Manifolds, Motivated

Daniel K. Rui - June 11, 2021

Abstract

This paper seeks to provide motivation for many important concepts in introductory topology; we will start with some basic definitions, and quickly move into motivating the development of abstract topological spaces by thinking about ways of defining local and global continuity of functions without the notion of distance. We will follow a narrative that meanders around ideas like partitions of unity, some theorems of Urysohn, the countability and separation axioms, metrization, compactness, topological dimension, and manifolds, ultimately culminating in a proof of a version of Whitney's embedding theorem. The recommended level of background is the completion of Math 334-335 at UW (preliminary sections could be Math 13X level).

Contents

1 Basic Notions

Before we begin getting into the meat of the paper, it is best to introduce some basic notions. Hopefully, the reader recalls (from say a high school algebra) class that in $\mathbb{R}$, we can write intervals with either parentheses, or with brackets, e.g. $[0, 1)$ or $(-\infty, 7)$ or $[e, \pi]$, where the difference is that the parentheses means "does NOT include the endpoint" and the bracket means "DOES include the endpoint". Also, intervals with parentheses on both ends were called *open intervals*, and intervals with brackets on both ends were called *closed intervals*. We can generalize these notions (to say sets in the plane, $\mathbb{R}^2$) by generalizing the notion of "endpoint". Let's consider the interval $[0, 1)$:

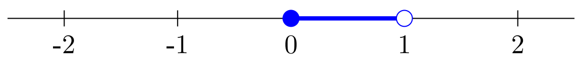

What makes an endpoint (i.e. $\{0, 1\}$) different from any other point? Suppose we pick another point, like $\frac{1}{2}$ or π. Zooming in very close on $\frac{1}{2}$, eventually all that will fill our screen will be blue. Similarly, zooming in on π, eventually we will just see the thin black number line. In other words, there is some small enough radius $r > 0$ such that (henceforth abbreviated "s.t.") the ball (really just interval, since we are in $\mathbb{R}$) centered at $\frac{1}{2}$ with radius r, which we denote as $B(\frac{1}{2}, r)$ that is entirely colored blue, i.e. contained within $[0, 1)$. Similarly, there is $r > 0$ s.t. $B(\pi, r)$ is colored black, i.e. contained in $\mathbb{R} \setminus [0, 1)$ (the backslash is the "set minus" symbol, i.e. $\mathbb{R} \setminus [0, 1)$ is just $\mathbb{R}$ but delete the points of $[0, 1)$; another notation is the "complement", $[0, 1)^{\complement}$, where the ambient space $\mathbb{R}$ is inferred). However, something very different happens when we zoom into 0 or 1. There, no matter how far we zoom in, we will see a little bit of blue and a little bit of black; phrasing this like we did above, we say that for ANY $r > 0$, the ball/interval $B(0, r)$ shares points with $[0, 1)$ AND $[0, 1)^{\complement}$ (and of course same with $B(1, r)$).

These observations allow us to generalize concepts of openness, closedness, and endpoints (more generally called "boundary points") to other spaces, namely the Euclidean spaces $\mathbb{R}^n$. Given a set $S \subseteq \mathbb{R}^n$, we define the interior of S to be the set of points $x \in S$ (equivalently $x \in \mathbb{R}^n$) s.t. there is a radius $r > 0$ s.t. $B(x, r) \subseteq S$ (i.e. the ball of radius r centered at x is COMPLETELY contained within S); similarly, the exterior of S is the set of $x \in \mathbb{R}^n$ s.t. there is $r > 0$ s.t. $B(x, r) \subseteq S^{\complement} := \mathbb{R}^n \setminus S$ (I will use the notation ":=" to mean "which is/was defined to be"); and lastly, the boundary of S, denoted ∂S, is the set of $x \in \mathbb{R}^n$ s.t. for EVERY $r > 0$, $B(x, r) \cap S \neq \varnothing$ and $B(x, r) \cap S^{\complement} \neq \varnothing$ (i.e. the intersections of the ball and $S, S^{\complement}$ are both nonempty).

Note that every point of $\mathbb{R}^n$ is exactly one of: an interior point of S, a boundary point of S, or an exterior point of S. We can then define $S \subseteq \mathbb{R}^n$ to be open exactly when S does NOT contain any points of its boundary ∂S (equivalently, S equals its interior), and we can define S to be closed exactly when S contains ALL points of its boundary ∂S (equivalently, S equals $S \cup \partial S = \operatorname{int}(S) \cup \partial S$). Note also that S is open $\iff$ $S = \operatorname{int}(S)$ $\iff$ for every $x \in S$, there is some ball $B(x, r) \subseteq S$, and that the exterior of S is $\operatorname{int}(S^{\complement})$, and that $\partial S = \partial(S^{\complement})$ (last two claims follow by the symmetry in the definitions).

Let us now prove that all balls $B(x,r)$ are open. The key ingredient is the triangle inequality: given any $y \in B(x,r)$, the triangle inequality tells us that $B(y, \frac{r-|y-x|}{2}) \subseteq B(x,r)$, and so by the definition of open at the end of the preceding paragraph, $B(x,r)$ is open (for any point x and radius $r > 0$).

Defining the *closure* of S, denoted $\overline{S}$, to be the union of S with its boundary, i.e. $\overline{S} := S \cup \partial S$, another equivalent formulation for closedness is: S is closed if and only if (denoted $\iff$) $S = \overline{S}$. Observe that $x \in \overline{S} \iff x \in S$ or $x \in \partial S \iff$ every ball $B(x,r)$ intersects nontrivially (i.e. shares at least one point in common) with S (since either $x \in S$, in which case $x \in B(x,r) \cap S$, or $x \in \partial S$, in which case by definition of ∂S there is nontrivial intersection).

As a couple warm-up exercises, let's prove that the closure of any set S is closed, and moreover is the smallest closed set containing S (in the sense that if F is a closed set containing F, then $\overline{S} \subseteq F$). *Sidenote:* I will often use the letters U, V and sometimes O to denote open sets, and the letters F and sometimes C to denote closed sets ("F" because closed in French is "fermé"). Also, I will often use the word "neighborhood" to refer to an open set (generally, I will say "neighborhood of the point x" and one should visualize this as a small open set around x, sort of a setting on which to observe behavior "near" x — yes, "near" is also a key word that should alert the reader to something going on related to neighborhoods, as is the word "local").

Ok, warm-up exercises: given any set $S \subseteq \mathbb{R}^n$, we want to prove that $\overline{S} := S \cup \partial S$ is closed. It suffices to prove that $\partial\overline{S} \subseteq \overline{S}$. Suppose $x \in \partial\overline{S}$; then ANY ball $B(x,r)$ intersects with $\overline{S}$, say at a point $y \in \overline{S}$. If $y \in \partial S$, then we have that $B(x,r)$ is a neighborhood of y (using the fact we proved above that balls are open!), so there is $B(y,r') \subseteq B(x,r)$, and by definition of $y \in \partial S$, there is some point $y' \in B(y,r') \cap S \subseteq B(x,r) \cap S$. If $y \in S$ (recall $y \in \overline{S} \iff y \in \partial S$ or $y \in S$), then $y \in B(x,r) \cap S$. We have just proven that ANY ball $B(x,r)$ intersects S, so indeed $x \in \overline{S}$. But $x \in \partial\overline{S}$ was chosen arbitrarily, so indeed $\partial\overline{S} \subseteq \overline{S}$. This proves that S is closed $\iff$ $S^{\complement}$ is open, because S closed $\implies S = \overline{S} = \partial S \cup \operatorname{int}(S) = \operatorname{ext}(S)^{\complement}$ where $\operatorname{ext}(S)$ is open; and conversely $S = U^{\complement}$ for some open U implies that $S = \partial U \cup \operatorname{ext}(U) = \partial(U^{\complement}) \cup \operatorname{int}(U^{\complement}) = \overline{U^{\complement}}$, hence S is closed.

Second warm-up exercise: suppose S is a set in $\mathbb{R}^n$ and F is a closed set containing S. We want to show that $\overline{S} \subseteq F$. It suffices to show that $\partial S \subseteq F$. Suppose $x \in \partial S$; then every ball $B(x,r)$ intersects nontrivially with S, and hence F because $S \subseteq F$. Reading over the paragraph in which I defined "closure" above, we see that this is exactly what it means for $x \in \overline{F}$. But F is closed, so $F = \overline{F}$. This in particular shows that the closure $\overline{S}$ of a set Suals the intersection of ALL closed sets containing S, and also shows that for any sets $A \subseteq B$, $\overline{A} \subseteq \overline{B}$.

1.1 Metric Spaces

Ok, enough of that. Above, we did everything in $\mathbb{R}^n$, and defined everything in terms of balls $B(x,r)$ defined using the metric/distance function in $\mathbb{R}^n$. That is to say, for vectors $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$, we have a distance function $d(\mathbf{x}, \mathbf{y}) := |\mathbf{x} - \mathbf{y}| = \sqrt{(x_1 - y_1)^2 + \ldots + (x_n - y_n)^2}$ defined from $\mathbb{R}^n \times \mathbb{R}^n \to [0, \infty)$. Recall that we used a very key property of this distance function d, the triangle inequality, to prove that balls were open. Now $\mathbb{R}^n$ is not special; in fact for any space X with such a distance function, we can make the above definitions (as well as talk about things like continuity, using the standard $\epsilon - \delta$ definition). In fact, whenever we use a distance function (more formally called a *metric*) d, we only ever use the following three properties, so we make a general definition using exactly those three properties:

Definition 1.1: Metric space

A set (of points) X equipped with a metric (informally, a "distance function") $d(\bullet, \bullet) : X \times X \to [0, \infty) \subseteq \mathbb{R}$ (also denoted $|\bullet - \bullet|$) satisfying the three properties:

- Non-negativity and identity: $d(x, y) \geq 0$ for all $x, y \in X$ and $d(x, y) = 0 \iff x = y$
- Symmetry: $d(x, y) = d(y, x)$ for all $x, y \in X$
- Triangle inequality (usually most important!): $d(x, z) \leq d(x, y) + d(y, z)$ for any $x, y, z \in X$

2 Continuity

In metric spaces, we can define (local) continuity in the $\epsilon - \delta$ sense (a function $f : X \to Y$ is continuous at x_0, iff for all $\epsilon > 0$, there is some $\delta > 0$ s.t. $|x - x_0| < \delta \implies |f(x) - f(x_0)| < \epsilon$), or using sequences where convergence is defined in terms of the metric (a function is continuous at x_0, iff all sequences $x_n \to x_0$ have that $f(x_n) \to f(x_0)$). The $\epsilon - \delta$ definition actually has some nice visual intuition, that we can make more apparent by rephrasing the definition slightly: for all $\epsilon > 0$, there is $\delta > 0$ s.t. $x \in B(x_0, \delta) \implies f(x) \in B(f(x_0), \epsilon)$. This in turn can be rephrased in terms of images or preimages of functions:

$$f(B(x_0, \delta)) := \{f(x) : x \in B(x_0, \delta)\} \subseteq B(f(x_0), \epsilon),$$

or

$$f^{-1}(B(f(x_0), \epsilon)) := \{x \in X : f(x) \in B(f(x_0), \epsilon)\} \supseteq B(x_0, \delta).$$

So far, this has only dealt with local continuity, i.e. continuity at a point. Global continuity is not so much harder; a function $f : X \to Y$ is continuous (globally) iff at every $x_0 \in X$, for every $\epsilon > 0$, we can find $\delta > 0$ s.t. $f(B(x_0, \delta)) \subseteq B(f(x_0), \epsilon)$ or $B(x_0, \delta) \subseteq f^{-1}(B(f(x_0), \epsilon))$. But take a closer look at the latter definition — it looks an awful lot like the definition of an open set: we have something about a ball around x_0 being contained in a set, and something about "every x_0".

Let us now formalize this intuition: for any $y_0 \in Y$, $f^{-1}(B(y_0, R))$ is an open set, because for any x_0 in this set, by definition of preimage, $f(x_0) \in B(y_0, R)$. But because $B(y_0, R)$ is open, that must mean that there is some $\epsilon > 0$ s.t. $B(f(x_0), \epsilon) \subseteq B(y_0, R)$. The preimage definition of continuity then gives that there is some $\delta > 0$ s.t. $B(x_0, \delta) \subseteq f^{-1}(B(f(x_0), \epsilon)) \subseteq f^{-1}(B(y_0, R))$ (because preimages are monotone: $A \subseteq B \implies f^{-1}(A) \subseteq f^{-1}(B)$), and so $f^{-1}(B(y_0, R))$ is open. Note that this proof did not use any properties of $B(y_0, R)$ other than the fact that it was open, so in fact if f is (globally) continuous, the preimage of *every* open set is open. Moreover, if the preimage of every open set is open, then trivially the previous preimage definition of continuity is satisfied, so it is true that a function $f : X \to Y$ is globally continuous *if and only if* the preimage of every open set is open.

This change of perspective shifts the focus of continuity from distances to open sets — now, as long as we have some special collection of subsets ("open sets") on X and Y, we can define the notion of continuity for functions $f : X \to Y$. However, the collections should have some sort of structure, in order to align better with our intuitions for continuity.

2.0.1 Hunt for Properties

Now how did we define "open" in $\mathbb{R}^n$? $U \subseteq \mathbb{R}^n$ was open iff for all $x \in U$, there was some B_x open, containing x s.t. $B_x \subseteq U$, where B_x was a ball. More generally, U is open if and only if each $x \in U$ has some open set U_x s.t. $U_x \subseteq U$. This property is possibly the **most fundamental property** of an open set (in fact, we literally used it several paragraphs ago, in the proof equating two definitions of continuity), and so we require that any such "special collection" $\mathscr{U}$ of "open sets" be subject to the condition that $\mathscr{U} = \{\bigcup_{a \in A} U_a : U_a \in \mathscr{U}, A \text{ some indexing set}\}$, i.e. $\mathscr{U}$ is closed under arbitrary unions.

What other properties should we want? Well, we build these collections for the purpose of defining global continuity, but in $\mathbb{R}$, recall that we started with local continuity first. So let's see if we can capture the idea of local continuity here, piggybacking off of our definition of global continuity: f is continuous at some $x \in X$ if there is some open set $U \subseteq X$ s.t. $x \in U$, and $f\big|_U : U \to Y$ is continuous. This seems ok at first glance, but looking more closely, what exactly does a restriction map being continuous even mean?

For example, taking $X = \mathbb{R}$, and the function $f : \mathbb{R} \to \mathbb{R}$ defined by $f(x) = x$. This is obviously continuous as the inverse image of any open set is exactly that open set, but taking its restriction $f\big|_{[0,1]} : [0,1] \to \mathbb{R}$, we see that $(f\big|_{[0,1]})^{-1}((-1/2, 3/2)) = [0,1]$ and $(f\big|_{[0,1]})^{-1}((-1/2, 1/2)) = [0, 1/2)$, which are both not open in the usual sense on $\mathbb{R}$. Although this is not so good, it should nonetheless not come as any surprise — $[0,1]$ has boundary points 0 and 1, at which every neighborhood doesn't lie completely in $[0,1]$; i.e. at 0 or 1, no matter how far we "zoom in", we will always be able to tell whether or not we are on the restriction or not.

In contrast, if we were to restrict to any open set $U \subseteq \mathbb{R}$, zooming in far enough, we will not be able to tell, which reflects the fact that on $\mathbb{R}$, the restriction to an open set $f\big|_U$ is continuous with respect to the same open sets as f originally. More formally, this is because $(f\big|_U)^{-1}(V) = \{x \in U : f(x) \in V\} = U \cap f^{-1}(V)$, which if $U, V \subseteq \mathbb{R}$ are open, is open too because an intersection of two open sets is open (if we have $B(x, r_1) \subseteq U_1$, and $B(x, r_2) \subseteq U_2$, then $B(x, \min\{r_1, r_1\}) \subseteq U_1 \cap U_2$). Thus, the continuity of any local maps with respect to the same open sets as the global map hinges on the requirement that the intersection of two open sets remains open. We call this the intersection axiom, and we ask that any "special collection" $\mathscr{U}$ of open sets satisfy this axiom.

Anyways, if we have $f : X \to Y$ continuous at every point, i.e. at every $x \in X$ there is some open U_x s.t. $f\big|_{U_x}$ is continuous, then f is continuous: for any open $V \in Y$, at every $x \in f^{-1}(V)$, local continuity gives that $(f\big|_{U_x})^{-1}(V)$ is an open neighborhood of x in $f^{-1}(V)$, and so $f^{-1}(V) = \bigcup_{x \in f^{-1}(V)} (f\big|_{U_x})^{-1}(V)$ is open.

2.0.2 Summary

To summarize, we needed the intersection axiom to have that restrictions of continuous functions to any open set remain continuous with respect to the same collection of open sets (i.e. going from global to local), and we needed the union axiom to have that continuous local restrictions remain globally continuous (i.e. going from local to global). Viewed in this way, the union and intersection axioms are really just two sides of the same coin. Finally, to ensure that constant functions are continuous, we ask that X and $\varnothing$ be open too, i.e. elements of $\mathcal{U}$. Any collection $\mathcal{U}$ of subsets of X with all three of these properties/axioms will be called a *topology* on X. Alternatively, we can call the pair $(X, \mathcal{U})$ a *topological space*. If the reader takes just one thing from this section, it should be this: topological spaces are those equipped with the barest minimum needed to guarantee that continuity and continuous functions are "nice"; or more succinctly: the heart of topology is continuity.

2.1 Subspace Topology

We talked about restricting continuous functions on X to open sets of X, but what about other sets? Well for any continuous function $f : X \to Y$ (with respective topologies $\mathcal{T}_X, \mathcal{T}_Y$) and subset $S \subseteq X$, the restriction $f|_S : S \to Y$ should of course remain continuous. But the inverse image of an arbitrary open set $V \in \mathcal{T}_Y$ is $f|_S^{-1}(V) = S \cap f^{-1}(V)$, which may not be an open set in $\mathcal{T}_X$. Thus, we will need to equip S with a new topology $\mathcal{T}_S$, not exactly just the open sets of $\mathcal{T}_X$, but such that $f|_S : S \to Y$ is continuous with respect to the new topology on S.

Well, given the form of the inverse images via $f|_S$, an obvious idea is to define $\mathcal{T}_S := \{S \cap U : U \in \mathcal{T}_X\}$. Clearly, this fulfills our goal of having $f|_S$ be continuous w.r.t. $\mathcal{T}_S$, but we do need to check that it is actually a topology:

- arbitrary unions: we have $\bigcup_{i \in I}(S \cap U_i) = S \cap (\bigcup_{i \in I} U_i)$, where indeed $(\bigcup_{i \in I} U_i) \in \mathcal{T}_X$;
- finite intersections: we have $(S \cap U_1) \cap (S \cap U_2) = S \cap (U_1 \cap U_2)$, where indeed $(U_1 \cap U_2) \in \mathcal{T}_X$;
- and $\varnothing$ and S: indeed $\varnothing = S \cap \varnothing$ and $S = S \cap X$.

Thus, for any subset $S \subseteq X$ for a topological space $(X, \mathcal{T}_X)$, $\mathcal{T}_X$ induces the subspace topology $\mathcal{T}_S := \{S \cap U : U \in \mathcal{T}_X\}$ on S.

2.2 Homeomorphisms

Let us now turn our attention to considering when two topological spaces are equivalent. I will provide two explicit examples to guide our inquiry: consider $B(0,1)$ and $B(0,2)$ in $\mathbb{R}^2$, and also $[0,1)$ and the unit circle $S^1 \subseteq \mathbb{R}^2$. It should be intuitively clear that $B(0,1)$ and $B(0,2)$ are extremely similar, while $[0,1)$ and S^1 are extremely different. Note that we can easily provide a bijection between each pair of spaces, namely $f : B(0,1) \to B(0,2)$ defined by $x \mapsto 2x$ and $g : [0,1) \to S^1$ defined by $t \mapsto (\cos 2\pi t, \sin 2\pi t)$.

However, the difference between f and g is that g somehow "adds in different structure", in that it "connects" the "endpoints" of the interval into a closed loop. More formally, g does not preserve the topology, because $[0, \frac{1}{2})$ is open in the subspace topology of $[0,1) \subseteq \mathbb{R}^1$, but "wrapping" this interval around the circle, $[0, \frac{1}{2})$ is mapped to $\{(\cos 2\pi t, \sin 2\pi t) : t \in [0, \frac{1}{2})\}$, or the open upper half circle union the point $(1,0)$, which is NOT open in the subspace topology of $S^1 \subseteq \mathbb{R}^2$. Basically, this reflects the fact that we fundamentally changed the shape going from an interval to a circle, because what used to be the endpoint is now connected into a loop, and hence has neighbors "on both sides".

In contrast, $f : B(0,1) \to B(0,2)$ (henceforth referred to respectively as X and Y) defined above DOES preserve the topology, in the sense that the bijection between the sets X, Y induces a bijection between the topologies on the sets $\mathcal{T}_X, \mathcal{T}_Y$; i.e. $\tilde{f} : \mathcal{T}_X \to \mathcal{T}_Y$ defined by $U \mapsto f(U)$ is a bijection between $\mathcal{T}_X, \mathcal{T}_Y$. To prove that $\tilde{f}$ is a bijection, we just need to find a (two-sided) inverse. Consider $\tilde{g} : \mathcal{T}_Y \to \mathcal{T}_X$ defined by $V \mapsto f^{-1}(V)$. We check that $\tilde{g} \circ \tilde{f} = \mathrm{id}_{\mathcal{T}_X}$ and $\tilde{f} \circ \tilde{g} = \mathrm{id}_{\mathcal{T}_Y}$. Well, we have $f^{-1}(f(U)) := \{x \in X : f(x) \in f(U)\}$; obviously this contains U, and the $(\subseteq)$ direction comes from the fact that $f(x) \in f(U) \implies$ there is $u \in U$ s.t. $f(x) = f(u)$, which implies that $x = u$ by injectivity of f. Similarly, $f(f^{-1}(V)) := \{f(x) : x \in f^{-1}(V)\}$, where $x \in f^{-1}(V) \iff f(x) \in V$, and so indeed $\{f(x) : x \in f^{-1}(V)\} = V$.

We actually missed a very important step above — we have to check that $\tilde{f}$ indeed goes from $\mathcal{T}_X$ to $\mathcal{T}_Y$, i.e. that $f(U) \in \mathcal{T}_Y$ for $U \in \mathcal{T}_X$. In other words, we want to check that f is an open map (maps open sets to open sets), which is equivalent to checking that f^{-1} is continuous (because f continuous $\iff f^{-1}$ is an open map)! Our explicit map f above has the obviously continuous inverse $x \mapsto \frac{1}{2}x$, but anyways in general we define:

Definition 2.1: Homoemorphic

We say that two topological spaces are equivalent, or *homeomorphic*, exactly when there is a bijection between the sets that induces a bijection between the topologies, which happens if and only if that bijection is continuous and has a continuous inverse. I will denote it by $(X, \mathcal{T}) \simeq (Y, \mathcal{U})$, or if the topologies are understood from context, $X \simeq Y$.

Last remarks: although it was true that our $g : [0,1) \to S^1$ was not a homeomorphism, that does not prove that $[0,1)$ and S^1 are not homeomorphic. There are a variety of ways to show this is true, but I will leave one as an exercise in Section 5.

2.3 Metrizability and Bases

It should be noted again that continuity and open sets were originally defined on metric spaces. That is to say, for every metric space, we have an induced topological space where the topology is generated by the open balls of the metric space. Of course, it goes without saying that metric spaces are much more intuitive and much easier to work with than topological spaces, so an interesting question to consider is the following: which topological spaces could be generated by the open balls of some metric? That is to say, which topological spaces $(X, \mathscr{T})$ can we equip with a metric $\rho : X \times X \to [0, \infty)$ s.t. the topology induced by ρ, i.e. $\mathscr{M} = \{\bigcup B_\rho(x, r) : x \in X, r \in [0, \infty)\}$ is equal to $\mathscr{T}$?

To prove that $\mathscr{M} = \mathscr{T}$, we must show that any open set in $\mathscr{M}$ is in $\mathscr{T}$, and that any open set in $\mathscr{T}$ is also in $\mathscr{M}$. But because all open sets in $\mathscr{M}$ are of the form $\bigcup_{i \in I} B_\rho(x_i, r_i)$ for some index set I, it suffices to show that (1) all $B_\rho(x, r)$ are in $\mathscr{T}$; and (2) that every open set U in $\mathscr{T}$ can be written as $\bigcup_{x \in U} B_\rho(x, r_x)$ — i.e. that for every $x \in U$, we can find some $r_x > 0$ s.t. $B_\rho(x, r_x) \subseteq U$. Showing that (1) all $B_\rho(x_0, r)$ are in $\mathscr{T}$ can in turn be shown by proving that for any $x \in B_\rho(x_0, r)$, there is some $U_x \in \mathscr{T}$ containing x s.t. $U_x \subseteq B_\rho(x_0, r)$.

This simplification to only considering the open balls (w.r.t. the metric ρ) is due to the fact that the open balls from what we call a *basis* of $\mathscr{M}$ — all open sets of $\mathscr{M}$ can be written as some union of such open balls. In fact, if $\mathscr{T}$ also has some basis, say $\{U_b\}_{b \in B}$ for some index set B, then we can simplify to that basis like we simplified to the open balls: we only need to show that (1) for any ball $B_\rho(x_0, r)$, for any x in that ball, there is some $U_x \in \{U_b\}_{b \in B}$ containing x s.t. $U_x \subseteq B_\rho(x_0, r)$; and (2) any $U_b \in \{U_b\}_{b \in B}$, for any $x \in U_b$, there is $r_x > 0$ s.t. $B_\rho(x, r_x) \subseteq U_b$. Also, note that the balls with radii $\frac{1}{n}$, $n \in \mathbb{N}$ also form a basis of $\mathscr{M}$, and so we can simplify a bit further by only considering radii of the form $\frac{1}{n}$.

It may also be interesting to consider what kinds of sets can form a basis. Let's say we have a collection $\mathscr{B}$ that is a basis for the topology $\mathscr{T}$. That means that $\mathscr{T} = \{\bigcup_{B \in \mathscr{B}'} B : \mathscr{B}' \subseteq \mathscr{B}\}$. If $\mathscr{T}$ truly is a topology, then $\mathscr{T}$ satisfies the three axioms, and so

- $\varnothing \in \mathscr{T}$ (obviously satisfied), and $X \in \mathscr{T} \iff \bigcup_{B \in \mathscr{B}} B = X$;
- arbitrary unions of open sets in $\mathscr{T}$ are still in $\mathscr{T}$ (obviously satisfied);
- and the intersection of two open sets in $\mathscr{T}$ to still be open: $\bigcup_{i \in I} B_i \cap \bigcup_{j \in J} B'_j = \bigcup_{(i,j) \in I \times J} (B_i \cap B'_j)$. Unfortunately, $B_i \cap B'_j$ may not be in $\mathscr{B}$; like in $\mathbb{R}^n$, the intersection of two balls may not be a ball. However, we don't need it to be! The intersection is open if all the $B_i \cap B'_j$ are the union of some B's. In fact, if it is the union of some $B's$, it must be $= \bigcup_{B \subseteq B_i \cap B'_j} B$, which happens if and only if all $x \in B_i \cap B'_j$ have some $B_x \in \mathscr{B}$ containing x s.t. $B_x \subseteq B_i \cap B'_j$.

So we found that $\{\bigcup_{B \in \mathscr{B}'} B : \mathscr{B}' \subseteq \mathscr{B}\}$ is a topology if $\bigcup_{B \in \mathscr{B}} B = X$, and if for all $B_1, B_2 \in \mathscr{B}$ and for all $x \in B_1 \cap B_2$, there is some $B_3 \in \mathscr{B}$ s.t. $B_3 \subseteq B_1 \cap B_2$ and $x \in B_3$. The other direction is trivial.

2.4 Product Topology

We first talk about finite products. The obvious candidate for the topology on a finite product $\prod_{i=1}^n X_i$ is the topology generated by basis of Cartesian products of open sets, i.e. sets of the form $\prod_{i=1}^n U_i$ for open $U_i \subseteq X_i$. Note that maps to a product space can always be decomposed into maps to each of the components — think about how in calculus any vector-valued function f from say $\mathbb{R}^3 \to \mathbb{R}^3$ can always be written as $f = (f_1, f_2, f_3)$ where $f_i(\mathbf{x}) \in \mathbb{R}$ is simply the ith component of $f(\mathbf{x}) \in \mathbb{R}^3$. It was true in the $\mathbb{R}^3$ example that f is continuous if and only if the component functions f_i were all continuous, and in fact that remains true in these general topological spaces.

Question 2.1: Exercise

Prove this above assertion, that for any space X, $f = (f_1, \ldots, f_n) : X \to \prod_{i=1}^n X_i$ is continuous if and only if $f_i : X \to X_i$ is continuous for all $i \in [n]$. See Theorem 2.3 below for a solution.

2.4.1 Infinite Products

Unfortunately in the infinite case, if we continue defining the topology in terms of the basis of Cartesian products of open sets (called the *box topology*), this nice property about continuity of functions with a product as codomain being equivalent to continuity of each component function does not hold for infinite products.

Example 2.2: Continuity fails for infinite products under box topology

Defining $f : \mathbb{R} \to \prod_{i=1}^\infty \mathbb{R}$ by $x \mapsto (x, x, \ldots)$, i.e. $f = (f_1, f_2, \ldots)$ where $f_i : \mathbb{R} \to \mathbb{R}$ is the identity function (obviously continuous), it is not true that f is continuous, since $f^{-1}(\prod_{i=1}^\infty(-\frac{1}{n}, \frac{1}{n}))$ (where $\prod_{i=1}^\infty(-\frac{1}{n}, \frac{1}{n})$ is an open set in the box topology) is $\{0\}$.

