MATH 6540
General Topology

Paul Skoufranis

August 3, 2021

Preface:

These are the second edition of these lecture notes for MATH 6540 (General Topology). Consequently, there may be several typographical errors. However, the goal of these notes is to be fairly self-contained provided one has the necessary background. If you come across any typos, errors, omissions, or unclear explanations, please feel free to contact me so that I may continually improve these notes.

Contents

Chapter 1

Topological Spaces

Topology, from the Greek $\tau \acute{o} \pi o \sigma$ meaning place and $\lambda \acute{o} \gamma o \sigma$ meaning study, is the study of properties of spaces and their deformations. As there are many different branches of topology, such as differential topology, geometric topology, and algebraic topology, this course will study what is known as point-set (or general) topology as it is the foundation of all branches of topology.

Point-set topology has been in existence in various forms for centuries and has been worked out in great detail and generality. Consequently, although the base concepts are not difficult by today's mathematical standards, the generality and plethora of examples does add some complications as ones intuition based on Euclidean spaces may fail.

In this chapter, we will begin by defining what this course pertains to; that is, what is a topology on a space? Subsequently, we will provide several examples of topological spaces. Of course, this list is nowhere near exhaustive as one can easily produce various exotic topological spaces if one requires. We will then proceed by describing simplifications one can use to understand the topology on a space and how one can construct new topologies from old topologies. The notion of topology then leads us to the notions of convergence where we observe sequences are no longer sufficient and the notions of various types of sets and points in a topological space.

1.1 Topologies

In order to study topologies in this course, we must first mathematically define what a topology is of course. We refer the reader to Appendix A for definitions and notation regarding sets (e.g. given a set X, $\mathcal{P}(X)$ denotes the power set of X).

Definition 1.1.1. Let X be a set. A set $\mathcal{T} \subseteq \mathcal{P}(X)$ is said to be a *topology* on X if

(1) $\emptyset, X \in \mathcal{T}$,

(2) (closed under unions) if $\{U_\alpha\}_{\alpha \in I} \subseteq \mathcal{T}$, then $\bigcup_{\alpha \in I} U_\alpha \in \mathcal{T}$, and

(3) (closed under finite intersections) if $\{U_\alpha\}_{\alpha \in I} \subseteq \mathcal{T}$ and I is finite, then $\bigcap_{\alpha \in I} U_\alpha \in \mathcal{T}$.

The pair $(X, \mathcal{T})$ is called a *topological space* and elements of $\mathcal{T}$ are called the *open sets* of $(X, \mathcal{T})$.

As we will see, there are many topologies we can place on a given set X so by saying that $(X, \mathcal{T})$ is a topological space means we have fixed $\mathcal{T}$ to be the topology on X. Once a topology is fixed on a set, one can think of the open sets as the sets that describe how points are related to one another. In particular, open sets provide some notion of whether two points are 'close' together; that is, given two points $x, y \in X$ and a $U \in \mathcal{T}$ such that $x \in U$, then y is close to x with respect to U only if $y \in U$. Thus we can see the above definition and thoughts are motivated by undergraduate real analysis where the 'open sets' on $\mathbb{R}$ were the sets that were unions of open intervals and that two points were 'close' only if there was a 'small' open interval around open point which contained the other. Consequently, one hope in this course is to generalize the nice analytical properties of $\mathbb{R}$ seen in undergraduate real analysis (such as convergence and continuous functions) to as general a setting as possible.

As we desire to study topological spaces, it is useful to have some examples to keep in mind. Of course the examples presented in this section are not all the examples in existence and we will continually encounter new topologies through the course.

Example 1.1.2. Let X be a set. Then $\mathcal{T} = \{\emptyset, X\}$ is a topology on X known as the *trivial topology*. This name derives from the fact that the open sets do not distinguish any two elements of X and most topological results become trivial if we consider this topology. We remark that it is trivial to verify the trivial topology is a topology.

Example 1.1.3. Let X be a set. Then $\mathcal{T} = \mathcal{P}(X)$ is a topology on X known as the *discrete topology*. This name derives from the fact that every set is open so singleton sets are open and thus every point is separated from the others. We remark that it is trivial to verify the discrete topology is a topology.

Example 1.1.4. Consider the set X consisting of three distinct points $\{a, b, c\}$. Up-to relabelling these three points to avoid symmetry, there are 20 possible intrinsically different subsets of $\mathcal{P}(X)$ containing X and $\emptyset$ for which only 9 are topologies. Below we represent these subsets of $\mathcal{P}(X)$ using both set notation and a diagram where each circle represents an element of the subset by the points it its interior.

	Diagram	**Set Notation**	**A Topology?**
1)	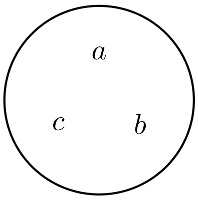	$\{\emptyset, X\}$	Yes
2)	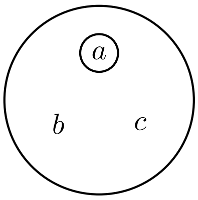	$\{\emptyset, \{a\}, X\}$	Yes
3)	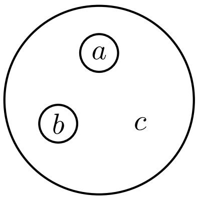	$\{\emptyset, \{a\}, \{b\}, X\}$	No
4)	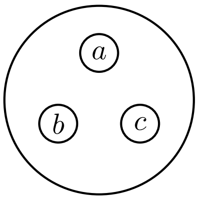	$\{\emptyset, \{a\}, \{b\}, \{c\}, X\}$	No
5)	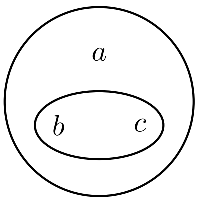	$\{\emptyset, \{b, c\}, X\}$	Yes
6)	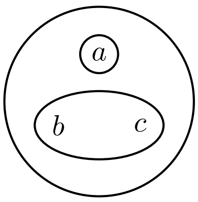	$\{\emptyset, \{a\}, \{b, c\}, X\}$	Yes
7)	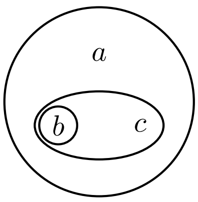	$\{\emptyset, \{b\}, \{b, c\}, X\}$	Yes

	Diagram	**Set Notation**	**A Topology?**
8)	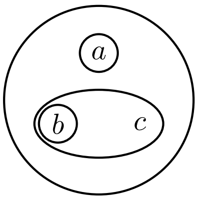	$\{\emptyset, \{a\}, \{b\}, \{b, c\}, X\}$	No
9)	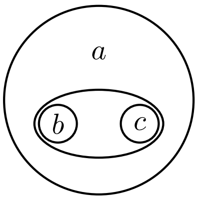	$\{\emptyset, \{b\}, \{c\}, \{b, c\}, X\}$	Yes
10)	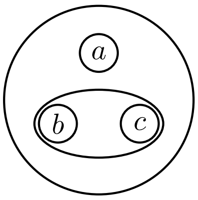	$\{\emptyset, \{a\}, \{b\}, \{c\}, \{b, c\}, X\}$	No
11)	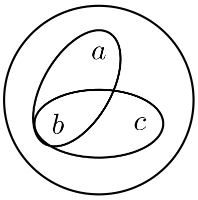	$\{\emptyset, \{a, b\}, \{b, c\}, X\}$	No
12)	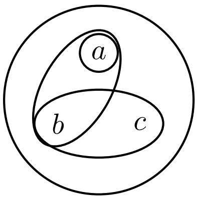	$\{\emptyset, \{a\}, \{a, b\}, \{b, c\}, X\}$	No
13)	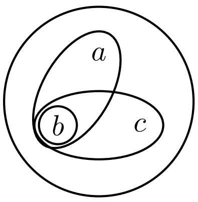	$\{\emptyset, \{b\}, \{a, b\}, \{b, c\}, X\}$	Yes
14)	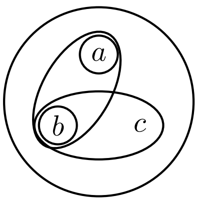	$\{\emptyset, \{a\}, \{b\}, \{a, b\}, \{b, c\}, X\}$	Yes

15)	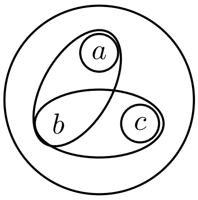	$\{\emptyset, \{a\}, \{c\}, \{a,b\}, \{b,c\}, X\}$	No
16)	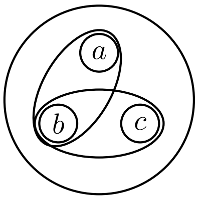	$\{\emptyset, \{a\}, \{b\}, \{c\}, \{a,b\}, \{b,c\}, X\}$	No
17)	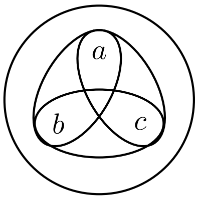	$\{\emptyset, \{a,b\}, \{a,c\}, \{b,c\}, X\}$	No
18)	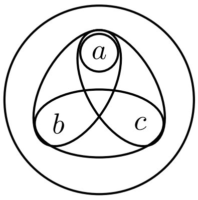	$\{\emptyset, \{a\}, \{a,b\}, \{a,c\}, \{b,c\}, X\}$	No
19)	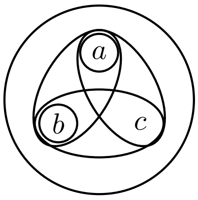	$\{\emptyset, \{a\}, \{b\}, \{a,b\}, \{a,c\}, \{b,c\}, X\}$	No
20)	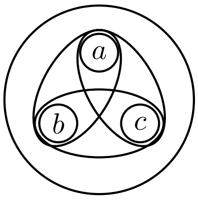	$\{\emptyset, \{a\}, \{b\}, \{c\}, \{a,b\}, \{a,c\}, \{b,c\}, X\}$	Yes

Note sets (3), (4), (8), and (10) are not topologies as $\{a\}$ and $\{b\}$ are in the sets but $\{a,b\} = \{a\} \cup \{b\}$ is not in the sets. Note (11), (12), (15), (17), and (18) are not topologies as $\{a,b\}$ and $\{b,c\}$ are in the sets but $\{b\} = \{a,b\} \cap \{b,c\}$ are not in the sets. Note (16) is not a topology as $\{a\}$ and $\{c\}$ is in the set but $\{a,c\} = \{a\} \cup \{c\}$ is not in the set. Finally (19) is not a topology as $\{a,c\}$ and $\{b,c\}$ are in the set but $\{c\} = \{a,c\} \cap \{b,c\}$ is not in the set. The remaining 9 sets can be verified to be topologies as they are closed under unions and (finite) intersections.

Example 1.1.5. Let X be any set and let

$$\mathcal{T} = \{\emptyset\} \cup \{A \subseteq X \mid X \setminus A \text{ is finite}\}.$$

Then $\mathcal{T}$ is a topology on X. To see this, we note that clearly $\emptyset \in \mathcal{T}$ and that $X \in \mathcal{T}$ as $X \setminus X = \emptyset$. Next, to see that $\mathcal{T}$ is closed under unions, let $\{A_\alpha\}_{\alpha \in I} \subseteq \mathcal{T}$ be arbitrary. Thus $X \setminus A_\alpha$ is finite for all $\alpha \in I$. Since

$$X \setminus \left(\bigcup_{\alpha \in I} A_\alpha \right) = \bigcap_{\alpha \in I} \left(X \setminus A_\alpha \right),$$

we see that $X \setminus (\bigcup_{\alpha \in I} A_\alpha)$ is a subset of a finite set and thus finite. Hence $\bigcup_{\alpha \in I} A_\alpha \in \mathcal{T}$ by definition. Finally, to see that $\mathcal{T}$ is closed under finite intersections, let $\{A_\alpha\}_{\alpha \in I} \subseteq \mathcal{T}$ with I finite be arbitrary. Thus $X \setminus A_\alpha$ is finite for all $\alpha \in I$. Since

$$X \setminus \left(\bigcap_{\alpha \in I} A_\alpha \right) = \bigcup_{\alpha \in I} \left(X \setminus A_\alpha \right),$$

we see that $X \setminus (\bigcap_{\alpha \in I} A_\alpha)$ a finite union of finite sets and thus finite. Hence $\bigcap_{\alpha \in I} A_\alpha \in \mathcal{T}$ by definition.

The topology $\mathcal{T}$ on X is called the *cofinite topology on X*.

Example 1.1.6. Let X be any set and let

$$\mathcal{T} = \{\emptyset\} \cup \{A \subseteq X \mid X \setminus A \text{ is countable}\}.$$

Then $\mathcal{T}$ is a topology on X. To see this, one need to simply repeat the proof of Example 1.1.5 with 'finite' replaced with 'countable' in the appropriate places.

The topology $\mathcal{T}$ on X is called the *cocountable topology on X*.

Example 1.1.7. Let X be any set and let

$$\mathcal{T} = \{A \subseteq X \mid A \text{ is finite}\}.$$

Notice if X is finite then $\mathcal{T} = \mathcal{P}(X)$ so that $\mathcal{T}$ is the discrete topology on X. However, if X is infinite then $\mathcal{T}$ is not a topology on X as $\mathcal{T}$ is not closed under unions (i.e. a countable union of finite sets is not finite).

Since we have seen that there are many possible topologies on a given set, it is useful to be able to compare the size of these topologies. The simplest way to compare topologies is based on inclusion and this notion of comparison has many analytic implications that will be seen throughout the course.

Definition 1.1.8. Let $\mathcal{T}$ and $\mathcal{T}'$ be two topologies on a set X. It is said that $\mathcal{T}$ is *finer* that $\mathcal{T}'$ or, equivalently, that $\mathcal{T}'$ is *coarser* than $\mathcal{T}$ if $\mathcal{T}' \subseteq \mathcal{T}$. In the case the inclusion is strict (i.e. $\mathcal{T}' \subsetneq \mathcal{T}$), it is said that $\mathcal{T}$ is *strictly finer* that $\mathcal{T}'$ or, equivalently, that $\mathcal{T}'$ is *coarser strictly* than $\mathcal{T}$. Finally, it is said that $\mathcal{T}$ and $\mathcal{T}'$ are *comparable* if $\mathcal{T} \subseteq \mathcal{T}'$ or $\mathcal{T}' \subseteq \mathcal{T}$.

The above terminology is derived from the fact that “if you have more sets in your topology, you can ‘divide up your space’ more finely”. That is, the more pixels per square inch, the finer the image.

Example 1.1.9. The discrete topology on a set is always finer than any other topology on the set and the trivial topology is always coarser than any other topology on the set. Provided the set is non-empty and not a singleton, the discrete topology is strictly finer than the trivial topology.

Example 1.1.10. The cofinite topology is coarser than the cocountable topology and will be strictly coarser provided the set is infinite.

Example 1.1.11. Consider the set X consisting of three distinct points $\{a, b, c\}$ and the following topologies on X exhibited in Example 1.1.4:

$$\begin{aligned} \mathcal{T}_1 &= \{\emptyset, \{b\}, \{c\}, \{b, c\}, X\} \\ \mathcal{T}_2 &= \{\emptyset, \{b\}, \{a, b\}, \{b, c\}, X\}. \end{aligned}$$

As $\mathcal{T}_1 \not\subseteq \mathcal{T}_2$ and $\mathcal{T}_2 \not\subseteq \mathcal{T}_1$, $\mathcal{T}_1$ and $\mathcal{T}_2$ are not comparable topologies on X. Hence it is possible to have topologies that are not comparable.

Of course, we will see later in this course how having finer/coarse topologies affect topological properties and object we desire to study.

1.2 Metric Spaces

All of the above topologies are useful examples of topologies. However, the motivation for the notion of topology was structure of open subsets of $\mathbb{R}$ derived from open intervals. As each finite open interval can be described as the set of points within a certain distance of the centre of the interval, perhaps by generalizing the notion of a distance function from $\mathbb{R}$ to other sets will yield particularly nice topologies to study in this course.

Definition 1.2.1. Let X be a non-empty set. A *metric* on X is a function $d : X \times X \to [0, \infty)$ such that

(1) if $x, y \in X$ then $d(x, y) = 0$ if and only if $x = y$,

(2) if $x, y \in X$ then $d(x, y) = d(y, x)$, and

(3) (Triangle Inequality) if $x, y, z \in X$, then $d(x, y) \leq d(x, z) + d(z, y)$.

A *metric space* is a pair (X, d) where X is a non-empty set and d is a metric on X.

Of course, our motivating example is a metric space and there are many other metrics. As many of our results work both for real and complex numbers, we will use $\mathbb{K}$ to denote both $\mathbb{R}$ and $\mathbb{C}$.

Example 1.2.2. Consider the function $d : \mathbb{K} \times \mathbb{K} \to [0, \infty)$ defined by $d(x, y) = |x - y|$. Then it is elementary to verify that d is a metric on $\mathbb{K}$. Unless otherwise specified, whenever $\mathbb{K}$ is considered as a metric space, we will assume that $\mathbb{K}$ is equipped with this metric.

Example 1.2.3. Let X be any non-empty set and define $d : X \times X \to [0, \infty)$ by

$$d(x, y) = \begin{cases} 0 & \text{if } x = y \\ 1 & \text{if } x \neq y \end{cases}$$

for any $x, y \in X$. It is elementary to verify that d is a metric on X. We call d the *discrete metric on* X.

Example 1.2.4. Let X be a non-empty set, let (Y, d_Y) be a metric space (e.g. $Y = \mathbb{K}$ with the canonical metric), and let $\mathcal{F}(X, Y)$ denote the set of all Y-valued functions with domain X. Define the map $d : \mathcal{F}(X, Y) \times \mathcal{F}(X, Y) \to [0, \infty)$ by

$$d(f, g) = \sup_{x \in X} \min(\{d_Y(f(x), g(x)), 1\})$$

for all $f, g \in \mathcal{F}(X, Y)$. It is not difficult to verify that d is a metric on $\mathcal{F}(X, Y)$. We call d the *uniform metric on* $\mathcal{F}(X, Y)$.

Of course there are many other examples of metrics. Some of the most important examples of metrics come from the following structure.

Definition 1.2.5. Let V be a vector space over $\mathbb{K}$. A *norm* on V is a function $\|\cdot\| : V \to [0, \infty)$ such that

(1) if $\vec{v} \in V$ then $\|\vec{v}\| = 0$ if and only if $\vec{v} = \vec{0}$,

(2) if $\alpha \in \mathbb{K}$ and $\vec{v} \in V$ then $\|\alpha\vec{v}\| = |\alpha| \, \|\vec{v}\|$, and

(3) (Triangle Inequality) if $\vec{v}, \vec{w} \in V$ then $\|\vec{v} + \vec{w}\| \leq \|\vec{v}\| + \|\vec{w}\|$.

A *normed linear space* is a pair $(V, \|\cdot\|)$ where V is a vector space over $\mathbb{K}$ and $\|\cdot\|$ is a norm on V.

Given a normed linear space $(V, \|\cdot\|)$, it is easy to use the norm to construct a metric on V.

Proposition 1.2.6. *Let $(V, \|\cdot\|)$ be a normed linear space and define $d : V \times V \to [0, \infty)$ by*

$$d(\vec{x}, \vec{y}) = \|\vec{x} - \vec{y}\|$$

for all $\vec{x}, \vec{y} \in V$. Then d is a metric on V.

Proof. First notice that d is clearly well-defined and maps into $[0, \infty)$ by the definition and properties of a normed linear space. To see that d is a metric on V, we will verify Definition 1.2.1.

First notice for all $\vec{x}, \vec{y} \in V$ that $d(\vec{x}, \vec{y}) = 0$ if and only if $\|\vec{x} - \vec{y}\| = 0$ if and only if $\vec{x} - \vec{y} = \vec{0}$ if and only if $\vec{x} = \vec{y}$. Hence d satisfies the first defining property from Definition 1.2.1.

Next, notice for all $\vec{x}, \vec{y} \in V$ that

$$d(\vec{x}, \vec{y}) = \|\vec{x} - \vec{y}\| = \|(-1)(\vec{y} - \vec{x})\| = |-1| \, \|\vec{y} - \vec{x}\| = d(\vec{y}, \vec{x})$$

by the second property of a norm from Definition 1.2.5 and by vector space properties. Hence d satisfies the second defining property from Definition 1.2.1.

Finally, notice for all $\vec{x}, \vec{y}, \vec{z} \in V$ that

$$\begin{aligned} d(\vec{x}, \vec{y}) &= \|\vec{x} - \vec{y}\| \\ &= \|(\vec{x} - \vec{z}) + (\vec{z} - \vec{y})\| \\ &\leq \|\vec{x} - \vec{z}\| + \|\vec{z} - \vec{y}\| \\ &= d(\vec{x}, \vec{z}) + d(\vec{z}, \vec{y}) \end{aligned}$$

by the third property of a norm from Definition 1.2.5 and by vector space properties. Hence d satisfies the Triangle Inequality and thus is a metric on V. ∎

Definition 1.2.7. Given a normed linear space $(V, \|\cdot\|)$, the metric d from Proposition 1.2.6 is called the *norm induced metric*. Unless otherwise specified, we will always consider a normed linear space a metric space equipped with the norm induced metric.

Of course there are many examples of norms. In particular, the canonical metric on $\mathbb{K}$ comes from a norm.

Example 1.2.8. The absolute value function $|\cdot| : \mathbb{K} \to [0, \infty)$ is a norm on $\mathbb{K}$.

Example 1.2.9. For a natural number n, the map $\|\cdot\|_1 : \mathbb{K}^n \to [0, \infty)$ defined by

$$\|(x_1, x_2, \ldots, x_n)\|_1 = \sum_{k=1}^{n} |x_k|$$

is a norm on $\mathbb{K}^n$ known as the 1-*norm*. It is not difficult to verify that the 1-norm is indeed a norm on $\mathbb{K}^n$.

Example 1.2.10. For a natural number n, the map $\|\cdot\|_\infty : \mathbb{K}^n \to [0, \infty)$ defined by

$$\|(x_1, x_2, \ldots, x_n)\|_\infty = \max\{|x_1|, |x_2|, \ldots, |x_n|\}$$

is a norm on $\mathbb{K}^n$ known as the *∞-norm*. It is not difficult to verify that the ∞-norm is indeed a norm on $\mathbb{K}^n$.

Example 1.2.11. For a natural number n, the map $\|\cdot\|_2 : \mathbb{K}^n \to [0, \infty)$ defined by

$$\|(x_1, x_2, \ldots, x_n)\|_2 = \left(\sum_{k=1}^n |x_k|^2\right)^{\frac{1}{2}}$$

is a norm on $\mathbb{K}^n$ known as the *Euclidean norm*. The fact that the Euclidean norm is a norm can be verified using the Cauchy-Scwarz Inequality on $\mathbb{K}^n$ or by verifying the following examples of norms are indeed norms.

Example 1.2.12. For a natural number n and $p \in [1, \infty)$, the map $\|\cdot\|_p : \mathbb{K}^n \to [0, \infty)$ defined by

$$\|(x_1, x_2, \ldots, x_n)\|_p = \left(\sum_{k=1}^n |x_k|^p\right)^{\frac{1}{p}}$$

is a norm on $\mathbb{K}^n$ known as the *p-norm*. The fact that the p-norm is actually a norm is left to Appendix B.

In fact, the above norms can be extended to certain types of sequences.

Example 1.2.13. For $p \in [1, \infty)$ let

$$\ell_p(\mathbb{K}) = \left\{ (x_n)_{n \geq 1} \,\middle|\, \sum_{k=1}^\infty |x_k|^p < \infty \right\}.$$

Then $\ell_p(\mathbb{K})$ is a vector space over $\mathbb{K}$ and the map $\|\cdot\|_p : \ell_p(\mathbb{K}) \to [0, \infty)$ defined by

$$\|(x_n)_{n\geq 1}\|_p = \left(\sum_{k=1}^\infty |x_k|^p\right)^{\frac{1}{p}}$$

is a norm on $\ell_p(\mathbb{K})$ known as the *p-norm*. The fact $\ell_p(\mathbb{K})$ is a vector space over $\mathbb{K}$ and that the p-norm is actually a norm is left to Appendix B for $p \neq 1$ and is simple when $p = 1$.

Example 1.2.14. Let

$$\ell_\infty(\mathbb{K}) = \left\{ (x_n)_{n \geq 1} \,\middle|\, \sup_{n\geq 1} |x_n| < \infty \right\}.$$

Then $\ell_\infty(\mathbb{K})$ is a vector space over $\mathbb{K}$ and the map $\|\cdot\|_\infty : \ell_\infty(\mathbb{K}) \to [0,\infty)$ defined by

$$\|(x_n)_{n\geq 1}\|_\infty = \sup_{n\geq 1} |x_n|$$

is a norm on $\ell_\infty(\mathbb{K})$ known as the ∞-*norm*. It is not difficult to verify that $\ell_\infty(\mathbb{K})$ is a vector space over $\mathbb{K}$ and that the ∞-norm is indeed a norm on $\ell_\infty(\mathbb{K})$.

Of course, not all metrics on vector spaces come from norms.

Example 1.2.15. The uniform metric and discrete metric are not norm induced metrics on any non-zero vector space. To see this, notice that if d is either the uniform or discrete metric, then the range of d is contained in $[0,1]$. However, given any non-zero normed linear space $(V, \|\cdot\|)$ and the norm induced metric d on V, the range of d must be all of $[0,\infty)$. To see this, fix a non-zero vector $\vec{v} \in V$. Hence Definition 1.2.5 implies that $\|\vec{v}\| \neq 0$. Since for all $t \in \mathbb{R}$

$$d(t\vec{v}, \vec{0}) = \|t\vec{v}\| = |t|\,\|\vec{v}\|$$

and since $\|\vec{v}\| \neq 0$, we obtain that the range of d must be all of $[0,\infty)$ as claimed.

Of course, we have defined metrics and norms in order to generalize the topological structure of $\mathbb{R}$ used in undergraduate real analysis. Now that we have all of these examples, we must understand how we get a topology on these spaces. To do this, we generalize the idea of an open interval.

Definition 1.2.16. Let (X,d) be a metric space. Given an $x \in X$ and an $r > 0$, the *open d-ball of radius r centred at x*, denoted $B_d(x,r)$, is the set

$$B_d(x,r) = \{y \in X \mid d(x,y) < r\}.$$

Using the open balls, we generalize the idea on $\mathbb{R}$ that the open sets are the unions of open intervals.

Theorem 1.2.17. *Let (X,d) be a metric space. Let $\mathcal{T}_d$ be the set of all subsets U of X such that for each $x \in U$ there exists an $\epsilon > 0$ such that $B_d(x,\epsilon) \subseteq U$. Then $\mathcal{T}_d$ is a topology on X.*

Proof. To see that $\mathcal{T}_d$ is a topology, we must verify the three properties in Definition 1.1.1. It is clear by definition that $\emptyset, X \in \mathcal{T}_d$.

Suppose $\{U_\alpha\}_{\alpha\in I}$ is a set of elements of $\mathcal{T}_d$. To see that $\bigcup_{\alpha\in I} U_\alpha \in \mathcal{T}_d$, let $x \in \bigcup_{\alpha\in I} U_\alpha$ be arbitrary. Then there must be an $\alpha_0 \in I$ such that $x \in U_{\alpha_0}$. Since $U_{\alpha_0} \in \mathcal{T}_d$, there exists an $\epsilon > 0$ such that $B_d(x,\epsilon) \subseteq U_{\alpha_0}$. Hence $B_d(x,\epsilon) \subseteq U_{\alpha_0} \subseteq \bigcup_{\alpha\in I} U_\alpha$. As $x \in \bigcup_{\alpha\in I} U_\alpha$ was arbitrary, $\bigcup_{\alpha\in I} U_\alpha \in \mathcal{T}_d$.

Finally, suppose $U_1, \ldots, U_n \in \mathcal{T}_d$. To see that $\bigcap_{k=1}^n U_k \in \mathcal{T}_d$, let $x \in \bigcap_{k=1}^n U_k$ be arbitrary. Hence $x \in U_k$ for all k so, as each $U_k \in \mathcal{T}_d$, there exists an $\epsilon_k > 0$ such that $B_d(x, \epsilon_k) \subseteq U_k$ for all k. Let $\epsilon = \min_{1 \leq k \leq n} \epsilon_k > 0$. Since

$$B_d(x, \epsilon) \subseteq B_d(x, \epsilon_k) \subseteq U_k$$

for all k, we see that $B_d(x, \epsilon) \subseteq \bigcap_{k=1}^n U_k$. As $x \in \bigcap_{k=1}^n U_k$ was arbitrary, we obtain that $\bigcap_{k=1}^n U_k \in \mathcal{T}_d$ as desired. ∎

As the topologies from Theorem 1.2.17 are incredibly important examples, it is useful to give them a name.

Definition 1.2.18. Let (X, d) be a metric space. The topology $\mathcal{T}_d$ from Theorem 1.2.17 is called the *metric topology on X induced by d*. Unless otherwise specified, we will always view a metric space as a topological space with the metric topology.

Specifically, as the topology on $\mathbb{K}$ induced by the absolutely value is the most basic example of a topology and thus will be used often, it is also worth of a name.

Definition 1.2.19. The metric topology on $\mathbb{K}$ induced by the absolute value metric is called the *canonical topology on $\mathbb{K}$*.

Of course, now that we have a topology on each metric space, we should actually verify that each open ball is indeed an open set.

Proposition 1.2.20. *Let (X, d) be a metric space. For all $x \in X$ and $r > 0$, $B_d(x, r) \in \mathcal{T}_d$.*

Proof. Fix $x \in X$ and $r > 0$. To see that $B_d(x, r)$ is open, let $y \in B_d(x, r)$ be arbitrary. Thus $d(x, y) < r$.

Let $\delta = r - d(x, y) > 0$. We claim that $B_d(y, \delta) \subseteq B_d(x, r)$. To see this, let $z \in B_d(y, \delta)$ be arbitrary. Then $d(z, y) < \delta$ so, by the Triangle Inequality,

$$d(z, x) \leq d(z, y) + d(y, x) < \delta + d(y, x) = r.$$

Therefore, since $z \in B_d(y, \delta)$ was arbitrary, $B_d(y, \delta) \subseteq B_d(x, r)$. Hence $B_d(x, \epsilon)$ is open as $y \in B_d(x, \epsilon)$ was arbitrary. ∎

Now that we have a nice topology on each metric spaces, we desire to understand what these topologies look like. For some metric spaces, the topology is easy to understand.

Example 1.2.21. Let X be a non-empty set and let d be the discrete metric on X. Then for all $x \in X$ we have that

$$\{x\} = B_d\left(x, \frac{1}{2}\right).$$

Hence every singleton point is an open set. Since $\mathcal{T}_d$ is closed under arbitrary unions, we obtain that $\mathcal{T}_d = \mathcal{P}(X)$ so that $\mathcal{T}_d$ is the discrete topology. Hence the name of the metric!

However, some metric induced topologies are more difficult to determine. For example, on $\mathbb{R}^n$ we have several topologies; indeed we have one topology induced by each p-norm for each $p \in [1, \infty]$. How can we better understand what the topologies look like and how can we compare these topologies?

1.3 Bases

In order to have a better understanding and control over topologies, we desire to describe smaller collections of open sets that determine the entire topology. For example, we have seen with metric topologies that all open sets are unions of open balls, so provided we can understand the open balls we should be able to understand the entire topology. In particular, by distilling down the proof of Theorem 1.2.17 and Proposition 1.2.20 we obtain the following method of constructing a topology from a collection of sets with a specific properties.

Theorem 1.3.1. *Let X be a non-empty set and let $\mathcal{B} \subseteq \mathcal{P}(X)$ be such that*

(1) if $x \in X$ then there exists a $B \in \mathcal{B}$ such that $x \in B$, and

(2) if $x \in X$ and $B_1, B_2 \in \mathcal{B}$ are such that $x \in B_1 \cap B_2$, then there exists a $B_3 \in \mathcal{B}$ such that $x \in B_3$, $B_3 \subseteq B_1$, and $B_3 \subseteq B_2$.

Let $\mathcal{T}_\mathcal{B}$ be the set of all subsets U of X such that for all $x \in U$ there exists a $B \in \mathcal{B}$ such that $x \in B$ and $B \subseteq U$. Then $\mathcal{T}_\mathcal{B}$ is a topology on X such that $\mathcal{B} \subseteq \mathcal{T}$.

Proof. To see that $\mathcal{T}_\mathcal{B}$ is a topology, we must verify the three properties in Definition 1.1.1. It is clear by definition of $\mathcal{T}_\mathcal{B}$ that $\emptyset \in \mathcal{T}_\mathcal{B}$.

To see that $X \in \mathcal{T}_\mathcal{B}$ recall by property (1) that for each $x \in X$ there exists an $B_x \in \mathcal{B}$ such that $x \in B_x$. As $B_x \subseteq X$ by definition, we obtain that $X \in \mathcal{T}_\mathcal{B}$ by the definition of $\mathcal{T}_\mathcal{B}$.

Next suppose $\{U_\alpha\}_{\alpha \in I}$ is a set of elements of $\mathcal{T}_\mathcal{B}$. To see that $\bigcup_{\alpha \in I} U_\alpha \in \mathcal{T}_\mathcal{B}$, let $x \in \bigcup_{\alpha \in I} U_\alpha$ be arbitrary. Then there must be an $\alpha_0 \in I$ such that $x \in U_{\alpha_0}$. Since $U_{\alpha_0} \in \mathcal{T}_\mathcal{B}$, there exists a $B \in \mathcal{B}$ such that $x \in B$ and $B \subseteq U_{\alpha_0}$. Hence $B \subseteq U_{\alpha_0} \subseteq \bigcup_{\alpha \in I} U_\alpha$. As $x \in \bigcup_{\alpha \in I} U_\alpha$ was arbitrary, we obtain that $\bigcup_{\alpha \in I} U_\alpha \in \mathcal{T}_\mathcal{B}$ by definition.

To complete the proof that $\mathcal{T}_\mathcal{B}$ is topology, suppose $U_1, \ldots, U_n \in \mathcal{T}_\mathcal{B}$. To see that $\bigcap_{k=1}^n U_k \in \mathcal{T}_\mathcal{B}$, let $x \in \bigcap_{k=1}^n U_k$ be arbitrary. Hence $x \in U_k$ for all k so, as each $U_k \in \mathcal{T}_\mathcal{B}$, there exists a $B_k \in \mathcal{B}$ such that $x \in B_k$ and $B_k \subseteq U_k$. By applying property (2) recursively $n-1$ times, there exists a $B \in \mathcal{B}$ such that $x \in B$ and $B \subseteq B_k$ for all k. Hence $B \in \mathcal{B}$, $x \in B$, and $B \subseteq B_k \subseteq U_k$ for all k so that $B \subseteq \bigcap_{k=1}^n U_k$. Therefore, as $x \in X$ was arbitrary, $\bigcap_{k=1}^n U_k \in \mathcal{T}_\mathcal{B}$ as desired.

Finally, the fact that $\mathcal{B} \subseteq \mathcal{T}_\mathcal{B}$ follows from the definition of $\mathcal{T}_\mathcal{B}$. ∎

As subsets of the power set of a given set as described in Theorem 1.3.1 are useful in constructing topologies, we define the following.

Definition 1.3.2. Let X be a non-empty set. A *basis for a topology on* X is a collection of subsets $\mathcal{B} \subseteq \mathcal{P}(X)$ such that

(1) if $x \in X$ then there exists a $B \in \mathcal{B}$ such that $x \in B$, and

(2) if $x \in X$ and $B_1, B_2 \in \mathcal{B}$ are such that $x \in B_1 \cap B_2$, then there exists a $B_3 \in \mathcal{B}$ such that $B_3 \subseteq B_1$ and $B_3 \subseteq B_2$.

The topology $\mathcal{T}_\mathcal{B}$ on X from Theorem 1.3.1 is called the *topology generated by the basis* $\mathcal{B}$. Note that a set $U \subseteq X$ is open with respect to $\mathcal{T}_\mathcal{B}$ if and only if for every $x \in U$ there exists a $B \in \mathcal{B}$ such that $x \in B$ and $B \subseteq U$. Consequently $\mathcal{B} \subseteq \mathcal{T}_\mathcal{B}$.

Here is one example of how we can construct topologies via bases that turns out not to be a topology we have previously seen.

Example 1.3.3. Let

$$\mathcal{B} = \{[a,b) \mid a, b \in \mathbb{R}, a < b\}.$$

We claim that $\mathcal{B}$ is a basis for a topology on $\mathbb{R}$. To see this, it suffices to verify the two defining properties of being a basis from Definition 1.3.2. To being, notice if $x \in \mathbb{R}$ then $x \in [x, x+1) \in \mathcal{B}$. Hence the first property is satisfied. To see the second property, let $[a_1, b_1), [a_2, b_2) \in \mathcal{B}$ and $x \in \mathbb{R}$ such that $x \in [a_1, b_1) \cap [a_2, b_2)$ be arbitrary. Let

$$a = \max(\{a_1, a_2\}) \qquad \text{and} \qquad b = \min(\{b_1, b_2\})$$

and let $B = [a, b)$. Since $x \in [a_1, b_1) \cap [a_2, b_2)$, we see that $a \leq x < b$ so $B \in \mathcal{B}$ and $x \in B$. Furthermore, by construction, $B \subseteq [a_1, b_1) \cap [a_2, b_2)$. Hence, since $[a_1, b_1), [a_2, b_2) \in \mathcal{B}$ and $x \in \mathbb{R}$ were arbitrary, $\mathcal{B}$ is a basis for a topology on $\mathbb{R}$.

The topology $\mathcal{T}_L$ on $\mathbb{R}$ generated by the basis $\mathcal{B}$ is called the *lower limit topology on* $\mathbb{R}$.

The fact that $\mathcal{T}_L$ is not the same as the canonical topology on $\mathbb{R}$ will come from material in Section 1.5 where we show that 'limits' behave different in these topologies. In particular, we will see why we call $\mathcal{T}_L$ the lower limit topology. Alternatively, we know that $[a, b)$ is open in the lower limit topology, but is not open in the canonical topology.

Topologies generated by a basis are particularly nice since it is very simple to completely understand the entire topology via the basis elements based on the above and below descriptions of open sets.

Theorem 1.3.4. *Let X be a non-empty set and let $\mathcal{B}$ be a basis for a topology on X. Then*

$$\mathcal{T}_{\mathcal{B}} = \left\{ \bigcup_{B \in \mathcal{B}_0} B \,\middle|\, \mathcal{B}_0 \subseteq \mathcal{B} \right\}.$$

Proof. Notice, since $\mathcal{T}_{\mathcal{B}}$ is a topology and since $\mathcal{B} \subseteq \mathcal{T}_{\mathcal{B}}$, we know that for all $\mathcal{B}_0 \subseteq \mathcal{B}$ that

$$\bigcup_{B \in \mathcal{B}_0} B$$

is a union of elements of $\mathcal{T}_{\mathcal{B}}$ and thus in $\mathcal{T}_{\mathcal{B}}$. Hence

$$\mathcal{T}_{\mathcal{B}} \supseteq \left\{ \bigcup_{B \in \mathcal{B}_0} B \,\middle|\, \mathcal{B}_0 \subseteq \mathcal{B} \right\}.$$

To see the other inclusion, let $U \in \mathcal{T}_{\mathcal{B}}$ be arbitrary. By the definition of $\mathcal{T}_{\mathcal{B}}$, for each $x \in U$ there exists a $B_x \in \mathcal{B}$ such that $x \in B_x$ and $B_x \subseteq U$. Hence we see that

$$U = \bigcup_{x \in U} B_x.$$

Therefore, as $U \in \mathcal{T}_{\mathcal{B}}$ was arbitrary, we obtain that

$$\mathcal{T}_{\mathcal{B}} = \left\{ \bigcup_{B \in \mathcal{B}_0} B \,\middle|\, \mathcal{B}_0 \subseteq \mathcal{B} \right\}$$

as claimed. ∎

As bases for a topology give us multiple nice descriptions of the open sets for that topology, it is useful when given a topology to have a basis that generates the given topology. Thus to simplify this terminology, we define the following.

Definition 1.3.5. Let $(X, \mathcal{T})$ be a topological space. A set $\mathcal{B} \subseteq \mathcal{P}(X)$ is said to be a *basis for* $(X, \mathcal{T})$ if $\mathcal{B}$ is a basis for a topology on X and $\mathcal{T}_{\mathcal{B}} = \mathcal{T}$.

Remark 1.3.6. Of course, as $\mathcal{B} \subseteq \mathcal{T}_{\mathcal{B}}$ for any basis of a topology $\mathcal{B}$, for a set $\mathcal{B} \subseteq \mathcal{P}(X)$ to be a basis for a topology $\mathcal{T}$, it must be the case that $\mathcal{B} \subseteq \mathcal{T}$. Furthermore, by Theorem 1.3.1 and Theorem 1.3.4, we see that if $\mathcal{B}$ is a basis for $(X, \mathcal{T})$ then

(1) a set $U \subseteq X$ is open if and only if for every $x \in U$ there exists a $B \in \mathcal{B}$ such that $x \in B$ and $B \subseteq U$, and

(2) the open sets in $(X, \mathcal{T})$ are exactly the union of elements of $\mathcal{B}$.

Furthermore, it is not difficult to see that every topology is generated by a basis as the following example shows.

Example 1.3.7. Let $(X, \mathcal{T})$ be a topological space. Then $\mathcal{T}$ is a basis for $(X, \mathcal{T})$. Of course this is not the most useful basis as our goal is to better understand $\mathcal{T}$ by using a basis with as few elements as possible.

Of course, many of our previously discussed topologies have far nicer bases.

Example 1.3.8. Let X be a non-empty set and let $\mathcal{T}$ be the discrete topology on X. Then
$$\mathcal{B} = \{\{x\} \mid x \in X\}$$
is a basis for $(X, \mathcal{T})$. Indeed clearly $\mathcal{B} \subseteq \mathcal{T}$ as $\mathcal{T}$ is the discrete topology. Next clearly the first property of Definition 1.3.2 holds and the second property also clearly holds since the only way $x \in X$ and $B_1, B_2 \in \mathcal{B}$ are such that $x \in B_1 \cap B_2$ is if $B_1 = B_2 = \{x\} \in \mathcal{B}$. Hence $\mathcal{B}$ is a basis for $(X, \mathcal{T})$.

Of course, we have our motivating example.

Example 1.3.9. Let (X, d) be a metric space. Then the set $\mathcal{B}$ of all open balls forms a basis for $(X, \mathcal{T}_d)$. Indeed clearly $\mathcal{B} \subseteq \mathcal{T}$ and if $x \in X$ then $B_d(x, 1) \in \mathcal{T}_d$ so the first property of Definition 1.3.2 is satisfied. To see the second property of Definition 1.3.2 is satisfied, let $x \in X$ and $B_1, B_2 \in \mathcal{B}$ be arbitrary such that $x \in B_1 \cap B_2$. Then there exists points $x_1, x_2 \in X$ and $r_1, r_2 > 0$ such that $B_1 = B_d(x_1, r_1)$ and $B_2 = B_d(x_2, r_2)$. Thus, as $x \in B_1 \cap B_2$, we see that
$$d(x, x_1) < r_1 \qquad \text{and} \qquad d(x, x_2) < r_2.$$
Let
$$r = \min\{r_1 - d(x, x_1), r_2 - d(x, x_2)\}.$$
Then $r > 0$. By the same argument as in Proposition 1.2.20, we see that $B_d(x, r) \subseteq B_d(x_1, r_1)$ and $B_d(x, r) \subseteq B_d(x_2, r_2)$. Hence, as $x \in X$ and $B_1, B_2 \in \mathcal{B}$ were arbitrary, the second property of Definition 1.3.2 is satisfied so that $\mathcal{B}$ is a basis for $(X, \mathcal{T}_d)$.

Example 1.3.10. Let (X, d) be a metric space and let $\epsilon > 0$. The set $\mathcal{B}$ of all open balls with radius at most ϵ forms a basis for $(X, \mathcal{T}_d)$. Indeed the proof is identical to that of Example 1.3.9 with the additional restraint that all radii involved are at most ϵ.

Example 1.3.11. Let (X, d) be a metric space. The set $\mathcal{B}$ of all open balls with radius positive rational radii forms a basis for $(X, \mathcal{T}_d)$. Indeed the proof is identical to that of Example 1.3.9 with the additional restraint that all radii involved are rational. This is advantageous over Examples 1.3.9 and 1.3.10 as this basis only has a countable number of elements centred at each point.

Of course the above examples were expected as the metric topologies were our motivating example for how to construct a topology from a basis. However, when we constructed the metric topologies in Theorem 1.2.17 there was no reference to bases and the topologies were constructed by saying the open sets were those that for each point in them there was a 'ball' around that point contained in the open set. This leads us to an alternate characterization for a basis for a topological space. In particular, the following is superior to Definition 1.3.2 in checking that a collection of sets is a basis for a specific topology and is the converse to fact (1) in Remark 1.3.6.

Proposition 1.3.12. *Let $(X, \mathcal{T})$ be a topological space. Suppose $\mathcal{B} \subseteq \mathcal{T}$ has the property that for all $U \in \mathcal{T}$ and for all $x \in U$ there exists a $B \in \mathcal{B}$ such that $x \in B$ and $B \subseteq U$. Then $\mathcal{B}$ is a basis for $(X, \mathcal{T})$.*

Proof. To see that $\mathcal{B}$ is a basis for a topology on X, we will simply verify the two properties in Definition 1.3.2. To begin, let $x \in X$ be arbitrary. Then, as $X \in \mathcal{T}$, the assumptions of the proposition imply there exists a $B \in \mathcal{B}$ such that $x \in B$ and $B \subseteq X$. Hence, as $x \in X$ was arbitrary, the first assumption of Definition 1.3.2 has been verified.

To see the second property of Definition 1.3.2 holds, let $x \in X$ and $B_1, B_2 \in \mathcal{B}$ such that $x \in B_1 \cap B_2$ be arbitrary. As $\mathcal{B} \subseteq \mathcal{T}$, we see that $B_1, B_2 \in \mathcal{T}$ and thus $B_1 \cap B_2 \in \mathcal{T}$. Therefore, by the assumptions of the proposition there exists a $B_3 \in \mathcal{B}$ such that $x \in B_3$ and $B_3 \subseteq B_1 \cap B_2$. Hence, as $x \in X$ and $B_1, B_2 \in \mathcal{B}$ were arbitrary, the second property of Definition 1.3.2 has been verified. Thus $\mathcal{B}$ is a basis for a topology on X.

To see that $\mathcal{T} = \mathcal{T}_{\mathcal{B}}$, we first note that as $\mathcal{B} \subseteq \mathcal{T}$ and as $\mathcal{T}$ is closed under unions, Theorem 1.3.4 implies that

$$\mathcal{T}_{\mathcal{B}} = \left\{ \bigcup_{B \in \mathcal{B}_0} B \;\middle|\; \mathcal{B}_0 \subseteq \mathcal{B} \right\} \subseteq \mathcal{T}.$$

Conversely, if $U \in \mathcal{T}$ then the assumptions of the proposition imply that $U \in \mathcal{T}_{\mathcal{B}}$ by Definition 1.3.2. Hence $\mathcal{T} = \mathcal{T}_{\mathcal{B}}$ as desired. ∎

Using the above, we obtain our final characterization of a basis for a topological space which acts as a converse to Theorem 1.3.4.

Corollary 1.3.13. *Let $(X, \mathcal{T})$ be a topological space. Suppose $\mathcal{B} \subseteq \mathcal{T}$ has the property that for every $U \in T$ there exists a subset $\mathcal{B}_0 \subseteq \mathcal{B}$ such that $U = \bigcup_{B \in \mathcal{B}_0} B$. Then $\mathcal{B}$ is a basis for $(X, \mathcal{T})$.*

Proof. To prove this result, we will verify that the assumption of Proposition 1.3.12 holds. To see this, let $U \in \mathcal{T}$ and $x \in U$ be arbitrary. Then, by the assumptions of this corollary, there exists a subset $\mathcal{B}_0 \subseteq \mathcal{B}$ such that $U = \bigcup_{B \in \mathcal{B}_0} B$. Hence, as $x \in U$, there exists a $B_x \in \mathcal{B}_0$ such that $x \in B_x$ and $B_x \subseteq \bigcup_{B \in \mathcal{B}_0} B = U$. Therefore, as $U \in \mathcal{T}$ and $x \in U$ were arbitrary, the assumption of Proposition 1.3.12 holds. Hence the result follows. ∎

Throughout this course it will often be useful and more convenient to work with a basis for a topological space than the topology itself. For example, the following demonstrates how to use bases to determine when one topology is finer or coarser than another. In particular, we will often use the case that one of the bases for one of the topologies is the topology itself, which is valid by Example 1.3.7.

Theorem 1.3.14. *Let $\mathcal{T}$ and $\mathcal{T}'$ be topologies on a set X and let $\mathcal{B}$ and $\mathcal{B}'$ be bases for $\mathcal{T}$ and $\mathcal{T}'$ respectively. Then the following are equivalent:*

(i) $\mathcal{T}'$ is finer than $\mathcal{T}$.

(ii) For every $x \in X$ and $B \in \mathcal{B}$ such that $x \in B$ there exists a $B' \in \mathcal{B}'$ such that $x \in B'$ and $B' \subseteq B$.

Proof. First suppose that $\mathcal{T}'$ is finer that $\mathcal{T}$. Thus $\mathcal{T} \subseteq \mathcal{T}'$. To see that (ii) holds, let $x \in X$ and $B \in \mathcal{B}$ such that $x \in B$ be arbitrary. Since $\mathcal{B}$ is a basis for $\mathcal{T}$, $\mathcal{B} \subseteq \mathcal{T} \subseteq \mathcal{T}'$. Thus $B \in \mathcal{T}'$. Therefore, as $x \in B$, as $B \in \mathcal{T}'$, and as $\mathcal{B}'$ is a basis for $\mathcal{T}'$, we obtain that there exists a $B' \in \mathcal{B}'$ such that $x \in B'$ and $B' \subseteq B$. Therefore as $x \in X$ and $B \in \mathcal{B}$ were arbitrary, (ii) follows.

Conversely, suppose that (ii) holds. To see that T' is finer than $\mathcal{T}$, let $U \in \mathcal{T}$ be arbitrary. To see that $U \in \mathcal{T}'$, let $x \in U$ be arbitrary. As $U \in \mathcal{T}$, there exists a $B \in \mathcal{B}$ such that $x \in B$ and $B \subseteq U$. Then, by assumption (ii), there exist a $B' \in \mathcal{B}'$ such that $x \in B'$ and $B' \subseteq B \subseteq U$. Therefore $U \in \mathcal{T}'$ as $\mathcal{B}'$ is a basis for $\mathcal{T}'$. Hence, as $U \in \mathcal{T}$ was arbitrary, $\mathcal{T} \subseteq \mathcal{T}'$ so $\mathcal{T}'$ is finer than $\mathcal{T}$. ∎

One important use of Theorem 1.3.14 is that we can now compare topologies using bases.

Example 1.3.15. Fix $n \in \mathbb{N}$ and for each $p \in [1, \infty]$ let $\mathcal{T}_p$ denoted the topology on $\mathbb{K}^n$ induced by the p-norm. We claim that all the topologies $\mathcal{T}_p$ for $p \in [1, \infty]$ are equal. To see this, we will show $\mathcal{T}_p = \mathcal{T}_\infty$ for all $p \in [1, \infty)$. Thus, to proceed, let use fix $p \in [1, \infty)$.

Let $\vec{x} = (x_1, x_2, \ldots, x_n) \in \mathbb{K}^n$ be arbitrary. Then clearly

$$\max(\{|x_1|, |x_2|, \ldots, |x_n|\}) \leq \left(\sum_{k=1}^{n} |x_k|^p\right)^{\frac{1}{p}} \leq n^{\frac{1}{p}} \max(\{|x_1|, |x_2|, \ldots, |x_n|\}).$$

Hence, for all $\vec{x} \in \mathbb{K}^n$ we obtain that

$$\|\vec{x}\|_\infty \leq \|\vec{x}\|_p \leq n^{\frac{1}{p}} \|\vec{x}\|_\infty.$$

Therefore, for any $r > 0$ and $\vec{y} \in \mathbb{K}^n$ we see that

$$B_{\|\cdot\|_\infty}(\vec{y}, r) \subseteq B_{\|\cdot\|_p}\left(\vec{y}, n^{\frac{1}{p}} r\right) \subseteq B_{\|\cdot\|_\infty}\left(\vec{y}, n^{\frac{1}{p}} r\right).$$

Hence, as the above holds for any $r > 0$ and $\vec{y} \in \mathbb{K}^n$, as for any point x in any open ball B_1 in a metric space there exists a ball B_2 centred at x contained in B_1, and as the balls for each norm form a basis for their respective metric topologies, the first of the above inclusions together with Theorem 1.3.14 implies that $\mathcal{T}_p \subseteq \mathcal{T}_\infty$ and the second of the above inclusions together with Theorem 1.3.14 implies that $\mathcal{T}_\infty \subseteq \mathcal{T}_p$. Hence the claim follows.

Of course, to generate a topology from a basis, we need a basis. This leads to the question about how can we construct collections of sets that satisfy the assumptions of Definition 1.3.2. As the second property required in Definition 1.3.2 relates to the intersection of basis elements containing basis elements, one way to avoid this problem is by taking all intersections of the sets we want to use to form a basis. Thus we define the following object.

Definition 1.3.16. Let $(X, \mathcal{T})$ be a topological space. A *subbasis* for $(X, \mathcal{T})$ is a collection of subsets $\mathcal{S} \subseteq \mathcal{T}$ such that the set of all finite intersections of elements of $\mathcal{S}$ is a basis for $(X, \mathcal{T})$.

Of course, for a set $\mathcal{S}$ to be a subbasis of some topology $\mathcal{T}$ on X, it is necessary that

$$X = \bigcup_{S \in \mathcal{S}} S.$$

as for each $x \in X$ there must be a basis element containing X. In fact, this is the only restriction for a collection of sets to be a subbasis for some topology on X.

Theorem 1.3.17. *Let X be a non-empty set and let $\mathcal{S} \subseteq \mathcal{P}(X)$ be such that*

$$X = \bigcup_{S \in \mathcal{S}} S.$$

Let $\mathcal{B} \subseteq \mathcal{P}(X)$ be the set of all finite intersections of elements of $\mathcal{S}$. Then $\mathcal{B}$ is a basis for a topology on X for which $\mathcal{S}$ is a subbasis.

Proof. To see that $\mathcal{B}$ is a basis for a topology on X, we need only check the two conditions on a basis from Definition 1.3.2. For the first, let $x \in X$ be arbitrary. Since $X = \bigcup_{S \in \mathcal{S}} S$ there exists an $S_x \in \mathcal{S}$ such that $x \in S_x$. As $\mathcal{S} \subseteq \mathcal{B}$, the first property of being a basis holds for $\mathcal{B}$.

For the second property, let $x \in X$ and let $B_1, B_2 \in \mathcal{B}$ such that $x \in B_1 \cap B_2$ be arbitrary. Since $B_1, B_2 \in \mathcal{B}$, B_1 and B_2 are finite intersections of elements of $\mathcal{S}$. Hence $B_1 \cap B_2$ is a finite intersection of elements of $\mathcal{S}$ so that $B_1 \cap B_2 \in \mathcal{B}$. Therefore, as $x \in X$ and let $B_1, B_2 \in \mathcal{B}$ were arbitrary, the first property of being a basis holds for $\mathcal{B}$. Hence $\mathcal{B}$ is a basis for a topology on X. The fact that that $\mathcal{S}$ is a subbasis is a subbasis for $\mathcal{T}_{\mathcal{B}}$ is then trivial. ■

Subbases are not as desirable as bases as the description of the entire topology is far more difficult using subbases than bases and thus make subbases far more difficult to use thereby limiting their applications. However, subbases are excellent for constructing topologies as the conditions that are required are far simpler.

1.4 Constructing Topologies

Since bases and subbases are so great for constructing topologies, let us examine how we can construct new topologies from old topologies. The first such example of this comes from restricting a topology on a set to a subset of the set.

Lemma 1.4.1. *Let $(X, \mathcal{T})$ be a topological space and let $Y \subseteq X$ be non-empty. The set*

$$\mathcal{T}_Y = \{Y \cap U \mid U \in \mathcal{T}\}$$

is a topology on Y.

Proof. To verify that $\mathcal{T}_Y$ is a topology on Y, we verify Definition 1.1.1. First clearly $\mathcal{T}_Y \subseteq \mathcal{P}(Y)$ by construction. Next, as $\emptyset, X \in \mathcal{T}$, we obtain that $\emptyset = Y \cap \emptyset \in \mathcal{T}_Y$ and $Y = Y \cap X \in \mathcal{T}_Y$. The fact that $\mathcal{T}_Y$ is closed under unions and finite intersections then follows from the facts that

$$\bigcup_{\alpha \in I} (Y \cap U_\alpha) = Y \cap \left(\bigcup_{\alpha \in I} U_\alpha \right) \text{ and}$$
$$\bigcap_{\alpha \in I} (Y \cap U_\alpha) = Y \cap \left(\bigcap_{\alpha \in I} U_\alpha \right)$$

for all $\{U_\alpha\}_{\alpha \in I} \subseteq \mathcal{T}$. ∎

Definition 1.4.2. Let $(X, \mathcal{T})$ be a topological space and let $Y \subseteq X$ be non-empty. The *subspace topology on Y* is the topology

$$\mathcal{T}_Y = \{A \cap U \mid U \in \mathcal{T}\}.$$

In addition, the pair $(Y, \mathcal{T}_Y)$ is called a *subspace* of $(X, \mathcal{T})$.

Remark 1.4.3. The subspace topology is very useful when one only wants to consider a portion of a topological space. For example, we often want to consider subspaces of $\mathbb{R}$ such as $Y = [0, 1]$ for analytical reasons. However, one should be careful as open subsets of Y need not be open subset of $\mathbb{R}$. Indeed since $[0, 1) = [0, 1] \cap (-1, 1)$ we see that $[0, 1)$ is an open subset of Y in the subspace topology but is not an open subset of $\mathbb{R}$ as $0 \in [0, 1)$ yet no open interval centred at 0 is contained in $[0, 1)$. Thus we really do need to specify the topology and space we are looking at when talking about open sets!

Of course, it is not surprising that a basis for a topological space yields a basis for any subspace.

