x such that $f(x) = 0$.

It is a good exercise to think about why these two statements are equivalent, but we will only use the last formulation (in bold) for the Ham Sandwich theorem proof. Let's explore what the last formulation really means using an example.

Below are two examples of different functions f with the property described in the assumption of the theorem for $\mathbb{R}^1$. You can imagine the situation below as somebody throwing a string on a line so that two points on opposite sides of the string end up having opposite values. The picture below illustrates how there will have to be a pair of antipodal points that map to 0.

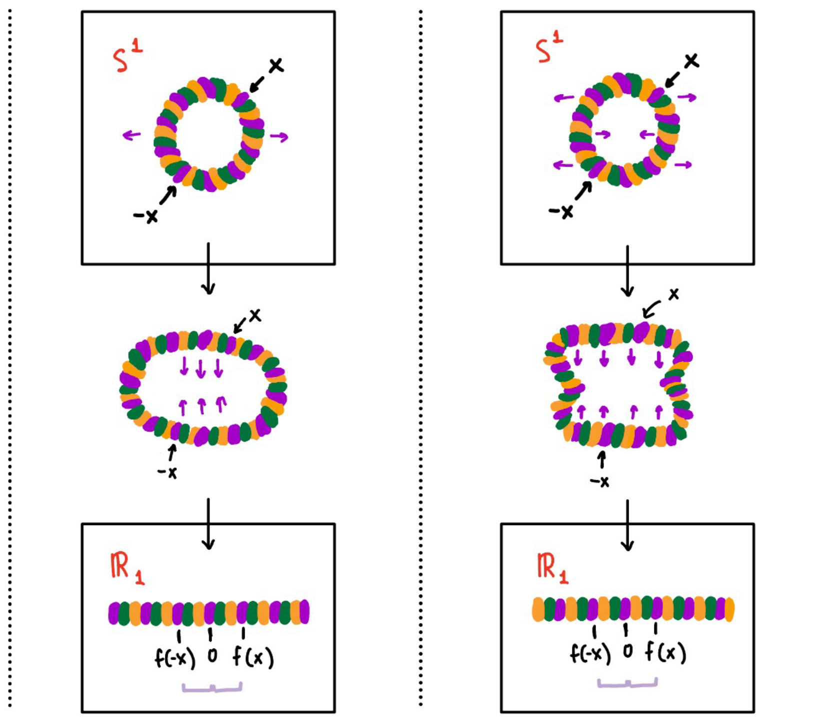

Isn't it interesting how this statement about throwing a string on a line so that there is a point that maps to 0 and the statement about squishing a rubber ball implying having points that end up on top of one another are equivalent statements?

Now that we have explored the Borsuk-Ulam theorem, we have all the blocks to finally prove the Ham Sandwich theorem.

Actual proof of Ham Sandwich Theorem:

We assume all of the notation defined above. So, just to summarize, we are cutting our bounded open sets using hyperplanes, and each hyperplane corresponds to some degree one polynomial $h(x_1, ...x_n) = h_0 + h_1x_1 + ...h_nx_n = 0$. The hyperplane, by definition, splits the space into two subspaces (positive and negative), which we call H^+ and H^-.

We discovered at the Ross program, we can represent any polynomial using a list of its coefficients. Since we only need to know about the sign of h to make the proper cut, we can represent the hyperplane as a scaled version of its coefficients where the coefficients are scaled to form a unit vector $v_h = (h_0, h_1, ...h_n)$. This means $V_h \in S^n \subset \mathbb{R}^{n+1}$. We define a function $f : S^n \mapsto \mathbb{R}^n$ where $f(v_h) = (y_1, y_2, ..y_n)$ such that

$$f(v_h) = (|H^+ \cap U_i| - |H^- \cap U_i|)_{i \in [n]}$$

where $|H^+ \cap U_i|$ is the measure of the part of U_i on the positive side of the hyperplane and $|H^- \cap U_i|$ is the measure of the part of U_i on the negative side of the hyperplane.