One obvious solution is to define a new topology (the *product topology*) where the open sets are just the FINITE Cartesian products of open sets (i.e. all but finitely many sets in the product ranging over $i \in I$ are exactly X_i, not any proper open subset of it); this avoids any "infinity problems". But how do we know this is the "right" definition, or the most "general"? Well, in fact (Theorem 3.37 in Lee [6]):

Theorem 2.3: Unique topology s.t. continuous iff component functions are continuous

The product topology is the unique topology on $\prod_{i \in I} X_i$ s.t. for any topological space X, $f : X \to \prod_{i \in I} X_i$ is continuous if and only if the component functions $f_i : X \to X_i$ are continuous

Proof: we first check that it IS a topology for which the if and only if condition holds. ($\implies$): essentially boils down to showing that the projection maps $\pi_j : \prod_{i \in I} X_i \to X_j$ are all continuous (for all $j \in I$), since $f_i := \pi_i \circ f$, and a composition of continuous functions remains continuous. To check continuity of the π_j, we just see that the inverse image $\pi_j^{-1}(U_j)$ for open $U_j \subseteq X_j$ is $\prod_{i \in I} S_i$ where $S_i := X_i$ for $i \neq j$ and $S_j := U_j$, which is indeed open in the product topology, as it has all but finitely many (in fact, exactly one) sets in the product equal to X_i.

($\impliedby$): to prove that f is continuous, we show that the inverse image $f^{-1}(\prod_{i \in I} S_i)$ is open in X, where $S_i := X_i$ for $i \in I \setminus \tilde{I}$ and $S_i := U_i$ open in X_i for $i \in \tilde{I}$, where $\tilde{I} \subseteq I$ is a finite set. This suffices since such finite Cartesian products form a basis of the product topology, i.e. all open sets in the product topology are just unions of the basis sets, and an inverse image of a union is the union of the inverse images. Well, we see that $f^{-1}(\prod_{i\in I} S_i) = \{x \in X : f(x) \in \prod_{i\in I} S_i\} = \{x \in X : f_i(x) \in S_i, i \in I\} = \bigcap_{i\in I} f_i^{-1}(S_i) = \bigcap_{i\in I\setminus\tilde{I}} f_i^{-1}(X_i) \cap \bigcap_{i\in\tilde{I}} f_i^{-1}(U_i) = X \cap \bigcap_{i\in\tilde{I}} f_i^{-1}(U_i) = \bigcap_{i\in\tilde{I}} f_i^{-1}(U_i)$. Because $\tilde{I}$ is a finite set and each $f_i^{-1}(U_i)$ is open in X by the continuity of the component functions f_i, and finite intersections of open sets remain open, $f^{-1}(\prod_{i\in I} S_i)$ is indeed open in X.

As for uniqueness, denote $\Pi := \prod_{i\in I} X_i$ with the product topology $\mathscr{P}$, and suppose $\mathscr{P}'$ is another topology on Π satisfying the "if and only if" condition. Defining the identity map $\mathrm{id}_\Pi : \Pi \to \Pi$ mapping $(x_i)_{i\in I} \mapsto (x_i)_{i\in I}$, we can talk about continuity by specifying which topology to consider. Note that $\mathrm{id}_{\Pi,\mathscr{P}} : (P, \mathscr{P}) \to (P, \mathscr{P})$ and $\mathrm{id}_{\Pi,\mathscr{P}'} : (P, \mathscr{P}') \to (P, \mathscr{P}')$ are continuous (in fact homeomorphisms since $U \in \mathscr{P}$ is send to $U \in \mathscr{P}$, similar with $\mathscr{P}'$), and so by the "if and only if" condition, all components $(\mathrm{id}_{\Pi,\mathscr{P}})_i, (\mathrm{id}_{\Pi,\mathscr{P}'})_i$ are continuous. But since $(P, \mathscr{P})$ satisfies the "if and only if" condition and we just proved that $(\mathrm{id}_{\Pi,\mathscr{P}'})_i : (P, \mathscr{P}') \to X_i$ are all continuous, we have that $\mathrm{id}_{\Pi,\mathscr{P}'\to\mathscr{P}} : (P, \mathscr{P}') \to (P, \mathscr{P})$ is continuous. Switching "$\mathscr{P}$" and "$\mathscr{P}'$", we get that $\mathrm{id}_{\Pi,\mathscr{P}\to\mathscr{P}'} : (P, \mathscr{P}) \to (P, \mathscr{P}')$ is also continuous. These two maps are inverses of each other, so indeed we have shown that $(P, \mathscr{P}) \simeq (P, \mathscr{P}')$ are homeomorphic, thereby giving us the desired uniqueness. ■

We also have the following:

Theorem 2.4: Unique topology s.t. projections are continuous

The product topology is the unique weakest (i.e. with fewest open sets) topology on $\prod_{i\in I} X_i$ s.t. the projection maps $\pi_j : \prod_{i\in I} X_i \to X_j$ are all continuous (for all $j \in I$).

Proof: see ($\implies$) of the above theorem. The "unique weakest" part is just proving that the projection maps being continuous force finite Cartesian products of open sets to be open. Well, if π_j is continuous, then $\pi_j^{-1}(U_j)$ for any open $U_j \subseteq X_j$ must be open, but by the definition of the projection map $\pi_j^{-1}(U_j) = \prod_{i\in I} S_i$ where $S_i := X_i$ for $i \neq j$ and $S_j := U_j$. Since a topology is closed under finite intersections, finite intersections of these sets must also be in the topology, so indeed all finite Cartesian products of open sets MUST be in the topology. ■

3 Manifolds in $\mathbb{R}^n$

In the above section, we built up the idea of topological spaces from a couple core tenets of continuity, a concept originally introduced on the setting of a metric space. The topological reformulation/redefinition of continuity allowed us a new perspective, telling us that continuity is not fundamentally a property about distances (something we intuitively think of as rigid, and geometric), but rather about open sets (things that we consider more "fuzzy" or more "blobby"). This ability to translate rigid, geometric notions to more "fuzzy" or "amorphous" ones is in fact topology's greatest strength (and the reason why it's interesting), but unfortunately it means that topology must become a great deal more abstract (which makes it harder to learn). The aim of the rest of this paper is to try to come up with a topological reformulation of another idea that is originally extremely rigid and geometric: the idea of a manifold in Euclidean space.

We start this discussion on manifolds by assuming the reader is already familiar with the definition of a smooth k-manifold in $\mathbb{R}^n$ ($k \leq n$) given at the end of Chapter 3.3 of Folland's *Advanced Calculus*. I will give a quick review of this definition, and extend upon it a little bit, borrowing heavily from pgs. 3-7 of Arun Debray's notes for the class Math 382d at the University of Texas at Austin: a smooth k-manifold in $\mathbb{R}^n$ is a subset $X \subseteq \mathbb{R}^n$ that satisfies the following equivalent properties:

(a) (Locus definition) Locally, X is the level set of a smooth map $\mathbf{f} : \mathbb{R}^n \to \mathbb{R}^{n-k}$.

(b) (Parametric definition) X is locally the image of a smooth map, i.e. for every $\mathbf{p} \in X$, there is a neighborhood U of $\mathbf{p}$ and a smooth $\mathbf{f} : \mathbb{R}^k \to \mathbb{R}^n$ with full rank such that $(\operatorname{im} \mathbf{f}) \cap U = X \cap U$.

(c) (Implicit function definition) For every $\mathbf{p} \in X$, there is a neighborhood U of $\mathbf{p}$ where one can write $n-k$ variables as smooth functions of the remaining k variables, i.e. there is a neighborhood $V \subseteq \mathbb{R}^k$ and a smooth $\mathbf{g} : V \to \mathbb{R}^{n-k}$ such that $X \cap U = \{(\mathbf{x}, \mathbf{g}(\mathbf{x})) : \mathbf{x} \in V\}$.

As for why these properties are equivalent or where they come from, it may help to consider the following central/fundamental theorems regarding smooth functions $\mathbf{f} : \mathbb{R}^n \to \mathbb{R}^m$ and $\mathbf{a} \in \mathbb{R}^n$ such that the derivative matrix evaluated at $\mathbf{a}$, i.e. $d\mathbf{f}|_{\mathbf{a}}$, has full rank:

(a) (Inverse function theorem) If $n = m$, then there is a neighborhood U of $\mathbf{a}$ such that $\mathbf{f}|_U$ is invertible with a smooth inverse $(\mathbf{f}|_U)^{-1} : V \to U$, where $V \subseteq \mathbb{R}^m$ is some neighborhood of $\mathbf{f}(\mathbf{a})$.

(b) (Implicit function theorem) Let $\mathbf{c} = \mathbf{f}(\mathbf{a})$, and define the level set L to be $L = \{\mathbf{x} \in \mathbb{R}^n : \mathbf{f}(\mathbf{x}) = \mathbf{c}\}$. Said another way, $L = f^{-1}(f(\mathbf{a}))$. If $n \geq m$, there is a neighborhood U of $\mathbf{a}$ such that $U \cap L$ is the graph of some smooth function $\mathbf{g} : \mathbb{R}^{n-m} \to \mathbb{R}^m$(or more specifically, $\mathbf{g} : S \to \mathbb{R}^m$ where $S \subseteq \mathbb{R}^{n-m}$ is a neighborhood of the projection of $U \cap L$ onto $n - m$ of the n variables, so we get that $U \cap L = \{(\mathbf{x}, \mathbf{g}(\mathbf{x})) : \mathbf{x} \in S\}$, up to a permutation of the indices).

(c) (Immersion theorem) If $n \leq m$, there is a neighborhood U of $\mathbf{a}$ such that $\mathbf{f}(U) \subseteq \mathbb{R}^m$ is the graph of a smooth $\mathbf{g} : \mathbb{R}^n \to \mathbb{R}^{m-n}$, again up to a permutation of the indices (as above in (b), $\mathbf{g}$ is not really expected to be defined on all of $\mathbb{R}^n$, but rather just some neighborhood in $\mathbb{R}^n$).

Of course, the reader is encouraged to take a look at Folland's text and Debray's notes for further information. It may also be of interest to the reader to recall why Folland was studying such manifolds anyways — he later works with manifolds in his chapters on integration (more specifically, line integrals and surface integrals). This may give some context to such objects; smooth manifolds in $\mathbb{R}^n$ are settings on which integration can be done.

Anyways, with this definition in mind, allow me to restate the aim of this paper: we want to come up with a topological reformulation of the idea of a manifold, i.e. the hope is to find some topological properties necessary and sufficient for a topological space to be deserving of the title of "manifold". Now the first major problem is this: the above discussion on manifolds was on *smooth* manifolds, where derivatives played a huge role. Topological spaces however are built for continuity, NOT differentiability. There is a way to talk about differentiability and derivatives with topological spaces, but that's a discussion for another time altogether. But, even if we can not talk about differentiability for now, there is still much we *can* talk about regarding manifolds; we just have to take inspiration from the above definition of a smooth k-manifold in $\mathbb{R}^n$, and define a more general class of manifolds (i.e. ones that may not be smooth) without relying on derivatives.

3.1 Locally Euclidean

Our path to a more general definition of manifold based solely on continuity starts with an elaboration on definition (b) of a smooth k-manifold X, i.e. that X is locally the image of smooth map from $\mathbb{R}^k$ — (b) can be strengthened by noting that not only is there a smooth map $\mathbb{R}^k \to U \cap X$, there's also an smooth inverse $U \to \mathbb{R}^k$ (not $U \cap X \to \mathbb{R}^k$, because we need U an open set in $\mathbb{R}^n$ to define smoothness). This is proven in page 8 of Arun Debray's notes mentioned above, a proof that I will copy here for ease of reference and sake of completeness.

Let X be a k-dimensional manifold in $\mathbb{R}^n$, so for any $p \in X$, there's a map ϕ from the neighborhood of the origin in $\mathbb{R}^k$ to a neighborhood of p in X, where $\phi(0) = p$ and $\mathrm{d}\phi|_0$ has rank k. We'd like a local inverse to ϕ, which we'll call F; it's a map from a neighborhood of $\mathbb{R}^n$ to a neighborhood of $\mathbb{R}^k$. We'd like F to be smooth, and we want $F \circ \phi = \mathrm{id}|_{\mathbb{R}^k}$ By permuting coordinates, we can assume that the first k rows of $\mathrm{d}\phi$ are linearly independent. That is, $\mathrm{d}\phi|_0$ has block form $\binom{A}{B}$ where A is invertible. Then, define $\tilde{\phi} : \mathbb{R}^k \times \mathbb{R}^{n-k} \to \mathbb{R}^n$ sending $(x, y) \mapsto \phi(x) + (0, y)$, so that $\tilde{\phi}((\mathbf{x}, \mathbf{0})) = \phi(\mathbf{x})$, ϕ and $\widetilde{\phi}$ fit into the following diagram:

$$\mathbb{R}^k \xhookrightarrow{\mathbf{x} \mapsto (\mathbf{x}, \mathbf{0})} \mathbb{R}^n \xrightarrow{\tilde{\phi}} \mathbb{R}^n, \qquad \mathbb{R}^k \xrightarrow{\phi} \mathbb{R}^n$$

Thus, by the chain rule,

$$\mathrm{d}\tilde{\phi}\big|_0 = \left(\begin{array}{c|c} A & 0 \\ \hline B & I \end{array}\right)$$

so $\mathrm{d}\tilde{\phi}|_0$ has full rank! Thus, in a neighborhood of p, it has an inverse, and certainly the inclusion $\mathbb{R}^k \hookrightarrow \mathbb{R}^n$ has a left inverse π (projection onto the first k coordinates), so we can let $F = \pi \circ \widetilde{\phi}^{-1}$, because

$$F \circ \phi(x) = F \circ \widetilde{\phi}(x, 0) = \pi \circ \widetilde{\phi}^{-1} \circ \widetilde{\phi}((x, 0)) = \pi(x, 0) = x$$

Likewise, $\phi \circ F = \mathrm{id}|_X$, since every point in our neighborhood is in the image of ϕ.

This all means that smooth k-manifolds are "like" $\mathbb{R}^k$ locally in a very strong sense; in fact, if we "forget" about smoothness and instead focus solely on continuity (in particular we can focus solely on the restriction to X $F|_X$ instead of F on a neighborhood of p in $\mathbb{R}^n$), we see that for local patches of X, $U \cap X$, we have a bi-continuous function between $U \cap X$ and $\mathbb{R}^k$ — that is, $U \cap X$ and $\mathbb{R}^k$ are homeomorphic. As we've said before, being homeomorphic means that two spaces are topologically equivalent. This makes very good ground for the definition of a general k-manifold in $\mathbb{R}^n$: a subset $X \subseteq \mathbb{R}^n$ is a k-manifold if locally, X is homeomorphic to $\mathbb{R}^k$. We call this property *locally k-Euclidean.*

3.2 Goal

Goal

Thus to again restate the **goal of this paper** — given some arbitrary topological space with some properties, is it homeomorphic (i.e. topologically equivalent) to some k-manifold (i.e. locally k-Euclidean subset) in $\mathbb{R}^N$? By the definition of homeomorphic, proving something like this requires two main parts: (1) we need to find a continuous map $f : X \to \mathbb{R}^N$ that is injective (so that we have a continuous bijection from X to $\operatorname{im} f \subseteq \mathbb{R}^N$); and (2) we need f to have a continuous inverse from $\operatorname{im} f \to X$.

4 Extension Theorems and Partitions of Unity

As perhaps you've noticed by now, a central focus of topology are the neighborhoods around some point; that is studying local properties of sets is a core tenet of topology. However, global phenomenon (and perhaps more importantly, their relations to local phenomenon) do also play a large role in the subject; recall how our global definition of continuity was based on the local phenomenon of neighborhoods of the codomain pulling back via f to neighborhoods in the domain. It therefore should not be be unexpected that it will be useful (or at least interesting) to try to tackle the goal of finding an injective continuous function $X \to \mathbb{R}^n$ by somehow finding some local continuous functions (defined on say a collection of open subsets forming an open cover), and stitching them together into some global function that is "close", if not equal, to the local functions on their respective open sets of definition.

Ok. So with this in mind, how might we go about doing this? Well, on each open set, we have a continuous function, but it is not defined outside that set, and it seems hard to try to extend it in some way. However, if we give up the condition that the global function agrees exactly with the local functions on the entire open set, we can make it easier to extend (to a global continuous function) by zeroing the function "close to" the boundary of the set so it can be easily extended outside the original open set of definition (just 0 everywhere outside), in such a manner that the function remains the same on the "interior" of the open set, and that the transition between "function" and "0" is continuous. We can accomplish this by multiplying by a bump function that is continuous and 1 on the "interior" of the open set and 0 "close to" the boundary. Of course, this is quite hand-wavy, so we flesh out the details below. First, I want to address why I want it to be 0 "close to" the boundary, and not just *on* the boundary — consider the following bump function and continuous function $-\frac{1}{x-2}$ on $(-2, 2)$, and their product:

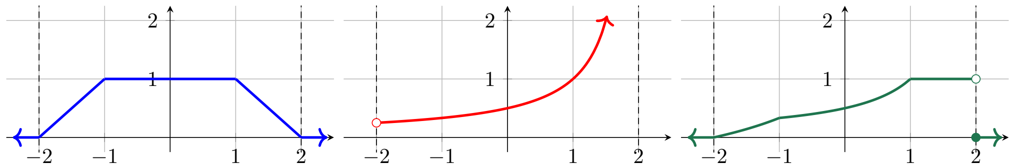

As you can see, although $-\frac{1}{x-2}$ is continuous on $(-2, 2)$, the bump function is not 0 "close to" the boundary, and consequently their product does not transition continuously to 0 as it goes outside $(-2, 2)$. That is to say, having the bump function be 0 "close to" the boundary guarantees a continuous transition to 0 outside of U.

Now if a function on U (open) is continuous and equal to some constant c on some subset $A \subseteq U$, like in the picture on the left below (where the regularly dashed yellow line indicates that U is open, and the irregularly dashed/dotted red line means that A could be open, close, or anything in between):

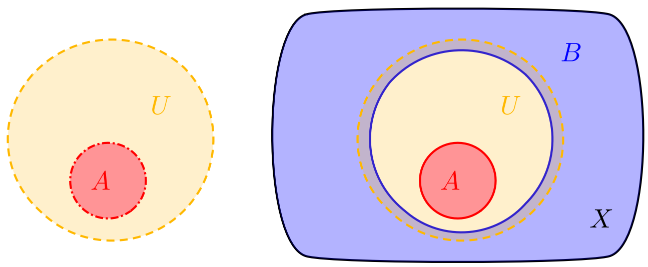

then hopefully it's intuitively clear that the function must $= c$ on $\overline{A}$ as well (as long as $\overline{A}$ doesn't exceed U; if it does we can't say anything because we only know the function is continuous on U). This "intuitively clear" observation can be proven easily by observing that the definition of continuity forces the inverse image of a closed subset $f^{-1}(\{c\})$ to also be closed, and $\overline{A}$ is the smallest closed set containing A. With this in mind, we can more precisely state that the bump function needs to be 1 on A, a closed subset $\subseteq U$, and 0 on a closed set B s.t. for all $x \in \partial U$, there is some $U_x \supseteq \{x\}$ open s.t. $U_x \subseteq B$ (satisfying 0 "close to" boundary, not just on boundary). This is equivalent to saying that B is closed and contains an open set containing $X \setminus U$. This is pictured above on the right.

Notation: going to all this trouble about requiring f defined on U to be 0 "close to" the boundary is equivalent to saying that the CLOSURE of $f^{-1}((0, \infty))$, i.e. the set $\overline{\{x \in X : f(x) > 0\}}$, also known as the **support** denoted $\operatorname{supp} f$, is contained in U; this is because if $\operatorname{supp} f \subseteq U$, we can just take our "$B$" to be $\overline{X \setminus \operatorname{supp} f}$ (a closed set that contains the open set $X \setminus \operatorname{supp} f$ that contains $X \setminus U$), and if f is 0 on B closed s.t. $B \supseteq O \supseteq X \setminus U$, then $\{x \in X : f(x) > 0\} \subseteq X \setminus B \subseteq X \setminus O$, and taking closures yields $\operatorname{supp} f \subseteq \overline{X \setminus O} = X \setminus O \subseteq U$.

On what spaces/what properties do we need to guarantee/prove the existence of such a function? Because we are trying to find properties that imply this, pretend we are in Euclidean space and keep track of what properties we use.

4.1 Urysohn's Lemma

We want to construct some continuous function f that "starts" at 1 on A and decreases to 0 as you get away from A. Let's try to do this by having a sequence of open sets indexed by rationals U_r where $r \in \mathbb{Q} \cap (0,1)$, where if $p > q$, then $U_p \subseteq U_q$ and the function at some x is $> r$ if $x \in U_r$, but $< r$ if it is not in U_r:

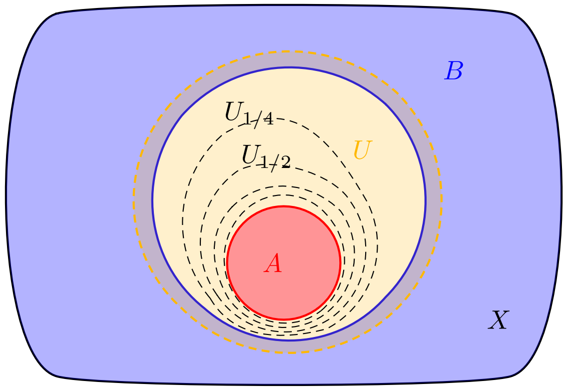

Actually, we should be a little more careful in our inclusion condition because we don't want all the sets to be the same, or have this sort of scenario:

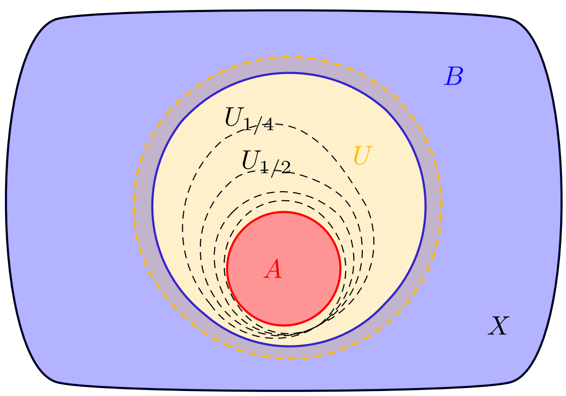

where the sets do have some sort of "gradient" at the top, but at the bottom are so "bunched up" that there's a sharp jump. Let's revise it so that $p > q \implies \overline{U_p} \subseteq U_q$ to actually allow some "breathing room". Ok. Now let's try to prove this is continuous: we just need to prove that the inverse image of the basis sets in $[0,1]$ (i.e. the intervals of the form $(a,b), [0,b), (a,1], [0,1]$) are open in X.

- $f^{-1}([0,1]) = X$
- $f^{-1}((a,1]) = \bigcup_{r \geq a} U_r$. At first glance, one may think that the $r \geq a$ should correspond with $[a,1]$, so let us dispel that myth before we move on. Say $x \notin \bigcup_{r \geq a} U_r$. Can $f(x) \in [a,1]$? Well, yes! If x is in all the U_α for $\alpha < a$, then $x > \alpha$ for all $\alpha < a$, so $x \geq a$. This in particular gives that $f^{-1}([a,1]) = \bigcap_{\alpha < a} U_\alpha$.
- $f^{-1}([0,b))$ we know from the previous bullet point that this is equal to $X \setminus \bigcap_{\alpha < b} U_\alpha = \bigcup_{\alpha < b} U_\alpha^{\complement}$. Unfortunately, because $U_\alpha^{\complement}$ is closed (complement of an open set), it is not clear that $f^{-1}([0,b))$ is open. If only we could replace $U_\alpha^{\complement}$ by $\overline{U_\alpha}^{\complement}$! Well, it turns out we can, because obviously

$x \in \bigcap_{\alpha<b} U_\alpha \implies x \in \bigcap_{\alpha<b} \overline{U_\alpha}$; and $x \in \overline{U_\beta}$ for $\beta < b \implies x \in \bigcap_{\alpha<\beta} U_\alpha$, which implies that $x \in \bigcap_{\beta<b} \overline{U_\beta} \implies x \in \bigcap_{\beta<b} \bigcap_{\alpha<\beta} U_\alpha = \bigcap_{\alpha<b} U_\alpha$.

- Finally, $f^{-1}((a,b)) = f^{-1}([0,b)) \cap f^{-1}((a,1])$, and the intersection of two open sets is open.

Thoughts: technically we could have indexed by all real numbers — with Euclidean intuition it doesn't matter. But the countable case works, so that's something to keep in mind. Also, if it seems daunting to construct open sets indexed by rationals like this (I mean how do you even start? There's no smallest positive rational!), dyadic rationals work too. Just find $U_{1/2}$ first, then $U_{1/4}$ and $U_{3/4}$, and keep subdividing.

Also, this proof relied on this "chain" of open sets "between" A and B satisfying $p > q \implies \overline{U_p} \subseteq U_q$. Intuitively, it seems like an if and only if statement (between the existence of such a chain and the existence of Urysohn's function): if we have such a continuous function, 1 on A and 0 on B, then $U_a = f^{-1}([0,a))$ is open and $\overline{U_a} = f^{-1}([0,a])$ because $f^{-1}([0,a])$ is a closed set containing U_a, and is contained in all $f^{-1}([0,a+\epsilon))$ (recall that the closure of a set is the smallest closed set containing it).

So given closed A, B, we have that: there is a countable sequence of open sets around A, increasing s.t. closure contained in next open sets $\iff$ there is a continuous function defined on X s.t. it is 1 on A and 0 on B. We'll call the first condition the "open chain condition", and the second the "Urysohn function condition".

It does seem kind of hard to verify the countable sequence proposition, but there is a bit of recursion here! First, we have closed A and B, and we construct $U_{1/2}$ s.t. $A \subseteq U_{1/2} \subseteq X \setminus B$. Then, to do $U_{1/4}$, we do the same thing but the "A" this time is $\overline{U_{1/2}}$ and "B" is still B. For $U_{3/4}$, take "A" to be A and "B" to be $X \setminus U_{1/2}$. Of course, we have to be careful in allowing some "breathing room"; e.g. we can not have $U_{1/2} = X \setminus B$. That is to say, we would like that $\overline{U_{1/2}} \subseteq B$, and similarly for every "A" and "B", we want "A" $\subseteq U \subseteq \overline{U} \subseteq X \setminus$"$B$". In this manner, given consecutive dyadic rationals $\frac{a}{2^n}, \frac{a+1}{2^n}$, we can construct $U_{(2a+1)/2^{n+1}}$ s.t. $\overline{U_{(a+1)/2^n}} \subseteq U_{(2a+1)/2^{n+1}} \subseteq \overline{U_{(2a+1)/2^{n+1}}} \subseteq U_{a/2^n}$, where we use the notation $U_1 = A$ and $U_0 = X \setminus B$.

Thus, we see that if our topological space X has the property that for any closed sets $A, B \subseteq X$, there is some open set $U \subseteq X$ s.t. $A \subseteq U$ and $\overline{U} \subseteq X \setminus B \iff X \setminus \overline{U} \supseteq B$, then we have the "open chain condition", or equivalently the "Urysohn function condition". The existence of an open set with this condition is equivalent to the existence of disjoint open U_A and U_B s.t. $A \subseteq U_A$ and $B \subseteq U_B$ (the ($\implies$) direction is obvious, and for the ($\impliedby$) direction, we have $U_A \subseteq X \setminus U_B \implies \overline{U_A} \subseteq \overline{X \setminus U_B} = X \setminus U_B$). We call this condition *normality*, and a topological space satisfying this property *normal*.

And of course, if we have the "Urysohn function condition", then $f^{-1}([0, \frac{1}{2}))$ and $f^{-1}((\frac{1}{2}, 1])$ are open sets that contain A and B respectively. Thus, normal $\iff$ "open chain condition" $\iff$ "Urysohn function condition".

Definition 4.1: Normal topological space

A topological space X being normal means that for any closed sets $A, B \subseteq X$, there exist open sets $U_A, U_B \subseteq X$ s.t. $A \subseteq U_A$ and $B \subseteq U_B$.
We proved that normality is equivalent to the "Urysohn function condition", i.e. for any closed sets $A, B \subseteq X$, there is a continuous function $f : X \to [0, 1]$ s.t. $f|_A = 1$ and $f|_B = 0$.

4.2 Partitions and Paracompactness

This section on partitions of unity is based around the corresponding section "Partitions of Unity" in Lee [6], pg. 114. We now define what a (pointwise!) partition of unity subordinate to an open cover $\{U_i\}_{i\in I}$ is:

Definition 4.2: Partition of unity (pointwise)

Given an open cover $\{U_i\}_{i\in I}$ of a topological space X, a partition of unity subordinate to this cover is a collection of functions $\{\varphi_i\}_{i\in I}$, $\varphi_i : X \to [0, 1]$, s.t.

- $\text{supp}(\varphi_i) \subseteq U_i$ (like discussed in the paragraphs prior to Section 4.1);
- for all $x \in X$, only a finite number of $\varphi_i(x)$ are positive (we do NOT want to deal with infinite sums and convergence issues!);
- and for all $x \in X$, $\sum_{i\in I} \varphi_i(x) = 1$.