Proposition 1.4.4. *Let* $(X, \mathcal{T})$ *be a topological space, let* $\mathcal{B}$ *be a basis for* $(X, \mathcal{T})$*, and let* $Y \subseteq X$ *be non-empty. Then*

$$\mathcal{B}_Y = \{Y \cap B \mid B \in \mathcal{B}\}$$

is a basis for $(Y, \mathcal{T}_Y)$*.*

Proof. To prove this result, we will verify Proposition 1.3.12. As such, first notice that $\mathcal{B}_Y \subseteq \mathcal{T}_Y$ by the definition of the subspace topology. Next, let $U \in \mathcal{T}_Y$ and $x \in U$ be arbitrary. Thus $x \in Y$ and, by the definition of the subspace topology, there exists a $V \in \mathcal{T}$ such that $U = Y \cap V$. Since $x \in U$ we see that $x \in V$. Therefore, as $\mathcal{B}$ is a basis for $(X, \mathcal{T})$, Remark 1.3.6 implies that there exists a $B \in \mathcal{B}$ such that $x \in B$ and $B \subseteq V$. Therefore $Y \cap B \in \mathcal{B}_Y$, $x \in Y \cap B$, and $Y \cap B \subseteq Y \cap V = U$. Thus, as $U \in \mathcal{T}_Y$ and $x \in U$ were arbitrary, $\mathcal{B}_Y$ is a basis for $(T, \mathcal{T}_Y)$ by Proposition 1.3.12. ∎

Unsurprisingly, we can use Proposition 1.4.4 along with our knowledge of open balls in metric spaces to understand subspaces of metric spaces.

Proposition 1.4.5. *Let* (X, d) *be a metric space and let* $Y \subseteq X$ *be non-empty. Then the subspace topology on* Y *is induced by the metric* $d_Y : Y \times Y \to [0, \infty)$ *defined by*

$$d_Y(y_1, y_2) = d(y_1, y_2)$$

for all $y_1, y_2 \in Y$*.*

Proof. Since d is a metric and restricting the domain of d will yield a metric, d_Y is a metric. Notice for all $y \in Y$ and $r > 0$ that

$$B_{d_Y}(y, r) = Y \cap B_d(y, r).$$

Therefore, since $\{Y \cap B_d(y, r) \mid y \in Y, r > 0\}$ is a basis for the subspace topology on Y induced by (X, d) by Proposition 1.4.4, since $\{B_{d_Y}(y, r) \mid y \in Y, r > 0\}$ is a basis for (Y, d_Y) by definition, and since each basis completely determines the topology by Remark 1.3.6, the result follows. ∎

Furthermore, a subspace of a subspace is a subspace. More accurately put, we have the following.

Proposition 1.4.6. *Let* $(X, \mathcal{T})$ *be a topological space and let* $A \subseteq B \subseteq X$ *be arbitrary non-empty sets. Let* $\mathcal{T}_B$ *be the subspace topology on* B *inherited from* $(X, \mathcal{T})$*, let* $\mathcal{T}_A$ *be the subspace topology on* A *inherited from* $(X, \mathcal{T}_X)$*, and let* $\mathcal{T}_{A,B}$ *be the subspace topology on* A *inherited from* $(B, \mathcal{T}_B)$*. Then* $\mathcal{T}_{A,B} = \mathcal{T}_A$*.*

Proof. By definitions and since $B \cap A = A$, we have that

$$\begin{aligned} \mathcal{T}_{A,B} &= \{A \cap U \mid U \in \mathcal{T}_B\} \\ &= \{A \cap (B \cap V) \mid V \in \mathcal{T}_X\} \\ &= \{A \cap V \mid V \in \mathcal{T}_X\} = \mathcal{T}_A \end{aligned}$$

as desired. ∎

Since subspaces create a topology on a smaller set from a topology on a larger set, it is useful to think of the opposite; that is, can we construct topologies on larger sets from topologies on smaller sets? The simplest way to construct a larger set from two sets is to take their product. Unsurprisingly perhaps, taking the product of the topologies then yields a (basis for a) topology.

Proposition 1.4.7. *Let $(X, \mathcal{T})$ and $(Y, \mathcal{T}')$ be topological spaces. Then*

$$\mathcal{B}_\times = \{U \times V \mid U \in \mathcal{T}, V \in \mathcal{T}'\} \subseteq \mathcal{P}(X \times Y)$$

is a basis for a topology on $X \times Y$.

Proof. To see that $\mathcal{B}_\times$ is a basis for a topology on $X \times Y$, we simply verify Definition 1.3.2. Indeed clearly $\mathcal{B}_\times \subseteq \mathcal{P}(X \times Y)$. Moreover, notice for all $x \in X$ and $y \in Y$ that $x \times y \in X \times Y$ and $X \times Y \in \mathcal{B}_\times$. Hence the first condition of Definition 1.3.2 holds.

To see the second condition, let $x \times y \in X \times Y$ and $B_1, B_2 \in \mathcal{B}_\times$ such that $x \times y \in B_1 \cap B_2$ be arbitrary. By the definition of $\mathcal{B}_\times$ there exists $U_1, U_2 \in \mathcal{T}$ and $V_1, V_2 \in \mathcal{T}'$ such that $B_1 = U_1 \times V_1$ and $B_2 = U_2 \times V_2$. Since $U_1 \cap U_2 \in \mathcal{T}$ and $V_1 \cap V_2 \in \mathcal{T}'$ as $\mathcal{T}$ and $\mathcal{T}'$ are topologies, and since $B_1 \cap B_2 = (U_1 \cap U_2) \times (V_1 \cap V_2)$, we see that $B_1 \cap B_2 \in \mathcal{B}_\times$ so we may take $B_3 = B_1 \cap B_2$ in Definition 1.3.2. Therefore, as $x \times y \in X \times Y$ and $B_1, B_2 \in \mathcal{B}_\times$ were arbitrary, $\mathcal{B}_\times$ is a basis for a topology on $X \times Y$. ∎

Since we are taking the set of Cartesian products of the two topologies to form a topology on the product, this topology has an unsurprising name.

Definition 1.4.8. Let $(X, \mathcal{T})$ and $(Y, \mathcal{T}')$ be topological spaces. The *product topology* is the topology generated by the basis

$$\{U \times V \mid U \in \mathcal{T}, V \in \mathcal{T}'\} \subseteq \mathcal{P}(X \times Y).$$

Of course, the product of $\mathbb{R}$ with itself yields a topology on $\mathbb{R}^2$ that we have seen before.

Example 1.4.9. As $\mathbb{K}^2 = \mathbb{K} \times \mathbb{K}$, we can consider the product topology on $\mathbb{K}^2$ where each copy of $\mathbb{K}$ is equipped with the canonical topology. In this case, we know a basis of $\mathbb{K} \times \mathbb{K}$ consists of open sets of the form

$$U_1 \times U_2$$

where U_1 and U_2 are open subset of $\mathbb{K}$ with respect to the canonical topology. As each point in each such product contains a $\|\cdot\|_\infty$-ball in the product, and as each $\|\cdot\|_\infty$-ball is such a product, we easily obtain that the product topology on $\mathbb{K}^n$ is the same as the metric topologies seen in Example 1.3.15 by Theorem 1.3.14.

Perhaps unsurprisingly, we can repeat the proof of Proposition 1.4.7 to simplify the basis for the product topology.

Proposition 1.4.10. *Let $(X, \mathcal{T})$ and $(Y, \mathcal{T}')$ be topological spaces with bases $\mathcal{B}$ and $\mathcal{B}'$ respectively. Then the set*

$$\mathcal{B}_\times = \{B \times B' \mid B \in \mathcal{B}, B' \in \mathcal{B}'\}$$

is a basis for the product topology on $X \times Y$.

Proof. To see that $\mathcal{B}_\times$ is a basis for a topology on $X \times Y$, we will apply Proposition 1.3.12. To see this, let U be an arbitrary open subset of $X \times Y$ with respect to the product topology and let $x \times y \in U$ be arbitrary. By the definition of the product topology (Definition 1.4.8) there exists sets $U_X \in \mathcal{T}$ and $U_Y \in \mathcal{T}'$ such that $x \times y \in U_X \times U_Y$ and $U_X \times U_Y \subseteq U$. Thus $x \in U_X$ and $y \in U_Y$. Since $\mathcal{B}$ and $\mathcal{B}'$ are bases for $(X, \mathcal{T})$ and $(Y, \mathcal{T}')$ respectively, there exists $B \in \mathcal{B}$ and $B' \in \mathcal{B}'$ such that $x \in B$, $y \in B'$, $B \subseteq U_X$ and $B' \subseteq U_Y$. Hence $x \times y \in B \times B'$ and $B \times B' \subseteq U_X \times U_Y \subseteq Y$. Therefore, as U and $x \times y$ were arbitrary, $\mathcal{B}_\times$ is a basis for the product topology on $X \times Y$ by Proposition 1.3.12. ∎

Alternatively, we can consider the product topology via a subbasis.

Proposition 1.4.11. *Let $(X, \mathcal{T})$ and $(Y, \mathcal{T}')$ be topological spaces. Then*

$$\mathcal{S} = \{U \times Y \mid U \in \mathcal{T}\} \cup \{X \times V \mid V \in \mathcal{T}'\}$$

is a subbasis for the product topology on $X \times Y$.

Proof. Since finite intersections of elements of $\mathcal{S}$ yields the set

$$\mathcal{B}_\times = \{U \times V \mid U \in \mathcal{T}, V \in \mathcal{T}'\}$$

as $\mathcal{T}$ and $\mathcal{T}'$ are topologies and thus closed under finite intersections, and since $\mathcal{B}_\times$ is a basis for the product topology on $X \times Y$ by Definition 1.4.8, the result follows. ∎

The real reason we are considering both the basis and subbasis approaches to the product topology on $X \times Y$ above is that if we move to infinite products of sets, these two approaches actually differ. Therefore, we will obtain two different topologies on infinite products to study in this course. Thus we embark on describing these topologies. As such, we refer a reader unfamiliar with infinite products of sets is to Appendix A.2.

Lemma 1.4.12. *Let I be a non-empty set, let $\{(X_\alpha, \mathcal{T}_\alpha)\}_{\alpha\in I}$ be a non-empty indexed family of topological spaces, and let*

$$\mathcal{B} = \left\{ \prod_{\alpha\in I} U_\alpha \,\middle|\, U_\alpha \in \mathcal{T}_\alpha \right\}.$$

Then $\mathcal{B}$ is a basis for a topology on $\prod_{\alpha\in I} X_\alpha$.

Proof. To see that $\mathcal{B}$ is a basis for a topology on $\prod_{\alpha\in I} X_\alpha$, we simply verify Definition 1.3.2. Indeed clearly $\mathcal{B} \subseteq \mathcal{P}(\prod_{\alpha\in I} X_\alpha)$. Moreover, notice for all $(x_\alpha)_{\alpha\in I} \in \prod_{\alpha\in I} X_\alpha$ that $(x_\alpha)_{\alpha\in I} \in \prod_{\alpha\in I} X_\alpha$ and $\prod_{\alpha\in I} X_\alpha \in \mathcal{B}$. Hence the first condition of Definition 1.3.2 holds.

To see the second condition, let $(x_\alpha)_{\alpha\in I} \in \prod_{\alpha\in I} X_\alpha$ and $B_1, B_2 \in \mathcal{B}$ such that $(x_\alpha)_{\alpha\in I} \in B_1 \cap B_2$ be arbitrary. By the definition of $\mathcal{B}$ there exists $U_\alpha, V_\alpha \in \mathcal{T}_\alpha$ for all $\alpha \in I$ such that $B_1 = \prod_{\alpha\in I} U_\alpha$ and $B_2 = \prod_{\alpha\in I} V_\alpha$. Since $U_\alpha, V_\alpha \in \mathcal{T}_\alpha$ and since $\mathcal{T}_\alpha$ is a topology, we see that $U_\alpha \cap V_\alpha \in \mathcal{T}_\alpha$ for all $\alpha \in I$. Hence as $B_1 \cap B_2 = \prod_{\alpha\in I}(U_\alpha \cap V_\alpha)$, we see that $B_1 \cap B_2 \in \mathcal{B}$ so we may take $B_3 = B_1 \cap B_2$ in Definition 1.3.2. Therefore, as $(x_\alpha)_{\alpha\in I} \in \prod_{\alpha\in I} X_\alpha$ and $B_1, B_2 \in \mathcal{B}$ were arbitrary, $\mathcal{B}$ is a basis for a topology on $\prod_{\alpha\in I} X_\alpha$. ∎

Based on the above, we define the following.

Definition 1.4.13. Let I be a non-empty set and let $\{(X_\alpha, \mathcal{T}_\alpha)\}_{\alpha\in I}$ be a non-empty indexed family of topological spaces. The *box topology* on $\prod_{\alpha\in I} X_\alpha$ is the topology generated by the basis

$$\left\{ \prod_{\alpha\in I} U_\alpha \,\middle|\, U_\alpha \in \mathcal{T}_\alpha \right\}.$$

Of course, the above was derived from the basis approach to the product topology on $X \times Y$ and is equal to the product topology when I has two elements. We call the above the box topology as we are specifying a basis of generalized boxes (i.e. the product of open intervals in three dimensions is box-shaped). However, the subbasis approach to the product topology on $X \times Y$ is the correct approach to generalizing the product topology to infinitely products.

Definition 1.4.14. Let I be a non-empty set and let $\{(X_\alpha, \mathcal{T}_\alpha)\}_{\alpha\in I}$ be a non-empty indexed family of topological spaces. The *product topology* on $\prod_{\alpha\in I} X_\alpha$ is the topology generated by the subbasis

$$\mathcal{S} = \{\mathcal{S}_\beta \mid \beta \in I\}$$

where

$$\mathcal{S}_\beta = \left\{ \prod_{\alpha\in I} Y_\alpha \,\middle|\, Y_\alpha = X_\alpha \text{ if } \alpha \neq \beta, Y_\beta \in \mathcal{T}_\beta \right\}.$$

Of course, we should note that the set $\mathcal{S}$ described in Definition 1.4.14 is actually a subbasis for some topology on $\prod_{\alpha\in I} X_\alpha$, but this simply follows from Theorem 1.3.17. Furthermore, defining the subbasis immediately tells us a basis for the product topology.

Corollary 1.4.15. *Let I be a non-empty set and let $\{(X_\alpha, \mathcal{T}_\alpha)\}_{\alpha\in I}$ be a non-empty indexed family of topological spaces. The product topology on $\prod_{\alpha\in I} X_\alpha$ has as a basis the set of all sets of the form $\prod_{\alpha\in I} U_\alpha$ where $U_\alpha \in \mathcal{T}_\alpha$ and $U_\alpha = X_\alpha$ for all but a finite number of $\alpha \in I$.*

Proof. As the set of all finite intersections of the subbasis for the product topology described in Definition 1.4.14 is exactly the sets described here as $\mathcal{T}_\alpha$ is closed under finite intersections for all $\alpha \in I$, the result follows by the definition of a subbasis (Definition 1.3.16). ∎

Before getting to examples, we immediately note that can compare the box and product topologies and we can form bases for the box and product topologies in the fashion one would expect.

Corollary 1.4.16. *Let I be a non-empty set and let $\{(X_\alpha, \mathcal{T}_\alpha)\}_{\alpha\in I}$ be a non-empty indexed family of topological spaces. In general, the box topology on $\prod_{\alpha\in I} X_\alpha$ is finer than the product topology. In the case that I is finite, these two topologies coincide.*

Proof. As each basis element for the product topology is a basis element for the box topology, the box topology on $\prod_{\alpha\in I} X_\alpha$ is finer than the product topology by Theorem 1.3.14. In the case that I is finite, the bases for the product and box topologies agree so clearly these two topologies then coincide by Definition 1.3.5. ∎

Corollary 1.4.17. *Let I be a non-empty set, let $\{(X_\alpha, \mathcal{T}_\alpha)\}_{\alpha\in I}$ be a non-empty indexed family of topological spaces, and for each $\alpha \in I$ let $\mathcal{B}_\alpha$ be a basis for $(X_\alpha, \mathcal{T}_\alpha)$. Then $\prod_{\alpha\in I} \mathcal{B}_\alpha$ is a basis for the box topology on $\prod_{\alpha\in I} X_\alpha$. Similarly, the set of all sets of the form $\prod_{\alpha\in I} B_\alpha$ where $B_\alpha = X_\alpha$ for all but a finite number of $\alpha \in I$ and $B_\alpha \in \mathcal{B}_\alpha$ for all remaining indices is a basis for the product topology on $\prod_{\alpha\in I} X_\alpha$.*

Proof. To see that $\prod_{\alpha\in I} \mathcal{B}_\alpha$ is a basis for the box topology on $\prod_{\alpha\in I} X_\alpha$, we simply verify Definition 1.3.2. Indeed clearly $\prod_{\alpha\in I} \mathcal{B}_\alpha \subseteq \mathcal{P}(\prod_{\alpha\in I} X_\alpha)$. Moreover, notice for all $(x_\alpha)_{\alpha\in I} \in \prod_{\alpha\in I} X_\alpha$ that $x_\alpha \in X_\alpha$ so the fact that $\mathcal{B}_\alpha$ is a basis for $(X_\alpha, \mathcal{T}_\alpha)$ yields a $B_\alpha \in \mathcal{B}_\alpha$ such that $x_\alpha \in B_\alpha$ for all $\alpha \in I$. Hence $(x_\alpha)_{\alpha\in I} \in \prod_{\alpha\in I} B_\alpha$ and $\prod_{\alpha\in I} B_\alpha \in \prod_{\alpha\in I} \mathcal{B}_\alpha$. Hence the first condition of Definition 1.3.2 holds.

To see the second condition, let $(x_\alpha)_{\alpha\in I} \in \prod_{\alpha\in I} X_\alpha$ and $B_1, B_2 \in \prod_{\alpha\in I} \mathcal{B}_\alpha$ such that $(x_\alpha)_{\alpha\in I} \in B_1 \cap B_2$ be arbitrary. By the definition of $\prod_{\alpha\in I} \mathcal{B}_\alpha$ there exists $B_{1,\alpha}, B_{2,\alpha} \in \mathcal{B}_\alpha$ for all $\alpha \in I$ such that $B_1 =$

$\prod_{\alpha\in I} B_{1,\alpha}$ and $B_2 = \prod_{\alpha\in I} B_{2,\alpha}$. Thus, as $(x_\alpha)_{\alpha\in I} \in B_1 \cap B_2$, we see that $x_\alpha \in B_{1,\alpha} \cap B_{2,\alpha}$ for all $\alpha \in I$. Therefore, as $B_{1,\alpha}, B_{2,\alpha} \in \mathcal{B}_\alpha$ and since $\mathcal{B}_\alpha$ is a basis for $(X_\alpha, \mathcal{T}_\alpha)$, there exists a $B_{3,\alpha} \in \mathcal{B}_\alpha$ such that $x_\alpha \in B_{3,\alpha}$ and $B_{3,\alpha} \subseteq B_{1,\alpha} \cap B_{2,\alpha}$ for all $\alpha \in I$. Hence $B_3 = \prod_{\alpha\in I} B_{3,\alpha} \in \prod_{\alpha\in I} \mathcal{B}_\alpha$, $(x_\alpha)_{\alpha\in I} \in B_3$, and $B_3 \subseteq B_1 \cap B_2$. Therefore, as $(x_\alpha)_{\alpha\in I} \in \prod_{\alpha\in I} X_\alpha$ and $B_1, B_2 \in \prod_{\alpha\in I} \mathcal{B}_\alpha$ were arbitrary, $\prod_{\alpha\in I} \mathcal{B}_\alpha$ is a basis for the box topology on $\prod_{\alpha\in I} X_\alpha$.

For the product topology, let $\mathcal{B}$ be the set described in the statement. To see that $\mathcal{B}$ is a basis for the product topology on $\prod_{\alpha\in I} X_\alpha$, we simply verify Definition 1.3.2. Indeed clearly $\mathcal{B} \subseteq \mathcal{P}(\prod_{\alpha\in I} X_\alpha)$. Moreover, notice for all $(x_\alpha)_{\alpha\in I} \in \prod_{\alpha\in I} X_\alpha$ that $(x_\alpha)_{\alpha\in I} \in \prod_{\alpha\in I} X_\alpha \in \mathcal{B}$. Hence the first condition of Definition 1.3.2 holds.

To see the second condition, let $(x_\alpha)_{\alpha\in I} \in \prod_{\alpha\in I} X_\alpha$ and $B_1, B_2 \in \mathcal{B}$ such that $(x_\alpha)_{\alpha\in I} \in B_1 \cap B_2$ be arbitrary. By the definition of $\mathcal{B}$ there exists $B_{1,\alpha}, B_{2,\alpha} \in \mathcal{B}_\alpha$ for all $\alpha \in I$ such that $B_1 = \prod_{\alpha\in I} B_{1,\alpha}$, $B_2 = \prod_{\alpha\in I} B_{2,\alpha}$, and only a finite number of $B_{1,\alpha}$ and $B_{2,\alpha}$ are not equal to X_α over all $\alpha \in I$. Thus, as $(x_\alpha)_{\alpha\in I} \in B_1 \cap B_2$, we see that $x_\alpha \in B_{1,\alpha} \cap B_{2,\alpha}$ for all $\alpha \in I$. If $B_{1,\alpha} = X_\alpha$ or $B_{2,\alpha} = X_\alpha$, let $B_{3,\alpha} = B_{1,\alpha} \cap B_{2,\alpha}$ so that either $B_{3,\alpha} = X_\alpha$ or $B_{3,\alpha} \in \mathcal{B}_\alpha$. Otherwise $B_{1,\alpha}, B_{2,\alpha} \in \mathcal{B}_\alpha$ so, since $\mathcal{B}_\alpha$ is a basis for $(X_\alpha, \mathcal{T}_\alpha)$, there exists a $B_{3,\alpha} \in \mathcal{B}_\alpha$ such that $x_\alpha \in B_{3,\alpha}$ and $B_{3,\alpha} \subseteq B_{1,\alpha} \cap B_{2,\alpha}$ for all $\alpha \in I$. Hence $B_3 = \prod_{\alpha\in I} B_{3,\alpha} \in \mathcal{B}$ as $B_\alpha \in \mathcal{B}_\alpha$ for all but a finite number of $\alpha \in I$ and $B_\alpha = X_\alpha$ for all remaining indices, $(x_\alpha)_{\alpha\in I} \in B_3$, and $B_3 \subseteq B_1 \cap B_2$. Therefore, as $(x_\alpha)_{\alpha\in I} \in \prod_{\alpha\in I} X_\alpha$ and $B_1, B_2 \in \mathcal{B}$ were arbitrary, $\mathcal{B}$ is a basis for the product topology on $\prod_{\alpha\in I} X_\alpha$. ∎

In fact, we have already encountered some product and box topologies in this course.

Example 1.4.18. Let $n \in \mathbb{N}$ be arbitrary. Then $\mathbb{K}^n = \prod_{k\in\{1,\ldots,n\}} \mathbb{K}$ so we can consider the product topology on $\mathbb{K}^n$. In this case, we know from Corollary 1.4.17 that a basis for the product topology on $\mathbb{K}^n$ is

$$I_1 \times I_2 \times \cdots \times I_n$$

where each I_k is an open ball in $\mathbb{K}$ with respect to the absolute value. As each point in each such product is contained in a $\|\cdot\|_\infty$-ball that is contained in the product, and as each $\|\cdot\|_\infty$-ball is such a product, we easily obtain that the product topology on $\mathbb{K}^n$ is the same as the metric topologies seen in Example 1.3.15 by Theorem 1.3.14.

In fact, we can show that the box and product topologies can differ in some cases.

Example 1.4.19. Consider the set $\mathcal{F}(\mathbb{N}, \mathbb{R})$ of all functions from $\mathbb{N}$ to $\mathbb{R}$. Recall by Example 1.2.4 that we can equip $\mathcal{F}(\mathbb{N}, \mathbb{R})$ with the uniform metric and thus has a metric topology $\mathcal{T}_m$. However, $\mathcal{F}(\mathbb{N}, \mathbb{R}) = \prod_{n\in\mathbb{N}} \mathbb{R}$ so we can consider also consider the box topology $\mathcal{T}_b$ and the product topology $\mathcal{T}_p$ on $\mathcal{F}(\mathbb{N}, \mathbb{R})$. We will show that

$$\mathcal{T}_p \subsetneq \mathcal{T}_m \subsetneq \mathcal{T}_b.$$

To see that $\mathcal{T}_p \subseteq \mathcal{T}_m$, let $x = (x_n)_{n\geq 1} \in \prod_{n\in\mathbb{N}} \mathbb{R}$ and $U \in \mathcal{T}_p$ such that $x \in U$ be arbitrary. Then there exists an $\epsilon \in (0, 1)$ and a set of the form $B = \prod_{n\in\mathbb{N}} I_n$ such that $I_n = \mathbb{R}$ for all but a finite number of indices and $I_n = (x_n - \epsilon, x_n + \epsilon)$ for all other n. However, clearly the uniform ball centred at x of radius $\frac{\epsilon}{2}$ is contained B. Hence Theorem 1.3.14 implies that $\mathcal{T}_p \subseteq \mathcal{T}_m$.

To see that $\mathcal{T}_p \neq \mathcal{T}_m$, consider the set Y which is the open ball of uniform metric radius 1 around $(0)_{n\in\mathbb{N}}$. Thus $Y \in \mathcal{T}_m$. If Y was an element of $\mathcal{T}_p$, then there must be an element of $\mathcal{B}_p$ that is contained in Y by Remark 1.3.6. However, it is clear that no element of $\mathcal{B}_p$ is contained in Y as every element of $\mathcal{B}_p$ is of the form $\prod_{n\in\mathbb{N}} I_n$ with at least one I_n equal to $\mathbb{R}$ which yields an element of $\prod_{n\in\mathbb{N}} I_n$ that is distance 1 with respect to the uniform metric away from $(0)_{n\in\mathbb{N}}$. Hence $\mathcal{T}_p \subsetneq \mathcal{T}_m$.

To see that $\mathcal{T}_m \subseteq \mathcal{T}_b$, recall from Example 1.3.10 that the balls with respect to the uniform norm of radius at most 1 are a basis for $\mathcal{T}_m$. As a ball of uniform metric radius at most $\epsilon \in (0, 1)$ centred at $(x_n)_{n\in\mathbb{N}}$ contains $\prod_{n\in\mathbb{N}} \left(x_n - \frac{\epsilon}{2}, x_n + \frac{\epsilon}{2}\right) \in \mathcal{T}_b$, we obtain by Theorem 1.3.14 that $\mathcal{T}_m \subseteq \mathcal{T}_b$.

To see that $\mathcal{T}_m \neq \mathcal{T}_b$, consider the set $Z = \prod_{n\in\mathbb{N}}(-1, 1)$. Clearly $Z \in \mathcal{T}_b$ by the definition of the box topology (Definition 1.4.13). However, $Z \notin \mathcal{T}_m$. To see this, consider the element $x = \left(1 - \frac{1}{n}\right)_{n\in\mathbb{N}}$. Clearly $x \in Z$. However, there is no uniform metric ball around x that is contained in Z. Indeed for every $\epsilon > 0$, the 2ϵ uniform metric ball around x contains the point $\left(1 - \frac{1}{n} + \epsilon\right)_{n\in\mathbb{N}}$ which is not in Z as $\lim_{n\to\infty} 1 - \frac{1}{n} + \epsilon = 1 + \epsilon > 1$. Hence $Z \notin \mathcal{T}_m$ by Remark 1.3.6 so $\mathcal{T}_m \neq \mathcal{T}_b$.

Of course, a natural question is, "Which is better; the box topology or the product topology?" The answer unequivocally is the product topology. The reasons for this will be seen throughout the course in that the product topology will have incredibly nice properties whereas the box topology will mainly be used for counter examples. This is because the box topology is somewhat too fine to exhibit desirable topological properties. Properties of the product and box topologies will be investigated throughout the course.

1.5 Nets and Limits

Now that we have seen several topologies and how to study them, we return to the notion that the open sets should yield some information about how close points are in topological space. In particular, we can ask what it means for a collection of points to get ‘closer and closer’ to a given point in a topological space.

In metric spaces, the answer is the well-known concept of convergent sequences. In particular, the ϵ-N notion of a limit of a sequence of real numbers easily generalizes to metric spaces.

Definition 1.5.1. Let (X, d) be a metric space. It is said that a sequence $(x_n)_{n\geq 1}$ of elements of X *converges* to an element $x \in X$ if for every $\epsilon > 0$ there exists an $N \in \mathbb{N}$ such that $d(x, x_n) < \epsilon$ for all $n \geq N$; that is, for any threshold $\epsilon > 0$ eventually all of the x_n must be within ϵ of x. In this case x is called *a limit* of $(x_n)_{n\geq 1}$.

Although the notion of a limit of a sequence in metric spaces is stated in terms of distances, we can easily translated into topology.

Proposition 1.5.2. *Let (X, d) be a metric space, let $(x_n)_{n\geq 1}$ be a sequence in X, and let $x \in X$. The following are equivalent:*

(i) The sequence $(x_n)_{n\geq 1}$ converges to x in (X, d).

(ii) For every open set U containing x, there exists an $N \in \mathbb{N}$ such that $x_n \in U$ for all $n \geq N$.

Proof. First, suppose that $(x_n)_{n\geq 1}$ converges to x in (X, d). To see (ii), let U be an arbitrary open subset of (X, d) such that $x \in U$. By Theorem 1.2.17, there exists an $\epsilon > 0$ such that $x \in B_d(x, \epsilon) \subseteq U$. Therefore, since $(x_n)_{n\geq 1}$ converges to x, there exists an $N \in \mathbb{N}$ such that $d(x, x_n) < \epsilon$ for all $n \geq N$. Hence $x_n \in B_d(x, \epsilon) \subseteq U$ for all $n \geq N$. Therefore, as U was arbitrary, (ii) holds.

Conversely, suppose that (ii) holds. To see that $(x_n)_{n\geq 1}$ converges to x in (X, d), let $\epsilon > 0$ be arbitrary. Since $B_d(x, \epsilon)$ is an open set in X containing x, (ii) implies there exists an $N \in \mathbb{N}$ such that $x_n \in B_d(x, \epsilon)$ for all $n \geq N$. Hence $d(x, x_n) < \epsilon$ for all $n \geq N$. Therefore, as $\epsilon > 0$ was arbitrary, the proof is complete. ∎

In fact, convergent sequences completely determine the metric topology on a metric space as the following result shows.

Theorem 1.5.3. *Let X be a non-empty set, let d_1 and d_2 be two metrics on X, and let $\mathcal{T}_1$ and $\mathcal{T}_2$ be the topologies induced by d_1 and d_2 respectively. Then the following are equivalent:*

(i) If a sequence $(x_n)_{n\geq 1}$ converges to x in (X, d_1), then $(x_n)_{n\geq 1}$ converges to x in (X, d_2).

(ii) $\mathcal{T}_1$ is finer than $\mathcal{T}_2$.

Consequently, $\mathcal{T}_1 = \mathcal{T}_2$ if and only if (X, d_1) and (X, d_2) have the same convergent sequences with the same limits.

Proof. First, suppose that (i) holds. To see that (ii) holds, suppose to the contrary that $\mathcal{T}_1$ is not finer that $\mathcal{T}_2$. By Theorem 1.3.14 there must exists an $x \in X$ and a $U \in \mathcal{T}_2$ such that for all $V \in \mathcal{T}_1$ such that $x \in V$ it must be the case that $V \nsubseteq U$. Hence for all $n \in \mathbb{N}$ there must exists an $x_n \in B_{d_1}\left(x, \frac{1}{n}\right) \setminus U$ since $x \in B_{d_1}\left(x, \frac{1}{n}\right) \in \mathcal{T}_1$ for all n. As $d_1(x, x_n) < \frac{1}{n}$ for all $n \in \mathbb{N}$, we easily see that $(x_n)_{n\geq 1}$ converges to x in (X, d_1). Hence, by the assumption of (i), it must be the case that $(x_n)_{n\geq 1}$ converges to x in (X, d_2). However, since $U \in \mathcal{T}_2$, since $x \in U$, and since $x_n \notin U$ for all $n \in \mathbb{N}$, we have a contradiction to the fact that $(x_n)_{n\geq 1}$ converges to x in (X, d_2) by Proposition 1.5.2. Hence it must have been the case that (ii) holds.

Conversely, suppose that (ii) holds. To see that (i) holds, let $(x_n)_{n\geq 1}$ be a sequence that converges to x in (X, d_1). To see that $(x_n)_{n\geq 1}$ converges to x in (X, d_2), let $U \in \mathcal{T}_2$ be an arbitrary set such that $x \in U$. Since $\mathcal{T}_1$ is finer than $\mathcal{T}_2$, Theorem 1.3.14 implies that there exists a $V \in \mathcal{T}_1$ such that $x \in V \subseteq U$. Since $(x_n)_{n\geq 1}$ converges to x in (X, d_1), by Proposition 1.5.2 there exists an $N \in \mathbb{N}$ such that $x_n \in V \subseteq U$ for every $n \geq N$. Hence, as U was arbitrary, Proposition 1.5.2 implies that $(x_n)_{n\geq 1}$ converges to x in (X, d_2). ∎

As (ii) of Proposition 1.5.2 easily extends to the topological setting, we can easily discuss convergence of sequences in general topological spaces. Unfortunately, sequences are not enough to really tackle properties and results for arbitrary topological spaces. In particular, unlike Theorem 1.5.3 for metric spaces, sequences are not enough to characterize the topology. This follows as the proof of Theorem 1.5.3 breaks down in that we needed to construct a sequence when showing (i) implies (ii) and to do this we needed to use the metric. To be specific, we needed to use something like Example 1.3.11 which implies that there is a basis for any metric space topology that contains a countable number of balls centred at each point. It is this countability and the countable number of elements in a sequence that allows sequences to be useful in metric spaces.

Thus, in order to have a similar notion of convergence in an arbitrary topological space that is sufficient to deduce properties of the space, we need to generalize the notion of a sequence. To do this, we first need to generalize the structure and ordering on the natural numbers.

Definition 1.5.4. A *directed set* is a pair $(\Lambda, \leq)$ where Λ is a non-empty set and $\leq$ is a relation on Λ such that

(1) (reflexivity) $\lambda \leq \lambda$ for all $\lambda \in \Lambda$,

(2) (transitivity) if $\lambda_1, \lambda_2, \lambda_3 \in \Lambda$ are such that $\lambda_1 \leq \lambda_2$ and $\lambda_2 \leq \lambda_3$, then $\lambda_1 \leq \lambda_3$, and

(3) (existence of upper bounds) if $\lambda_1, \lambda_2 \in \Lambda$, then there exists a $\lambda_3 \in \Lambda$ such that $\lambda_1 \leq \lambda_3$ and $\lambda_2 \leq \lambda_3$.

The relation $\leq$ is sometimes called the *direction* of Λ.

As we are generalizing the order structure of the natural numbers, our first example is no surprise.

Example 1.5.5. The pair $(\mathbb{N}, \leq)$ where $\leq$ is the natural ordering on the natural numbers is easily seen to be a directed set.

Example 1.5.6. The pair $(\mathbb{R}, \leq)$ where $\leq$ is the natural ordering on the real numbers is easily seen to be a directed set.

Example 1.5.7. Let X be any non-empty set and let $\mathcal{F} \subseteq \mathcal{P}(X)$ be non-empty and closed under finite unions. For two sets $A, B \in \mathcal{F}$, we define $A \leq B$ if and only if $A \subseteq B$. Then $(\mathcal{F}, \leq)$ is a directed set. Indeed it is clear that $\leq$ is reflexive and transitive. Furthermore, if $A, B \in \mathcal{F}$, then $A \cup B \in \mathcal{F}$ has the property that $A \subseteq A \cup B$ so $A \leq A \cup B$, and $B \subseteq A \cup B$ so $B \leq A \cup B$. Hence $(\mathcal{F}, \leq)$ is a directed set by Definition 1.5.4.

Example 1.5.8. Let X be any non-empty set and let $\mathcal{F} \subseteq \mathcal{P}(X)$ be non-empty and closed under finite intersections. For two sets $A, B \in \mathcal{F}$, we define $A \leq B$ if and only if $B \subseteq A$. Then $(\mathcal{F}, \leq)$ is a directed set. Indeed it is clear that $\leq$ is reflexive and transitive. Furthermore, if $A, B \in \mathcal{F}$, then $A \cap B \in \mathcal{F}$ has the property that $A \cap B \subseteq A$ so $A \leq A \cap B$, and $A \cap B \subseteq B$ so $B \leq A \cap B$. Hence $(\mathcal{F}, \leq)$ is a directed set by Definition 1.5.4.

Of course, there are many more directed sets. For notational convenience, instead of writing $(\Lambda, \leq)$ for a direct set, we will often just say that Λ is a directed set provided there is no ambiguity for the direction relation which will then be denoted by $\leq$.

With the generalization of the ordering on $\mathbb{N}$, we can describe a generalization of the notion of a sequence.

Definition 1.5.9. A *net* is a function $F : \Lambda \to X$ where Λ is a direct set and X is a non-empty set. For notational convenience, we will use $(x_\lambda)_{\lambda \in \Lambda}$ to denote the net $F : \Lambda \to X$ where $F(\lambda) = x_\lambda$.

There are many examples of nets, some of which we are quite familiar with.

Example 1.5.10. Every sequence is a net. Indeed a sequence $(x_n)_{n \geq 1}$ can be realized as a net by taking the directed set $(\mathbb{N}, \leq)$ (where $\leq$ is the usual ordering of the natural numbers) and defining F on $\mathbb{N}$ by $F(n) = x_n$.

Example 1.5.11. Consider a closed interval $[a, b]$ and the collection $\mathcal{P}$ of all finite partitions of $[a, b]$; that is, all finite subsets $P = \{t_k\}_{k=0}^n \subseteq [a, b]$ such that

$$a = t_0 < t_1 < \cdots < t_n = b.$$

For two sets $P_1, P_2 \in \mathcal{P}$, if we define $P_1 \leq P_2$ if and only if $P_1 \subseteq P_2$, then $(\mathcal{P}, \leq)$ is a directed set by Example 1.5.7 as the collection of all finite partitions is closed under finite unions.

Let $f : [a, b] \to \mathbb{R}$ be a function. For each partition $P = \{t_k\}_{k=0}^n \in \mathcal{P}$ and each $1 \leq k \leq n$, choose $c_k \in [t_{k-1}, t_k]$ and define

$$S_P = \sum_{k=1}^{n} f(c_k)(t_k - t_{k-1}).$$

Then $(S_P)_{P\in\mathcal{P}}$ is a net of Riemann sums.

The next example is motivated by trying to take sums over uncountable sets.

Example 1.5.12. Let I be any infinite (and not necessarily countable) set. Let $\mathcal{F}$ be the set of all finite (non-empty) subsets of I. For two sets $F_1, F_2 \in \mathcal{F}$, if we define $F_1 \leq F_2$ if and only if $F_1 \subseteq F_2$, then $(\mathcal{F}, \leq)$ is a directed set by Example 1.5.7 as finite unions of finite sets are finite.

For each $\alpha \in I$, let $x_\alpha \in \mathbb{R}$ be non-negative. For each $F \in \mathcal{F}$, define

$$S_F = \sum_{\alpha\in F} x_\alpha,$$

which is well-defined as F is finite. Then $(S_F)_{F\in\mathcal{F}}$ is a net of all finite sums of $\{x_\alpha \mid \alpha \in I\}$.

Of course, our interest does not stem from the existence of nets as generalizations of sequences, but the properties and results that the notion of the convergence of a net will yield. Thus, building on the idea of using open sets to describe convergence in metric spaces, we generalize the notion of a convergent sequence for nets in arbitrary topological spaces.

Definition 1.5.13. Let $(X, \mathcal{T})$ be a topological space. A net $(x_\lambda)_{\lambda\in\Lambda}$ in X is said to *converge* to a point $x_0 \in X$ (or equivalently, x_0 is *a limit* of $(x_\lambda)_{\lambda\in\Lambda}$) if for every $U \in \mathcal{T}$ such that $x_0 \in U$ there exists a $\lambda_0 \in \Lambda$ such that $x_\lambda \in U$ for all $\lambda \geq \lambda_0$.

Before we get to examples, the notion of taking a set from the topology containing a specified point is occurring in greater and greater frequency. Thus, at this point, it is about time we gave it a name.

Definition 1.5.14. Let $(X, \mathcal{T})$ be a topological space. A subset $U \subseteq X$ is said to be a *neighbourhood of a point* $x \in X$ if $x \in U$ and $U \in \mathcal{T}$.

Remark 1.5.15. The term ‘neighbourhood’ comes from the notion that an open set containing a point x contains all points that are ‘geographically’ close to x. However, one must be careful with the term ‘neighbourhood’ in topology as many authors do not require a neighbourhood of a point to be open; they just require that a neighbourhood contains an open set containing the specified point. As we will want to be working with mainly open sets in this course, our definition is preferable.

Now onto examples. Of course, this is nowhere near an exhaustive list.

Example 1.5.16. It is clear that a sequence in a metric space converges to a point as a net if and only if it converges as a sequence to the same point.

Example 1.5.17. Consider the net $(S_P)_{P\in\mathcal{P}}$ from Example 1.5.11. If f is Riemann integrable, then $(S_P)_{P\in\mathcal{P}}$ converges and converges to $\int_a^b f(x)\,dx$. Indeed suppose f is integrable and let U neighbourhood of $\int_a^b f(x)\,dx$. Hence there exists an $\epsilon > 0$ such that

$$\left(\int_a^b f(x)\,dx - \epsilon, \int_a^b f(x)\,dx + \epsilon\right) \subseteq U.$$

By the definition of the Riemann integral, there exists a partition $P_0 \in \mathcal{P}$ such that if $U(f, P_0)$ is the upper Riemann sum of f corresponding to P_0 and $L(f, P_0)$ is the lower Riemann sum of f corresponding to P_0, then

$$L(f, P_0) \leq \int_a^b f(x)\,dx \leq U(f, P_0) < L(f, P_0) + \epsilon.$$

If $P \in \mathcal{P}$ and $P \geq P_0$, then P is a refinement of P_0 so

$$L(f, P_0) \leq L(f, P) \leq S_P \leq U(f, P) \leq U(f, P_0).$$

Hence

$$S_P \in \left(\int_a^b f(x)\,dx - \epsilon, \int_a^b f(x)\,dx + \epsilon\right) \subseteq U.$$

Therefore, as U was arbitrary, $(S_P)_{P\in\mathcal{P}}$ converges to $\int_a^b f(x)\,dx$.

Somewhat conversely, if every net from Example 1.5.11 converges and converges to the same number, then f is Riemann integrable. In fact, it is only required that the net of upper Riemann sum $(U_P)_{P\in P}$ and the net of lower Riemann sums $(L_P)_{P\in P}$ converge to the same number I. To see this, suppose $(U_P)_{P\in P}$ and $(L_P)_{P\in P}$ both converge to I and let $\epsilon > 0$ be arbitrary. Since $(U_P)_{P\in P}$ converges to I, there exists a $P_1 \in \mathcal{P}$ such that

$$U_P \in \left(I - \frac{\epsilon}{2}, I + \frac{\epsilon}{2}\right)$$

for all $P \geq P_1$. Similarly, since $(L_P)_{P\in P}$ converges to I, there exists a $P_2 \in \mathcal{P}$ such that

$$L_P \in \left(I - \frac{\epsilon}{2}, I + \frac{\epsilon}{2}\right)$$

for all $P \geq P_2$. Thus, if $P_0 = P_1 \cup P_2$, then $P_0 \in \mathcal{P}$, $P_0 \geq P_1$ and $P_0 \geq P_2$ so

$$U_{P_0}, L_{P_0} \in \left(I - \frac{\epsilon}{2}, I + \frac{\epsilon}{2}\right).$$

Hence, as $L_{P_0} \leq U_{P_0}$, we obtain that $U_{P_0} - L_{P_0} < \epsilon$. Therefore, as $\epsilon > 0$ was arbitrary, f is Riemann integrable.

Example 1.5.18. Consider the net $(S_F)_{F\in\mathcal{F}}$ from Example 1.5.12. Then $(S_F)_{F\in\mathcal{F}}$ converges if and only if

$$L = \sup\{S_F \mid F \in \mathcal{F}\}$$

is finite, in which case $(S_F)_{F\in\mathcal{F}}$ converges to L. Indeed suppose L is finite and let U be a neighbourhood of L. Then there exists an $\epsilon > 0$ such that

$$(L - \epsilon, L + \epsilon) \subseteq U.$$

By the definition of the supremum, there exists an $F_0 \in \mathcal{F}$ such that

$$L - \epsilon < S_{F_0} \leq L.$$

Hence, as $x_\alpha \geq 0$ for all $\alpha \in I$, we see that for all $F \in \mathcal{F}$ with $F \geq F_0$ that

$$L - \epsilon < S_{F_0} \leq S_F \leq L.$$

Hence $S_F \in U$ for all $F \geq F_0$. Therefore, as U was arbitrary, $(S_F)_{F\in\mathcal{F}}$ converges to L.

Conversely suppose that $L = \infty$. Hence for any $M \in \mathbb{R}$ there exists an $F_M \in \mathcal{F}$ such that $S_{F_M} \geq M$. To proceed by contradiction, suppose $(S_F)_{F\in\mathcal{F}}$ converges to some point $K \in \mathbb{R}$. Then there exists an $F_0 \in \mathcal{F}$ such that $S_F \in (K-1, K+1)$ for all $F \geq F_0$. Hence, as $F_0 \cup F_{K+1} \in \mathcal{F}$ and $F_0 \cup F_{K+1} \geq F_0$, we must have that

$$S_{F_0 \cup F_{K+1}} \in (K-1, K+1).$$

However

$$S_{F_0 \cup F_{K+1}} \geq S_{F_{K+1}} \geq K + 1$$

as $x_\alpha \geq 0$ for all $\alpha \in I$ so $S_{F_0 \cup F_{K+1}} \notin (K-1, K+1)$. Hence we have a contradiction as desired.

The above is quite useful in summing over uncountable sets. In particular, we define the *sum of* $\{x_\alpha \mid \alpha \in I\}$, denoted $\sum_{\alpha\in I} x_\alpha$, to be

$$\sum_{\alpha\in I} x_\alpha = \sup\{S_F \mid F \in \mathcal{F}\} \in [0, \infty].$$

Furthermore, if $\sum_{\alpha\in I} x_\alpha < \infty$ then for all $n \in \mathbb{N}$ we must have that $F_n = \left\{\alpha \in I \mid x_\alpha \geq \frac{1}{n}\right\}$ is finite for otherwise for each $m \in \mathbb{N}$ we can find a finite subset $F_{n,m} \subseteq F_n$ with m elements so that $S_{F_{n,m}} \geq \frac{m}{n}$ thereby yielding $\sum_{\alpha \in I} x_\alpha = \infty$. Therefore if $\sum_{\alpha \in I} x_\alpha < \infty$ then, each F_n is a finite set so

$$\bigcup_{n \geq 1} F_n = \{\alpha \in I \mid x_\alpha > 0\},$$

is countable. Thus, after removing all x_α that take the value 0, we can simply add a countable sum of non-negative numbers to determine the value of $\sum_{\alpha \in I} x_\alpha$.

Of course, as bases determine a topology, we need only check neighbourhoods of a point that come from a basis.

Lemma 1.5.19. *Let $(X, \mathcal{T})$ be a topological space and let $\mathcal{B}$ be a basis for $(X, \mathcal{T})$. A net $(x_\lambda)_{\lambda \in \Lambda}$ in X converges to $x_0 \in X$ if and only if for every $B \in \mathcal{B}$ such that $x_0 \in B$ there exists a $\lambda_0 \in \Lambda$ such that $x_\lambda \in B$ for all $\lambda \geq \lambda_0$.*

Proof. If $(x_\lambda)_{\lambda \in \Lambda}$ converges to x_0, then Definition 1.5.13 implies that for every $B \in \mathcal{B}$ such that $x_0 \in B$ there exists a $\lambda_0 \in \Lambda$ such that $x_\lambda \in B$ for all $\lambda \geq \lambda_0$ since $\mathcal{B} \subseteq \mathcal{T}$.

Conversely, suppose for every $B \in \mathcal{B}$ such that $x_0 \in B$ there exists a $\lambda_0 \in \Lambda$ such that $x_\lambda \in B$ for all $\lambda \geq \lambda_0$. To see that $(x_\lambda)_{\lambda \in \Lambda}$ converges to x_0, let U be an arbitrary neighbourhood of x_0. Then, as $\mathcal{B}$ is a basis for $(X, \mathcal{T})$, there exists a $B \in \mathcal{B}$ such that $x \in B \subseteq U$. Thus, by assumption, a $\lambda_0 \in \Lambda$ such that $x_\lambda \in B \subseteq$ for all $\lambda \geq \lambda_0$. Therefore, as U was arbitrary, the proof is complete. ∎

Of course, in Lemma 1.5.19, we need only information about the neighbourhoods of x_0. Consequently, we do not need to consider a basis for the entire space. In particular, we need only consider the following.

Definition 1.5.20. Let $(X, \mathcal{T})$ be a topological space and let $x \in X$. A set $\mathcal{B} \subseteq \mathcal{T}$ is said to be a *neighbourhood basis of* x if $x \in B$ for all $B \in \mathcal{B}$ and for all neighbourhoods U of x there exists a $B \in \mathcal{B}$ such that $x \in B \subseteq U$.

Theorem 1.5.21. *Let $(X, \mathcal{T})$ be a topological space, let $x_0 \in X$, and let $\mathcal{B}$ be a neighbourhood basis for x_0. A net $(x_\lambda)_{\lambda \in \Lambda}$ in X converges to x_0 if and only if for every $B \in \mathcal{B}$ such that $x_0 \in B$ there exists a $\lambda_0 \in \Lambda$ such that $x_\lambda \in B$ for all $\lambda \geq \lambda_0$.*

Proof. The proof of this result is identical to the proof of Lemma 1.5.19. ∎

Unsurprisingly, we can construct a basis from neighbourhood bases.

Proposition 1.5.22. *Let $(X, \mathcal{T})$ be a topological space and for each $x \in X$ let $\mathcal{B}_x$ be a neighbourhood basis for x. Then $\bigcup_{x \in X} \mathcal{B}_x$ is a basis for $(X, \mathcal{T})$.*

Proof. This follows immediately from the definition of a neighbourhood basis and Proposition 1.3.12. ∎

Returning to the notion of convergent nets, we can easily use bases to describe convergence in the lower limit, subspace, and product topologies. In particular, the following is the reason the lower limit topology has its name.

Proposition 1.5.23. *Let $\mathcal{T}_L$ be the lower limit topology on $\mathbb{R}$. A net $(x_\lambda)_{\lambda\in\Lambda}$ in $\mathbb{R}$ converges to a point x in $(\mathbb{R}, \mathcal{T}_L)$ if and only if for every $\epsilon > 0$ there exists an $\lambda_0 \in \Lambda$ such that $x \leq x_\lambda < x + \epsilon$ for all $\lambda \geq \lambda_0$.*

Proof. To begin, suppose a net $(x_\lambda)_{\lambda\in\Lambda}$ in $\mathbb{R}$ converges to a point x in $(\mathbb{R}, \mathcal{T}_L)$. To see the result, let $\epsilon > 0$ be arbitrary. Since $[x, x+\epsilon)$ is a neighbourhood of x, the definition of a convergent net implies there exists an $\lambda_0 \in \Lambda$ such that $x_\lambda \in [x, x+\epsilon)$ (that is, $x \leq x_\lambda < x + \epsilon$) for all $\lambda \geq \lambda_0$. Therefore, as $\epsilon > 0$ was arbitrary, the result holds.

Conversely, suppose $(x_\lambda)_{\lambda\in\Lambda}$ is an net in $\mathbb{R}$ and $x \in \mathbb{R}$ are such that for every $\epsilon > 0$ there exists an $\lambda_0 \in \Lambda$ such that $x \leq x_\lambda < x+\epsilon$ for all $\lambda \geq \lambda_0$. To see that $(x_\lambda)_{\lambda\in\Lambda}$ converges to x in $(\mathbb{R}, \mathcal{T}_L)$, let $B = [a, b) \in \mathcal{B}$ be an arbitrary element such that $x \in B$. Hence $a \leq x$ and $x < b$ so there exists an $\epsilon > 0$ such that $x < x + \epsilon < b$. Therefore, by the assumptions on $(x_\lambda)_{\lambda\in\Lambda}$, there exists a $\lambda_0 \in \Lambda$ such that $x \leq x_\lambda < x + \epsilon$ for all $\lambda \geq \lambda_0$. Hence

$$x_\lambda \in [x, x+\epsilon) \subseteq [a, b) = B \in \mathcal{B}.$$

Therefore, as $B \in \mathcal{B}$ was arbitrary, $(x_\lambda)_{\lambda\in\Lambda}$ converges to x in $(\mathbb{R}, \mathcal{T}_L)$ as desired. ∎

Proposition 1.5.24. *Let $(X, \mathcal{T})$ be a topological space, let $A \subseteq X$ be non-empty, let $\mathcal{T}_A$ be the subspace topology on A, let $(a_\lambda)_{\lambda\in\Lambda}$ be a net in A, and let $a \in A$. Then $(a_\lambda)_{\lambda\in\Lambda}$ converges to a in $(A, \mathcal{T}_A)$ if and only if $(a_\lambda)_{\lambda\in\Lambda}$ converges to a in $(X, \mathcal{T})$*

Proof. Since $a_\lambda \in A$ for all $\lambda \in A$, the result follows immediately by Definition 1.5.13 as the neighbourhoods of a in $(A, \mathcal{T}_A)$ are precisely the neighbourhoods of a in $(X, \mathcal{T})$ intersected with A. ∎

Theorem 1.5.25. *Let I be a non-empty set, let $\{(X_\alpha, \mathcal{T}_\alpha)\}_{\alpha\in I}$ be a non-empty indexed family of topological spaces, let $(f_\lambda)_{\lambda\in\Lambda}$ be a net in $\prod_{\alpha\in I} X_\alpha$, and let $(x_\alpha)_{\alpha\in I} \in \prod_{\alpha\in I} X_\alpha$. Then $(f_\lambda)_{\lambda\in\Lambda}$ converges to $(x_\alpha)_{\alpha\in I}$ when $\prod_{\alpha\in I} X_\alpha$ is equipped with the product topology if and only if $(f_\lambda(\alpha))_{\lambda\in\Lambda}$ converges to x_α in $(X_\alpha, \mathcal{T}_\alpha)$ for all $\alpha \in I$.*

Proof. Suppose that $(f_\lambda)_{\lambda\in\Lambda}$ converges to $(x_\alpha)_{\alpha\in I}$ when $\prod_{\alpha\in I} X_\alpha$ is equipped with the product topology. To see that $(f_\lambda(\alpha))_{\lambda\in\Lambda}$ converges to x_α in $(X_\alpha, \mathcal{T}_\alpha)$ for all $\alpha \in I$, fix $\alpha_0 \in I$ and let U_{α_0} be an arbitrary neighbourhood of x_{α_0} in $(X_{\alpha_0}, \mathcal{T}_{\alpha_0})$. For each $\alpha \in I \setminus \{\alpha_0\}$, let $U_\alpha = X_\alpha$. As $\prod_{\alpha\in I} U_\alpha$ is an element of the subbasis for the product topology on $\prod_{\alpha\in I} X_\alpha$ by Definition 1.4.14 and thus open, we easily see that $\prod_{\alpha\in I} U_\alpha$ is a neighbourhood of $(x_\alpha)_{\alpha\in I}$. Therefore, as $(f_\lambda)_{\lambda\in\Lambda}$ converges to $(x_\alpha)_{\alpha\in I}$ when $\prod_{\alpha\in I} X_\alpha$ is equipped with the product topology, there exists a $\lambda_0 \in \Lambda$ such that $f_\lambda \in \prod_{\alpha\in I} U_\alpha$ for all $\lambda \geq \lambda_0$. Hence $f_\lambda(\alpha_0) \in U_{\alpha_0}$ for all $\lambda \geq \lambda_0$. Therefore, as $\alpha_0 \in I$ and U_{α_0} where arbitrary, $(f_\lambda(\alpha))_{\lambda\in\Lambda}$ converges to x_α in $(X_\alpha, \mathcal{T}_\alpha)$ for all $\alpha \in I$.

Conversely, suppose $(f_\lambda(\alpha))_{\lambda\in\Lambda}$ converges to x_α in (X_α, T_α) for all $\alpha \in I$. Recall from Corollary 1.4.15 that the product topology on $\prod_{\alpha\in I} X_\alpha$ has as a basis $\mathcal{B}$ consisting of all sets of the form $\prod_{\alpha\in I} U_\alpha$ where $U_\alpha \in \mathcal{T}_\alpha$ and $U_\alpha = X_\alpha$ for all but a finite number of $\alpha \in I$. To see that $(f_\lambda)_{\lambda\in\Lambda}$ converges to $(x_\alpha)_{\alpha\in I}$, let $\prod_{\alpha\in I} U_\alpha$ be an arbitrary element of $\mathcal{B}$ that is a neighbourhood of $(x_\alpha)_{\alpha\in I}$. Hence U_α is a neighbourhood of x_α for all $\alpha \in I$ and

$$\{\alpha \in I \mid U_\alpha \neq X_\alpha\} = \{\alpha_1, \alpha_2, \ldots, \alpha_n\}$$

for some $n \in \mathbb{N}$. Since for each $k \in \{1, \ldots, n\}$ we know that $(f_\lambda(\alpha_k))_{\lambda\in\Lambda}$ converges to x_{α_k} in $(X_{\alpha_k}, T_{\alpha_k})$, there exists a $\lambda_k \in \Lambda$ such that $f_\lambda(\alpha_k) \in U_{\alpha_k}$ for all $\lambda \geq \lambda_k$. Luckily, by the properties of a direct set, there exists a $\lambda' \in \lambda$ such that $\lambda' \geq \lambda_k$ for all $k \in \{1, \ldots, n\}$. Hence $f_\lambda(\alpha_k) \in U_{\alpha_k}$ for all $\lambda \geq \lambda'$ and for all $k \in \{1, \ldots, n\}$. Since $f_\lambda(\alpha) \in X_\alpha = U_\alpha$ for all $\alpha \in I \setminus \{\alpha_1, \alpha_2, \ldots, \alpha_n\}$, we obtain that $(f_\lambda)_{\lambda\in\Lambda} \in \prod_{\alpha\in I} U_\alpha$ for all $\lambda \geq \lambda'$. Therefore, as $\prod_{\alpha\in I} U_\alpha$ was arbitrary, Lemma 1.5.19 implies the result. ∎

Remark 1.5.26. One reason the product topology is superior to the box topology is that 'if' direction of Theorem 1.5.25 fails when the box topology is used. In fact, it even fails if sequences are used! Indeed consider $\prod_{m\in\mathbb{N}} \mathbb{R}$ equipped with the box topology and consider the sequence $(f_n)_{n\geq 1} \in \prod_{m\in\mathbb{N}} \mathbb{R}$ where

$$f_n(m) = \begin{cases} 0 & \text{if } m < n \\ 1 & \text{if } m \geq n \end{cases}$$

and the element $f \in \prod_{m\in\mathbb{N}} \mathbb{R}$ where $f(m) = 0$ for all $m \in \mathbb{N}$. It is elementary to see that $(f_n(m))_{n\geq 1}$ converges to $f(m) = 0$ for every $m \in \mathbb{N}$. However, we claim that $(f_n)_{n\geq 1}$ does not converge f in the box topology. To see this, consider the set

$$U = \prod_{m\in\mathbb{N}} \left(-\frac{1}{2}, \frac{1}{2}\right),$$

which is open in the box topology and clearly then a neighbourhood of f. However, as $f_n(n) = 1$ for all $n \in \mathbb{N}$, we see that $f_n \notin U$ for all $n \in \mathbb{N}$. Hence $(f_n)_{n\geq 1}$ does not converge f in the box topology.