Expanding, we can see what the function actually does for each U_i.

$$\begin{aligned}
y_1 &= |H^+ \cap U_1| - |H^- \cap U_1| \\
y_2 &= |H^+ \cap U_2| - |H^- \cap U_2| \\
&\vdots \\
y_n &= |H^+ \cap U_n| - |H^- \cap U_n|
\end{aligned}$$

Notice that f is continuous. It is also clear that $f(-v_h) = -f(v_h)$ because $-v_h = (-h_0, -h_1, ... -h_n)$, so $f(-v_h) = (-y_1, y_2, ..., y_n)$, which defines the same space just switching the location of H^+ and H^- $(|H^- \cap U_1| - |H^+ \cap U_1|)$.

Since f is continuous and $f(-v_h) = -f(v_h)$, all of the conditions of the Borsuk-Ulam theorem are met, and we can apply it to conclude that there exists a zero of f. In other words, there exists a hyperplane h(x) such that $y_i = |H^+ \cap U_i| - |H^- \cap U_i| = 0$ for all i's. This means that there exists a hyperplane that bisects each of the U_i's into sets of equal Lebesgue measure. This completes the proof the the Classical Ham Sandwich theorem.

Polynomial Ham Sandwich Theorem

The Polynomial Ham Sandwich theorem is an extension of the Ham Sandwich Theorem that enables us to cut a higher number of bounded open sets at the same time. However, in order to do this, hyperplanes are not enough. We must employ another tool to split our space into subspaces that allows for more degrees of freedom, and this tool is called a "hypersurface."

Hypersurface

A hypersurface is a set $H = \{x \in \mathbb{R}^n | h(x) = 0\}$, where now h can be a polynomial of arbitrary degree d.

Now that we know what we are cutting with, we can formally state the Polynomial Ham Sandwich Theorem.

Polynomial Ham Sandwich Theorem

Let $U_1, ..., U_t \in \mathbb{R}^n$ be bounded open sets with $t < \binom{n+d}{d}$. Then there exists a degree d hypersurface H that bisects each of the sets $U_i, i \in [t]$.

If $t = n$, then we are in the case of the classical theorem and use hyperplanes to make the bisection. If $n < t < \binom{n+d}{d}$, then we need to use hypersurfaces.

Proof. The proof is identical to the degree one proof. We identify each degree d hypersurface with its (unit) vector of coefficients and apply the Borsuk-Ulam theorem on the function f mapping to the differences of the measures of the

The $\binom{n+d}{d}$ bound arises from the fact that there are $\binom{n+d}{d}$ coefficients in a polynomial of degree d with n variables and each term of the hyperspace corresponds to an equation $y_i = |H^+ \cap U_i| - |H^- \cap U_1| = 0$. There need to be enough terms in the polynomial to have a non-trivial solution $y_i = 0\ \forall i$. If there are too many $U_i's$", the hypersurface doesn't have enough terms to have a nontrivial solution to all equations.

Formalized Definitions (Referred to above)

Open Set

In $\mathbb{R}^n$ an open set U is one that has the property that for any point x in U, there exists a positive radius r such that all points within the distance r from x (forming an open ball around x) are also contained within U.

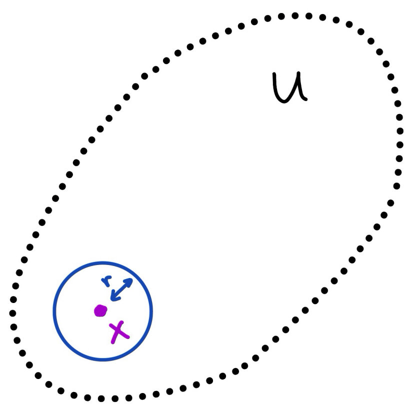

Bounded Set

In $\mathbb{R}^n$ a set A is bounded if there exists a real number M and a point x in the space such that the distance d(x,y) is less than M for all points y in A. In other words, all points in the set are within a finite distance M from some fixed point x.

Lebesgue Measure

The formal definition involves considering all possible countable coverings and taking the infimum of the total "volume" of these coverings. Specifically, for a set $A \subset \mathbb{R}^n$, its Lebesgue measure $\mu(A)$ is defined as

$\mu(A) = \inf\{\sum_{i=1}^{\infty} l_i : A \subseteq \bigcup_{i=1}^{\infty} S_i$ and each S_i is a rectangular box with measure $l_i\}$