Let me remark that this definition is not standard; in fact, most of the time, we ask for a stronger condition 2, to guarantee that stitching local continuous functions $f_i : U_i \to \mathbb{R}$ into a global CONTINUOUS function $\sum_{i\in I} f_i\varphi_i$ that is close, if not equal to the local functions on their respective open sets of definition (cf. the first paragraph of this section, Section 4):

Definition 4.3: Partition of unity (local)

Given an open cover $\{U_i\}_{i\in I}$ of a topological space X, a partition of unity subordinate to this cover is a collection of functions $\{\varphi_i\}_{i\in I}$, $\varphi_i : X \to [0, 1]$, s.t.

- $\text{supp}(\varphi_i) \subseteq U_i$;
- for all $x \in X$, there is a neighborhood N_x containing x s.t. only a finite number of φ_i are positive on N_x;
- and for all $x \in X$, $\sum_{i\in I} \varphi_i(x) = 1$.

Indeed, it is not necessarily true that with a pointwise partition, stitching local continuous functions will yield a global continuous function; however with local partitions, we see that at every $x \in X$, the global function $\sum_{i\in I} f_i\varphi_i$ is locally continuous (on the neighborhood N_x the global function is just a finite sum of continuous functions), and hence it is indeed globally continuous. It is moreover true that

Theorem 4.4: Partition of unity pointwise yields local

Given a pointwise partition of unity $\{\varphi_i\}_{i\in I}$ subordinate to an open cover $\{U_i\}_{i\in I}$ of a topological space X, we can in fact find a local partition of unity $\{\psi_i\}_{i\in I}$ subordinate to $\{U_i\}_{i\in I}$.

See Theorem 2.3 in https://www.math.uni-bielefeld.de/~tcutler/pdf/Partitions%20of%20Unity.pdf. Spoiler: let $\tilde{\psi}_j(x) := \max\{0, \varphi_j(x) - [\sup_{i\in I}\varphi_i(x)]/2\}$ and $\psi_j := \tilde{\psi}_j / \sum_{i\in I}\tilde{\psi}_i$.

From now on, referring to "partitions of unity" will mean only LOCAL partitions of unity. In order to apply Urysohn's lemma to prove the existence of such a partition of unity, we would first need X to be normal, and then for any U_i in the open cover $\{U_i\}_{i\in I}$, we would need to get closed sets A_i, B_i like in the pictures above in Section 4.1. In order to guarantee condition 3, i.e. $\sum_{i\in I}\varphi_i = 1$, we would need $\{X \setminus B_i\}_{i\in I}$ to also be an open cover of X. These issues boil down to figuring out how to show that any open cover $\{U_i\}_{i\in I}$ can be *refined* to a subcover $\{V_i\}_{i\in I}$ s.t. $\overline{V_i} \subseteq U_i$ for all $i \in I$. Moreover, condition 2 makes it so that we need the refinement $\{V_i\}_{i\in I}$ to be s.t. for any $x \in X$, there is a neighborhood N_x containing x s.t. there are only a finite number of indices $i \in I$ s.t. $N_x \cap V_i \neq \varnothing$; we call such a cover *locally finite*. The property that any open cover can be refined to a locally finite subcover we call *paracompact*; in fact it is easy to see that paracompactness is a necessary for having a partition of unity subordinate to any open cover $\{U_i\}_{i\in I}$ because $\{\varphi_i^{-1}((0,\infty))\}_{i\in I}$ is a locally finite refinement by conditions 1 and 2.

Lemma 4.5: Refinement lemma

Let X be a topological space with properties we'll explicitly identity in the proof (starting out, it should at the very least be normal, since we'll be using Urysohn eventually anyways, and paracompact since it's necessary for partitions of unity). For any open cover $\{U_i\}_{i\in I}$ of X, there is a locally finite open refinement $\{V_i\}_{i\in I}$ s.t. $\overline{V_i} \subseteq U_i$ for all $i \in I$.

Proof: well if we were trying to find such a cover but didn't necessarily need the index set I to be the same, we could try to use normality to find a neighborhood N_x of each point $x \in X$ s.t. $\overline{N_x} \subseteq U_i$ for some $i \in I$, and then use paracompactness to get a locally finite refinement $\{N_x\}_{x\in\tilde{X}} \subseteq \{N_x\}_{x\in X}$ for some subset $\tilde{X} \subseteq X$. We could find such N_x if we know that every singleton point $x \in X$ is closed (see paragraph above Definition 4.1 for equivalent definition of normal).

To guarantee that singleton points x are closed sets, we need to guarantee that $X \setminus \{x\}$ is open (an iff statement). This in turn happens if and only if for every $y \in X \setminus \{x\}$ has some open $U_y \subseteq X \setminus \{x\}$ containing y (from "the most fundamental property of an open set" discussed in Section 2.0.1). Thus, we have that all singleton points are closed if and only if for all $x, y \in X$, there is some open set containing y but not x. This property is called "being a T_1 space".

Ok, so we ask now that X be T_1 in addition to normal and paracompact. Now, we try to amalgamate sets in the cover $\{N_x\}_{x\in\tilde{X}}$ so that we have exactly one set for every index $i \in I$. Well, by definition of N_x, we know that for each $x \in X$, there is at least one $i \in I$ s.t. $\overline{N_x} \subseteq U_i$. So let us define the indexer function $\iota : \tilde{X} \to I$ s.t. $\overline{N_x} \subseteq U_{\iota(x)}$ for all $x \in \tilde{X}$, and let us define $O_i := \bigcup_{x\in\iota^{-1}(i)} N_x$. The major problem now is determining if $\overline{O_i} \subseteq U_i$; it is true that the O_i are open and form an open cover of X (since $\{N_x\}_{x\in\tilde{X}}$ was an open cover of X), and $O_i \subseteq U_i$ — we just don't know about the closure/boundary.

Well, all we know is that $\overline{N_x} \subseteq U_i$ for all $x \in \iota^{-1}(i)$. It'd be awfully nice if $\overline{O_i} \subseteq \bigcup_{x\in\iota^{-1}(i)} \overline{N_x}$, wouldn't it? It is true that for any collection of sets $\{S_i\}_{i\in I}$, $\bigcup_{i\in I} \overline{S_i} \subseteq \overline{\bigcup_{i\in I} S_i}$ because $y \in \overline{S_i} \iff$ any open set containing y intersects non-trivially with $S_i \implies$ any open set containing y intersects non-trivially with $\bigcup_{i\in I} S_i \iff y \in \overline{\bigcup_{i\in I} S_i}$, and so our desired claim would only possibly be "barely" true. In any case, let us try to prove it.

Let us suppose for sake of contradiction that we have some $y \in \overline{O_i}$ not in $\overline{N_x}$ for any $x \in \iota^{-1}(i)$. Because $\{N_x\}_{x\in\tilde{X}}$ was a locally finite cover of X, there must be some open set U_y containing y s.t. only finitely many $x \in \tilde{X}$, say $x \in \{x_1, \ldots, x_n\} \subseteq \iota^{-1}(i)$ are s.t. N_x and U_y intersect non-trivially. Then, $V_y := U_y \setminus (\bigcup_{k=1}^{n} \overline{N_{x_i}}) \subseteq U_y$ is an open set containing y (a FINITE union of closed sets remains closed!), that does not intersect any N_x for $x \in \iota^{-1}(i)$! This is because $V_y \subseteq U_y$ and hence can not intersect N_x for $x \in \iota^{-1}(i) \setminus \{x_1, \ldots, x_n\}$, and by definition V_y does not intersect $\overline{N_x}$ for any $x \in \{x_1, \ldots, x_n\}$. Thus, we have found an open set containing y that does not intersect O_i, and hence y can not possibly be in the closure $\overline{O_i}$; contradiction. ■

4.3 Metrization

In this subsection, we continue to explore the incredible applications of the Urysohn functions above (i.e. continuous functions f that are 1 on some closed A and 0 on some closed B disjoint from A). Indeed, this entire section could be thought of as a love letter to Urysohn's lemma — it really is that powerful. For instance, given a (normal) topological space $(X, \mathscr{T})$, such Urysohn functions (and some other minor properties) can be used to construct a metric $\rho : X \times X \to [0, \infty)$ on X that generates the topology $\mathscr{T}$. This section is heavily inspired from pg. 242-243 of Royden's *Real Analysis* [9]. Recall from our section above on Metrizability (Section 2.3) that this is saying that we can construct the metric ρ s.t. (i) for $B_\rho(x_0, r)$ and any $x \in B_\rho(x_0, r)$, there is some $U_x \in \mathscr{T}$ containing x s.t. $U_x \subseteq B_\rho(x_0, r)$; and (ii) for every $U \in \mathscr{T}$ and $x \in U$, we can find some $r_x > 0$ s.t. $B_\rho(x, r_x) \subseteq U$.

4.3.1 (ii) $\mathscr{M} \leq \mathscr{T}$

Inklings of (ii) already begin to show up in the simplest metric-like function we can define using the Urysohn lemma functions: consider $\delta(x, y) = |f(x) - f(y)|$ where $f = 1$ on closed A and 0 on closed B; then (even without using any properties of f) δ is symmetric and satisfies the triangle inequality (and so is "almost a metric"), and moreover (using the properties of f) has that for all $x \in O := f^{-1}([1, \frac{1}{2}))$ (an open set containing A), $B_\delta(x, \frac{1}{2}) = \{x' \in X : \delta(x, x') < \frac{1}{2}\}$ is contained completely within $V := X \setminus B$ (an open set). Side note: mathematicians have actually given a name to functions like

δ (i.e. functions that would be metrics except that they do not satisfy $d(x,y)=0 \implies x=y$) — *pseudometrics*.

This is promising, but not every $x \in V$ is in O, and V is just one open set. Let's try to solve the first problem first. As a baby step, let's say that there are two closed sets $A_1, A_2 \subseteq V$ with f_1 and f_2 respectively being 1 on A_1, A_2 and 0 on V, each defining a pseudometric δ_1 and δ_2. Of course, $A_1 \cup A_2$ is a closed set in V and so technically this can be reduced to the previous case above, but let's try to think of a way to combine the two δ_i $(i=1,2)$ into some pseudometric δ with desirable properties. Well, if we have that for any $x_i \in O_i = f_i^{-1}([1,\frac{1}{2}))$, $B_{\delta_i}(x_i,\frac{1}{2}) \subseteq V$ for $i=1,2$, then for ANY $x \in O_1 \cup O_2$, $B_{\delta_1+\delta_2}(x,\frac{1}{2}) \subseteq V$, because $\delta_1+\delta_2$ is always $\geq$ each of δ_1, δ_2 individually (and so if $x \in O_i$, then $B_{\delta_1+\delta_2}(x,\frac{1}{2}) \subseteq B_{\delta_i}(x,\frac{1}{2}) \subseteq V$).

The two-closed-sets case is easily extended to any finite collection of closed sets $A_1,\ldots,A_n$, so let us move to the countable case. With infinitely many δ_i, we can no longer simply add them up, since we might get ∞. However, because all the δ_i are bounded by 1, if we scale each δ_i by $\frac{1}{2^i}$, the sum will converge. Thus, we define $\delta_V(x,y) = \sum_{i=1}^{\infty} \frac{1}{2^i}\delta_i(x,y)$. Then for any $x \in \bigcup_{i=1}^{\infty} O_i$ where $O_i := f_i^{-1}([1,\frac{1}{2}))$ (for $i \in \mathbb{N}$), we know that x is in say O_i, and so $B_{\delta_V}(x,\frac{1}{2}\cdot\frac{1}{2^i}) \subseteq B_{2^{-i}\delta_i}(x,\frac{1}{2}\cdot\frac{1}{2^i}) = \{x' \in X : \frac{1}{2^i}\delta_i(x,x') < \frac{1}{2}\cdot\frac{1}{2^i}\} = \{x' \in X : \delta_i(x,x') < \frac{1}{2}\} \subseteq V$.

If $\bigcup_{i=1}^{\infty} O_i = V$, we would in fact have that for every $x \in V$, there is some $r_x > 0$ s.t. $B_{\delta_V}(x,r_x) \subseteq V$. One way we could guarantee that $\bigcup_{i=1}^{\infty} O_i = V$ is if X has some countable dense subset of points $Q = \{x_1,\ldots\}$ (like the rationals $\mathbb{Q}^n$ in $\mathbb{R}^n$) — then, if we have that singleton points are closed sets, then the above analysis tells us that taking $Q' = \{x'_1,\ldots\}$ to be the points of Q in V (could be finite), we could set $A_i = x'_i$ and get that $\bigcup_{i=1}^{\infty} O_i \subseteq V$ is an open set containing Q', a dense subset of V, and hence must be equal to V itself. The property that a space X has a countable dense subset is called *separable*.

Recall from our proof of Lemma 4.5 above (the refinement lemma) that singleton points are closed $\iff$ the topological space is T_1.

Now, we can tackle the second problem; i.e. the problem that V is just one open set. From our experience in thinking about the first problem, we see that we can extend to countably many sets $V_1, V_2, \ldots$ as follows: we know that on each V_j, we have a pseudometric δ_{V_j} (defined as above) such that for every $x \in V_j$, there is some $r_x > 0$ s.t. $B_{\delta_{V_j}}(x,r_x) \subseteq V_j$. Defining $\delta(x,y) = \sum_{j=1}^{\infty} \frac{1}{2^j}\delta_{V_j}(x,y)$, we have that for any $x \in \bigcup_{j=1}^{\infty} V_j$, x is in say V_j, and so $B_\delta(x, r_x \cdot \frac{1}{2^j}) \subseteq B_{2^{-j}\delta_{V_j}}(x, r_x\cdot\frac{1}{2^j}) = \{x' \in X : \frac{1}{2^j}\delta_{V_j}(x,x') < r_x \cdot \frac{1}{2^j}\} = \{x' \in X : \delta_{V_j}(x,x') < r_x\} = B_{\delta_{V_j}}(x,r_x) \subseteq V$. Thus, if $\bigcup_{j=1}^{\infty} V_j = X$, (2) is almost complete (the only thing left to resolve is δ being only a pseudometric, not a metric).

4.3.2 (i) $\mathcal{T} \leq \mathcal{M}$

Let us now see if we can prove (i), i.e. that each $B_\delta(x_0,r)$ is open in the $\mathcal{T}$ topology. This turns out in fact to be very easy — the key insight is that $B_\delta(x_0,r) = \{x \in X : \delta(x_0,x) < r\} = g^{-1}([0,r))$, where $g(x) := \delta(x_0, x)$. Thus, if we prove that $g(x)$ is continuous, we have that $B_\delta(x_0, r)$ is a preimage of an open set (in the $[0, 1]$ subspace topology), and hence is open in the $\mathscr{T}$ topology.

To prove that $g(x)$ is continuous (where recall $g(x) := \delta(x_0, x)$ and $\delta(x, y) := \sum_{j=1}^{\infty} \frac{1}{2^j} \delta_{V_j}(x, y)$ and $\delta_V(x, y) = \sum_{i=1}^{\infty} \frac{1}{2^i} \delta_i(x, y)$, where the δ_i are of the form $|f(x) - f(y)|$ for Urysohn functions f), note that because each f given by Urysohn's Lemma is continuous, $f(x) - f(x_0)$ is continuous; and because $x \mapsto |x|$ is continuous and composition of continuous functions is continuous, $|f(x) - f(x_0)|$ is continuous; and finally recall that a uniform limit of continuous functions is continuous, where the series here is a uniform limit of continuous functions (i.e. it is uniformly convergent) by the Weierstrass M-test. This uniform convergence result is important enough that I'll give a nice box:

Lemma 4.6: Limits of (sums of) continuous functions

Uniform limit lemma: for any topological space X and metric space Y, if $f_n : X \to Y$ is a sequence of continuous functions converging uniformly to some $f : X \to Y$, then f is continuous as well.

Weierstrass M-test: if Y is furthermore (real/complex) Euclidean space, and $|f_n(x)| \leq M_n$ for all $x \in X$, where $\sum_{n=1}^{\infty} M_n < \infty$, then $S(x) := \sum_{n=1}^{\infty} f_n(x)$ is absolutely and uniformly convergent on X (i.e. $S_n(x) := \sum_{k=1}^{n} f_k(x)$ converges uniformly to $S(x)$, so by the uniform limit lemma above if the f_n are continuous, S is too).

4.3.3 Metric, not Pseudometric

The last thing for us to resolve is the conundrum of how to ensure that δ is a metric. The only condition we have not shown is true is: $\delta(x, y) = 0 \implies x = y$. The definition of δ is $\delta(x, y) = \sum_{j=1}^{\infty} \frac{1}{2^j} \delta_{V_j}(x, y)$, and so if $\delta(x, y) = 0$, it must be that for every open set V_j ($j \in \mathbb{N}$), $\delta_{V_j}(x, y) = 0$ as well. Now we know that if $x \in V_j$ and $y \notin V_j$, $\delta_{V_j}(x, y) > 0$, because all the f_i used in the definition of $\delta_{V_j}(x, y) = \sum_{i=1}^{\infty} \frac{1}{2^i} \delta_i(x, y) = \sum_{i=1}^{\infty} \frac{1}{2^i} |f_i(x) - f_i(y)|$ are 0 outside of V_j, and we know that at least one f_i is positive on $x \in V_j$ (by construction/choice of the f_i from above). Thus, if we choose our countable collection $\{V_j\}_{j=1}^{\infty}$ s.t. $\bigcup_{j=1}^{\infty} V_j = X$ (as above), and for any $x, y \in X$, there is some V_j s.t. $x \in V_j$ and $y \in X \setminus V_j$, then δ is a bona-fide metric satisfying (1) and (2), as desired.

One way to guarantee this is if X has a countable basis, and we choose $\{V_j\}_{j=1}^{\infty}$ to be that countable basis. This is because T_1-ness gives that the point y is closed, so $X \setminus \{y\}$ is an open set in $\mathscr{T}$, and because $\{V_j\}_{j=1}^{\infty}$ is a basis, we have that $X \setminus \{y\} = \bigcup_{j \in J} V_j$ for some index set $J \subseteq \mathbb{N}$, and so $x \in X \setminus \{y\}$ must be in one of the V_j for $j \in J$.

4.3.4 Summary

The property that $(X, \mathscr{T})$ has a countable dense subset is called *separable*. The property that $(X, \mathscr{T})$ has a countable basis is called *second countable*. Thus, we've proved that

Theorem 4.7: Urysohn metrization theorem

If $(X, \mathscr{T})$ is separable (with countable dense subset $Q = \{x_i\}_{i=1}^\infty$), second countable (with countable basis $\{V_j\}_{j=1}^\infty$), T_1, and normal (so that we can use Urysohn's lemma), it is metrizable with metric

$$\rho(x,y) = \sum_{j=1}^\infty \frac{1}{2^j}\left(\sum_{i=1, x_i \in V_j}^\infty \frac{1}{2^i} \cdot |f_{j,i}(x) - f_{j,i}(y)|\right)$$

where $f_{j,i}$ is a continuous function that is 1 at the point $x_i \in Q$ and 0 on $X \setminus V_j$, guaranteed to exist by Urysohn's lemma. *Sidenote:* sometimes, being T_1 is included in the definition of *normal*, but I will not do that. Instead, the property of being both T_1 and normal I will call T_4.

Last remarks: assuming the axiom of choice, separability is easily seen to be implied by second countability, because given a countable basis $\mathscr{B} := \{B_n\}_{n=1}^\infty$, we can use the axiom of choice to pick $x_n \in B_n$, which is then a countable dense subset because any nonempty open set U contains some $x \in X$ and hence by openness of U and the fact that $\mathscr{B}$ is a basis, there must be some $B_n \in \mathscr{B}$ s.t. $x \in B_n \subseteq U$, implying of course that $x_n \in U$.

Note also that we have a partial converse to the :

Lemma 4.8: Metric spaces are T_4

Metric spaces are T_4

Proof: given any closed $A, B \subseteq X$ in a metric space, we can define $U_A = \bigcup_{x\in A} B(x, \frac{1}{3}d(x,B))$ and similarly $U_B = \bigcup_{y \in B} B(y, \frac{1}{3}d(y,A))$, where $d(A,B) := \inf\{d(x,y) : x \in A, y \in B\}$. These are obviously open sets that contain A, B respectively, and they are disjoint, because if not, then there is some z in the intersection hence satisfying $d(x,z) < \frac{1}{2}d(x,B)$ and $d(y,z) < \frac{1}{2}d(y,A)$ for some $x \in A, y \in B$. The triangle inequality then gives $d(x,y) < \frac{1}{2}(d(x,B) + d(y,A))$, but $d(x,y)$ is $\geq$ both $d(x,B)$ and $d(A,y)$, i.e. $d(x,B), d(A,y) < \frac{1}{2}(d(x,B) + d(y,A))$. Adding the two inequalities yields $d(x,B) + d(y,A) < d(x,B) + d(y,A)$, a contradiction. ■

4.4 Tietze Extension Theorem

Urysohn's lemma above allows us to extend a very specific type of $[0,1]$-valued continuous function on a closed subset $F = A \cup B$ of a normal topological space X to a continuous function on all of X. A very clever argument using Urysohn's lemma and some approximation techniques can be used to extend this extension result to ANY continuous real-valued function defined on a closed subset F. We will actually prove this for $[-1,1]$-valued continuous functions first; as for why not $[0,1]$-valued, the symmetry around 0 will be important when we make statements like $|f| \leq 1$, and moreover our approximations may sometimes overshoot or undershoot, so it's no use to start off ≥ 0 anyways. This section is heavily inspired by §35 of Munkres [8].

Theorem 4.9: Tietze extension theorem

Let X be a normal topological space and let $F \subseteq X$ be closed. Suppose we have a continuous function $f : F \to [-1, 1]$ (continuous w.r.t. the subspace topology). Then, we can extend to a continuous function $g : X \to [-1, 1]$, i.e. g is continuous and $g|_F = f$.

As I hinted at in the introduction to this section, the main idea of the proof is to prove that we can approximate f with a continuous function $g_1 : X \to [-?, ?] \subsetneq [-1, 1]$ so that for all $x \in F$, $|f(x) - g_1(x)|$ is $\leq$ some approximation factor $\alpha < 1$, and then approximating $f - g_1$ (as a function on F) by a continuous function $g_2 : X \to [-??, ??] \subsetneq [-?, ?]$ so that for all $x \in F$, $|f(x) - g_1(x) - g_2(x)| \leq \alpha^2$, and so on until we can define $f := \sum_{i=1}^{\infty} g_i$ which hopefully is continuous and satisfies for all $x \in F$, $|f(x) - f(x)| = 0$.

Lemma 4.10: Approximation lemma

The setting (normal X, closed $F \subseteq X$) is the same. Suppose $f : F \to [-r, r]$ is continuous w.r.t. the subspace topology of F; then there is a continuous $\tilde{g} : X \to [-\beta r, \beta r]$ for some $\beta \in (0, 1)$ s.t. for all $x \in F$, $|f(x) - \tilde{g}(x)| < \alpha r$ for some $\alpha \in (0, 1)$. To be more clear, the reason why we have these β, α is so that after successive approximations, we get convergence to desirable quantities by geometric series.

Proof of approximation lemma: note that we can replace r with 1 in the above statement without weakening it, because we can just scale $f : F \to [-1, 1]$ by r to get $rf : F \to [-r, r]$ and so on. Well if we require the output of $\tilde{g}$ to lie in $[-\beta, \beta]$ and we know that the output of f lies in $[-1, 1]$, the best and easiest approximation we can do for the maximal and minimal values of f is to try to let $\tilde{g}$ be $-\beta$ on $A := f^{-1}([-1, -\beta])$ and β on $B := f^{-1}([\beta, 1])$. Because $[-1, -\beta]$ and $[\beta, 1]$ are closed in $[-1, 1]$ and f is continuous, A, B are both closed sets w.r.t. the subspace topology of F. But because F is closed, A, B are closed in X, and so indeed Urysohn's lemma gives the existence of $\tilde{g} : X \to [-\beta, \beta]$ s.t. $\tilde{g}|_A = -\beta$ and $\tilde{g}|_B = \beta$.

Now given this (albeit sparse) information about $\tilde{g}$, what can we conclude about the approximation factor α regarding "closeness" with f? Well, we know $\alpha \geq 1 - \beta$, since we could have some $x \in A \cup B$ s.t. $f(x) = \pm 1$ and $\tilde{g}(x) = \pm\beta$. In the other case $x \in F \setminus (A \cup B)$, both $f(x), \tilde{g}(x) \in [-\beta, \beta]$, and so can take $\alpha = \max\{2\beta, 1 - \beta\}$. Note that the value on the RHS is minimized when $2\beta = 1 - \beta \iff \beta = \frac{1}{3} \iff \alpha = \frac{2}{3}$; keep these concrete values in mind, but know that I will continue to use β, α in the next parts. ■

Proof of Theorem 4.9: so we have $f : F \to [-1, 1]$, i.e. $|f(x)| \leq 1$ on F. The approximation lemma yields $g_1 : X \to [-\beta \cdot 1, \beta \cdot 1]$ (i.e. $|g_2(x)| \leq \beta \cdot 1$ on X) s.t. $|f(x) - g_1(x)| \leq \alpha \cdot 1$ on F. Applying the approximation lemma again for $[f - g_1]$, we get that there is $g_2 : X \to [-\beta \cdot \alpha, \beta \cdot \alpha]$ (i.e. $|g_2(x)| \leq \beta \cdot \alpha$ on X) s.t. $|f(x) - g_1(x) - g_2(x)| \leq \alpha \cdot \alpha = \alpha^2$ on F. Continuing on like this, we can define for every $n \in \mathbb{N}$ a function $g_n : X \to [-\beta \cdot \alpha^{n-1}, \beta \cdot \alpha^{n-1}]$. (i.e. $|g_n(x)| \leq \beta \cdot \alpha^{n-1}$ on X) s.t. $|f(x) - \sum_{k=1}^{n} g_k(x)| \leq \alpha^n$ on F (let us denote $S_n(x) := \sum_{k=1}^{n} g_k(x)$).

Then, defining $g(x) := \sum_{n=1}^{\infty} g_n(x)$, Lemma 4.6 above gives that g is continuous on X (more explicitly, we have $|g_n(x)| \le \beta\alpha^{n-1}$ on X, and $\sum_{n=1}^{\infty} \beta\alpha^{n-1} = \frac{\beta}{1-\alpha} < \infty$ so the Weierstrass M-test applies). And, we have that f and g agree on F because

$$|f(x) - g(x)| \le |f(x) - S_n(x)| + |S_n(x) - g(x)| \le \alpha^n + |S_n(x) - g(x)|$$

for any $x \in F$ and $n \in \mathbb{N}$, and because $S_n \to g$ uniformly on X and $\alpha^n \to 0$ as $n \to \infty$, for any $\epsilon > 0$, we can find n large enough so that the RHS becomes $< \epsilon$, yielding $f(x) = g(x)$ for all $x \in F$, as desired.

Finally, we if use the values $\alpha = \frac{2}{3}$ and $\beta = \frac{1}{3}$, we get that $|g(x)| \le \frac{\beta}{1-\alpha} = 1$, so indeed $g : X \to [-1, 1]$, but even without those values of α, β, we can just rescale g to be $g(x)$ when $g(x) \in [-1, 1]$ and $\frac{g(x)}{|g(x)|}$ when $g(x) \notin [-1, 1]$. ■

By rescaling and translating, the above theorem holds for any functions $f : F \to [a, b]$. We can in fact prove the theorem for $f : F \to \mathbb{R}$:

Theorem 4.11: Tietze extension theorem (unbounded)

Let X be a normal topological space and let $F \subseteq X$ be closed. Suppose we have a continuous function $f : F \to \mathbb{R}$ (continuous w.r.t. the subspace topology). Then, we can extend to a continuous function $h : X \to \mathbb{R}$, i.e. h is continuous and $h|_F = f$.

Proof: because $(-1, 1)$ and $\mathbb{R}$ are homeomorphic, we can replace "$\mathbb{R}$" in the statement above by "$(-1, 1)$. In this manner we see a more obvious duality between Theorem 4.9 (about closed intervals) and Theorem 4.11 (about open intervals). To start, note that Theorem 4.9 already gives the existence of $g : X \to [-1, 1]$ extending $f : A \to (-1, 1)$. Defining $D = g^{-1}(\{-1, 1\})$ (which is closed because g is continuous and $\{-1, 1\}$ is a closed subset of $\mathbb{R}$), observe that $D \cap A = \varnothing$ (because $g(x) = f(x) \in (-1, 1)$ for all $x \in A$). We now have two disjoint closed sets in a normal space X, so by Urysohn's lemma (Section 4.1), there is a continuous function $\phi : X \to [0, 1]$ that is 0 on D and 1 on A. Then, $h := g \cdot \phi$ is a continuous function (product of two continuous functions is continuous), equals f on A (because for $x \in A$, $h(x) = g(x)\phi(x) = g(x) \cdot 1 = f(x)$), 0 on D, and $|h(x)| \le |g(x)| \cdot |\phi(x)| \le |g(x)| < 1$ for $x \in X \setminus D$, meaning that $h : X \to (-1, 1)$ is a continuous extension of $f : A \to (-1, 1)$, as desired. ■

5 Compactness

Before we continue further on our quest to embed spaces into $\mathbb{R}^n$, let us take a break, and introduce a very key concept: compactness. Now, I would have preferred to introduce compactness in a more natural way (as in something we encounter naturally on the trail, instead of a something akin to taking a detour up to a mountain village to buy supplies), but it just didn't fit in particularly well anywhere, and it's important enough that it's better to nail down the basics before getting into the more complicated weeds.