An alternative way to show that Theorem 1.5.25 fails when the product topology is changed to the box topology is to analyze what convergent nets imply about the topology. In particular, convergent nets completely determine the topology thereby generalizing Theorem 1.5.3 to nets and non-metric topological spaces.

Theorem 1.5.27. *Let X be a non-empty set and let $\mathcal{T}$ and $\mathcal{T}'$ be two topologies on X. Then $\mathcal{T}$ is finer than $\mathcal{T}'$ if and only if whenever $(x_\lambda)_{\lambda\in\Lambda}$ is a net that converges to x in $(X, \mathcal{T})$, then $(x_\lambda)_{\lambda\in\Lambda}$ converges to x in $(X, \mathcal{T}')$.*

Consequently, if $(X, \mathcal{T})$ and $(X, \mathcal{T}')$ have exactly the same nets converging to the same points, then $\mathcal{T} = \mathcal{T}'$.

Proof. If $\mathcal{T}$ is finer than $\mathcal{T}'$, then $\mathcal{T}' \subseteq \mathcal{T}$. It is then clear that if $(x_\lambda)_{\lambda\in\Lambda}$ is a net that converges to x in $(X, \mathcal{T})$, then $(x_\lambda)_{\lambda\in\Lambda}$ converges to x in $(X, \mathcal{T}')$ by the definition of a convergent net.

Conversely, suppose whenever $(x_\lambda)_{\lambda\in\Lambda}$ is a net that converges to x in $(X, \mathcal{T})$, then $(x_\lambda)_{\lambda\in\Lambda}$ converges to x in $(X, \mathcal{T}')$. To proceed by contradiction, suppose there exists a set $U \in \mathcal{T}'$ such that $U \notin \mathcal{T}$. By Theorem 1.3.14 there exists an $x_0 \in U$ such that for each $\mathcal{T}$-neighbourhood V of x_0, $V \setminus U$ is non-empty.

Let

$$\Lambda = \{V \subseteq X \mid V \text{ is a } \mathcal{T}\text{-neighourhood of } x_0\}.$$

As Λ is closed under finite intersections, if for $V_1, V_2 \in \Lambda$ we define $V_1 \leq V_2$ if $V_2 \subseteq V_1$, then $(\Lambda, \leq)$ is a direct set by Example 1.5.8.

For each $V \in \Lambda$, let $x_V \in V \setminus U$ (note we are using the Axiom of Choice A.2.4 here). We claim that $(x_V)_{V\in\Lambda}$ is a net that converges to x_0 in $(X, \mathcal{T})$ but does not converge to x_0 in $(X, \mathcal{T}')$ thereby yielding a contradiction. To see that $(x_V)_{V\in\Lambda}$ is a net that converges to x_0 in $(X, \mathcal{T})$, let V_0 be an arbitrary $\mathcal{T}$-neighbourhood x_0. Then for all $V \geq V_0$ we have that $x_V \in V \subseteq V_0$. Hence $(x_V)_{V\in\Lambda}$ is a net that converges to x_0 in $(X, \mathcal{T})$ by Definition 1.5.13. To see that $(x_V)_{V\in\Lambda}$ does not converge to x_0 in $(X, \mathcal{T}')$, we simply note that U is a $\mathcal{T}'$-neighbourhood of x_0 but $x_V \notin U$ for all $V \in \Lambda$. Hence we have obtained a contradiction thereby finishing the proof. ∎

Of course Theorem 1.5.27 immediately implies once we know the box and product topologies differ (one example of which was demonstrated in Example 1.4.19) that there exists a net that converges in the product topology but not in the box topology. This then implies that the 'if' direction of Theorem 1.5.25 must fail for the box topology.

Now that we have seen that the convergence of nets determine the topology, we return the our initial claim that sequences are not enough to understand topological spaces. In the proof Theorem 1.5.27 we really needed to use a net based on the directed set of all neighbourhoods of a point. Of course, if we were dealing with a metric space we could have used the $\frac{1}{n}$-balls centred at the point in the construction and thus used a sequence instead of a net. In particular, the following demonstrates that it is enough to consider sequence in 'nice' topological spaces thereby generalizing Theorem 1.5.3 as much as we can.

Theorem 1.5.28. *Let X be a non-empty set and let $\mathcal{T}$ and $\mathcal{T}'$ be two topologies on X. Suppose that for each point $x \in X$ that both $\mathcal{T}$ and $\mathcal{T}'$ have* *a countable neighbourhood basis of x. Then $\mathcal{T}$ is finer than $\mathcal{T}'$ if and only if whenever $(x_n)_{n\geq 1}$ is a sequence that converges to x in $(X,\mathcal{T})$, then $(x_n)_{n\geq 1}$ converges to x in $(X,\mathcal{T}')$. Consequently, if $(X,\mathcal{T})$ and $(X,\mathcal{T}')$ have exactly the same sequences converging to the same points, then $\mathcal{T}=\mathcal{T}'$.*

Proof. If $\mathcal{T}$ is finer than $\mathcal{T}'$, then $\mathcal{T}'\subseteq\mathcal{T}$. It is then clear that if $(x_n)_{n\geq 1}$ is a sequence that converges to x in $(X,\mathcal{T})$, then $(x_n)_{n\geq 1}$ converges to x in $(X,\mathcal{T}')$ by the definition of a convergent sequence.

Conversely, suppose whenever $(x_n)_{n\geq 1}$ is a sequence that converges to x in $(X,\mathcal{T})$, then $(x_n)_{n\geq 1}$ converges to x in $(X,\mathcal{T}')$. To proceed by contradiction, suppose there exists a set $U\in\mathcal{T}'$ such that $U\notin\mathcal{T}$. By Theorem 1.3.14 there exists an $x_0\in U$ such that for each $\mathcal{T}$-neighbourhood V of x_0, $V\setminus U$ is non-empty.

By assumption there exists a countable $\mathcal{T}$-neighbourhood basis $\{W_n\}_{n=1}^{\infty}$ of x_0. For each $n\in\mathbb{N}$, let

$$V_n=\bigcap_{k=1}^{n} W_k.$$

Then clearly V_n is a $\mathcal{T}$-neighbourhood of x_0 and $V_{n+1}\subseteq V_n$ for all $n\in\mathbb{N}$. Furthermore, we claim that $\{V_n\}_{n=1}^{\infty}$ is a $\mathcal{T}$-neighbourhood basis for x_0. To see this, suppose that V is a $\mathcal{T}$-neighbourhood of x_0. Since $\{W_n\}_{n=1}^{\infty}$ is a $\mathcal{T}$-neighbourhood basis of x_0, there exists an $m\in\mathbb{N}$ such that $x_0\in W_m\subseteq V$. Hence $x_0\in V_m\subseteq W_m\subseteq V$. Thus $\{V_n\}_{n=1}^{\infty}$ is a $\mathcal{T}$-neighbourhood basis for x_0.

For each $n\in\mathbb{N}$, let $x_n\in V_n\setminus U$, which exists as $V_n\setminus U$ is non-empty from above. We claim that $(x_n)_{n\geq 1}$ converges to x_0 in $(X,\mathcal{T})$ but does not converge to x_0 in $(X,\mathcal{T}')$ thereby yielding a contradiction. To see that $(x_n)_{n\geq 1}$ converges to x_0 in $(X,\mathcal{T})$, let V_0 be an arbitrary $\mathcal{T}$-neighbourhood x_0. Since $\{V_n\}_{n=1}^{\infty}$ is a $\mathcal{T}$-neighbourhood basis for x_0, there exists an $N\in\mathbb{N}$ such that $V_N\subseteq V_0$. Thus, for all $n\geq N$ we have that

$$x_n\in V_n\subseteq V_N\subseteq V_0.$$

Therefore, as V_0 was arbitrary, $(x_n)_{n\geq 1}$ converges to x_0 in $(X,\mathcal{T})$. To see that $(x_n)_{n\geq 1}$ does not converge to x_0 in $(X,\mathcal{T}')$, we simply note that U is a $\mathcal{T}'$-neighbourhood of x_0 but $x_n\notin U$ for all $n\in\mathbb{N}$. Hence we have obtained a contradiction thereby proving the result. ∎

However, in general, the level of abstraction of nets is required as there are two different topologies on a single space that have the same convergent sequences.

Example 1.5.29. Consider $\ell_1(\mathbb{R})$ and recall that $\|\cdot\|_1$ is a norm on $\ell_1(\mathbb{R})$ and thus induces a topology, which will be denoted $\mathcal{T}_n$. There is another topology on $\ell_1(\mathbb{R})$ known as the weak topology, which is a topology based on the dual space of $\ell_1(\mathbb{R})$ (namely $\ell_\infty(\mathbb{R})$). It is possible to show that the weak topology is not a topology induced by a norm and thus $\mathcal{T}_w \neq \mathcal{T}_n$. In addition, it is possible to show that $\mathcal{T}_n \subsetneq \mathcal{T}_w$ and that every sequence that converges in $\mathcal{T}_n$ converges in $\mathcal{T}_w$. Thus sequences are not enough to determine a topology! We leave the details of these facts to Appendix B.2

A clever observer at this point would have likely noticed that we have only defined 'a limit and not 'the limit' of a net when we defined when a net converges to a point. This is because, in a general topological space, a net can converge to multiple points so the 'the' in 'the limit' no longer make sense. This is even true if we consider sequences in general topological spaces as the following example demonstrates.

Example 1.5.30. Consider the set $X = \{a, b, c\}$ and the topology

$$\mathcal{T} = \{\emptyset, \{a\}, \{b, c\}, X\}$$

(i.e. see diagram 6 of Example 1.1.4). It is not difficult to see that a sequence $(x_n)_{n\geq 1}$ in X converges to a if and only if there exists an $N \in \mathbb{N}$ such that $x_n = a$ for all $n \geq N$ as $\{a\}$ is a neighbourhood of a. However, $(x_n)_{n\geq 1}$ in X converges to b if and only if there exists an $N \in \mathbb{N}$ such that $x_n \in \{b, c\}$ for all $n \geq N$ as the only open sets containing b are X and $\{b, c\}$. Similarly $(x_n)_{n\geq 1}$ in X converges to c if and only if there exists an $N \in \mathbb{N}$ such that $x_n \in \{b, c\}$ for all $n \geq N$. Thus there are several sequences in X that converge to both b and c.

The reason the above example does not yields unique limits is that there are not enough open sets to distinguish the points. The correct notion in order for there to be unique limits is the following.

Definition 1.5.31. A topological space $(X, \mathcal{T})$ is said to be *Hausdorff* (equivalently $(X, \mathcal{T})$ is a *Hausdorff space*) if for all $x, y \in X$ where $x \neq y$ there exists sets $U, V \in \mathcal{T}$ such that $x \in U$, $y \in V$, and $U \cap V = \emptyset$.

Example 1.5.32. The trivial topology on a set with at least two points is not Hausdorff as the only open sets are the empty set and the entire set.

Example 1.5.33. The discrete topology on any set is Hausdorff as every singleton is an open set.

Example 1.5.34. Let X be finite. The only topology on X that is Hausdorff is the discrete topology. Indeed suppose $\mathcal{T}$ is a Hausdorff topology on X and fix a point $x \in X$. For each point $y \in X \setminus \{y\}$ there exists an open set $U_y \in \mathcal{T}$ such that $x \in U_y$ but $y \notin U_y$. Then, as $X \setminus \{x\}$ is finite, we see that

$$\{x\} = \bigcap_{y \in X \setminus \{x\}} U_y \in \mathcal{T}.$$

Therefore, as $x \in X$ was arbitrary, every singleton from X is in $\mathcal{T}$. Therefore, as X is finite, $\mathcal{T}$ must be the discrete topology.

Example 1.5.35. The cofinite topology on an infinite set is not Hausdorff as the intersection of any two non-empty open sets in the cofinite topology on an infinite set must contain an infinite number of points. Similarly the cocountable topology on an uncountable set is not Hausdorff as the intersection of any two non-empty open sets in the cocountable topology on an uncountable set must contain an uncountable number of points. However, the cofinite topology on a finite set and the cocountable topology on a countable set are Hausdorff as every singleton is open (and thus the topologies are discrete in this case).

Example 1.5.36. The metric topology on a metric space (X, d) is Hausdorff. Indeed given two points $x, y \in X$ with $x \neq y$, let $\delta = \frac{1}{2}d(x, y)$. Then $B_d(x, \delta)$ and $B_d(y, \delta)$ are disjoint open sets one of which contains x and the other of which contains y. Hence the topology is Hausdorff by definition. Consequently, any non-Hausdorff topology is not induced by a metric.

Example 1.5.37. The lower limit topology on $\mathbb{R}$ is Hausdorff. To see this, let $a, b \in \mathbb{R}$ be such that $a < b$. Then $U = [a, b)$ and $V = [b, \infty)$ are open sets in the lower limit topology such that $a \in U$, $b \in V$, and $U \cap V = \emptyset$. Thus, as $a, b \in \mathbb{R}$ were arbitrary, the lower limit topology is Hausdorff.

Example 1.5.38. A subspace of any Hausdorff space is Hausdorff. This follows directly from the definition of a Hausdorff space and the description of the open subsets in the subspace topology (i.e. the open sets are simple the intersection of open sets with the subspace).

Example 1.5.39. The product and box topologies of Hausdorff spaces are Hausdorff. This follows directly from the description of the open sets in these topologies. To be specific, given two elements of the product $\prod_{\alpha \in I} X_\alpha$ of Hausdorff spaces, they differ at one value of α, say $\alpha_0 \in I$. Thus we can find disjoint open sets in $(X_{\alpha_0}, \mathcal{T}_{\alpha_0})$ that separate these two values and by taking the product of these open sets with X_α for all $\alpha \neq \alpha_0$, the desired open sets separating the two elements of the product have been found.

As advertised, Hausdorff spaces have unique limits.

Theorem 1.5.40. *Let $(X, \mathcal{T})$ be a Hausdorff space. If a net $(x_\lambda)_{\lambda \in \Lambda}$ converges to two points $x_1, x_2 \in X$, then $x_1 = x_2$.*

Proof. Suppose to the contrary that there exists a net $(x_\lambda)_{\lambda \in \Lambda}$ that converges to two points $x_1, x_2 \in X$ where $x_1 \neq x_2$. As $(X, \mathcal{T})$ is Hausdorff, there exist $U, V \in \mathcal{T}$ such that $x_1 \in U$, $x_2 \in V$, and $U \cap V = \emptyset$. As $(x_\lambda)_{\lambda \in \Lambda}$ converges to x_1, there exists a $\lambda_1 \in \Lambda$ such that $x_\lambda \in U$ for all $\lambda \geq \lambda_1$. Similarly as $(x_\lambda)_{\lambda \in \Lambda}$ converges to x_2, there exists a $\lambda_2 \in \Lambda$ such that $x_\lambda \in V$ for all $\lambda \geq \lambda_2$. However, by the properties of directed sets, there exists a $\lambda_3 \in \Lambda$ such that $\lambda_1 \leq \lambda_3$ and $\lambda_2 \leq \lambda_3$. Hence the above yields $x_{\lambda_3} \in U \cap V$ which contradicts the fact that $U \cap V = \emptyset$. Hence the result follows. ∎

In particular, for Hausdorff spaces, we can define limits.

Definition 1.5.41. Let $(X, \mathcal{T})$ be a Hausdorff space and let $(x_\lambda)_{\lambda\in\Lambda}$ be a net in X that converges in X. The unique point that $(x_\lambda)_{\lambda\in\Lambda}$ converges to in X is called the *limit of* $(x_\lambda)_{\lambda\in\Lambda}$ and is denoted $\lim_{\lambda\in\Lambda} x_\lambda$.

In fact, the only topological spaces that have unique limits for every converging net are Hausdorff spaces.

Theorem 1.5.42. *Let $(X, \mathcal{T})$ be a topological space such that every convergent net in $(X, \mathcal{T})$ converges to exactly one point. Then $(X, \mathcal{T})$ is Hausdorff.*

Proof. Let $(X, \mathcal{T})$ be a topological space such that every convergent net in $(X, \mathcal{T})$ converges to exactly one point. Suppose to the contrary that that $(X, \mathcal{T})$ is not Hausdorff. Then there exist points $x, y \in X$ such that for every neighbourhood U of x and neighbourhood V of y, $U \cap V \neq \emptyset$.

Consider the set

$$\Lambda = \{(U, V) \mid U, V \in \mathcal{T} \text{ are such that } x \in U \text{ and } y \in V\}.$$

For $(U_1, V_1), (U_2, V_2) \in V$, we define $(U_1, V_1) \leq (U_2, V_2)$ if and only if $U_2 \subseteq U_1$ and $V_2 \subseteq V_1$. We claim that $(\Lambda, \leq)$ is a directed set. Indeed, clearly $\leq$ is reflexive and transitive. Furthermore, if $(U_1, V_1), (U_2, V_2) \in V$, then by taking $U_3 = U_1 \cap U_2$ and $V_3 = V_1 \cap V_2$, we easily see that $(U_3, V_3) \in \Lambda)$, $(U_1, V_1) \leq (U_3, V_3)$, and $(U_2, V_2) \leq (U_3, V_3)$. Hence $(\Lambda, \leq)$ is a directed set.

For each $(U, V) \in \Lambda$, choose a $z_{(U,V)} \in U \cap V$, which exists by assumption (note we are using the Axiom of Choice A.2.4 here). Hence $(z_{(U,V)})_{(U,V)\in\Lambda}$ is a net. We claim that $(z_{(U,V)})_{(U,V)\in\Lambda}$ converges to both x and y. Indeed if U is an arbitrary neighbourhood of x, then for all $(U', V') \geq (U, X)$ we see that

$$z_{(U',V')} \in U' \cap V' \subseteq U \cap X = U.$$

Hence $(z_{(U,V)})_{(U,V)\in\Lambda}$ converges to x. Similarly, if V is an arbitrary neighbourhood of x, then for all $(U', V') \geq (X, V)$ we see that

$$z_{(U',V')} \in U' \cap V' \subseteq X \cap V = V.$$

Hence $(z_{(U,V)})_{(U,V)\in\Lambda}$ converges to y. As this contradicts the fact that every convergent net in $(X, \mathcal{T})$ converges to exactly one point, the proof is complete. ∎

In general, asking that a space is Hausdorff is a very mild condition in that it is simply asking that we can separate any two distinct points with open sets. However, the fact that nets in Hausdorff spaces have unique limits is very useful for the study of spaces that are Hausdorff. In particular, trying to prove results for arbitrary topological spaces can often be difficult or impossible as they need not have enough structure. Thus, we will often impose conditions like being Hausdorff on certain topological spaces in order to be able to prove certain results, which will then only apply to certain collections of topological spaces.

If we add the additional assumption that we are in a normed linear space, then we can even see that limits behave well with respect to the vector space operations.

Proposition 1.5.43. *Let $(V, \|\cdot\|)$ be a normed linear space. Suppose $(\vec{v}_\lambda)_{\lambda\in\Lambda}$ and $(\vec{w}_\lambda)_{\lambda\in\Lambda}$ are two nets indexed by the same directed set such that*

$$\lim_{\lambda\in\Lambda} \vec{v}_\lambda = \vec{v} \qquad \textit{and} \qquad \lim_{\lambda\in\Lambda} \vec{w}_\lambda = \vec{w}$$

for some $\vec{v}, \vec{w} \in V$. Then for all $\alpha \in \mathbb{K}$,

$$\lim_{\lambda\in\Lambda} \alpha\vec{v}_\lambda + \vec{w} = \alpha\vec{v} + \vec{w}.$$

Proof. Fix an $\alpha \in \mathbb{K}$ and let $\epsilon > 0$. Since

$$\lim_{\lambda\in\Lambda} \vec{v}_\lambda = \vec{v} \qquad \text{and} \qquad \lim_{\lambda\in\Lambda} \vec{w}_\lambda = \vec{w}$$

there exists $\lambda_1, \lambda_2 \in \Lambda$ such that

$$\|\vec{v} - \vec{v}_\lambda\| < \epsilon \text{ for all } \lambda \geq \lambda_1 \qquad \text{and} \qquad \|\vec{w} - \vec{w}_\lambda\| < \epsilon \text{ for all } \lambda \geq \lambda_2.$$

By the properties of directed sets, there exists a $\lambda_0 \in \Lambda$ such that $\lambda_0 \geq \lambda_1$ and $\lambda_0 \geq \lambda_2$. Hence, for all $\lambda \geq \lambda_0$, we obtain that

$$\begin{aligned} \|(\alpha\vec{v} + \vec{w}) - (\alpha\vec{v}_\lambda + \vec{w}_\lambda)\| &= \|\alpha(\vec{v} - \vec{v}_\lambda) + (\vec{w} - \vec{w}_\lambda)\| \\ &\leq |\alpha|\, \|\vec{v} - \vec{v}_\lambda\| + \|\vec{w} - \vec{w}_\lambda\| \\ &< (|\alpha| + 1)\epsilon. \end{aligned}$$

Therefore, as $\alpha \in \mathbb{K}$ is fixed, the result follows. ∎

To finish off this section, we recall one useful tool in undergraduate analysis is the ability to take subsequences. For nets, things are a little more delicate, but will be equally useful.

Definition 1.5.44. Let X be a non-empty set, let $(\Lambda, \leq)$ and $(M, \leq_0)$ be two directed sets, and let $F : \Lambda \to X$ be a net. A *subnet* of F directed by $(M, \leq_0)$ is the composition $F \circ \varphi : M \to X$ where $\varphi : M \to \Lambda$ is such that

(1) (increasing) if $\mu_1, \mu_2 \in M$ are such that $\mu_1 \leq_0 \mu_2$, then $\varphi(\mu_1) \leq \varphi(\mu_2)$, and

(2) (cofinal) for each $\lambda \in \Lambda$ there exists a $\mu \in M$ such that $\lambda \leq \varphi(\mu)$.

Remark 1.5.45. Note it is elementary to see that a subnet of a net is a net. In particular, if the net F is denoted by $(x_\lambda)_{\lambda\in\Lambda}$, we will often use $(x_{\lambda_\mu})_{\mu\in M}$ to denote a subnet where $\varphi(\mu) = \lambda_\mu$.

Subnets can be a little tricky.

Example 1.5.46. A subnet of a sequence need not be a subsequence. Indeed consider the sequence $(x_n)_{n\geq 1}$ and consider the directed set $(\mathbb{R}, \leq)$. Then if we define $\varphi : \mathbb{R} \to \mathbb{N}$ by

$$\varphi(x) = \begin{cases} 1 & \text{if } x < 1 \\ n & \text{if } x \in (n-1, n] \end{cases},$$

then φ is increasing and cofinal. However, clearly $(x_{n_\mu})_{\mu\in\mathbb{R}}$ is not a subsequence of $(x_n)_{n\geq 1}$.

Example 1.5.47. Let $(x_\lambda)_{\lambda\in\Lambda}$ be a net in a topological space. Choose $\lambda_1 \in \Lambda$. Then there exists a $\lambda_2 \in \Lambda$ such that $\lambda_2 \geq \lambda_1$. By repetition, we can obtain a sequence $(\lambda_n)_{n\geq 1}$ of elements of Λ that are increasing. However $(x_{\lambda_n})_{n\geq 1}$ need not be a subnet of $(x_\lambda)_{\lambda\in\Lambda}$ as it need not be cofinal (even if $\lambda_k \neq \lambda_{k+1}$). Note this is quite different than the situation with sequences.

However, as with sequences, subnets of convergent nets still converge.

Proposition 1.5.48. *Let $(X, \mathcal{T})$ be a topological space, let $(x_\lambda)_{\lambda\in\Lambda}$ be a net in X, and let $(x_{\lambda_\mu})_{\mu\in M}$ be a subnet of $(x_\lambda)_{\lambda\in\Lambda}$. If $(x_\lambda)_{\lambda\in\Lambda}$ converges to some point $x \in (X, \mathcal{T})$, then $(x_{\lambda_\mu})_{\mu\in M}$ converges to x in $(X, \mathcal{T})$.*

Proof. Suppose $(x_\lambda)_{\lambda\in\Lambda}$ converges to some point $x \in (X, \mathcal{T})$. To see that $(x_{\lambda_\mu})_{\mu\in M}$ converges to x in $(X, \mathcal{T})$, let U be an arbitrary neighbourhood of x. As $(x_\lambda)_{\lambda\in\Lambda}$ converges to x, there exists an $\lambda_0 \in \Lambda$ such that $x_\lambda \in U$ for all $\lambda \geq \lambda_0$. Due to the properties of subnets, there exists a $\mu_0 \in M$ such that $\lambda_{\mu_0} \geq \lambda_0$. Hence, by the properties of subnets, if $\mu \geq \mu_0$ then $\lambda_\mu \geq \lambda_{\mu_0} \geq \lambda_0$ and thus $x_{\lambda_\mu} \in U$. Therefore, as U was arbitrary, $(x_{\lambda_\mu})_{\mu\in M}$ converges to x in $(X, \mathcal{T})$. ∎

1.6 Sets and Points

With the completion of our basic understanding of nets, we can now turn out attention to types of points and sets inside topological spaces. These various types of points and sets occur regularly throughout topology and will be of incredible use in this course. Most of these notions are generalizations of known types of sets and points in metric spaces. In particular, the following type of sets are well known.

Definition 1.6.1. Let $(X, \mathcal{T})$ be a topological space. A set $F \subseteq X$ is said to be *closed* if $X \setminus F \in \mathcal{T}$.

Example 1.6.2. Every closed interval $[a, b]$ is a closed subset of $\mathbb{R}$ with its canonical topology as $\mathbb{R} \setminus [a, b] = (-\infty, a) \cup (b, \infty)$ is the union of two open sets and thus is open.

Example 1.6.3. The set $[0, 1)$ is neither open nor closed when $\mathbb{R}$ is equipped with its canonical topology. Indeed $[0, 1)$ is not open as there is no neighbourhood of 0 that is contained in $[0, 1)$. Similarly $[0, 1)$ is not closed as $\mathbb{R} \setminus [0, 1) = (-\infty, 0) \cup [1, \infty)$ is not open as there is no neighbourhood of 1 that is contained in $\mathbb{R} \setminus [0, 1)$. Thus it is possible that sets are neither open nor closed.

Example 1.6.4. Given a topological space $(X, \mathcal{T})$, the sets $\emptyset$ and X are always closed as $X \setminus \emptyset = X$ and $X \setminus X = \emptyset$ are open.

Example 1.6.5. In the discrete topology, every set is closed as every set is open so the complement of every set is open.

Example 1.6.6. In the cofinite topology, the closed sets are exactly the entire space and the set of finite subsets. Similarly, in the cocountable topology, the closed sets are exactly the entire space and the set of countable subsets.

Example 1.6.7. Let (X, d) be a metric space. Given an $x \in X$ and an $r > 0$, the *closed d-ball of radius r centred at x*, denoted $B_d[x, r]$, is the set

$$B_d[x, r] = \{y \in X \mid d(x, y) \leq r\}.$$

Any closed ball in any metric space is closed. Indeed to see that $B_d[x, r]$ is closed, let $y \in X \setminus B_d[x, r]$ be arbitrary. Then $d(x, y) > r$. Thus $B_d(y, d(x, y) - r)$ is an open set in (X, d). Furthermore, notice if $z \in B_d(y, d(x, y) - r)$ then $d(z, y) < d(x, y) - r$ so

$$d(x, z) \geq d(x, y) - d(y, z) > d(x, y) - (d(x, y) - r) = r$$

and thus $z \notin B_d[x, r]$. Hence $B_d(y, d(x, y) - r)$ is an open set containing y that is contained in $X \setminus B_d[x, r]$. Therefore, as $y \in X \setminus B_d[x, r]$ was arbitrary, $X \setminus B_d[x, r]$ is open and thus $B_d[x, r]$ is closed.

Example 1.6.8. If $(X, \mathcal{T})$ is a Hausdorff space, then every singleton is closed. Indeed let $x \in X$ be arbitrary. As $(X, \mathcal{T})$ is Hausdorff, for each $y \in Y$ there exists a $U_y \in \mathcal{T}$ such that $y \in U_y$ but $x \notin U_y$. Thus

$$X \setminus \{x\} = \bigcup_{y \in X \setminus \{x\}} U_y \in \mathcal{T}.$$

Hence $\{x\}$ is closed.

As the notion of a topological space immediately invokes properties on open sets, we immediately have the following properties of closed sets by taking complements and using De Morgan's Laws.

Proposition 1.6.9. *Let $(X, \mathcal{T})$ be a topological space. Then:*

(1) $\emptyset$ and X are closed sets.

(2) If $\{F_\alpha\}_{\alpha\in I}$ are closed sets in $(X, \mathcal{T})$, then $\bigcap_{\alpha\in I} F_\alpha$ is closed in $(X, \mathcal{T})$.

(3) If $\{F_\alpha\}_{\alpha\in I}$ are closed sets in $(X, \mathcal{T})$ and I is finite, then $\bigcup_{\alpha\in I} F_\alpha$ is closed in $(X, \mathcal{T})$.

Proof. Simply apply Definition 1.1.1 and De Morgan's Laws. ∎

Example 1.6.10. In any Hausdorff space, any finite union of points is closed as Example 1.6.8 shows singleton points are closed and Proposition 1.6.9 concludes finite unions of closed sets are closed.

Example 1.6.11. Let $P_0 = [0, 1]$. Construct P_1 from P_0 by removing the open interval of length $\frac{1}{3}$ from the middle of P_0 (i.e. $P_1 = [0, \frac{1}{3}] \cup [\frac{2}{3}, 1]$). Then construct P_2 from P_1 by removing the open intervals of length $\frac{1}{3^2}$ from the middle of each closed subinterval of P_1. Subsequently, having constructed P_n, construct P_{n+1} by removing the open intervals of length $\frac{1}{3^{n+1}}$ from the middle of each of the 2^n closed subintervals of P_n. Specifically, P_n is the union of the 2^n closed intervals of the form

$$\left[\sum_{k=1}^{n} \frac{a_k}{3^k}, \frac{1}{3^n} + \sum_{k=1}^{n} \frac{a_k}{3^k}\right]$$

where $a_1, \ldots, a_n \in \{0, 2\}$.

The set

$$\mathcal{C} = \bigcap_{n\geq 1} P_n$$

is known as the *Cantor set*. The Cantor set is closed by Proposition 1.6.9 being the intersection of finite unions of closed sets. In fact, it can be shown that $\mathcal{C}$ is uncountable.

When we restrict to subspaces of topological spaces, the closed subsets are easy to understand.

Lemma 1.6.12. *Let $(Y, \mathcal{T}_Y)$ be a subspace of a topological space $(X, \mathcal{T})$. A subset $A \subseteq Y$ is closed in $(Y, \mathcal{T}_Y)$ if and only if $A = Y \cap F$ where F is a closed set in $(X, \mathcal{T})$.*

Proof. First suppose $A = Y \cap F$ where F is a closed set in $(X, \mathcal{T})$. As F is closed in $(X, \mathcal{T})$, $V = X \setminus F \in \mathcal{T}$. Hence $U = Y \cap V$ is open in $(Y, \mathcal{T}_Y)$ so

$$Y \setminus U = \{y \in Y \mid y \notin U\} = \{y \in Y \mid y \notin V\} = A$$

is closed.

Conversely, suppose $A \subseteq Y$ is closed in $(Y, \mathcal{T}_Y)$. Then $U = Y \setminus A$ is open in $(Y, \mathcal{T}_Y)$. By the definition of the open subsets of a subspace, there exists a $V \in \mathcal{T}$ such that $U = Y \cap V$. Hence $F = X \setminus V$ is closed in $(X, \mathcal{T})$ and

$$Y \cap F = \{y \in Y \mid y \notin V\} = \{y \in Y \mid y \notin U\} = A.$$

Hence the result is complete. ∎

Example 1.6.13. Consider $Y = (0, 2)$ with the subspace topology inherited from the canonical topology on $\mathbb{R}$. Then

$$(0, 1] = Y \cap [-1, 1]$$

is closed in the subspace topology even though $(0, 1]$ is not closed in $\mathbb{R}$.

The reason closed sets are awesome is due to their relations with limits of nets.

Theorem 1.6.14. *Let $(X, \mathcal{T})$ be a topological space and let $F \subseteq X$. Then the following are equivalent:*

(i) F is a closed set in $(X, \mathcal{T})$.

(ii) Whenever $(x_\lambda)_{\lambda \in \Lambda}$ is a net such that $x_\lambda \in F$ for all $\lambda \in \Lambda$ that converges to a point $x_0 \in X$, then $x_0 \in F$.

Proof. To begin, suppose F is a closed set in $(X, \mathcal{T})$ and that $(x_\lambda)_{\lambda \in \Lambda}$ is a net such that $x_\lambda \in F$ for all $\lambda \in \Lambda$ that converges to a point $x_0 \in X$. Suppose to the contrary that $x_0 \notin F$. Then $x_0 \in X \setminus F$. As F is closed, $X \setminus F$ is open so $x_0 \in X \setminus F$ and the definition of a convergent net implies there exists a $\lambda_0 \in \Lambda$ such that $x_\lambda \in X \setminus F$ for all $\lambda \geq \lambda_0$. As this contradicts the fact that $x_\lambda \in F$ for all $\lambda \in \Lambda$, we have a contradiction. Hence $x_0 \in F$ as desired.

Conversely, suppose that whenever $(x_\lambda)_{\lambda \in \Lambda}$ is a net such that $x_\lambda \in F$ for all $\lambda \in \Lambda$ that converges to a point $x_0 \in X$, then $x_0 \in F$. To see that F must be closed, suppose to the contrary that F is not closed. Then $X \setminus F$ is not open. Hence there exists a point $x_0 \in X \setminus F$ such that for every neighbourhood U of x_0, $U \cap F \neq \emptyset$.

Let

$$\Lambda = \{U \subseteq X \mid U \text{ is a neighourhood of } x_0\}.$$

As Λ is closed under finite intersections, if for $U_1, U_2 \in \Lambda$ we define $U_1 \leq U_2$ if $U_2 \subseteq U_1$, then $(\Lambda, \leq)$ is a direct set by Example 1.5.8.

For each $U \in \Lambda$, let $x_U \in F \cap U$ (note we are using the Axiom of Choice A.2.4 here). We claim that $(x_U)_{U \in \Lambda}$ is a net that converges to x_0. This then leads to a contradiction as $x_U \in F$ for all $U \in \Lambda$ but $x_0 \notin F$ thereby completing the proof. To see that $(x_U)_{U \in \Lambda}$ is a net that converges to x_0 in $(X, \mathcal{T})$, let U_0 be an arbitrary neighbourhood x. Then for all $U \geq U_0$ we have that $x_U \in U \subseteq U_0$. Hence $(x_U)_{U \in \Lambda}$ is a net that converges to x_0 in $(X, \mathcal{T})$ as claimed. ∎

Example 1.6.15. Let I be a non-empty set, let $\{(X_\alpha, \mathcal{T}_\alpha)\}_{\alpha\in I}$ be a non-empty indexed family of topological spaces, and let F_α be a closed subset of $(X_\alpha, \mathcal{T}_\alpha)$ for all $\alpha \in I$. We claim that $\prod_{\alpha\in I} F_\alpha$ is closed in $\prod_{\alpha\in I} X_\alpha$ when equipped with either the product or box topology. Indeed let $(f_\lambda)_{\lambda\in\Lambda}$ be an arbitrary net in $\prod_{\alpha\in I} F_\alpha$ that converges to some element $f \in \prod_{\alpha\in I} X_\alpha$. By the 'if'-direction of Theorem 1.5.25, for each $\alpha \in I$ the net $(f_\lambda(\alpha))_{\lambda\in\Lambda}$ converges to $f(\alpha)$ in $(X_\alpha, \mathcal{T}_\alpha)$. Therefore, as $f_\lambda(\alpha) \in F_\alpha$ for all $\lambda \in \Lambda$ and F_α is closed in $(X_\alpha, \mathcal{T}_\alpha)$, Theorem 1.6.14 implies $f(\alpha) \in F_\alpha$ for all α. Hence $f \in \prod_{\alpha\in I} F_\alpha$. Therefore, as $(f_\lambda)_{\lambda\in\Lambda}$ was arbitrary, Theorem 1.6.14 implies $\prod_{\alpha\in I} F_\alpha$ is closed.

Given a subset of a topological space, there will be potentially lots of convergent nets contained in a given subset. It would be nice to find a closed set that contains all the possible points of convergence. In particular, it would be nice to find the smallest possible set with this property.

Construction 1.6.16. Let $(X, \mathcal{T})$ be a topological space and let $A \subseteq X$. Note the set

$$\mathcal{F} = \{F \subseteq X \mid A \subseteq F \text{ and } F \text{ is a closed set in } (X, \mathcal{T})\}$$

is non-empty as $X \in \mathcal{F}$. Consequently, Proposition 1.6.9 implies the set

$$\overline{A} = \bigcap_{F\in\mathcal{F}} F$$

is a closed set in $(X, \mathcal{T})$ that contains A. As clearly $A \subseteq \overline{A}$, we obtain that $\overline{A} \in \mathcal{F}$ and thus $\overline{A}$ is the smallest closed set in $(X, \mathcal{T})$ that contains A. This causes us to define the following.

Definition 1.6.17. The *closure* of a set A in a topological space $(X, \mathcal{T})$ is the set $\overline{A}$ obtained by taking the intersection of all closed subsets of $(X, \mathcal{T})$ that contain A.

Example 1.6.18. Given $\mathbb{R}$ with its canonical topology and $a, b \in \mathbb{R}$ with $a < b$, the closure of each of (a, b), $[a, b]$, $(a, b]$, and $[a, b)$ is $[a, b]$. Indeed $[a, b]$ is a closed set containing each of these sets. Furthermore, as every other close subset of $\mathbb{R}$ containing one of these sets must also contain a and b by Theorem 1.6.14, $[a, b]$ is the smallest closed subset of $\mathbb{R}$ containing each of these sets and thus must be the closure of each of these sets.

Example 1.6.19. Let (X, d) be a metric space with at least two points, let $x \in X$, and let $r > 0$. It is possible that $\overline{B_d(x, r)} \neq B_d[x, r]$. Indeed let d be the discrete metric on X. Then $\{x\} = B_d(x, 1)$ is closed and thus equal to its own closure. However $B_d[x, 1] = X$ which does not equal $B_d(x, 1)$.

Of course, closures of sets behave well with respect to subspaces and products.

Lemma 1.6.20. *Let $(Y, \mathcal{T}_Y)$ be a subspace of a topological space $(X, \mathcal{T})$ and let $A \subseteq Y$. The closure of A in $(Y, \mathcal{T}_Y)$ is the intersection of Y and the closure of A in $(X, \mathcal{T})$.*

Proof. Let B denote the closure of A in $(Y, \mathcal{T}_Y)$ and let C denote the closure of A in $(X, \mathcal{T})$. As C is closed in $(X, \mathcal{T})$, $Y \cap C$ is closed in $(Y, \mathcal{T}_Y)$ by Lemma 1.6.12. Therefore $B \subseteq Y \cap C$ by definition.

To see the other inequality, recall since B is a closed set in $(Y, \mathcal{T}_Y)$ that Lemma 1.6.12 implies there exists a closed set F in $(X, \mathcal{T})$ such that $B = Y \cap F$. However as $A \subseteq B = Y \cap F \subseteq F$, and as F is a closed subset in $(X, \mathcal{T})$, the definition of the closure of a set implies $C \subseteq F$. Hence

$$Y \cap C \subseteq Y \cap F = B$$

as desired. ∎

Before we show how closures work for the product and box topologies, we demonstrate the following useful tool.

Theorem 1.6.21. *Let $(X, \mathcal{T})$ be a topological space, let $A \subseteq X$, and let $x \in X$. The following are equivalent:*

(i) $x \in \overline{A}$.

(ii) There exists a net $(x_\lambda)_{\lambda \in \Lambda}$ of points in A that converges to x.

(iii) For every neighbourhood $U \in \mathcal{T}$ of x, $U \cap A \neq \emptyset$.

Furthermore, if $\mathcal{B}$ is a basis for $(X, \mathcal{T})$ or a neighbourhood basis for x, then the above are equivalent to

(iv) For every neighbourhood $U \in \mathcal{B}$ of x, $U \cap A \neq \emptyset$.

Proof. First suppose $x \in \overline{A}$. To see that (iii) holds, suppose to the contrary that there exists a neighbourhood $U \in \mathcal{T}$ of x such that $U \cap A = \emptyset$. Then $X \setminus U$ is a closed set containing A so $\overline{A} \subseteq X \setminus U$. However $x \in U$ and $x \in \overline{A}$ contradict the fact that $\overline{A} \subseteq X \setminus U$. Hence (i) implies (iii).

Next, suppose (iii) holds. To see that (ii) holds, let

$$\Lambda = \{U \subseteq X \mid U \text{ is a neighourhood of } x\}.$$

As Λ is closed under finite intersections, if for $U_1, U_2 \in \Lambda$ we define $U_1 \leq U_2$ if $U_2 \subseteq U_1$, then $(\Lambda, \leq)$ is a direct set by Example 1.5.8.

For each $U \in \Lambda$, let $x_U \in A \cap U$ (note we are using the Axiom of Choice A.2.4 here). We claim that $(x_U)_{U \in \Lambda}$ is a net that converges to x. To see this, let U_0 be an arbitrary neighbourhood x. Then for all $U \geq U_0$ we have that $x_U \in U \subseteq U_0$. Hence $(x_U)_{U \in \Lambda}$ is a net that converges to x in $(X, \mathcal{T})$ and, as

$x_U \in A$ for all $U \in \Lambda$, we have constructed an acceptable net. Hence (iii) implies (ii).

To see that (ii) implies (i), we note that if exists a net $(x_\lambda)_{\lambda\in\Lambda}$ of points in A that converges to x, then x must be in every closed subset of X containing A by Theorem 1.6.14. Thus $x \in \overline{A}$ by the definition of the closure of a set. Hence (ii) implies (i).

Finally, in the case $\mathcal{B}$ is a basis for $(X, \mathcal{T})$ or a neighbourhood basis for x, i) implies iv) by identical arguments. Furthermore to see that (iv) implies (ii), consider

$$\Lambda = \{U \in \mathcal{B} \mid U \text{ is a neighourhood of } x\}.$$

Then Λ is a net with the same ordering as above by the properties of bases and neighbourhood bases. A net $(x_U)_{U\in\Lambda}$ is constructed as above and still converges to x by the properties of bases and neighbourhood bases. ∎

Using Theorem 1.6.21, we can describe closure in the box and product topologies.

Proposition 1.6.22. *Let I be a non-empty set, let $\{(X_\alpha, \mathcal{T}_\alpha)\}_{\alpha\in I}$ be a non-empty indexed family of topological spaces, and let $A_\alpha \subseteq X_\alpha$ for all $\alpha \in I$. Then, when $\prod_{\alpha\in I} X_\alpha$ is equipped with either the box or the product topology,*

$$\overline{\prod_{\alpha\in I} A_\alpha} = \prod_{\alpha\in I} \overline{A_\alpha}.$$

Proof. By Example 1.6.15 we see that $\prod_{\alpha\in I} \overline{A_\alpha}$ is a closed set containing $\prod_{\alpha\in I} A_\alpha$. Hence

$$\overline{\prod_{\alpha\in I} A_\alpha} \subseteq \prod_{\alpha\in I} \overline{A_\alpha}$$

by definition.

To see the other inequality we will use Theorem 1.6.21. Let $x \in \prod_{\alpha\in I} \overline{A_\alpha}$ be arbitrary and write $x = (x_\alpha)_{\alpha\in I}$. To see that $x \in \overline{\prod_{\alpha\in I} A_\alpha}$, let V be a neighbourhood of x. Thus, by our knowledge of bases, we can find a set $U = \prod_{\alpha\in I} U_\alpha$ where $U_\alpha \in \mathcal{T}_\alpha$ (with $U_\alpha = X_\alpha$ for all but finitely many $\alpha \in I$ in the case we are using the product topology) such that $x \in U \subseteq V$. Since $x \in U$ we have that $x_\alpha \in U_\alpha$ for all $\alpha \in I$. Moreover, since $x \in \prod_{\alpha\in I} \overline{A_\alpha}$ we know that $x_\alpha \in \overline{A_\alpha}$ for all $\alpha \in I$ by Theorem 1.6.21. Therefore, since $x_\alpha \in U_\alpha$ and since $x_\alpha \in \overline{A_\alpha}$ there exists an $a_\alpha \in A_\alpha \cap U_\alpha$ for all $\alpha \in I$ by a property of the closure. Thus

$$(a_\alpha)_{\alpha\in I} \in U \cap \left(\prod_{\alpha\in I} A_\alpha\right) \subseteq V \cap \left(\prod_{\alpha\in I} A_\alpha\right).$$

Therefore, as V was an arbitrary neighbourhood of x, we obtain that $x \in \overline{\prod_{\alpha\in I} A_\alpha}$ by Theorem 1.6.21 as desired. ∎

Of course, when studying closures and closed sets via limits, for each element x in a set A there is clearly a net with elements from A that converges to x; namely a constant net where every value is x. Thus when taking a closure or asking whether a set is closed, we are more interested in nets that do not take the value of a specific point in a set. In particular, analysing the proof of Theorem 1.6.21, we easily see the following.

Corollary 1.6.23. *Let $(X, \mathcal{T})$ be a topological space, let $A \subseteq X$, and let $x \in X$. The following are equivalent:*

(i) There exists a net $(x_\lambda)_{\lambda \in \Lambda}$ of points in $A \setminus \{x\}$ that converges to x.

(ii) For every neighbourhood U of x, $U \cap (A \setminus \{x\}) \neq \emptyset$.

Proof. The proof that (ii) implies (i) is identical to the proof that (iii) implies (ii) in Theorem 1.6.21. Conversely, the proof that (i) implies (ii) follows directly from the definition of a convergent net. ∎

As being able to determining points of convergence from non-constant nets is useful in many scenarios, we give said object a name.

Definition 1.6.24. Let $(X, \mathcal{T})$ be a topological space and let $A \subseteq X$. A point $x \in X$ is said to be a *cluster point* of A if one of the two equivalent conditions from Corollary 1.6.23 hold for x, A, and $(X, \mathcal{T})$.

The set of cluster points of A is denoted $\mathrm{cluster}(A)$.

Remark 1.6.25. Note some authors use the term ‘limit points’ instead of cluster points. However a disjoint set of authors use the term ‘limits points’ to mean the set of all points of convergence. Thus we endeavour to ignore this ambiguity.

Example 1.6.26. Given $\mathbb{R}$ equipped with its canonical topology and $a, b \in \mathbb{R}$ with $a < b$, it is not difficult to see that the set of cluster points for $[a, b]$, (a, b), $[a, b)$, and $(a, b]$ are all $[a, b]$ as every point in $[a, b]$ is a point of convergence for some non-constant net from (a, b).

Example 1.6.27. Let $A = \left\{\frac{1}{n} \mid n \in \mathbb{N}\right\}$ viewed as a subset of $\mathbb{R}$ with its canonical topology. Then the only cluster point of A is 0. Indeed clearly the sequence $\left(\frac{1}{n}\right)_{n \geq 1}$ converges to but never equals 0 and thus 0 is a cluster point. To see that 0 is the only cluster point of A, we first claim that $\overline{A} = A \cup \{0\}$. To see this, we note that $A \cup \{0\}$ is closed as its complement is a countable union of open intervals and thus is open. However A is not closed by Theorem 1.6.14 as $\left(\frac{1}{n}\right)_{n \geq 1}$ is a sequence from A that converges to 0, which is not in A. Hence $\overline{A} = A \cup \{0\}$. Therefore, by Theorem 1.6.14, the set of possible cluster points must be contained in $A \cup \{0\}$. However, it is clear that no point in A can be a cluster point of A since the distance between $\frac{1}{n}$ and any other point in A is at least $\frac{1}{n} - \frac{1}{n-1}$ so it is impossible for a net from $A \setminus \left\{\frac{1}{n}\right\}$ to converge to $\frac{1}{n}$. Hence the only cluster point of A is 0.

Example 1.6.28. Let $\mathcal{C}$ be the Cantor set from Example 1.6.11. Then the set of cluster points of $\mathcal{C}$ is precisely $\mathcal{C}$. To see this, note as $\mathcal{C}$ is closed that the cluster points of $\mathcal{C}$ are contained in $\mathcal{C}$. To see that every point in $\mathcal{C}$ is a cluster point of $\mathcal{C}$, let $x \in \mathcal{C}$ be arbitrary. Thus, by the definition of $\mathcal{C}$, for each $n \in \mathbb{N}$ there exists a unique closed interval I_n of the form $\left[\frac{2k_n}{3^n}, \frac{2k_n+1}{3^n}\right]$ where $k_n \in \{0, 1, \ldots, \frac{1}{2}(3^n - 1)\}$. Choose x_n to be one of the endpoints of I_n that is not equal to x (as there are two distinct endpoints, such a point exists). As it is elementary to verify that the endpoints of I_n are elements of $\mathcal{C}$, we see that $x_n \in \mathcal{C}$ and $|x - x_n| < \frac{1}{3^n}$. Hence $(x_n)_{n\geq 1}$ is a sequence in $\mathcal{C} \setminus \{x\}$ that converges to x. Hence $\mathcal{C}$ is equal to its cluster points.

Example 1.6.29. Let A be a non-empty subset of a topological space $(X, \mathcal{T})$. Then the closure of A in the subspace $(A, \mathcal{T}_A)$ is A as A is closed in the subspace topology.

Perhaps unsurprisingly, the only thing that prevents a set from begin closed is it not containing its cluster points.

Theorem 1.6.30. *Let $(X, \mathcal{T})$ be a topological space and let $A \subseteq X$. Then*

$$\overline{A} = A \cup \mathrm{cluster}(A).$$

Proof. First, it is clear that $A \subseteq \overline{A}$ and that $\mathrm{cluster}(A) \subseteq \overline{A}$ by Theorem 1.6.21 and the definition of a cluster point. Hence

$$\overline{A} \supseteq A \cup \mathrm{cluster}(A).$$

To see the other inequality, let $x \in \overline{A}$ be arbitrary. If $x \in A$ then $x \in A \cup \mathrm{cluster}(A)$ and there is nothing left to show. Thus we may suppose that $x \notin A$. Since $x \in \overline{A}$, Theorem 1.6.21 implies that $U \cap A \neq \emptyset$ for every neighbourhood U of x. As $x \notin A$, $U \cap (A \setminus \{x\}) \neq \emptyset$ for every neighbourhood U of x. Hence Corollary 1.6.23 implies that $x \in \mathrm{cluster}(A)$. Therefore, in either case $x \in A \cup \mathrm{cluster}(A)$. Hence, as $x \in \overline{A}$ was arbitrary, $\overline{A} = A \cup \mathrm{cluster}(A)$ as desired. ∎

Corollary 1.6.31. *Let $(X, \mathcal{T})$ be a topological space and let $A \subseteq X$. Then A is closed if and only if* $\mathrm{cluster}(A) \subseteq A$.

Proof. If A is closed, then $A = \overline{A} = A \cup \mathrm{cluster}(A)$ by Theorem 1.6.30 and hence $\mathrm{cluster}(A) \subseteq A$. Conversely, if $\mathrm{cluster}(A) \subseteq A$, then Theorem 1.6.30 implies that $\overline{A} = A \cup \mathrm{cluster}(A) = A$. Therefore, as A is equal to its closure and the closure of a set is a closed set, A is closed. ∎

All of the above has been focused on closures and closed sets via describing points of convergence for nets based on a set. However, it is often useful to understand just the points inside a set. In particular, it is useful to understand the set of points in a set that are 'far away' from the complement of the set. These points are described based on the following, which is constructed in a similar fashion to how we constructed the closure of a set.

Construction 1.6.32. Let $(X, \mathcal{T})$ be a topological space and let $A \subseteq X$. Note the set

$$\mathcal{U} = \{U \subseteq X \mid U \subseteq A \text{ and } U \in \mathcal{T}\}$$

is non-empty as $\emptyset \in A$. Consequently, Definition 1.1.1 implies the set

$$\operatorname{int}(A) = \bigcup_{U \in \mathcal{U}} U$$

is an open set in $(X, \mathcal{T})$ that is contained A. As clearly $\operatorname{int}(A) \subseteq A$, we obtain that $\operatorname{int}(A) \in \mathcal{U}$ and thus $\operatorname{int}(A)$ is the largest open set in $(X, \mathcal{T})$ that is contained in A. This causes us to define the following.

Definition 1.6.33. The *interior* of a set A in a topological space $(X, \mathcal{T})$ is the set $\operatorname{int}(A)$ obtained by taking the union of all open subsets of $(X, \mathcal{T})$ contained A.

Example 1.6.34. Given $\mathbb{R}$ equipped with its canonical topology and $a, b \in \mathbb{R}$ with $a < b$, it is not difficult to see that interior of $[a, b]$, (a, b), $[a, b)$, and $(a, b]$ is (a, b) as clearly (a, b) is open, is contained in these sets, contains all points in these sets for except possible a and b, and as the addition of a or b to (a, b) creates a set that is not open.

Example 1.6.35. Let $A = \left\{\frac{1}{n} \mid n \in \mathbb{N}\right\}$ viewed as a subset of $\mathbb{R}$ with its canonical topology. Then the interior of A is the empty set as no open interval is contained in A.

Example 1.6.36. The Cantor set $\mathcal{C}$, viewed as a subset of $\mathbb{R}$, has no interior. Indeed suppose to the contrary that $\operatorname{int}(\mathcal{C})$ is non-empty. Hence there exists an open interval $(a, b) \subseteq \operatorname{int}(\mathcal{C}) \subseteq \mathcal{C}$ by the definition of the interior. By the elementary properties of real numbers, we can choose $N \in \mathbb{N}$ such that $\frac{1}{3^N} < b - a$. This then implies that (a, b) cannot be contained in P_N as defined in Example 1.6.11 as none of the separated intervals in P_N have length greater than $\frac{1}{3^N}$. Hence the Cantor set has no interior (even though it is uncountable and every point is a cluster point).

Example 1.6.37. Let A be the x-axis in $\mathbb{R}^2$ equipped with its topology from the Euclidean norm. Then the interior of A is empty as A contains no open balls from $\mathbb{R}^2$.

Example 1.6.38. Let A be a non-empty subset of topological space $(X, \mathcal{T})$. Then the interior of A in the subspace $(A, \mathcal{T}_A)$ is A as A is open in the subspace topology.

Although we do not have any theory here related to the interior like we did with the closure results seen above, the interior will be useful later in the course.

As we can see based on these examples, the set of interior points to a set are those that are 'far away' from the complement of the set as there is an open set containing these points that does not intersect the complement. To formalize this, we define another type of point for a given set.

Definition 1.6.39. Let $(X, \mathcal{T})$ be a topological space and let $A \subseteq X$. A point $x \in X$ is said to be a *boundary point of* A if $A \cap U \neq \emptyset$ and $(X \setminus A) \cap U \neq \emptyset$ for every neighbourhood $U \in T$ of x.

The set of boundary points of A is denoted $\mathrm{bdy}(A)$.

Before we look at examples of boundary points, we first prove two results which completely describe the set of boundary points based on objects we have previously studied.

Corollary 1.6.40. *Let $(X, \mathcal{T})$ be a topological space and let $A \subseteq X$. Then $\mathrm{bdy}(A) = \overline{A} \cap \overline{(X \setminus A)}$.*

Proof. This result easily follows from Theorem 1.6.21 and the definition of a boundary point. ∎

Theorem 1.6.41. *Let $(X, \mathcal{T})$ be a topological space and let $A \subseteq X$. Then $\mathrm{int}(A)$ and $\mathrm{bdy}(A)$ are disjoint sets such that*

$$\overline{A} = \mathrm{int}(A) \cup \mathrm{bdy}(A).$$

Furthermore

$$\mathrm{int}(A) = A \setminus \overline{(X \setminus A)}.$$

Proof. Clearly if $x \in \mathrm{int}(A)$, then there exists an open set U containing x (namely $\mathrm{int}(A)$) such that $U \cap (X \setminus A) = \emptyset$ and thus $x \notin \mathrm{bdy}(A)$. Hence $\mathrm{int}(A)$ and $\mathrm{bdy}(A)$ are disjoint sets.

To see that $\overline{A} \subseteq \mathrm{int}(A) \cup \mathrm{bdy}(A)$, let $x \in \overline{A}$ be arbitrary. Hence Theorem 1.6.21 implies that for every neighbourhood U of x, $U \cap A \neq \emptyset$. If there exists a neighbourhood U of x such that $U \cap (X \setminus A) = \emptyset$, then $U \subseteq A$ and hence $x \in \mathrm{int}(A)$. Otherwise for every neighbourhood U of x, $U \cap A \neq \emptyset$ and $U \cap (X \setminus A) \neq \emptyset$ so $x \in \mathrm{bdy}(A)$. Therefore, as $x \in \overline{A}$ was arbitrary, $\overline{A} \subseteq \mathrm{int}(A) \cup \mathrm{bdy}(A)$. For the reverse inequality, we note that $\mathrm{bdy}(A) \subseteq \overline{A}$ by the definition of a boundary point and Theorem 1.6.21, and similarly $\mathrm{int}(A) \subseteq A \subseteq \overline{A}$ trivially.