Personally, I think the best motivation for/most natural way of encountering compactness is from developing measure theory on $\mathbb{R}$ — driven by the question of defining some sort of "measure" for subsets of $\mathbb{R}$, one may come up with the rather intuitive definition of $\mu^*(A) = \inf\{\sum_{k=1}^{\infty}(b_k - a_k) : A \subseteq \bigcup_{k=1}^{\infty}(a_k, b_k)\}$ (i.e. saying that the "length" of A is the sum of the lengths of the intervals that make up the "tightest" cover of A), which then one may try to use to prove that the measure of $[0, 1]$ is 1, by first noticing that the claim is easy to prove via induction if one restricts to only considering finite (open) covers of $[0, 1]$; see my introductory measure theory paper, *Measure and Measurability*, for more details.

However, since this is neither the time nor place to get into measure theory, I will propose a "philosophical definition" of compactness from which we will build. You, the reader, have already been subject to several of my "philosophical definitions" in this paper; e.g. that topological spaces are "general settings on which we can talk reasonably about continuity", and that partitions of unity are "tools used to go from local to global/to break global down into local". In a similar vibe, let us think of **compact spaces** as "*those on which we can not run away from everyone (i.e. every point)*". This section is heavily inspired by Maxime Ramzi's answer to my MSE question "Could *I* have come up with the definition of compactness?" [2].

Now what do I mean by "running away" from every point? I can "run away" from every point if I can "move" in such a way that for any point $p \in X$, past some point "in time", I will never be close to p, i.e. there is some neighborhood of p (because in a topological space without a metric, "closeness" is best thought of in terms of "neighborhoods") in which I will never "step foot". This "moving" may sound like we are taking a sequence of moves $\{x_n\}_{n=1}^{\infty}$, but I was careful not to use that word because that carries with it the "baggage" of countability; it may be the case that we want to consider "moves" $\{x_i\}_{i \in I}$ where the index set I is uncountable. I'll start off with some easy examples ("snacks for thought") and then say a little bit more about "uncountable moves".

Consider $\mathbb{R}^2 \setminus \{0\}$. An obvious way to "run away" is to just to take a leisurely but steady walk off to the right: $(1, 0), (2, 0), (3, 0), \ldots$ and so on. We can think of this walk as walking away from the neighborhoods $B(0, 1), B(0, 2), B(0, 3), \ldots$ and not looking back (i.e. never returning to those neighborhoods again). In contrast, if we take a "sequence of hops" $(1, 0), (2, 0), (1, 0), (3, 0), (1, 0), (4, 0), \ldots$ and so on, where we return to $(1, 0)$ every other hop, we can not consider this "running away", since we will always return to $(1, 0)$. In a reciprocal manner, consider $(0, -1), (0, -\frac{1}{2}), (0, -\frac{1}{3}), \ldots$ and so on. If 0 was in our set, then this sequence would be "running towards" it, in the sense that for every neighborhood of 0, the sequence eventually enters the neighborhood, and in fact KEEPS REVISITING it. However, 0 is not in $\mathbb{R}^2 \setminus \{0\}$, so this sequence can also be thought of as "running away" (from the neighborhoods $[|x| > 1], [|x| > \frac{1}{2}], [|x| > \frac{1}{3}], \ldots$), like looking at someone walking down a long hallway, watching them shrink against the vanishing point. The point of that last example is to illustrate how "running away" does not necessarily mean anything related to the metric, like "the distance increases" or anything.

In those two examples above, I exhibited a sequence of moves "running away", and found corresponding neighborhoods that I was "running away" from. Note that the neighborhoods covered the entire space, so indeed I was "running away" from EVERY point in a sense. Now let's go the other way, i.e. supposing we are given a collection of neighborhoods, we ask ourselves if it is possible to "run away". Consider again $\mathbb{R}^2 \setminus \{0\}$, and consider the neighborhoods U_L, U_M, U_R defined to be the left half plane, a middle strip containing the y-axis, and the right half plane (huge neighborhoods, but just run with it). Well, any "moves" I take, since there are only three neighborhoods, I will have to keep revisiting at least one of them. That is to say, no matter how I "run", I will always eventually again be "close" (in the sense of being in a neighborhood with) to some point.

Well, obviously the same outcome will occur given any finite covering of neighborhoods. That means, to have a chance at "running away", there needs to be some infinite covering of neighborhoods. Of course, that's not a sufficient condition, since given our U_L, U_M, U_R example above, I can just add random (open) blobs in the plane, but of course those are superfluous/extraneous; they don't change that U_L, U_M, U_R is still a finite covering of neighborhoods. That is to say, if I have a space s.t. given any collection of neighborhoods covering it, I can always clear away superfluous/extraneous ones leaving behind a finite subcollection, I will never be able to "run away".

Of course, this yields a sufficient condition, but it does seem rather strong. Is it necessary? That is to say, given a space and a collection of neighborhoods covering it s.t. NO finite subcollection covers it, MUST it always be possible to "run away"? Indeed, yes! Suppose the collection $\{U_i\}_{i\in I}$ is indexed by I (infinite, countable or uncountable). Now considering a sequence of finite index sets $I_1 \subseteq I_2 \subseteq I_3 \subseteq \dots$ where I_1 has one index of I, and I_2 is I_2 with one more index of I, and so on, we can define x_n to be some element not in I_n (where such an element exists because we assumed that NO finite subcollection covers the space). Then, as long as our subcollections do not cover the space, we can continue this and produce a (countable or uncountable) "list" which we can follow to "run away" successfully. Basically, the "NO finite subcollection covering" property is just to guarantee the existence of an infinite "list" of moves; we can go on to uncountable "lists" as long as there are more points not covered by a subcollection we pick. We therefore arrive at the following definition:

Definition 5.1: Compactness

A topological space X is called *compact* exactly when every open cover $\{U_i\}_{i\in I}$ (i.e. each U_i is an open set in the topology, and $\bigcup_{i\in I} U_i = X$) has a finite subcover (i.e. there is a finite subset $\tilde{I}$ of the index set I s.t. $X = \bigcup_{i\in\tilde{I}} U_i$). The "philosophical definition" is that a compact space is one in which I will never able to "run away", where being able to "run away" means that for every point $p \in X$, my "running" will be s.t. eventually, I will NEVER be near p (there will be a neighborhood of p that I will eventually NEVER be inside).

Just to give you a feeling as to why we were so picky about countable/uncountable moves above, consider a space consisting of "points" S where S is some countable set of real numbers. In such a space, we can talk about increasing and decreasing sequences by saying that $S_1, S_2, S_3, \dots$ is an increasing sequence $\iff S_1 \subseteq S_2 \subseteq \dots$. Now in this space, every increasing sequence $\{S_n\}_{n=1}^{\infty}$ does converge, to $\bigcup_{n=1}^{\infty} S_n$, because a countable union of countable sets is still countable. You see, we engineered this scenario to ensure that countable increasing sequences would always converge, unlike in $\mathbb{R}$. But perhaps right about now you should be feeling a wiggle in the back of your mind, about considering "uncountable" sequences. Unfortunately, the way we've defined the ordering, there are no "uncountable increasing sequences". In a sense, our ordering of using $\subseteq$ "throws away" a lot of sets when we think about increasing. The following paragraph is inspired by Brian M. Scott's answer to my MSE question.

We will now define something similar, but with a different ordering to fully utilize the "power" of uncountability. Suppose now we could arrange the elements of $\mathbb{R}$ in some sort of "uncountable list", i.e. there would be a "first real number", "second real number" and so on, for an uncountable amount of indices (the reason why I phrase it as "uncountable list" is so the reader can have intuitive ideas about this ordering, like transitivity, anticommutativity, and comparability for every pair of elements). This may be hard to wrap your head around, so perhaps an easier concept to understand that ultimately is equivalent, is that of a function $\ell : 2^{\mathbb{R}} \to \mathbb{R}$ that tells you the "least element" of any subset $S \subseteq \mathbb{R} \iff S \in 2^{\mathbb{R}}$, i.e. ℓ takes S to some element of S (ℓ for "least"). In a sense, ℓ encodes the "uncountable list" we chose, since knowing the ordering we can construct ℓ, and knowing ℓ, we can construct the ordering (let the "first real number" r_1 be $\ell(\mathbb{R})$, and then $r_2 = \ell(\mathbb{R} \setminus \{r_1\}), \dots$. The existence of such a function ℓ is guaranteed by the axiom of choice. Let us denote this ordering as $\prec$. Now for any $x \in \mathbb{R}$, we can define $P_x := \{y \in \mathbb{R} : y \prec x\}$, and X to be the space of all $x \in \mathbb{R}$ s.t. P_x is countable. We can "push" the ordering we had for $\mathbb{R}$ to X, by saying $P_x \prec P_y \iff x \prec y$. Finally, exactly like we figured in the above paragraph, any increasing countable sequence will converge (to $\bigcup_{n=1}^{\infty} P_{x_n}$), but since we can have "uncountable increasing sequences" now, you may again feel a wiggle in the back of your mind saying that "obviously there's room to continue running away!". This has hopefully given some taste as to why I use the words "move" or "running away" instead of something more precise like "sequence".

5.0.1 Ordinal Numbers

Above, we saw a crude way of "counting past infinity", by using the axiom of choice to define a "least element" and constructing a "list" that "transcends" the natural numbers/countable ordering. However, the problem is that we have no easy way of telling which is the "lesser" element (in terms of the "uncountable list") given two arbitrary real numbers. For example, we can not just apply ℓ to the two element set, because that may not agree with the "uncountable list". Thus, we would have to "write out" the "uncountable list" until we got to one of the two real numbers and then that would officially be the lesser element. So, if we could create some "canonical uncountable list" (perhaps even more infinite than uncountable) in which it was easier to see which elements are less than some other elements, we could put that in bijection with our "uncountable list"/total-ordering of $\mathbb{R}$, and maybe feel that the construction is on a bit more concrete ground.

Well, if we want to make it easy to tell which elements are less than some element, the most obvious thing to do is to try to "tag" the element with a list of the elements smaller than it, so to check if $a < b$, we just check if a is in b's "tag". For example, if we let 0 just be 0, we could think of 1 as the set $\{0\}$, and then 2 as the set $\{0, 1\}$, and 3 as the set $\{0, 1, 2\}$, and so on. Of course, when I write the symbols "1" or "2" in the sets, we can again identify them with their own "set definition", i.e. we would have $1 = \{0\}$, $2 = \{0, \{0\}\}$, $3 = \{0, \{0\}, \{0, \{0\}\}\}$, and so on. In other words, in this framework of identifying numbers with their "tag"/list of lesser numbers, we have a very orderly pattern happening, namely that the successor function s (from 0 to 1, 1 to 2, 2 to 3, and so on) is simply $s(x) = x \cup \{x\}$.

This successor function provides us a "canonical" way to count past infinity: we can define $\mathbb{N}$ exactly as I explicitly started out doing in the previous paragraph, and then define ω (the classical infinity at the "end" of the real line) to be $0 \cup 1 \cup 2 \cup 3 \cup \ldots = \bigcup_{n \in \mathbb{N}} n = \mathbb{N}$ (recall that we identify a "number" with the set of all the "numbers" less than it), and then "$\omega + 1$" as $s(\omega) = \omega \cup \{\omega\} = \mathbb{N} \cup \{\mathbb{N}\}$, and so on until $2\infty := 0 \cup 1 \cup \ldots \cup \omega \cup \omega + 1 \cup \ldots$, and still onward. These are called the "Von-Neumann ordinals". To create an uncountable ordinal (in fact the "first uncountable ordinal" ω_1), define ω_1 to be the set of all countable ordinals; see https://math.stackexchange.com/questions/15638/simple-example-of-uncountable-ordinal.

I hope this provides a clearer picture of what "uncountable lists" look like. Another more commonly used terminology for such "lists" are *chains*. For more, take a look at https://en.wikipedia.org/wiki/Ordinal_number#Von_Neumann_definition_of_ordinals and https://en.wikipedia.org/wiki/Total_order.

5.1 Applications

Let us first establish rigorously that such sets exist and in fact are quite common "everyday" objects.

Theorem 5.2: Heine-Borel theorem

The sets in $\mathbb{R}$ for which every open cover can be reduced to a finite open cover (i.e. the compact sets of $\mathbb{R}$) are exactly the sets which are closed and bounded

Proof: ($\Longrightarrow$): let S be compact. First, S must be bounded: consider the open cover $\{(-k, k)\}_{k=1}^{\infty}$, which necessarily covers S because $\bigcup_{k=1}^{\infty} (-k, k) = \mathbb{R}$. By compactness, S is covered by $(-1, 1) \cup \ldots \cup (-N, N)$, and so $S \subseteq (-N, N)$ and so it's bounded.

S must also be closed: suppose not. Then there is $p \notin S$ that is on the boundary. Consider the open cover $\{\mathbb{R} \setminus \overline{B(p, 1/k)}\}_{k=1}^{\infty}$ where $B(p, 1/k)$ is the open ball centered at p with radius $1/k$ (and the overline is "closure of"), which must necessarily cover S because

$$\bigcup_{k=1}^{\infty} I_k = \bigcup_{k=1}^{\infty} \mathbb{R} \setminus \overline{B(p, 1/k)} = \mathbb{R} \setminus \{p\} \supseteq S$$

By compactness, S can be covered by $I_1 \cup \ldots \cup I_M$, and we know $(\mathbb{R} \setminus \overline{B(p, 1/M)})$ and $B(p, 1/M)$ are disjoint, so we know S and $B(p, 1/M)$ must be disjoint. This means that p is actually not on the boundary of S (it is in the exterior); contradiction.

($\Longleftarrow$): for this direction, we prove the claim first for closed intervals by doing an "inchworm" argument, and then extend to the general case. Suppose F is a closed bounded set in $\mathbb{R}$ and $\mathscr{C}$ an open cover of F. First consider the case where $F = [a,b]$. Define D to be $D = \{d \in [a,b] : [a,d] \text{ has a finite subcover from } \mathscr{C}\}$. D is not empty because D must contain at least a, because a must be contained in at least one open set in $\mathscr{C}$.

We want to prove that $s = \sup D$ is b. Well, $[a,b]$ is closed so the supremum s is also in $[a,b]$. Suppose that $s < b$. Then, there must be some open set G from $\mathscr{C}$ that covers $(s-\delta, s+\delta)$ for some $\delta > 0$. But we know by definition of s that $[a, s-\frac{\delta}{2}]$ is covered by finitely many open sets from $\mathscr{C}$, say $G_1 \cup \ldots \cup G_N$. But then $[a, s+\frac{\delta}{2}]$ is also covered by finitely many open sets from $\mathscr{C}$, i.e. $G_1 \cup \ldots G_N \cup G$ and so s is not the supremum; contradiction; thus $s \geq b$, so $[a,b]$ can be covered by finitely many sets from $\mathscr{C}$.

Now let F be any arbitrary closed and bounded set in $\mathbb{R}$ with open cover $\mathscr{C}$. By boundedness, $F \subset [a,b]$ for some $a, b \in \mathbb{R}$. Then $\mathscr{C} \cup \{\mathbb{R} \setminus F\}$ is an open cover covering $\mathbb{R}$ and hence $[a,b]$. From the first part, we know that for some M, $F \subset [a,b] \subset G_1 \cup \ldots \cup G_M \cup (\mathbb{R} \setminus F)$. F obviously is not in $\mathbb{R} \setminus F$, so $F \subset G_1 \cup \ldots \cup G_M$. ■

5.1.1 Closedness and Boundedness in General

One may wonder given the Heine-Borel theorem for $\mathbb{R}$ whether or not boundedness and closedness characterizes compactness in general topological spaces. Although boundedness is not topological property ($\rho(x,y) = \min\{1, |x-y|\}$ is equivalent metric in $\mathbb{R}^n$) and hence definitely not a property of general topological spaces, closedness is a general topological property, and in fact it is very easy to see that

Lemma 5.3: Closed subset of compact space also compact

A closed subset F of a compact space X is also compact

Proof: given any open cover $F \subseteq \bigcup_{i \in I} U_i$, $X \setminus F \cup \{U_i\}_{i \in I}$ is an open cover of X, and so by compactness of X we have a finite subcover, and after removing $X \setminus F$ if it's still around, yields a finite subcover of F. ■

Moreover, our proof that compact sets must be closed can be extended to a quite general class of topological spaces. Let's repeat the proof above with a tiny bit of terminology changed:

"Suppose we have S a compact set in a topological space X, and suppose for sake of contradiction that it is not closed, i.e. there is $p \notin S$ that is on the boundary. Consider the open cover $\{X \setminus \overline{U_i}\}_{i \in I}$ where $\{U_i\}_{i \in I}$ is the collection/cover of all the open sets containing p. This must necessarily cover S because

$$\bigcup_{i \in I} X \setminus \overline{U_i} = X \setminus \{p\} \supseteq S.$$

By compactness, S can be covered by $\{X\setminus\overline{U_i}\}_{i\in\tilde{I}}$ for a finite set $\tilde{I}\subseteq I$, and we know the sets $\{X\setminus\overline{U_i}\}_{i\in\tilde{I}}$ are disjoint from $\bigcap_{i\in\tilde{I}} U_i$, an open set containing p (a finite intersection of open sets, hence open). This means that p is actually not on the boundary of S (it is in the exterior); contradiction."

The only thing that needs further justification is the step that $\bigcup_{i\in I} X\setminus\overline{U_i} = X\setminus\{p\}$. In order for this step to hold, we need that for any point $x\in X$ not equal to p, there is some open set around p s.t. $x\notin\overline{U_i}$. Note that this is a (small, but actually major) strengthening of the T_1 condition, which just asks that there exist an open set around p s.t. $x\notin U_i$. This defines what's known as the T_2 condition:

Definition 5.4: T_2 topological space, otherwise known as Hausdorff

A topological space X is called T_2/Hausdorff when for any two points $x,y\in X$ there exists an open set U containing x s.t. $y\notin\overline{U_i}$. Equivalently, for any two points $x,y\in X$ there exists disjoint open U_x, U_y containing x,y respectively.

In other words, we have just proven that

Lemma 5.5: Compact implies closed in Hausdorff

For T_2/Hausdorff spaces X, compact implies closed

It is furthermore no surprise that we proved this theorem for metric spaces first (the above proof holds essentially verbatim in any arbitrary metric space), since

Lemma 5.6: Metric spaces are Hausdorff

Metric spaces are T_2/Hausdorff

Proof: given any $x,y\in X$ a metric space, let $\delta=d(x,y)$ and then $B(x,\frac{\delta}{2})$ and $B(y,\frac{\delta}{2})$ are disjoint open balls containing x,y respectively. Or note that $T_4 \implies T_2$ and use Lemma 4.8. ■

5.1.2 Functions on Compact Space

You may recall some theorems such as "continuous functions on a closed interval $[a,b]$ attain their extremums" or "continuous functions on a closed interval $[a,b]$ are uniformly continuous on $[a,b]$" from an introductory real analysis class. The following theorems are simply generalizations of those.

Theorem 5.7: Continuous functions map compact sets to compact sets

The image of a continuous function on compact set/space is compact. I.e. continuous function on space that forbids running away yields a space that forbids running away (which is in particular bounded)

Proof: suppose we have a continuous function $f: X\to Y$ and compact $K\subseteq X$ and suppose we have an arbitrary open cover $\{V_i\}_{i\in I}$ of $f(K)$. Then, we can pullback $U_i := f^{-1}(V_i)$ to get an open cover $\{U_i\}_{i\in I}$ of K. Because K is compact, there is a finite set $\tilde{I}\subseteq I$ s.t. $\{U_i\}_{i\in\tilde{I}}$ is an open cover of K. Then, $\{V_i\}_{i\in\tilde{I}}$ remains an open cover of $f(K)$ because $y \in f(K) \implies$ there is $x \in K$ s.t. $f(x) = y$, but because $\{U_i\}_{i\in\tilde{I}}$ covers K, there is some $i \in \tilde{I}$ s.t. $x \in U_i = f^{-1}(V_i) \implies y \in V_i$. ■

> *Theorem 5.8: Continuous functions on compact metric space is uniformly continuous*
>
> A continuous function on a compact metric space X (to a metric space Y) is uniformly continuous.

Proof: since uniform continuity is phrased in the $\epsilon - \delta$ way, we work with ϵ and δ. Let $\epsilon > 0$ be arbitrary. Then we know for any $x \in X$ there is $\delta_x > 0$ s.t. $x' \in B(x, \delta_x) \implies |f(x') - f(x)| < \frac{\epsilon}{2}$. Then, $\{B(x, \frac{\delta_x}{2})\}_{x\in X}$ is an open cover of X, and so by compactness, we have a finite subcover $\{B(x_i, \frac{\delta_{x_i}}{2})\}_{i=1}^n$ for some $\{x_1, \dots, x_n\} \subseteq X$. Defining $\delta := \min_{i\in[n]} \frac{\delta_{x_i}}{2}$, I claim that for any $x, x' \in X$, $d(x, x') < \delta \implies |f(x') - f(x)| < \frac{\epsilon}{2}$. This is because $x \in X \implies x \in B(x_i, \frac{\delta_{x_i}}{2})$ for some $i \in [n]$, and so $d(x, x') < \delta \implies x' \in B(x_i, \delta_{x_i})$, and so $x, x' \in B(x_i, \delta_{x_i}) \implies |f(x) - f(x')| \leq |f(x) - f(x_i)| + |f(x_i) + f(x')| < \frac{\epsilon}{2} + \frac{\epsilon}{2} = \epsilon$. ■

5.1.3 Goal (2)

The above propositions about closed subsets of a compact space/set being compact, compact sets being closed in Hausdorff spaces, and continuous functions taking compact sets to compact sets addresses (2) from the **goal statement**, Section 3.2 above.

> *Theorem 5.9: Injective continuous maps on compact spaces admit continuous inverse*
>
> Given an injective continuous function $f : X \to \mathbb{R}^N$, if we have that X is compact and Hausdorff (or a metric space, by Lemma 5.6), then the inverse map $f^{-1} : \operatorname{im} f \to X$ is continuous.

Proof: to prove f^{-1} is continuous, we just need to prove that for any open set $U \subseteq X$, the inverse image $\{y \in \operatorname{im} f : f^{-1}(y) \in U\} = \{f(x) : x \in U\} = f(U)$ is open (the first equality is true $[f^{-1}(y) \in U$ for $y \in \operatorname{im} f] \iff [\exists! x \in U$ s.t. $f(x) = y]$ by injectivity). This is equivalent to proving that for any closed set $K \subseteq X$, $f(K)$ is closed. But because closed sets in X are compact and we know continuous functions take compact sets to compact sets, and compact sets are closed in Hausdorff spaces, $f(K)$ is closed and we are done. ■

5.1.4 Products

The above Heine-Borel theorem (Theorem 5.2) for $\mathbb{R}$ hold more generally on $\mathbb{R}^n$ for any $n \in \mathbb{N}$, and since the extension from closed and bounded intervals $[a, b]$ to arbitrary closed and bounded sets in $\mathbb{R}$ in the above proof is easily generalized to higher dimensions, we only need to prove the claim for closed and bounded rectangles $[a_1, b_1] \times \ldots \times [a_n, b_n]$. We show that indeed ANY product of compact sets is compact. The contents of this section stem from Lee's Proposition 4.36 (pg. 96) [6], and David G. Wright's 1994 article in the AMS "TYCHONOFF'S THEOREM" [13] (the whole article is well worth checking out). It may interest the reader in knowing that in fact Wright's proofs seem to be non-standard and actually remarkably simple, and according to Wright may be a University of Wisconsin rarity (yes, the other UW! How coincidental).

Let us first prove the claim for the product of two compact spaces (and hence for finite products by induction):

Lemma 5.10: Finite product of compact spaces is compact

For compact spaces X, Y, their Cartesian product $X \times Y$ is compact.

Proof using standard definition of compactness: suppose we have a collection $\mathscr{C}$ of open sets that covers $X \times Y$. Then, because X is compact and $X \times \{y\}$ is homeomorphic to X, $X \times \{y\}$ is compact and so there for any $y \in Y$ there is a finite collection $\mathscr{C}_y$ of open sets in $X \times Y$ s.t. their intersection with the subspace $X \times \{y\}$ (of $X \times Y$) forms a finite open cover of $X \times \{y\}$. I claim that this finite cover $\mathscr{C}_y$ of $X \times \{y\}$ has enough breathing room to actually cover $X \times V_y$ for some neighborhood V_y of y. This claim about "breathing room" is commonly known as the *tube lemma*.

Defining W_y to be the open set in $X \times Y$ that is the union of the open sets in $\mathscr{C}_y$, we have that because the product topology is generated by a basis of Cartesian products of open sets, any $(x, y) \in X \times \{y\}$ is contained in $U_{(x,y)} \times V_{(x,y)} \subseteq W_y$, so $\{U_{(x,y)} \times V_{(x,y)}\}_{x \in X}$ is an open cover of $X \times \{y\}$, and again by compactness there is a finite subcover, say $\{U_{(x,y)} \times V_{(x,y)}\}_{x \in \tilde{X}}$ for some finite $\tilde{X}$ (dependent on y). Then, $X \times \bigcap_{x \in \tilde{X}} V_{(x,y)} \subseteq \bigcup_{x \in \tilde{X}} U_{(x,y)} \times V_{(x,y)} \subseteq W_y$.

With this "tube lemma", we see that for every $y \in Y$, we have a finite open cover $\mathscr{C}_y$ for $X \times V_y$ where V_y is an open set containing y. Because Y is compact and $\{V_y\}_{y \in Y}$ is an open cover of Y, there is finite subcover, say $\{V_y\}_{y \in \tilde{Y}}$ for some finite $\tilde{Y} \subseteq Y$, and so $\bigcup_{y \in \tilde{Y}} \mathscr{C}_y$ is a finite subcover of $X \times Y$. ■

This proof does not generalize beyond finiteness, because the last step we basically end up saying that by cutting $X \times Y$ (or in the general finite case $X_1 \times \ldots \times X_n$) into strips (finitely many, by compactness) , finite covers of each of these strips can be gathered together into a finite cover the the whole product; this last step of course of "gathering finite covers into a finite cover" does not work when we have infinitely many strips due to there being infinitely many components in the product.

In a sense, the compactness condition is just a bit too "difficult" to work with. To prove something is compact, we need to show the existence of a finite subcollection of open sets that contains every point of the space! Perhaps it'd be easier if we could prove something about there being an open cover such that every finite subcollection is missing a point (proving existence of missing point may be easier than proving all points are covered). Hmm, but this is still proving something for every finite subcollection. We have so far that compactness $\iff$ no open cover s.t. every finite subcollection misses some point. That means that for a compact space X, if every finite subcollection misses a point, then the collection could not have been an open cover of X. The converse of this is also true, because if X satisfies this condition and we have an open cover $\mathscr{C}$ of X, then there must exist a finite subcover of $\mathscr{C}$ because if there wasn't, the condition would tell us that $\mathscr{C}$ wasn't actually an open cover of X. This almost-trivially equivalent reformulation of compactness gives us much more substantial material to start with; before we were just given an open cover of X and asked to whittle it away somehow someway, but now we are assuming that any finite subcollection misses a point to try to prove that the collection itself must miss a point.

We now repeat the proof Lemma 5.10 using this new definition of compactness. I will use ideas from the old proof though so it is important to read that one first.

Proof of Lemma 5.10 using "easier" but almost-trivially equivalent definition of compactness: suppose we have a collection $\mathscr{C}$ of open sets in $X \times Y$ with the property that no finite subcollection covers $X \times Y$. We would like to find some point $(x, y) \in X \times Y$ that is not covered by $\mathscr{C}$. I claim that there must be some $x \in X$ s.t. for any neighborhood $U \subseteq X$ containing x, the strip $U \times Y$ is not covered by any finite subcollection of $\mathscr{C}$; indeed if there wasn't, then for every $x \in X$, there is a neighborhood $U_x \subseteq X$ containing x s.t. the strip $U_x \times Y$ is covered by a finite subcollection, meaning that since $\{U_x\}_{x \in X}$ is an open cover of X, compactness of X gives a finite subcover and hence there is a finite subcollection covering $X \times Y$, a contradiction. Thus, we have such an x, say x_0, such that any strip $U \times Y$ $(x_0 \in U)$ is not covered by any finite subcollection of $\mathscr{C}$. Running the same argument again but this time using compactness of Y, we see that there is some $y_0 \in Y$ s.t. any open V containing y_0 is s.t. $(x_0, y_0) \in U \times V$ is not covered by any finite subcollection of $\mathscr{C}$. Because U, V are arbitrary open sets containing x_0, y_0 respectively, it must be that no set in $\mathscr{C}$ contains (x_0, y_0) because if it did we would have some small enough open basis set $U \times V$ covered by $\mathscr{C}$. ■

Theorem 5.11: Tychonoff's theorem; arbitrary product of compact spaces is compact

For an arbitrary collection of compact spaces $\{X_i\}_{i \in I}$, their product $\prod_{i \in I} X_i$ (with the product topology, Section 2.4) is also compact.

Proof: given any collection $\{X_i\}_{i \in I}$ of compact spaces, we use the axiom of choice to form a bijection with the ordinals to enforce a well-ordering of I. Then, we inductively define $x_i \in X_i$ s.t. for any open basis set W containing $\prod_{\alpha \leq i}\{x_\alpha\} \times \prod_{\beta > i} X_\beta$, no finite subcollection of $\mathscr{C}$ covers W (the existence of such $x_i \in X_i$ is just the proof above, where instead of using compactness of X we use the induction hypothesis, which is compactness of $\prod_{\alpha \leq i} X_\alpha$). Finally, the point $\prod_{i \in I}\{x_i\} \in \prod_{i \in I} X_i$ is not covered by any element of $\mathscr{C}$ (again if it were, we would find some open basis set covered by $\mathscr{C}$, but we have just proven that we can not). ■

Remarks: for alternative perspectives on Tychonoff, see `http://umu.diva-portal.org/smash/get/diva2:847896/FULLTEXT02.pdf` (for equivalence with the axiom of choice) and `https://people.clas.ufl.edu/kees/files/AlexanderTychonoff.pdf` (for a proof via the Alexander subbase lemma).