Finally, to see that $\mathrm{int}(A) = A \setminus \overline{(X \setminus A)}$, note if $x \in \mathrm{int}(A)$ then there exists a neighbourhood U of x contained in A and thus $x \notin \overline{(X \setminus A)}$ by Theorem 1.6.21. Furthermore, as $\mathrm{int}(A) \subseteq A$ by construction, $\mathrm{int}(A) \subseteq A \setminus \overline{(X \setminus A)}$. To see the reverse inequality, note as

$$A \setminus \overline{(X \setminus A)} \subseteq X \setminus \overline{(X \setminus A)} \subseteq X \setminus (X \setminus A) = A,$$

and as $X \setminus \overline{(X \setminus A)}$ is the complement of an closed set and thus is open, $A \setminus \overline{(X \setminus A)} \subseteq \mathrm{int}(A)$ by definition. Hence $\mathrm{int}(A) = A \setminus \overline{(X \setminus A)}$ as desired. ■

Example 1.6.42. Given $\mathbb{R}$ equipped with its canonical topology and $a, b \in \mathbb{R}$ with $a < b$, it is not difficult to see that boundary of $[a, b]$, (a, b), $[a, b)$, and $(a, b]$ is $\{a, b\}$. Indeed as the closure of each of these sets is $[a, b]$ by Example 1.6.18 and the interior of each of these sets is (a, b) by Example 1.6.34, the claim follows from Theorem 1.6.41.

Example 1.6.43. Given $\mathbb{R}$ equipped with its canonical topology, the boundary of $\mathbb{Q}$ is $\mathbb{R}$ as every neighbourhood of some point from $\mathbb{Q}$ contains an interval, which must contain a rational and irrational number.

Example 1.6.44. Let (X, d) be a metric space with at least two points, let $x \in X$, and let $r > 0$. It is possible that $\mathrm{bdy}(B_d(x, r))$ and $\mathrm{bdy}(B_d[x, r])$ are not equal to

$$\{y \in Y \in \mid \, d(x, y) = r\}.$$

Indeed let d be the discrete metric on X. Then $B_d(x, 1) = \{x\}$ and $B_d[x, 1] = X$ have empty boundary sets as they are open sets and thus equal to their own interior. However the above set is X which is not equal to $\emptyset$.

The notions of points and sets observed in this section will be seen throughout the course (less so with the boundary points). As with undergraduate real analysis, the notions related to closed sets and closures of sets will be of vital importance when discussing continuous functions and compact sets; which happen to be the next two chapters.

Chapter 2

Continuous Functions

With the completion of our introduction to topological spaces, it is natural to consider the maps between topological spaces. This is a common theme in mathematics; first one defines objects, then one studies the morphisms between them. The correct morphisms between topological spaces are known as continuous functions. These generalizations of the continuous functions studied in calculus have wide reaching applications throughout mathematics and are some of the best behave functions in all of mathematics. Consequently, this chapter is devoted to the development of the theory of continuous functions. After developing the numerous characterizations of continuous functions and some basic properties in the first section, we will use continuous functions to study quotient spaces, connected sets in topological spaces thereby enabling us to generalize the Intermediate Value Theorem, and other forms of connectedness.

2.1 Continuous Functions

To begin our study of continuous functions on topological spaces, we recall a function $f : \mathbb{R} \to \mathbb{R}$ is continuous if for each $x_0 \in \mathbb{R}$ and each $\epsilon > 0$ there exists a $\delta > 0$ such that if $x \in \mathbb{R}$ and $|x - x_0| < \delta$, then $|f(x) - f(x_0)| < \epsilon$. Alternatively, this can be rewritten as

$$(x_0 - \delta, x_0 + \delta) \subseteq f^{-1}((f(x_0) - \epsilon, f(x_0) + \epsilon))$$

(see Definition A.3.4 for the definition of $f^{-1}(A)$ for a set A). As open intervals form a basis for the canonical topology on $\mathbb{R}$, it is elementary to generalize the above idea of a continuous function to topological spaces.

Definition 2.1.1. Let $(X, \mathcal{T}_X)$ and $(Y, \mathcal{T}_Y)$ be topological spaces. A function $f : X \to Y$ is said to be *continuous* if $f^{-1}(U) \in \mathcal{T}_X$ for every $U \in \mathcal{T}_Y$; that is, the inverse image of every open set (from Y) is open (in X).

To verify that Definition 2.1.1 is indeed a generalization of a continuous function on $\mathbb{R}$, we demonstrate the following equivalent characterization of continuous functions between metric spaces.

Proposition 2.1.2. *Let (X, d_x) and (Y, d_Y) be metric spaces and let $f : X \to Y$. Then f is a continuous function from X to Y when equipped with their metric topologies if and only if for all $x_0 \in X$ and $\epsilon > 0$ there exists a $\delta > 0$ such that if $x \in X$ and $d_X(x, x_0) < \delta$ then $d_Y(f(x), f(x_0)) < \epsilon$.*

Proof. Let $f : X \to Y$ be continuous. To see that f has the desired property, let $x_0 \in X$ and $\epsilon > 0$ be arbitrary. Consider the set $U = B_{d_Y}(f(x_0), \epsilon)$. As U is open in Y and as f is continuous, $f^{-1}(U)$ is open in X. Therefore, since $x_0 \in f^{-1}(U)$ as $f(x_0) \in B_{d_Y}(f(x_0), \epsilon) = U$, there exists a $\delta > 0$ such that $B_{d_X}(x_0, \delta) \subseteq f^{-1}(U)$ by the definition of the metric topology. Thus if $x \in X$ and $d_X(x, x_0) < \delta$, then $x \in B_{d_X}(x_0, \delta) \subseteq f^{-1}(U)$ so $f(x) \in U = B_{d_Y}(f(x_0), \epsilon)$ hence $d_Y(f(x), f(x_0)) < \epsilon$. Therefore, as $x_0 \in X$ and $\epsilon > 0$ were arbitrary, f has the desired property.

Conversely, suppose for all $x_0 \in X$ and $\epsilon > 0$ there exists a $\delta > 0$ such that if $x \in X$ and $d_X(x, x_0) < \delta$ then $d_Y(f(x), f(x_0)) < \epsilon$. To see that f is continuous, let U be an arbitrary open set in Y. To see that $f^{-1}(U)$ is open in X, let $x_0 \in f^{-1}(U)$ be arbitrary. Then $f(x_0) \in U$ so, as U is open, the definition of the metric topology implies there exists an $\epsilon > 0$ such that $B_{d_Y}(f(x_0), \epsilon) \subseteq U$. By the assumption of this direction of the proof, there exists a $\delta > 0$ such that if $x \in X$ and $d_X(x, x_0) < \delta$ then $d_Y(f(x), f(x_0)) < \epsilon$. Thus, if $x \in B_{d_X}(x_0, \delta)$ then $d_X(x, x_0) < \delta$ so $d_Y(f(x), f(x_0)) < \epsilon$ and thus $f(x) \in B_{d_Y}(f(x_0), \epsilon) \subseteq U$. Therefore $B_{d_X}(x_0, \delta) \subseteq f^{-1}(U)$. Therefore, as $x_0 \in f^{-1}(U)$ was arbitrary, $f^{-1}(U)$ is open by the definition of the metric topology. Hence, as U was arbitrary, f is continuous as desired. ∎

Thus Proposition 2.1.2 implies that all the continuous functions on $\mathbb{R}$ seen in undergraduate calculus are continuous functions between topological spaces. Thus we have already have several examples of continuous functions. Of course there are many more. Before we get to examples of continuous functions between general topological spaces, we mention one useful example in the metric space setting. To do so, we must define the following.

Definition 2.1.3. Let (X, d) be a metric space and let $A \subseteq X$ be a non-empty set. Given $x \in X$, the *distance from x to A*, denoted $\mathrm{dist}(x, A)$, is defined to be

$$\mathrm{dist}(x, A) = \inf\{d(x, a) \mid a \in A\}.$$

Example 2.1.4. Let (X, d) be a metric space and let $A \subseteq X$ be non-empty. The function $f : X \to \mathbb{R}$ defined by

$$f(x) = \mathrm{dist}(x, A)$$

is a continuous function from X equipped with the metric topology to $\mathbb{R}$ equipped with its canonical topology.

To see that f is continuous, first let $x_1, x_2 \in X$ be arbitrary. If $\delta > 0$, then by the definition of the distance there exists an $a \in A$ such that $d(x_1, a) \leq \mathrm{dist}(x_1, A) + \delta$. Therefore

$$\mathrm{dist}(x_2, A) \leq \mathrm{dist}(x_2, a) \leq d(x_1, x_2) + d(x_1, a) \leq d(x_1, x_2) + \mathrm{dist}(x_1, A) + \delta.$$

As the above inequality holds for all $\delta > 0$, we obtain that $f(x_2) \leq f(x_1) + d(x_1, x_2)$. By reversing the roles of x_1 and x_2, we obtain that $f(x_1) \leq f(x_2) + d(x_1, x_2)$ and hence $|f(x_1) - f(x_2)| \leq d(x_1, x_2)$.

To see now that f is continuous, let $x_0 \in X$ and $\epsilon > 0$ be arbitrary. Let $\delta = \epsilon > 0$. Therefore, if $x \in X$ is such that $d(x, x_0) < \delta$ then $|f(x) - f(x_0)| \leq d(x, x_0) < \delta = \epsilon$. Thus, as $x_0 \in X$ and $\epsilon > 0$ were arbitrary, f is continuous by Proposition 2.1.2.

Now onto some non-metric related examples.

Example 2.1.5. Let $(X, \mathcal{T}_X)$ and $(Y, \mathcal{T}_Y)$ be topological spaces and let $y_0 \in Y$. The constant function $f : X \to Y$ defined by $f(x) = y_0$ for all $x \in X$ is a continuous function. Indeed for every open set U in Y we have that

$$f^{-1}(U) = \begin{cases} X & \text{if } y_0 \in U \\ \emptyset & \text{if } y_0 \notin U \end{cases}.$$

Therefore, as $\emptyset, X \in \mathcal{T}_X$ by the definition of a topology, f is continuous.

Example 2.1.6. Let $(X, \mathcal{T}_X)$ and $(Y, \mathcal{T}_Y)$ be topological spaces. If $\mathcal{T}_X$ is the discrete topology, then every function $f : X \to Y$ is continuous as $\mathcal{T}_X = \mathcal{P}(X)$ implies $f^{-1}(U) \in \mathcal{T}_X$ for every $U \in \mathcal{T}_Y$.

Example 2.1.7. Let $(X, \mathcal{T}_X)$ and $(Y, \mathcal{T}_Y)$ be topological spaces. If $\mathcal{T}_Y$ is the trivial topology, then every function $f : X \to Y$ is continuous as $f^{-1}(Y) = X$, $f^{-1}(\emptyset) = \emptyset$, and $\mathcal{T}_Y = \{\emptyset, Y\}$.

Example 2.1.8. Let I be a non-empty set and let $\{(X_\alpha, \mathcal{T}_\alpha)\}_{\alpha \in I}$ be a set of topological spaces. For a fixed $\alpha_0 \in I$, consider the map

$$\pi_{\alpha_0} : \prod_{\alpha \in I} X_\alpha \to X_{\alpha_0}$$

defined by

$$\pi_{\alpha_0}((x_\alpha)_{\alpha \in I}) = x_{\alpha_0}$$

for all $(x_\alpha)_{\alpha \in I} \in \prod_{\alpha \in I} X_\alpha$. The map π_{α_0} is called the *projection map onto the α_0^{th} coordinate*.

Every projection map is continuous when $\prod_{\alpha\in I} X_\alpha$ is equipped with either the product topology or the box topology. Indeed notice for all $U \in \mathcal{T}_{\alpha_0}$ that

$$\pi_{\alpha_0}^{-1}(U) = \prod_{\alpha\in I} V_\alpha$$

where $V_\alpha = X_\alpha$ if $\alpha \neq \alpha_0$ and $V_{\alpha_0} = U$. Thus, as $\prod_{\alpha\in I} V_\alpha$ is open in both the product and box topologies and as $U \in \mathcal{T}_{\alpha_0}$ was arbitrary, π_{α_0} is continuous.

In fact, as the collection $\{\pi_\alpha(U_\alpha) \mid \alpha \in I, U_\alpha \in \mathcal{T}_\alpha\}$ is a subbasis for the product topology, the product topology is the coarsest topology for which each projection map is continuous.

Of course, there are many ways to test whether a function on $\mathbb{R}$ is continuous. In particular, one characterization of continuous functions on $\mathbb{R}$ that is often used as the definition of continuity due to its viability is the characterization that a function is continuous if and only if it maps convergent sequences to convergent sequences. In the following result, we extend all of these characterizations to arbitrary topological spaces.

Theorem 2.1.9. *Let $(X, \mathcal{T}_X)$ and $(Y, \mathcal{T}_Y)$ be topological spaces and let $f : X \to Y$. The following are equivalent:*

(i) f is continuous.

(ii) For every $x \in X$ and every $\mathcal{T}_Y$-neighbourhood U of $f(x)$ there exists a $\mathcal{T}_X$-neighbourhood V of x such that $V \subseteq f^{-1}(U)$.

(iii) For any bases $\mathcal{B}_X$ of $(X, \mathcal{T}_X)$ and $\mathcal{B}_Y$ of $(Y, \mathcal{T}_Y)$, for every $x \in X$ and every neighbourhood $U \in \mathcal{B}_Y$ of $f(x)$ there exists a neighbourhood $V \in \mathcal{B}_X$ of x such that $V \subseteq f^{-1}(U)$.

(iv) For some bases $\mathcal{B}_X$ of $(X, \mathcal{T}_X)$ and $\mathcal{B}_Y$ of $(Y, \mathcal{T}_Y)$, for every $x \in X$ and every neighbourhood $U \in \mathcal{B}_Y$ of $f(x)$ there exists a neighbourhood $V \in \mathcal{B}_X$ of x such that $V \subseteq f^{-1}(U)$.

(v) For every net $(x_\lambda)_{\lambda\in\Lambda}$ in X that converges to some x_0 in $(X, \mathcal{T}_X)$, the net $(f(x_\lambda))_{\lambda\in\Lambda}$ converges to $f(x_0)$ in $(Y, \mathcal{T}_Y)$.

(vi) For every $A \subseteq X$, $f\left(\overline{A}\right) \subseteq \overline{f(A)}$, where the closures are taken in their appropriate spaces.

(vii) For every closed set F in $(Y, \mathcal{T}_Y)$, $f^{-1}(F)$ is closed in $(X, \mathcal{T}_X)$.

Proof. To see that (i) implies (ii), let f be continuous and let $x \in X$ and U a $\mathcal{T}_Y$-neighbourhood of $f(x)$ be arbitrary. Then, as f is continuous, $V = f^{-1}(U)$ is clearly a $\mathcal{T}_X$-neighbourhood of x such that $V \subseteq f^{-1}(U)$. Therefore, as x and U were arbitrary, (i) implies (ii).

To see that (ii) implies (iii), let $x \in X$ and $U \in \mathcal{B}_Y$ a $\mathcal{T}_Y$-neighbourhood of $f(x)$ be arbitrary. By ii) there exists a $\mathcal{T}_X$-neighbourhood V_0 of x such that $V_0 \subseteq f^{-1}(U)$. Since $\mathcal{B}_X$ is a basis for $(X, \mathcal{T}_X)$ there exists $V \in \mathcal{B}_X$ such that $x \in V \subseteq V_0$. Hence $V \in \mathcal{B}_X$ is a neighbourhood of x such that $V \subseteq V_0 \subseteq f^{-1}(U)$. Therefore, as x and U were arbitrary, (ii) implies (iii).

Note (iii) trivially implies (iv).

To see that (iv) implies (v), suppose $\mathcal{B}_X$ is a basis for $(X, \mathcal{T}_X)$ and $\mathcal{B}_Y$ is a basis for $(Y, \mathcal{T}_Y)$ such that for every $x \in X$ and every neighbourhood $U \in \mathcal{B}_Y$ of $f(x)$ there exists a neighbourhood $V \in \mathcal{B}_X$ of x such that $V \subseteq f^{-1}(U)$. Let $(x_\lambda)_{\lambda\in\Lambda}$ be an arbitrary net in X that converges to some x_0 in $(X, \mathcal{T}_X)$. To see that $(f(x_\lambda))_{\lambda\in\Lambda}$ converges to $f(x_0)$ in $(Y, \mathcal{T}_Y)$, let $U \in \mathcal{B}_Y$ such that $f(x_0) \in U$ be arbitrary. By assumption there exists a $V \in \mathcal{B}_X$ such that $x_0 \in V$ and $V \subseteq f^{-1}(U)$. Thus, as V is an open set containing x_0 and as $(x_\lambda)_{\lambda\in\Lambda}$ converges to x_0 in $(X, \mathcal{T}_X)$, there exists a $\lambda_0 \in \Lambda$ such that $x_\lambda \in V$ for all $\lambda \geq \lambda_0$. Hence $f(x_\lambda) \in f(V) \subseteq U$ for all $\lambda \geq \lambda_0$. Therefore, as U was arbitrary, $(f(x_\lambda))_{\lambda\in\Lambda}$ converges to $f(x_0)$ in $(Y, \mathcal{T}_Y)$ by Lemma 1.5.19. Hence (iv) implies (v).

To see that (v) implies (vi), fix $A \subseteq X$ and let $x_0 \in \overline{A}$ be arbitrary. As $x_0 \in \overline{A}$ there exists a net $(x_\lambda)_{\lambda\in\Lambda}$ of points in A that converges to x_0 by Theorem 1.6.21. Therefore, by v), $(f(x_\lambda))_{\lambda\in\Lambda}$ is a net of points in $f(A)$ that converges to $f(x_0)$ in $(Y, \mathcal{T}_Y)$. Hence Theorem 1.6.21 $f(x_0) \in \overline{f(A)}$. Therefore, as $x_0 \in \overline{A}$ was arbitrary, $f\left(\overline{A}\right) \subseteq \overline{f(A)}$. Hence (v) implies (vi).

To see that (vi) implies (vii), let F be an arbitrary closed subset of $(Y, \mathcal{T}_Y)$ and let $A = f^{-1}(F \cap \overline{f(X)})$. Thus $F \cap \overline{f(X)} = f(A)$. Since $A \subseteq \overline{A}$, (vi) implies that

$$F \cap \overline{f(X)} = f(A) \subseteq f\left(\overline{A}\right) \subseteq \overline{f(A)} = \overline{F \cap \overline{f(X)}} = F \cap \overline{f(X)}$$

as $F \cap \overline{f(X)}$ is closed. Hence $f\left(\overline{A}\right) = F \cap \overline{f(X)}$ so $\overline{A} \subseteq f^{-1}(F \cap \overline{f(X)}) = A \subseteq \overline{A}$ so $A = \overline{A}$. Thus A is closed. Therefore, as F was arbitrary, (vi) implies (vii).

Finally, to see that (vii) implies (i), let $U \in \mathcal{T}_Y$ be arbitrary. Then $Y \setminus U$ is closed in $(Y, \mathcal{T}_Y)$. By assuming vii) we know that

$$f^{-1}(Y \setminus U) = f^{-1}(Y) \setminus f^{-1}(U) = X \setminus f^{-1}(U)$$

is closed in $(X, \mathcal{T}_X)$. Thus $f^{-1}(U) \in \mathcal{T}_X$. Hence, as $U \in \mathcal{T}_Y$ was arbitrary, f is continuous. Thus (vii) implies (i). ∎

Of course, alternate characterizations of continuous functions are always useful in proving results and obtaining examples of continuous functions.

Theorem 2.1.10. *Let $(X, \mathcal{T})$ be a topological space, let I be a non-empty set, let $\{(Y_\alpha, \mathcal{T}_\alpha)\}_{\alpha\in I}$ be a set of topological spaces, and, for each $\alpha \in I$, let*

$f_\alpha : X \to Y_\alpha$. The function $f : X \to \prod_{\alpha\in I} Y_\alpha$ defined by

$$f(x) = (f_\alpha(x))_{\alpha\in I}$$

for all $x \in X$ is continuous when $Y = \prod_{\alpha\in I} Y_\alpha$ is equipped with the product topology if and only if f_α is continuous for all $\alpha \in I$.

Furthermore, if $\{(X_\alpha, \mathcal{T}'_\alpha)\}_{\alpha\in I}$ is a set of topological spaces, if for each $\alpha \in I$ $f_\alpha : X_\alpha \to Y_\alpha$, if $X = \prod_{\alpha\in I} X_\alpha$ is equipped with the product topology, and if $f : X \to Y$ is defined by $f((x_\alpha)_{\alpha\in I}) = (f_\alpha(x_\alpha))_{\alpha\in I}$, then f is continuous if and only if f_α is continuous for all $\alpha \in I$.

Proof. For the first part, let $(x_\lambda)_{\lambda\in\Lambda}$ be an arbitrary net in X that converges to some point x_0 in $(X, \mathcal{T})$. By Theorem 1.5.25, $(f(x_\lambda))_{\lambda\in\Lambda}$ converges to $f(x_0)$ when $\prod_{\alpha\in I} X_\alpha$ is equipped with the product topology if and only if $(f_\alpha(x_\lambda))_{\lambda\in\Lambda}$ converges to $f_\alpha(x_0)$ in $(Y_\alpha, \mathcal{T}_\alpha)$ for all $\alpha \in I$. Hence the result follows from Theorem 2.1.9.

Similarly, for the second part, let $(x_\lambda)_{\lambda\in\Lambda}$ be an arbitrary net in X that converges to some point x_0 in $(X, \mathcal{T})$. Hence $(x_\lambda(\alpha))_{\lambda\in\Lambda}$ converges to $x_0(\alpha)$ in $(X_\alpha, \mathcal{T}'_\alpha)$ for all $\alpha \in I$. By Theorem 1.5.25, $(f(x_\lambda))_{\lambda\in\Lambda}$ converges to $f(x_0)$ when $\prod_{\alpha\in I} X_\alpha$ is equipped with the product topology if and only if $(f_\alpha(x_lambda(\alpha)))_{\lambda\in\Lambda}$ converges to $f_\alpha(x_0(\alpha))$ in $(X_\alpha, \mathcal{T}'_\alpha)$ for all $\alpha \in I$. Hence the result follows from Theorem 2.1.9. ∎

Remark 2.1.11. Unsurprisingly, Theorem 2.1.10 fails when the product topology is replaced with the box topology. To see this, consider $\mathbb{R}$ equipped with its canonical topology and consider $f : \mathbb{R} \to \prod_{n\in\mathbb{N}} \mathbb{R}$ defined by

$$f(x) = (x)_{n\in\mathbb{N}}$$

for all $x \in \mathbb{R}$. Clearly each entry of f constitutes a continuous function from $\mathbb{R}$ to $\mathbb{R}$ (i.e. $f_n(x) = x$ for all $n \in \mathbb{N}$ is clearly continuous). However, we claim that f is not continuous when $\prod_{n\in\mathbb{N}} \mathbb{R}$ is equipped with the box topology. To see this, for each $n \in \mathbb{N}$ let $I_n = \left(-\frac{1}{n}, \frac{1}{n}\right)$. Then

$$U = \prod_{n\in\mathbb{N}} I_n$$

is open in the box topology. However, it is elementary to verify that $f^{-1}(U) = \{0\}$ which is not open in $\mathbb{R}$. Hence f is not continuous when $\prod_{n\in\mathbb{N}} \mathbb{R}$ is equipped with the box topology.

However, the other direction holds. Indeed if f is continuous, the proof of Theorem 2.1.10 implies that each f_α is continuous as the 'if' direction of Theorem 1.5.25 holds when the box topology is used. Alternatively, to see that if f is continuous when $\prod_{\alpha\in I} X_\alpha$ is equipped with the box topology, let $\alpha_0 \in I$ and $U \in \mathcal{T}_\alpha$ be arbitrary. By defining $U_{\alpha_0} = U$ and $U_\alpha = X_\alpha$ for all

$\alpha \in I \setminus \{\alpha_0\}$, we know that $\prod_{\alpha\in I} U_\alpha$ is open in the box topology. Therefore, as f is continuous,

$$f^{-1}\left(\prod_{\alpha\in I} U_\alpha\right) = f_{\alpha_0}^{-1}(U_{\alpha_0})$$

is open in $(X, \mathcal{T})$. Therefore, as $\alpha_0 \in I$ and $U \in \mathcal{T}_\alpha$ were arbitrary, each f_α must be continuous.

Of course, in generality, we are interested in continuous functions as they will behave well with respect to any topological property we are interested in studying. On occasion, it is useful to study a more local property with respect to continuity. In particular, analyzing the proof of Theorem 2.1.9 yields the following.

Theorem 2.1.12. *Let $(X, \mathcal{T}_X)$ and $(Y, \mathcal{T}_Y)$ be topological spaces, let $x_0 \in X$, and let $f : X \to Y$. The following are equivalent:*

(i) For every $\mathcal{T}_Y$-neighbourhood U of $f(x_0)$ there exists a $\mathcal{T}_X$-neighbourhood V of x_0 such that $V \subseteq f^{-1}(U)$.

(ii) For any bases $\mathcal{B}_X$ of $(X, \mathcal{T}_X)$ and $\mathcal{B}_Y$ of $(Y, \mathcal{T}_Y)$, every neighbourhood $U \in \mathcal{B}_Y$ of $f(x_0)$ there exists a neighbourhood $V \in \mathcal{B}_X$ of x_0 such that $V \subseteq f^{-1}(U)$.

(iii) For some bases $\mathcal{B}_X$ of $(X, \mathcal{T}_X)$ and $\mathcal{B}_Y$ of $(Y, \mathcal{T}_Y)$, for every neighbourhood $U \in \mathcal{B}_Y$ of $f(x_0)$ there exists a neighbourhood $V \in \mathcal{B}_X$ of x_0 such that $V \subseteq f^{-1}(U)$.

(iv) For every net $(x_\lambda)_{\lambda\in\Lambda}$ in X that converges to x_0 in $(X, \mathcal{T}_X)$, the net $(f(x_\lambda))_{\lambda\in\Lambda}$ converges to $f(x_0)$ in $(Y, \mathcal{T}_Y)$.

Proof. The fact that (i) $\Rightarrow$ (ii) $\Rightarrow$ (iii) $\Rightarrow$ (iv) can be obtained by repeating (ii) $\Rightarrow$ (iii) $\Rightarrow$ (iv) $\Rightarrow$ (v) of Theorem 2.1.9 verbatim.

To see that (iv) implies (i), assume that (iv) holds. Suppose to the contrary that there exists a $\mathcal{T}_Y$-neighbourhood U of $f(x_0)$ such that for every $\mathcal{T}_X$-neighbourhood V of x_0 that $V \setminus f^{-1}(U) \neq \emptyset$. Thus for every $\mathcal{T}_X$-neighbourhood V of x_0 there exists a $x_V \in V$ such that $f(x_V) \notin U$ (note we are using the Axiom of Choice A.2.4 here).

Let

$$\Lambda = \{V \subseteq X \mid V \text{ is a } \mathcal{T}_X\text{-neighourhood of } x_0\}.$$

As Λ is closed under finite intersections, if for $V_1, V_2 \in \Lambda$ we define $V_1 \leq V_2$ if $V_2 \subseteq V_1$, then $(\Lambda, \leq)$ is a direct set by Example 1.5.8.

We claim that $(x_V)_{V\in\Lambda}$ converges to x_0 in $(X, \mathcal{T}_X)$ but $(f(x_V))_{V\in\Lambda}$ does not converge to $f(x_0)$ in $(Y, \mathcal{T}_Y)$. To see that $(x_V)_{V\in\Lambda}$ is a net that converges to x_0 in $(X, \mathcal{T})$, let V_0 be an arbitrary $\mathcal{T}$-neighbourhood x_0. Then for all $V \geq V_0$ we have that $x_V \in V \subseteq V_0$. Hence $(x_V)_{V\in\Lambda}$ is a net that converges to x_0 in $(X, \mathcal{T})$ by Definition 1.5.13. Thus $(f(x_V))_{V\in\Lambda}$ does not converge to $f(x_0)$ in $(Y, \mathcal{T}_Y)$, we simply note that U is a $\mathcal{T}_Y$-neighbourhood of $f(x_0)$ but $f(x_V) \notin U$ for all $V \in \Lambda$. Hence we have obtained a contradiction thereby finishing the proof. ∎

Due to the above, we define the following.

Definition 2.1.13. Let $(X, \mathcal{T}_X)$ and $(Y, \mathcal{T}_Y)$ be topological spaces, let $x_0 \in X$, and let $f : X \to Y$. It is said that f is *continuous at* x_0 if one of the four equivalent characterizations in Theorem 2.1.12 hold

Of course, global continuity is exactly local continuity at each point.

Corollary 2.1.14. *Let $(X, \mathcal{T}_X)$ and $(Y, \mathcal{T}_Y)$ be topological spaces and let $f : X \to Y$. Then f is continuous if and only if f is continuous at each point in X.*

Proof. Combine Theorem 2.1.9 and Theorem 2.1.12. ∎

As mentioned earlier, it is on occasion useful to consider this local property of continuity due to all of the equivalent characterizations produced in Theorem 2.1.12. Another useful ability is to be able to construct continuous functions from other continuous functions. The most well-known way to do this is the following whose proof trivially follows from the definition of continuity.

Theorem 2.1.15. *Let $(X, \mathcal{T}_X)$, $(Y, \mathcal{T}_Y)$, and $(Z, \mathcal{T}_Z)$ be topological spaces. If $f : X \to Y$ and $g : Y \to Z$ are continuous functions, then $g \circ f : X \to Z$ is a continuous function.*

Proof. To see that $g \circ f$ is a continuous function, let $U \in \mathcal{T}_Z$ be arbitrary. Then $g^{-1}(U) \in \mathcal{T}_Y$ as g is continuous thus $f^{-1}(g^{-1}(U)) \in \mathcal{T}_X$ as f is continuous. Hence $(g \circ f)^{-1}(U) = f^{-1}(g^{-1}(U)) \in \mathcal{T}_X$. Therefore, as $U \in \mathcal{T}_Z$ was arbitrary, $g \circ f$ is continuous by Definition 2.1.1. ∎

Another way to obtain continuous functions is by taking limits of a specific form.

Definition 2.1.16. Let $(X, \mathcal{T})$ be a topological space and let (Y, d) be a metric space. A net $(f_\lambda)_{\lambda\in\Lambda}$ in $\mathcal{F}(X, Y)$ is said to *converge uniformly* to a function $f : X \to Y$ if $(f_\lambda)_{\lambda\in\Lambda}$ converges to f with respect to the uniform metric; that is, for every $\epsilon > 0$ (without loss of generality, $\epsilon < 1$) there exists a $\lambda_0 \in \Lambda$ such that

$$d_{\text{unif}}(f_\lambda, f) = \sup_{x\in X} \min(\{d_Y(f_\lambda(x), f(x)), 1\}) < \epsilon$$

for all $\lambda \geq \lambda_0$.

Of course, as Definition 2.1.16 is a notion of convergence with respect to a metric, sequences are enough to determine the topology of convergence. However, as we often need to deal with nets in topological spaces, we do not restrict ourselves to sequences in the following essential result.

Theorem 2.1.17. *Let $(X, \mathcal{T})$ be a topological space, let (Y, d) be a metric space, and let $(f_\lambda)_{\lambda \in \Lambda}$ be a net of continuous functions from X to Y. If $(f_\lambda)_{\lambda \in \Lambda}$ converges uniformly to a function $f : X \to Y$, then f is continuous.*

Proof. To see that f is continuous, let $x_0 \in X$ and U a neighbourhood of $f(x_0)$ be arbitrary. Thus there exists an $\epsilon > 0$ such that $B_d(f(x_0), \epsilon) \subseteq U$. Without loss of generality by choosing a smaller ϵ if necessary, we may assume that $0 < \epsilon < 1$. Since $(f_\lambda)_{\lambda \in \Lambda}$ converges uniformly to f, there exists a $\lambda_0 \in \Lambda$ such that $d_{\mathrm{unif}}(f_{\lambda_0}, f) < \frac{\epsilon}{3}$. Hence $d(f_{\lambda_0}(x), f(x)) < \frac{\epsilon}{3}$ for all $x \in X$ by the definition of the uniform metric. However, as f_{λ_0} is continuous and as $B_d\left(f_{\lambda_0}(x_0), \frac{\epsilon}{3}\right)$ is an neighbourhood of $f_{\lambda_0}(x_0)$, there exists a neighbourhood V of x_0 such that $f_{\lambda_0}(V) \subseteq B_d\left(f_{\lambda_0}(x_0), \frac{\epsilon}{3}\right)$. Hence for all $x \in V$,

$$\begin{aligned} d(f(x), f(x_0)) &\leq d(f(x), f_{\lambda_0}(x)) + d(f_{\lambda_0}(x), f_{\lambda_0}(x_0)) + d(f_{\lambda_0}(x_0), f(x_0) \\ &< \frac{\epsilon}{3} + \frac{\epsilon}{3} + \frac{\epsilon}{3} = \epsilon. \end{aligned}$$

Thus $f(V) \subseteq B_d(f(x_0), \epsilon) \subseteq U$. Therefore, as x_0 and U were arbitrary, Theorem 2.1.9 implies that f is continuous. ∎

Another way to construct continuous function is to use inclusions and restrictions together with the subspace topology.

Lemma 2.1.18. *Let $(X, \mathcal{T}_X)$ and $(Y, \mathcal{T}_Y)$ be topological spaces, let $A \subseteq X$, and let $B \subseteq Y$. The following hold:*

(1) If A is equipped with the subspace topology, then the inclusion map $i : A \to X$ defined by $i(a) = a$ for all $a \in A$ is continuous.

(2) If A is equipped with the subspace topology and $f : X \to Y$ is continuous, then the restriction $f|_A : A \to Y$ defined by $f|_A(a) = f(a)$ for all $a \in A$ is continuous.

(3) If B is equipped with the subspace topology and $f : X \to B$ is continuous, then $f : X \to Y$ is continuous.

(4) If B is equipped with the subspace topology, $f : X \to Y$ is continuous, and $f(X) \subseteq B$, then $f : X \to B$ is continuous.

Proof. To see that (1) holds, notice for all open subsets U of X that $i^{-1}(U) = A \cap U$ is open in the subspace topology on A. Hence i is continuous by Definition 2.1.1.

To see that (2) holds, notice for all open subsets U of X that $f|_A^{-1}(U) = A \cap f^{-1}(U)$ is open in the subspace topology on A as $f^{-1}(U)$ is an open subset of X since f is continuous. Hence $f|_A$ is continuous by Definition 2.1.1.

To see that (3) holds, notice for all open subset V of Y that $f^{-1}(V) = f^{-1}(B \cap V)$ which must be open since $f : X \to B$ is continuous and $B \cap V$ is open in the subspace topology on B by definition. Hence $f : X \to Y$ is continuous by Definition 2.1.1.

Finally, to see that (4) holds, recall that if V is an open subset of B in the subspace topology that $V = B \cap V_0$ for some open subset V_0 in Y. Therefore, since $f(X) \subseteq B$, we see that $f^{-1}(V) = f^{-1}(B \cap V_0) = f^{-1}(V_0)$ is open in X as $f : X \to Y$ is continuous and V_0 is open in Y. Hence $f : X \to B$ is continuous by Definition 2.1.1. ∎

Instead of trying to restrict or compress a continuous function to obtain a continuous function, we can combine continuous functions to get continuous functions. Indeed the first of the following two results says that if we can cover a topological space with open sets and we have a function that is continuous on each of these open sets, then the function on the whole space must be continuous. The second result does the same for closed sets provided we have a finite number of closed sets with union all of X. Both of these results have uses in differential geometry.

Lemma 2.1.19. *Let $(X, \mathcal{T}_X)$ and $(Y, \mathcal{T}_Y)$ be topological spaces and let $f : X \to Y$. Suppose there exist $\{U_\alpha\}_{\alpha \in I} \subseteq \mathcal{T}$ be such that $X = \bigcup_{\alpha \in I} U_\alpha$ and $f|_{U_\alpha}$ is continuous for all $\alpha \in I$, then f is continuous.*

Proof. To see that f is continuous, let $V \in \mathcal{T}_Y$ be arbitrary. Notice for all $\alpha \in I$ that $f|_{U_\alpha}^{-1}(V)$ is open in U_α equipped with the subspace topology from X as $f|_{U_\alpha}$ is continuous. Hence, by the definition of the subspace topology, there exists a $V_\alpha \in \mathcal{T}_X$ such that

$$f|_{U_\alpha}^{-1}(V) = U_\alpha \cap V_\alpha.$$

However, since $U_\alpha \in \mathcal{T}_X$, we obtain that $f|_{U_\alpha}^{-1}(V) \in \mathcal{T}_X$ being the intersection of two elements of $\mathcal{T}_X$. Therefore, since

$$f^{-1}(V) = \bigcup_{\alpha \in I} f|_{U_\alpha}^{-1}(V),$$

we obtain that $f^{-1}(V) \in \mathcal{T}_X$ as $\mathcal{T}_X$ is closed under unions. Hence, as $V \in \mathcal{T}_Y$ was arbitrary, f is continuous by Definition 2.1.1. ∎

Theorem 2.1.20 (The Pasting Lemma). *Let $(X, \mathcal{T}_X)$ and $(Y, \mathcal{T}_Y)$ be a topological spaces, let A, B be closed subsets of X such that $X = A \cup B$, and let $f : A \to Y$ and $g : B \to Y$ be continuous functions such that $f(x) = g(x)$ for all $x \in A \cap B$. Then the function $h : X \to Y$ such that $h(a) = f(a)$ for all $a \in A$ and $h(b) = g(b)$ for all $b \in B$ is continuous.*

Proof. To see that h is continuous, let F be an arbitrary closed subset of Y. Notice by construction that

$$h^{-1}(F) = f^{-1}(F) \cup g^{-1}(F).$$

However, as f and g are continuous functions, Theorem 2.1.9 implies that $f^{-1}(F)$ is a closed subset of A when A is equipped with the subspace topology and $g^{-1}(F)$ is a closed subset of B when B is equipped with the subspace topology. Thus Lemma 1.6.12 implies that there exist closed subsets F_1 and F_2 in X such that $f^{-1}(F) = A \cap F_1$ and $g^{-1}(F) = B \cap F_2$. Therefore, as A and B are closed in X, $f^{-1}(F)$ and $g^{-1}(F)$ are closed in X. Thus $h^{-1}(F) = f^{-1}(F) \cup g^{-1}(F)$ is closed in X. Therefore, as F was arbitrary, h is continuous by Theorem 2.1.9. ∎

2.2 Homeomorphisms

With the construction of the objects and morphisms studied in this course complete, the next natural progression in mathematics is to define using ones morphisms when two objects are the same. As topological spaces are the objects in this course and continuous functions are the morphisms in this course, we study the following concept in order to determine the notion of when two topological spaces are the same.

Definition 2.2.1. Let $(X, \mathcal{T}_X)$ and $(Y, \mathcal{T}_Y)$ be topological spaces. A function $f : X \to Y$ is said to be a *homeomorphism* if f is bijective and both f and f^{-1} are continuous. Equivalently, a function $f : X \to Y$ is a homeomorphism if f is bijective and $U \in \mathcal{T}_X$ if and only if $f(U) \in \mathcal{T}_Y$.

Due to the above, we define the following notion.

Definition 2.2.2. Two topological spaces $(X, \mathcal{T}_X)$ and $(Y, \mathcal{T}_Y)$ are said to be *homeomorphic* if there is a homeomorphism from X to Y.

The reason why two homeomorphic topological spaces are 'the same' is because a bijection means the sets are the same (so $X = Y$ upto relabelling) and the continuity of the homeomorphism and its inverse implies the open sets are the same. This probably causes a modern mathematician to ask why we do not call homeomorphisms isomorphisms and why we do not call homeomorphic topological spaces isomorphic topological spaces. The only reason for this is tradition.

Of course, any notion of equality in mathematics must be an equivalence relation (see Section A.4), we verify the following.

Proposition 2.2.3. *Consider a set Φ of topological spaces and define a relation $\sim$ on Φ by $(X, \mathcal{T}_X) \sim (Y, \mathcal{T}_Y)$ if and only if $(X, \mathcal{T}_X)$ and $(Y, \mathcal{T}_Y)$ are homeomorphic. Then $\sim$ is an equivalence relation.*

Proof. First, clearly $(X, \mathcal{T}_X) \sim (X, \mathcal{T}_X)$ via the identity map. Secondly, if $(X, \mathcal{T}_X) \sim (Y, \mathcal{T}_Y)$, then there is a homeomorphism $f : X \to Y$. As $f^{-1} : Y \to X$ is then a homeomorphism by definition, $(Y, \mathcal{T}_Y) \sim (X, \mathcal{T}_X)$.

Finally suppose $(X, \mathcal{T}_X) \sim (Y, \mathcal{T}_Y)$ and $(Y, \mathcal{T}_Y) \sim (Z, \mathcal{T}_Z)$. Thus there exists homeomorphisms $f : X \to Y$ and $g : Y \to Z$. Consider the map $h = g \circ f : X \to Z$. We claim that h is a homeomorphism. Indeed as the composition of bijections is a bijection, h is a bijection. Furthermore, by Theorem 2.1.15 h is continuous being the composition of continuous functions. Finally, as $h^{-1} = f^{-1} \circ g^{-1}$, h^{-1} is the composition of continuous functions (as f and g are homeomorphisms) and thus continuous. Hence h is a homeomorphism so $(X, \mathcal{T}_X) \sim (Z, \mathcal{T}_Z)$ as desired. ∎

Now onto some examples.

Example 2.2.4. Let $\mathbb{R}$ be equipped with its canonical topology and let $A = \left(-\frac{\pi}{2}, \frac{\pi}{2}\right)$ be equipped with the subspace topology inherited from $\mathbb{R}$. Then $\mathbb{R}$ and A are homeomorphic. Indeed consider the function $f : A \to \mathbb{R}$ defined

$$f(x) = \tan(x)$$

for all $x \in \mathbb{R}$. It is well-known that f is a continuous bijective function on A whose inverse, namely $f^{-1}(x) = \arctan(x)$ is also continuous. Hence $\mathbb{R}$ and A are homeomorphic.

As often a topological space is only homeomorphic to a subspace, we define the following.

Definition 2.2.5. Let $(X, \mathcal{T}_X)$ and $(Y, \mathcal{T}_Y)$ be topological spaces. An function $f : X \to Y$ is said to be a *embedding* if $f : X \to f(X)$ is a homeomorphism when $f(X)$ is equipped with the subspace topology.

Example 2.2.6. Let $\mathbb{R}^2$ and $\mathbb{R}^3$ be equipped with their Euclidean topologies, and let

$$S^2 = \{(x, y, z) \in \mathbb{R}^3 \mid x^2 + y^2 + z^2 = 1\}$$

(that is, S^2 is the boundary of the unit ball in $\mathbb{R}^3$) equipped with the subspace topology from $\mathbb{R}^3$. Since the Euclidean topologies on $\mathbb{R}^2$ and $\mathbb{R}^3$ are product topologies by Example 1.4.18, Theorem 1.5.25 implies a net converges in either of these spaces if and only if it converges entry-wise. Hence Proposition 1.5.24 implies that a net converges in S^2 if and only if it converges entry-wise.

Consider the map $f : \mathbb{R}^2 \to S^2$ defined by

$$f(x, y) = \left(\frac{2x}{x^2 + y^2 + 1}, \frac{2y}{x^2 + y^2 + 1}, \frac{x^2 + y^2 - 1}{x^2 + y^2 + 1}\right)$$

for all $(x, y) \in \mathbb{R}^2$. It is not difficult to see that f is continuous by the net characterization of continuity from Theorem 2.1.9. However, f is not bijective. Indeed the point $(0, 0, 1) \in S^2$ is not in the range of f.

Consider the function

$$g : S^2 \setminus \{(0,0,1)\} \to \mathbb{R}^2$$

defined by

$$g(x,y,z) = \left(\frac{x}{1-z}, \frac{y}{1-z}\right)$$

for all $(x,y,z) \in S^2 \setminus \{(0,0,1)\}$. It is not difficult to see that if the codomain of f is restricted to $S^2 \setminus \{(0,0,1)\}$, then f and g are inverses to each other. Furthermore g is continuous by the net characterization of continuity from Theorem 2.1.9. Hence f is an embedding of $\mathbb{R}^2$ into S^2 and $\mathbb{R}^2$ and $S^2 \setminus \{(0,0,1)\}$ are homeomorphic.

Example 2.2.7. Let $\mathbb{R}$ and $\mathbb{R}^2$ be equipped with their Euclidean topologies, let $A = [0, 2\pi)$ equipped with the subspace topology induced by $\mathbb{R}$, and let

$$S^1 = \{(x,y) \in \mathbb{R}^2 \mid x^2 + y^2 = 1\}$$

equipped with the subspace topology induced by $\mathbb{R}^2$. Since the Euclidean topologies on $\mathbb{R}$ and $\mathbb{R}^2$ are product topologies by Example 1.4.18, Theorem 1.5.25 implies a net converges in either of these spaces if and only if it converges entry-wise. Hence Proposition 1.5.24 implies that a net converges in S^1 if and only if it converges entry-wise.

Consider the map $f : A \to S^1$ defined by

$$f(x) = (\cos(x), \sin(x))$$

for all $x \in A$. It is elementary to see that f is a bijection. It is not difficult to see that f is continuous by the net characterization of continuity from Theorem 2.1.9. However f^{-1} is not continuous. Indeed consider the set $U = [0, 1)$. Since U is an open subset of A as $U = A \cap (-\infty, 1)$, if f^{-1} were continuous, then $(f^{-1})^{-1}(U)) = f(U)$ would be open in S^1, so $S^1 \setminus f(U)$ would be closed in S^1. However, the sequence

$$\left(\left(\cos\left(2\pi - \frac{1}{n}\right), \sin\left(2\pi - \frac{1}{n}\right)\right)\right)_{n\geq 1}$$

is a net in $S^1 \setminus f(U)$ that converges to $(1,0) \in f(U)$ thereby contradicting the fact that $S^1 \setminus f(U)$ was closed. Hence f cannot be continuous.

The reason that the map f in Example 2.2.7 fails is that we have not placed the correct topology on the circle. If one wants a bijective map from a topological space to be a homeomorphism, we know exactly what topology to put on the codomain to ensure as the following result demonstrates.

Proposition 2.2.8. *Let $(X, \mathcal{T}_X)$ be a topological space, let Y be a non-empty set, and let $q : X \to Y$ be a surjective map. Let*

$$\mathcal{T}_Y = \{A \subseteq Y \mid q^{-1}(A) \in \mathcal{T}_X\}.$$

Then $\mathcal{T}_Y$ is the finest topology on Y such that q is continuous. If q is bijective, then q is a homeomorphism.

Proof. First, we claim that $\mathcal{T}_Y$ is a topology. To see this, we note that $\emptyset, Y \in \mathcal{T}_Y$ since $q^{-1}(\emptyset) = \emptyset \in \mathcal{T}_X$ and $q^{-1}(Y) = X \in \mathcal{T}_X$ as q is surjective and as $\mathcal{T}_X$ is a topology. Moreover, since for all $\{U_\alpha\}_{\alpha \in I} \subseteq \mathcal{P}(Y)$ we have that

$$q^{-1}\left(\bigcup_{\alpha \in I} U_\alpha\right) = \bigcup_{\alpha \in I} q^{-1}(U_\alpha) \qquad \text{and} \qquad q^{-1}\left(\bigcap_{\alpha \in I} U_\alpha\right) = \bigcap_{\alpha \in I} q^{-1}(U_\alpha),$$

it is elementary to see that $\mathcal{T}_Y$ is closed under unions and finite intersections since $\mathcal{T}_X$ is. Hence $\mathcal{T}_Y$ is a topology.

To see that $q : (X, \mathcal{T}_X) \to (Y, \mathcal{T}_Y)$ is continuous, we know that $q^{-1}(U) \in \mathcal{T}_X$ for all $U \in \mathcal{T}_Y$ by the definition of $\mathcal{T}_Y$. Hence q is continuous by definition. To see that $\mathcal{T}_Y$ is the finest topology on T so that q is continuous, suppose $\mathcal{T}$ is a topology on Y for which q is continuous. Then, by the definition of continuity, $q^{-1}(U) \in \mathcal{T}_X$ for all $U \in \mathcal{T}$. Therefore, by the definition of $\mathcal{T}_Y$, we obtain that $\mathcal{T} \subseteq \mathcal{T}_Y$. Hence $\mathcal{T}_Y$ is the finest topology on Y such that q is continuous.

Finally, to see that q is a homeomorphism when q is bijective, we need simply check that q^{-1} is continuous. To see this, let $V \in \mathcal{T}_X$ be arbitrary. Then $(q^{-1})^{-1}(V) = q(V)$ will be an element of $\mathcal{T}_Y$ by definition as $q^{-1}(q(V)) = V \in \mathcal{T}_X$. Therefore, as $V \in \mathcal{T}_X$ was arbitrary, q is a homeomorphism as desired. ∎

Due to the importance and usefulness of the above topology, we name this topology as follows.

Definition 2.2.9. Let $(X, \mathcal{T}_X)$ be a topological space, let Y be a non-empty set, and let $q : X \to Y$ be a surjective map. The topology

$$\mathcal{T}_Y = \{A \subseteq Y \mid q^{-1}(A) \in \mathcal{T}_X\}$$

from Proposition 2.2.8 is called the *quotient topology on Y induced by q*.

The reason we call the above the quotient topology is that one is really identifying all of the points in $q^{-1}(\{y\})$ as a single point for all $y \in Y$ and placing a topology on these collections of points based on the original topology; that is, we are taking a 'quotient' of a topological space by identify points. This idea is also motivated from geometry by 'cutting-and-pasting' to identify points to create new geometric objects. Before we formalize this and explore some examples, we first demonstrate, like with all things, how bases work in the quotient topology.

Proposition 2.2.10. *Let $(X, \mathcal{T}_X)$ be a topological space, let Y be a non-empty set, let $q : X \to Y$ be a surjective map, and let $\mathcal{T}_Y$ be the quotient topology on Y induced by q. If $\mathcal{B}_X$ is a basis for $(X, \mathcal{T}_X)$, then*

$$\mathcal{B}_Y = \{A \subseteq Y \mid q^{-1}(A) \in \mathcal{B}_X\}$$

is a basis for $(Y, \mathcal{T}_Y)$.

Proof. To see that $\mathcal{B}_Y$ is a basis for $(Y, \mathcal{T}_Y)$, let $y \in Y$ and $U \in \mathcal{T}_Y$ be arbitrary. By the definition of the quotient topology, $q^{-1}(U) \in \mathcal{T}_X$. Therefore, as $q^{-1}(y) \in q^{-1}(U)$, the fact that $\mathcal{B}_X$ is a basis for $(X, \mathcal{T}_X)$ implies that there exists a $B \in \mathcal{B}_X$ such that $q^{-1}(y) \subseteq B \subseteq q^{-1}(U)$. Therefore, if $B' = q(B) \subseteq Y$, then $B' \in \mathcal{B}_Y$ by definition and $y \in q(B) = B' \subseteq U$. Therefore, as y and U were arbitrary, $\mathcal{B}_Y$ is a basis for $(Y, \mathcal{T}_Y)$. ∎

Example 2.2.11. Let $A = [0, 2\pi)$ and let

$$S^1 = \{(x, y) \in \mathbb{R}^2 \mid x^2 + y^2 = 1\} \subseteq \mathbb{R}^2$$

equipped with the subspace topology induced by $\mathbb{R}^2$. As the intersection of open balls in $\mathbb{R}^2$ with S^1 yield open arcs on S^1, the open arcs on S^1 are a basis for the subspace topology on S^1.

Consider the map $q : S^1 \to A$ defined by

$$q\left((\cos(x), \sin(x))\right) = x$$

for all $x \in [0, 2\pi)$ and let $\mathcal{T}$ be the quotient topology on A induced by q. By the description of the basis of S^1 given above and since only arcs of arbitrarily small length around a point matter in forming a neighbourhood basis, we see for all $x \in (0, 2\pi)$ that

$$\{(x - \epsilon, x + \epsilon) \mid 0 < \epsilon < \min\{x, 2\pi - x\}\}$$

form a neighbourhood basis of x and that

$$\{[0, \epsilon) \cup (2\pi - \epsilon, 2\pi) \mid 0 < \epsilon < 2\pi\}$$

for a neighbourhood basis of 0 in the quotient topology.

Example 2.2.12. Let $\mathbb{R}$ be equipped with its usual topology, let $Y = \{a, b, c\}$, and let $q : \mathbb{R} \to Y$ be defined by

$$q(x) = \begin{cases} a & \text{if } x < 0 \\ b & \text{if } x = 0 \\ c & \text{if } x > 0 \end{cases}.$$

Let $\mathcal{T}$ be the quotient topology on Y induced by q. Notice for all $A \subseteq Y$ that

$$q^{-1}(A) = \begin{cases} \emptyset & \text{if } A = \emptyset \\ (-\infty, 0) & \text{if } A = \{a\} \\ \{0\} & \text{if } A = \{b\} \\ (0, \infty) & \text{if } A = \{c\} \\ (-\infty, 0] & \text{if } A = \{a, b\} \\ (-\infty, 0) \cup (0, \infty) & \text{if } A = \{a, c\} \\ [0, \infty) & \text{if } A = \{b, c\} \\ \mathbb{R} & \text{if } A = Y \end{cases}.$$

Therefore, by our knowledge of the open subsets of $\mathbb{R}$, we see that

$$\mathcal{T} = \{\emptyset, \{a\}, \{c\}, \{a, c\}, \{a, b, c\}\}.$$

Diagrammatically, the topology $\mathcal{T}$ is the following.

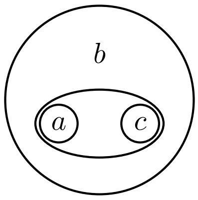

In Example 2.2.12, one can think of a as $(-\infty, 0) \subseteq \mathbb{R}$, b as $\{0\} \subseteq \mathbb{R}$, and c as $(0, \infty) \subseteq \mathbb{R}$. That is, we can think of Y as a partition of $\mathbb{R}$ and the quotient topology on Y is then induced by this partition. To formalize this, we recall the definition a partition and how a partition produces a topological space.

Definition 2.2.13. Let X be a non-empty set. A *partition of* X is a collection $\{X_\alpha\}_{\alpha \in I} \subseteq \mathcal{P}(X)$ such that $X = \bigcup_{\alpha \in I} X_\alpha$ and $X_\alpha \cap X_\beta = \emptyset$ if $\alpha, \beta \in I$ and $\alpha \neq \beta$.

Definition 2.2.14. Let $(X, \mathcal{T})$ be a topological space, let X^* be a partition of X, and let $q : X \to X^*$ be the surjective map that maps each element $x \in X$ to the unique element in X^* containing x. The pair $(X^*, \mathcal{T}^*)$ where $\mathcal{T}^*$ is quotient topology on X^* induced by q is called a *quotient space*.

Example 2.2.15. Let $\mathbb{R}$ be equipped with its usual topology and let

$$\mathcal{P} = \{\{x + 2\pi n \mid n \in \mathbb{Z}\} \mid x \in [0, 2\pi)\}.$$

If $(\mathbb{R}^*, \mathcal{T}^*)$ is the quotient space induced by $\mathcal{P}$, then $\mathbb{R}^*$ is in canonical bijective correspondence with $[0, 2\pi)$ by identifying $\{x + 2\pi n \mid n \in \mathbb{Z}\}$ for $x \in [0, 2\pi)$ with x. Under this identification, if $q : X \to X^*$ is the surjective map that maps each element $x \in \mathbb{R}$ to the unique element in $\mathbb{R}^*$ containing x, then we recall that

$$\mathcal{T}^* = \{A \subseteq [0, 2\pi) \mid q^{-1}(A) \text{ is open in } \mathbb{R}\}.$$

As the inverse image of every basis element exhibited in Example 2.2.11 is a union of a countable number of open intervals (each of which is a translate of one fixed open interval by an integer multiple of 2π), we see that the topology from Example 2.2.11 must be coarser than $\mathcal{T}^*$. Furthermore, given a subset $A \subseteq [0, 2\pi)$ we see that $q^{-1}(A)$ is 2π-periodic and will be open if and only if it is a union of open intervals and closed under 2π-periodicity. Hence $\mathcal{T}^*$ is precisely the topology on $[0, 2\pi)$ exhibited in Example 2.2.11.

Example 2.2.16. Let $\mathbb{R}^2$ be equipped with the Euclidean topology, let $A = \{(x, y) \in \mathbb{R}^2 \mid x^2 + y^2 \leq 1\}$ equipped with the subspace topology induced by $\mathbb{R}^2$, and let

$$\mathcal{P} = \{\{(x, y)\} \mid x^2 + y^2 < 1\} \cup \{(x, y) \mid x^2 + y^2 = 1\}.$$

Consider the quotient space $(A^*, \mathcal{T}^*)$ and let $q : A \to A^*$ be the canonical surjective map. Then A^* is canonically in bijective correspondence with the shell

$$S^2 = \{(x, y, z) \in \mathbb{R}^3 \mid x^2 + y^2 + z^2 = 1\} \subseteq \mathbb{R}^3$$

via the the map $f : A^* \to S^2$ defined by

$$f(r\cos(\theta), r\sin(\theta)) = (\sin(r\pi)\cos(\theta), \sin(r\pi)\sin(\theta), \cos(r\pi))$$

for all $r \in [0, 1]$ and $\theta \in [0, 2\pi)$. If S^2 is equipped with the subspace topology inherited from $\mathbb{R}^3$, then f is a homeomorphism from A^* to S^2. To see this, first notice that a subset of $\{(x, y) \mid x^2 + y^2 < 1\}$ is open in A^* if and only if it is open in A and thus open in $\mathbb{R}^2$ by definition. Next, suppose U is an open set in A^* that contains $S^1 = \{(x, y) \mid x^2 + y^2 = 1\}$. Thus $q^{-1}(U)$ is open in the subspace topology on A and contains S^1. Note for each point $(x, y) \in S^1$ there must exist a $\delta_{(x,y)} > 0$ such that $B((x, y), \delta_{(x,y)}) \cap A \subseteq q^{-1}(U)$. As S^1 is compact (see Chapter 3), we may cover S^1 with a finite number of these balls in which case an $\epsilon > 0$ may be found so that $\{(x, y) \mid \epsilon < x^2 + y^2 \leq 1\} \subseteq q^{-1}(U)$. Consequently, we see that $q^{-1}(U)$ is a union of a set of the form $\{(x, y) \mid \epsilon < x^2 + y^2 \leq 1\}$ and a subset of $\{(x, y) \mid x^2 + y^2 < 1\}$ that is open in A^*. It is then not difficult to see that the open sets in A^* are in bijective correspondence with those of S^2 via f. Hence f is a homeomorphism from A^* to S^2.

Example 2.2.17. Let $\mathbb{R}^2$ be equipped with the Euclidean topology, let $A = [0, 1]^2 \subseteq \mathbb{R}^2$ equipped with the subspace topology induced by $\mathbb{R}^2$, and let $\mathcal{P}$ be the union of

$$\begin{aligned}&\{\{(x,y)\} \mid x,y \in (0,1)\},\\&\{\{(x,0),(x,1)\} \mid x \in (0,1)\},\\&\{\{(0,y),(1,y)\} \mid y \in (0,1)\}, \text{ and}\\&\{(0,0),(1,0),(0,1),(1,1)\}.\end{aligned}$$

Consider the quotient space $(A^*, \mathcal{T}^*)$ and let $q : A \to A^*$ be the canonical surjective map. Then A^* is canonically in bijective correspondence to a torus in $\mathbb{R}^3$ in such a way that that $\mathcal{T}^*$ corresponds to the subspace topology on the torus inherited from $\mathbb{R}^3$. The details are similar to Example 2.2.16.

To better understand functions, continuous functions, and homeomorphisms on quotient spaces, we give a name to the maps under consideration when definition a quotient topology.

Definition 2.2.18. Let $(X, \mathcal{T}_X)$ and $(Y, \mathcal{T}_Y)$ be topological spaces. A map $q : X \to Y$ is said to be a *quotient map* if q is surjective and a set $U \subseteq Y$ is open if and only if $p^{-1}(U) \in \mathcal{T}_X$; that is, if $\mathcal{T}_Y$ is the quotient topology on Y induced by q.

Clearly quotient maps are continuous maps by the definition of a quotient map and by the definition of a continuous function. In addition, of course a quotient of a quotient is still a quotient.

Lemma 2.2.19. *Let $(X, \mathcal{T}_X)$, $(Y, \mathcal{T}_Y)$, and $(Z, \mathcal{T}_Z)$ be topological spaces, let $q : X \to Y$ be a quotient map, and let $p : Y \to Z$ be a quotient map. Then $p \circ q : X \to Z$ is a quotient map.*

Proof. To see that $p \circ q$ is a quotient map, we first note that quotient maps are continuous by definition. Therefore $p \circ q$ is a composition of continuous maps and thus continuous. Hence if $U \in \mathcal{T}_Z$, then $(p \circ q)^{-1}(U)$ is open in $\mathcal{T}_X$ by continuity.

Conversely, let $U \subseteq Z$ such that $(p \circ q)^{-1}(U) = q^{-1}(p^{-1}(U))$ is open in $(X, \mathcal{T}_X)$ be arbitrary. Since q is a quotient map, $q^{-1}(p^{-1}(U))$ being open in $(X, \mathcal{T}_X)$ implies that $p^{-1}(U)$ is open in $(Y, \mathcal{T}_Y)$ by the definition of a quotient map. Therefore, since p is a quotient map, $p^{-1}(U)$ being open in $(Y, \mathcal{T}_Y)$ implies that U is open in $(Z, \mathcal{T}_Z)$ by the definition of a quotient map. Therefore, as U was arbitrary, $p \circ q$ is a quotient map. ∎

One of the main reason quotient spaces are nice is that certain maps factor over quotients and preserve topological properties.

Theorem 2.2.20. *Let $(X, \mathcal{T}_X)$, $(Y, \mathcal{T}_Y)$, and $(Z, \mathcal{T}_Z)$ be topological spaces, let $q : X \to Z$ be a quotient map, and let $g : X \to Y$ be a map that is constant* *on $q^{-1}(\{z\})$ for each $z \in Z$. Then there exists a unique map $f : Z \to Y$ such that $g = f \circ q$.*

$$\begin{tikzcd} X \arrow[d, "q"'] \arrow[dr, "g"] & \\ Z \arrow[r, "f"'] & Y \end{tikzcd}$$

The map f is continuous if and only if g is continuous. Furthermore, f is a quotient map if and only if g is a quotient map.

Proof. First, since g is constant on $q^{-1}(\{z\})$ for each $z \in Z$, we define $f : Z \to Y$ by setting $f(z)$ for each $z \in Z$ to be the unique value of g obtained on $q^{-1}(\{z\})$, then f is a well-defined function such that $g = f \circ q$ as desired. Furthermore, as this is clearly the only way to define f so that $g = f \circ q$ as q is surjective, uniqueness has been obtained.

Next, clearly if f is a continuous function than g is a continuous function since quotient maps are continuous and the composition of continuous functions is continuous. Conversely, suppose that g is continuous and let U be an arbitrary open set in $(Y, \mathcal{T}_Y)$. Hence $g^{-1}(U) = (f \circ q)^{-1}(U) = q^{-1}(f^{-1}(U))$ must be an open set in X. However, as q is a quotient map, $q^{-1}(f^{-1}(U)) \in \mathcal{T}_X$ implies $f^{-1}(U) \in \mathcal{T}_Z$. Hence, as U was arbitrary f is continuous as desired.

Finally, if f is a quotient map, then g is a quotient map as the composition of quotient maps is a quotient map. Conversely, suppose that g is a quotient map. Thus a set $U \subseteq Y$ is such that $U \in \mathcal{T}_Y$ if and only if $g^{-1}(U) = (f \circ q)^{-1}(U) = q^{-1}(f^{-1}(U)) \in \mathcal{T}_X$. However, as q is a quotient map, $q^{-1}(f^{-1}(U)) \in \mathcal{T}_X$ if and only if $f^{-1}(U) \in \mathcal{T}_Z$. Hence f is a quotient map by definition. ∎

Using Theorem 2.2.20, we obtain a better understanding of quotient spaces obtained by partitioning based on a surjective continuous linear map. In particular, every surjective continuous linear map factors through a quotient space.

Corollary 2.2.21. *Let $(X, \mathcal{T}_X)$ and $(Y, \mathcal{T}_Y)$ be topological spaces, let $g : X \to Y$ be a surjective continuous linear map, let*

$$X^* = \{g^{-1}(\{y\}) \mid y \in Y\}$$

equipped with the quotient topology, and let $q : X \to X^*$ be the surjective map from Definition 2.2.14 that maps each element $x \in X$ to the unique element in X^* containing x. Then there exists a unique bijective, continuous map

$f : X^* \to Y$ such that $g = f \circ q$.

$$\begin{tikzcd} X \arrow[d, "q"'] \arrow[dr, "g"] & \\ X^* \arrow[r, "f"'] & Y \end{tikzcd}$$

Furthermore X^* is Hausdorff if $(Y, \mathcal{T}_Y)$ is Hausdorff. Finally f is a homeomorphism if and only if g a quotient map.

Proof. First, the fact that f exists is unique, and is continuous follows from Theorem 2.2.20 as g is continuous. Furthermore, since g is surjective and $g = f \circ q$, f is surjective. To see that f is injective, suppose $x_1, x_2 \in X^*$ are such that $f(x_1) = f(x_2)$. As q is surjective, there exists $x_1', x_2' \in X$ such that $q(x_1') = x_1$ and $q(x_2') = x_2$. Hence

$$g(x_1') = f(q(x_1')) = f(x_1) = f(x_2) = f(q(x_2')) = g(x_2').$$

Therefore, by the definition of X^* we must have $x_1 = q(x_1') = q(x_2') = x_2$. Thus f is bijective.

Next, suppose $(Y, \mathcal{T}_Y)$ is Hausdorff. To see that X^* is Hausdorff, let $x_1, x_2 \in X^*$ be arbitrary points such that $x_1 \neq x_2$. Then, as f is bijective, $f(x_1) \neq f(x_2)$. Hence, as $(Y, \mathcal{T}_Y)$ is Hausdorff, there exists open sets $U_1, U_2 \in \mathcal{T}_Y$ such that $f(x_1) \in U_1$, $f(x_2) \in U_2$, and $U_1 \cap U_2 = \emptyset$. Therefore, since f is a continuous bijection, $V_1 = f^{-1}(U_1)$ and $V_2 = f^{-1}(U_2)$ are open sets in X^* such that $x_1 \in V_1$, $x_2 \in V_2$, and $V_1 \cap V_2 = \emptyset$. Therefore, as x_1 and x_2 were arbitrary, X^* is Hausdorff.

To see the last part of the statement, we note that if g is a quotient map, then f is a quotient map by Theorem 2.2.20. Therefore, as f is a bijective quotient map, f is a homeomorphism by definition.

Finally, suppose f is a homeomorphism. To see that g is a quotient map, we first notice g is surjective and, since g is continuous, that if $U \in \mathcal{T}_Y$ then $g^{-1}(U) \in \mathcal{T}_X$. Thus, to complete the proof that g is a quotient map, let $U \subseteq Y$ be an arbitrary set such that $g^{-1}(U) \in \mathcal{T}_X$. Hence $g^{-1}(U) = q^{-1}(f^{-1}(U)) \in \mathcal{T}_X$. Therefore, as q is a quotient map, $f^{-1}(U)$ is open in X^*. However, since f is a homeomorphism, this implies that $U \in \mathcal{T}_Y$. Therefore, as U was arbitrary, g is a quotient map. ∎

Of course, one may ask, "Is every quotient space is Hausdorff?" The answer to this question is of course no as Example 2.2.12 is an example of a quotient space of $\mathbb{R}$ that is not Hausdorff. Thus Corollary 2.2.21 is the best we can do to ensure a quotient space is Hausdorff.

2.3 Connectedness

Unsurprisingly, continuous functions preserve and play an important role in many topological properties one may desire to study. The first topological property we desire to study together with continuous functions is a property that enables a generalization of the Intermediate Value Theorem from calculus. Indeed it is not a specific property of the real numbers that enables the Intermediate Value Theorem but a specific topological property of intervals in $\mathbb{R}$.

Definition 2.3.1. A topological space $(X, \mathcal{T})$ is said to be *connected* if there does not exists $U, V \in \mathcal{T} \setminus \{\emptyset\}$ such that $U \cap V = \emptyset$ and $X = U \cap V$. Equivalently, X is connected if the only subsets of X that are both open and closed are $\emptyset$ and X.

For our first example, we describe all connected subspaces of $\mathbb{R}$.

Theorem 2.3.2. *Let $\mathbb{R}$ be equipped with its canonical topology and let $A \subseteq \mathbb{R}$ be equipped with its subspace topology. Then A is connected if and only if A is an interval (singletons count as intervals here).*

Proof. Suppose that A is not an interval. To see that A is not connected, note since A is not an interval that there exists $x, y \in A$ and $z \in \mathbb{R} \setminus A$ such that $x < z < y$. Therefore, since $U = (-\infty, z) \cap A$ and $V = (z, \infty) \cap A$ are open sets in A such that $x \in U$ so $U \neq \emptyset$, $y \in V$ so $V \neq \emptyset$, $U \cap V = \emptyset$, and $U \cup V = A \setminus \{z\} = A$, A is not connected by definition.

To see the converse, let A be an interval in $\mathbb{R}$. Suppose to the contrary that A is not connected. Hence there exists non-empty open subsets U and V of A such that $U \cap V = \emptyset$ and $U \cup V = A$. As U and V are non-empty, select $a \in U$ and $b \in V$. As $U \cap V \neq \emptyset$, it must be the case that $a \neq b$. By exchanging the labelling of U and V if necessary, we may assume that $a < b$. Since A is an interval, we must have that $[a, b] \subseteq A$. Therefore, if

$$U' = U \cap [a, b] \qquad \text{and} \qquad V' = V \cap [a, b],$$

then U' and V' are non-empty open subsets $[a, b]$ equipped with the subspace topology (as a subspace of a subspace is a subspace) such that $a \in U'$, $b \in V'$, $U' \cap V' = \emptyset$ and $U' \cup V' = [a, b]$.

Since $U' \neq \emptyset$ and $U' \subseteq [a, b]$, the scalar

$$\alpha = \sup(U')$$

is an element of $[a, b]$. Thus, as $[a, b] = U' \cup V'$, either $\alpha \in U'$ or $\alpha \in V'$. If $\alpha \in U'$, then by the definition of an open subsets of $[a, b]$ there exists an $\epsilon > 0$ such that $(\alpha - \epsilon, \alpha + \epsilon) \cap [a, b] \subseteq U$. If $\alpha + \epsilon < b$ then $\alpha + \epsilon \in U'$. However, this contradicts the fact that $\alpha = \sup(U')$. Hence it must be the case that $b \in (\alpha-\epsilon, \alpha+\epsilon) \cap [a, b] \subseteq U'$ which contradicts the fact that $b \in V'$ and $U' \cap V' = \emptyset$. Hence $\alpha \notin U'$ so $\alpha \in V'$ by the above paragraph. However, if $\alpha \in V'$, then by the definition of an open subsets of $[a, b]$ there exists an $\epsilon > 0$ such that $(\alpha-\epsilon, \alpha+\epsilon) \cap [a, b] \subseteq V'$. Therefore, as $U' \cap V' = \emptyset$, it must be the case that $(\alpha-\epsilon, \alpha] \cap U' = \emptyset$ thereby contradicting the fact that $\alpha = \sup(U)$. Hence $\alpha \notin V'$. Therefore we have a contradiction. Hence A is connected. ∎

Based on the above, it is important to consider connected subspaces of topological spaces.

Proposition 2.3.3. *Let $(X, \mathcal{T})$ be a topological space and let $A \subseteq X$ be non-empty. Then A is connected when equipped with the subspace topology if and only if there does not exist $U, V \in \mathcal{T}$ such that $U \cap A \neq \emptyset$, $V \cap A \neq \emptyset$, $A \subseteq U \cup V$, and $U \cap V \cap A = \emptyset$.*

Proof. This follows trivially from the definition of a connected topological space (Definition 2.3.1) and the definition of the subspace topology (Definition 1.4.2). ∎

Example 2.3.4. The condition $U \cap V \cap A = \emptyset$ in Proposition 2.3.3 cannot be replaced with simply $U \cap V = \emptyset$ in general. Indeed consider $X = \{a, b, c\}$ together with the topology $\mathcal{T} = \{\emptyset, X, \{b\}, \{a, b\}, \{b, c\}\}$.

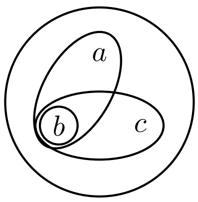

The set $A = \{a, c\}$ is not a connected subspace of $(X, \mathcal{T})$. Indeed notice the subspace topology on A is $\mathcal{T}_A = \{\emptyset, Y, \{a\}, \{c\}\}$ so clearly $U = \{a\}$ and $V = \{c\}$ imply that $(A, \mathcal{T}_A)$ is not connected.

Suppose $U, V \in \mathcal{T}$ are such that $U \cap A \neq \emptyset$, $V \cap A \neq \emptyset$, and $U \cap V = \emptyset$. As $A = \{a, c\}$, by interchanging U and V if necessary, it must be the case that $a \in U$ and $c \in V$. But by the definition of $\mathcal{T}$, either $U = \{a, b\}$ or $U = A$, and either $V = \{b, c\}$ or $V = \mathcal{A}$. Hence $b \in U \cap V$ thereby contradicting the fact that $U \cap V = \emptyset$. Hence the $U \cap V \cap A = \emptyset$ in Proposition 2.3.3 cannot be replaced with simply $U \cap V = \emptyset$ in general.

Before we exhibit more examples of connected and not connected topologies, we emphasize the fact that the notion of connected topological spaces behaves well with respect to continuous maps and yields the most general version of the Intermediate Value Theorem possible.

Theorem 2.3.5 (The Intermediate Value Theorem). *Let $(X, \mathcal{T}_X)$ and $(Y, \mathcal{T}_Y)$ be topological spaces and let $f : X \to Y$ be continuous. If $(X, \mathcal{T}_X)$ is connected, $f(X)$ is connected when equipped with the subspace topology inherited from Y.*

Proof. Suppose to the contrary that $f(X)$ is not connected. Hence, by the definition of the subspace topology, there exists $U, V \in \mathcal{T}_Y$ such that $U \cap f(X) \neq \emptyset$, $V \cap f(X) \neq \emptyset$, $U \cap V \cap f(X) = \emptyset$, and $f(X) \subseteq U \cup V$. Since f is a continuous function, $U' = f^{-1}(U) \in \mathcal{T}_X$ and $V' = f^{-1}(V) \in \mathcal{T}_X$. Moreover, by the assumptions on U and V, we must have that U' and V' are non-empty, $U' \cap V' = \emptyset$, and $U' \cup V' = X$. However, this contradicts the fact that $(X, \mathcal{T}_X)$ is connected. Hence $f(X)$ must be connected. ∎

Using the Intermediate Value Theorem (Theorem 2.3.5), we immediately obtain the usual Intermediate Value Theorem studied in undergraduate calculus. In particular, the proof of the undergraduate version of the Intermediate Value Theorem is very similar to the proof of Theorem 2.3.2 together with properties of continuous functions.

Corollary 2.3.6. *Let $(X, \mathcal{T})$ be a connected topological space (for example $X = [a, b]$) and let $f : X \to \mathbb{R}$ be a continuous map with respect to the canonical topology $\mathbb{R}$. If $a, b \in X$ and $y \in \mathbb{R}$ are such that $f(a) < y < f(b)$ or $f(b) < y < f(a)$, then there exists a $c \in X$ such that $f(c) = y$.*

Proof. By the Intermediate Value Theorem (Theorem 2.3.5) the set $f(X)$ is connected in $\mathbb{R}$. Hence Theorem 2.3.2 implies that $f(X)$ is an interval. Therefore, $a, b \in X$ and $y \in \mathbb{R}$ are such that $f(a) < y < f(b)$ or $f(b) < y < f(a)$, then as $f(a), f(b) \in f(X)$ and $f(X)$ is an interval, $y \in f(X)$ as desired. ∎

Although the following Intermediate Value Theorem is stated for continuous functions from connected topological spaces into $\mathbb{R}$, we can replace $\mathbb{R}$ with any totally ordered set $(Y, <)$ provided $(Y, <)$ has the Least Upper Bound property and the sets $\{y \in Y \mid a < y < b\}$ are non-empty for any $a, b \in Y$ and are bases for the topology (i.e. generalize Theorem 2.3.2 under these assumptions).

Using Corollary 2.3.6 in connection with Theorem 2.3.2, it is elementary to develop additional examples of connected sets.

Example 2.3.7. Let $\mathbb{R}^2$ be equipped with its Euclidean topology and let

$$A_1 = \{(x, 0) \mid x \in \mathbb{R}\} \quad \text{and} \quad A_2 = \left\{ \left(x, \frac{1}{x}\right) \,\middle|\, x \in \mathbb{R}, x > 0 \right\}.$$

Then A_1 is connected as it is the image of the continuous function $f_1 : \mathbb{R} \to \mathbb{R}^2$ defined by $f_1(x) = (x, 0)$ and $\mathbb{R}$ is connected. Similarly, A_2 is connected as it is the image of the continuous function $f_2 : (0, \infty) \to \mathbb{R}^2$ defined by $f_2(x) = \left(x, \frac{1}{x}\right)$ and $(0, \infty)$ is connected.

However, $A = A_1 \cup A_2$ is not a connected subsets of $\mathbb{R}^2$ even though

$$\operatorname{dist}(A_1, A_2) = \inf(\{\|\vec{x} - \vec{y}\|_2 \mid \vec{x} \in A_1, \vec{y} \in A_2\}) = 0.$$

To see that A is not connected, notice that

$$F = \{(x, y) \mid x \in \mathbb{R}, y \leq 0\}$$

is closed in $\mathbb{R}^2$ so $A_1 = A \cap F$ is closed in A by the definition of the subspace topology. However

$$U = \left\{(x, y) \,\middle|\, x \in \mathbb{R}, y < \frac{1}{|x|}\right\}$$

is open in $\mathbb{R}^2$ so $A_1 = A \cap U$ is open in A by the definition of the subspace topology. Therefore A_1 is a non-trivial subset of A that is both open and closed. Hence A is not connected.

One way to rectify the above in order to combine connected sets to obtain a new connected set is the following.

Proposition 2.3.8. *Let $\{(A_\alpha, \mathcal{T}_\alpha)\}_{\alpha \in I}$ be a collection of connected subspaces of a topological space $(X, \mathcal{T})$. If $\bigcap_{\alpha \in I} A_\alpha \neq \emptyset$, then $\bigcup_{\alpha \in I} A_\alpha$ is a connected subspace of $(X, \mathcal{T})$.*

Proof. Let $Y = \bigcup_{\alpha \in I} A_\alpha$ and choose $y \in \bigcap_{\alpha \in I} A_\alpha$. Suppose to the contrary that Y is not connected. Hence there exists $U, V \in \mathcal{T}$ such that $U \cap Y \neq \emptyset$, $V \cap Y \neq \emptyset$, $U \cap V \cap Y = \emptyset$, and $Y \subseteq U \cup V$. By exchanging the labelling of U and V if necessary, we may assume that $y \in V$. Notice for all $\alpha \in I$ that the sets $U_\alpha = U \cap A_\alpha$ and $V_\alpha = V \cap A_\alpha$ are open subsets of A_α when equipped with the subspace topology (a subspace of a subspace is a subspace) such that $U_\alpha \cap V_\alpha = \emptyset$ and $U_\alpha \cup V_\alpha = A_\alpha$. Therefore, as A_α is connected, it must be the case that $U_\alpha = \emptyset$ or $V_\alpha = \emptyset$. However, as $y \in A_\alpha$ for all $\alpha \in I$ and $y \in V$, it must be the case that $y \in V_\alpha$ for all $\alpha \in I$. Thus we must have that $U_\alpha = \emptyset$ for all $\alpha \in I$ and thus

$$U = U \cap A = \bigcup_{\alpha \in I} U \cap A_\alpha = \bigcup_{\alpha \in I} U_\alpha = \emptyset$$

thereby contradicting the fact that $U \neq \emptyset$. Hence Y is connected. ∎

Example 2.3.9. Let $\mathbb{R}^2$ be equipped with its canonical topology. For each $q \in \mathbb{Q}$, let

$$X_q = \{(x, qx) \mid x \in \mathbb{R}\}.$$

Clearly X_q is a connected subspace of $\mathbb{R}^2$ for each $q \in \mathbb{Q}$ as X_q is the continuous image of $\mathbb{R}$. Therefore, as $(0,0) \in X_q$ for all $q \in \mathbb{Q}$, the set

$$X = \bigcup_{q \in \mathbb{Q}} X_q$$

is a connected subset of $\mathbb{R}$.

Requiring a singleton in multiply topological spaces is not the only way the union of connected topological spaces is connected. As the following example shows, even the union of two disjoint connected sets may be connected.

Example 2.3.10. The *topologist's sine curve* is the set

$$X = \{(0,y) \mid -1 \leq y \leq 1\} \cup \left\{ \left(x, \sin\left(\frac{1}{x}\right)\right) \,\middle|\, 0 < x \leq \frac{1}{\pi} \right\}$$

equipped with the subspace topology inherited from the canonical topology on $\mathbb{R}^2$. Clearly X is the union of the disjoint sets

$$A_1 = \{(0,y) \mid -1 \leq y \leq 1\} \quad \text{and} \quad A_2 = \left\{ \left(x, \sin\left(\frac{1}{x}\right)\right) \,\middle|\, 0 < x \leq \frac{1}{\pi} \right\}.$$

Furthermore A_1 and A_2 are connected being continuous images of connected sets.

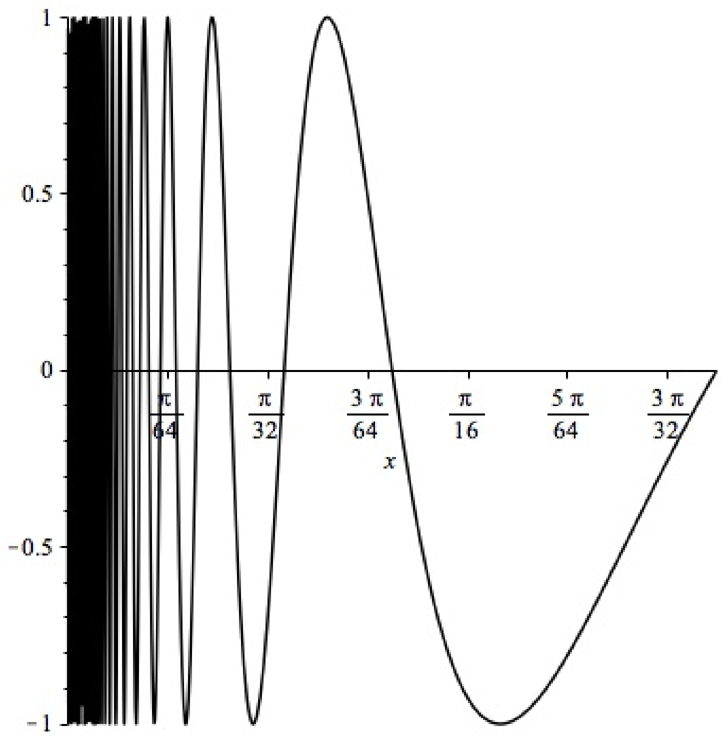

We claim that X is connected. However, this does not follow from Proposition 2.3.8 as $A_1 \cap A_2 = \emptyset$. To see that X is connected, we first claim that $X = \overline{A_2}$ where the closure is computed in $\mathbb{R}^2$. Indeed it is trivial that $A_2 \subseteq \overline{A_2}$. To see that $A_1 \subseteq \overline{A_2}$, for each $y \in [-1, 1]$ choose $z \in [0, \pi]$ such that $\sin(z) = y$. Then $\left(\left(\frac{1}{z+2\pi n}, y\right)\right)_{n\geq 1}$ is a sequence that is eventually in X that converges to $(0, y)$. Hence $A_1 \subseteq \overline{A_2}$. Finally, the fact that $\overline{A_2} \subseteq A_1 \cup A_2$ follows from the fact that the only possible cluster points of A_2 are contained in $A_1 \cup A_2$ as the x-terms must converge to a point x_0 in $\left[0, \frac{1}{\pi}\right]$ for which the y-term must converge to $\sin\left(\frac{1}{x_0}\right)$ if x_0 is non-zero and otherwise must be in $[-1, 1]$. Hence $X = \overline{A_2}$ as desired.

The fact that X is connected then follows from the following proposition (Proposition 2.3.11).

Proposition 2.3.11. *Let $(X, \mathcal{T})$ be a topological space and let $A \subseteq X$ be a connected subspace of $(X, \mathcal{T}_X)$. If $A \subseteq B \subseteq \overline{A}$ (where $\overline{A}$ is computed in $(X, \mathcal{T})$), then B is connected subspace of $(X, \mathcal{T}_X)$.*

Proof. Suppose that B is not connected. Thus there exists open subsets $U, V \in \mathcal{T}$ such that $U \cap B \neq \emptyset$, $V \cap B \neq \emptyset$, $U \cap V \cap B = \emptyset$, and $B \subseteq U \cup V$. Since $A \subseteq B$, these imply that $U \cap V \cap A = \emptyset$ and $A \subseteq U \cup V$. Therefore, as A is a connected subsets of $(X, \mathcal{T})$, it must be the case that $U \cap A = \emptyset$ or $V \cap A = \emptyset$. By exchanging the labelling of U and V if necessary, we may assume that $U \cap A = \emptyset$. Hence $A \subseteq V$.

Recall that Lemma 1.6.20 implies that the closure of A in B is the intersection of the closure of A in X with B. Hence

$$B = \overline{A} \cap B = \overline{A \cap B} \subseteq \overline{V \cap B}$$

where the first closure is computed in $(X, \mathcal{T})$ and the other two closures are computed in B. However, $V \cap B$ is closed in B as $(B \setminus (V \cap B)) = B \cap U$ is open in B. Hence $B \subseteq (V \cap B)$. However, as $U \cap V \cap B = \emptyset$, it must be the case that $U \cap B = \emptyset$ thereby contradicting the fact that $U \cap B \neq \emptyset$. Hence, as we have obtained a contradiction, B must be connected as desired. ∎

Going back to Example 2.3.7, it is possible for a topological space to not be connected but have connected portions. In order for us to find the largest possible connected portions of a topological space, we use the following which lets us compare points by the connected sets containing them.

Lemma 2.3.12. *Let $(X, \mathcal{T})$ be a topological space and define a relation $\sim$ on X by defining $x_1 \sim x_2$ for two points $x_1, x_2 \in X$ if and only if there exists a connected subspace A of X such that $x_1, x_2 \in A$. Then $\sim$ is an equivalence relation on X.*

Proof. To see that $\sim$ is an equivalence relation, we first notice for all $x \in X$ that $x \sim x$ since $\{x\}$ is a connected set by definition. Next, suppose that $x, y \in X$ are such that $x \sim y$. Thus there exists a connected subspace A of X such that $x, y \in A$. Hence $y \sim x$ by definition.

Finally, suppose that $x, y, z \in X$ are such that $x \sim y$ and $y \sim z$. Hence there exists connected subspaces A_1 and A_2 of X such that $x, y \in A_1$ and $y, z \in A_2$. Therefore, since $y \in A_1$, $y \in A_2$, and A_1 and A_2 are connected, Proposition 2.3.8 implies that $A_1 \cup A_2$ is a connected subspace of $(X, \mathcal{T})$. Therefore, as $x, z \in A_1 \cup A_2$, $x \sim z$. Hence, as we have demonstrated that $\sim$ is reflexive, symmetric, and transitive, $\sim$ is an equivalence relation. ∎

Using the above equivalence relation, we can define what turns out to be the largest connected subspaces of a topological space.

Definition 2.3.13. Let $(X, \mathcal{T})$ be a topological space and let $\sim$ be the equivalence relation from Lemma 2.3.12. The equivalence classes of $\sim$ are called the *connected components* of $(X, \mathcal{T})$.

As advertised, the following result shows us that the connected components of a topological space are connected and thus the largest possible connected subspaces as there is no connected subspace containing points from different connected components.

Proposition 2.3.14. *Let $(X, \mathcal{T})$ be a topological space. Then the connected components of $(X, \mathcal{T})$ are non-empty, pairwise disjoint, connected subspaces of X whose union is X. Furthermore, if A is a connected subspace of X, then A lies entirely inside one connected component.*

Proof. First, clearly the connected components are non-empty, pairwise disjoint, and have union X by the properties of equivalence classes. To see that every connected component of $(X, \mathcal{T})$ is a connected subspace of X, let Y be an arbitrary connected component of X. Thus there exists an $x \in Y$. For each $y \in Y$ we know that $x \sim y$ so there exists a connected subspace X_y of X such that $x, y \in X_y$. Since $z \sim x$ for all $z \in X_y$, we have that $X_y \subseteq Y$ for all $y \in Y$. Hence

$$Y = \bigcup_{y \in Y} X_y.$$

However, as $x \in X_y$ for all $y \in Y$, Proposition 2.3.8 implies that Y is connected. Therefore, as Y was arbitrary, every connected component of Y is connected.

Finally, suppose that A is as connected subspace of X. If there were two distinct connected components X_1 and X_2 of X containing points in A, then there would exist $x_1 \in X_1 \cap A$ and $x_2 \in X_2 \cap A$. Hence $x_1, x_2 \in A$ so, as A is connected, $x_1 \sim x_2$. As $x_1 \in X_1$ and $x_2 \in X_2$, and as X_1 and X_2 were two distinct connected components of X so no element of X_1 is equivalent to an element of X_2, this is a contradiction. Hence A must lie entirely inside one connected component of X as X is the union of its connected components by the definition of an equivalence relation. ■

One incredible use of Proposition 2.3.14 is that we can decompose any topological space into disjoint connected topological spaces. Therefore, if we want to completely understand a general topological space, we need only understand connected topological spaces and how the property we wish to study behaves under taking unions.

For some examples of decomposing topological spaces into their connected components, we turn to the following.

Example 2.3.15. Let $X = [0,1] \cup [2,3]$ equipped with the subspace topology inherited from $\mathbb{R}$. Then there are two connected components of X, namely $[0,1]$ and $[2,3]$. Indeed clearly $[0,1]$ and $[2,3]$ are connected subsets of X and thus contained in connected components. Furthermore, if $x \in [0,1]$, $y \in [2,3]$, and $x, y \in A$, then A must be disconnected via the open sets $U = \left(-1, \frac{3}{2}\right)$ and $V = \left(\frac{3}{2}, 4\right)$. Hence $[0,1]$ and $[2,3]$ must be the connected components of X.

Example 2.3.16. Let $X = \mathbb{Q}$ equipped with the subspace topology inherited from $\mathbb{R}$. Then the connected components of $\mathbb{Q}$ are

$$\{\{q\} \mid q \in \mathbb{Q}\}.$$

Indeed clearly $\{q\}$ is a connected set for all $q \in \mathbb{Q}$. However, $A \subseteq \mathbb{Q}$ is such that there exist $q_1, q_2 \in A$ such that $q_1 < q_2$, then there exists an irrational number r such that $q_1 < r < q_2$ so the open sets $(-\infty, r)$ and (r, ∞) show that A is not connected. Hence no subset of $\mathbb{Q}$ with two points is connected, so it must be the case that the connected components of $\mathbb{Q}$ are $\{\{q\} \mid q \in \mathbb{Q}\}$.

Using the idea of connected components, we can show the following.

Theorem 2.3.17. *Let $\{(X_\alpha, \mathcal{T}_\alpha)\}_{\alpha \in I}$ be a collection of connected topological spaces. Then $\prod_{\alpha \in I} X_\alpha$ is a connected topological space when equipped with the product topology.*

Proof. First, we claim that if $x = (x_\alpha)_{\alpha \in I} \in \prod_{\alpha \in I} X_\alpha$ and $y = (y_\alpha)_{\alpha \in I} \in \prod_{\alpha \in I} X_\alpha$ are two elements that differ in exactly one entry, then x and y must lie in the same connected component of $\prod_{\alpha \in I} X_\alpha$. Indeed, suppose there exists an $\alpha_0 \in I$ such that $x_\alpha = y_\alpha$ for all $\alpha \in I \setminus \{\alpha_0\}$. Consider the function $f : X_{\alpha_0} \to \prod_{\alpha \in I} X_\alpha$ defined by $f(z) = (z_\alpha)_{\alpha \in I}$ where

$$z_\alpha = \begin{cases} x_\alpha & \text{if } \alpha \neq \alpha_0 \\ z & \text{if } \alpha = \alpha_0 \end{cases}.$$

It is elementary based on the definition of continuity and the product topology on $\prod_{\alpha\in I} X_\alpha$ to see that f is continuous. Therefore, since $(X_{\alpha_0}, \mathcal{T}_{\alpha_0})$ is connected, $f(X_{\alpha_0})$ is a connected subset of $\prod_{\alpha\in I} X_\alpha$ by the Intermediate Value Theorem (Theorem 2.3.5). Therefore, as $x, y \in f(X_{\alpha_0})$, x and y are in the same connected component of $\prod_{\alpha\in I} X_\alpha$.

Consequently, if $x = (x_\alpha)_{\alpha\in I} \in \prod_{\alpha\in I} X_\alpha$ and $y = (y_\alpha)_{\alpha\in I} \in \prod_{\alpha\in I} X_\alpha$ are two elements that differ in a finite number of entries, then x and y must lie in the same connected component of $\prod_{\alpha\in I} X_\alpha$. Indeed if x and y differ in a finite number of entries, here is a sequence of elements $x_0 = x$, x_1, x_2, …, $x_n = y$ such that x_k and x_{k+1} differ in exactly one entry for each $k \in \{0, \ldots, n-1\}$. Hence x_k and x_{k+1} are in the same connected component for each $k \in \{0, \ldots, n-1\}$ so Lemma 2.3.12 implies x and y must be in the same connected component of $\prod_{\alpha\in I} X_\alpha$.

Now suppose to the contrary that $\prod_{\alpha\in I} X_\alpha$ is disconnected. Then there exist non-empty open sets U and V in the product topology on $\prod_{\alpha\in I} X_\alpha$ such that $U \cap V = \emptyset$ yet $U \cup V = \prod_{\alpha\in I} X_\alpha$. Let $x \in U$ and $y \in V$ be arbitrary. If x and y were in the same connected component A of $\prod_{\alpha\in I} X_\alpha$, then $U \cap A$ and $V \cap A$ would imply that A is disconnected, which is a contradiction. Hence x and y must be in different connected components of $\prod_{\alpha\in I} X_\alpha$.

Since U is a non-empty open set in the product topology on $\prod_{\alpha\in I} X_\alpha$, there exists an open set U' of the form $\prod_{\alpha\in I} U_\alpha$ where $U_\alpha = X_\alpha$ for all but finitely many α and $U_\alpha \in \mathcal{T}$ for all $\alpha \in I$ such that $U' \subseteq U$. However, as U' only has restrictions on the values of its elements at a finite number of $\alpha \in I$, for each $y \in V$ there exists a $x \in U'$ such that x and y differ at only a finite number of entries (namely at those $\alpha \in I$ such that $U_\alpha \neq X_\alpha$). Hence x and y must then be in the same connected component, which contradicts the above paragraph since $y \in V$ and $x \in U' \subseteq U$. Thus we have a contradiction to the existence of U and V, so $\prod_{\alpha\in I} X_\alpha$ is connected. ∎

Of course, the box topology is not so nice.

Example 2.3.18. Consider $X = \prod_{n\in\mathbb{N}} \mathbb{R}$ equipped with the box topology. Clearly $\mathbb{R}$ is a connected topological space by Theorem 2.3.2. Let

$$U = \left\{ (x_n)_{n\geq 1} \in X \,\middle|\, \sup_{n\in\mathbb{N}} |x_n| < \infty \right\}$$

so that if $V = X \setminus U$ then V is the set of all unbounded elements of X. Clearly $U \cap V = \emptyset$ and $U \cup V = X$.

We claim that U is open in X. To see this, let $(x_n)_{n\geq 1} \in U$ be arbitrary. Then

$$\prod_{n\in\mathbb{N}} (x_n - 1, x_n + 1) \subseteq U$$

as every element of this product of intervals is bounded since $(x_n)_{n\geq 1}$ is bounded Therefore, since $\prod_{n\in\mathbb{N}}(x_n - 1, x_n + 1)$ is open in the box topology and since $(x_n)_{n\geq 1} \in U$ was arbitrary, U is open in X.

Similarly, we claim that V is open in X. To see this, let $(x_n)_{n\geq 1} \in V$ be arbitrary. Then

$$\prod_{n\in\mathbb{N}} (x_n - 1, x_n + 1) \subseteq V$$

as every element of this product of intervals is unbounded since $(x_n)_{n\geq 1}$ is unbounded. Therefore, since $\prod_{n\in\mathbb{N}}(x_n-1, x_n+1)$ is open in the box topology and since $(x_n)_{n\geq 1} \in V$ was arbitrary, V is open in X. Hence U and V imply that X is not connected by definition.

Not only are connected components useful in showing that the product topology on the product of connected spaces is connected, but connected components can be a useful tool in showing that topological spaces are not homeomorphic via the following result.

Proposition 2.3.19. *Let $(X, \mathcal{T}_X)$ and $(Y, \mathcal{T}_Y)$ be topological spaces and let $f : X \to Y$ be a homeomorphism. For any subset A of X, the cardinality of the set of connected components of $X \setminus A$ equipped with the subspace topology from $(X, \mathcal{T}_X)$ and the cardinality of the set of connected components of $Y \setminus f(A)$ equipped with the subspace topology from $(Y, \mathcal{T}_Y)$ must be equal.*

Proof. As $f : X \to Y$ is a homeomorphism, it is clear that $f|_{X\setminus A} : X \setminus A \to Y \setminus f(A)$ is a homeomorphism by Lemma 2.1.18. By the Intermediate Value Theorem (Theorem 2.3.5), $f|_{X\setminus A}$ must be a bijection between the connected components of $X \setminus A$ and $Y \setminus f(A)$ thereby yielding the result. ∎

Here is a couple of examples of how to use Proposition 2.3.19.

Example 2.3.20. Let $X = [0, 1]$ and $Y = (0, 1)$ be equipped with the subspace topologies inherited from $\mathbb{R}$. Then X and Y are not isomorphic. To see this, suppose to the contrary that there exists a homeomorphism $f : X \to Y$. It is clear that $X \setminus \{0\}$ is connected. Therefore, by Proposition 2.3.19, $Y \setminus f(0)$ must be connected. However, one can easily verify that $Y \setminus \{z\}$ has two connected components, namely $(0, z)$ and $(z, 1)$ thereby yielding a contradiction.

Example 2.3.21. Let $X = [0, 1]$ equipped with its subspace topology inherited from $\mathbb{R}$ and let

$$S^1 = \{(x, y) \mid x^2 + y^2 = 1\}$$

equipped with its subspace topology inherited from $\mathbb{R}^2$. Then X and S^1 are not isomorphic. To see this, suppose to the contrary that there exists a homeomorphism $f : X \to S^1$. It is clear that $X \setminus \left\{\frac{1}{2}\right\}$ has two connected components, namely $\left[0, \frac{1}{2}\right)$ and $\left(\frac{1}{2}, 1\right]$. Therefore, by Proposition 2.3.19, $S^1 \setminus f\left(\frac{1}{2}\right)$ must have two connected components. However, one can easily verify that $S^1 \setminus \{z\}$ is connected for any $z \in S^1$ thereby yielding a contradiction.

In the above example, we can verify that $S^1 \setminus \{z\}$ is connected for any $z \in S^1$ by explicitly showing $S^1 \setminus \{z\}$ is homeomorphic to $(0, 1)$ by describing the homeomorphism. This exact type of function leads us towards other forms of connectedness.

2.4 Other Forms of Connectedness

As we have seen in the previous section, the notion of a connected topological spaces has a lot of interesting ideas, such as the Intermediate Value Theorem and the ability to decompose any topological space into its connected components. As the notion of being connected has several uses in the theory of topology, it is not surprising that there are several variations of this notion that we will study in this section.

For our first notion, we try to develop an notion of connectedness by being able to connect any two points in the space via some analogue of a line.

Definition 2.4.1. Let $(X, \mathcal{T})$ be a topological space and let $x_1, x_2 \in X$. A *path from x_1 to x_2 in X* is a continuous function $f : [a, b] \to X$ such that $f(a) = x_1$ and $f(b) = x_2$.

Remark 2.4.2. Since the composition of continuous function is continuous and since for any two closed intervals $[a, b]$ and $[c, d]$ of $\mathbb{R}$ there exists a homeomorphism $h : [a, b] \to [c, d]$ (i.e. h is a linear function), if there is a path from x_1 to x_2 in $(X, \mathcal{T})$ then for any closed interval $[a, b]$ in $\mathbb{R}$ there exists a continuous function $f : [a, b] \to X$ such that $f(a) = x_1$ and $f(b) = x_2$. Thus we will often take $[a, b] = [0, 1]$ but we can take any other interval we desired as needed.

Definition 2.4.3. A topological space $(X, \mathcal{T})$ is said to be *path connected* if for every $x_1, x_2 \in X$ there exists a path from x_1 to x_2 in X.

The first immediate question that must pop into a mathematician's head is how does the notion of a path connected topological space relate to the notion of a connected topological space. One direction is immediate.

Proposition 2.4.4. *If $(X, \mathcal{T})$ is a path connected topological space, then $(X, \mathcal{T})$ is a connected topological space.*

Proof. Let $(X, \mathcal{T})$ be path connected. Suppose that $(X, \mathcal{T})$ is not connected. Hence there exists non-empty sets $U, V \in \mathcal{T}$ such that $U \cap V = \emptyset$ and $U \cup V = X$. Choose $x_1 \in U$ and $x_2 \in V$. Since $(X, \mathcal{T})$ is path connected, there exists a continuous function $f : [0, 1] \to X$ such that $f(0) = x_1$ and $f(1) = x_2$. However, since f is continuous, the assumptions on U and V imply that $f^{-1}(U)$ and $f^{-1}(V)$ must be non-empty pairwise disjoint open subsets of $[a, b]$ whose union is $[a, b]$. Hence $[a, b]$ is not connected, which contradicts Theorem 2.3.2. Thus $(X, \mathcal{T})$ must be connected. ∎

In terms of a converse to Proposition 2.4.4, we return to a previous example.

Example 2.4.5. Let $X \subseteq \mathbb{R}^2$ be the topologist's sine curve from Example 2.3.10. Recall that Example 2.3.10 demonstrated that X is a connected space. We claim that X is not path connected. To see this, suppose $f : [0, 1] \to X$ is a continuous function such that $f(0) = \left(\frac{1}{\pi}, 0\right)$ and $f(1) = (0, 0)$. For each $t \in [0, 1]$, write $f(t) = (f_1(t), f_2(t))$. Since f is continuous and since the subspace topology on X is inherited from the topology on $\mathbb{R}^2 = \mathbb{R} \times \mathbb{R}$, which is a product topology, it must be the case that $f_1, f_2 : [0, 1] \to [-1, 1]$ must be continuous. As $f_1(1) = 0$ and $f_1(0) = \frac{1}{\pi}$, we must have that

$$t_0 = \inf\{t \in [0, 1] \mid f_1(t) = 0\} \in (0, 1]$$

and $f_1(t_0) = 0$ by the continuity of f_1. As $f_1(t) > 0$ for all $t \in [0, t_0)$, the description of X implies that $f_2(t) = \sin\left(\frac{1}{f_1(t)}\right)$ for all $t \in [0, t_0)$. However, as f_1 is continuous, as t converges to t_0 from the left, $f_1(t)$ must obtains all of the values of $\frac{1}{\frac{\pi}{2}+\pi n}$ for sufficiently large $n \in \mathbb{N}$ so $f_2(t)$ fluctuates between -1 and 1. Hence $\lim_{t\to t_0} f_2(t)$ does not exist thereby contradicting the fact that f_2 was continuous. Hence there is no path in $(X, \mathcal{T})$ that connects $\left(\frac{1}{\pi}, 0\right)$ and $(0, 0)$ so $(X, \mathcal{T})$ is not path connected.

Often it is significantly easier to verify a topological space is path connected and use Proposition 2.4.4 to verify the topological spaces is connected.

Example 2.4.6. Consider the subspace $A = \mathbb{R}^n \setminus \{\vec{0}\}$ of $\mathbb{R}^n$. Then A is path connected if and only if $n \neq 1$. Indeed for $n = 1$ we see that $A = (-\infty, 0) \cup (0, \infty)$ is not path connected by Proposition 2.4.4 as A is not connected by Theorem 2.3.2. Otherwise, if $n \neq 1$ and $x_1, x_2 \in A$, then either the straight line from x_1 to x_2 does not pass through $\vec{0}$ and thus exhibits a path from x_1 to x_2, or else there exists a third point $x_3 \in A$ such that the lines from x_1 to x_3 and from x_3 to x_2 do not pass through the origin thereby giving a path from x_1 to x_2 via x_3 (see Lemma 2.4.10). Hence A is path connected when $n \neq 1$.

Example 2.4.7. Every open ball in every normed linear space $(V, \|\cdot\|)$ is path connected. Indeed consider the ball $B_{\|\cdot\|}(\vec{x}, r)$ for some $r > 0$ and $\vec{x} \in V$. If $\vec{y}_1, \vec{y}_2 \in B_{\|\cdot\|}(\vec{x}, r)$, then the function $f : [0, 1] \to V$ defined by

$$f(t) = t\vec{y}_1 + (1 - t)\vec{y}_2$$

is such that

$$\begin{aligned}\|f(t) - \vec{x}\| &= \|t(\vec{y}_1 - \vec{x}) + (1 - t)(\vec{y}_2 - \vec{x})\| \\ &\leq t\,\|\vec{y}_1 - \vec{x}\| + (1 - t)\,\|\vec{y}_2 - \vec{x}\| \\ &< tr + (1 - t)r = r\end{aligned}$$

for all $t \in [0, 1]$. Thus f is a path from $\vec{y}_1$ to $\vec{y}_2$ inside of $B_{\|\cdot\|}(\vec{x}, r)$. Hence $B_{\|\cdot\|}(\vec{x}, r)$ is path connected.

Example 2.4.8. Consider the unit circle

$$S^1 = \{(x, y) \in \mathbb{R}^2 \mid x^2 + y^2 = 1\} \subseteq \mathbb{R}^2.$$

Then S^1 is path connected since the function $f : [0, 4\pi] \to S^1$ defined by

$$f(t) = (\cos(t), \sin(t))$$

is a path that passes through any two points of S^1 in any order one chooses. Similarly, the unit sphere

$$S^2 = \{(x, y, z) \in \mathbb{R}^3 \mid x^2 + y^2 + z^2 = 1\} \subseteq \mathbb{R}^3$$

is path connected as any two points in S^2 lie on a great circle and thus one may use a path like those used for S^1.

As alluded to in the above examples, the notion of points being path connected is very well behaved. In particular, it has some common connections with the notion of connected topological spaces, such as the Intermediate Value Theorem.

Proposition 2.4.9. *Let $(X, \mathcal{T}_X)$ and $(Y, \mathcal{T}_Y)$ be topological spaces and let $f : X \to Y$ be continuous. If $(X, \mathcal{T}_X)$ is path connected, $f(X)$ is path connected when equipped with the subspace topology inherited from Y.*

Proof. Let $y_1, y_2 \in f(X)$ be arbitrary. Hence there exists $x_1, x_2 \in X$ such that $f(x_1) = y_1$ and $f(x_2) = y_2$. Since X is path connected, there exists a continuous function $g : [0, 1] \to X$ such that $g(0) = x_1$ and $g(1) = x_2$. Therefore, if $h = f \circ g : [0, 1] \to f(X)$, then $h(0) = y_1$, $h(1) = y_2$, and h is continuous being the composition of continuous functions. Therefore, since $y_1, y_2 \in f(X)$ were arbitrary, $f(X)$ is path connected. ∎

Lemma 2.4.10. *Let $(X, \mathcal{T})$ be a topological space and define a relation $\sim$ on X by defining $x_1 \sim x_2$ for two points $x_1, x_2 \in X$ if and only if there exists a path from x_1 to x_2. Then $\sim$ is an equivalence relation on X.*

Proof. Notice if $x \in X$, then $x \sim x$ via the path $f : [0, 1] \to X$ defined by $f(t) = x$ for all $t \in [0, 1]$.

Next, suppose $x_1, x_2 \in X$ are such that $x_1 \sim x_2$. Thus there exists a path $f : [0, 1] \to X$ such that $f(0) = x_1$ and $f(1) = x_2$. Define $g : [0, 1] \to X$ by $g(t) = f(1 - t)$ for all $t \in [0, 1]$. As g is a composition of continuous functions, g is continuous. Therefore since $g(0) = x_2$ and $g(1) = x_1$, $x_2 \sim x_1$ as desired.

Finally, suppose $x_1, x_2, x_3 \in X$ are such that $x_1 \sim x_2$ and $x_2 \sim x_3$. Thus there exist paths $f : [0, 1] \to X$ and $g : [1, 2] \to X$ such that $f(0) = x_1$,

$f(1) = x_2 = g(1)$, and $g(2) = x_3$. By the Pasting Lemma (Theorem 2.1.20) the function $h : [0, 2] \to X$ defined by

$$h(t) = \begin{cases} f(t) & \text{if } t \in [0,1] \\ g(t) & \text{if } t \in [1,2] \end{cases}$$

is continuous. Hence as $h(0) = x_1$ and $h(2) = x_3$ we have that $x_1 \sim x_3$. Therefore, as $\sim$ is symmetric, reflective, and transitive, $\sim$ is an equivalence relation as desired. ∎

Definition 2.4.11. Let $(X, \mathcal{T})$ be a topological space and let $\sim$ be the equivalence relation from Lemma 2.4.10. The equivalence classes of $\sim$ are called the *path connected components* of $(X, \mathcal{T})$.

Proposition 2.4.12. *Let $(X, \mathcal{T})$ be a topological space. Then the path connected components of $(X, \mathcal{T})$ are non-empty, pairwise disjoint, path connected subspaces of X. Furthermore, if A is a path connected subspace of X, then A lies entirely inside one path connected component.*

Proof. First, clearly the path connected components are non-empty, pairwise disjoint, and have union X by the properties of equivalence classes. To see that every path connected component of $(X, \mathcal{T})$ is a path connected subspace of X, let Y be an arbitrary path connected component of X and let $x_1, x_2 \in Y$ be arbitrary. Thus $x_1 \sim x_2$ so there exists a continuous function $f : [0, 1] \to X$ such that $f(0) = x_1$ and $f(1) = x_2$. However, for all $t \in (0, 1)$ we see that $f|_{[0,t]} : [0, t] \to X$ is a path from x_1 to $f(t)$ so $x_1 \sim f(t)$. Hence $f(t) \in Y$ for all $t \in (0, 1)$. Hence f is a path from x_1 to x_2 in Y. Therefore, as $x_1, x_2 \in Y$ were arbitrary, Y is path connected.

Finally, suppose that A is as path connected subspace of X. If there were two distinct path connected components X_1 and X_2 of X containing points in A, then there would exist $x_1 \in X_1 \cap A$ and $x_2 \in X_2 \cap A$. Hence $x_1, x_2 \in A$ so, as A is path connected, $x_1 \sim x_2$. As $x_1 \in X_1$ and $x_2 \in X_2$, and as X_1 and X_2 were two distinct connected components of X so no element of X_1 is equivalent to an element of X_2, this is a contradiction. Hence A must lie entirely inside one path connected component of X as X is the union of its path connected components by the definition of an equivalence relation. ∎

Example 2.4.13. The path connected components of $\mathbb{Q}$ (equipped with the subspace topology) are

$$\{\{q\} \mid q \in \mathbb{Q}\}.$$

Indeed clearly each set $\{q\}$ for $q \in \mathbb{Q}$ is path connected. Furthermore, if $A \subseteq \mathbb{Q}$ has two points, then A is not connected as it has elements from two different connected components of $\mathbb{Q}$ by Example 2.3.16 and thus cannot be path connected by Proposition 2.4.4.

Example 2.4.14. Recall the topologist's sine curve is the set

$$X = \{(0, y) \mid -1 \leq y \leq 1\} \cup \left\{ \left(x, \sin\left(\frac{1}{x}\right)\right) \,\middle|\, 0 < x \leq \frac{1}{\pi} \right\}$$

equipped with the subspace topology inherited from the canonical topology on $\mathbb{R}^2$. Clearly X is the union of the disjoint sets

$$A_1 = \{(0, y) \mid -1 \leq y \leq 1\} \quad \text{and} \quad A_2 = \left\{ \left(x, \sin\left(\frac{1}{x}\right)\right) \,\middle|\, 0 < x \leq \frac{1}{\pi} \right\}.$$

Furthermore A_1 and A_2 are path connected being continuous images of intervals of $\mathbb{R}$. We claim that A_1 and A_2 are the path connected components of X. Indeed A_1 and A_2 are path connected and the proof in Example 2.4.5 shows that no element of A_1 is path connected to an element in A_2 in X. Hence A_1 and A_2 are the path connected components of the topologist's sine curve.

Analyzing the notion of path connectedness for product spaces is far easier than analyzing the notion of connectedness was.

Proposition 2.4.15. *Let $\{(X_\alpha, \mathcal{T}_\alpha)\}_{\alpha \in I}$ be a collection of path connected topological spaces. Then $\prod_{\alpha \in I} X_\alpha$ is a connected topological space when equipped with the product topology.*

Proof. Let $x = (x_\alpha)_{\alpha \in I} \in \prod_{\alpha \in I} X_\alpha$ and $y = (y_\alpha)_{\alpha \in I} \in \prod_{\alpha \in I} X_\alpha$ be arbitrary. Since $(X_\alpha, \mathcal{T}_\alpha)$ is path connected for each $\alpha \in I$, there exists a continuous function $f_\alpha : [0, 1] \to X_\alpha$ such that $f_\alpha(0) = x_\alpha$ and $f_\alpha(1) = y_\alpha$. Define $f : [0, 1] \to \prod_{\alpha \in I} X_\alpha$ by

$$f(t) = (f_\alpha(t))_{\alpha \in I}$$

for all $t \in [0, 1]$. Then $f(0) = x$, $f(1) = y$, and f is continuous by Theorem 2.1.10. Hence, as x and y were arbitrary, $\prod_{\alpha \in I} X_\alpha$ is path connected in the product topology. ∎

Remark 2.4.16. Unsurprisingly, Proposition 2.4.15 fails when the product topology is replaced with the box topology. Indeed consider $X = \prod_{n \in \mathbb{N}} \mathbb{R}$ equipped with the box topology. Clearly $\mathbb{R}$ is path connected as the identity function restricted to any closed interval yields a path in $\mathbb{R}$ between the endpoints of the interval. However X cannot be path connected. Indeed X is not connected by Example 2.3.18 and thus cannot be path connected by Proposition 2.4.4.

To conclude our discussion of the notions of connectedness in topological spaces, it often is enough for a topological space to only have connected neighbourhoods bases at each point. Thus the following concept is defined.

Definition 2.4.17. Let $(X, \mathcal{T})$ be a topological space and let $x \in X$. It is said that $(X, \mathcal{T})$ is *locally connected at x* if there exists a neighbourhood basis of x consisting of sets that are connected topological spaces when equipped with the subspace topology.

Definition 2.4.18. A topological space $(X, \mathcal{T})$ is said to be *locally connected* if $(X, \mathcal{T})$ is locally connected at each point.

Perhaps surprisingly, the notions of connected and locally connected topological spaces have no relations to one another.

Example 2.4.19. Consider $X = [0, 1] \subseteq \mathbb{R}$ equipped with the subspace topology. Then X is connected by Theorem 2.3.2. In addition, X is locally connected. Indeed if $x \in X$ is arbitrary, then the sets

$$\{X \cap (x - \epsilon, x + \epsilon) \mid \epsilon > 0\}$$

are a neighbourhood basis of x consisting of intervals, which then are connected by Theorem 2.3.2. Therefore, as $x \in X$ was arbitrary, X is locally connected.