Question 5.1: Exercise

Using the Heine-Borel theorem for Euclidean space (really only $\mathbb{R}^2$), and that compactness is a topological property (preserved by homeomorphisms), prove that $[0,1)$ and S^1 are not homeomorphic. See Section 2.2.

5.2 Compactification

Above, we had the "philosophical definition" of compactness as "those on which we can not run away from everyone". We have seen just how useful compactness is, so here we consider the question of how to tinker with a non-compact space X a little bit and turn it into a compact space $\hat{X}$. The answer is that we add a point called ∞, so $\hat{X} = X \cup \{\infty\}$, and define the open sets around this point ∞ to be exactly $(X \setminus K) \cup \{\infty\}$ for any compact set $K \subseteq X$. That way any point "running way" from any compact K in the original X ends up "running toward" ∞. This is referred to as the "one-point compactification" of the space X.

Let us verify some properties of this construction [4]. Let us check that adding this new point and these new open sets makes things remain a topological space. Note that there are two "classes" of open sets here, "class one" which consists of all the original open sets, and "class two" which consists of all the sets containing the point ∞, i.e. all the sets $(X \setminus K) \cup \{\infty\}$ for compact $K \subseteq X$. It's true that "class one" open sets are closed under arbitrary unions and finite intersection, and it's true that "class two" open sets are closed under arbitrary unions — $\bigcup_{i \in I}(X \setminus K_i) \cup \{\infty\} = X \setminus (\bigcap_{i \in I} K_i) \cup \{\infty\}$, where $(\bigcap_{i \in I} K_i)$ is compact (**IF** we are in a space where compact $\implies$ closed) because it's closed (arbitrary intersection of closed sets is closed) and a subset of a compact set (any of the K_i), and so it's compact by Lemma 5.3 — and under finite intersections ($(X \setminus K_1) \cup \{\infty\}$ intersected with $(X \setminus K_2) \cup \{\infty\}$ is $(X \setminus (K_1 \cup K_2)) \cup \{\infty\}$, where $(K_1 \cup K_2)$ is compact because any open cover of it is also an open cover of K_1, K_2, so putting both finite subcovers together yields a finite subcover of $K_1 \cup K_2$).

Thus, any arbitrary union/finite intersection of open sets (both "class one" and "class two") can be simplified to just a union/intersection of a "class one" open set and "class two" open set. And indeed, $U \cap [(X \setminus K) \cup \{\infty\}]$ for a "class one" open set $U \subseteq X$ is just $U \cap (X \setminus K)$, where $(X \setminus K)$ is also a "class one" open set in X (**IF** we are in a space where compact $\implies$ closed), implying that the intersection is a "class one" open set of X; and dually $U \cup [(X \setminus K) \cup \{\infty\}] = (X \setminus (K \cap U^{\complement})) \cup \{\infty\}$, where $K \cap U^{\complement}$ is compact because it's a closed (intersection of closed sets, **IF** we are in a space where compact $\implies$ closed) subset of the compact set K.

Regarding the big "**IF**"s above: note that in the above analysis, we needed that our compact sets K were also closed in a couple different places. Recall that this is true in Hausdorff/T_2 spaces (Def. 5.4), but not necessarily true in general. Fortunately, this is not a big deal; we simply remedy the situation by asking that the "K" in the definition of "one-point compactification" be BOTH closed

AND compact.

We also check that the "one-point compactification" is actually compact, as the name suggests. Any open cover of $\hat{X}$ must cover ∞, and the only open sets covering ∞ are of the form $(X \setminus K) \cup \{\infty\}$. As K is compact in X, any open cover of K can be reduced to a finite subcover, so tacking back $(X \setminus K) \cup \{\infty\}$, we get overall a finite subcover, so indeed $\hat{X}$ is compact.

6 Embedding

This section is strongly inspired by §50 of Munkres [8]. This section tackles (1) from the **goal statement** above. We are looking for topological spaces homeomorphic to subsets of Euclidean space, but such subsets are metrizable (they inherit the metric on the Euclidean space in which it resides), and metrizability is a topological property, we should obviously consider only metrizable topological spaces. Since we already have (2) for compact spaces, and of course the abundance of nice properties we found out about compact spaces, let us start our exploration dealing with a compact metric space X. Basically, we consider the set $\mathcal{C}(X, \mathbb{R}^N)$ of continuous functions $X \to \mathbb{R}^N$ (for some $N \in \mathbb{N}$ that we will specify later), and we want to show that there exists some injective function in this set.

We use the square metric $|\mathbf{x} - \mathbf{y}| = \max\{|x_i - y_i|\}$ on $\mathbb{R}^N$ which induces the metric $|f - g| = \sup_{x \in X}\{|f(x) - g(x)|\}$ on $\mathcal{C}(X, \mathbb{R}^N)$. We define the following "measure of injectivity": $\Delta(f) := \sup\{\text{diam}(f^{-1}(\{\mathbf{p}\})) : \mathbf{p} \in f(X)\}$. Notice that $\Delta(f) = 0 \iff f$ is injective. We now define the sets $O_\epsilon := \{f \in \mathcal{C}(X, \mathbb{R}^N) : \Delta(f) < \epsilon\}$ (yes, the choice of the letter O is foreshadowing). We study these sets, and hope that we can prove that $\bigcap_{n=1}^{\infty} O_{1/n}$ is non-empty.

6.1 O_ϵ is Non-Empty

First, we should find out if these sets are even non-empty! So let us fix some $\epsilon > 0$. The clever idea is that we can use the fact that for some functions $\{\varphi_i\}_{i \in I}$ that make up a partition of unity (subordinate to some open cover $\{U_i\}_{i \in I}$), $\varphi_i(x), \varphi_i(y) > 0 \implies x, y \in U_i$. Thus, if we have an open cover s.t. $\text{diam}(U_i) < \epsilon$, then you can already see inklings of the argument coalesce.

Alright. Now let us do just this. Let $\eta > 0$ be some small quantity (less than ϵ) that **we will precisely determine later** (see Section 6.4.1). Compactness of X tells us that we can cover X by finitely many U_i (say $i \in [n]$) of diameter $\leq \eta$, since $\{B(x, \frac{\eta}{2})\}_{x \in X}$ is an open cover of X and compact means we can take a finite collection that still covers X. Let $\{\phi_i\}_{i=1}^n$ be a partition of unity subordinate to this open cover $\{U_i\}_{i=1}^n$. Let us now take n points $\{\mathbf{p}_i\}_{i=1}^\infty \subset \mathbb{R}^N$ and define the function $g(x) := \sum_{i=1}^n \phi_i(x)\mathbf{p}_i$. We will try to show that $g \in O_\epsilon$.

MAIN: So we want to show that $g(x) = g(y) \implies x, y \in U_i$ for some $i \in [n]$. Observe immediately that $g(x) = g(y) \iff \sum_{i=1}^n [\phi_i(x) - \phi_i(y)]\mathbf{p}_i = 0$. Obviously, we need some sort of linear independence to make it so that that this sum being 0 implies that for all $i \in [n]$, $\phi_i(x) = \phi_i(y)$, which helps us because we already know that there is some $i \in [n]$ s.t. $\phi_i(x) > 0$ (recall that $\sum_{i=1}^n \phi_i(x) = 1!$), and as

I said before $\phi_i(x) = \phi_i(y) > 0 \implies x, y$ in some $U_i \implies |x - y| \le \eta \implies \Delta(g) \le \eta < \epsilon \implies g \in O_\epsilon$ (x, y could have been in many other U_i, i.e. many $i \in [n]$ s.t. $\phi_i(x) = \phi_i(y) > 0$, but we just needed one).

Unfortunately, it is true that in $\mathbb{R}^N$, any set of $> N$ points will not be linearly independent. We have seemlingly come to a standstill — n is likely much much larger than N, and so no it is probably not the case that the $\{\mathbf{p}_i\}_{i=1}^n$ are linearly independent.

But, this is not really what's going on, since we do not actually have n terms in the sum because MOST of them are 0, assuming the U_i are mostly not overlapping $\iff$ most $\phi_i(x)$ are 0 for any given $x \in X$! So in fact we see that $v = \max\{k \in \mathbb{N} : \exists S \subseteq [n] \text{ s.t. } |S| = k \text{ and } \bigcap_{i \in S} U_i \neq \varnothing\}$ is an important quantity (v for o<u>**v**</u>erlap), since if somehow we could prove that we could manipulate any open cover of X in such a way that v is bounded, then we could set N at the very beginning to be $2v$ so that (for above fixed $x, y \in X$) our sums $\sum_{i=1}^n \phi_i(x)\mathbf{p}_i$, $\sum_{i=1}^n \phi_i(y)\mathbf{p}_i$ would each only have $\le v$ non-zero terms (for our previously fixed $x, y \in X$), implying that $\sum_{i=1}^n a_i\mathbf{p}_i$ would have $\le 2v = N$ non-zero terms.

It may help if you, the reader, pause here and consider arbitrary blobs in $\mathbb{R}$, $\mathbb{R}^2$, and maybe $\mathbb{R}^3$ and open covers of these blobs... you may find that you can always refine the open cover to get $\le \dim + 1$ overlap... a fact you find very interesting/suspicious. We will explore this further in Section 6.3.

We can in fact do ever so slightly better, taking advantage of the fact that the $a_i = \phi_i(x) - \phi_i(y)$ are not completely arbitrary, as they have to satisfy $\sum_{i=1}^n a_i = 1 - 1 = 0$. Thus, $\sum_{i=1}^n a_i\mathbf{p}_i = \sum_{i=1}^{n-1} a_i\mathbf{p}_i + \mathbf{p}_n(-\sum_{i=1}^{n-1} a_i) = \sum_{i=1}^{n-1} a_i(\mathbf{p}_i - \mathbf{p}_n)$ equaling 0 would imply $a_i = 0$ for all $i \in [n-1]$ (hence also for $i = n$) IF WE ALSO HAD that all the $(\mathbf{p}_i - \mathbf{p}_n)$ were linearly independent. Since this will be an important concept, we will define a collection of points $\{\mathbf{p}_i\}_{i=0}^k \subseteq \mathbb{R}^N$ being *geometrically independent* to mean that $\{(\mathbf{p}_i - \mathbf{p}_0)\}_{i=1}^k$ are linearly independent.

See pg. 309 in Munkres for visualization tips. Basically, any two distinct points are geometrically independent (g.i.); any three points that are not colinear form g.i. set; four points in $\mathbb{R}^3$ are g.i. if they are not coplanar; and in general, a set of $k+1$ points $\{\mathbf{p}_i\}_{i=0}^k$ are g.i. if they do not lie in the same k-plane (the affine subspace $\mathbf{p}_0 + \text{span}\{\mathbf{p}_1, \ldots, \mathbf{p}_k\}$).

By taking advantage of this extra condition $\sum_{i=1}^n a_i = 0$, we only need the sum $\sum_{i=1}^n a_i\mathbf{p}_i$ to have $\le N+1$ non-zero terms (instead of the above $\le N$ non-zero terms) in order to use a linear independence argument to get that all the $a_i = 0$ for $i \in [n]$, so in fact we can instead set $2v = N+1 \iff N = 2v-1$ at the beginning, giving us a reduction by 1 dimension. Not earth-shattering, but we'll take it.

6.2 General Position and the Baire Category Theorem

So now we are at the point where we want $\{\mathbf{p}_i\}_{i=1}^n$ to satisfy that any collection of size $\le N+1$ is geometrically independent. We say that that such a set is in *general position*. To prove the existence of a set in general position in $\mathbb{R}^N$, it is actually easier to construct one via perturbing an existing set very slightly than to actually give some kind of formula describing n points s.t. any collection of size $N+1$ is geometrically independent (we would have to check a pretty messy algebraic/computational condition on ALL collections of size $\leq N+1$!). The intuition behind this is that lower dimensional affine subspaces of $\mathbb{R}^N$ are super "thin", and so given some set of n points we should always be able to nudge the points very slightly to ensure that there is no co-planarity.

In light of this intuition about linear subspaces being "thin" and being able to nudge an arbitrarily small amount, we see that we want to prove that the complement of a a union of affine subspaces is dense in $\mathbb{R}^N$. Well, let us consider the case of a complement of a single affine subspace first. We can translate this affine subspace so that it contains the origin and becomes a linear subspace, so we only need to prove that linear subspaces of $\mathbb{R}^N$ have empty interior. So suppose we have a linear subspace; we can consider it to be the span of some vectors $\{\mathbf{v}_1, \ldots, \mathbf{v}_k\}$ for $\mathbf{v}_i \in \mathbb{R}^N$. Then, there must be some basis vector $\mathbf{e}_i \in \mathbb{R}^N$ that is not in this subspace (because if it was, then the subspace would be all of $\mathbb{R}^N$). Now suppose that the interior is non-empty; then there is $\mathbf{x}$ and $r > 0$ s.t. $B(\mathbf{x}, r)$ is contained within the linear subspace. But note that $\mathbf{x} + \frac{r}{2}\mathbf{e}_i \in B(\mathbf{x}, r)$, and hence in the linear subspace, meaning that $\frac{2}{r}[(\mathbf{x} + \frac{r}{2}\mathbf{e}_i) - \mathbf{x}] = \mathbf{e}_i$ is also in the subspace; contradiction. To extend this result to the complement of a union of (finite, in our scenario) affine subspaces, we prove the following theorem:

Theorem 6.1: Baire category theorem

A countable intersection of open dense sets in Euclidean space is dense.

Proof: let $U_1, \ldots$ be our open dense sets and let $X = \mathbb{R}^N$. To prove that $\bigcap_{n=1}^{\infty} U_n$ is dense, we need to show that for every (non-empty) open set $V \subseteq \mathbb{R}^N$, there is some point $x \in V \cap \bigcap_{n=1}^{\infty} U_n$. Well, because U_1 is dense and open, we know that there exists $x_1 \in U_1 \cap V$ and $r_1 \in (0,1)$ s.t. $B(x_1, r_1) \subseteq U_1 \cap V$. We can continue inductively, at each step finding $x_n \in U_n \cap B(x_{n-1}, r_{n-1})$ and $r_n \in (0, \frac{1}{n})$ s.t. $B(x_n, r_n) \subseteq U_n \cap B(x_{n-1}, r_{n-1})$.

Now since $n \geq m \implies x_n \in B(x_m, r_m)$ and $r_m \to 0$, we see that the $\{x_n\}_{n=1}^{\infty}$ form a Cauchy sequence, and because $\mathbb{R}^N$ is complete ($\mathbb{R}$ is complete by construction, and a Cauchy sequence in $\mathbb{R}^N$ "projects down" to a Cauchy sequence in each component, and easy triangle inequality arguments yield that the $\mathbb{R}^N$ Cauchy sequence converges pointwise to the point in $\mathbb{R}^N$ where each component is the limit of each component's "projected" Cauchy sequence), we see that the $x_n \to x$ for some $x \in X$. For every $m \in \mathbb{N}$, since all the x_n for $n > m$ are in $B(x_m, r_m)$, we know that $x \in \overline{B(x_m, r_m)}$. Although in the setup we have, it is not necessarily true that $\overline{B(x_m, r_m)} \subseteq B(x_{m-1}, r_{m-1}) \cap U_{m-1}$, we can easily go back through the 1st paragraph of the proof and put this condition in place, using the fact that $\overline{B(x, \frac{r}{2})} \subseteq B(x, r)$ for any $x \in X$, $r > 0$. ■

Remarks: it is easy to easy that the above theorem holds more generally for all **complete pseudometric** spaces X, not just $X = \mathbb{R}^N$. **Reminder:** if you need a refresher on completeness, check out the below subsubsection, Section 6.2.1.

With this theorem in mind, we can now rigorously prove by induction our construction of a set in general position by nudging I introduced above. Let $\{\tilde{\mathbf{p}}_i\}_{i=1}^n$ be a set of n points in $\mathbb{R}^N$ and $\delta > 0$. Then, we can construct a set $\{\mathbf{p}_i\}_{i=1}^n$ points in general position by induction s.t. $|\tilde{\mathbf{p}}_i - \mathbf{p}_i| < \delta$ for all $i \in [n]$. The base case of one element is trivial. Now suppose we are given $\{\mathbf{p}_i\}_{i=1}^k$ in general position and $\tilde{\mathbf{p}}_{k+1}$. Since $\{\mathbf{p}_i\}_{i=1}^k$ is in general position, any subset (say of size j) containing $\leq N$ elements is geometrically independent and determines a $(j-1)$-dimensional affine subspace, and so the BCT tells us that the complement of the union of all such affine subspaces is dense in $\mathbb{R}^N$. That means that there is some point $\mathbf{p}_{k+1} \in B(\tilde{\mathbf{p}}_{k+1}, \delta)$ also in this complement, and indeed $\{\mathbf{p}_i\}_{i=1}^{k+1}$ is a set in general position.

6.2.1 Completeness Refresher

A space X is complete if every Cauchy sequence (i.e. every sequence $\{x_n\}_{n=1}^\infty$ s.t. for every $\epsilon > 0$, there is $N \in \mathbb{N}$ s.t. $n, m \geq N \implies d(x_n, x_m) < \epsilon$) converges (i.e. there is $x \in X$ s.t. the Cauchy sequence $\{x_n\}_{n=1}^\infty$ satisfies $\lim_{n\to\infty} x_n = x \iff \forall \epsilon > 0, \exists N \in \mathbb{N}$ s.t. $n \geq N \implies d(x_n, x) < \epsilon$). The standard completeness axiom in $\mathbb{R}$, the least upper bound property that states that for any (non-empty) set $S \subseteq \mathbb{R}$ with an upper bound, there is a LEAST upper bound/supremum (i.e. smallest $s \in \mathbb{R}$ s.t. all elements of S are $\leq s$) is equivalent to the Cauchy sequence definition, as follows.

($\impliedby$), proof from Wiki: for any non-empty set $S \subseteq \mathbb{R}$ ($a_1 \in S$) with an upper bound b_1, we define sequences $\{a_n\}_{n=1}^\infty$ and $\{b_n\}_{n=1}^\infty$ inductively by checking if $\frac{a_n+b_n}{2}$ is an upper bound of S, and defining $a_{n+1} = a_n, b_{n+1} = \frac{a_n+b_n}{2}$ if it is, and $a_{n+1} > \frac{a_n+b_n}{2}, = b_{n+1} = b_n$ if it's not (such an element $a_{n+1} \in S$ exists because in the second case, $\frac{a_n+b_n}{2}$ is NOT an upper bound of S). Then, $a_1 \leq a_2 \leq \ldots \leq b_2 \leq b_1$ satisfy $|a_n - b_n| \to 0$ as $n \to \infty$, and hence by the Cauchy property both converge to the same limit L, which is indeed the least upper bound of S.

($\implies$): for a Cauchy sequence $\{x_n\}_{n=1}^\infty$, we can use the Cauchy property to prove that the limsup, $\limsup_{n\to\infty} x_n := \inf_{N\in\mathbb{N}}\{\sup_{n\geq N} x_n\}$, and and liminf, $\liminf_{n\to\infty} x_n := \sup_{N\in\mathbb{N}}\{\inf_{n\geq N} x_n\}$, are equal, to say a quantity we denote L (where sup and inf are defined because of the least upper bound property), and L is indeed the limit of the Cauchy sequence.

6.3 Topological Dimension

Looking back at the core of the argument (see the paragraph labeled **MAIN** in Section 6.1), it looks like with the above section on general position (Section 6.2), we are almost done. Not completely, though, because we still have that mysterious quantity v. Recall that we wanted then to refine our open cover $\{U_i\}_{i=1}^n$ to get this "overlap" measure v to be bounded. Another word we will use for "overlap measure" is "overlap order". Some playing around will yield the conjecture that we can always get $v \leq \dim(\mathbb{R}^N) + 1 = N + 1$. Let us first prove that there *exists* an open cover of $\mathbb{R}^N$ that has order $v = N + 1$ (let's further require that the open sets be small; if our goal is to refine $\{U_i\}_{i=1}^n$ which is already a cover of small open sets, it will do us no good to consider partitions consisting of unbounded sets). This section is taken from Example 3 and Theorem 50.6 in §50 of Munkres [8].

Consider the following open cover of $\mathbb{R}^2$ (split across three panels so one can see the different types of open sets):

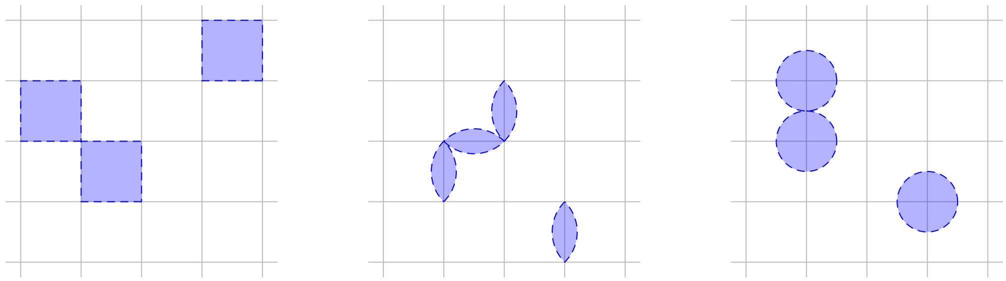

We don't want the cover of the edges (in $\mathbb{R}^2$) or faces and edges (in $\mathbb{R}^3$) and so on to overlap, so for one of these subspaces $S \subseteq \mathbb{R}^N$ (picture a face or edge of cube in $\mathbb{R}^3$), take the $U_S \supseteq S$ to be $U_S = \bigcup_{x \in S} B(x, r_x)$ where $r_x = \frac{1}{2} \min\{x_1 - \lfloor x_1 \rfloor, \lceil x_1 \rceil - x_N, \ldots, x_N - \lfloor x_N \rfloor, \lceil x_N \rceil - x_N\}$. In $\mathbb{R}^2$, the subspaces were $\mathscr{S}_2$ (open squares), $\mathscr{S}_1$ (edges of squares covered by lens-shaped open sets in panel 2 of above figure), and $\mathscr{S}_0$ (vertices of squares covered by circles in panel 3 of above figure). In $\mathbb{R}^3$, $\mathscr{S}_3$ will denote all the open 3-cubes, $\mathscr{S}_2$ will denote all the open faces (2-cubes) of the open 3-cubes ("open" when thought of as in a 2-plane), $\mathscr{S}_1$ will denote all the open edges (1-cubes) of the open 3-cubes, and lastly $\mathscr{S}_0$ will denote all the vertices of the open 3-cubes. In general in $\mathbb{R}^N$, letting $\mathscr{I}$ denote the set of open intervals $\{(n, n+1) : n \in \mathbb{Z}\}$ and $\mathscr{P}$ denote the set of points $\{\{n\} : n \in \mathbb{Z}\}$, we have $\mathscr{S}_i$ ($i \in \{0, \ldots, N\}$) to be the set of all products $A_1 \times \ldots \times A_N$ where exactly i sets out of $\{A_i\}_{i=1}^n$ are in $\mathscr{I}$ and the rest are in $\mathscr{P}$.

Now is it possible that for some open covers $\{U_i\}_{i \in I}$ of some subsets $X \subseteq \mathbb{R}^N$, we can make scale this open cover $\mathscr{U}$ of $\mathbb{R}^N$ small enough that each individual open set in the cover lies completely within U_i for some $i \in I$ (so that for each $U_i \in \{U_i\}_{i \in I}$, we can just take the refinement to be the union of all sets in $\mathscr{U}$ contained strictly within U_i not already chosen by other U_i)? That is to say, since we know that each U_i is open, so for any U_i and for each $x \in X$, we have some r_x s.t. $B(x, r_x) \subseteq U_i$; we can then define $r(x, U_i) > 0$ to be the supremum $\sup\{r > 0 : B(x, r) \subseteq U_i\}$, and $r(x)$ to be $\sup\{r(x, U_i) : i \in I\}$. I claim that this is in fact a positive continuous function on X!

Hence, if X is compact, then because continuous functions attain their extreme values (we know from Theorem 5.7 that the image of r on X is a compact subset of $(0, \infty)$), $r(x)$ is bounded below by some $\delta > 0$ for all $x \in X$, so in fact we have shown that if X is compact, then scaling the open cover of $\mathbb{R}^N$ by $\frac{\delta}{2}$ (call this open cover $\mathscr{U}$), we can refine the open cover $\{U_i\}_{i=1}^n$ (finite because we are taking X compact) to have overlap order $v \leq N + 1$ (let $V_1 = \bigcup\{U \in \mathscr{U} : U \subseteq U_1\}$, and then $V_j = \bigcup\{U \in \mathscr{U} : U \subseteq U_j, U \not\subseteq U_i \text{ for } i < j\}$). This is just spelling out more explicitly what is in the parentheses in the previous paragraph.

The proof of the continuity of $r(x)$ is as follows: let $\epsilon > 0$ be arbitrary. Now for any $\delta > 0$ and $i \in I$, we see that for any $x' \in B(x, \delta)$, $B(x', r(x, U_i) - \delta) \subseteq B(x, r(x, U_i)) \subseteq U_i \implies r(x', U_i) \geq r(x, U_i) - \delta$

and simultaneously $B(x, r(x', U_i) - \delta) \subseteq B(x', r(x', U_i)) \subseteq U_i \implies r(x, U_i) \geq r(x', U_i) - \delta$, and so $|r(x', U_i) \geq r(x, U_i)| \leq \delta$. Then, $r(x') = \sup\{r(x', U_i) : i \in I\} \geq \sup\{r(x, U_i) - \delta : i \in I\} = r(x) - \delta$, and similarly, $r(x) = \sup\{r(x, U_i) : i \in I\} \geq \sup\{r(x', U_i) - \delta : i \in I\} = r(x') - \delta$, and so taking $\delta = \frac{\epsilon}{2}$, we see that $|x - x'| < \delta \implies |r(x) - r(x')| \leq \delta < \epsilon$, and we are done. This result is known as the Lebesgue number lemma:

Lemma 6.2: Lebesgue number lemma

For any compact metric space X and open cover $\{U_i\}_{i \in I}$ of X, there is $\delta > 0$ (called the **Lebesgue number** of the cover) s.t. for all $x \in X$, there is $U \in \{U_i\}_{i \in I}$ s.t. $B(x, \delta) \subseteq U$ (i.e. there is $\delta > 0$ s.t. any δ-ball fits ENTIRELY within some open set of the cover)

6.3.1 Summary So Far

We have proven that:

Lemma 6.3: Compact subsets of $\mathbb{R}^m$ have topological dimension $\leq m$

Any open cover of a compact subset of $\mathbb{R}^m$ has a refinement that has overlap order $v \leq m + 1$

and so indeed we begin to see rigorously a connection between the minimal overlap order of an open cover of a set and its dimension. Of course, there's also the previous connection with the minimal overlap order v of an open covering of X and the dimension of the Euclidean space we could embed X in (see Section 6.1), so indeed this quantity has earned its name:

Definition 6.4: Topological dimension

The topological dimension of a space X is the smallest number $m \in \mathbb{N}$ such that any open cover of X can be refined to an open cover with overlap order $v \leq m + 1$ (where again the overlap order v of a cover is the smallest $v \in \mathbb{N}$ s.t. for any $x \in X$, x lies in $\leq v$ sets of the cover). If there is no such number, we say the topological dimension is infinite. Denoted "$\dim X$"

6.3.2 Topological Dimension of Subspaces

Knowing the topological dimension of a space X, we do know some information about the topological dimension of the subspaces of X. The following is Theorem 50.1 (pg. 306) in Munkres [8].

Lemma 6.5: $\dim Y \leq \dim X$ for closed subspaces $Y \subseteq X$

For closed subspaces $Y \subseteq X$ where X has finite topological dimension, $\dim Y \leq \dim X$

Proof: suppose X has topological dimension m. Let $\mathcal{U}_Y$ be an open covering of Y with sets open in the subspace topology of Y (i.e. elements of $\mathcal{U}_Y$ look like $(U \cap Y)$ for some open U in X; let the collection of such U be denoted as $\mathcal{U}_X$). Then, $\mathcal{U}_X \cup \{X \setminus Y\}$ is an open cover of X, so we can refine this so that the overlap order of this cover is $\leq m + 1$. $\{X \setminus Y\}$ does not intersect with Y

(hence neither will any refinement of it), so the refinement $\mathscr{U}'_X$ of $\mathscr{U}_X$ still covers Y, meaning intersecting Y with all the $U' \in \mathscr{U}'_X$ yields a refinement of $\mathscr{U}_Y$ that covers Y with overlap order $\leq m+1$. ■

There is sort of a converse to this — knowing the topological dimensions of closed subspaces Y, Z s.t. $X = Y \cup Z$ gives us information about the topological dimension of X. This is Theorem 50.2/Corollary 50.3 (pg. 307) in Munkres [8].

Lemma 6.6: $\dim X = \max\{\dim Y, \dim Z\}$ *for closed subspaces* $Y \cup Z = X$

If $X = Y \cup Z$ for closed subspaces $Y, Z \subseteq X$ with finite topological dimension, $\dim X = m := \max\{\dim Y, \dim Z\}$. By induction, one can extend to finitely many closed subspaces $X = Y_1 \cup \ldots \cup Y_n \implies \dim X = \max\{\dim Y_1, \ldots, \dim Y_n\}$.