Example 2.4.20. Consider $X = [0, 1] \cup [2, 3] \subseteq \mathbb{R}$ equipped with the subspace topology. Then X is not connected by Theorem 2.3.2. However, X is locally connected. Indeed if $x \in X$ is arbitrary, then the sets

$$\{X \cap (x - \epsilon, x + \epsilon) \mid \epsilon \in (0, 1)\}$$

are a neighbourhood basis of x consisting of intervals, which then are connected by Theorem 2.3.2. Therefore, as $x \in X$ was arbitrary, X is locally connected.

Example 2.4.21. Consider $\mathbb{Q} \subseteq \mathbb{R}$ equipped with the subspace topology. Then $\mathbb{Q}$ is not connected by Example 2.3.16. Furthermore $\mathbb{Q}$ is not locally connected since for every $\epsilon > 0$, the neighbourhood $(-\epsilon, \epsilon) \cap \mathbb{Q}$ of 0 is not connected (by Theorem 2.3.2)) so no neighbourhood of 0 will be connected. Hence $\mathbb{Q}$ is not locally connected.

Example 2.4.22. Let X be the topologist's sine curve in $\mathbb{R}^2$. Then X is connected by Example 2.3.10. However, X is not locally connected. Indeed note that $(0, 1) \in X$. However, for any $\epsilon \in (0, 1)$ we see that $B_2((0, 1), \epsilon) \cap X$ is not connected as there will exists a $0 < \delta < \epsilon$ such that

$$\begin{aligned} A_1 &= B_2((0, 1), \epsilon) \cap X \cap \{(x, y) \in \mathbb{R}^2 \mid x < \delta\} \quad \text{and} \\ A_2 &= B_2((0, 1), \epsilon) \cap X \cap \{(x, y) \in \mathbb{R}^2 \mid x > \delta\} \end{aligned}$$

are non-empty pairwise disjoint open subsets of X. Hence $B_2((0, 1), \epsilon) \cap X$ is not connected for all $\epsilon \in (0, 1)$ so no neighbourhood of $(0, 1)$ in X can be connected. Hence X is not locally connected.

The reason locally connected topological spaces are so nice is the following result illustrating the relation of connectedness to the connected components of open sets.

Theorem 2.4.23. *A topological space $(X, \mathcal{T})$ is locally connected if and only if for every $U \in \mathcal{T}$, the connected components of U when equipped with the subspace topology are open in $(X, \mathcal{T})$.*

Proof. First assume that $(X, \mathcal{T})$ is locally connected. To see the claim, let $U \in \mathcal{T}$ be arbitrary and let C be an arbitrary connected component of U. Notice for all $x \in C$ that $x \in U$ so the definition of a locally connected space implies there exists an open connected subset $U_x \in \mathcal{T}$ such that $x \in U_x \subseteq U$. However, since $x \in U_x$, U_x is connected, $U_x \subseteq U$, and C is the connected component of U containing x, we must have that $U_x \subseteq C$. Hence

$$C = \bigcup_{x \in C} U_x \in \mathcal{T}$$

as desired. Therefore, since U and C were arbitrary, the claim follows.

To see the converse, suppose for every $U \in \mathcal{T}$, the connected components of U when equipped with the subspace topology are open in $(X, \mathcal{T})$. To see that $(X, \mathcal{T})$ is locally connected, let $x \in X$ and U a neighbourhood of x be arbitrary. If C is the connected component of U containing x, then C is open in $(X, \mathcal{T})$ by assumption, C is connected, and $x \in C \subseteq U$. Therefore, since U was arbitrary, x has a neighbourhood basis consisting of connected sets. Hence, as $x \in X$ was arbitrary, X is locally connected. ∎

Chapter 3

Compact Topological Spaces

As we have seen in Chapter 2, the notion of connectedness was intimately tied to properties of continuous functions via the Intermediate Value Theorem. In this section, we will develop the property of topological spaces necessary for an extension of the Extreme Value Theorem; namely the notion of compactness.

Compact topological spaces are some of the nicest topological spaces in existence. This stems from the fact that compact spaces enable one to use only a finite number of open set when describing the entire space. It is this finiteness that is desirable for many reasons, such as being able to find a element of a net further along than representatives of the net chosen from a (now finite) covering a topological space.

In this section, we will delve into the notion of compactness for general topological spaces. Excluding discussions of compact subsets of $\mathbb{K}^n$, the discussion of compactness in metric spaces is worth of its own chapter and will be postponed until Chapter 4. After developing the basic notion of compact topological spaces, we will analyze various methods and consequences of a topological space being compact. In addition, we will prove the well-known theorem of Tychonoff, which is equivalent to the Axiom of Choice (Axiom A.2.4), and further emphasizes why the product topology is superior to the box topology. Finally, we will analyze topological spaces that are locally compact and demonstrate these spaces are compact up to adding a single point thereby emphasizing the strength of compactness.

3.1 Compact Topological Spaces

To begin our study of compact topological spaces, we must define the notion of compactness. As our goal is to be able to cover our topological spaces with nice finite collections of sets, we define the following.

Definition 3.1.1. Let $(X, \mathcal{T})$ be a topological space and let $A \subseteq X$. Sets $\{U_\alpha\}_{\alpha \in I} \subseteq \mathcal{P}(X)$ are said to be an *open cover* of A if each $U_\alpha \in \mathcal{T}$ for all

$\alpha \in I$ and $A \subseteq \bigcup_{\alpha\in I} U_\alpha$.

A *subcover of A from* $\{U_\alpha\}_{\alpha\in I}$ is any collection $\{U_\alpha\}_{\alpha\in J}$ where $J \subseteq I$ such that $A \subseteq \bigcup_{\alpha\in J} U_\alpha$

The notion of compactness then follows by asking that we can extract a finite open cover of our space from any open cover we may wish to consider.

Definition 3.1.2. A topological space $(X, \mathcal{T})$ is said to be *compact* if every open cover of $(X, \mathcal{T})$ contains a finite subcover; that is, if $\{U_\alpha\}_{\alpha\in I} \subseteq \mathcal{T}$ are such that $X = \bigcup_{\alpha\in I} U_\alpha$, then there exists $J \subseteq I$ such that J is finite and $X = \bigcup_{\alpha\in J} U_\alpha$.

Of course, we have some trivial examples

Example 3.1.3. Technically the empty set is compact as every open cover has a subcover consisting of one element.

Example 3.1.4. The trivial topology on a set X is always compact as the only open covers of X will be $\{X\}$ and $\{X, \emptyset\}$.

Example 3.1.5. Let $(X, \mathcal{T})$ be a topological space with X finite. Then $(X, \mathcal{T})$ is compact as $\mathcal{T} \subseteq \mathcal{P}(X)$ is finite.

Example 3.1.6. Let X be an infinite set and let $\mathcal{T}$ be the discrete topology on X. Then $(X, \mathcal{T})$ is not compact as $\{\{x\}\}_{x\in X}$ is an open cover with no finite subcovers.

To obtain more examples of compact topological spaces, we turn our attention to subsets of $\mathbb{R}$. Of course we have the following.

Example 3.1.7. If $\mathbb{R}$ is equipped with its canonical topology, then $\mathbb{R}$ is not compact. Indeed $\mathcal{U} = \{(n-1, n+1) \mid n \in \mathbb{Z}\}$ is an open cover of $\mathbb{R}$ with no finite subcovers as each element of $\mathbb{Z}$ is covered by a unique element of $\mathcal{U}$.

In order to determine which subsets of $\mathbb{R}$ are compact when equipped with the subspace topology, we note the following.

Lemma 3.1.8. *Let $(X, \mathcal{T})$ be a topological space and let Y be a subspace of $(X, \mathcal{T})$. Then Y is compact if and only if every open cover of Y in $(X, \mathcal{T})$ has a finite subcover.*

Proof. For simplicity, let $\mathcal{T}_Y$ denote the subspace topology on Y inherited from $(X, \mathcal{T})$.

To begin, suppose $(Y, \mathcal{T}_Y)$ is compact. To see the result, let $\{U_\alpha\}_{\alpha\in I}$ be an arbitrary open cover of Y in $(X, \mathcal{T})$. Hence $\{Y \cap U_\alpha\}_{\alpha\in I}$ is an open cover of Y in $(Y, \mathcal{T}_Y)$ so the fact that $(Y, \mathcal{T}_Y)$ is compact implies there exists a finite subset $J \subseteq I$ such that $\{Y \cap U_\alpha\}_{\alpha\in J}$ is an open cover of Y in $(Y, \mathcal{T}_Y)$. Hence clearly $\{U_\alpha\}_{\alpha\in J}$ is a finite open subcover of Y from $\{U_\alpha\}_{\alpha\in I}$. Therefore, as $\{U_\alpha\}_{\alpha\in I}$ was arbitrary, the claim follows.

Conversely, suppose that every open cover of Y in $(X, \mathcal{T})$ has a finite subcover. To see that $(Y, \mathcal{T}_Y)$ is compact, let $\{V_\alpha\}_{\alpha \in I}$ be an arbitrary open cover of Y in $(Y, \mathcal{T}_Y)$. By the definition of the subspace topology there exists $\{U_\alpha\}_{\alpha \in I} \subseteq \mathcal{T}$ such that $V_\alpha = Y \cap U_\alpha$ for all $\alpha \in I$. Hence $\{U_\alpha\}_{\alpha \in I}$ is an open cover of Y in $(X, \mathcal{T})$, which then must have a finite subcover $\{U_\alpha\}_{\alpha \in J}$ of Y by assumption. Hence $\{V_\alpha\}_{\alpha \in J}$ is a finite open subcover of Y from $\{V_\alpha\}_{\alpha \in I}$. Therefore, since $\{V_\alpha\}_{\alpha \in I}$ was arbitrary, $(Y, \mathcal{T}_Y)$ is compact. ∎

Example 3.1.9. The subset $X = \{0\} \cup \left\{ \frac{1}{n} \mid n \in \mathbb{N} \right\} \subseteq \mathbb{R}$ is a compact subspace of $\mathbb{R}$. To see this, suppose that $\{U_\alpha\}_{\alpha \in I}$ is an open cover of X using open subsets from $\mathbb{R}$. Hence there exists an $\alpha_0 \in I$ such that $0 \in U_{\alpha_0}$. Since U_{α_0} is open, there exists an $\epsilon > 0$ such that $(-\epsilon, \epsilon) \subseteq U_{\alpha_0}$. Since X contains only a finite number of elements outside of $(-\epsilon, \epsilon)$, X contains only a finite number of element outside of U_{α_0}. Thus we can write $X \setminus U_{\alpha_0} = \{x_1, \ldots, x_m\}$ for some $m \in \mathbb{N}$. Since $\{U_\alpha\}_{\alpha \in I}$ is an open cover of X, for each $k \in \{1, \ldots, m\}$ there exists an $\alpha_k \in I$ such that $x_k \in U_{\alpha_k}$. Hence $\{U_{\alpha_k}\}_{k=0}^m$ is a finite subcover of X from $\{U_\alpha\}_{\alpha \in I}$. Hence, as $\{U_\alpha\}_{\alpha \in I}$ was arbitrary, X is compact.

Example 3.1.10. The subset $X = \left\{ \frac{1}{n} \mid n \in \mathbb{N} \right\} \subseteq \mathbb{R}$ of $\mathbb{R}$ is not compact. Indeed if $U_n = \left(\frac{1}{n}, 1\right)$ for all $n \in \mathbb{N}$, then $\{U_n\}_{n=1}^\infty$ is an open cover of X. However, clearly $\{U_n\}_{n=1}^\infty$ does not have a finite subcover as if $n_1, \ldots, n_m \in \mathbb{N}$ then $\bigcup_{k=1}^m U_{n_k} = \left(\frac{1}{\max\{n_1, \ldots, n_m\}}, 1\right)$ which does not contain all of X since $\lim_{n \to \infty} \frac{1}{n} = 0$.

Remark 3.1.11. It is clear in the above example that the reason why X is not compact was that X was not closed. However, for a general topological space $(X, \mathcal{T})$, a subspace A of $(X, \mathcal{T})$ may still be compact even if A is not closed in $(X, \mathcal{T})$. Indeed, if X is finite, then every subspace of $(X, \mathcal{T})$ is compact as every subspace topology consists only of a finite number of sets. As there are clearly examples of topologies on finite sets such that not every set is closed (see Example 1.1.4), we have demonstrated our claim.

The reason why the set in Example 3.1.10 is not compact because it was not closed follows from the fact that $\mathbb{R}$ is Hausdorff. In particular, mixing the notions of Hausdorff and compactness yields some powerful results. Recall that a topological space $(X, \mathcal{T})$ is Hausdorff means that $(X, \mathcal{T})$ has a lot of open sets to separate points (i.e. it is "close" to the discrete topology) whereas $(X, \mathcal{T})$ is compact means that $(X, \mathcal{T})$ does not have too many open sets as every open cover has a finite subcover (i.e. it is "close" to the trivial topology). It is this Goldilocks zone that makes compact Hausdorff topological spaces some of the nicest topological spaces to study.

In order to study compact Hausdorff topological spaces, we note the following incredibly useful lemma that lets us separate points from compact subsets.

Lemma 3.1.12. *Let $(X, \mathcal{T})$ be a Hausdorff space, let Y be a compact subspace of X, and let $x_0 \in X \setminus Y$. Then there exists $U, V \in \mathcal{T}$ such that $x_0 \in U$, $Y \subseteq V$, and $U \cap V = \emptyset$.*

Proof. Since $(X, \mathcal{T})$ is Hausdorff and $x_0 \in X \setminus Y$, for each $y \in Y$ there exists $U_y, V_y \in \mathcal{T}$ such that $x_0 \in U_y$, $y \in V_y$, and $U_y \cap V_y = \emptyset$. Hence $\{V_y\}_{y \in Y}$ is an open cover of Y. Therefore, as Y is a compact subspace of X, Lemma 3.1.8 implies there exists an $n \in \mathbb{N}$ and $y_1, y_2, \ldots, y_n \in Y$ such that $\{V_{y_k}\}_{k=1}^n$ is an open cover of Y. Let

$$U = \bigcap_{k=1}^{n} U_{y_k} \qquad \text{and} \qquad V = \bigcup_{k=1}^{n} V_k.$$

Clearly $Y \subseteq V$ by construction. Furthermore, as $x_0 \in U_{y_k}$ for all $k \in \{1, \ldots, n\}$, $x_0 \in U$. Finally, since $U_y \cap V_y = \emptyset$ for all $y \in Y$, we obtain that $U \cap V = \emptyset$ as desired. ∎

Using Lemma 3.1.12, we can formalize the problem with Example 3.1.10.

Theorem 3.1.13. *Every compact subspace of a Hausdorff topological space is closed.*

Proof. Let $(X, \mathcal{T})$ be a Hausdorff topological space and let Y be a compact subspace of X. To see that Y is closed in $(X, \mathcal{T})$, it will be demonstrated that $X \setminus Y$ is open. To see that $X \setminus Y$ is open, let $x_0 \in X \setminus Y$ be arbitrary. By Lemma 3.1.12, there exists open sets $U, V \in \mathcal{T}$ such that $x_0 \in U$, $Y \subseteq V$, and $U \cap V = \emptyset$. Hence U is a neighbourhood of x_0 that is contained in $X \setminus Y$. Therefore, as $x_0 \in X \setminus Y$ was arbitrary, $X \setminus Y$ is open in $(X, \mathcal{T})$. Hence Y is closed in $(X, \mathcal{T})$ as desired. ∎

In fact, Theorem 3.1.13 has somewhat of a converse in compact topological spaces.

Theorem 3.1.14. *Every closed subspace of a compact topological space is compact.*

Proof. Let $(X, \mathcal{T})$ be a compact topological space and let F be a closed subspace of $(X, \mathcal{T})$. To see that F is compact, we will verify the conditions of Lemma 3.1.8. Thus let $\{U_\alpha\}_{\alpha \in I}$ be an arbitrary open cover of F from $(X, \mathcal{T})$. Hence, as F is closed in $(X, \mathcal{T})$, $\{X \setminus F\} \cup \{U_\alpha\}_{\alpha \in I}$ is an open cover of $(X, \mathcal{T})$. Therefore, as $(X, \mathcal{T})$ is compact, there exists a finite subset $J \subseteq I$ such that $\{X \setminus F\} \cup \{U_\alpha\}_{\alpha \in J}$ is an open cover of $(X, \mathcal{T})$. Therefore, since $X \setminus F$ is disjoint from F, $\{U_\alpha\}_{\alpha \in I}$ is a finite subcover F from $\{U_\alpha\}_{\alpha \in I}$. Hence Lemma 3.1.8 implies that F is a compact subspace of $(X, \mathcal{T})$. ∎

Combining Theorem 3.1.13 and Theorem 3.1.14, we can construct new compact subspaces from other compact subspaces.

Corollary 3.1.15. *The arbitrary non-empty intersection of compact subspaces of a Hausdorff topological space is compact.*

Proof. Let $(X, \mathcal{T})$ be a Hausdorff topological space and let $\{K_\alpha\}_{\alpha \in I}$ be compact subspaces of $(X, \mathcal{T})$ with I non-empty. By Theorem 3.1.13, K_α is closed in $(X, \mathcal{T})$ for all $\alpha \in I$. Hence $K = \bigcap_{\alpha \in I} K_\alpha$ is closed in $(X, \mathcal{T})$. As I is non-empty, K is a closed subset of K_α for all $\alpha \in I$ and thus a compact subspace K_α for all $\alpha \in I$ by Theorem 3.1.14. ∎

Remark 3.1.16. Note Corollary 3.1.15 does not extend to (even the finite intersection) of compact subspaces of non-Hausdorff topological spaces. For such an example, let $X = \mathbb{N}$ and let

$$\mathcal{T} = \{A \mid A \subseteq \mathbb{N} \setminus \{1, 2\}\} \cup \{\mathbb{N}\} \cup \{\mathbb{N} \setminus \{1\}\} \cup \{\mathbb{N} \setminus \{2\}\}.$$

It is not difficult to verify that $\mathcal{T}$ is a topology on X. Furthermore, if $K_1 = \mathbb{N} \setminus \{1\}$ and $K_2 = \mathbb{N} \setminus \{2\}$, it is not difficult to verify that K_1 and K_2 are compact subspaces of X as any open cover of K_1 must include either K_1 or $\mathbb{N}$ (both of which are finite subcovers of K_1) and any open cover of K_2 must include either K_2 or $\mathbb{N}$ (both of which are finite subcovers of K_2). However, $K_1 \cap K_2 = \mathbb{N} \setminus \{1, 2\}$ is clearly not compact as $\{\{n\} \mid n \in \mathbb{N} \setminus \{1, 2\}\}$ is an open cover of $K_1 \cap K_2$ with no finite subcovers.

Corollary 3.1.17. *The finite union of compact subspaces of a topological space is compact.*

Proof. Let $\{K_k\}_{k=1}^n$ be compact subspaces of a topological space $(X, \mathcal{T})$ and let $K = \bigcup_{k=1}^n K_k$. To see that K is compact in $(X, \mathcal{T})$, let $\mathcal{U} = \{U_\alpha\}_{\alpha \in I}$ be an arbitrary open cover of K. Hence $\mathcal{U}$ is an open cover of K_k for all $k \in \{1, \ldots, n\}$. Therefore, since K_k is compact for all $k \in \{1, \ldots, n\}$, there exists a finite subset $J_k \subseteq I$ such that $\{U_\alpha\}_{\alpha \in J_k}$ is an open cover off K_k. Thus if $J = \bigcup_{k=1}^n J_k$, then J is a finite subset of I and $\{U_\alpha\}_{\alpha \in J}$ is an open cover of K. Therefore, as $\mathcal{U}$ was arbitrary, K is compact as desired. ∎

Remark 3.1.18. Note Corollary 3.1.17 does not extend to arbitrary unions of compact subspaces. Indeed clearly $X = \left\{\frac{1}{n} \mid n \in \mathbb{N}\right\}$ is a union of compact subsets of $\mathbb{R}$ as every singleton in $\mathbb{R}$ is trivially compact. However X is not compact by Example 3.1.10.

However, there more obstructions for a subset of $\mathbb{R}$ to be compact.

Example 3.1.19. Let $X = \mathbb{Z} \subseteq \mathbb{R}$ equipped with the subspace topology inherited from the canonical topology on $\mathbb{R}$. Clearly X is a closed subset of $\mathbb{R}$ since any convergent net from X must eventually be constant as $\mathrm{dist}(n, X \setminus \{n\}) = 1$ for all $n \in \mathbb{N}$. However, we claim that X is not compact. Indeed if $U_n = (-n, n)$ for each $n \in \mathbb{N}$, then $\mathcal{U} = \{U_n\}_{n=1}^\infty$ is an open cover of X. However, $\mathcal{U}$ does not have a finite subcover since $U_n \subseteq U_{n+1}$ for all

$n \in \mathbb{N}$ so that $\mathcal{U}$ is closed under unions, and since each element of $\mathcal{U}$ contains only a finite number of points in the infinite set X.

It is not difficult to see that the set in Example 3.1.19 is not compact as its elements get arbitrary far away from 0. To give a name to this issue, we define the following property for metric spaces.

Definition 3.1.20. A subset A of metric space (X, d) is said to be *bounded* if there exists an $M \geq 0$ such that

$$\{d(a_1, a_2) \mid a_1, a_2 \in A\} \subseteq [0, M].$$

There are many ways to characterize boundedness in a metric space.

Lemma 3.1.21. *Let (X, d) be a metric space and let $A \subseteq X$ be non-empty. The following are equivalent:*

(i) A is bounded.

(ii) For each $a_0 \in A$, $A \subseteq B_d(a_0, R)$ for some $R > 0$.

(iii) For an $a_0 \in A$, $A \subseteq B_d(a_0, R)$ for some $R > 0$.

Proof. To see that (i) implies (ii), suppose A is bounded. Thus there exists an $M > 0$ such that

$$\{d(a_1, a_2) \mid a_1, a_2 \in A\} \subseteq [0, M].$$

Hence (ii) follows by taking $R = M$ for each $a_0 \in A$.

Clearly (ii) implies (iii). To see that (iii) implies (i), let $a_0 \in A$ and $R > 0$ be such that $A \subseteq B_d(a_0, R)$. Hence for all $a_1, a_2 \in A$,

$$d(a_1, a_2) \leq d(a_1, a_0) + d(a_0, a_2) \leq R + R = 2R.$$

Hence

$$\{d(a_1, a_2) \mid a_1, a_2 \in A\} \subseteq [0, 2R]$$

so A is bounded by definition. ∎

Using the same idea as Example 3.1.19, we have the following.

Theorem 3.1.22. *Every compact metric space is bounded.*

Proof. Let (X, d) be a compact metric space. To see that (X, d) is bounded, fix a point $x_0 \in X$. For each $n \in \mathbb{N}$, consider the open set $U_n = B_d(x_0, n)$. Since for all $x \in X$ there exists an $m \in \mathbb{N}$ such that $d(x, x_0) < m$, we see that $\bigcup_{n=1}^{\infty} U_n = X$. Hence $\{U_n\}_{n=1}^{\infty}$ is an open cover of (X, d). Therefore, since (X, d) is compact, there exists $n_1, \ldots, n_q \in \mathbb{N}$ such that $X = \bigcup_{k=1}^{q} U_{n_k}$. If $N = \max\{n_1, \ldots, n_q\}$, we clearly obtain that $X = B_d(x_0, N)$. Thus (X, d) is bounded by Lemma 3.1.21. ∎

Combining Theorem 3.1.13 and Theorem 3.1.22, every compact subspace of $\mathbb{R}$ must be closed and bounded. We desire to prove the converse to this statement. To simplify notation, we define the following.

Definition 3.1.23. Let (X, d) be a metric space and let $A \subseteq X$ be non-empty. The *diameter of* A, denoted $\mathrm{diam}(A)$, is defined to be

$$\mathrm{diam}(A) = \sup(\{d(a_1, a_2) \mid a_1, a_2 \in A\}) \subseteq [0, \infty].$$

Example 3.1.24. In $\mathbb{R}$ equipped with its canonical metric

$$\mathrm{diam}((0,1)) = \mathrm{diam}([0,1]) = 1$$

whereas $\mathrm{diam}(\mathbb{R}) = \infty$.

Theorem 3.1.25 (The Heine-Borel Theorem). *Let $K \subseteq \mathbb{K}^n$. Then K is compact in $(\mathbb{K}^n, \|\cdot\|_\infty)$ if and only if K is closed and bounded.*

Proof. First, suppose K is compact. As any subspace of a metric space (X, d) has topology induced by a metric that was induced from d by Proposition 1.4.5, Theorem 3.1.22 implies that K is bounded in $(\mathbb{K}^n, \|\cdot\|_\infty)$. Furthermore, as the subspace of any Hausdorff topological space is Hausdorff, Theorem 3.1.13 implies that K is closed.

Conversely, let K be closed and bounded subspace of $(\mathbb{K}^n, \|\cdot\|_\infty)$. Suppose to the contrary that K is not compact. Hence there exists an open cover $\{U_\alpha\}_{\alpha\in I}$ of K that has no finite subcover.

Since K is bounded, there exists an $M \in \mathbb{R}$ such that

$$K \subseteq [-M, M] \times \cdots \times [-M, M]$$

when $\mathbb{K} = \mathbb{R}$, and

$$K \subseteq \{(a_1 + b_1 i, \ldots, a_n + b_n i) \mid a_i, b_j \subseteq [-M, M]\}$$

when $\mathbb{K} = \mathbb{C}$. We will proceed with the proof where $\mathbb{K} = \mathbb{R}$ as the case where $\mathbb{K} = \mathbb{C}$ follows by the same arguments using $2n$ in place of n.

Divide $[-M, M]^n$ into 2^n closed balls with side-lengths M. To be specific, for all $q_1, \ldots, q_n \in \{0, 1\}$ let

$$J_{q_1,\ldots,q_n} = [-M + Mq_1, Mq_1] \times \cdots \times [-M + Mq_n, Mq_n].$$

Clearly each $J_{q_1,\ldots,q_n}$ is closed and the union of all possible $J_{q_1,\ldots,q_n}$s contains K. Therefore, since $\{U_\alpha\}_{\alpha\in I}$ does not have a finite subcover of K, there must exist one of these $J_{q_1,\ldots,q_n}$s such that $\{U_\alpha\}_{\alpha\in I}$ does not have a finite subcover of $K \cap J_{q_1,\ldots,q_n}$ (as there are a finite number of $J_{q_1,\ldots,q_n}$s). Denote this $J_{q_1,\ldots,q_n}$ by B_1 and notice $\mathrm{diam}(B_1) = M$.

Suppose for each $k \in \mathbb{N}$ we have constructed closed balls $B_1, \ldots, B_k$ such that $B_{j+1} \subseteq B_j$, $\mathrm{diam}(B_j) = \frac{1}{2^j}M$, and $\{U_\alpha\}_{\alpha\in I}$ does not have a finite subcover of $B_j \cap K$ for all $j \in \{1, \ldots, k-1\}$. By repeating the above process on B_k, there exists a closed ball $B_{k+1} \subseteq B_k$ such that $\mathrm{diam}(B_{k+1}) = \frac{1}{2^{k+1}}M$ and such that $\{U_\alpha\}_{\alpha \in I}$ does not have a finite subcover of $B_{k+1} \cap K$. Thus, by repeating this process ad infinitum, we obtain a collection $\{B_k\}_{k=1}^\infty$ of closed balls of $\mathbb{K}^n$ such that $B_{k+1} \subseteq B_k$, $\mathrm{diam}(B_k) = \frac{1}{2^k}M$, and $\{U_\alpha\}_{\alpha \in I}$ does not have a finite subcover of $V_k \cap K$ for all $k \in \mathbb{N}$ (and thus $B_k \cap K \neq \emptyset$ for all $k \in \mathbb{N}$).

For each $k \in \mathbb{N}$, let $x_k \in B_k \cap K$. Then, as $\mathrm{diam}(B_k \cap K) \subseteq \mathrm{diam}(B_k) \leq \frac{1}{2^k}M$, we see for all $n \geq m \geq N$ that $x_n, x_m \in B_N \cap K$ so

$$d(x_n, x_m) \leq \frac{1}{2^N}M$$

Therefore, since $\lim_{N\to\infty} M\frac{1}{2^N} = 0$, we see that $(x_n)_{n\geq 1}$ is a Cauchy sequence. Hence, as $\mathbb{K}^n$ is complete, there exists an $x_0 \in \mathbb{K}^n$ such that $\lim_{n\to\infty} x_n = x_0$ (see Section 4.1 for more detail if necessary). Moreover, since $x_n \in B_m \cap K$ for all $n \geq m$, the fact that $B_m \cap K$ is closed implies that $x_0 \in B_m \cap K$ for all $m \in \mathbb{N}$. Hence

$$Y = \bigcap_{k=1}^\infty (I_k \cap K) \neq \emptyset.$$

We claim that Y has exactly one element. Indeed if $x, y \in Y$ then $x, y \in B_k$ for all $k \in \mathbb{N}$ so $d(x, y) \leq \mathrm{diam}(B_k) = \frac{1}{2^k}M$ for all $k \in \mathbb{N}$ which implies $d(x, y) = 0$, or, equivalently, $x = y$. Hence Y contains exactly one point, say z.

By construction $z \in K$. Therefore, as $\{U_\alpha\}_{\alpha \in I}$ is an open cover of K, there exists an $\alpha_0 \in I$ such that $z \in U_{\alpha_0}$. Thus, since U_{α_0} is open, there exists an $\epsilon > 0$ such that $B(z, \epsilon) \subseteq U_{\alpha_0}$. Since $\mathrm{diam}(B_k) = \frac{1}{2^k}M$ for all $k \in \mathbb{N}$, there exists a $k_0 \in \mathbb{N}$ such that $\mathrm{diam}(B_{k_0}) < \epsilon$. Therefore, as $z \in B_{k_0}$ we obtain for all $x \in B_{k_0}$ that $d(z, x) < \epsilon$ so $x \in B(z, \epsilon) \subseteq U_{\alpha_0}$ for all $x \in B_{k_0}$. This implies $B_{k_0} \cap K \subseteq I_{k_0} \subseteq B(z, \epsilon) \subseteq U_{\alpha_0}$ which contradicts the fact that $\{U_\alpha\}_{\alpha \in I}$ did not have a finite subcover of $B_{k_0} \cap K$. As we have obtained a contradiction, it must be the case that K is compact. ■

Now that we have the Heine-Borel Theorem (Theorem 3.1.25) and thus a plethora of examples of compact topological spaces, we return to our initial motivation for compact topological spaces; namely a generalization to the Extreme Value Theorem to topological spaces. To obtain this characterization knowing that every finite closed interval in $\mathbb{R}$ is compact, we note the following exemplary property of compact topological spaces.

Theorem 3.1.26 (The Extreme Value Theorem). *Let $(X, \mathcal{T}_X)$ and $(Y, \mathcal{T}_Y)$ be topological spaces and let $f : X \to Y$ be continuous. If $(X, \mathcal{T}_X)$ is compact, then $f(X)$ is a compact subspace of $(Y, \mathcal{T}_Y)$.*

Proof. To see that $f(X)$ is compact, let $\{U_\alpha\}_{\alpha \in I}$ be an arbitrary open cover of $f(X)$ in $(Y, \mathcal{T}_Y)$. Therefore $\{f^{-1}(U_\alpha)\}_{\alpha \in I}$ is an open cover of $(X, \mathcal{T}_X)$.

Hence, as $(X, \mathcal{T}_X)$ is compact, there exists a finite subset $J \subseteq I$ such that $\{f^{-1}(U_\alpha)\}_{\alpha\in J}$ is an open cover of $(X, \mathcal{T}_X)$. Therefore $f(X) \subseteq \bigcup_{\alpha\in J} U_\alpha$ so $\{U_\alpha\}_{\alpha\in J}$ is a finite subcover of $f(X)$ from $\{U_\alpha\}_{\alpha\in I}$ Therefore, as $\{U_\alpha\}_{\alpha\in I}$ was arbitrary, $f(X)$ is compact. ∎

Theorem 3.1.26 has some wide-reaching implications.

Theorem 3.1.27 (The Extreme Value Theorem). *Let $(X, \mathcal{T})$ be a compact topological space and let $f : X \to \mathbb{R}$ be continuous. Then there exists points $x_1, x_2 \in X$ such that $f(x_1) \leq f(x) \leq f(x_2)$ for all $x \in X$.*

Proof. Since f is continuous and X is compact, Theorem 3.1.26 implies that $f(X)$ is a compact subset of $\mathbb{R}$. Hence $f(X)$ is closed and bounded by the Heine-Borel Theorem (Theorem 3.1.25). Since $f(X)$ is non-empty and bounded, $\sup(f(X))$ and $\inf(f(X))$ are finite and we can construct sequences of elements of $f(X)$ converging to $\sup(f(X))$ and $\inf(f(X))$ respectively. Since $f(X)$ is also closed, this implies $\sup(f(X)), \inf(f(X)) \in f(X)$. Hence there exists $x_1, x_2 \in X$ such that $f(x_1) = \inf(f(X))$ and $f(x_2) = \sup(f(X))$ so $f(x_1) \leq f(x) \leq f(x_2)$ for all $x \in X$ as desired. ∎

Of course, perhaps the notion of being 'closed and bounded' would be a nice metric space property to extend to the general topological setting and obtain nice results with respect to continuous maps. The following example shows this is not the case.

Example 3.1.28. Consider the metric space $(\mathbb{Z}, d)$ where $d : \mathbb{Z}\times\mathbb{Z} \to [0, \infty)$ is defined by

$$d(n, m) = \frac{|n - m|}{1 + |n - m|}$$

for all $n, m \in \mathbb{Z}$. Clearly d is well-defined. To see that d is a metric on $\mathbb{Z}$, notice $d(n, m) = d(m, n)$ for all $n, m \in \mathbb{Z}$ and $d(n, m) = 0$ if and only if $|n - m| = 0$ if and only if $n = m$. To see that d satisfies the triangle inequality, notice for all $a, b \in \mathbb{Z}$ that

$$\begin{aligned}
|a + b| &\leq |a| + |b| \\
\Rightarrow |a + b| &\leq |a| + |b| + 2|a||b| + |a||b||a + b| \\
\Rightarrow |a + b|(1 + |a|)(1 + |b|) &\leq |a|(1 + |b|)(1 + |a + b|) \\
&\quad + |b|(1 + |a|)(1 + |a + b|) \\
\Rightarrow \frac{|a + b|}{1 + |a + b|} &\leq \frac{|a|(1 + |b|) + |b|(1 + |a|)}{(1 + |a|)(1 + |b|)} \\
\Rightarrow \frac{|a + b|}{1 + |a + b|} &\leq \frac{|a|}{1 + |a|} + \frac{|b|}{1 + |b|}.
\end{aligned}$$

Therefore, if $n, m, q \in \mathbb{Z}$, by letting $a = n - q$ and $b = q - m$, we obtain that

$$d(n, m) \leq d(n, q) + d(q, m)$$

so d satisfies the triangle inequality. Hence d is a metric.

Notice that $d(n,m) \in [0,1)$ for all $n, m \in \mathbb{Z}$. Hence $\mathbb{Z}$ is a closed, bounded subset of $(\mathbb{Z}, d)$. Furthermore, for each $n \in \mathbb{Z}$,

$$\inf\{d(n,m) \mid m \in \mathbb{Z} \setminus \{n\}\} = \min\{d(n,n+1), d(n,n-1)\} > 0$$

as the function $x \mapsto \frac{x}{1+x}$ is increasing on $[0,\infty)$. Hence d induces the discrete topology on $\mathbb{Z}$ and thus every function from $\mathbb{Z}$ to a metric space must be continuous.

Define $f : \mathbb{Z} \to \mathbb{R}$ by

$$f(n) = \begin{cases} n+1 & \text{if } n \geq 0 \\ -\frac{1}{n} & \text{if } n < 0 \end{cases}$$

for all $n \in \mathbb{N}$. Thus f is continuous. However, as

$$f(\mathbb{Z}) = \mathbb{N} \cup \left\{\frac{1}{n} \mid n \in \mathbb{N}\right\}$$

we clearly see that $f(\mathbb{Z})$ is neither closed nor bounded in $\mathbb{R}$.

Another application of Theorem 3.1.26 is the following.

Corollary 3.1.29. *Every topological space homeomorphic to a compact topological space is compact.*

Proof. The result easily follows from Theorem 3.1.26. ∎

To finish off our preliminary study of compact topological spaces, we further note that the notions of compact and Hausdorff topological spaces intertwine nicely.

Theorem 3.1.30. *Let $(X, \mathcal{T}_X)$ be a compact topological space and let $(Y, \mathcal{T}_Y)$ be a Hausdorff space. If $f : X \to Y$ is a continuous bijection, then f is a homeomorphism. Thus $(X, \mathcal{T}_X)$ is Hausdorff and $(Y, \mathcal{T}_Y)$ is compact.*

Proof. Suppose $f : X \to Y$ is a continuous bijection. To see that $f^{-1} : Y \to X$ is continuous, let $U \in \mathcal{T}_X$ be arbitrary. Then $K = X \setminus U$ is a closed subset of $(X, \mathcal{T}_X)$ and thus compact by Theorem 3.1.14. Hence $(f^{-1})^{-1}(K) = f(K)$ is a compact subset of $(Y, \mathcal{T}_Y)$ by Theorem 3.1.26. Therefore, since $(Y, \mathcal{T}_Y)$ is Hausdorff, $f(K)$ is closed by Theorem 3.1.14. Thus

$$(f^{-1})^{-1}(U) = f(U) = f(X \setminus K) = f(X) \setminus f(K) = Y \setminus f(K)$$

is open in $(Y, \mathcal{T}_Y)$. Therefore, as $U \in \mathcal{T}_X$ was arbitrary, f^{-1} is continuous. Hence f is a homeomorphism.

The facts that $(X, \mathcal{T}_X)$ is Hausdorff and $(Y, \mathcal{T}_Y)$ is compact then follow as homeomorphisms clearly preserve these topological properties. ∎

3.2 Other Characterizations of Compactness

Now that we have some knowledge of the basics and some examples of compact topological spaces, we turn our attention to equivalent characterizations of compact topological spaces in the general topological setting. For example, in $\mathbb{R}$, the Bolzano-Weierstrass Theorem states (upto a reformulation) that a set $A \subseteq \mathbb{R}$ is closed and bounded if and only every sequence of elements of A has a convergent subsequence to an element of A. As the Heine-Borel Theorem (Theorem 3.1.25) implies the closed, bounded subsets of $\mathbb{R}$ are exactly the compact subsets of $\mathbb{R}$, it is natural to ask whether 'every sequence having a convergent subsequence' is a characterization of compact topological spaces. Of course, as sequences are insufficient in general topological spaces, it is more useful to ask whether 'every net has a convergent subnet' is an equivalent characterization of compact topological spaces.

In this section, our main (and only) result will show this is indeed the case. In addition, there is another incredibly useful characterization of compact topological spaces using intersections of certain collections of subsets.

Definition 3.2.1. Let $(X, \mathcal{T})$ be a topological space. A collection $\{F_\alpha\}_{\alpha\in I}$ is said to have the *finite intersection property* if $\bigcap_{\alpha\in J} F_\alpha \neq \emptyset$ for every finite subset $J \subseteq I$.

Theorem 3.2.2. *Let $(X, \mathcal{T})$ be a topological. The following are equivalent:*

(i) $(X, \mathcal{T})$ *is compact.*

(ii) *Whenever $\{F_\alpha\}_{\alpha\in I}$ is a collection of closed subsets of $(X, \mathcal{T})$ with the finite intersection property, $\bigcap_{\alpha\in I} F_\alpha \neq \emptyset$.*

(iii) *For every net $(x_\lambda)_{\lambda\in\Lambda}$, there exists an $x_0 \in X$ such that*

$$x_0 \in \overline{\{x_\lambda \mid \lambda \in \Lambda \text{ such that } \lambda \geq \lambda_0\}}$$

for every $\lambda_0 \in \Lambda$ (often x_0 is called a cluster point of the net*).*

(iv) *Every net in $(X, \mathcal{T})$ has a convergent subnet.*

Proof. To see that (i) implies (ii), let $(X, \mathcal{T})$ be a compact topological space. Suppose $\{F_\alpha\}_{\alpha\in I}$ is a set of closed subsets of $(X, \mathcal{T})$ with the finite intersection property yet $\bigcap_{\alpha\in I} F_\alpha = \emptyset$. For each $\alpha \in I$, let $U_\alpha = X \setminus F_\alpha$. Hence $\{U_\alpha\}_{\alpha\in I}$ are open subsets of $(X, \mathcal{T})$ as $\{F_\alpha\}_{\alpha\in I}$ is a set of closed subsets of $(X, \mathcal{T})$. Since

$$\bigcup_{\alpha\in I} U_\alpha = \bigcup_{\alpha\in I} X \setminus F_\alpha = X \setminus \left(\bigcap_{\alpha\in I} F_\alpha\right) = X \setminus \emptyset = X,$$

we see that $\{U_\alpha\}_{\alpha\in I}$ is an open cover of X. Therefore, as $(X, \mathcal{T})$ is compact, there exists a finite subset $J \subseteq I$ such that

$$X = \bigcup_{\alpha\in J} U_\alpha.$$

Hence

$$\emptyset = X \setminus X = X \setminus \left(\bigcup_{\alpha\in J} U_\alpha\right) = \bigcap_{\alpha\in J} X \setminus U_\alpha = \bigcap_{\alpha\in J} F_\alpha$$

thereby contradicting the fact that $\{F_\alpha\}_{\alpha\in I}$ has the finite intersection property. Thus, as we have obtained our contradiction, (i) implies (ii).

To see that (ii) implies (iii), suppose (ii) holds. Let $(x_\lambda)_{\lambda\in\Lambda}$ be an arbitrary net in $(X, \mathcal{T})$. For each $\lambda \in \Lambda$, let

$$A_\lambda = \{x_{\lambda'} \mid \lambda' \geq \lambda\} \qquad \text{and} \qquad F_\lambda = \overline{A_\lambda} \subseteq X.$$

Clearly $\{F_\lambda\}_{\lambda\in\Lambda}$ are closed subsets of X. We claim that $\{F_\lambda\}_{\lambda\in\Lambda}$ has the finite intersection property. Indeed suppose $J \subseteq \Lambda$ is finite. Since J is finite, the existence of upper bounds in directed sets implies there exists a $\lambda_0 \in \Lambda$ such that $\lambda_0 \geq \lambda$ for all $\lambda \in \Lambda$. Hence

$$x_{\lambda_0} \in \bigcap_{\lambda\in J} F_\lambda \quad \text{so} \quad \bigcap_{\lambda\in J} F_\lambda \neq \emptyset.$$

Therefore, as $J \subseteq I$ was an arbitrary finite subset, $\{F_\lambda\}_{\lambda\in\Lambda}$ has the finite intersection property. By the assumption of (ii), we know that there exists an $x_0 \in \bigcap_{\lambda\in\Lambda} F_\lambda$ as desired. Hence, as $(x_\lambda)_{\lambda\in\Lambda}$ was arbitrary, (ii) implies (iii).

To see that (iii) implies (iv), let $(x_\lambda)_{\lambda\in\Lambda}$ be an arbitrary net in $(X, \mathcal{T})$. For each $\lambda \in \Lambda$, let

$$A_\lambda = \{x_{\lambda'} \mid \lambda' \geq \lambda\} \qquad \text{and} \qquad F_\lambda = \overline{A_\lambda} \subseteq X.$$

By the assumption of (iii), there exists an $x_0 \in \bigcap_{\lambda\in\Lambda} F_\lambda$.

We claim that there exists a subnet of $(x_\lambda)_{\lambda\in\Lambda}$ that converges to x_0. To see this, first notice for an arbitrary neighbourhood U of x_0 that $A_\lambda \cap U \neq \emptyset$ by Theorem 1.6.21 as $x_0 \in F_\lambda$. Hence, for every neighbourhood U of x_0 and for every $\lambda \in \Lambda$ there exists a $\lambda' \in \Lambda$ such that $\lambda' \geq \lambda$ and $x_{\lambda'} \in U$.

Let

$$M = \{(U, \lambda) \mid U \text{ a neighbourhood of } x_0 \text{ and } \lambda \in \Lambda \text{ such that } x_\lambda \in U\},$$

which is non-empty by previous discussions. For two pairs $(U_1, \lambda_1), (U_2, \lambda_2) \in M$, define $(U_1, \lambda_1) \leq (U_2, \lambda_2)$ if and only if $U_1 \supseteq U_2$ and $\lambda_1 \leq \lambda_2$. We claim that $(M, \leq)$ is a directed set. Indeed it is clear $(M, \leq)$ is reflexive and transitive since reverse inclusion and the ordering on Λ are. Finally let

$(U_1, \lambda_1), (U_2, \lambda_2) \in M$ be arbitrary. Let $U_3 = U_1 \cap U_2$. Clearly U_3 is a neighbourhood of x_0 as U_1 and U_2 are. As $(\Lambda, \leq)$ is a directed set, there exists a $\lambda' \in \Lambda$ such that $\lambda' \geq \lambda_1$ and $\lambda' \geq \lambda_2$. By the previous paragraph there exists a $\lambda_3 \in \Lambda$ such that $\lambda_3 \geq \lambda'$ (so $\lambda_3 \geq \lambda_1$ and $\lambda_3 \geq \lambda_2$) and $x_{\lambda_3} \in U_3$. Hence $(U_3, \lambda_3) \in M$, $(U_3, \lambda_3) \geq (U_1, \lambda_1)$, and $(U_3, \lambda_3) \geq (U_2, \lambda_2)$. Therefore, as $(U_1, \lambda_1), (U_2, \lambda_2) \in M$ were arbitrary, $(M, \leq)$ is a directed set.

We claim that $(x_\lambda)_{(U,\lambda)\in M}$ is a subnet of $(x_\lambda)_{\lambda\in\Lambda}$. To see this, define $\varphi : M \to \Lambda$ by $\varphi((U, \lambda)) = \lambda$. Clearly φ is increasing by the definition of the ordering on M. To see that φ is cofinal, let $\lambda \in \Lambda$ be arbitrary. Then clearly $(X, \lambda) \in M$ and $\varphi((X, \lambda)) = \lambda \geq \lambda$. Hence $(x_\lambda)_{(U,\lambda)\in M}$ is a subnet of $(x_\lambda)_{\lambda\in\Lambda}$ by Definition 1.5.44

Finally, we claim that $(x_\lambda)_{(U,\lambda)\in M}$ converges to x_0. To see this, let U be an arbitrary neighbourhood of x_0. From previous discussions there exists a $\lambda \in \Lambda$ such that $(U, \lambda) \in M$. Thus for all $(U', \lambda') \geq (U, \lambda)$ we have that $x_{\lambda'} \in U' \subseteq U$. Therefore, as U was arbitrary, $(x_\lambda)_{(U,\lambda)\in M}$ is a subnet of $(x_\lambda)_{\lambda\in\Lambda}$ that converges to x_0. Therefore, as $(x_\lambda)_{\lambda\in\Lambda}$ was arbitrary, (iii) implies (iv).

To see that (iv) implies (i), suppose (iv) holds. To see that $(X, \mathcal{T})$ is compact, suppose to the contrary that there exists an open cover $\{U_\alpha\}_{\alpha\in I}$ of $(X, \mathcal{T})$ that has no finite subcover. We will use $\{U_\alpha\}_{\alpha\in I}$ to construct a net that has no convergent subnets.

Let

$$\Lambda = \left\{ U \subseteq X \,\middle|\, U = \bigcup_{\alpha\in J} U_\alpha \text{ for some finite subset } J \subseteq U \right\}.$$

For two sets $U_1, U_2 \in \Lambda$, define $U_1 \leq U_2$ if and only if $U_1 \subseteq U_2$. Since Λ is closed under finite unions, Example 1.5.7 implies $(\Lambda, \leq)$ is a directed set.

Since $\{U_\alpha\}_{\alpha\in I}$ has no finite subcover of $(X, \mathcal{T})$, $X \setminus \left(\bigcup_{\alpha\in J} U_\alpha\right) \neq \emptyset$ for each finite set $J \subseteq I$. Hence, for each $U \in \Lambda$ we may choose a point $x_U \in X \setminus U$. Thus $(x_U)_{U\in\Lambda}$ is a net in $(X, \mathcal{T})$.

We claim that $(x_U)_{U\in\Lambda}$ has no convergent subnets thereby contradicting the assumption that (iv) holds and yielding (iv) implies (i). To see this, suppose to the contrary that $(x_U)_{U\in\Lambda}$ has a subnet $(x_{\lambda_\mu})_{\mu\in M}$ that converges to some point $x_0 \in X$. Since $\{U_\alpha\}_{\alpha\in I}$ is an open cover of $(X, \mathcal{T})$, there exists an $\alpha_0 \in I$ such that $x_0 \in U_{\alpha_0}$. Thus, as $(x_{\lambda_\mu})_{\mu\in M}$ is a subnet of $(x_U)_{U\in\Lambda}$, there exists a μ_0 such that $\lambda_{\mu_0} \geq U_{\alpha_0}$. Moreover, since $(x_{\lambda_\mu})_{\mu\in M}$ converges to x_0, there exists an $\mu_1 \in M$ such that $x_{\lambda_\mu} \in U_{\alpha_0}$ for all $\mu \geq \mu_1$. By the definition of a directed set, there exists a $\mu_2 \in M$ such that $\mu_2 \geq \mu_0$ and $\mu_2 \geq \mu_1$. Hence $\lambda_{\mu_2} \geq U_{\alpha_0}$ and $x_{\lambda_{\mu_2}} \in U_{\alpha_0}$. However, $\lambda_{\mu_2} \geq U_{\alpha_0}$ implies that $\lambda_{\mu_2} = U$ for some open set U in $(X, \mathcal{T})$ such that $U \supseteq U_{\alpha_0}$ so the definition of $x_{\lambda_{\mu_2}}$ implies that

$$x_{\lambda_{\mu_2}} \notin U \qquad \text{so} \qquad x_{\lambda_{\mu_2}} \notin U_{\alpha_0}.$$

As this contradicts the fact that $x_{\lambda_{\mu_2}} \in U_{\alpha_0}$, $(x_U)_{U\in\Lambda}$ has no convergent subnets. Hence (iv) implies (i). ∎

3.3 Tychonoff's Theorem

Theorem 3.2.2 is incredibly useful in verifying topological spaces are compact using either the finite intersection property or the 'every net has a convergent subnet characterization'. Furthermore, once compactness is established, we can use these characterizations for several useful applications. Consequently, our current goal is to use Theorem 3.2.2 to increase our repertoire of compact topological spaces.

One common theme throughout this course has been to look at products of topological spaces either equipped with the box or product topology. The following elementary result shows us what must be true for a product of topological spaces to be compact.

Proposition 3.3.1. *Let $\{(X_\alpha, \mathcal{T}_\alpha)\}_{\alpha\in I}$ be a non-empty collection of non-empty topological spaces. If the product $\prod_{\alpha\in I} X_\alpha$ is a compact topological space when equipped with the either the box or product topology, then $(X_\alpha, \mathcal{T}_\alpha)$ is compact for each $\alpha \in I$.*

Proof. Suppose $X = \prod_{\alpha\in I} X_\alpha$ is compact when equipped with either the box or product topology. To see that $(X_\alpha, \mathcal{T}_\alpha)$ is compact for each $\alpha \in I$, fix an arbitrary $\alpha_0 \in I$. Let $\pi_{\alpha_0} : X \to X_\alpha$ be the projection map from Example 2.1.8. Hence π_{α_0} is a continuous surjective map when X is equipped with the product topology by Example 2.1.8. Therefore, since X is compact, $\pi_{\alpha_0}(X) = X_{\alpha_0}$ is compact as desired. ∎

Perhaps unsurprising at this point, the box topology is the incorrect topology to place on a product in order to make the converse of Proposition 3.3.1 hold.

Example 3.3.2. The box topology on a product of compact topological spaces need not be compact. To see this, for each $n \in \mathbb{N}$, let $X_n = [0, 1]$ equipped with the subspace topology inherited from $\mathbb{R}$. Clearly X_n is a compact subset of $\mathbb{R}$ for each $n \in \mathbb{N}$ by the Heine-Borel Theorem (Theorem 3.1.25). Let $X = \prod_{n\in\mathbb{N}} X_n$. To see that X is not compact when equipped with the box topology, consider the set

$$\mathcal{U} = \left\{ \prod_{n\in\mathbb{N}} I_n \,\middle|\, I_n \in \left\{ \left[0, \frac{2}{3}\right), \left(\frac{1}{3}, 1\right] \right\} \text{ for all } n \in \mathbb{N} \right\}.$$

Since

$$\left[0, \frac{2}{3}\right) = \left(-\frac{2}{3}, \frac{2}{3}\right) \cap [0, 1] \qquad \text{and} \qquad \left(\frac{1}{3}, 1\right] = \left(\frac{1}{3}, \frac{4}{3}\right) \cap [0, 1]$$

are open subsets of X_n with union X_n for all $n \in \mathbb{N}$, we see that $\mathcal{U}$ is an open cover of X when equipped with the box topology. However, since each element of the infinite set

$$\prod_{n\in\mathbb{N}} \{0,1\} \subseteq X$$

is contained in exactly one element of $\mathcal{U}$, it is impossible that $\mathcal{U}$ has a finite subcover of X. Hence the box topology on a product of compact topological spaces need not be compact.

However, the product of compact topological spaces will be compact when equipped with the product topology! To show this amazing theorem, we will invoke the finite intersection portion of Theorem 3.2.2. This will be accomplished by using a maximal collection of sets with the finite intersection property. Thus we present the following lemma.

Lemma 3.3.3. *Let $(X, \mathcal{T})$ be a topological space and let $\mathcal{F} \subseteq \mathcal{P}(X)$ be a non-empty collection of sets with the finite intersection property. Then there exists an $\mathcal{M} \subseteq \mathcal{P}(X)$ such that*

(1) $\mathcal{F} \subseteq \mathcal{M}$,

(2) $\mathcal{M}$ has the finite intersection property,

(3) if $F \in \mathcal{P}(X) \setminus \mathcal{M}$, then $\mathcal{M} \cup \{F\}$ does not have the finite intersection property,

(4) if $\{F_k\}_{k=1}^n \subseteq \mathcal{M}$ for some $n \in \mathbb{N}$, then $\bigcap_{k=1}^n F_k \in \mathcal{M}$, and

(5) if $Y \subseteq X$ and $Y \cap M \neq \emptyset$ for all $M \in \mathcal{M}$, then $Y \in \mathcal{M}$.

Proof. Let

$$\mathcal{C} = \{\mathcal{S} \subseteq \mathcal{P}(X) \mid \mathcal{F} \subseteq \mathcal{S} \text{ and } \mathcal{S} \text{ has the finite intersection property}\}.$$

Clearly $\mathcal{C} \neq \emptyset$ since $\mathcal{F} \subseteq \mathcal{C}$. For $\mathcal{S}_1, \mathcal{S}_2 \in \mathcal{C}$, define $\mathcal{S}_1 \preceq \mathcal{S}_2$ if and only if $\mathcal{S}_1 \subseteq \mathcal{S}_2$. Clearly $(\mathcal{C}, \preceq)$ is a partially ordered set.

We claim that every chain in $(\mathcal{C}, \preceq)$ has an upper bound. To see this, suppose that $\{\mathcal{S}_\alpha\}_{\alpha\in I}$ is a chain in $(\mathcal{C}, \preceq)$. Let

$$\mathcal{S} = \bigcup_{\alpha\in I} \mathcal{S}_\alpha.$$

We claim that $\mathcal{S} \in \mathcal{C}$ from which it trivially follows that $\mathcal{S}$ is an upper bound for $\{\mathcal{S}_\alpha\}_{\alpha\in I}$. To see that $\mathcal{S} \in \mathcal{C}$, first notice since $\mathcal{F} \subseteq \mathcal{S}_\alpha$ for all $\alpha \in I$ that $\mathcal{F} \subseteq \mathcal{S}$ by construction. To see that $\mathcal{S}$ has the finite intersection property, suppose that $n \in \mathbb{N}$ and $S_1, S_2, \ldots, S_n \in \mathcal{S}$. By the properties of a chain, there exists an $\alpha_0 \in I$ such that $S_k \in \mathcal{S}_{\alpha_0}$ for all $k \in \{1, \ldots, n\}$. Therefore, since $\mathcal{S}_{\alpha_0}$ has the finite intersection property as $\mathcal{S}_{\alpha_0} \in \mathcal{C}$, we obtain that

$\bigcap_{k=1}^n S_k \neq \emptyset$. Therefore, as $n \in \mathbb{N}$ and $S_1, \ldots, S_n \in \mathcal{S}$ were arbitrary, $\mathcal{S}$ has the finite intersection property and thus $\mathcal{S} \in \mathcal{C}$.

Since $(\mathcal{C}, \preceq)$ is a non-empty partially ordered set such that every chain has an upper bound, Zorn's Lemma (Axiom A.5.10) implies that there exists an $\mathcal{M} \in \mathcal{C}$ such that if $\mathcal{S} \in \mathcal{C}$ and $\mathcal{M} \preceq \mathcal{S}$, then $\mathcal{S} = \mathcal{M}$ (i.e. $\mathcal{M}$ is maximal in $(\mathcal{C}, \preceq)$). We claim that $\mathcal{M}$ has the desired properties. Indeed $\mathcal{F} \subseteq \mathcal{M}$ and $\mathcal{M}$ has the finite intersection property since $\mathcal{M} \in \mathcal{C}$. Thus (1) and (2) hold.

To see that (3) holds, let $F \in \mathcal{P}(X) \setminus \mathcal{M}$ be arbitrary. If $\mathcal{M}_0 = \mathcal{M} \cup \{F\}$ had the finite intersection property, then since $\mathcal{F} \subseteq \mathcal{M} \subseteq \mathcal{M}_0$ we would have that $\mathcal{M}_0 \in \mathcal{C}$, $\mathcal{M}_0 \neq \mathcal{M}$ as $F \in \mathcal{M}_0 \setminus \mathcal{M}$, and $\mathcal{M} \preceq \mathcal{M}_0$ thereby contradicting the maximality of $\mathcal{M}$. Therefore (3) holds.

To see that (4) holds, let $n \in \mathbb{N}$ and $\{F_k\}_{k=1}^n \subseteq \mathcal{M}$ be arbitrary. If $F = \bigcap_{k=1}^n F_k$, then clearly $\mathcal{M} \cup \{F\}$ has the finite intersection property as $\mathcal{M}$ has the finite intersection property. Thus (3) implies that $F \in \mathcal{M}$. Therefore, as $n \in \mathbb{N}$ and $\{F_k\}_{k=1}^n \subseteq \mathcal{M}$ were arbitrary, (4) follows.

Finally, to see that (5) holds, suppose $Y \subseteq X$ is such that $Y \cap M \neq \emptyset$ for all $M \in \mathcal{M}$. Therefore (4) implies that for all $n \in \mathbb{N}$ and $\{F_k\}_{k=1}^n \subseteq \mathcal{M}$ that $Y \cap (\bigcap_{k=1}^n F_k) \neq \emptyset$. Thus $\mathcal{M} \cup \{Y\}$ has the finite intersection property so (3) implies that $Y \in \mathcal{M}$ as desired. ∎

Theorem 3.3.4 (Tychonoff's Theorem). *Let $\{(X_\alpha, \mathcal{T}_\alpha)\}_{\alpha \in I}$ be compact topological spaces. Then $\prod_{\alpha \in I} X_\alpha$ is a compact topological space when equipped with the product topology.*

Proof. Let $X = \prod_{\alpha \in I} X_\alpha$ and let $\mathcal{T}$ denote the product topology on X. To see that $(X, \mathcal{T})$ is compact, we will apply Theorem 3.2.2 and verify that any set of closed subsets of $(X, \mathcal{T})$ with finite intersection property has non-empty intersection.

Let $\mathcal{F}$ be an arbitrary set of closed subsets of $(X, \mathcal{T})$ with the finite intersection property. Let $\mathcal{M}$ be a set with the finite intersection property containing $\mathcal{F}$ as created via Lemma 3.3.3. Since

$$\bigcap_{F \in \mathcal{F}} F \supseteq \bigcap_{A \in \mathcal{M}} \overline{A},$$

it suffices to show that $\bigcap_{A \in \mathcal{M}} \overline{A} \neq \emptyset$.

For each $\alpha \in I$, let $\pi_\alpha : X \to X_\alpha$ be the projection map from X to X_α from Example 2.1.8. Since $\mathcal{M}$ has the finite intersection property, it is clear that

$$\{\pi_\alpha(A) \mid A \in \mathcal{M}\}$$

has the finite intersection property in (X_α, T_α) so

$$\left\{\overline{\pi_\alpha(A)} \mid A \in \mathcal{M}\right\}$$

is a collection of closed sets in $(X_\alpha, \mathcal{T}_\alpha)$ with the finite intersection property. Therefore, since $(X_\alpha, \mathcal{T}_\alpha)$ is compact, Theorem 3.2.2 implies for all $\alpha \in I$ there exists an $x_\alpha \in X_\alpha$ such that

$$x_\alpha \in \bigcap_{A \in \mathcal{M}} \overline{\pi_\alpha(A)}.$$

Let $x = (x_\alpha)_{\alpha \in I} \in X$. We claim that $x \in \bigcap_{A \in \mathcal{M}} \overline{A}$ thereby completing the proof that $\bigcap_{A \in \mathcal{M}} \overline{A} \neq \emptyset$.

To begin, let $\alpha_0 \in I$ and $U \in \mathcal{T}_{\alpha_0}$ be such that $x_{\alpha_0} \in U$. Since $x_{\alpha_0} \in \overline{\pi_{\alpha_0}(A)}$ for all $A \in \mathcal{M}$, Theorem 1.6.21 implies that $\pi_{\alpha_0}(A) \cap U \neq \emptyset$ for all $A \in \mathcal{M}$. Hence $A \cap \pi_{\alpha_0}^{-1}(U) \neq \emptyset$ for all $A \in \mathcal{M}$. Therefore, the properties of $\mathcal{M}$ from Lemma 3.3.3 imply that $\pi_{\alpha_0}^{-1}(U) \in \mathcal{M}$ for all $\alpha_0 \in I$ and $U \in \mathcal{T}_{\alpha_0}$ such that $x_{\alpha_0} \in U$.

Since $\mathcal{M}$ is closed under finite intersections from Lemma 3.3.3,

$$\left\{ \bigcap_{\alpha \in J} \pi_\alpha^{-1}(U_\alpha) \,\middle|\, \substack{J \subseteq I \text{ finite and} \\ U_\alpha \text{ a } \mathcal{T}_\alpha\text{-neighbourhood of } x_\alpha \text{ for all } \alpha \in J} \right\}$$

is both contained in $\mathcal{M}$ and is a neighbourhood basis of x in $(X, \mathcal{T})$. Therefore, as $\mathcal{M}$ has the finite intersection property, every element of $\mathcal{M}$ has non-empty intersection with each element of a neighbourhood basis of x. Hence Theorem 1.6.21 implies that $x \in \overline{A}$ for all $A \in \mathcal{M}$. Thus $x \in \bigcap_{A \in \mathcal{M}} \overline{A}$ thereby completing the proof. ∎

Of course, the proof of Tychonoff's Theorem (Theorem 3.3.4) relies on Zorn's Lemma (Axiom A.5.10). However, it is well-known that Zorn's Lemma is equivalent to the Axiom of Choice (Axiom A.2.4). It is perhaps unsurprisingly that the Axiom of Choice is required to prove Tychonoff's Theorem for otherwise we would not know the product of compact topological spaces contains a point (although the empty set is technically compact). However, it is perhaps surprising that the Axiom of Choice is implied by Tychonoff's Theorem.

Theorem 3.3.5. *Suppose that Tychonoff's Theorem holds; that is, the product of compact topological spaces is compact when equipped with the product topology. Then for any non-empty set I and any set $\{X_\alpha\}_{\alpha \in I}$ of non-empty sets, the product $\prod_{\alpha \in I} X_\alpha$ is non-empty.*

Proof. Let I be an non-empty set and let $\{X_\alpha\}_{\alpha \in I}$ be a set of non-empty sets. For each $\alpha \in I$, let $Y_\alpha = X_\alpha \cup \{\infty_\alpha\}$ for some symbol ∞_α and let $Y = \prod_{\alpha \in I} Y_\alpha$. We note that Y is automatically non-empty without the use of the Axiom of Choice (Axiom A.2.4). Indeed we already know for all $\alpha \in I$ that $\infty_\alpha \in Y_\alpha$; that is, we do not need to choose an element of Y_α for each $\alpha \in I$ as we already know (i.e. have assigned) an element of Y_α for each

$\alpha \in I$. Hence the element $\infty = (\infty_\alpha)_{\alpha \in I}$ is an element of Y without the use of the Axiom of Choice (Axiom A.2.4).

For each $\alpha \in I$, let $\mathcal{T}_\alpha = \{\emptyset, Y_\alpha, X_\alpha, \{\infty_\alpha\}\}$. Clearly $\mathcal{T}_\alpha$ is a topology on Y_α. Furthermore, since $\mathcal{T}_\alpha$ only has a finite number of sets, every $\mathcal{T}_\alpha$-open cover of Y_α has a finite subcover (namely the original open cover) so $(Y_\alpha, \mathcal{T}_\alpha)$ is compact. Hence Tychonoff's Theorem implies that $Y = \prod_{\alpha \in I} Y_\alpha$ is compact when equipped with the product topology.

For each $\alpha_0 \in I$, let

$$U_{\alpha_0} = \prod_{\alpha \in I} U_{\alpha_0, \alpha}$$

where

$$U_{\alpha_0,\alpha} = \begin{cases} Y_\alpha & \text{if } \alpha \neq \alpha_0 \\ \{\infty_\alpha\} & \text{if } \alpha = \alpha_0 \end{cases}.$$

Again, the construction of U_{α_0} does not require the Axiom of Choice (Axiom A.2.4) since we do not need to choose an element of $\mathcal{T}_\alpha$ for each $\alpha \in I$ as we already know an element of $\mathcal{T}_\alpha$ for each $\alpha \in I$. Clearly U_{α_0} is open in the product topology on Y by definition.

We claim that $\mathcal{U} = \{U_\alpha\}_{\alpha \in I}$ cannot cover Y. To see this, suppose to the contrary that $\mathcal{U}$ is an open cover of Y. Since Y is compact, there exists a finite subset $J \subseteq I$ such that $Y = \bigcup_{\alpha \in J} U_\alpha$. For each $\alpha \in J$, choose $x_\alpha \in X_\alpha$. Note this does not require the Axiom of Choice (Axiom A.2.4) since J is finite. For each $\alpha \in I \setminus J$, let $x_\alpha = \infty_\alpha$. Again, this does not require the Axiom of Choice (Axiom A.2.4). Thus $x = (x_\alpha)_{\alpha \in I} \in Y$ by definition. However $x \notin \bigcup_{\alpha \in J} U_\alpha$ by construction since $x_\alpha \notin \{\infty_\alpha\}$ for all $\alpha \in J$. Hence we have a contradiction so $\mathcal{U}$ is not a cover of Y.

Since $\mathcal{U}$ is not a cover of Y, there must exist an element $y = (y_\alpha)_{\alpha \in I} \in Y$ such that $y \notin U_\alpha$ for all $\alpha \in I$. Then, by the definition of U_α, we see that $y_\alpha \notin \{\infty_\alpha\}$ for all $\alpha \in I$. Hence $y_\alpha \in X_\alpha$ for all $\alpha \in I$ so that $y \in \prod_{\alpha \in I} X_\alpha$. Hence Tychonoff's Theorem implies the Axiom of Choice. ∎

3.4 Local Compactness

Before moving onto discussing the notion of compactness in metric spaces, we study the notion of being compact near each point. In particular, we study the notion of a locally compact topological space. Perhaps surprisingly considering that connected and locally connected topological spaces behave very differently, the notions of compact and locally compact for Hausdorff topological spaces differ only by the inclusion of a single point!

Before we get to that, we define what we mean by a locally compact topological space. The following definition may not seem the correct one considering how we defined a locally connected topological space. In fact, there are some discrepancies in the literature on how to define a locally compact topological space. This is often the case in topology in that different authors can use different definitions. However, the definition we provide is the 'easiest' to work with when dealing with topologies and, if we restrict to Hausdorff topological spaces, Theorem 3.4.11 shows the definitions are equivalent.

Definition 3.4.1. Let $(X, \mathcal{T})$ be a topological space and let $x \in X$. It is said that $(X, \mathcal{T})$ is *locally compact at* x if there exists a compact subspace K of $(X, \mathcal{T})$ and a $U \in \mathcal{T}$ such that $x \in U \subseteq K$.

A topological space $(X, \mathcal{T})$ is said to be *locally compact* if $(X, \mathcal{T})$ is locally compact at each point of X.

Before we discuss examples, we note that the notion of a locally compact topological space is a good notion to consider for topological spaces since it is invariant under homeomorphisms.

Theorem 3.4.2. *Every topological space homeomorphic to a locally compact topological space is locally compact.*

Proof. As homeomorphisms map open sets to open sets and compact sets to compact sets by Theorem 3.1.26, the result trivially follows. ∎

Example 3.4.3. Clearly every compact topological space is locally compact by definition.

Example 3.4.4. The real numbers equipped with their canonical topology is locally compact. Indeed if $x \in \mathbb{R}$, then $x \in (x-1, x+1) \subseteq [x-1, x+1]$ where $(x-1, x+1)$ is open and $[x-1, x+1]$ is compact by Theorem 3.1.25. A similar argument immediately implies that $\mathbb{K}^n$ is locally compact when equipped with its Euclidean topology.

Example 3.4.5. Consider $\mathbb{R}^2$ equipped with its Euclidean topology and the subspace X of $\mathbb{R}^2$ given by

$$X = \{(0,0)\} \cup \{(x,y) \in \mathbb{R}^2 \mid y > 0\}.$$

We claim that X is not locally compact. To see this, suppose to the contrary that there exists $U, K \subseteq X$ such that U is open in X, K is compact in X, and $(0,0) \in U \subseteq K$. Thus there exists an $\epsilon > 0$ such that $B_2((0,0), \epsilon) \subseteq U$. Therefore there exists an $x \neq 0$ and a $\delta > 0$ such that

$$\{(x,y) \mid 0 < y < \delta\} \subseteq B_2((0,0), \epsilon) \subseteq U \subseteq K.$$

Choose $N \in \mathbb{N}$ such that $\frac{1}{N} < \delta$ and consider the sequence $(z_n)_{n \geq N}$ in $U \subseteq K$ where

$$z_n = \left(x, \frac{1}{n}\right)$$

for all $n \geq N$. We claim that $(z_n)_{n \geq N}$ has no convergent subnets in K and thus K cannot possibly be compact by Theorem 3.2.2. Indeed, if $(z_n)_{n \geq N}$

has a subnet that converged to $z \in K$, then for every $\epsilon > 0$ there must exist arbitrary large $n \geq N$ such that $d_2(z, z_n) < \epsilon$ as the subspace topology on X is the restriction of the Euclidean metric topology from $\mathbb{R}^2$ to X by Proposition 1.4.5. This implies that the subsequence of $(z_n)_{n\geq N}$ converges to z in $\mathbb{R}^2$. However, $(z_n)_{n\geq N}$ converges to $(x, 0)$ in $\mathbb{R}^2$ and thus so does every subnet as $\mathbb{R}^2$ is Hausdorff. However, $(x, 0) \notin X$ as $x \neq 0$ thereby yielding a contradiction. Hence X is not locally compact.

Example 3.4.5 raises an interesting question; what subspaces of locally compact topological spaces are locally compact? The following is not too surprising considering the connection between compactness and closed sets.

Proposition 3.4.6. *Let $(X, \mathcal{T})$ be a locally compact topological space and let Y be a subspace of $(X, \mathcal{T})$. If Y is closed in $(X, \mathcal{T})$, then Y is locally compact.*

Proof. To see that Y is locally compact, let $y \in Y$ be arbitrary. Since X is locally compact, there exists a compact subset K of $(X, \mathcal{T})$ and a $U \in \mathcal{T}$ such that $x \in U \subseteq K$. Let $V = Y \cap U \subseteq Y$ and let $K_0 = Y \cap K \subseteq Y$ so that $x \in V \subseteq K_0$. Notice V is open in Y by the definition of the subspace topology. Thus, to complete the proof, it suffices to show that K_0 is compact.

Notice K_0 is a closed subset of K when K is equipped with the subspace topology as Y is closed in X. Hence Theorem 3.1.14 implies that K_0 is a compact set. Therefore, as $y \in Y$ was arbitrary, Y is locally compact. ∎

In fact, Proposition 3.4.6 will hold if 'closed' is replaced with 'open' provided $(X, \mathcal{T}_X)$ is Hausdorff (see Corollary 3.4.12). However, adding the condition of Hausdorff to a locally compact topological space is incredibly nice and 'almost compact' as the following result demonstrates.

Theorem 3.4.7. *A Hausdorff topological space $(X, \mathcal{T}_X)$ is locally compact if and only if there exists a compact Hausdorff topological space $(Y, \mathcal{T}_Y)$ such that $(X, \mathcal{T}_X)$ is a subspace of $(Y, \mathcal{T}_Y)$ and $Y \setminus X$ consists of a single point. Furthermore, $(Y, \mathcal{T}_Y)$ is unique topological space (upto homeomorphism) with the above properties.*

Proof. First, suppose $(X, \mathcal{T}_X)$ is a Hausdorff space that is a subspace of a compact Hausdorff topological space $(Y, \mathcal{T}_Y)$ such that $Y \setminus X = \{\infty\}$ for some point ∞. To see that $(X, \mathcal{T}_X)$ is locally compact, let $x \in X$ be arbitrary. Since $(Y, \mathcal{T}_Y)$ is Hausdorff, there exist $U, V \in \mathcal{T}_Y$ such that $x \in U$, $\infty \in V$, and $U \cap V = \emptyset$. Consider $K = Y \setminus V$. Since $V \in \mathcal{T}_Y$, K is a closed subset of $(Y, \mathcal{T}_Y)$ and thus is a compact set by Theorem 3.1.14 as $(Y, \mathcal{T}_Y)$ is compact. Since $\infty \in V$, we see that $K \subseteq X$ and thus is a compact subset of $(X, \mathcal{T}_X)$ as $(X, \mathcal{T}_X)$ is a subspace of $(Y, \mathcal{T}_Y)$. Furthermore, as $U \cap V = \emptyset$, we obtain that $U \subseteq K$. Hence, as $x \in X$ was arbitrary, $(X, \mathcal{T}_X)$ is locally compact.

Conversely suppose that $(X, \mathcal{T}_X)$ is locally compact, Hausdorff topological space. Let $Y = X \cup \{\infty\}$ for some symbol ∞ and let

$$\mathcal{T}_Y = \mathcal{T}_X \cup \{Y \setminus K \mid K \text{ a compact subset of } (X, \mathcal{T}_X)\}.$$

We claim that $\mathcal{T}_Y$ is a topology on Y. To see this, first notice that $\emptyset \in \mathcal{T}_X \subseteq \mathcal{T}_Y$ and that $Y \in \mathcal{T}_Y$ since $Y \setminus Y = \emptyset$ is a compact subset of $(X, \mathcal{T}_X)$. Next, to see that $\mathcal{T}_Y$ is closed under arbitrary unions, let $\{U_\alpha\}_{\alpha \in I} \subseteq \mathcal{T}_X$ and $\{K_\beta\}_{\beta \in J}$ a set of compact subset of $(X, \mathcal{T}_X)$ be arbitrary. Clearly $U = \bigcup_{\alpha \in I} U_\alpha \in \mathcal{T}_X \subseteq \mathcal{T}_Y$. Notice

$$\bigcup_{\beta \in J} (Y \setminus K_\beta) = Y \setminus \left(\bigcap_{\beta \in J} K_\beta \right).$$

However, since $(X, \mathcal{T})$ is Hausdorff, $K = \bigcap_{\beta \in J} K_\beta$ is compact by Corollary 3.1.15. Thus the above implies that $\bigcup_{\beta \in J} (Y \setminus K_\beta) \in \mathcal{T}_Y$. Finally, we see that

$$\left(\bigcup_{\alpha \in I} U_\alpha \right) \cup \left(\bigcup_{\beta \in J} (Y \setminus K_\beta) \right) = U \cup (Y \setminus K) = Y \setminus (K \setminus U).$$

However, as we saw above, U is an open set in $(X, \mathcal{T}_X)$ and K is a compact subspace of $(X, \mathcal{T}_X)$ so $K \setminus U$ is a closed subset of K and thus compact by Theorem 3.1.14. Hence $Y \setminus (K \setminus U) \in \mathcal{T}_Y$ by definition. Thus $\mathcal{T}_Y$ is closed under arbitrary unions.

To see that $\mathcal{T}_Y$ is closed under finite intersections, let $\{U_\alpha\}_{\alpha \in I} \subseteq \mathcal{T}_X$ and $\{K_\beta\}_{\beta \in J}$ a set of compact subset of $(X, \mathcal{T}_X)$ be arbitrary with I and J finite. Clearly $U = \bigcap_{\alpha \in I} U_\alpha \in \mathcal{T}_X \subseteq \mathcal{T}_Y$. Notice

$$\bigcap_{\beta \in J} (Y \setminus K_\beta) = Y \setminus \left(\bigcup_{\beta \in J} K_\beta \right).$$

Since J is finite, $K = \bigcup_{\beta \in J} K_\beta$ is compact by Corollary 3.1.17. Thus the above implies that $\bigcup_{\beta \in J} (Y \setminus K_\beta) \in \mathcal{T}_Y$. Finally, we see that

$$\left(\bigcap_{\alpha \in I} U_\alpha \right) \cap \left(\bigcap_{\beta \in J} (Y \setminus K_\beta) \right) = U \cap (Y \setminus K).$$

However, as K is compact and $(X, \mathcal{T}_X)$ is Hausdorff, Theorem 3.1.13 implies K is closed. Hence $U \cap (Y \setminus K) = U \cap (X \setminus K) \in \mathcal{T}_X$ as $\mathcal{T}_X$ is closed under finite intersections. Hence $\mathcal{T}_Y$ is closed under finite unions and thus $\mathcal{T}_Y$ is a topology on Y.

We claim $(Y, \mathcal{T}_Y)$ is the topological space we are looking for. First, to see that $(X, \mathcal{T}_X)$ is a subspace of $(Y, \mathcal{T}_Y)$, we must show that

$$\mathcal{T}_X = \{X \cap V \mid V \in \mathcal{T}_Y\}.$$

Indeed if $V \in \mathcal{T}_X \subseteq \mathcal{T}_Y$, then $X \cap V = V \in \mathcal{T}_X$. Furthermore, if $K \subseteq X$ is compact then K is closed in $(X, \mathcal{T}_X)$ by Theorem 3.1.13 so $X \cap (Y \setminus K) = X \setminus K \in \mathcal{T}_X$. Hence the equality of the topologies follows so $(X, \mathcal{T}_X)$ is a subspace of $(Y, \mathcal{T}_Y)$.

To see that $(Y, \mathcal{T}_Y)$ is Hausdorff, let $y_1, y_2 \in Y$ be arbitrary points such that $y_1 \neq y_2$. We desire to find $U, V \in \mathcal{T}_Y$ such that $y_1 \in U$, $y_2 \in V$, and $U \cap V = \emptyset$. If $y_1, y_2 \in X$, then the fact that $(X, \mathcal{T}_X)$ is Hausdorff implies there exist $U, V \in \mathcal{T}_X \subseteq \mathcal{T}_Y$ with the desired properties. Thus we may assume that $y_1 = \infty$ or $y_2 = \infty$. By symmetry we may assume that $y_2 = \infty$ without loss of generality. Since $(X, \mathcal{T}_X)$ is locally compact, there exists a compact subspace K of $(X, \mathcal{T}_X)$ and a $U \in \mathcal{T}_X \subseteq \mathcal{T}_Y$ such that $y_1 \in U \subseteq K$. Hence if $V = Y \setminus K \in \mathcal{T}_Y$, then $y_2 \in V$ and $U \cap V = \emptyset$ as desired. Therefore, as $y_1, y_2 \in Y$ were arbitrary, $(Y, \mathcal{T}_Y)$ is Hausdorff.

Finally, to see that $(Y, \mathcal{T}_Y)$ is compact, let $\{V_\alpha\}_{\alpha \in I}$ be an arbitrary open cover of $(Y, \mathcal{T}_Y)$. Since $\infty \in Y$ and every element of $\mathcal{T}_X$ does not contain ∞, there exists an $\alpha_0 \in I$ such that $\infty \in V_{\alpha_0} = Y \setminus K$ for some compact subset K of $(X, \mathcal{T}_X)$.

We claim that $\{X \cap V_\alpha\}_{\alpha \in I \setminus \{\alpha_0\}}$ is an open cover of K. To see this, recall that since $(X, \mathcal{T}_X)$ is a subspace of $(Y, \mathcal{T}_Y)$ that $X \cap V_\alpha \in \mathcal{T}_X$ for all $\alpha \in I$. Furthermore, as $\{V_\alpha\}_{\alpha \in I}$ is an open cover of $(Y, \mathcal{T}_Y)$ and $\infty \in V_{\alpha_0} = Y \setminus K$, it must be the case that

$$K \subseteq \bigcup_{\alpha \in I \setminus \{\alpha_0\}} X \cap V_\alpha.$$

Thus $\{X \cap V_\alpha\}_{\alpha \in I \setminus \{\alpha_0\}}$ is an open cover of K.

Since K is compact, there exists a finite subset $J \subseteq I \setminus \{\alpha_0\}$ such that

$$K \subseteq \bigcup_{\alpha \in J} X \cap V_\alpha.$$

Hence, as $V_{\alpha_0} = Y \setminus K$, $\{V_\alpha\}_{\alpha \in J \cup \{\alpha_0\}}$ is a finite subcover of $(Y, \mathcal{T}_Y)$ from $\{V_\alpha\}_{\alpha \in I}$. Therefore, as $\{V_\alpha\}_{\alpha \in I}$ was arbitrary, $(Y, \mathcal{T}_Y)$ is compact as desired.

Finally, to see the uniqueness, suppose $(Y, \mathcal{T}_Y)$ and $(Z, \mathcal{T}_Z)$ are compact Hausdorff topological spaces such that $(X, \mathcal{T}_X)$ is a subspace of $(Y, \mathcal{T}_Y)$ and $(Z, \mathcal{T}_Z)$, $Y \setminus X = \{\infty_Y\}$, and $Z \setminus X = \{\infty_Z\}$. Define $f : Y \to Z$ by that $f(x) = x$ for all $x \in X$ and $f(\infty_Y) = \infty_Z$. Clearly f is bijective. To see that f is a homeomorphism, we claim that $f(U) \in \mathcal{T}_Z$ for all $U \in \mathcal{T}_Y$. Symmetry will then imply that $f^{-1}(V) \in \mathcal{T}_Y$ for all $V \in \mathcal{T}_Z$ thereby proving that $(Y, \mathcal{T}_Y)$ and $(Z, \mathcal{T}_Z)$ are homeomorphic.

To see the claim, let $U \in \mathcal{T}_Y$ be arbitrary. If $U \subseteq X$, then U is open in $(X, \mathcal{T}_X)$ as $(X, \mathcal{T}_X)$ is a subspace of $(Y, \mathcal{T}_Y)$. Hence $f(U) = U$ is open in $(X, \mathcal{T}_X)$ as a subspace of $(Z, \mathcal{T}_Z)$. Thus there exists a $V \in \mathcal{T}_Z$ such that $U = X \cap V$. However, since $(Z, \mathcal{T}_Z)$ is Hausdorff, $\{\infty_Z\}$ is closed in $(Z, \mathcal{T}_Z)$ by Example 1.6.8 and thus $X = Z \setminus \{\infty_Z\}$ is open in $(Z, \mathcal{T}_Z)$.

Hence $U = X \cap V \in \mathcal{T}_Z$ as desired. Thus, to complete the proof, we may assume that $\infty_Y \in U$. However, this implies that $C = Y \setminus U \subseteq X$ is closed in $(Y, \mathcal{T}_Y)$ and thus compact as $(Y, \mathcal{T}_Y)$ is compact by Theorem 3.1.14. Therefore, since $(X, \mathcal{T}_X)$ is a subspace of $(Y, \mathcal{T}_Y)$ and $C \subseteq X$, C is a compact subset of $(X, \mathcal{T}_X)$. Hence $f(C) = C$ is compact in $(X, \mathcal{T}_X)$ as a subspace of $(Z, \mathcal{T}_Z)$. Therefore, since $(Z, \mathcal{T}_Z)$ is Hausdorff, C is closed in $(Z, \mathcal{T}_Z)$. Thus $Z \setminus C = f(Y) \setminus f(C) = f(Y \setminus C) = f(U)$ is open in $(Z, \mathcal{T}_Z)$ as desired. Thus we have completed the uniqueness portion of the proof and thus the proof. ∎

Due to the uniqueness in Theorem 3.4.7, it makes sense to give these objects a name.

Definition 3.4.8. Let $(X, \mathcal{T}_X)$ be a locally compact, Hausdorff topological space. The unique (up to homeomorphism) compact Hausdorff topological space $(Y, \mathcal{T}_Y)$ from Theorem 3.4.7 is called the *one-point compactification of* $(X, \mathcal{T}_X)$.

Example 3.4.9. Consider $X = (-\pi, \pi)$ equipped with the subspace topology inherited from the canonical topology on $\mathbb{R}$. Clearly X is locally compact as for each $x \in (-\pi, \pi)$ there exists an $\epsilon > 0$ so that

$$x \in (x - \epsilon, x + \epsilon) \subseteq [x - \epsilon, x + \epsilon] \subseteq X.$$

As $[x - \epsilon, x + \epsilon]$ is compact in $\mathbb{R}$ by the Heine-Borel Theorem (Theorem 3.1.25), $[x - \epsilon, x + \epsilon]$ is compact. Hence X is locally compact.

Let

$$S^1 = \{(x, y) \in \mathbb{R}^2 \mid x^2 + y^2 = 1\}$$

equipped with the subspace topology inherited from the Euclidean topology on $\mathbb{R}^2$. Then S^1 is Hausdorff as $\mathbb{R}^2$ is Hausdorff (see Example 1.5.38) and is a compact subset of $\mathbb{R}^2$ by the Heine-Borel Theorem. Since the function $f : X \to S^1$ defined by

$$f(x) = (\cos(x), \sin(x))$$

for all $x \in X$ is an embedding of X, S^1 is the one-point compactification of X by definition. Hence, since X and $\mathbb{R}$ are homeomorphic by the function $g : X \to \mathbb{R}$ defined by

$$g(x) = \tan\left(\frac{x}{2}\right),$$

we see that the one-point compactification of $\mathbb{R}$ is S^1.

Example 3.4.10. By Example 2.2.6, we see that the one-point compactification of $\mathbb{R}^2$ equipped with its Euclidean topology is

$$S^2 = \{(x, y, z) \in \mathbb{R}^3 \mid x^2 + y^2 + z^2 = 1\}$$

equipped with the subspace topology inherited from the Euclidean topology on $\mathbb{R}^3$.

Using one-point compactifications, we can obtain a stronger characterization of locally compact topological spaces, which of course is more useful for applications, but often more difficult to verify by hand. However, this characterization only works for Hausdorff spaces.

Theorem 3.4.11. *Let $(X, \mathcal{T})$ be a Hausdorff space. Then $(X, \mathcal{T})$ is locally compact if and only if for every $x \in X$ and neighbourhood U of x there exists a neighbourhood U_0 of x such that $\overline{U_0}$ is a compact subspace of $(X, \mathcal{T})$ and $\overline{U_0} \subseteq U$.*

Proof. Suppose $(X, \mathcal{T})$ is locally compact and let $(Y, \mathcal{T}_Y)$ be the one-point compactification of $(X, \mathcal{T}_X)$ with $Y \setminus X = \{\infty\}$. To see the result, let $x \in X$ and $U \in \mathcal{T}_X$ such that $x \in U$ be arbitrary. By our description of the topology $\mathcal{T}_Y$ from Theorem 3.4.7, we know that $U \in \mathcal{T}_Y$. Hence $K = Y \setminus U$ is a closed subset of $(Y, \mathcal{T}_Y)$ and thus compact by Theorem 3.1.14 as $(Y, \mathcal{T}_Y)$ is compact.

Since $x \notin K$, K is a compact subspace of $(Y, \mathcal{T}_Y)$, and $(Y, \mathcal{T}_Y)$ is Hausdorff, Lemma 3.1.12 implies there exists $U_0, V_0 \in \mathcal{T}_Y$ such $x \in U_0$, $K \subseteq V_0$, and $U_0 \cap V_0 = \emptyset$. To complete the proof of this direction, it suffices to show that $U_0 \in \mathcal{T}_X$, $\overline{U_0}$ is compact (where the closure is in $(X, \mathcal{T}_X)$), and $\overline{U_0} \subseteq U$.

As $\infty \in K \subseteq V_0$ and $U_0 \cap V_0 = \emptyset$, $U_0 \subseteq X$ and thus $U_0 \in \mathcal{T}_Y$ implies $U_0 \in \mathcal{T}_X$ as $(X, \mathcal{T}_X)$ is a subspace of $(Y, \mathcal{T}_Y)$.

Let $K_0 = Y \setminus V_0$. Then K_0 is a closed subset of $(Y, \mathcal{T}_Y)$ and thus compact by Theorem 3.1.14. Moreover, since $\infty \in V_0$ and $U_0 \cap V_0 = \emptyset$, $U_0 \subseteq K_0 \subseteq X$. Hence, as $(X, \mathcal{T}_X)$ is a subspace of $(Y, \mathcal{T}_Y)$, K_0 is a compact subspace of $(X, \mathcal{T}_X)$ and thus closed by Theorem 3.1.13 as $(X, \mathcal{T}_X)$ is Hausdorff. Consequently $\overline{U_0} \subseteq K_0$ so $\overline{U_0}$ is a closed subset of a compact set in $(X, \mathcal{T}_X)$ and thus compact by Theorem 3.1.14.

Finally, as $\overline{U_0} \subseteq K_0$, $K_0 = Y \setminus V_0$, $K \subseteq V_0$, and $K = Y \setminus U$, we obtain that $\overline{U_0} \subseteq U$. thereby completing the proof of this direction since x and U were arbitrary.

To see that the converse holds, let $x \in X$ be arbitrary. By taking $U = X$ there exists a neighbourhood U_0 of x such that $K = \overline{U_0}$ is compact. As $x \in U_0 \subseteq K$, we obtain that $(X, \mathcal{T})$ is locally compact at x. Therefore, as x was arbitrary, $(X, \mathcal{T})$ is locally compact as desired. ∎

Using Theorem 3.4.11, we obtain the 'open set' version of Proposition 3.4.6 for Hausdorff spaces. Note the following also shows that removing a single point from a compact Hausdorff space yields a locally compact Hausdorff space as every singleton is closed in a Hausdorff space and thus the complements of singletons are open.

Corollary 3.4.12. *Let $(X, \mathcal{T})$ be a locally compact Hausdorff topological space and let Y be a subspace of $(X, \mathcal{T})$. If Y is open in $(X, \mathcal{T})$, then Y is locally compact.*

Proof. To see that Y is locally compact, let $y \in Y$ be arbitrary. Since Y is a neighbourhood of y in $(X, \mathcal{T})$ and $(X, \mathcal{T})$ is locally compact, Theorem 3.4.11 implies there exists a $U \in \mathcal{T}$ such that $\overline{U}$ is compact and $\overline{U} \subseteq Y$. Thus U is an open set in Y, $\overline{U}$ is a compact set in Y, and $y \in U \subseteq \overline{U}$. Therefore, since $y \in Y$ was arbitrary, Definition 3.4.1 implies that Y is locally compact. ∎

Using Corollary 3.4.12 along with the one-point compactification, we can obtain another characterization of locally compact Hausdorff topological spaces.

Corollary 3.4.13. *A Hausdorff topological space $(X, \mathcal{T})$ is a locally compact if and only if $(X, \mathcal{T})$ is homeomorphic to an open subspace of a compact Hausdorff topological space.*

Proof. First, suppose that $(X, \mathcal{T})$ is locally compact. Let $(Y, \mathcal{T}_Y)$ be the one-point compactification of $(X, \mathcal{T})$ with $Y \setminus X = \{\infty\}$. As $(Y, \mathcal{T}_Y)$ is Hausdorff, $\{\infty\} = Y \setminus X$ is closed by Example 1.6.8 so X is an open subset of $(Y, \mathcal{T}_Y)$. Hence, as $(Y, \mathcal{T}_Y)$ is a compact Hausdorff topological space, this direction of the result follows.

Conversely, suppose $(X, \mathcal{T})$ is homeomorphic to an open subspace of a compact Hausdorff topological space. since every compact space is locally compact and since Corollary 3.4.12 implies open subspaces of locally compact Hausdorff topological spaces are locally compact, $(X, \mathcal{T})$ is homeomorphic to a locally compact space and thus is locally compact by Theorem 3.4.2. ∎

Moreover, we can completely characterize the locally compact subsets of a locally compact Hausdorff topological space. Before we prove this result, we need a quick lemma.

Lemma 3.4.14. *Let $(X, \mathcal{T})$ be a topological space, let Y be a subspace of X, and let $V \in \mathcal{T}$. Then $\overline{V \cap \overline{Y}} = \overline{V \cap Y}$ where the closures are computed in $(X, \mathcal{T})$.*

Proof. Clearly $V \cap Y \subseteq V \cap \overline{Y}$ implies $\overline{V \cap Y} \subseteq \overline{V \cap \overline{Y}}$ by definition. To see the reverse inclusion, let $x_0 \in \overline{V \cap \overline{Y}}$. Then for every $U \in \mathcal{T}$ such that $x_0 \in U$ there exists an $x \in V \cap \overline{Y}$ such that $x \in U$. Hence $U \cap V$ is a $\mathcal{T}$-neighbourhood of x and $x \in \overline{Y}$ there must exist a $y \in Y$ such that $y \in U \cap V$. Hence $y \in U$ and $y \in V \cap Y$. Therefore, as U was arbitrary, $x_0 \in \overline{V \cap Y}$. Hence for all $V \in \mathcal{T}$ we have that $\overline{V \cap \overline{Y}} \subseteq \overline{V \cap Y}$ where the closures are computed in $(X, \mathcal{T})$. ∎

Corollary 3.4.15. *Let $(X, \mathcal{T})$ be a locally compact, Hausdorff topological space and let Y be a subspace of $(X, \mathcal{T})$. Then Y is locally compact if and only if $Y = U \cap F$ where U is an open subset of $(X, \mathcal{T})$ and F is a closed subset of $(X, \mathcal{T})$.*

Proof. First, suppose Y is a subspace of X such that $Y = U \cap F$ where U is an open subset of $(X, \mathcal{T})$ and F is a closed subset of $(X, \mathcal{T})$. To see that Y is locally compact, by Proposition 3.4.6 and Corollary 3.4.12 we know that F and U are locally compact subspaces of $(X, \mathcal{T})$. To see that $Y = U \cap F$ is locally compact, let $x \in U \cap F$ be arbitrary. Since U and F are locally compact, there exist $V_1, V_2 \in \mathcal{T}$, $K_1 \subseteq U$, and $K_2 \subseteq F$ such that $x \in V_1 \cap U \subseteq K_1$, $x \in V_2 \cap F \subseteq K_2$, K_1 is a compact subspace of U, and K_2 is a compact subspace of F. Let $V = V_1 \cap V_2 \in \mathcal{T}$ and let $K = K_1 \cap K_2$. Clearly $V \cap (U \cap F)$ is an open set in $U \cap F$, $K \subseteq U \cap F$, and $x \in V \cap (U \cap F) \subseteq K_1 \cap K_2$. Thus it remains only to show that K is a compact subspace of $U \cap F$.

To see that K is compact, note since K_1 and K_2 are compact subspaces of U and F respectively, they are compact subspaces of $(X, \mathcal{T})$. Therefore, since $(X, \mathcal{T})$ is Hausdorff, K_1 and K_2 are closed subsets of $(X, \mathcal{T})$. Hence $K = K_1 \cap K_2$ is a closed subset of $(X, \mathcal{T})$. However, $K \subseteq K_1$ so K is a closed subset of the compact subspace K_1 and thus compact as desired. Hence one direction of the result is complete.

Conversely, suppose Y be a locally compact subspace of $(X, \mathcal{T})$. Notice that Y is a Hausdorff space since Y is a subspace of $(X, \mathcal{T})$ and since $(X, \mathcal{T})$ is Hausdorff. Since Y is locally compact Hausdorff subspace of $(X, \mathcal{T})$, for each $y \in Y$ there exists a $U_y \in \mathcal{T}$ such that $y \in U_y \cap Y \subseteq \overline{U_y \cap Y}$ where the closure is taken in Y, and $\overline{U_y \cap Y}$ is a compact subspace of $(X, \mathcal{T})$. Therefore, since $(X, \mathcal{T})$ is Hausdorff, $\overline{U_y \cap Y}$ is a closed subset of $(X, \mathcal{T})$. Furthermore, by the properties of the closure in a subspace, we know that $\overline{U_y \cap Y} = \overline{U_y \cap Y} \cap Y$ where the first closure is computed in Y and the second is computed in $(X, \mathcal{T})$. Hence $\overline{U_y \cap Y} \cap Y$ where $\overline{U_y \cap Y}$ is computed in $(X, \mathcal{T}_Y)$ is a closed subset of $(X, \mathcal{T}_Y)$ containing $U_y \cap Y$. Hence for all $y \in Y$, $\overline{U_y \cap Y} \subseteq \overline{U_y \cap Y} \cap Y$ where both closures are computed in $(X, \mathcal{T})$.

Let $F = \overline{Y}$, which is a closed subset of $(X, \mathcal{T})$, and let

$$U = \bigcup_{y \in Y} U_y,$$

which is clearly an open subset of $(X, \mathcal{T})$. We claim that $Y = U \cap F$. To see this, we first note that $Y \subseteq U \cap F$ since $Y \subseteq U$ and $Y \subseteq F$ by construction. For the converse, notice for all $y \in Y$ that

$$U_y \cap \overline{Y} \subseteq \overline{U_y \cap \overline{Y}} \subseteq \overline{U_y \cap Y} \subseteq \overline{U_y \cap Y} \cap Y \subseteq Y$$

where all of the closures are computed in $(X, \mathcal{T})$, the first inclusion is trivial, the second inclusion was demonstrated above, the third inclusion follows as $\overline{U_y} \cap Y$ is closed in $(X, \mathcal{T})$, and the fourth inclusion is trivial. Hence, as this holds for all $y \in Y$, we obtain that $U \cap F \subseteq Y$. Hence $Y = U \cap F$ as desired. ∎

As all of the above demonstrates, locally compact Hausdorff topological spaces are nice since they can be studied inside compact Hausdorff topological spaces and, as compact topological spaces are very nice, many properties can be obtained. This raises the question of which topological spaces can be embedded into compact topological spaces. One answer to this question will be studied in Section 5.4. For now, we turn our attention to studying compactness in metric spaces.

Chapter 4

Compact Metric Spaces

With Chapter 3 complete, we have a basic understanding of compact topological spaces. However, a far deeper study persists if we further restrict our topological spaces. In particular, this chapter will be devoted to studying compact metric spaces where a far richer and deeper theory occurs.

We begin this chapter by determining a further characterization of compact metric spaces by generalizing the Heine-Borel Theorem. This requires correcting the notions of 'closed' and 'bounded' to a general metric space. This will also lead us to a sequential characterization of compactness in metric spaces as the metric permits the use of sequences instead of nets.

This study then leads us to examining compactness in functions spaces. In particular, using the Extreme Value Theorem (Theorem 3.1.27) we can place a supremum norm on the continuous functions from a compact topological space into the real numbers. Thus we may ask for conditions to determine when collections of functions from the spaces form compact subspaces. This also leads to the question of which collections of functions can be used to approximate other such functions on these spaces.

4.1 Complete Metric Spaces

In this section, we analyze the correct generalization of being a 'closed set' to correctly generalize the Heine-Borel Theorem (Theorem 3.1.25) to arbitrary metric spaces. As the notion of compact topological spaces is connected to the notion of nets having convergent subnets, and since metric spaces are characterized by their convergent sequences (see Theorem 1.5.3), perhaps it is not surprising that we desired a similar notion to that of a closed set to aid in verifying sequences converge. Thus we begin by analyzing a particular type of sequence which one would expect to converge as the terms of the sequence are eventually not too far away from one another.

Definition 4.1.1. Let (X, d) be a metric space. A sequence $(x_n)_{n\geq 1}$ in X is said to be *Cauchy* if for all $\epsilon > 0$ there exists an $N \in \mathbb{N}$ such that

$d(x_n, x_m) < \epsilon$ for all $n, m \geq N$.

Remark 4.1.2. It is important to note that there exists sequences $(x_n)_{n\geq 1}$ such that

$$\lim_{n\to\infty} d(x_n, x_{n+1}) = 0$$

that are not Cauchy. Indeed consider $\mathbb{R}$ with the canonical metric and let $x_n = \sum_{k=1}^{n} \frac{1}{k}$ for all $n \in \mathbb{N}$. Clearly $d(x_n, x_{n+1}) = \frac{1}{n+1}$ which tends to 0 as n tends to infinity, yet $(x_n)_{n\geq 1}$ is not Cauchy as for all $m \in \mathbb{N}$

$$\sup_{m\to\infty} d(x_n, x_m) = \sup_{m\to\infty} \sum_{k=n}^{m} \frac{1}{k} = \infty.$$

In terms of examples, lots of sequences are Cauchy.

Lemma 4.1.3. *Every convergent sequence in a metric space is Cauchy.*

Proof. Let $(x_n)_{n\geq 1}$ be a convergent sequence in a metric space (X, d). Let $x_0 = \lim_{n\to\infty} x_n$. To see that $(x_n)_{n\geq 1}$ is Cauchy, let $\epsilon > 0$ be arbitrary. Since $x_0 = \lim_{n\to\infty} x_n$, there exists an $N \in \mathbb{N}$ such that $d(x_n, x_0) < \frac{\epsilon}{2}$ for all $n \geq N$. Therefore, for all $n, m \geq N$,

$$d(x_n, x_m) \leq d(x_n, x_0) + d(x_0, x_m) < \frac{\epsilon}{2} + \frac{\epsilon}{2} = \epsilon.$$

Thus, as $\epsilon > 0$ was arbitrary, $(x_n)_{n\geq 1}$ is Cauchy by definition. ∎

Of course, there are some obvious restrictions for a sequence to be Cauchy.

Lemma 4.1.4. *Every Cauchy sequence in a metric space is bounded.*

Proof. Let $(x_n)_{n\geq 1}$ be a Cauchy sequence in a metric space (X, d). Since $(x_n)_{n\geq 1}$ is Cauchy, there exists an $N \in \mathbb{N}$ such that $d(x_n, x_m) < 1$ for all $n, m \geq N$.

Let $M = \max\{d(x_1, x_N), \ldots, d(x_{N-1}, x_N), 1\}$. Using the above paragraph, we see that $d(x_n, x_N) \leq M$ for all $n \in \mathbb{N}$. Hence $(x_n)_{n\geq 1}$ is bounded. ∎

Of course, it would be nice if the converse Lemma 4.1.3 were true as this would enable us to deduce the convergence of a sequence by checking it is Cauchy without any knowledge of the limit. Thus we make the following definition.

Definition 4.1.5. A metric space (X, d) is said to be *complete* if every Cauchy sequence converges.

As most readers of these notes should already know, $\mathbb{R}$ equipped with the canonical metric is a complete metric space. We quickly remind the reader of the ingredients of the proof of this fact.

Lemma 4.1.6. *Let $(x_n)_{n\geq 1}$ be a Cauchy sequence in a metric space (X, d). If a subsequence of $(x_n)_{n\geq 1}$ converges, then $(x_n)_{n\geq 1}$ converges.*

Proof. Let $(x_n)_{n\geq 1}$ be a Cauchy sequence with a convergent subsequence $(x_{k_n})_{n\geq 1}$ and let $x_0 = \lim_{n\to\infty} x_{k_n}$. We claim that $\lim_{n\to\infty} x_n = x_0$. To see this, let $\epsilon > 0$ be arbitrary. Since $(x_n)_{n\geq 1}$ is Cauchy, there exists an $N \in \mathbb{N}$ such that $d(x_n, x_m) < \frac{\epsilon}{2}$ for all $n, m \geq N$. Furthermore, since $x_0 = \lim_{n\to\infty} x_{k_n}$, there exists an $k_j \geq N$ such that $d(x_{k_j}, x_0) < \frac{\epsilon}{2}$. Hence, if $n \geq N$ then

$$d(x_n, x_0) \leq d(x_n, x_{k_j}) + d(x_{k_j}, x_0) < \frac{\epsilon}{2} + \frac{\epsilon}{2} = \epsilon.$$

Thus, as $\epsilon > 0$ was arbitrary, $(x_n)_{n\geq 1}$ is converges to x_0 by definition. ∎

In addition, recall the following theorem.

Theorem 4.1.7 (Bolzano-Weierstrass Theorem). *Every bounded sequence of real numbers has a convergent sequence.*

Proof. The least upper bound property of $\mathbb{R}$ implies the Monotone Convergence Theorem for $\mathbb{R}$ which states that every bounded monotone sequence in $\mathbb{R}$ converges. As every sequence in $\mathbb{R}$ can be shown to have a monotone subsequence, the result follows. ∎

Consequently, we have our first example of a metric spaces that is complete but not compact.

Theorem 4.1.8 (Completeness of the Real Numbers). *The real numbers are complete with their canonical metric.*

Proof. Let $(x_n)_{n\geq 1}$ be a Cauchy sequence of real numbers. Thus $(x_n)_{n\geq 1}$ is bounded by Lemma 4.1.4. Therefore $(x_n)_{n\geq 1}$ has a convergent sequence by the Bolzano-Weierstrass Theorem (Theorem 4.1.7). Hence $(x_n)_{n\geq 1}$ converges by Lemma 4.1.6. ∎

Corollary 4.1.9. *For every $p \in [1, \infty]$ and $n \in \mathbb{N}$, $(\mathbb{K}^n, \|\cdot\|_p)$ is complete.*

Proof. To see that $(\mathbb{R}^n, \|\cdot\|_p)$ is complete, let $(\vec{x}_k)_{k\geq 1}$ be an arbitrary Cauchy sequence in $(\mathbb{R}^n, \|\cdot\|_p)$. Write $\vec{x}_k = (x_{k,1}, \ldots, x_{k,n})$. Since for all $k, m \in \mathbb{N}$ we have

$$|x_{k,j} - x_{m,j}| \leq \|\vec{x}_k - \vec{x}_m\|_p\,,$$

it is elementary to see that $(x_{k,j})_{k\geq 1}$ is a Cauchy sequence in $\mathbb{R}$ for all $j \in \{1, \ldots, n\}$. Since $\mathbb{R}$ is complete, for each $j \in \{1, \ldots, n\}$ there exists an $x_j \in \mathbb{R}$ such that $x_j = \lim_{k\to\infty} x_{k,j}$. If $\vec{x} = (x_1, \ldots, x_n)$, then $\vec{x} = \lim_{k\to\infty} \vec{x}_k$ in $(\mathbb{R}^n, \|\cdot\|_p)$ by elementary properties of the norm. Therefore, as $(\vec{x}_k)_{k\geq 1}$ was arbitrary, $(\mathbb{R}^n, \|\cdot\|_p)$ is complete.

To see that $(\mathbb{C}^n, \|\cdot\|_p)$, it suffices by the same arguments to show that $(\mathbb{C}, |\cdot|)$ is complete. To see that $(\mathbb{C}, |\cdot|)$ is complete, let $(z_k)_{k\geq 1}$ be an arbitrary Cauchy sequence in $\mathbb{C}$. For each k, write $z_k = a_k + ib_k$ where $a_k, b_k \in \mathbb{R}$. Since for all $k, m \in \mathbb{N}$ we have

$$|a_k - a_m|, |b_k - b_m| \leq |z_k - z_m|,$$

it is elementary to see that $(a_k)_{k\geq 1}$ and $(b_k)_{k\geq 1}$ are Cauchy sequences in $\mathbb{R}$. Since $\mathbb{R}$ is complete, $a = \lim_{k\to\infty} a_k$ and $b = \lim_{k\to\infty} b_k$ exist. Hence $z = a + bi$, then $z = \lim_{k\to\infty} z_k$ by elementary properties of the norm. Hence, as $(z_k)_{k\geq 1}$ was arbitrary, $(\mathbb{C}, |\cdot|)$ is complete. ∎

For the sake of completeness, we quickly present a simple example of a metric space that is not complete. In particular $\mathbb{R}$ equipped with another metric need not be complete even if the metric induces the canonical topology on $\mathbb{R}$. That is, completeness is a metric space property; not a topological property.

Example 4.1.10. Define $d : \mathbb{R} \times \mathbb{R} \to [0, \infty)$ by

$$d(x, y) = \left|e^{-x} - e^{-y}\right|$$

for all $x, y \in \mathbb{R}$. Since the exponential function is injective and since the absolute value is a metric on $\mathbb{R}$, it is elementary to see that d is a metric on $\mathbb{R}$.

We claim that $(\mathbb{R}, d)$ is not complete. To see this, consider the sequence of natural numbers $(n)_{n\geq 1}$. We claim that $(n)_{n\geq 1}$ is Cauchy in $(\mathbb{R}, d)$. To see this, let $\epsilon > 0$. Since $\lim_{n\to\infty} e^{-n} = 0$, there exists an $N \in \mathbb{N}$ such that $0 < e^{-n} < \frac{\epsilon}{2}$ for all $n \geq N$. Hence for all $n, m \geq N$ we have that

$$d(n, m) = \left|e^{-n} - e^{-m}\right| \leq e^{-n} + e^{-m} < \frac{\epsilon}{2} + \frac{\epsilon}{2} = \epsilon.$$

Therefore, as $\epsilon > 0$ was arbitrary, $(n)_{n\geq 1}$ is Cauchy in $(\mathbb{R}, d)$.

However, we claim that $(n)_{n\geq 1}$ does not converge in $(\mathbb{R}, d)$. To see this, we note that a sequence $(x_n)_{n\geq 1}$ converges in $(\mathbb{R}, d)$ to $x_0 \in \mathbb{R}$ if and only if $e^{-x_0} = \lim_{n\to\infty} e^{-x_n}$ in the canonical topology on $\mathbb{R}$ if and only if $x_0 = \lim_{n\to\infty} x_n$ in the canonical topology on $\mathbb{R}$. Therefore, as $(n)_{n\geq 1}$ does not converge in the canonical topology on $\mathbb{R}$, $(n)_{n\geq 1}$ does not converge in $(\mathbb{R}, d)$ as claimed. Hence $(\mathbb{R}, d)$ is not complete even though the topology on $\mathbb{R}$ is the canonical topology. Thus the metric is important when it comes to completeness!

Of course, $\mathbb{R}$ equipped with its canonical metric is not a compact metric space. Thus the notions of compactness and completeness for a metric space need not agree. However, completeness will be a necessary ingredient for a metric space to be compact. Although we can prove this directly, we take a slightly indirect approach to develop a tool needed much later in these notes.

Theorem 4.1.11 (Cantor's Theorem). *Let (X,d) be a metric space. Then the following are equivalent:*

(i) *(X,d) is a complete metric space.*

(ii) *If $(F_n)_{n\geq 1}$ is a sequence of non-empty closed subsets of X such that $F_{n+1} \subseteq F_n$ for all $n \in \mathbb{N}$ and $\lim_{n\to\infty} \operatorname{diam}(F_n) = 0$, then $\bigcap_{n=1}^{\infty} F_n \neq \emptyset$.*

Proof. First let (X,d) is a complete metric space. Suppose $(F_n)_{n\geq 1}$ is an arbitrary sequence of non-empty closed subsets of X such that $F_{n+1} \subseteq F_n$ for all $n \in \mathbb{N}$ and $\lim_{n\to\infty} \operatorname{diam}(F_n) = 0$. To see that $\bigcap_{n=1}^{\infty} F_n \neq \emptyset$, for each $n \in \mathbb{N}$ choose $x_n \in F_n$. We claim that $(x_n)_{n\geq 1}$ is a Cauchy sequence. To see this, let $\epsilon > 0$ be arbitrary. Since $\lim_{n\to\infty} \operatorname{diam}(F_n) = 0$, there exists an $N \in \mathbb{N}$ such that $\operatorname{diam}(F_N) < \epsilon$. As $F_{n+1} \subseteq F_n$ for all $n \in \mathbb{N}$, we obtain that $x_n \in F_N$ for all $n \geq N$. Hence $d(x_n, x_m) \leq \operatorname{diam}(F_N) < \epsilon$ for all $n, m \geq N$. Hence, as $\epsilon > 0$ was arbitrary, $(x_n)_{n\geq 1}$ is a Cauchy sequence.

Since (X,d) is complete, $x = \lim_{n\to\infty} x_n$ exists. Since for each $m \in \mathbb{N}$ we have $x_n \in F_m$ for all $n \geq m$, we obtain from Theorem 1.6.14 together with the fact that F_m is closed that $x \in F_m$ for all $m \in \mathbb{N}$. Hence $x \in \bigcap_{n=1}^{\infty} F_n$ so $\bigcap_{n=1}^{\infty} F_n \neq \emptyset$.

For the converse direction, let (X,d) has property (ii). To see that (X,d) is complete, suppose $(x_n)_{n\geq 1}$ is an arbitrary Cauchy sequence. For each $n \in \mathbb{N}$, let

$$F_n = \overline{\{x_k \mid k \geq n\}}.$$

Clearly each F_n is a non-empty closed subset of X such that $F_{n+1} \subseteq F_n$ for all $n \in \mathbb{N}$.

We claim that $\lim_{n\to\infty} \operatorname{diam}(F_n) = 0$. To see this, let $\epsilon > 0$ be arbitrary. Since $(x_n)_{n\geq 1}$ is Cauchy, there exists an $N \in \mathbb{N}$ such that $d(x_n, x_m) < \frac{\epsilon}{3}$ for all $n, m \geq N$. We claim that $\operatorname{diam}(F_n) \leq \epsilon$ whenever $n \geq N$. To see this, fix $n \geq N$ and let $x, y \in F_n$ be arbitrary. By the definition of F_n, there exists $k, j \geq n \geq N$ such that

$$d(x, x_j) < \frac{\epsilon}{3} \qquad \text{and} \qquad d(y, x_k) < \frac{\epsilon}{3}.$$

Hence

$$d(x,y) \leq d(x, x_j) + d(x_j, x_k) + d(x_k, y) < \frac{\epsilon}{3} + \frac{\epsilon}{3} + \frac{\epsilon}{3} = \epsilon$$

as $k, j \geq N$ and by our choice of N. Hence $\operatorname{diam}(F_n) \leq \epsilon$ whenever $n \geq N$ by the definition of the diameter of a set. Thus the claim is complete.

By the assumption of property (ii), the above implies that $\bigcap_{n=1}^{\infty} F_n \neq \emptyset$. Let $x \in \bigcap_{n=1}^{\infty} F_n$. We claim that $(x_n)_{n\geq 1}$ converges to x. To see this, let $\epsilon > 0$ be arbitrary. Since $\lim_{n\to\infty} \operatorname{diam}(F_n) = 0$, there exists an $N \in \mathbb{N}$ such that $\operatorname{diam}(F_n) < \epsilon$ for all $n \geq N$. Since $x, x_n \in F_n$ for all $n \in \mathbb{N}$, we obtain that

$$d(x, x_n) \leq \operatorname{diam}(F_n) < \epsilon$$

for all $n \geq N$. Therefore, as $\epsilon > 0$ was arbitrary, $x = \lim_{n\to\infty} x_n$. Hence, as $(x_n)_{n\geq 1}$ was an arbitrary Cauchy sequence, (X, d) is complete. ∎

Theorem 4.1.12. *Every compact metric spaces is complete.*

Proof. Let (X, d) be a compact metric spaces. To see that (X, d) is complete, let $(F_n)_{n\geq 1}$ be a sequence of non-empty closed subsets of X such that $F_{n+1} \subseteq F_n$ for all $n \in \mathbb{N}$ and $\lim_{n\to\infty} \mathrm{diam}(F_n) = 0$. Clearly $\{F_n\}_{n\geq 1}$ has the finite intersection property as $F_{n+1} \subseteq F_n$ for all $n \in \mathbb{N}$. Hence Theorem 3.2.2 implies that $\bigcap_{n=1}^{\infty} F_n \neq \emptyset$. Thus Cantor's Theorem (Theorem 4.1.11) implies that (X, d) is complete. ∎

Of course, not every complete metric spaces is compact as $\mathbb{R}$ with its canonical metric is complete but not compact. However, it is possible to construct new examples of complete metric spaces from other examples in a similar fashion to how we constructed new example of compact topological spaces from other examples. This also provides us with additional examples of metric spaces that are not complete.