Proof: it suffices to prove that $\dim X \leq m$ (the other direction follows from the above lemma, Lemma 6.5). Here's the plan: we prove that for any open cover $\mathscr{A}$ of X, we can refine $\mathscr{A}$ to an open cover $\mathscr{B}$ of X s.t. the overlap order at any point of Y is $\leq m+1$ (where for every $A \in \mathscr{A}$ there is exactly one $B \in \mathscr{B}$ satisfying $B \subseteq A$, and moreover all $B \in \mathscr{B}$ are one of these such B's); then treating "$\mathscr{B}$" as the new "$\mathscr{A}$" and "Z" as the new "Y", we can refine $\mathscr{B}$ to an open cover $\mathscr{C}$ s.t. the overlap order at any point of Z is $\leq m+1$ (where for every $B \in \mathscr{B}$ there is exactly one $C \in \mathscr{C}$ satisfying $C \subseteq B$, and moreover all $C \in \mathscr{C}$ are one of these such C's), meaning $\mathscr{C}$ has a total overlap order of $\leq m+1$ ($\mathscr{C}$ must have overlap order at any point of Y bounded by $\leq m+1$ because otherwise, we would have $> m+1$ sets of $\mathscr{C}$ covering a point $y \in Y$, but because $\mathscr{C}$ is refinement of $\mathscr{B}$ and there is a 1-1 correspondence between $\mathscr{C}$ and $\mathscr{B}$, i.e. every $C \in \mathscr{C}$ corresponds to a distinct $B \in \mathscr{B}$ s.t. $C \subseteq B$, there would be $> m+1$ sets in $\mathscr{B}$ covering y, contradicting the properties of $\mathscr{B}$).

Ok, onward to the existence of $\mathscr{B}$ with the properties promised above. The collection $\mathscr{A}_Y := \{Y \cap A : A \in \mathscr{A}\}$ is an open (in subspace topology of Y) cover of Y, and because $\dim Y \leq m$, there is a refinement $\mathscr{B}_Y$ of $\mathscr{A}_Y$ with overlap order $\leq m+1$ (on Y). By the definition of subspace topology, every $B_Y \in \mathscr{B}_Y$ is of the form $B_Y = U_{B_Y} \cap Y$ for some U_{B_Y} open in X. Because $\mathscr{B}_Y$ is a refinement of $\mathscr{A}_Y$, there is some $A_{B_Y} \in \mathscr{A}$ s.t. $B_Y \subseteq A_{B_Y} \cap Y \subseteq A_{B_Y}$; then $B_Y = (U_{B_Y} \cap A_{B_Y}) \cap Y$ where $(U_{B_Y} \cap A_{B_Y})$ is a refinement of $A_{B_Y} \in \mathscr{A}$. Define $\tilde{\mathscr{B}} := \{(U_{B_Y} \cap A_{B_Y}) : B_Y \in \mathscr{B}_Y\} \cup \{A \setminus Y : A \in \mathscr{A}\}$, and note that it is an open cover of X and refinement of $\mathscr{A}$ which indeed has overlap order $\leq m+1$ at every point of Y.

Finally, we can put $\mathscr{B}$ with the 1-1 correspondence with $\mathscr{A}$ by defining the labeling $L : \tilde{\mathscr{B}} \to \mathscr{A}$ mapping each $\tilde{B} \in \tilde{\mathscr{B}}$ to an $A \in \mathscr{A}$ that contains it, and defining for each A the corresponding $B := \bigcup_{\tilde{B} \in L^{-1}(A)} \tilde{B}$, with $\mathscr{B}$ the collection of such B's (can extend L to be a function $\mathscr{B} \to \mathscr{A}$ by mapping $L(B) = A$ for $B := \bigcup_{\tilde{B} \in L^{-1}(A)} \tilde{B}$); observe that if $B_1 \neq B_2$ are in $\mathscr{B}$, their corresponding A_1, A_2 are not equal since if they were, we would have $B_1 = \bigcup_{\tilde{B} \in L^{-1}(A_1)} \tilde{B} = \bigcup_{\tilde{B} \in L^{-1}(A_2)} \tilde{B} = B_2$. ■

6.4 O_ϵ is Dense and Open

With our strategy of perturbing points slightly used in the proof above that O_ϵ is non-empty, we can actually show that O_ϵ is dense in $\mathcal{C}(X, \mathbb{R}^N)$.

6.4.1 Denseness

Suppose we have arbitrary $f \in \mathcal{C}(X, \mathbb{R}^N)$ and $\delta > 0$; we show that $|f - g| < \delta$ for some $g \in O_\epsilon$ (where recall we are using $|\cdot|$ to denote the supremum norm on $\mathcal{C}(X, \mathbb{R})$). We again consider a function $g(x) := \sum_{i=1}^n \phi_i(x)\mathbf{p}_i$ (where $\{\phi_i\}_{i=1}^n$ is a partition of unity subordinate to $\{U_i\}_{i=1}^n$, a open cover of X of order $v \leq N+1$ s.t. $\text{diam}(U_i) \leq \eta$), where this time $\tilde{\mathbf{p}}_i$ will be $f(x_i)$ for some arbitrarily chosen $x_i \in U_i$ and $\mathbf{p}_i$ will be a slight nudging of $\tilde{\mathbf{p}}_i$ s.t. $|\tilde{\mathbf{p}}_i - \mathbf{p}_i| < \frac{\delta}{2}$ for all $i \in [n]$ and $\{\mathbf{p}_i\}_{i=1}^n \subseteq \mathbb{R}^N$ is a set in general position.

Then, we have that for any $x \in X$,

$$\begin{aligned} |g(x) - f(x)| &= \left|\sum_{i=1}^n \phi_i(x)\mathbf{p}_i - f(x)\sum_{i=1}^n \phi_i(x)\right| = \left|\sum_{i=1}^n \phi_i(x)(\mathbf{p}_i - f(x))\right| \\ &\leq \sum_{i=1}^n \phi_i(x) \cdot |\mathbf{p}_i - f(x_i)| + \sum_{i=1}^n \phi_i(x) \cdot |f(x_i) - f(x)|. \end{aligned}$$

We already know that the first sum is $< \frac{\delta}{2}$ because $|\tilde{\mathbf{p}}_i - \mathbf{p}_i| < \frac{\delta}{2}$ for all $i \in [n]$. To get the second sum $< \frac{\delta}{2}$, we can just specify that η (which recall we specified earlier to just be $< \epsilon$) is furthermore small enough that $d(x_i, x) < \eta \implies |f(x_i) - f(x)| < \frac{\delta}{2}$ (which we can do because continuous f on compact metric space $X = (X, d)$ is uniformly continuous; see Section 5.1), which suffices since $\phi_i(x) > 0 \implies x \in U_i \implies d(x, x_i) < \eta$. Then, we just have the proof from Section 6.1, which I sketch again here: $g(x) = g(y) \iff \sum_{i=1}^n [\phi_i(x) - \phi_i(y)]\mathbf{p}_i = 0 \implies [\phi_i(x) - \phi_i(y)] = 0$ for all $i \in [n]$ because $\{\mathbf{p}_i\}_{i=1}^n$ is in general position, i.e. any collection of size $\leq N+1$ is geometrically independent (where for any given $x \in X$ the sum has $\leq 2v = N+1$ non-zero terms) — again, if that was confusing, refresh with another look at Section 6.1.

6.4.2 Openness

Finally, we show that O_ϵ is open, and then we can use the Baire category theorem (Theorem 6.1) to see that $\bigcap_{n=1}^\infty O_{1/n} \neq \varnothing$, thus proving the existence of an injective continuous function in $\mathcal{C}(X, \mathbb{R}^N)$ (our long sought after goal (1), Section 3.2, though only for compact metric spaces X as of now). Suppose we have $f \in O_\epsilon$. Then, we have $\Delta(f) < \eta < \epsilon$ for some $\eta > 0$, i.e. $d(x, y) < \eta$ for any $x, y \in X$ s.t. $f(x) = f(y)$. Then, the function $|f(x) - f(y)| : X \times X \to \mathbb{R}$ is strictly positive on the set $A := \{(x, y) \in X \times X : d(x, y) \geq \eta\} = d^{-1}([\eta, \infty))$. Because d is a continuous function on $X \times X$ ($d(x, y) < \epsilon \implies d(x, y) < \epsilon$), A is a closed set in the compact set $X \times X$, and hence is also compact, meaning that $|f(x) - f(y)| \geq \delta > 0$ on A.

I claim that $B(f, \frac{\delta}{2})$ (i.e. the set of all $g \in \mathcal{C}(X, \mathbb{R}^N)$ s.t. $|f - g|_\infty < \frac{\delta}{2}$) is contained in O_ϵ. This is because for any such g, $|f(x) - f(y)| \geq \delta$ implies that $|g(x) - g(y)| > 0$ (if $g(x) = g(y)$, then because $|f - g|_\infty < \frac{\delta}{2}$, $|f(x) - f(y)| \leq |f(x) - g(x)| + |g(y) - f(y)| < \frac{\delta}{2} + \frac{\delta}{2}$; contradiction), i.e. we have that $|g(x) - g(y)|$ is strictly positive on A, meaning that $g(x) = g(y) \implies d(x, y) < \eta < \epsilon$, i.e. $\Delta(g) \leq \eta < \epsilon$.

6.5 Beyond Compactness

To reiterate what we have done so far, we have proven that for compact (hence paracompact, for partitions of unity) metric spaces (hence normal for Urysohn, by Lemma 4.8) with topological dimension $\leq m$ (i.e. overlap number/order $v \leq m+1$), we can embed X into $\mathbb{R}^{2m+1}$. Let us see if we can extend this to to non-compact metric spaces X; a natural first step would be to **assume that** X has a metric d, that $X = \bigcup_{n=1}^{\infty} K_n$ for compact subsets of X (this is called being **σ-compact** — visualize them as nested sets, so like bigger and bigger "compact approximations" of X), and that the topological dimension of each K_n is $\leq m$ (with $N := 2m + 1$). We will add more specifications to X, but let us start of with these ideas. This section is inspired by Exercise 6 in §50 of Munkres [8].

Recall that we started with the assumption that X is compact because Theorem 5.9 tells us that the inverse of a injective continuous maps on compact spaces is continuous, hence fulfilling ② of the goal (Section 3.2). This is not true in arbitrary spaces, but we will need something similar for $f : X \to \mathbb{R}^N$ for X satisfying (at least) the conditions specified in the first paragraph of this subsection. Recall from the subsection on compactification (Section 5.2) that we can always add one point ∞_X to a topological space X and some new open sets to make a compact space $\hat{X}$. Well, if we could somehow extend $f : X \to \mathbb{R}^N$ to a continuous injective function $\hat{f} : \hat{X} \to \mathbb{R}^N$, then by Theorem 5.9, $\hat{f}$ would be a homeomorphism between $\hat{X}$ and $\operatorname{im}(\hat{f})$, and maybe that would yield that f is a homeomorphism between X and $\operatorname{im} f = \operatorname{im}(\hat{f}) \setminus \{f(\infty_X)\}$.

One problem is that we know from Theorem 5.7 that since $\hat{f}$ is a continuous map on a compact space, the image $\operatorname{im}(\hat{f}) = \operatorname{im} f \cup \{f(\infty_X)\} \subseteq \mathbb{R}^N$ must be compact as well. We know from Heine-Borel (Theorem 5.2) that the compact sets of $\mathbb{R}^N$ are exactly those that are closed and bounded, and so if $\operatorname{im} f$ is unbounded, then automatically $\operatorname{im}(\hat{f})$ can not be compact, and so there will not be a continuous extension $\hat{f} : \hat{X} \to \mathbb{R}^N$ of $f : X \to \mathbb{R}^N$. If we consider a continuous extension to the COMPACTIFICATION of $\mathbb{R}^N$, say $\hat{f} : \hat{X} \to \hat{\mathbb{R}^N}$, then $\operatorname{im}(f)$ can just be taken to be closed (by Lemma 5.3, which says that closed subsets of a compact space are also compact).

Given $f : X \to \mathbb{R}^N$, we can define any sort of INJECTIVE extension $\hat{f} : \hat{X} \to \hat{\mathbb{R}^N}$ by just letting $\hat{f}(\infty_X) \in \hat{\mathbb{R}^N} \setminus \operatorname{im} f$ where $\operatorname{im} f \subseteq \mathbb{R}^N$. It could be that $\operatorname{im} f = \mathbb{R}^N$, in which case the only possibility is $\hat{f}(\infty_X) = \infty_{\mathbb{R}^N}$. This possibility of course works in all cases, so let's take this and run with it for now.

We now have to ascertain whether or not $\hat{f}$ is continuous, where again $\hat{f} = f$ on X and $\hat{f}(\infty_X) = \infty_{\mathbb{R}^N}$. We know that $\hat{f}$ is continuous everywhere on X, so we just need to see if $\hat{f}$ is continuous at ∞_X; i.e. we have to see if $\hat{f}^{-1}((\mathbb{R}^N \setminus K') \cup \{f(\infty_X)\}) = f^{-1}(\mathbb{R}^N \setminus K') \cup \{\infty_X\}$ (compact $K' \subset \mathbb{R}^N$) is open in $\hat{X}$ (recall also from Section 5.2 that the only open sets of $\hat{X}$ containing ∞_X are of the form

$(X \setminus K) \cup \{\infty_X\}$ for compact $K \subseteq X$).

It suffices to prove this for $\hat{f}^{-1}((\mathbb{R}^N \setminus \overline{B(0,r)}) \cup \{f(\infty_X)\})$, since any compact $K' \subset \mathbb{R}^N$ is contained in some $B(0,r)$, and $(\mathbb{R}^N \setminus K') \cup \{f(\infty_X)\} = (B(0,r) \setminus K') \cup (\mathbb{R}^N \setminus \overline{B(0,r)}) \cup \{f(\infty_X)\}$, where we already know the inverse image of $(B(0,r) \setminus K')$ is open in X, and the inverse image of a union is the union of the inverse images. That is to say, we only need that $f^{-1}(\mathbb{R}^N \setminus \overline{B(0,r)}) = X \setminus K \iff f^{-1}(\overline{B(0,r)}) = K$ for some compact $K \subset X$. It furthermore suffices that $f^{-1}(\mathbb{R}^N \setminus \overline{B(0,r)}) \supseteq X \setminus K \iff f^{-1}(\overline{B(0,r)}) \subseteq K$ for some compact $K \subset X$, because f being continuous and $\overline{B(0,r)}$ being closed in $\mathbb{R}^N$ means that $f^{-1}(\overline{B(0,r)})$ is closed, and we know (Lemma 5.3) that a closed subset of compact K is also compact.

Thus, we have shown that if $f : X \to \mathbb{R}^N$ satisfies that

> *Definition 6.7: $f(x) \to \infty$ as $x \to \infty$, or "going to infinity"*
>
> For any $r > 0$, there is a compact $K \subseteq X$ s.t. $|f(x)| > r$ for all $x \in X \setminus K$

then the extension $\hat{f} : \hat{X} \to \hat{\mathbb{R}^N}$ defined by $\hat{f}(\infty_X) = \infty_{\mathbb{R}^N}$ is continuous.

This shows that $\hat{X}$ is homeomorphic to $\operatorname{im}(\hat{f})$ (via the homeomorphism $\hat{f}$). From here, it is easy to see that f must also be a homeomorphism (i.e. is an open map), because $\hat{f}$ being an open map means that for any open $U \subseteq X$ (which note is also open in $\hat{X}$), $\hat{f}(U) = f(U)$ is also open in $\hat{\mathbb{R}^N}$; but there are only two "classes" of open sets in $\hat{\mathbb{R}^N}$, the standard open sets of $\mathbb{R}^N$ and the sets of the form $K^\complement \cup \{\infty_{\mathbb{R}^N}\}$ for compact $K \subset \mathbb{R}^N$ (see Section 5.2), and of course $f(U)$ is a "class one" set (i.e. a standard open set of $\mathbb{R}^N$) as it does not contain $\infty_{\mathbb{R}^N}$.

6.5.1 Existence of Injective Function

We learned above that if $f : X \to \mathbb{R}^N$ is a continuous injective function that satisfies $f(x) \to \infty$ as $x \to \infty$ (Def. 6.7), then f is a homeomorphism between X and $\operatorname{im} f \subseteq \mathbb{R}^N$. This ② of the goal (Section 3.2), so now we seek to prove ① of the goal, i.e. that there is an injective $f \in \mathcal{C}(X, \mathbb{R}^N)$ that in addition satisfies $f(x) \to \infty$ as $x \to \infty$.

Let us define $O_\epsilon(K) := \{f \in \mathcal{C}(X, \mathbb{R}^N) : \Delta(f|_K) < \epsilon\}$ (so the same O_ϵ from Section 6 except restricted to K, where K is some compact set $K \subseteq X$). The proofs from Section 6.4 essentially hold over to prove that $O_\epsilon(K)$ is dense and open for all $\epsilon > 0$ and compact $K \subseteq K$. I write out some of the details more explicitly below.

Openness: suppose we have $f \in O_\epsilon(K)$ (i.e. $\Delta(f|_K) < \eta < \epsilon$), then $|f(x) - f(y)| : X \times X \to \mathbb{R}$ is strictly positive on $A := \{(x,y) \in K \times K : d(x,y) \geq \eta\} = d^{-1}([\eta, \infty))$, a closed set (by continuity of d) in the compact set $K \times K$, hence compact, meaning $|f(x) - f(y)| \geq \delta > 0$ on A.

I claim that $B(f, \frac{\delta}{2}) \subseteq \mathcal{C}(X, \mathbb{R}^N)$ is contained in $O_\epsilon(K)$, because for any such g, $|f(x) - f(y)| \geq \delta \implies |g(x) - g(y)| > 0$ (again, more detailed proof in Section 6.4.2), so $|g(x) - g(y)|$ is strictly positive on A, meaning for $x, y \in K$, $g(x) = g(y) \implies d(x,y) < \eta < \epsilon$, implying that $\Delta(g|_K) \leq \eta < \epsilon$.

Denseness: the proof from Section 6.4.1 (applied to K) tells us that for all $f \in \mathcal{C}(X, \mathbb{R}^N) \subseteq \mathcal{C}(K, \mathbb{R}^N)$ and $\delta > 0$, there is $\tilde{g} \in \mathcal{C}(K, \mathbb{R}^N)$ s.t. $\Delta(\tilde{g}|_K) < \epsilon$ (so $g \in O_\epsilon(K)$) and $\max_{x \in K} |f(x) - \tilde{g}(x)| < \delta$. Then, we can extend $f - \tilde{g} : K \to [-\delta, \delta]^N$ to X continuously by using Tietze (Theorem 4.9) on each coordinate (we can apply Tietze because X is a metric space hence normal hence Hausdorff (Lemma 4.8, 5.6), and K compact in Hausdorff space implies closed), and so defining $g \in \mathcal{C}(X, \mathbb{R}^N)$ to be the negative of the extension of $f - \tilde{g}$ plus f, we indeed get that $|f - g|_\infty \leq \delta$ (sup-norm on all of X), and $g \in O_\epsilon(K)$.

At the beginning of this section (Section 6.5), we asked that $X = \bigcup_{n=1}^{\infty} K_n$ for compact $K_n \subseteq X$. The Baire category theorem (Theorem 6.1) then gives that $\bigcap_{n=1}^{\infty} O_{1/n}(K_n)$ (the set of injective functions in $\mathcal{C}(X, \mathbb{R}^N)$) is dense in $\mathcal{C}(X, \mathbb{R}^N)$. Now, we just need to find one that satisfies $f(x) \to \infty$ as $x \to \infty$. Well, by denseness, we just need to explicitly find SOME "helper" continuous function h "going to infinity", use denseness to find injective $f \in \mathcal{C}(X, \mathbb{R}^N)$ nearby, and prove that nearby functions of ones that "go to infinity" also "go to infinity".

6.5.2 Existence of Function Going to Infinity

With help from [1]. Our next steps are exactly as I laid out at the end of the preceding paragraph. First, note that when we write $X = \bigcup_{n=1}^{\infty} K_n$ for compact K_n, we can always think of the K_n as increasing, i.e. $K_1 \subseteq K_2 \subseteq K_3 \subseteq \ldots$ and so on, because given $X = \bigcup_{n=1}^{\infty} \tilde{K_n}$ for arbitrary compact $\tilde{K_n}$, we can just define $K_k := \bigcup_{k=1}^{n} \tilde{K_k}$ (a finite union of compact sets is compact, because any open cover of such a union can be refined to a finite number of finite subcovers for each compact set, and the union of these is a finite subcover of the entire union). Now again as I advised in the first paragraph of this section (Section 6.5), visualize the $K_1 \subseteq K_2 \subseteq \ldots$ as nested sets, as bigger and bigger "compact approximations" of X).

With this image in mind, let us try to construct h "going to infinity": if we imagine the boundaries $\partial K_1, \partial K_2, \ldots$ looking like concentric rings, we can define $\tilde{h}$ to be 1 on K_1, and then n on ∂K_n for all $n \in \mathbb{N}$ (so $\tilde{h}$ is defined on $K_1 \cup \partial K_2 \cup \partial K_3 \cup \ldots$). Then, for any $n \in \mathbb{N}$, since ∂K_n is closed, we can use Urysohn's lemma (Section 4.1) to define a continuous function $h_n : K_{n+1} \setminus \mathrm{int}(K_n) \to [n, n+1]$ that is n on ∂K_n and $(n+1)$ on ∂K_{n+1}. Then, with h_n defined for all $n \in \mathbb{N}$, we can stitch these h_n together into a continuous function $h : X \to [1, \infty)$ (well defined because the domains of h_n intersect along a boundary ∂K_n, and the h_n agree with $\tilde{h}$ on all boundaries). More rigorously, let us prove h is continuous using the $\epsilon - \delta$ definition of continuity (X is metric space, and $[1, \infty)$ is too): for any $x \notin \bigcup_{n=1}^{\infty} \partial K_n$, h agrees with h_n locally at x for some h_n, and because h_n is continuous at x, h is too; for $x \in \bigcup_{n=1}^{\infty} \partial K_n$, say $x \in \partial K_n$, then x is in the domain of h_n and h_{n-1}, and so for any $\epsilon > 0$ we have δ_n, δ_{n-1} corresponding to h_n, h_{n-1} for that ϵ, so we can just take $\delta := \min\{\delta_n, \delta_{n-1}\}$ which suffices to show continuity of h at x.

Just a clarification: here we act like h is an $\mathbb{R}$-valued function, but we can easily make h into an $\mathbb{R}^N$-valued function by taking the first coordinate to be the $\mathbb{R}$ value, and the rest of the coordinates to be 0. Lastly, of course h "goes to infinity", because for any $n \in \mathbb{N}$, $|h(x)| > n$ for all $x \in X \setminus K_n$.

The last thing we need to do here is to ensure that we can think of the K_n as nested in this manner (i.e. we used very crucially that the ∂K_n were disjoint, which is equivalent to requiring that $K_n \subseteq \text{int}(K_{n+1})$, given that we already have $K_n \subseteq K_{n+1}$). Equivalently, we want to find some property/properties that make it so that given any nested sequence $\tilde{K_1} \subseteq \tilde{K_2} \subseteq \dots$ that unions to X, we can create a "strictly" nested sequence $K_1 \subseteq K_2 \subseteq \dots$ (again that unions to X) so that $K_n \subseteq \text{int}(K_{n+1})$ for all $n \in \mathbb{N}$.

Suppose we start with $K_1 := \tilde{K_1}$. How do I make K_2 s.t. $K_1 \subseteq \text{int}(K_2)$? Well, we can find an open set containing K_1 by just letting N_x be a neighborhood around each point $x \in K_1$ and defining $U_1 := \bigcup_{x \in K_1} N_x$. Actually, since K_1 is compact, there's a finite subcover $N_1, \dots, N_{n_1} \in \{N_x\}_{x \in K_1}$ that still covers K, so we can define $U_1 := \bigcup_{i=1}^{n_1} N_i$. But is it the case that $\overline{U_1}$ is compact? Note that $\overline{U_1} = \bigcup_{i=1}^{n_1} \overline{N_i}$ (we proved this for any locally finite collection of sets in the proof of Lemma 4.5, the "refinement lemma" from the section on partitions of unity; of course a finite collection is locally finite), and recall from a couple of paragraphs ago that a finite union of compact sets is still compact. That is to say, if for every $x \in X$, there was a neighborhood N_x containing x s.t. $\overline{N_x}$ was compact, then we would be able to construct such a "strictly" nested sequence: simply take $K_1 := \tilde{K_1}$, and inductively define U_n to be a finite union of neighborhoods (the special neighborhoods s.t. their closures are compact) covering the compact set $K_n \cup \tilde{K_n}$ (which we can do because $\{N_x\}_{x \in K_n \cup \tilde{K_n}}$, again for the special neighborhoods N_x, forms an open cover of $K_n \cup \tilde{K_n}$, which we can then refine by compactness to a finite subcover), and then $K_{n+1} := \overline{U_n}$ (which is equal to the union of the closures of finitely many "special neighborhoods", hence compact). In this manner, we construct $K_n \supseteq \tilde{K_n}$ (which implies that $X = \bigcup_{n=1}^{\infty} K_n$) which is furthermore a "strictly" nested sequence.

The property that every point $x \in X$ has a neighborhood N_x s.t. $\overline{N_x}$ is compact (called a "compact neighborhood") is referred to as **locally compact**. Thus in addition to all the properties of X requested in the first paragraph of Section 6.5, we ask that X is also locally compact. Do note that there are many (inequivalent) definitions of this term in different references; I take definition number 1 from the Wikipedia page for local compactness.

6.5.3 Close Functions Also Go to Infinity

As I said in the last paragraph of Section 6.5.1 above, denseness of the set of injective functions in $\mathcal{C}(X, \mathbb{R}^N)$ gives that there is some injective $f \in \mathcal{C}(X, \mathbb{R}^N)$ s.t. $|f - h|_\infty < 1 \iff$ for all $x \in X$, $|f(x) - h(x)| < 1$. In particular, we have $|f(x) - h(x)| < 1 \implies \big||f(x)| - |h(x)|\big| < 1$ by reverse triangle inequality, so $|f(x)| - |h(x)| > -1 \implies |f(x)| > |h(x)| - 1$ for all $x \in X$. Applying this to $x \in X \setminus K_{n+1}$, we see that for every $n \in \mathbb{N}$, $|f(x)| > (n+1) - 1 = n$ for all $x \in X \setminus K_{n+1}$, where the existence of a compact set with this property is exactly the definition of f "going to infinity". Thus, we have at last shown the existence of an injective $f \in \mathcal{C}(X, \mathbb{R}^N)$ that "goes to infinity", where such an f is in particular a homeomorphism between X and a subset of $\mathbb{R}^N$ (see the material from Section 6.5 prior to Section 6.5.1).

In summary, we have proven the following theorem:

Theorem 6.8: Embedding theorem

If X is a topological space that is metrizable with metric d (hence normal for Urysohn, by Lemma 4.8) that is also σ-compact and locally compact, such that every compact subset K has topological dimension $\leq m$ (and defining $N := 2m + 1$), then we can embed X into $\mathbb{R}^N$, i.e. there is a continuous injective map $f : X \hookrightarrow \mathbb{R}^N$ that is furthermore a homeomorphism between X and $\operatorname{im} f \subseteq \mathbb{R}^N$.

6.6 Adding in Locally Euclidean

To use the above theorems on locally m-Euclidean spaces (which is a property we wanted all manifolds in $\mathbb{R}^N$ to have, see Section 3.1), we first need to establish a link between locally m-Euclidean and topological dimension. In Section 6.5 ("Beyond Compactness"), we requested that every compact subspace of X have topological dimension $\leq m$, and our first result from Section 6 was about compact spaces X, so indeed we only need to prove that compact subsets of locally m-Euclidean spaces have topological dimension $\leq m$. We already know from Lemma 6.6 that a finite union of spaces of topological dimension $\leq m$ also has topological dimension $\leq m$, and we know from Lemma 6.3 that compact subsets of $\mathbb{R}^m$ have topological dimension $\leq m$.

As topological dimension is a topological property (if X_1, X_2 are homeomorphic via a homeomorphism h, open coverings of one can be pushed forward/pulled back via h to get an open cover, and the overlap order at $x \in X_1$ and $f(x) \in X_2$ are equal), we only need to prove that any compact subset of locally m-Euclidean space X can be written as finite union of subsets homeomorphic to the CLOSED unit ball in $\mathbb{R}^m$. The reason why I emphasized "CLOSED" is because we already know that any compact subset K of locally m-Euclidean space X is a finite union of subsets homeomorphic to the OPEN unit ball (i.e. homeomorphic to $\mathbb{R}^m$) because locally m-Euclidean means that for every $x \in X$ there is neighborhood $U_x \subseteq X$ s.t. $U_x \simeq \mathbb{R}^m \simeq B(0,1)$ are all homeorphic, and $\{U_x\}_{x \in X}$ is an open cover of K, so there is a finite subcover.

Well [3], suppose we have a compact subspace $K \subseteq X$, covered by $\{U_i\}_{i=1}^n$, a cover of open sets in X, where each $U_i \simeq B(0,1) \subseteq \mathbb{R}^m$ via the homeomorphisms $\{h_i\}_{i=1}^n : U_i \to B(0,1)$. Defining $A_1 := K \setminus (U_2 \cup \ldots \cup U_n)$, we see that A_1 is a closed, and so $h(A_1)$ is a closed subset of $B(0,1)$ (because h is a homeomorphism $\iff$ h is an open map $\iff$ h is a closed map). By the Heine-Borel theorem (Theorem 5.2), $h(A_1)$ is compact, and so there must be some $r \in (0,1)$ s.t. $h(A_1) \subseteq B(0,r)$ (because otherwise, $\{B(0, 1 - \frac{1}{n})\}_{n=2}^{\infty}$ would be an open cover with no finite subcover — this is a special case of the more general fact that a a compact set and disjoint closed set must be distant). Defining $V_1 := h^{-1}(B(0,r))$ and $B_1 := h^{-1}(\overline{B(0,r)})$, we see that V_1 covers $K \setminus (U_2 \cup \ldots \cup U_n)$, so $\{V_1, U_2, \ldots, U_n\}$ is an open cover of K, and B_1 is a closed set containing V_1 homeomorphic to the compact set $\overline{B(0,r)} \subseteq \mathbb{R}^m$).

We can then continue, by defining $A_2 := K \setminus (V_1 \cup U_3 \cup \ldots \cup U_n)$, and so on until we have $\{V_i\}_{i=1}^n$ a (refined) open cover of K, and $\{B_i\}_{i=1}^n$ a closed cover of K with all the B_i homeomorphic to some compact set in $\mathbb{R}^m$.

6.6.1 If and Only If

Looking again at Theorem 6.8, we see that the assumption that X is locally m-Euclidean takes care of the request that X satisfy that every compact subset has topological dimension $\leq m$ (the above material in this subsection, Section 6.6), and that X is locally compact, since for any point $x \in X$, we have a homeomorphism $h : B(0,1) \to X$, and the image $h(\overline{B(0, \frac{1}{2})})$ of a compact set via a continuous function is compact, where $h(\overline{B(0, \frac{1}{2})})$ (and hence its interior) contains the open set $h(B(0, \frac{1}{2}))$ because h is an open map.

The only remaining requests from Theorem 6.8 are that X is σ-compact and metrizable. Recall from Theorem 4.7 that a second countable T_4 topological space X is metrizable (where metrizable also implies T_4), and note that second countability on top of local compactness gives σ-compactness because second countability implies separable (see Section 2.3) with say countable dense set $Q := \{x_i\}_{i=1}^{\infty}$, and for the "special neighborhoods" N_{x_i} s.t. $\overline{N_{x_i}}$ is compact, we see that $X = \bigcup_{i=1}^{\infty} \overline{N_{x_i}}$ is indeed the union of countably many compact subsets.

Thus, assuming X is locally m-Euclidean, we see that X being second countable and T_4 implies that X embeds into $\mathbb{R}^N$ (for some N; if we want to be specific, we can choose $N := 2m+1$). And in fact, the converse is ALSO TRUE! Assuming X is locally m-Euclidean and X embeds into $\mathbb{R}^N$, we see that X is metrizable (using the metric on $\mathbb{R}^N$ just restricted to the subset), hence T_4 (by Lemma 4.8), and also second countable (because $\mathbb{R}^N$ is; just take balls with rational radii centered at points with rational coordinates)! Thus, we have characterized all locally m-Euclidean topological spaces that embed into Euclidean space, successfully generalizing our preliminary definition of manifold from Section 3.1! We now can make the definition:

Goal: Definition of Topological Manifold

A topological space X that satisfies the following three properties is called a **topological m-manifold**:

- X is locally m-Euclidean: at every $x \in X$, there is a neighborhood N_x containing x s.t. N_x is homeomorphic to $\mathbb{R}^m$ (or equivalently the unit ball $B(0,1) \subseteq \mathbb{R}^m$, since $B(0,1) \simeq \mathbb{R}^m$)
- X is second countable: there is a countable basis $\{B_n\}_{n=1}^{\infty}$ (basis is defined in Section 2.3).
- X is T_4: X is T_1 ($\iff$ every point $x \in X$ is closed) and normal (defined in Section 4.1).

Remarks: the standard definition replaces the third bullet with "X is T_2/Hausdorff", but it is true that any locally Euclidean second countable Hausdorff space is T_4, so the definitions are equivalent.

The main storyline for this paper is COMPLETELY FINISHED (and my goodness, what an amount of work it took!), but I do have some treats left over (essentially a "meditation on Brouwer's fixed point theorem") for those who still have the spirit in them to continue!

Challenge question for the reader. The proof should be actually pretty short! Much easier than any of the proofs in this section.

Question 6.1: Exercise

Suppose X is a COMPACT topological space that admits a partition of unity subordinate to any open cover (so say X is normal and paracompact) AND is locally k-Euclidean. Prove that X embeds into $\mathbb{R}^N$ for some $N \in \mathbb{N}$. See Theorem 4.86 in Lee [6] for a solution.

7 Bonus: Sperner's Lemma & Invariance of Dimension

Now for some fun! I do warn the reader that the proofs in these bonus sections are particularly clever, where I mean that in a somewhat derogatory way. In the course of this article, I have tried very hard to prove things in the "most natural way possible", by discovering concepts as we go along instead of as an info-dump in the exposition. I personally think I've done a pretty good job laying out the heart of the arguments and showing that they really are the "first idea someone could try"; at least I would be happy if someone had taught me topology the way I have explained it here (and I am always unhappy when I feel that a concept/proof is too much of a "miracle" or a "black box"). Unfortunately or fortunately depending on how much you like magic, that all ends here — the following proofs are all very elegant and beautiful, but also definitely not the "first idea someone could try". They are really still a bit "miraculous" to me, but hey, these are bonuses! We can have some miracles for bonus.

The Sperner lemma part is heavily inspired by Francis Su's excellent paper "Rental Harmony: Sperner's Lemma in Fair Division" [10] and a Youtube video of Mathologer (the Youtube alias of a math professor by the name of Burkard Polster) on the aforementioned paper, `https://www.youtube.com/watch?v=7s-YM-kcKME&ab_channel=Mathologer`.

Lemma 7.1: Sperner's lemma in 2-dimensions

Consider a triangle T (sitting in the plane, $\mathbb{R}^2$), triangulated into many smaller triangles, whose vertices are labeled by one of $\{1, 2, 3\}$ (or equivalently you can think of each vertex as colored red, blue, or green). If the labeling is such that the three vertices of T are all labeled/colored differently, and the labels of a vertex along any edge of T are one of the two on the endpoints of that edge (such a labeling is called a Sperner labeling/coloring), then the number of smaller sub-triangles with all three vertices labeled differently (one each of $\{1, 2, 3\}$) is odd (hence there's at least one).

The existence of one such smaller sub-triangle is in fact guaranteed by just having one edge of the triangle colored exclusively in two colors, with endpoints not the same color (this fact is obvious from the constructive walk-through-doorways proof seen in Su's paper/Mathologer's video).

Proof: Mathologer's video is easy/fun to watch, pretty short, and much better than anything I could write down, so I'll leave this one to him: `https://www.youtube.com/watch?v=7s-YM-kcKME&ab_channel=Mathologer`. If you're really antsy about not having a proof here, no worries, because I'll prove the n-dimension Sperner's lemma below anyways (this case is just to build up intuition). ■

This generalizes very well to higher dimensions, but we do still need to set the stage a little with some new definitions/notations.

7.1 Preliminary Definitions

We define an n-simplex informally to be the n-dimensional analogue of a triangle (so 0-simplex would be a point, a 1-simplex a line, a 2-simplex a bona-fide triangle, a 3-simplex a tetrahedron, and so on). More rigorously, an n-simplex can be thought of as the convex hull of $n+1$ geometrically independent points in $\mathbb{R}^m$ for $m \geq n$, where the convex hull of a set of points is the smallest convex set containing those points; for a finite set of points, the convex hull is the set of weighted averages of the points (this set is indeed convex, i.e. for every two points $\mathbf{x}, \mathbf{y}$ in the set the line connecting them, $t\mathbf{x} + (1-t)\mathbf{y}$ ($t \in [0,1]$) is contained within the set as well)

This fact can be proven by induction: if the convex hull of n points S_n is the set of weighted averages, then adding in a $(n+1)$th point $\mathbf{p}$, by convexity all $t\mathbf{x} + (1-t)\mathbf{p}$ ($t \in [0,1]$) must be in the convex hull of the $n+1$ points for $x \in S_n$, but all $x \in S_n$ are just weighted averages of the n points, so the set of all $t\mathbf{x}+(1-t)\mathbf{p}$ ($t \in [0,1]$, $\mathbf{x} \in S_n$) are just the set of weighted averages of the $n+1$ points.

A triangulation of an n-simplex S is a collection of smaller n-simplices contained in S (that union to S) s.t. two sub-n-simplices either do not intersect at all (i.e. are disjoint), or have a face common to both (a k-face of S is the convex hull of k out of the $n+1$ points defining S). A facet of S is the convex hull of exactly n out of the $n+1$ points defining S (i.e. a facet of S is exactly an n-face of S); e.g. the facets of a 1-simplex are the endpoints (0-simplices), the facets of a 2-simplex are the edges (1-simplices), and the facets of a 3-simplex are the faces (2-simplices).

One major explicit triangulation is the **barycentric** triangulation: for an n-simplex S that is the convex hull of $n+1$ points, the barycentric subdivision can be defined inductively by adding in the barycenter $\mathbf{b}$ (the unweighted average of all $n+1$ vertices), drawing the $n+1$ lines from each of the $n+1$ vertices to $\mathbf{b}$ and noting that the line intersects the facet opposite the vertex of origin at the barycenter of that facet, and defining the new sub-n-simplices to consist of an $(n-1)$-simplex from the barycentric subdivision of one of the facets plus the $(n+1)$th point $\mathbf{b}$. Of course, one can iterate this, so the first iteration is the first barycentric subdivison $\mathscr{B}_1$, and then one can perform barycentric subdivision on the sub-n-simplices in $\mathscr{B}_1$ to get $\mathscr{B}_2$, and so on. Note that the diameters of successive barycentric subdivisions goes to 0 (in fact I think the diameters more than halve each time).

7.2 Sperner's Lemma in n-dimensions

We are now ready to give the statement of Sperner's lemma in n-dimensions, and to give its proof (again, this is all from Su [10]).

Lemma 7.2: Sperner's lemma in n-dimensions

Consider an n-simplex S, triangulated into many smaller triangles, whose vertices are labeled by some $i \in [n+1]$. If the labeling is such that the $n+1$ vertices of S are all labeled differently, and the label of every point in the facet of S opposite the vertex labeled i is NOT equal to i (such a labeling is called a Sperner labeling/coloring), then the number of smaller sub-triangles with all $n+1$ vertices labeled differently (one each of $i \in [n+1]$) is odd (hence there's at least one).

Proof: from Francis Su's excellent paper [10]. We proceed by induction. The base case $n = 1$ is just the case in which we have a line segment with endpoints labeled 1 and 2, meaning any "triangulation" (i.e. just adding points on the line segment labeled 1 or 2) must have the labeling switch an odd number of times as we go from the endpoint labeled 1 to the endpoint labeled 2 (switching an even number of times means that the two endpoints would have the same label).

Now assume (the induction hypothesis) that the claim holds for any $(n-1)$-simplex (have in your mind $n = 3$, where you have intuition about Sperner's lemma being true for 2-simplices). Given an n-simplex S, and a triangulation of S with a Sperner labeling, think of each sub-n-simplex as a "room", and each facet as either a "door" or a "wall", where we specify that a facet is considered a "door" exactly when the n vertices of that facet are labeled $1, 2, \ldots, n$. Note that on facets of S, "doors" can only appear on the facet opposite the vertex labeled $n+1$ (because all other facets are opposite some other vertex labeled $i \in [n]$, hence the n vertices in those facets can not possibly be labeled $1, 2, \ldots, n$. Note furthermore that a room can have 0, 1, or 2 doors, where there is 1 door if and only if that room has all $n+1$ labels on its vertices (and 2 doors if and only if exactly two vertices of that room have the same label $i \in [n]$).

Thus, the statement of Sperner's lemma is equivalent to saying that there are an odd number of rooms with 1 door. We prove this by "taking a tour" of the rooms of the n-simplex S. By the induction hypothesis (the facet opposite the vertex labeled $n+1$, which we hereby refer to as the "$(n+1)$th facet", is an $(n-1)$-simplex with Sperner labeling), there are an odd number of doors "entering" S from the outside. Choosing any such door, let us walk through the rooms, closing the doors behind us as we walk. We can not loop back on a room we previously visited (since by walking in and walking out, we closed both doors), and because the number of rooms is finite, the walk must end. Well, how does such a walk end? Either we end up locking ourselves into a room with 1 door, or we walk back outside (no other possibilities because we can never enter room with 0 doors, and entering a room with 2 doors we can always exit out the other door). But again by the induction hypothesis, there are an odd number of doors entering S from the outside, meaning that every time we end up back outside (each "exit door" of S is paired with exactly one "entrance door"), there will still be doors open leading back into S.

Thus, starting on the outside of S and touring the rooms like this, we are GUARANTEED to eventually lock ourselves in a room with exactly 1 door (thereby showing the existence of at least one sub-n-simplex with all $n+1$ labels). Finally, the number of such rooms is odd as follows: because each room that is reachable from the outside corresponds to one door on the ($n+1$th facet (let's call such doors "trap doors" and such rooms with 1 door "trap rooms" because going in one we get trapped in a room), and the number of non-trap doors is even (because they can be paired into "entrance/exit" doors), the number of trap rooms reachable from the outside is odd; and of course the number of trap rooms not-reachable from the outside must be paired, because spawning in one such trap room, we can keep walking through doors until we end up in another trap room. ■

7.3 Invariance of Dimension

This subsection is heavily inspired from [11]. With Sperner's lemma, we are ready to give the proofs of two major theorems, the first of which establishes that open sets of $\mathbb{R}^m$ and $\mathbb{R}^n$ (for $m \neq n$) can not be homeomorphic. This in particular shows that if X is locally m-Euclidean, then it can not be locally n-Euclidean for $n \neq m$. We do this via a (also very important!) lemma telling us that the topological dimension of an n-simplex is $\geq n$ (i.e. the other direction of the inequality we established in Lemma 6.3).

Lemma 7.3: Topological dimension of n-simplex

The topological dimension of an n-simplex S is $\geq n$, and because we have already proven the reverse inequality in Lemma 6.3, the topological dimension of an n-simplex is indeed EXACTLY n.

Proof: We will prove that any finite closed cover of diameter $< \epsilon$ (for some $\epsilon > 0$) must have overlap order $\geq n+1$, which suffices because any open cover can be refined to a finite cover (by compactness), and Lemma 4.5 says that for any paracompact (implied by compact) and T_4 space (recall from Lemma 4.8 that metric spaces, like $\mathbb{R}^n$, are T_4), we can refine any open cover to a closed cover, and so if all open covers of diameter $< \epsilon$ can be refined to a finite closed cover of diameter $< \epsilon$, and if such closed covers must have $\geq n+1$ overlap order, so must the original open cover/any refinement of it.

Let $F_1, \ldots, F_{n+1}$ denote the facets of S (one corresponding to each missing vertex $v_1, \ldots, v_{n+1}$ of S). Because each F_i is closed, $S_n \setminus F_i$ is open in the subspace topology of S for each $i \in [n+1]$, so $\{S \setminus F_i\}_{i=1}^{n+1}$ is an open cover of S. As S is closed and bounded in $\mathbb{R}^m$ ($m \geq n$), it is compact by Heine-Borel (Theorem 5.2), so by the Lebesgue number lemma (Lemma 6.2), there is a Lebesgue number $\epsilon > 0$ for this cover. Let $\mathscr{C}$ be a finite closed cover of S by (closed) sets of diameter $< \epsilon$. As per the above paragraph, it suffices to show that $\mathscr{C}$ has overlap order $\geq n+1$.

Since $\text{diam}(C) < \epsilon$ for all (finitely many) $C \in \mathscr{C}$, we have that $C \subseteq B(x, \epsilon)$ for any $x \in C$, so by the definition of $\epsilon > 0$ as the Lebesgue number of the cover, there is some $i \in [n+1]$ s.t. $C \subseteq S \setminus F_i$. We can then define $\phi : \mathscr{C} \to [n+1]$ mapping each $C \in \mathscr{C}$ to a corresponding $i \in [n+1]$ s.t. $C \subseteq S \setminus F_i$, implying in particular that $C \cap F_i = \varnothing$ for all $C \in \mathscr{C}$ s.t. $\phi(C) = i$.

Defining $A_i := \bigcup_{\phi^{-1}(i)} C$, we see that $A_i \cap F_i = \varnothing$, and $\bigcup_{i=1}^{n+1} A_i = \bigcup_{C \in \mathscr{C}} C = S$, so in fact the vertex $v_i \in A_i$ (because no other A_j contains v_i, because $v_i \in F_j$ for $j \neq i$). We now label each $x \in S$ by $L(x) = \min\{i \in [n+1] : x \in A_i\}$. Because only A_i contains v_i for each vertex, $i \in [n+1]$, we see that $L(v_i) = i$; moreover on each face F_i, $x \in F_i$ is not contained in A_i, so indeed no point of F_i is labeled with the label i. Thus, for any triangulation of S, $L(x)$ is a Sperner labeling, and so we can apply Sperner's lemma. As for which triangulation we pick, we just go with the barycentric subdivisions $\mathscr{B}_1, \mathscr{B}_2, \ldots$ so that there is a simplex S_s for each $\mathscr{B}_s$ that has all $n+1$ labels on its vertices. We pick some $x_s \in S_s$, and form a sequence $\{x_s\}_{s=1}^{\infty}$.

Although this sequence does not necessarily converges, we have infinitely many points inside a bounded set S, so intuitively there must be some clustering going on, i.e. there must be some point $\mathbf{p}$ s.t. every neighborhood has infinite many elements from the sequence. If not, we would find a neighborhood of every point with only finitely many elements from the sequence, which together would form an open cover of the compact set S, meaning there is finite subcover of neighborhoods each with finitely many elements from the sequence, yielding only finitely many elements of the sequence total, which is a contradiction. In a metric space, having infinitely many points in every neighborhood of $\mathbf{p}$ is equivalent to having a subsequence of the original sequence $\{x_s\}_{s=1}^{\infty}$ converge to $\mathbf{p}$ (for every $k \in \mathbb{N}$, choose $s_k \in \mathbb{N}$ (larger than previously chosen s_{k-1}) s.t. $x_{s_k} \in B(\mathbf{p}, \frac{1}{k})$; then $\{x_{s_k}\}_{k=1}^{\infty} = \{x_s\}_{s \in S}$ for $S := \{s_1, s_2, \ldots\}$ is a subsequence of $\{x_s\}_{s=1}^{\infty}$ converging to $\mathbf{p}$). This is essentially the **Bolzano-Weierstrass theorem**.

I now claim that $\mathbf{p}$ is in ALL the A_i, $i \in [n+1]$. This is because $x_s \in S_s$, where S_s is a sub-n-simplex from $\mathscr{B}_s$ that has a vertex w_s labelled i, hence in A_i. Because the diameter of the S_s to to 0 as $s \to \infty$, we have that $|w_s - x_s| \to 0$ as $s \to \infty$, so indeed if the subsequence $\{x_s\}_{s \in Q}$ for some infinite index set $Q \subseteq \mathbb{N}$ converges to $\mathbf{p}$ as $s \in Q$ goes to ∞, then $\{w_s\}_{s \in Q}$ must also converge to $\mathbf{p}$ as $s \in Q$ goes to ∞. Because A_i is the finite union of closed sets, it itself is closed, so the limit point $\mathbf{p} \in A_i$. Because $i \in [n+1]$ was arbitrary, we have indeed proven that $\mathbf{p} \in A_i$ for all $i \in [n+1]$.

By construction of the A_i, there must be $C_i \in \phi^{-1}(i)$ s.t. $\mathbf{p} \in C_i$ for all $i \in [n+1]$. Because the $\{\phi^{-1}(i)\}_{i=1}^{n+1}$ is a disjoint partition of $\mathscr{C}$, all the C_i are distinct, and hence we have shown that $\mathbf{p}$ is in $(n+1)$ sets of $\mathscr{C}$. ■

Theorem 7.4: Invariance of dimension

If $U \subseteq \mathbb{R}^m$ and $V \subseteq \mathbb{R}^n$ are open and homeomorphic, then it must be that $m = n$

Proof: suppose we have open $U \subseteq \mathbb{R}^m$ and $V \subseteq \mathbb{R}^n$ and homeomorphism $h : U \to V$. Since $U \subseteq \mathbb{R}^m$ is open, it contains an open ball, so it contains an m-simplex S_m (just take the convex hull of $m + 1$ points in the ball, which will still lie in the ball because the balls is convex). Heine-Borel (Theorem 5.2 again tells us that S_m is compact, and since h is a continuous function, the image $h(S_m) \subseteq V \subseteq \mathbb{R}^n$ is compact (Theorem 5.7). Because $h(S_m)$ is in particular bounded, it is contained in $B(0, R)$ for some large enough $R \in \mathbb{N}$, and we can find a large enough n-simplex S_n that contains $B(0, R)$. We know from Lemma 6.5 that because $h(S_m)$ is a closed subspace of S_n (which has topological dimension n per the above lemma, Lemma 7.3), we must have that $\dim h(S_m) \leq \dim S_n = n$. But h is a homemorphism and topological dimension is a topological property (see the second paragraph of Section 6.6), we have that $m = \dim S_m = \dim h(S_m) \leq n$. Switching U and V and considering h^{-1} (also a homeomorphism and in particular continuous function), we get $n \leq m$ so indeed $n = m$. ■

We also summarize the Bolzano-Weierstrass theorem here, since we use in the next subsection:

Theorem 7.5: Bolzano-Weierstrass

In a metric space, a compact set has the property that any sequence has a convergent subsequence. This latter property is called **sequential compactness**.

7.4 Brouwer's Fixed-Point Theorem

This subsection is heavily inspired from [12]. Our second major theorem that can be proven by Sperner's lemma is the following:

Theorem 7.6: Brouwer's fixed-point theorem (BFPT)

For any Euclidean space $\mathbb{R}^n$, denoting by $\overline{B(0,1)}$ the closed unit disk in $\mathbb{R}^n$, it is true that any continuous self-map of the disk $f : \overline{B(0,1)} \to \overline{B(0,1)}$ must have a fixed point (i.e. there is $x \in D$ s.t. $f(x) = x$).

Note that this theorem also holds for any domain D homeomorphic to $\overline{B(0,1)}$ (say via a homeomorphism $h : \overline{B(0,1)} \to D$) because for any continuous $f : D \to D$, the function $h^{-1} \circ f \circ h : D \to D$ is a composition of continuous functions, hence continuous, and thus must have a fixed point; i.e. $h^{-1}(f(h(x)) = x \iff f(h(x)) = h(x)$, giving us that $h(x)$ is a fixed point of f.

Proof: because the theorem holds for any space homeomorphic to the closed disk, we prove it for an n-simplex S, using Sperner's lemma. First, let us introduce *barycentric coordinates*: recall from Section 7.1 ("Preliminary Definitions") that every point in the convex hull of $n + 1$ points (i.e. every point in the n-simplex) is a weighted average of those $n+1$ points, which means that every point in the n-simplex can be uniquely identified by the $n+1$ weights (each in $[0,1]$) in the weighted average of the $n+1$ points (where we pick an ordering of the $n+1$ points/vertices of the n-simplex, i.e. the weight in the ith coordinate corresponds to the weight of the ith vertex).

To use Sperner's lemma, we must first define a Sperner labeling of the points of S: supposing we have $n+1$ weights $\mathbf{w} := (w_1, \ldots, w_{n+1})$, and $f : S \to S$ maps $\mathbf{w}' := (w_1, \ldots, w_{n+1}) \mapsto (w_1', \ldots, w_{n+1}')$, we define the label $L(\mathbf{w})$ to be the maximum $j \in [n+1]$ s.t. $w_i' \geq w_i$ for all $i \in \{1, \ldots, j-1\}$ and $w_j' < w_j$. Informally, this labeling labels a point $\mathbf{w}$ by the number of a vertex of the n-simplex that f maps $\mathbf{w}$ farther away from (this is not entirely accurate, since it is possible to get farther away from vertex $\mathbf{v}_i$ without changing w_i, but it is just an intuition). Observe that it is possible that $w_i' \geq w_i$ for all $i \in [n+1]$ (and hence would get no label), but because all the weights must sum to 1, this happening would imply that $w_i' = w_i$ for all $i \in [n+1]$, i.e. $\mathbf{w}$ is a fixed point of f. So, we suppose by contradiction that there is no fixed point of $f : S \to S$, and then our labeling L is defined for all points of S.

We now check that L is a Sperner labeling. If $\mathbf{w}$ is a vertex $\mathbf{v}_i$ of S (the ith vertex in the previously fixed ordering of the vertices), the barycentric coordinates of $\mathbf{w}$ is just 0 everywhere except a 1 at the ith coordinate. In this case, since f does not fix $\mathbf{w}$, all other coordinates, i.e. w_j' for $j \neq i$ will be $\geq w_j$, and w_i' will be strictly $< w_i$, meaning $L(\mathbf{w}) := L(\mathbf{v}_i) = i$. And any point $\mathbf{w}$ on the facet opposite the vertex $\mathbf{v}_i$ will have ith coordinate equal to 0, so it must be that $w_i' \geq w_i = 0$, meaning that $L(\mathbf{w})$ will never be i (since $L(\mathbf{w}) = i$ implies in particular that $w_i' < w_i$).

Now, we choose a sequence of triangulations of S s.t. their diameters go to 0 (say the barycentric divisions $\{\mathscr{B}_s\}_{s=1}^{\infty}$ from Section 7.1), and apply Sperner's lemma to each. Well, we get a sub-n-simplex for each $s \in \mathbb{N}$, say with $n+1$ vertices $\{(w_{s,i,1}, \ldots, w_{s,i,n+1})\}_{i=1}^{n+1}$ where $(w_{s,i,1}, \ldots, w_{s,i,n+1})$ is labeled with the label i. Since S is compact, Bolzano-Weierstrass (Theorem 7.5) tells us that there is a subsequence (say with infinite indexing set $\tilde{S} \subseteq \mathbb{N}$) of $\{(w_{s,1,1}, \ldots, w_{s,1,n+1})\}_{s=1}^{\infty}$ that converges to say $\mathbf{x} := (x_1, \ldots, x_{n+1})$. Because the diameters of the sub-n-simplices go to 0 as $s \to \infty$, for all fixed $i \in [n+1]$ the subsequence of the ith vertex of the sub-n-simplices, $\{(w_{s,i,1}, \ldots, w_{s,i,n+1})\}_{s \in \tilde{S}}$, converges to $\mathbf{x}$.

Because the subsequence of 1st vertices of the sub-n-simplices $\{(w_{s,1,1}, \ldots, w_{s,1,n+1})\}_{s \in \tilde{S}}$ are all labeled 1 by L, we have that $w_{s,1,1}' < w_{s,1,1}$ for all $s \in \tilde{S}$. Because $w_{s,1,1} \to x_1$ as $s \to \infty$, and $w_{s,1,1}' \to x_1'$ as $s \to \infty$ (because f is continuous and hence convergent sequences get mapped to convergent sequences), we have that $x_1' \leq x_1$. Doing this for all other ith vertices for $i \in [n+1]$, we get that $x_i' \leq x_i$ for all $i \in [n+1]$, and so contradiction, we found a fixed point $\mathbf{x}$. ■

8 Bonus Bonus: Jordan Curve Theorem

This section is heavily inspired by [5], which itself was heavily inspired by [7]. The Jordan curve theorem is perhaps the most famous example of a statement that feels completely obvious but is inordinately difficult to prove. The following MSE threads may be of interest: What seemingly innocuous results in mathematics require advanced proofs?, "It looks straightforward, but actually it isn't", and Why is the Jordan Curve Theorem not "obvious"?. The statement is as follows:

Theorem 8.1: Jordan curve theorem (JCT)

Informally: any curve in $\mathbb{R}^2$ that "looks like" a circle (where "looks like" can be formalized as "homeomorphic to") has an inside and an outside.
Formally: for any Jordan curve $J \subseteq \mathbb{R}^2$, i.e. the image of an embedding $S^1 \hookrightarrow \mathbb{R}^2$, its complement $\mathbb{R}^2 \setminus J$ consists of two disjoint open, path-connected components, where one is bounded (the inside) and the other is unbounded (the outside), where moreover the boundary of each component is J itself.

Before we begin, let us prove another "obvious" fact that turns out to be extremely useful in our proof:

Lemma 8.2: Horizontal and vertical paths must intersect in a rectangle

We use the notation $\mathbb{I} := [-1, 1]$. Then, any continuous functions ("paths") $h = (h_1, h_2), v = (v_1, v_2) : \mathbb{I} \to \mathbb{I}^2$ s.t. $h_1(-1) = -1, h_1(1) = 1$ and $v_2(-1) = -1, v_2(1) = 1$ (i.e. h starts at the left edge of the square and ends at the right edge, and v starts at the bottom edge of the square and ends at the top edge) must intersect, i.e. there are $s, t \in \mathbb{I}$ s.t. $h(s) = v(t)$. Via translations and vertical/horizontal stretches, this result holds in any rectangle.

Proof: let $|\cdot|_\infty$ denote the ∞-norm (max of the differences in each individual coordinate), and let S^1_∞ denote the unit circle in this norm, i.e. the boundary of the square $\mathbb{I}^2$. The key ingredient that we use will be Brouwer's fixed point theorem. Suppose for sake of contradiction that $h(s) \neq v(t)$ for all $s, t \in \mathbb{I}$, and consider the mapping $f : \mathbb{I}^2 \to S^1_\infty \subseteq \mathbb{I}^2$ defined by $f(s,t) = \frac{h(s)-v(t)}{|h(s)-v(t)|_\infty}$, where the denominator is never 0 because $h(s) \neq v(t)$ for all $s, t \in \mathbb{I}$.

Brouwer's fixed point theorem tells us that there exist $s_0, t_0 \in \mathbb{I}$ s.t. $f(s_0, t_0) = (s_0, t_0)$. But because the image of f is in S^1_∞, we must have $s_0, t_0 = \pm 1$. If $t_0 = 1$, then $f_2(s_0, 1) = \frac{h_2(s_0)-v_2(1)}{|h(s_0)-v(1)|_\infty}$ is ≤ 0 because $h_2(s_0) \in [-1, 1]$ and $v_2(1) = 1$, and hence t_0 can not possibly equal 1. Similarly, if $t_0 = -1$, then $f_2(s_0, -1) \geq 0$, and so again $t_0 \neq -1$. Thus, in either case we result in contradiction, and so our initial assumption must have been wrong; it must be that $h(s) = v(t)$ for some $s, t \in \mathbb{I}$. ■

8.1 Proof of JCT

We begin with the assertion that every component of $\mathbb{R}^2 \setminus J$ is open and path-connected. Before we begin however, we have to figure out what I mean exactly by "component".

8.1.1 Connectedness

Think for example of J equal to the unit circle; $\mathbb{R}^2 \setminus J$ consists of $\{\mathbf{x} \in \mathbb{R}^2 : |\mathbf{x}| < 1\} \cup \{\mathbf{x} \in \mathbb{R}^2 : |\mathbf{x}| > 1\}$. Clearly, one would not say that $\mathbb{R}^2 \setminus J$ is one component, because it is the union of two disjoint open sets. Why do we stress open? Because any set S can be written as the union of two disjoint sets; namely if $A \subseteq S$, then A and $S \setminus A$ are disjoint and union to S. In essence, being a union of disjoint open sets means there's some "separation" or "disconnection" between the open sets. Thus, we might say that if S is a set in some topological space X, calling it connected means that there do not exist (non-empty) disjoint open subsets $U_1, U_2 \subseteq S$ s.t. $S = U_1 \cup U_2$. Unfortunately, we do know that the union of open sets is open, so our previous definition would only allow open sets to be disconnected; no worries, we can easily edit the definition as follows:

Definition 8.3: Connected subset and space

Calling a subset S of a topological space X means that there do not exist disjoint open sets in the subset topology $(S \cap U_1)$ and $(S \cap U_2)$ that union to S (we call such disjoint open sets a *disconnection* or *separation* of S). Of course, this is equivalent to defining "connectedness" for a topological space X to mean that there are no disjoint open sets U_1, U_2 in the topology on X s.t. $X = U_1 \cup U_2$, and then applying this definition to S and its subspace topology.

Another way of phrasing it is that X is disconnected $\iff$ there is an open set U_1 s.t. its complement $X \setminus U_1 =: U_2$ is also open, which leads us to the equivalent definition of: a space X is connected $\iff$ the only sets that are both open and closed ("clopen") are $\varnothing$ and X itself.

Back on track: usually, connectedness is what people mean when they say "component". However, for our problem we will take "components" to be path-connected components (since the theorem statement only concerns path-connectedness). Of course, it is possible to do it with only thinking of "components" as connected components, but we would have a prove a variety of small theorems, such as every interval in $\mathbb{R}$ is connected, path connected implies connected, and for open sets in $\mathbb{R}^n$, connected $\iff$ path-connected.

8.1.2 Components are Open

Since we are taking "components" to mean path-connected components, we just have to prove that such components are open. Suppose C is a component of $J^\complement = \mathbb{R}^2 \setminus J$ (we know we can decompose $J^\complement$ into a union of disjoint path-connected components by considering the equivalence classes of the equivalence relation $x \sim y \iff$ there is path between x and y). Because $J^\complement$ is open (J is compact because it is the image via continuous function of the compact = closed and bounded set $S^1 \subseteq \mathbb{R}^2$), for any $x \in C \subseteq J^{\complement}$, we have some ball $B_x \subseteq J^{\complement}$. We would like to prove that $B_x \subseteq C$.

Suppose not; then there is $y \in B_x$ s.t. $y \in \mathbb{R}^2 \setminus C$. Because we said $B_x \subseteq J^{\complement}$, $y \in J^{\complement} \setminus C$, and so it must in some other component $D \neq C$ of $J^{\complement}$. But because B_x is path-connected (it is convex in $\mathbb{R}^2$ so the straight line path $x(1-\lambda) + y\lambda$, $\lambda \in [0, 1]$ is a path between x and y), we have a path between $x \in C$ and $y \in D$, contradicting that D is a different component from C.

8.1.3 Boundary of Components is J

The following is a lemma that will be used once at the very end of the the proof of JCT (it is also the last assertion in the statement of the JCT, Theorem 8.1, above). I could postpone this lemma and its proof and insert it at the end, but I felt it disrupted the flow too much. Also, I find it interesting how this fact (phrased as a "moreover/trivia/tangential"-type fact in Theorem 8.1 above) actually plays a crucial role in proving the theorem (in particular, proving that there are only two components).

Lemma 8.4: Boundary of components is J

Supposing we have a Jordan curve J s.t. $\mathbb{R}^2 \setminus J$ has ≥ 2 components (exactly one is unbounded and the rest are bounded; this is because J is bounded $\implies J \subseteq B(0, Q)$ for Q large enough, so there is one component O containing $\mathbb{R}^2 \setminus B(0, Q)$ (containing one point of $\mathbb{R}^2 \setminus B(0, Q)$ means it contains all of them by path connectedness!), and any other component must be $\subseteq B(0, Q)$ since otherwise it would intersect with O, contradicting that they are distinct components), then each component has boundary exactly equal to J.

Proof: let U be an arbitrary component of $\mathbb{R}^2 \setminus J$. Then, because U is open, we have that $\partial U = \overline{U} \cap (\mathbb{R}^2 \setminus U)$, where $\mathbb{R}^2 \setminus U$ equals J union all the other components $\mathscr{C}$. But note that no other component W can intersect $\overline{U}$, because if it did, say $x \in W \cap \overline{U}$, W being open implies that there is a neighborhood $N_x \subseteq W$ containing x, but any neighborhood of an element of $\overline{U}$ contains elements of U, meaning of course that $W \cap U \neq \varnothing$, contradicting that W, U are different components. Thus, $\partial U = (\overline{U} \cap J) \cup (\bigcup_{W \in \mathscr{C}} \overline{U} \cap W) = (\overline{U} \cap J) \cup \varnothing = \overline{U} \cap J$ must be contained within J. Suppose now that this containment is strict; i.e. $\partial U \subsetneq J$.

Our technique is as follows: somehow, using the details of the situation, we define a function $f : D \to D \setminus \{\mathbf{p}\}$ for a large closed disk D (radius R) and point $\mathbf{p}$ in the interior of D. Then, we can define $\pi : D \setminus \{\mathbf{p}\} \to \partial D$ by projecting points in $D \setminus \{\mathbf{p}\}$ radially outwards from the point $\mathbf{p}$ onto the boundary circle ∂D (π is continuous by $\epsilon - \delta$ argument: for $\epsilon > 0$ and $\mathbf{x}_0 \in D \setminus \{\mathbf{p}\}$, a simple Euclidean geometry argument, i.e. projecting from $\mathbf{p}$ a small circle around $\mathbf{x}_0$ to a larger circle around $g(\mathbf{x}_0)$ and using similar triangles, shows that $\delta := \frac{\epsilon |\mathbf{x}_0 - \mathbf{p}|}{2R}$ suffices). Finally, we define $\rho : \partial D \to \partial D$ to be a rotation of the circle around its center by $180°$. With these definitions, we see that the function $h \circ g \circ f : D \to \partial D \subseteq D$ (a composition of continuous functions) has no fixed points (anything in the interior of D gets mapped to the boundary, and anything on the boundary stays the same until the rotation ρ, at which point it is also mapped away), contradicting Brouwer's fixed point theorem (Theorem 7.6). The only missing part of this proof is the construction of f (note: it has to be the identity on ∂D!).

Back to our assumption $\partial U \subsetneq J$: there must be a closed arc $A \subsetneq J$ containing ∂U, because if $\mathbf{x} \in J$ not in ∂U and $\theta_x \mapsto \mathbf{x}$ by the Jordan curve embedding $h : S^1 \hookrightarrow\!\!\!\!\rightarrow J \subseteq \mathbb{R}^2$, then because ∂U is closed and h is continuous, there is a neighborhood of θ_x, say $(\theta_x - \delta, \theta_x + \delta)$ that maps to outside ∂U, so ∂U is contained in the arc A corresponding to the closed arc $A_{S^1} := S^1 \setminus (\theta_x - \delta, \theta_x + \delta)$ of the circle — note furthermore that A is homeomorphic to the closed arc A_{S^1} (which is of course homemorphic to the interval $[0, 1]$) because the homeomorphism $h : S_1 \hookrightarrow\!\!\!\!\rightarrow J$ ($\epsilon - \delta$ continuous at every point of S^1) restricted to A_{S^1} is an injective and surjective (i.e. bijective) function between A_{S^1} and its image, i.e. A, where this restriction is still an $\epsilon - \delta$ continuous function at every point of A_{S^1}; as the same can be said of h^{-1}, we see that indeed h appropriately restricted is a homeomorphism between A_{S^1} and A.

Because we assumed that there are ≥ 2 components and we know that exactly one is unbounded, there is some component V that is bounded (if the U we chose above was bounded, just take $V := U$). Let us now choose $\mathbf{p} \in V$, and define $D := \overline{B(\mathbf{p}, R)}$ to be a closed disk centered with R large enough to contain $J \subseteq B(\mathbf{p}, R)$. Then, since we argued above that A is homeomorphic to $(0, 1)$, and we know that $\mathbb{R}^2$ is normal (Lemma 4.8) and A is closed, we can use Tietze's extension theorem (Theorem 4.9), which tells us defining $\tilde{r} : A \to A$ to be the identity (really a map $r : A \to [0, 1]$, which we then compose with the homeomorphism $[0, 1] \hookrightarrow\!\!\!\!\rightarrow A$ to get a map $A \to A$), we can extend it to a continuous function $r : D \to A$. In the two cases U is bounded (i.e. $V = U$, $\mathbf{p} \in U$) or U is unbounded (i.e. $V \neq U$, $\mathbf{p} \in V$), we can now define our desired $f : D \to D \setminus \{\mathbf{p}\}$:

$$f(\mathbf{x}) = \begin{cases} r(\mathbf{x}) & \text{if } \mathbf{x} \in \overline{U} \\ \mathbf{x} & \text{if } \mathbf{x} \in \mathbb{R}^2 \setminus U \end{cases} \qquad \text{or} \qquad f(\mathbf{x}) = \begin{cases} x & \text{if } \mathbf{x} \in \overline{U} \\ r(\mathbf{x}) & \text{if } \mathbf{x} \in \mathbb{R}^2 \setminus U \end{cases}$$

where f is well defined because $\overline{U} \cap (\mathbb{R}^2 \setminus U) \subseteq A$, and we specified that r is the identity on A. Also we should note that the places we ask for the value of r is within the domain D of r, because for any bounded U, we can always find a closed disk D large enough s.t. $\overline{U}$ is contained in the interior of D; and if U is unbounded, it must contain all of $\mathbb{R}^2 \setminus B(0, Q)$ (see Lemma 8.4 statement above), and so $\mathbb{R}^2 \setminus U \subseteq B(0, Q) \subseteq \overline{B(0, Q)} =: D$ — this in particular gives us that f is identity on ∂D. Moreover, because r is continuous and the identity is continuous, f is continuous by the pasting/gluing lemma; basically, for a closed set F in the codomain, $f^{-1}(F) = (r^{-1}(F) \cap \overline{U}) \cup (\mathrm{id}^{-1}(F) \cap \mathbb{R}^2 \setminus U)$ is a union of closed sets, hence closed, giving us continuity of f. ■.

8.1.4 Setup with Pictures, and Proving Unique Bounded Component

To study rigorously this curve J, we will need to set things up with clear notation and even clearer definitions (and maybe a picture or ~~two~~ five) to ensure that we know exactly what we're talking about. Again, J is the image of the circle S^1 via a continuous embedding $S^1 \hookrightarrow \mathbb{R}^2$, and so in particular because S^1 is compact (closed and bounded in $\mathbb{R}^2$, see Heine-Borel/Theorem 5.2) and the image of a compact set via a continuous map is compact (Theorem 5.7), J is compact in $\mathbb{R}^2$.

Then, the product space $J \times J$ is compact (Lemma 5.10), and because $d(\mathbf{x}, \mathbf{x}') : J \times J \to [0, \infty)$ is continuous (the metric $d(\mathbf{x}, \mathbf{x}') : \mathbb{R}^2 \times \mathbb{R}^2 \to [0, \infty)$ is continuous because $|\mathbf{x} - \mathbf{x}'| = d(\mathbf{x}, \mathbf{x}') < \epsilon \implies d(\mathbf{x}, \mathbf{x}') < \epsilon$), the image of $J \times J$ via d is compact in $[0, \infty)$ (again Theorem 5.7), so in particular d attains its maximum on $J \times J$, i.e. there are $\mathbf{a}, \mathbf{b} \in J$ s.t. $d(\mathbf{a}, \mathbf{b}) \geq d(\mathbf{x}, \mathbf{x}')$ for any $\mathbf{x}, \mathbf{x}' \in J$. We can now scale J (just using rigid transformations, like rotations, dilations, translations) to bring $\mathbf{a}, \mathbf{b}$ to $(-1, 0)$ and $(1, 0)$ respectively.

The condition $d(\mathbf{a}, \mathbf{b}) \geq d(\mathbf{x}, \mathbf{x}')$ gives in particular that $d(\mathbf{a}, \mathbf{b}) \geq d(\mathbf{a}, \mathbf{x})$ and $d(\mathbf{a}, \mathbf{b}) \geq d(\mathbf{b}, \mathbf{x})$ for any $\mathbf{x} \in J$, so the curve J lies in the intersection of the closed disks of radius 2 centered at $\mathbf{a} = (-1, 0)$ and $\mathbf{b} = (1, 0)$, a "lens" shape that sits in the rectangle $R := [-1, 1] \times [-2, 2]$, where ∂R intersects the coordinate axes at $\mathbf{a}, \mathbf{b}$, and the points $\mathbf{c} := (0, 2)$ and $\mathbf{d} := (0, -2)$. Note that the only place the "lens" shape and ∂R touch are at $\mathbf{a}, \mathbf{b}$, so indeed $J \cap \partial R = \{\mathbf{a}, \mathbf{b}\}$. Observe furthermore that $\mathbf{a}, \mathbf{b}$ split J into two "arcs": parameterizing S^1 using an angle parameter $\theta \in [0, 2\pi)$, if $\theta_a \mapsto \mathbf{a}$ and $\theta_b \mapsto \mathbf{b}$ (where without loss of generality we'll suppose $\theta_a < \theta_b$), then one arc will correspond to $\theta \in [\theta_a, \theta_b)$ and the other will correspond with $\theta \in [\theta_b, \theta_a + 2\pi) \pmod{2\pi}$. This is illustrated in the picture on the left below:

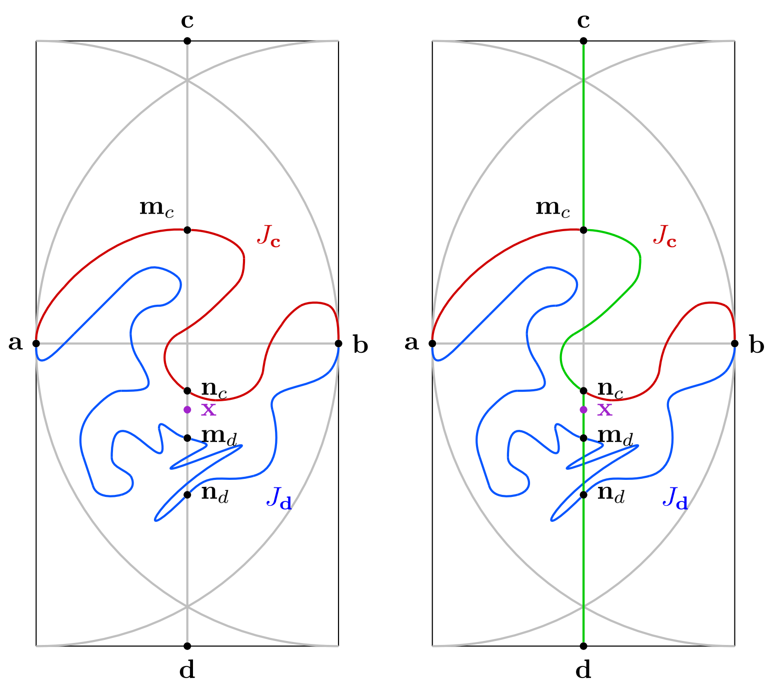

Now we precisely define what the labeled points are in the above pictures. The intersection of J and the line segment $\overline{\mathbf{cd}}$ must be nonempty (by Lemma 8.2), so let us define $\mathbf{m}_c$ to be the point in the set $J \cap \overline{\mathbf{cd}}$ with <u>**m**</u>aximum y-coordinate. Moreover, $\mathbf{m}_c$ is contained in this set, because J is closed and the line segment $\overline{\mathbf{cd}}$ is closed implies that their intersection is closed, and closed sets contain their limit point ($\mathbf{m}_c$ is the "supremum" of this set). Recall in the above paragraph we talked about $\mathbf{a}, \mathbf{b}$ splitting J into two arcs; let us denote by $J_\mathbf{c}$ the arc that contains $\mathbf{m}_c$, and denote by $J_\mathbf{d}$ the other arc. Since we know the set $J_\mathbf{c} \cap \overline{\mathbf{cd}}$ is nonempty (it contains $\mathbf{m}_c$) and closed, it also contains its "infimum", or the point $\mathbf{n}_c$ with mi<u>**n**</u>imum y-coordinate.

We do a similar process to construct $\mathbf{m}_d$ and $\mathbf{n}_d$. We know that the intersection $J_{\mathbf{d}} \cap \overline{\mathbf{n}_c\mathbf{d}}$ must be nonempty; this is because otherwise, the paths $J_{\mathbf{d}}$ and $\overline{\mathbf{cm}_c} + \widetilde{\mathbf{m}_c\mathbf{n}_c} + \overline{\mathbf{n}_c\mathbf{d}}$ would not intersect (the squiggle above the $\mathbf{m}_c\mathbf{n}_c$ indicates it's the path between $\mathbf{m}_c, \mathbf{n}_c$ inside the squiggly curve $J_{\mathbf{c}}$, drawn above on the right in green; more rigorously, we have $J_{\mathbf{c}}$ correspond to the angle parameter on S^1 in the range $[\theta_a, \theta_b]$, and $\mathbf{m}_c, \mathbf{n}_c$ correspond to $\theta_m, \theta_n \in [\theta_a, \theta_b]$, so $\widetilde{\mathbf{m}_c\mathbf{n}_c}$ corresponds to $[\theta_m, \theta_n]$) would not intersect (no intersection along the first piece $\overline{\mathbf{cm}_c}$ because $\mathbf{m}_c = \max\{J \cap \overline{\mathbf{cd}}\}$; no intersection along second piece because Jordan curve J embedding is injective; and no intersection along third piece by assumption), contradicting Lemma 8.2. Then as above, we define $\mathbf{m}_d := \max\{J_{\mathbf{d}} \cap \overline{\mathbf{n}_c\mathbf{d}}\}$ and $\mathbf{n}_d := \min\{J_{\mathbf{d}} \cap \overline{\mathbf{n}_c\mathbf{d}}\}$. Note that $\mathbf{n}_c > \mathbf{m}_d$ (comparing y-coordinates), where "$\geq$" comes from the fact that $\mathbf{m}_d \in J_{\mathbf{d}} \cap \overline{\mathbf{n}_c\mathbf{d}}$, and "$>$" is because the Jordan curve embedding is injective. We can now define $\mathbf{x}$ to be some arbitrary point on $\overline{\mathbf{cd}}$ s.t. $\mathbf{n}_c > \mathbf{x} > \mathbf{m}_d$. The pictures above are now completely explained, and the (very long!) setup is complete.

Component containing x is bounded: let us call this component U, and suppose that is is not bounded. Then, there is a path in U from $\mathbf{x}$ to some point outside the rectangle R (this is illustrated in the leftmost picture below). Let $\mathbf{e}$ denote the first point along this path that intersects ∂R ($\mathbf{e}$ for <u>e</u>dge). The y-coordinate of $\mathbf{e}$ can not be 0, because that would mean that $\mathbf{e} = \mathbf{a}$ or $\mathbf{b}$, but $\mathbf{a}, \mathbf{b} \in J$ and hence can not be in U because U is defined to be a component of $\mathbb{R}^2 \setminus J$. If the y-coordinate of $\mathbf{e}$ is < 0, then denoting by $\widehat{\mathbf{ed}}$ the path along ∂R from $\mathbf{e}$ to $\mathbf{d}$, we see that the paths $J_{\mathbf{d}}$ and $\overline{\mathbf{cm}_c} + \widetilde{\mathbf{m}_c\mathbf{n}_c} + \overline{\mathbf{n}_c\mathbf{x}} + \widetilde{\mathbf{xe}} + \widehat{\mathbf{ed}}$ would not intersect (first and second pieces discussed in preceding paragraph; no intersection along the third piece $\overline{\mathbf{n}_c\mathbf{x}}$ because $\mathbf{x} > \mathbf{m}_d$ where $\mathbf{m}_d := \max\{J_{\mathbf{d}} \cap \overline{\mathbf{n}_c\mathbf{d}}\}$; no intersection along fourth piece by assumption; and no intersection along fifth piece because $J \cap \partial R = \{\mathbf{a}, \mathbf{b}\}$), contradicting Lemma 8.2. This is illustrated in the center picture below. The exact same argument works for the "> 0" case (or, you could just flip across the x-axis and use the "< 0" case).

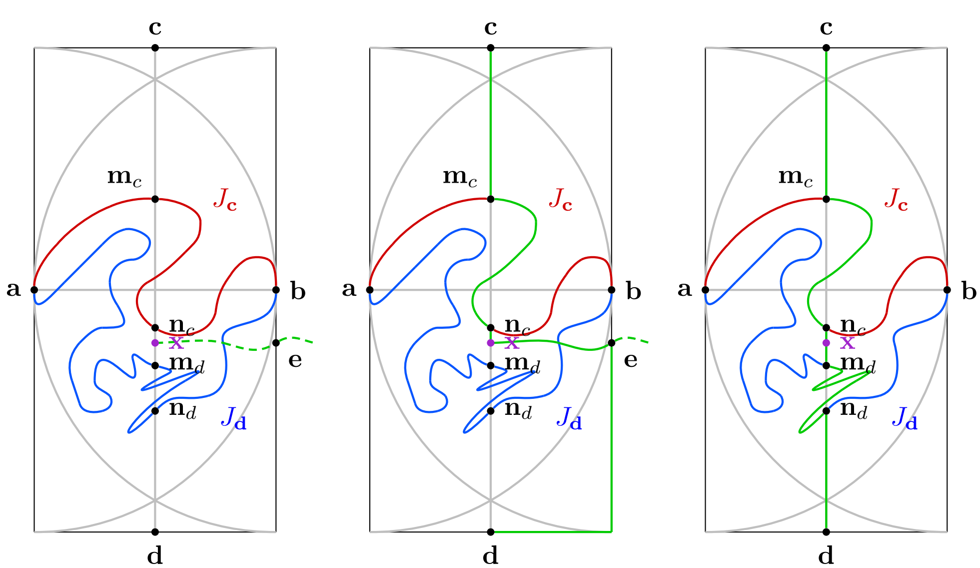

No other bounded component: suppose we had another bounded component V. Then, the path $\Gamma := \overline{\mathbf{c}\mathbf{m}_c} + \widetilde{\mathbf{m}_c\mathbf{n}_c} + \overline{\mathbf{n}_c\mathbf{m}_d} + \widetilde{\mathbf{m}_d\mathbf{n}_d} + \overline{\mathbf{n}_d\mathbf{d}}$ (illustrated in the rightmost picture above) would not intersect V (the first and fifth pieces are path connected with ∂R, which is path connected to any point of $\mathbb{R}^2 \setminus R$, which of course intersects the unbounded component O, meaning the first and fifth pieces are part of O, not V; the second and fourth pieces are contained in J, which is not in any component; and the third piece is path connected to $\mathbf{x}$, hence is part of the component U from the preceding paragraph, not V). Because Γ is closed and does not contain $\mathbf{a}, \mathbf{b}$, $\mathbb{R}^2 \setminus \Gamma$ is open, so there are open balls B_a and B_b containing $\mathbf{a}, \mathbf{b}$ respectively that do not intersect Γ. By Lemma 8.4, we know that $\mathbf{a}, \mathbf{b} \in \overline{V}$, so B_a, B_b intersect V nontrivially, say at $\mathbf{v}_a \in B_a \cap V$ and $\mathbf{v}_b \in B_b \cap V$. Then, since V is path connected we have a path $\widetilde{\mathbf{v}_a\mathbf{v}_b}$ in V (hence not intersecting Γ), and line segments $\overline{\mathbf{a}\mathbf{v}_a}$ and $\overline{\mathbf{v}_b\mathbf{b}}$ lying in B_a, B_b respectively (hence not intersecting Γ). We have just proven that these two paths, Γ and $\overline{\mathbf{a}\mathbf{v}_a} + \widetilde{\mathbf{v}_a\mathbf{v}_b} + \overline{\mathbf{v}_b\mathbf{b}}$, do not intersect, which of course contradicts Lemma 8.2.

In summary, we have shown that for our arbitrary Jordan curve J, $\mathbb{R}^2 \setminus J$ consists of exactly one unbounded component O, a bounded component U containing $\mathbf{x}$, and no other components (i.e. $\mathbb{R}^2 \setminus J$ consists of exactly two components, one bounded and the other unbounded), which furthermore satisfy (by Lemma 8.4) that their boundaries equal J exactly. This is exactly the statement of Theorem 8.1. ■

8.1.5 Concluding Remarks

Unfortunately, this proof does not generalize to higher dimensions, or at least I have not seen any way of doing it without much messier geometric arguments than the ones here (our proof was based on proving contradictions by intersecting paths in the plane, but in higher dimension, we would have to intersect planes, which are much more difficult to deal with).

9 Acknowledgements

Most figures made with TikZit (all with generous referencing to TEX.stackexchange.com, and commutative diagram made with Shen Yi Chuan's TikzCD editor. Great thanks to my mentor, Monty McGovern. Thus concludes a project one year in the making, to my great relief. It is done!

References

[1] Tyrone (https://math.stackexchange.com/users/258571/tyrone). *Munkres Topology section 50 exercise 6(e), continuous function that tends to infinity.* Mathematics Stack Exchange. (version: 2020-12-13). eprint: https://math.stackexchange.com/q/3947402. URL: https://math.stackexchange.com/q/3947402.

[2] Maxime Ramzi (https://math.stackexchange.com/users/408637/maxime-ramzi). *Could *I* have come up with the definition of Compactness (and Connectedness)?* Mathematics Stack Exchange. (version: 2019-08-07). eprint: https://math.stackexchange.com/q/3316531. URL: https://math.stackexchange.com/q/3316531.

[3] Daniel Fischer (https://math.stackexchange.com/users/83702/daniel-fischer). *How can 1-manifold can be written as a finite union of spaces homeomorphic to [0,1]*. Mathematics Stack Exchange. (version: 2014-05-09). eprint: https://math.stackexchange.com/q/788474. URL: https://math.stackexchange.com/q/788474.

[4] nLab authors. *one-point compactification*. http://ncatlab.org/nlab/show/one-point%20compactification. Revision 39. May 2021.

[5] Sina Greenwood and Jiling Cao. *Brouwer's Fixed Point Theorem and the Jordan Curve Theorem (class notes)*. University of Auckland, New Zealand, class "MATHS 750 - Topology". Spring 2006. URL: https://www.math.auckland.ac.nz/class750/section5.pdf.

[6] John M. Lee. *Introduction to Topological Manifolds*. 2nd edition. Graduate Texts in Mathematics. Springer, 2010. ISBN: 978-1-4419-7939-1.

[7] Ryuji Maehara. "The Jordan Curve Theorem Via the Brouwer Fixed Point Theorem". In: *The American Mathematical Monthly* 91.10 (1984), pp. 641–643. DOI: 10.1080/00029890.1984.11971517. eprint: https://doi.org/10.1080/00029890.1984.11971517. URL: https://doi.org/10.1080/00029890.1984.11971517.

[8] James R. Munkres. *Topology; a First Course*. 2nd edition. Prentice Hall, Inc. Upper Saddle River, NJ, 2000. ISBN: 0-13-181629-2.

[9] Halsey L. Royden and Patrick M. Fitzpatrick. *Real Analysis*. 4th edition. Prentice Hall, 2010. ISBN: 978-0131437470.

[10] Francis Edward Su. "Rental Harmony: Sperner's Lemma in Fair Division". In: *The American Mathematical Monthly* 106.10 (1999), pp. 930–942. DOI: 10.1080/00029890.1999.12005142. eprint: https://doi.org/10.1080/00029890.1999.12005142. URL: https://doi.org/10.1080/00029890.1999.12005142.

[11] Scott Taylor. *Invariance of Dimension (class notes)*. Colby College, class "MA331". Spring 2009. URL: http://personal.colby.edu/~sataylor/teaching/S09/MA331/InvarianceOfDomain.pdf.

[12] Alex Wright. *SPERNER'S LEMMA AND BROUWER'S FIXED POINT THEOREM*. University of Waterloo (undergraduate project). July 2005. URL: http://www-personal.umich.edu/~alexmw/BFPT.pdf.

[13] David G. Wright. "Tychonoff's theorem". In: *Proceedings of the American Mathematical Society* 120.3 (Mar. 1994). URL: https://www.ams.org/journals/proc/1994-120-03/S0002-9939-1994-1170549-2/S0002-9939-1994-1170549-2.pdf.