Before we get to this result, we note that the proof of Theorem 1.6.14 easily adapts from general topological spaces to metric space so that we need only consider sequences when checking a set is closed..

Theorem 4.1.13. *Let (X, d) be a metric and let $F \subseteq X$. Then the following are equivalent:*

(i) F is a closed set in (X, d).

(ii) If $(x_n)_{n\geq 1}$ is a sequence such that $x_n \in F$ for all $n \in \mathbb{N}$ that converges to a point $x_0 \in X$, then $x_0 \in F$.

Proof. To begin, suppose F is a closed set in (X, d) and that $(x_n)_{n\geq 1}$ is a sequence such that $x_n \in F$ for all $n \in \mathbb{N}$ that converges to a point $x_0 \in X$. Suppose to the contrary that $x_0 \notin F$. Then $x_0 \in X \setminus F$. As F is closed, $X \setminus F$ is open so $x_0 \in X \setminus F$ and the definition of a convergent sequence implies there exists a $N \in \mathbb{N}$ such that $x_n \in X \setminus F$ for all $n \geq N$. As this contradicts the fact that $x_n \in F$ for all $n \in \mathbb{N}$, we have a contradiction. Hence $x_0 \in F$ as desired.

Conversely, suppose that whenever $(x_n)_{n\geq 1}$ is a sequence such that $x_n \in F$ for all $n \geq N$ that converges to a point $x_0 \in X$, then $x_0 \in F$. To see that F must be closed, suppose to the contrary that F is not closed. Then $X \setminus F$ is not open. Hence there exists a point $x_0 \in X \setminus F$ such that for every neighbourhood U of x_0, $U \cap F \neq \emptyset$. Hence for every $n \in \mathbb{N}$ there exists an $x_n \in \mathcal{B}_d\left(x_0, \frac{1}{n}\right) \cap F$. Hence $(x_n)_{n\geq 1}$ is a sequence in F such that

$$0 \leq \limsup_{n\to\infty} d(x_0, x_n) \leq \lim_{n\to\infty} \frac{1}{n} = 0.$$

Hence $(x_n)_{n\geq 1}$ is a sequence in F that converges to x_0. Hence $x_0 \in F$ thereby contradicting the fact that $x_0 \in X \setminus F$. Thus F is closed as desired. ∎

Theorem 4.1.14. *Let (X, d) be a space and let $A \subseteq X$ be non-empty. Then*

(1) If (A, d) is complete, then A is closed in X.

(2) If (X, d) is complete and A is closed in X, then (A, d) is complete.

Proof. Suppose (A, d) is complete. To see that A is closed in (X, d), let $(a_n)_{n\geq 1}$ be an arbitrary sequence of elements from A that converges to some element $x \in X$. Since $(a_n)_{n\geq 1}$ converges in X, $(a_n)_{n\geq 1}$ is Cauchy in (X, d) by Lemma 4.1.3 and therefore is Cauchy in (A, d) as (X, d) and (A, d) have the same metric. Hence $(a_n)_{n\geq 1}$ converges in A to some element $a \in A$ as (A, d) is complete. Since metric spaces are Hausdorff so limits are unique (Theorem 1.5.40), $a = x$. Hence $x \in A$ so A is closed by Theorem 1.6.14.

For the converse, suppose A is closed in (X, d). To see that (A, d) is complete, let $(a_n)_{n\geq 1}$ be an arbitrary Cauchy sequence in (A, d). Hence $(a_n)_{n\geq 1}$ is a Cauchy sequence in (X, d). Since (X, d) is complete, $(a_n)_{n\geq 1}$ converges to some element $x \in X$. Since A is closed in X, Theorem 1.6.14 implies that $x \in A$. Hence as $(a_n)_{n\geq 1}$ was an arbitrary Cauchy sequence, (A, d) is complete. ∎

One class of complete metric spaces that are particularly nice are the following.

Definition 4.1.15. A *Banach space* is a complete normed linear space.

One reason Banach spaces are superior to complete metric spaces is the vector space structures allow for us to take sums. In particular, we are interested in the following notions of series and how they intertwine with completeness.

Definition 4.1.16. Let $(X, \|\cdot\|)$ be a normed linear space. A series $\sum_{n=1}^\infty \vec{x}_n$ is said to be *summable* if the sequence of partial sums $(s_n)_{n\geq 1}$ (where $s_n = \sum_{k=1}^n \vec{x}_k$) converges in $(X, \|\cdot\|)$.

A series $\sum_{n=1}^\infty \vec{x}_n$ is said to be *absolutely summable* if $\sum_{n=1}^\infty \|\vec{x}_n\| < \infty$.

The beauty of a Banach space is it is precisely the structure one requires for a result from undergraduate analysis to hold.

Theorem 4.1.17. *Let $(X, \|\cdot\|)$ be a normed linear space. Then $(X, \|\cdot\|)$ is complete (i.e. a Banach space) if and only if every absolutely summable series is summable.*

Proof. Suppose $(X, \|\cdot\|)$ is complete. To see the result, let $\sum_{n=1}^\infty \vec{x}_n$ be an arbitrary absolutely summable series. To see that $\sum_{n=1}^\infty \vec{x}_n$ is summable, let $\epsilon > 0$ be arbitrary. Since $\sum_{n=1}^\infty \|\vec{x}_n\| < \infty$, there exists an $N \in \mathbb{N}$ such that $\sum_{n=N}^{\infty} \|\vec{x}_n\| < \epsilon$. Therefore, if $k, m \geq N$ and, without loss of generality, $m \geq k$, then

$$\begin{aligned} \|s_m - s_k\| &= \left\| \sum_{n=1}^{m} \vec{x}_n - \sum_{n=1}^{k} \vec{x}_n \right\| \\ &= \left\| \sum_{n=k+1}^{m} \vec{x}_n \right\| \\ &\leq \sum_{n=k+1}^{m} \|\vec{x}_n\| \\ &\leq \sum_{n=N}^{\infty} \|\vec{x}_n\| < \epsilon. \end{aligned}$$

Therefore, as $\epsilon > 0$ was arbitrary, the sequence of partial sums $(s_n)_{n\geq 1}$ is Cauchy. Hence $(s_n)_{n\geq 1}$ converges as $(X, \|\cdot\|)$ is complete. Thus, as $\sum_{n=1}^{\infty} \vec{x}_n$ was arbitrary, every absolutely summable series in $(X, \|\cdot\|)$ is summable.

For the converse, suppose every absolutely summable sequence in $(X, \|\cdot\|)$ is summable. To see that $(X, \|\cdot\|)$ is complete, let $(\vec{x}_n)_{n\geq 1}$ be an arbitrary Cauchy sequence. Since $(\vec{x}_n)_{n\geq 1}$ is Cauchy, there exists an $n_1 \in \mathbb{N}$ such that $\|\vec{x}_m - \vec{x}_j\| < \frac{1}{2}$ for all $m, j \geq n_1$. Similarly, since $(\vec{x}_n)_{n\geq 1}$ is Cauchy, there exists an $n_2 \in \mathbb{N}$ such that $n_2 > n_1$ and $\|\vec{x}_m - \vec{x}_j\| < \frac{1}{2^2}$ for all $m, j \geq n_2$. By repeating the above process, for each $k \in \mathbb{N}$ there exists an $n_k \in \mathbb{N}$ such that $n_k < n_{k+1}$ for all k and $\|\vec{x}_m - \vec{x}_j\| < \frac{1}{2^k}$ for all $m, j \geq n_k$.

For each $k \in \mathbb{N}$ let $\vec{y}_k = \vec{x}_{n_{k+1}} - \vec{x}_{n_k}$. By the above paragraph, we see that

$$\sum_{k=1}^{\infty} \|\vec{y}_k\| \leq \sum_{k=1}^{\infty} \frac{1}{2^k} < \infty.$$

Hence $\sum_{k=1}^{\infty} \vec{y}_k$ is an absolutely summable series in $(X, \|\cdot\|)$. Therefore, by the assumptions on $(X, \|\cdot\|)$, $\sum_{k=1}^{\infty} \vec{y}_k$ is summable in $(X, \|\cdot\|)$.

Let $\vec{x} = \vec{x}_{n_1} + \sum_{k=1}^{\infty} \vec{y}_k$. We claim that $(\vec{x}_{n_k})_{k\geq 1}$ converges to $\vec{x}$. To see this, let $\epsilon > 0$ be arbitrary. Then there exists a $M \in \mathbb{N}$ such that if $m \geq M$ then

$$\left\| \sum_{k=1}^{\infty} \vec{y}_k - \sum_{k=1}^{m} \vec{y}_k \right\| < \epsilon.$$

Therefore, if $m \geq M$,

$$\begin{aligned} \|\vec{x} - \vec{x}_{n_{m+1}}\| &\leq \left\| \sum_{k=1}^{\infty} \vec{y}_k - \sum_{k=1}^{m} \vec{y}_k \right\| + \left\| \vec{x}_{n_1} - \vec{x}_{n_{m+1}} + \sum_{k=1}^{m} \vec{y}_k \right\| \\ &< \epsilon + \left\| \vec{x}_{n_1} - \vec{x}_{n_{m+1}} + \sum_{k=1}^{m} \vec{x}_{n_{k+1}} - \vec{x}_{n_k} \right\| \\ &= \epsilon. \end{aligned}$$

Therefore, as $\epsilon > 0$ was arbitrary, $(\vec{x}_{n_k})_{k\geq 1}$ converges to $\vec{x}$. Hence $(\vec{x}_n)_{n\geq 1}$ converges to $\vec{x}$ by Lemma 4.1.6. Therefore, as $(\vec{x}_n)_{n\geq 1}$ was an arbitrary Cauchy sequence, $(X, \|\cdot\|)$ is complete. ∎

To see uses of Theorem 4.1.17, we better have some examples.

Proposition 4.1.18. *For each $p \in [1, \infty]$, $(\ell_p(\mathbb{K}), \|\cdot\|_p)$ is a Banach space.*

Proof. Note the beginning of this proof appears very much similar to Corollary 4.1.9. However, there are complications due to the fact that $\ell_p(\mathbb{K})$ is infinite dimensional.

Fix $p \in [1, \infty]$ and let $(\vec{x}_n)_{n\geq 1}$ be an arbitrary Cauchy sequence in $(\ell_p(\mathbb{K}), \|\cdot\|_p)$. For each $n \in \mathbb{N}$, write $\vec{x}_n = (x_{n,k})_{k\geq 1}$. Since for all $m, j, k \in \mathbb{N}$,

$$|x_{m,k} - x_{j,k}| \leq \|\vec{x}_m - \vec{x}_j\|_p,$$

we see that for each $k \in \mathbb{N}$ the sequence $(x_{n,k})_{n\geq 1}$ is Cauchy. Therefore, as $\mathbb{K}$ is complete, $y_k = \lim_{n\to\infty} x_{n,k}$ exists for each $k \in \mathbb{N}$.

Let $\vec{y} = (y_n)_{n\geq 1}$. To complete the proof, it suffices to verify two things: that $\vec{y} \in \ell_p(\mathbb{K})$; and that $\lim_{n\to\infty} \|\vec{y} - \vec{x}_n\|_p = 0$. We will only discuss the case $p \neq \infty$ and the case $p = \infty$ is similar. For the former, notice for all $m \in \mathbb{N}$ that

$$\left(\sum_{k=1}^{m} |y_k - x_{1,k}|^p\right)^{\frac{1}{p}} = \lim_{n\to\infty} \left(\sum_{k=1}^{m} |x_{n,k} - x_{1,k}|^p\right)^{\frac{1}{p}} \leq \limsup_{n\to\infty} \|\vec{x}_n - \vec{x}_1\|_p.$$

Since $(x_{n,k})_{n\geq 1}$ is Cauchy, $(x_{n,k})_{n\geq 1}$ is bounded by Lemma 4.1.4. Hence $\limsup_{n\to\infty} \|\vec{x}_n - \vec{x}_1\|_p$ is finite. Therefore, by taking the limit as m tends to infinity, we obtain that

$$\left(\sum_{k=1}^{\infty} |y_k - x_{1,k}|^p\right)^{\frac{1}{p}} \leq \limsup_{n\to\infty} \|\vec{x}_n - \vec{x}_1\|_p.$$

Hence $\vec{z} = (y_k - x_{1,k})_{k\geq 1} \in \ell_p(\mathbb{K})$. Therefore, as $\vec{y} = \vec{z} + \vec{x}_1$, we obtain that $\vec{y} \in \ell_p(\mathbb{K})$ by the triangle inequality.

To see that $\lim_{n\to\infty} \|\vec{y} - \vec{x}_n\|_p = 0$, let $\epsilon > 0$ be arbitrary. Note the above proof also shows for all $j \in \mathbb{N}$ that

$$\|\vec{y} - \vec{x}_j\|_p \leq \limsup_{n\to\infty} \|\vec{x}_n - \vec{x}_j\|_p.$$

Since $(\vec{x}_n)_{n\geq 1}$ is Cauchy in $(\ell_p(\mathbb{K}), \|\cdot\|_p)$, there exists an $N \in \mathbb{N}$ such that $\|\vec{x}_m - \vec{x}_j\|_p \leq \epsilon$ for all $m, j \geq N$. Hence if $j \geq N$, the above implies $\|\vec{y} - \vec{x}_j\|_p \leq \epsilon$. Therefore, as $\epsilon > 0$ was arbitrary, we obtain that $\lim_{n\to\infty} \|\vec{y} - \vec{x}_n\|_p = 0$. Hence $(\vec{x}_n)_{n\geq 1}$ converges in $(\ell_p(\mathbb{K}), \|\cdot\|_p)$ so, as $(\vec{x}_n)_{n\geq 1}$ was arbitrary, $(\ell_p(\mathbb{K}), \|\cdot\|_p)$ is complete. ∎

It is not difficult to see that Proposition 4.1.18 arises as $\mathbb{K}$ is complete and a sequence being Cauchy in $\ell_p(\mathbb{K})$ implies the sequence formed by each entry of the original sequence is Cauchy in $\mathbb{K}$ and thus converges. In particular, this sounds very much like a product of metric spaces. In particular, as products are really function spaces, this leads us to the question of which nice functions spaces are complete metric spaces?

4.2 Complete Function Spaces

In this section, we will analyze which function spaces we can place metric on in order to obtain complete metric spaces. In particular, as functions spaces with codomain $\mathbb{R}$ will be seen to be Banach spaces, we will have a use of Theorem 4.1.17 as it immediately implies the well-known Weierstrass M-Test. As function spaces are subsets of product spaces, we prove the following as a step towards making the product of complete metric space a complete metric space.

Proposition 4.2.1. *Let I be a non-empty set, let $\{(X_\alpha, d_\alpha)\}_{\alpha\in I}$ be a collection of metric spaces, and let $X = \prod_{\alpha\in I} X_\alpha$. Define $d : X \times X \to [0,\infty)$ by*

$$d((x_\alpha)_{\alpha\in I}, (y_\alpha)_{\alpha\in I}) = \sup_{\alpha\in I} \min(\{d_\alpha(x_\alpha, y_\alpha), 1\})$$

for all $(x_\alpha)_{\alpha\in I}, (y_\alpha)_{\alpha\in I} \in X$. Then d is a metric on X.

Proof. It is not difficult to verify for each $\alpha \in I$ that the map $d'_\alpha : X_\alpha \times X_\alpha \to [0, 1]$ defined by

$$d'_\alpha(x, y) = \min(\{d_\alpha(x, y), 1\})$$

for all $x, y \in X_\alpha$ is a metric on X_α as d_α is a metric on X_α. It is then easy to verify that d is a metric as the supremum of bounded metrics is a metric. ∎

Due the future usefulness of the above metric on a product of metric spaces, we provide it with a name.

Definition 4.2.2. Let $\{(X_\alpha, d_\alpha)\}_{\alpha\in I}$ be a non-empty collection of metric spaces and let $X = \prod_{\alpha\in I} X_\alpha$. The *uniform metric on the product* is the metric $d : X \times X \to [0,\infty)$ from Proposition 4.2.1 defined by

$$d((x_\alpha)_{\alpha\in I}, (y_\alpha)_{\alpha\in I}) = \sup_{\alpha\in I} \min(\{d_\alpha(x_\alpha, y_\alpha), 1\})$$

for all $(x_\alpha)_{\alpha\in I}, (y_\alpha)_{\alpha\in I} \in X$.

Remark 4.2.3. It turns out we have already seen the uniform metric on a product of metric spaces. Indeed if X is any non-empty set and (Y, d_Y) is a metric space, then the uniform metric on $\mathcal{F}(X, Y)$ is the uniform metric on the product $\prod_{x\in X} Y$. In particular, by Example 1.4.19, we know that the topology on a product of metric spaces induced by the uniform metric on the product need not be the box topology nor the product topology.

However, there is a relation between the strength of the product topology and the uniform metric topology on a product of metric spaces.

Lemma 4.2.4. *Let $\{(X_\alpha, d_\alpha)\}_{\alpha \in I}$ be a non-empty collection of metric spaces and let $X = \prod_{\alpha \in I} X_\alpha$. Then the uniform metric topology on X is finer than the product topology on X.*

Proof. Let d denote the uniform metric on X. To see that the uniform metric topology on X is finer than the product topology on X, let $(f_\lambda)_{\lambda \in \Lambda}$ be an arbitrary net in X that converges to $f_0 \in X$ with respect to the uniform metric topology. Let $\epsilon > 0$ be arbitrary. Since $(f_\lambda)_{\lambda \in \Lambda}$ converges to f_0 with respect to the metric topology, there exists an $\lambda_0 \in \Lambda$ such that $d(f_\lambda, f_0) < \epsilon$ for all $\lambda \geq \lambda_0$. Hence for all $\alpha \in I$

$$d_\alpha(f_\lambda(\alpha), f_0(\alpha)) \leq d(f_\lambda, f_0) < \epsilon$$

for all $\lambda \geq \lambda_0$. Therefore, as $\epsilon > 0$ was arbitrary, we obtain that $(f_\lambda(\alpha))_{\lambda \in \Lambda}$ converges to $f_0(\alpha)$ in (X_α, d_α) for all $\alpha \in I$. Therefore $(f_\lambda)_{\lambda \in \Lambda}$ converges to f_0 with respect to the product topology by Theorem 1.5.25. Hence Theorem 1.5.27 implies that uniform metric topology on X is finer than the product topology on X. ∎

One reason the uniform metric on a product is preferred is precisely makes the product of complete metric spaces complete. Thus the following uses the same base idea as Proposition 4.1.18 together with some finesse.

Theorem 4.2.5. *Let $\{(X_\alpha, d_\alpha)\}_{\alpha \in I}$ be a non-empty collection of metric spaces, let $X = \prod_{\alpha \in I} X_\alpha$, and let d be the uniform metric on the product X. If $(X_\alpha, d_\alpha)_{\alpha \in I}$ is complete for all $\alpha \in I$, then (X, d) is complete.*

Proof. Let $(f_n)_{n \geq 1}$ be an arbitrary Cauchy sequence in (X, d). Since for all $m, n \in \mathbb{N}$ and $\alpha \in I$ we have that

$$\min(\{d_\alpha(f_n(\alpha), f_m(\alpha)), 1\}) \leq d(f_n, f_m),$$

we see that for each $\alpha \in I$ the sequence $(f_n(\alpha))_{n \geq 1}$ is Cauchy in (X_α, d_α) (note taking the minimum of $d_\alpha(f_n(\alpha), f_m(\alpha))$ and 1 does not affect the sequence being Cauchy). Therefore, as (X_α, d_α) is complete, $x_\alpha = \lim_{n \to \infty} f_n(\alpha)$ in (X_α, d_α) for all $\alpha \in I$.

Let $x = (x_\alpha)_{\alpha \in I} \in X$. To complete the proof, it suffices to verify that $(f_n)_{n \geq 1}$ converges to x. To begin, notice for all $\alpha \in I$ and $n \in \mathbb{N}$ that

$$\begin{aligned} \min(\{d_\alpha(x_\alpha, f_n(\alpha)), 1\}) &= \lim_{m \to \infty} \min(\{d_\alpha(f_m(\alpha), f_n(\alpha)), 1\}) \\ &\leq \limsup_{m \to \infty} d(f_m, f_n) \end{aligned}$$

since $x_\alpha = \lim_{n\to\infty} f_n(\alpha)$ in (X_α, d_α). However, since this holds for all $\alpha \in I$, we see that

$$d(x, f_n) \leq \limsup_{m\to\infty} d(f_m, f_n)$$

for all $n \in \mathbb{N}$. However, since $(f_n)_{n\geq 1}$ is Cauchy in (X, d), for every $\epsilon > 0$ there exists an $N \in \mathbb{N}$ such that

$$d(x, f_n) \leq \limsup_{m\to\infty} d(f_m, f_n) < \epsilon$$

for all $n \geq N$. Hence we obtain that $(f_n)_{n\geq 1}$ converges to x in (X, d). Therefore, as $(f_n)_{n\geq 1}$ was arbitrary, (X, d) is complete as desired. ∎

Using Theorem 4.2.5, we can construct some immensely important complete metric spaces. Thus we begin by defining these spaces.

Definition 4.2.6. Let $(X, \mathcal{T}_X)$ and $(Y, \mathcal{T}_Y)$ be topological spaces. The *space of continuous functions from X to Y*, denoted $\mathcal{C}(X, Y)$, is the set of all continuous functions from X to Y. In the case that $Y = \mathbb{R}$ equipped with the canonical topology, we will use $\mathcal{C}(X)$ instead of $\mathcal{C}(X, \mathbb{R})$. In the case that $X = [a, b]$ equipped with the canonical subspace topology, we will use $\mathcal{C}[a, b]$ instead of $\mathcal{C}([a, b])$.

Definition 4.2.7. Let $(X, \mathcal{T})$ be a topological space and let (Y, d) be a metric space. The *space of bounded functions from X to Y*, denoted $\mathcal{B}(X, Y)$, is the set of all functions $f : X \to Y$ such that $f(X)$ is a bounded subset of (Y, d).

Definition 4.2.8. Let $(X, \mathcal{T})$ be a topological space and let (Y, d) be a metric space. The *space of continuous bounded functions from X to Y*, denoted $\mathcal{C}_b(X, Y)$, is the set $\mathcal{C}(X, Y) \cap \mathcal{B}(X, Y)$.

Theorem 4.2.9. *Let $(X, \mathcal{T})$ be a topological space and let (Y, d_Y) be a metric space. Then $\mathcal{C}(X, Y)$, $\mathcal{B}(X, Y)$, and $\mathcal{C}_b(X, Y)$ are closed in $(\mathcal{F}(X, Y), d)$ where d is the uniform metric. Consequently, if (Y, d_Y) is complete, $\mathcal{C}(X, Y)$, $\mathcal{B}(X, Y)$, and $\mathcal{C}_b(X, Y)$ are complete with respect to the uniform metric.*

Proof. To see that $\mathcal{C}(X, Y)$ is closed in $(\mathcal{F}(X, Y), d)$, let $(f_\lambda)_{\lambda\in\Lambda}$ be a net in $\mathcal{C}(X, Y)$ that converges to some element $f \in \mathcal{F}(X, Y)$ with respect to d. Hence $(f_\lambda)_{\lambda\in\Lambda}$ converges uniformly to f so Theorem 2.1.17 implies that $f \in \mathcal{C}(X, Y)$. Hence $\mathcal{C}(X, Y)$ is a closed subset of $(\mathcal{F}(X, Y), d)$ by Theorem 1.6.14.

To see that $\mathcal{B}(X, Y)$ is closed in $(\mathcal{F}(X, Y), d)$, let $(f_\lambda)_{\lambda\in\Lambda}$ be a net in $\mathcal{B}(X, Y)$ that converges to some element $f \in \mathcal{F}(X, Y)$ with respect to d. Hence there exists an $\lambda_0 \in \Lambda$ such that

$$d(f, f_{\lambda_0}) < \frac{1}{2}.$$

Therefore $d_Y(f(x), f_{\lambda_0}(x)) < \frac{1}{2}$ for all $x \in X$. However, since $f_{\lambda_0} \in \mathcal{B}(X,Y)$ so $f_{\lambda_0}(X)$ is bounded in (Y, d_Y), we know there exists an $M \in \mathbb{R}$ such that

$$\{d_Y(f_{\lambda_0}(x_1), f_{\lambda_0}(x_2)) \mid x_1, x_2 \in X\} \subseteq [0, M].$$

Thus, for all $x_1, x_2 \in X$, we see that

$$\begin{aligned} &d_Y(f(x_1), f(x_2)) \\ &\leq d_Y(f(x_1), f_{\lambda_0}(x_1)) + d_Y(f_{\lambda_0}(x_1), f_{\lambda_0}(x_2)) + d_Y(f_{\lambda_0}(x_2), f(x_2)) \\ &\leq \frac{1}{2} + M + \frac{1}{2} = M + 1. \end{aligned}$$

Hence

$$\{d_Y(f(x_1), f(x_2)) \mid x_1, x_2 \in X\} \subseteq [0, M+1]$$

so $f(X)$ is bounded and thus $f \in \mathcal{B}(X,Y)$. Hence $\mathcal{B}(X,Y)$ is a closed subset of $(\mathcal{F}(X,Y), d)$ by Theorem 1.6.14.

Since $\mathcal{C}_b(X,Y) = \mathcal{C}(X,Y) \cap \mathcal{B}(X,Y)$, we obtain that $\mathcal{C}_b(X,Y)$ is closed by the above. Finally, if (Y, d_Y) is complete, then $(\mathcal{F}(X,Y), d)$ is complete by Theorem 4.2.5. Therefore, since $\mathcal{C}(X,Y)$, $\mathcal{B}(X,Y)$, and $\mathcal{C}_b(X,Y)$ are closed in $(\mathcal{F}(X,Y), d)$, $\mathcal{C}(X,Y)$, $\mathcal{B}(X,Y)$, and $\mathcal{C}_b(X,Y)$ are complete with respect to d by Theorem 4.1.14. ∎

Of course, having to take the minimum with 1 inside the uniform metric is sort of annoying. However, in most circumstances, we can avoid this minimum. In particular, we note the following.

Lemma 4.2.10. *Let $(X, \mathcal{T})$ be a topological space and let (Y, d_Y) be a metric space. Define $d : \mathcal{B}(X,Y) \times \mathcal{B}(X,Y) \to [0, \infty)$ by*

$$d(f, g) = \sup_{x \in X} d_Y(f(x), g(x)).$$

Then d is a well-defined metric on $\mathcal{B}(X,Y)$. Furthermore, d defines the same topology on $\mathcal{B}(X,Y)$ as the uniform metric d_{unif} and a sequence is Cauchy in $(\mathcal{B}(X,Y), d)$ if and only if it is Cauchy in $(\mathcal{B}(X,Y), d_{\mathrm{unif}})$. Thus $(\mathcal{B}(X,Y), d)$ is complete when (Y, d_Y) is complete.

Proof. To see that the range of d is contained in $[0, \infty)$, let $f, g \in \mathcal{B}(X,Y)$ be arbitrary. Hence there exists $M_1, M_2 \in \mathbb{R}$ such that

$$\begin{aligned} \sup(\{d_Y(f(x_1), f(x_2)) \mid x_1, x_2 \in X\}) &\subseteq [0, M_1] \text{ and} \\ \sup(\{d_Y(g(x_1), g(x_2)) \mid x_1, x_2 \in X\}) &\subseteq [0, M_2]. \end{aligned}$$

Fix an $x_0 \in X$. Then, for all $x \in X$ we see that

$$\begin{aligned} d_Y(f(x), g(x)) &\leq d_Y(f(x), f(x_0)) + d_Y(f(x_0), g(x_0)) + d_Y(g(x_0), g(x)) \\ &\leq M_1 + d_Y(f(x_0), g(x_0)) + M_2. \end{aligned}$$

Hence, as M_1, M_2, and x_0 are fixed, we see that

$$d(f, g) = \sup_{x \in X} d_Y(f(x), g(x)) < \infty.$$

Hence, as $f, g \in \mathcal{B}(X, Y)$ were arbitrary, the range of d is contained in $[0, \infty)$.

The fact that d is a metric then trivially follows from the fact that d_Y is a metric. Moreover, it is elementary to see that a net converges in $(\mathcal{B}(X, Y), d)$ if and only if it converges in $(\mathcal{B}(X, Y), d_{\text{unif}})$ and a sequence is Cauchy in $(\mathcal{B}(X, Y), d)$ if and only if it is Cauchy in $(\mathcal{B}(X, Y), d_{\text{unif}})$ as these properties are defined so that only arbitrary small $\epsilon > 0$ matter and since for $\epsilon \in (0, 1)$ we see that $d(f, g) < \epsilon$ if and only if $d_{\text{unif}}(f, g) < \epsilon$. ∎

As the metric from Lemma 4.2.10 is far easier to use than the uniform metric on a product, a name is required.

Definition 4.2.11. Let $(X, \mathcal{T})$ be a topological space and let (Y, d_Y) be a metric space. The *sup metric on* $\mathcal{B}(X, Y)$ is the metric $d : \mathcal{B}(X, Y) \times \mathcal{B}(X, Y) \to [0, \infty)$ from Lemma 4.2.10 defined by

$$d(f, g) = \sup_{x \in X} d_Y(f(x), g(x))$$

for all $f, g \in \mathcal{B}(X, Y)$.

Remark 4.2.12. Of course, since $\mathcal{C}_b(X, Y) \subseteq \mathcal{B}(X, Y)$, we easily see that the sup metric defines a metric on $\mathcal{C}_b(X, Y)$ and that $\mathcal{C}_b(X, Y)$ and $\mathcal{B}(X, Y)$ are complete with respect to the sup metric by Theorem 4.2.9 provided (Y, d_Y) is complete. However, it is not necessary that the sup metric is well-defined on $\mathcal{C}(X, Y)$ as there are easily seen to be examples of continuous functions that are not bounded.

There is one situation where the sup metric is well-defined on $\mathcal{C}(X, Y)$.

Theorem 4.2.13. *Let $(X, \mathcal{T})$ be a compact topological space and let (Y, d) be a metric space. The sup metric is a well-defined metric on $\mathcal{C}(X, Y)$ and $\mathcal{C}(X, Y)$ is complete with respect to the sup metric provided (Y, d) is complete.*

Proof. We claim that $\mathcal{C}(X, Y) = \mathcal{C}_b(X, Y)$ from which the theorem easily follows by Remark 4.2.12. To see this, we note that the inclusion $\mathcal{C}_b(X, Y) \subseteq \mathcal{C}(X, Y)$ is clearly. For the reverse inclusion, let $f \in \mathcal{C}(X, Y)$ be arbitrary. Since $(X, \mathcal{T})$ is compact, Theorem 3.1.26 implies that $f(X)$ is a compact subset of (Y, d). Hence Theorem 3.1.22 implies that $f(X)$ is bounded in (Y, d). Thus $f \in \mathcal{C}_b(X, Y)$ by definition. Therefore, as $f \in \mathcal{C}(X, Y)$ was arbitrary, the result follows. ∎

We can even add a nicer structure to these function spaces provided (Y, d) is actually a metric space induced by a norm.

Theorem 4.2.14. *Let $(X, \mathcal{T})$ be a topological space and let $(Y, \|\cdot\|_Y)$ be a normed linear space over $\mathbb{K}$. Then $\mathcal{B}(X,Y)$ is a vector space over $\mathbb{K}$ and the sup metric d on $\mathcal{B}(X,Y)$ is a metric induced by the norm $\|\cdot\|_\infty : \mathcal{B}(X,Y) \to [0,\infty)$ defined by*

$$\|f\|_\infty = \sup_{x \in X} \|f(x)\|_Y$$

for all $f \in \mathcal{B}(X,Y)$. Consequently, if $(Y, \|\cdot\|_Y)$ is a Banach space, then $(\mathcal{B}(X,Y), \|\cdot\|_\infty)$ and $(\mathcal{C}_b(X,Y), \|\cdot\|_\infty)$ are Banach spaces. Furthermore, if $(X, \mathcal{T})$ is compact and $(Y, \|\cdot\|_Y)$ is a Banach space, then $(\mathcal{C}(X,Y), \|\cdot\|_\infty)$ is a Banach space.

Proof. Recall $\mathcal{F}(X,Y)$ is a vector space over $\mathbb{K}$. To see that $\mathcal{B}(X,Y)$ is a vector subspace of $\mathcal{F}(X,Y)$, note by Lemma 3.1.21 that a set $A \subseteq Y$ is bounded in $(Y, \|\cdot\|_Y)$ if and only if

$$\sup(\{\|\vec{a}\|_Y \mid \vec{a} \in A\}) < \infty.$$

Thus it is elementary to see that if $A, B \subseteq Y$ are bounded and $\alpha \in \mathbb{K}$, then

$$A + B = \{\vec{a} + \vec{b} \mid \vec{a} \in A, \vec{b} \in B\} \quad \text{and} \quad \alpha A = \{\alpha\vec{a} \mid \vec{a} \in A\}$$

are bounded subsets of $(Y, \|\cdot\|_Y)$. Thus it is now elementary to verify that $\mathcal{B}(X,Y)$ is a vector subspace of $\mathcal{F}(X,Y)$ by definitions.

It is elementary to verify that $\|\cdot\|_\infty$ is indeed a norm on $\mathcal{B}(X,Y)$ that induces the sup metric. Hence the remainder of the statements follow from Remark 4.2.12, Theorem 4.2.13, and the fact that the addition and scalar multiple of continuous function is continuous by considering the net characterization of continuity (Theorem 2.1.9) and the continuity of addition and scalar multiplication in vector spaces (see Proposition 1.5.43). ∎

The importance of the above norm merits a name.

Definition 4.2.15. Let $(X, \mathcal{T})$ be a topological space and let $(Y, \|\cdot\|_Y)$ be a normed linear space over $\mathbb{K}$. The *infinity norm on $\mathcal{B}(X,Y)$* is the norm $\|\cdot\|_\infty : \mathcal{B}(X,Y) \to [0,\infty)$ from Theorem 4.2.14 defined by

$$\|f\|_\infty = \sup_{x \in X} \|f(x)\|_Y$$

for all $f \in \mathcal{B}(X,Y)$.

There is an additional excellent subset of the continuous functions into a normed linear space we may wish to consider in regards to the infinity norm.

Definition 4.2.16. Let $(X, \mathcal{T})$ be a topological space and let $(Y, \|\cdot\|)$ be a normed linear space. A continuous function $f \in \mathcal{C}(X,Y)$ is said to *vanish at infinity* if for all $\epsilon > 0$ there exists a compact subset $K \subseteq X$ such that $\|f(x)\| < \epsilon$ for all $x \in X \setminus K$.

The set of continuous functions from X to Y that vanish at infinity is denoted $\mathcal{C}_0(X, Y)$. In the case that $Y = \mathbb{R}$ with the canonical norm, we will use $\mathcal{C}_0(X)$ in lieu of $\mathcal{C}_0(X, \mathbb{R})$.

Example 4.2.17. Let $\mathbb{N}$ be equipped with the discrete topology. Hence every function from $\mathbb{N}$ to $\mathbb{R}$ is continuous so $\mathcal{C}(\mathbb{N})$ may be identified with all sequences in $\mathbb{R}$. To determine $\mathcal{C}_0(\mathbb{N})$, we note that the compact subsets of $\mathbb{N}$ are the finite subsets. Hence an $f \in \mathcal{C}(\mathbb{N})$ vanishes at infinity if and only if for all $\epsilon > 0$ the set $\{n \in \mathbb{N} \mid |f(n)| \geq \epsilon\}$ is finite. Thus $\mathcal{C}_0(\mathbb{N})$ is the set of all sequences of elements of $\mathbb{R}$ that converge to 0.

As the Heine-Borel Theorem implies the compact subsets of $\mathbb{R}$ are bounded, a similar examination of $\mathcal{C}_0(\mathbb{R})$ shows that $\mathcal{C}_0(\mathbb{R})$ consists of all $f \in \mathcal{C}(\mathbb{R})$ such that $\lim_{x\to\pm\infty} f(x) = 0$. Hence the notation $\mathcal{C}_0$ and the name 'vanish at infinity'.

Of course, to study $\mathcal{C}_0(X, Y)$, one would like to assume that $(X, \mathcal{T})$ is at least locally compact in order for there to be a plethora of compact subsets to consider. If in addition we assume that $(X, \mathcal{T})$ is Hausdorff, we obtain the following from which is the true reason we call these the continuous functions that 'vanishing at infinity'.

Theorem 4.2.18. *Let $(X, \mathcal{T}_X)$ be a locally compact, Hausdorff topological space and let $(Y, \|\cdot\|)$ be a normed linear space. Let $(Z, \mathcal{T}_Z)$ be the one-point compactification of $(X, \mathcal{T}_X)$ with $Z \setminus X = \{\infty\}$. Then*

$$\mathcal{C}_0(X, Y) = \left\{g|_X \mid g \in \mathcal{C}(Z, Y), f(\infty) = \vec{0}\right\}.$$

Proof. For notational simplicity, let

$$C = \left\{f|_X \mid f \in \mathcal{C}(Z, Y), f(\infty) = \vec{0}\right\}.$$

To see that $\mathcal{C}_0(X, Y) = C$, first let $f \in \mathcal{C}(X, Y)$ be arbitrary. Define $g : Z \to Y$ by

$$g(z) = \begin{cases} f(z) & \text{if } z \in X \\ \vec{0} & \text{if } z = \infty \end{cases}.$$

Clearly $g|_X = f$. Thus to show that $f \in C$ it suffices to show that g is continuous on Z. To see that g is continuous, let $(z_\lambda)_{\lambda\in\Lambda}$ be a net in Z that converges in $(Z, \mathcal{T}_Z)$ to some $z_0 \in Z$.

If $z_0 \neq \infty$ so $z_0 \in X$, then as $(Z, \mathcal{T}_Z)$ is Hausdorff there exists a $U \in \mathcal{T}_Z$ such that $z_0 \in U$ and $\infty \notin U$. Hence $U \in \mathcal{T}_X$ and, as $(z_\lambda)_{\lambda\in\Lambda}$ converges to z_0 in $(Z, \mathcal{T}_Z)$ there exists a $\lambda_0 \in \Lambda$ such that $z_\lambda \in U$ for all $\lambda \geq \lambda_0$. Hence $(z_\lambda)_{\lambda\geq\lambda_0}$ is a net in $(X, \mathcal{T}_X)$. Therefore, since $(z_\lambda)_{\lambda\in\Lambda}$ converges to z_0 in $(Z, \mathcal{T}_Z)$ and $(X, \mathcal{T}_X)$ is a subspace of $(Z, \mathcal{T}_Z)$, we obtain that $(z_\lambda)_{\lambda\geq\lambda_0}$ converges to z_0 in $(X, \mathcal{T})$. Hence $g(z_\lambda) = f(z_\lambda)$ for all $\lambda \geq \lambda_0$ so $(g(z_\lambda))_{\lambda\geq\lambda_0}$

converges to $f(z_0) = g(z_0)$ in $(Y, \|\cdot\|)$. The existence of upper bounds for nets then implies that $(g(z_\lambda))_{\lambda\in\Lambda}$ converges to $g(z_0)$ in $(Y, \|\cdot\|)$ as desired.

Otherwise suppose $z_0 = \infty$. To see that $(g(z_\lambda))_{\lambda\in\Lambda}$ converges to $g(z_0)$ in $(Y, \|\cdot\|)$, let $\epsilon > 0$ be arbitrary. As $f \in \mathcal{C}_0(X, Y)$, there exists a compact set $K \subseteq X$ such that $\|f(x)\| < \epsilon$ for all $x \in X \setminus K$. However, as K is compact and $(Z, \mathcal{T}_Z)$ is Hausdorff, K is closed in $(Z, \mathcal{T}_Z)$ by Theorem 3.1.13 so $Z \setminus K$ is a neighbourhood of z_0. Thus, as $(z_\lambda)_{\lambda\in\Lambda}$ converges in $(Z, \mathcal{T}_Z)$ to z_0, there exists a $\lambda_0 \in \Lambda$ such that $z_\lambda \in Z \setminus K$ for all $\lambda \geq \lambda_0$. Hence if $\lambda \geq \lambda_0$ then either $z_\lambda = \infty$ so $g(z_\lambda) = 0$ or $z_\lambda \in X \setminus K$ so $f(z_\lambda) = g(z_\lambda)$ so $\|g(z_\lambda)\| < \epsilon$. Hence

$$\|g(z_0) - g(z_\lambda)\| = \|g(z_\lambda)\| < \epsilon$$

for all $\lambda \geq \lambda_0$. Therefore, as $\epsilon > 0$ was arbitrary, $(g(z_\lambda))_{\lambda\in\Lambda}$ converges to $g(z_0)$ in $(Y, \|\cdot\|)$. Hence g is continuous on Z and the above implies that $\mathcal{C}_0(X, Y) \subseteq C$.

To see the reverse inclusion, suppose $f \in C$. Hence $f = g|_X$ where $g \in \mathcal{C}(Z, Y)$ is such that $g(\infty) = 0$. Clearly $f = g|_X$ is continuous by Lemma 2.1.18. To see that $f \in \mathcal{C}_0(X, Y)$, let $\epsilon > 0$ be arbitrary. Since $g \in \mathcal{C}(Z, Y)$ and since $g(\infty) = 0$, there exists a $U \in \mathcal{T}_Z$ such that $\|g(z)\| < \epsilon$ for all $z \in U$. However, due to the properties of one-point compactification, $K = X \setminus U$ is a compact subset of $(X, \mathcal{T}_X)$ such that $\|f(x)\| < \epsilon$ for all $x \in X \setminus K \subseteq U$. Hence, as $\epsilon > 0$ was arbitrary, $f \in C_0(X, Y)$ as desired. ∎

Remark 4.2.19. Theorem 4.2.18 easily implies that if $(X, \mathcal{T}_X)$ is a locally compact, Hausdorff topological space, $(Z, \mathcal{T}_Z)$ is the one-point compactification of $(X, \mathcal{T}_X)$, and $(Y, \|\cdot\|)$ is a normed linear space, then $\mathcal{C}_0(X, Y)$ is a $\|\cdot\|_\infty$-closed subset of $\mathcal{C}(Z, Y)$ and thus complete with respect to the infinity norm. In addition, it is not difficult to verify that $\mathcal{C}_0(X, Y)$ is then a vector subspace of $\mathcal{B}(X, Y)$ and hence a Banach space. In order to obtain this result without appealing to the one-point compactification argument required in Theorem 4.2.18, we present the following more general argument.

Theorem 4.2.20. *Let $(X, \mathcal{T})$ be a topological space and let $(Y, \|\cdot\|_Y)$ be a normed linear space over $\mathbb{K}$. The space $\mathcal{C}_0(X, Y)$ is a $\|\cdot\|_\infty$-closed vector subspace of $\mathcal{B}(X, Y)$. Hence if $(Y, \|\cdot\|)$ is a Banach space, then $\mathcal{C}_0(X, Y)$ is a Banach space.*

Proof. First, to see that $\mathcal{C}_0(X, Y) \subseteq \mathcal{B}(X, Y)$, let $f \in \mathcal{C}_0(X, Y)$ be arbitrary. Hence there exists a compact set $K \subseteq X$ such that $\|f(x)\|_Y < 1$ for all $x \in X \setminus K$. Since f is continuous and K is compact, $f(K)$ is compact in $(Y, \|\cdot\|_Y)$ by Theorem 3.1.26 and thus bounded in $(Y, \|\cdot\|_Y)$ by Theorem 3.1.22. Thus Lemma 3.1.21 implies there exists an $M \in \mathbb{R}$ such that $\|f(x)\|_Y \leq M$ for all $x \in K$. Thus

$$\|f(x)\|_Y \leq \max(\{M, 1\})$$

for all $x \in X$ so $f(X)$ is bounded by Lemma 3.1.21. Hence, as f was arbitrary, $\mathcal{C}_0(X, Y) \subseteq \mathcal{B}(X, Y)$.

To see that $\mathcal{C}_0(X,Y)$ is a vector subspace of $\mathcal{B}(X,Y)$, let $f,g \in \mathcal{C}_0(X,Y)$ and $\alpha \in \mathbb{K}$ be arbitrary. Clearly $f+g, \alpha f \in \mathcal{C}(X,Y)$ by considering the net characterization of continuity (Theorem 2.1.9) and the continuity of addition and scalar multiplication in vector spaces (see Proposition 1.5.43). To see that $f+g, \alpha f \in \mathcal{C}_0(X,Y)$, recall since $f,g \in \mathcal{C}_0(X,Y)$ that for all $\epsilon > 0$ there exist compact subsets $K_1, K_2 \subseteq X$ such that

$$\|f(x)\|_Y < \epsilon \text{ for all } x \in X \setminus K_1 \qquad \text{and} \qquad \|g(x)\|_Y < \epsilon \text{ for all } x \in X \setminus K_2.$$

Hence

$$\|(\alpha f)(x)\|_Y \leq |\alpha|\epsilon \text{ for all } x \in X \setminus K_1$$

immediately implies $\alpha f \in \mathcal{C}_0(X,Y)$ as α was fixed and ϵ was arbitrary. Moreover, if $K = K_1 \cup K_2$, then K is a compact subset of $(X,\mathcal{T})$ such that

$$\|f(x)\|_Y < \epsilon \text{ for all } x \in X \setminus K \qquad \text{and} \qquad \|g(x)\|_Y < \epsilon \text{ for all } x \in X \setminus K.$$

Hence

$$\|f(x)+g(x)\|_Y < 2\epsilon \text{ for all } x \in X \setminus K$$

so $f+g \in \mathcal{C}_0(X,Y)$ as desired. Hence $\mathcal{C}_0(X,Y)$ is a vector subspace of $\mathcal{B}(X,Y)$.

To see that $\mathcal{C}_0(X,Y)$ is a $\|\cdot\|_\infty$-closed subset of $\mathcal{B}(X,Y)$, let $(f_\lambda)_{\lambda\in\Lambda}$ be an arbitrary net in $\mathcal{C}_0(X,Y)$ that converges to some element $f \in \mathcal{B}(X,Y)$ with respect to $\|\cdot\|_\infty$. By Theorem 2.1.17, we obtain that $f \in \mathcal{C}(X,Y)$. To see that $f \in \mathcal{C}_0(X,Y)$, let $\epsilon > 0$ be arbitrary. Since $(f_\lambda)_{\lambda\in\Lambda}$ converges to f with respect to $\|\cdot\|_\infty$, there exists a $\lambda_0 \in \Lambda$ such that

$$\|f - f_{\lambda_0}\|_\infty < \frac{\epsilon}{2}.$$

Moreover, since $f_{\lambda_0} \in \mathcal{C}_0(X,Y)$, there exists a compact subset K of $(X,\mathcal{T})$ such that

$$\|f_{\lambda_0}(x)\|_Y < \frac{\epsilon}{2}$$

for all $x \in X \setminus K$. Hence for all $x \in X \setminus K$ we have that

$$\|f(x)\|_Y \leq \|f(x) - f_{\lambda_0}(x)\|_Y + \|f_{\lambda_0}(x)\|_Y < \|f - f_{\lambda_0}\|_\infty + \frac{\epsilon}{2} < \frac{\epsilon}{2} + \frac{\epsilon}{2} = \epsilon.$$

Hence $f \in \mathcal{C}_0(X,Y)$ as desired so $\mathcal{C}_0(X,Y)$ is a $\|\cdot\|_\infty$-closed subset of $\mathcal{B}(X,Y)$.

Finally, if $(Y, \|\cdot\|_Y)$ is a Banach space, then the fact that $\mathcal{C}_0(X,Y)$ is a Banach space with respect to the infinity norm then follows from Theorem 4.1.14 and Theorem 4.2.9. ∎

To end this section, we desire to address two remaining questions: which metric spaces are subspaces of the continuous functions on a topological space, and how far is every metric space from being complete? These questions are easily answered with the above knowledge of our function spaces. To discuss these questions, it is best to have a stronger version of the notion of homeomorphism for metric spaces.

Definition 4.2.21. Let (X, d_X) and (Y, d_Y) be metric spaces. A function $\varphi : X \to Y$ is said to be an *isometry* if $d_Y(\varphi(x_1), \varphi(x_2)) = d_X(x_1, x_2)$ for all $x_1, x_2 \in X$.

If in addition to being an isometry φ is a bijection, it is said that φ is an *isomorphism*. Finally, its is said that (X, d_X) and (Y, d_Y) are *isomorphic* if there exists an isomorphism from X to Y.

Using isomorphisms, we see that every metric space can be viewed inside a function space.

Theorem 4.2.22. *Let (X, d_X) be a metric space. Then (X, d_X) is isomorphic to a subset of $(\mathcal{C}_b(X, \mathbb{R}), \|\cdot\|_\infty)$.*

Proof. Fix a point $a \in X$. For each $z \in X$, define a function $f_z : X \to \mathbb{R}$ by

$$f_z(x) = d(x, z) - d(x, a)$$

for all $x \in X$. We claim that $f_z \in \mathcal{C}_b(X, \mathbb{R})$. To see this, notice for all $x \in X$ that

$$|f_z(x)| = |d(x, z) - d(x, a)| \leq d(z, a).$$

Hence f_z is bounded by $d(z, a)$. Furthermore, to see that f_z is continuous, we notice that the functions $x \mapsto d(x, z)$ and $x \mapsto d(x, a)$ are continuous by the triangle inequality (i.e. $|d(y, a) - d(x, a)| \leq d(x, y)$ for all $x, y \in X$). Hence $f_z \in \mathcal{C}_b(X, \mathbb{R})$.

Define the map $\varphi : X \to \mathcal{C}_b(X, \mathbb{R})$ by

$$\varphi(z) = f_z.$$

We claim that φ is an isomorphism. To see this, notice for all $z_1, z_2 \in X$ and $x \in X$ that

$$\begin{aligned} |f_{z_1}(x) - f_{z_2}(x)| &= |(d(x, z_1) - d(x, a)) - (d(x, z_2) - d(x, a))| \\ &= |d(x, z_1) - d(x, z_2)| \leq d(z_1, z_2) \end{aligned}$$

by the triangle inequality. Hence $\|\varphi(z_1) - \varphi(z_2)\|_\infty \leq d(z_1, z_2)$ for all $z_1, z_2 \in X$. However, since

$$|f_{z_1}(z_2) - f_{z_2}(z_2)| = |(d(z_2, z_1) - d(z_2, a)) - (d(z_2, z_2) - d(z_2, a))| = d(z_2, z_1)$$

we obtain that $\|\varphi(z_1) - \varphi(z_2)\|_\infty = d(z_1, z_2)$. Hence φ is an isometry as desired. ∎

Using Theorem 4.2.22 to embed a metric space (X, d) into a function space, we can use the completeness of function spaces to get a complete metric space containing (X, d). This is given the following name.

Definition 4.2.23. Let (X, d_X) be a metric space. A *completion* of (X, d_X) is a complete metric space (Y, d_Y) such that there exists an isometry $\varphi : X \to Y$ such that $\overline{\varphi(X)} = Y$.

Corollary 4.2.24. *Every metric space has a completion.*

Proof. Let (X, d) be a metric space. By Theorem 4.2.22, there exists a subset $A \subseteq (\mathcal{C}_b(X, \mathbb{R}), \|\cdot\|_\infty)$ such that X is isomorphic to A. As $(\mathcal{C}_b(X, \mathbb{R}), \|\cdot\|_\infty)$ is complete by Theorem 4.2.9, $\overline{A}$ is complete by Theorem 4.1.14. Hence $\overline{A}$ is a completion of X by definition. ∎

Of course, it would be nice if each metric space only had one completion. The following demonstrates this is the case.

Proposition 4.2.25. *Any two completions of a metric space are isomorphic.*

Proof. Let $(\mathcal{X}, d_\mathcal{X})$ be a metric space. Suppose that $(\mathcal{Y}, d_\mathcal{Y})$ and $(\mathcal{Z}, d_\mathcal{Z})$ are completions of $(\mathcal{X}, d_\mathcal{X})$. Therefore there exists isometries $\varphi_\mathcal{Y} : \mathcal{X} \to \mathcal{Y}$ and $\varphi_\mathcal{Z} : \mathcal{X} \to \mathcal{Z}$ such that $\overline{\varphi_\mathcal{Y}(\mathcal{X})} = \mathcal{Y}$ and $\overline{\varphi_\mathcal{Z}(\mathcal{X})} = \mathcal{Z}$. Our goal is to extend the identity map from $\mathcal{X} \subseteq \mathcal{Y}$ to $\mathcal{X} \subseteq \mathcal{Z}$ to obtain an isometry from $\mathcal{Y}$ to $\mathcal{Z}$. To do this, we will make use of the fact that $\mathcal{Y}$ and $\mathcal{Z}$ are complete and thus have convergent Cauchy sequences.

To define an isometry $\varphi : \mathcal{Y} \to \mathcal{Z}$, let $y \in \mathcal{Y}$ be arbitrary. Hence, as $\mathcal{Y}$ is the closure of $\mathcal{X}$ there exists a sequence $(x_n)_{n\geq 1}$ of elements of $\mathcal{X}$ such that $y = \lim_{n\to\infty} \varphi_\mathcal{Y}(x_n)$. However, as $(\varphi_\mathcal{Y}(x_n))_{n\geq 1}$ converges in $(\mathcal{Y}, d_\mathcal{Y})$, $(\varphi_\mathcal{Y}(x_n))_{n\geq 1}$ is Cauchy in $(\mathcal{Y}, d_\mathcal{Y})$. Therefore, $(x_n)_{n\geq 1}$ is Cauchy in $(\mathcal{X}, d_\mathcal{X})$ as $\varphi_\mathcal{Y}$ is an isometry. Hence $(\varphi_\mathcal{Z}(x_n))_{n\geq 1}$ also must be Cauchy as $\varphi_\mathcal{Z}$ is an isometry. Since $(\mathcal{Z}, d_\mathcal{Z})$ is complete, $(\varphi_\mathcal{Z}(x_n))_{n\geq 1}$ converges in $(\mathcal{Z}, d_\mathcal{Z})$. Let $z_y = \lim_{n\to\infty} \varphi_\mathcal{Z}(x_n)$. We would like to define $\varphi : \mathcal{Y} \to \mathcal{Z}$ such that $f(y) = z_y$.

There is one technical issue with this definition that we should get out of the way; that is, we desire to show that if $(x'_n)_{n\geq 1}$ is another sequence of elements of $\mathcal{X}$ such that $y = \lim_{n\to\infty} \varphi_\mathcal{Y}(x'_n)$, then $z_y = \lim_{n\to\infty} \varphi_\mathcal{Z}(x'_n)$. This will demonstrate that the sequence of elements of $\mathcal{X}$ we choose converging to $y \in \mathcal{Y}$ does not affect the limit in $(\mathcal{Z}, d_\mathcal{Z})$. To see this, notice by the triangle inequality and properties of limits that

$$\begin{aligned}
\lim_{n\to\infty} d_\mathcal{Z}(\varphi_\mathcal{Z}(x'_n), \varphi_\mathcal{Z}(x_n)) &= \lim_{n\to\infty} d_\mathcal{X}(x'_n, x_n) \\
&= \lim_{n\to\infty} d_\mathcal{Y}(\varphi_\mathcal{Y}(x'_n), \varphi_\mathcal{Y}(x_n)) \\
&= d_\mathcal{Y}(y, y) = 0.
\end{aligned}$$

Hence as $z_y = \lim_{n\to\infty} \varphi_{\mathcal{Z}}(x_n)$, the above easily implies $z_y = \lim_{n\to\infty} \varphi_{\mathcal{Z}}(x'_n)$. Hence the claim is complete.

Hence we may define $\varphi : \mathcal{Y} \to \mathcal{Z}$ as follows: for each $y \in \mathcal{Y}$ choose a sequence $(x_n)_{n\geq 1}$ of elements of $\mathcal{X}$ such that $y = \lim_{n\to\infty} \varphi_{\mathcal{Y}}(x_n)$ and define $\varphi(y) = \lim_{n\to\infty} \varphi_{\mathcal{Z}}(x_n)$. We claim that φ is an isometry. To see this, let $y, y' \in \mathcal{Y}$ be arbitrary. Choose sequence $(x_n)_{n\geq 1}$ and $(x'_n)_{n\geq 1}$ of elements of $\mathcal{X}$ such that $y = \lim_{n\to\infty} \varphi_{\mathcal{Y}}(x_n)$ and $y' = \lim_{n\to\infty} \varphi_{\mathcal{Y}}(x'_n)$. Then, by the triangle inequality and properties of limits,

$$\begin{aligned} d_{\mathcal{Z}}(\varphi(y), \varphi(y')) &= \lim_{n\to\infty} d_{\mathcal{Z}}(\varphi_{\mathcal{Z}}(x_n), \varphi_{\mathcal{Z}}(x'_n)) \\ &= \lim_{n\to\infty} d_{\mathcal{X}}(x_n, x'_n) \\ &= \lim_{n\to\infty} d_{\mathcal{Y}}(\varphi_{\mathcal{Y}}(x_n), \varphi_{\mathcal{Y}}(x'_n)) \\ &= d_{\mathcal{Y}}(y, y'). \end{aligned}$$

Hence φ is an isometry (and therefore injective).

To see that φ is surjective (and thus a bijection) let $z \in \mathcal{Z}$ be arbitrary. Note as $\mathcal{Z}$ is the completion of $\varphi_{\mathcal{Z}}(\mathcal{X})$, there exists a sequence $(x_n)_{n\geq 1}$ of elements of $\mathcal{X}$ such that $z = \lim_{n\to\infty} \varphi_{\mathcal{Z}}(x_n)$. By similar arguments to those above, $y = \lim_{n\to\infty} \varphi_{\mathcal{Y}}(x_n)$ exists and thus $\varphi(y) = z$. Hence, as $z \in \mathcal{Z}$ was arbitrary, φ is surjective. Hence $\mathcal{Y}$ and $\mathcal{Z}$ are isomorphic. ∎

The above demonstrates everything we could possibly want to know about completions for metric spaces. However, if we are dealing with normed linear spaces, we would like our maps to preserve the vector space structures. In particular, we would like our maps to be linear in order to obtain the appropriate notion of equality. Thus we make the following definitions.

Definition 4.2.26. Let $(\mathcal{X}, \|\cdot\|_{\mathcal{X}})$ and $(\mathcal{Y}, \|\cdot\|_{\mathcal{Y}})$ be normed linear spaces. A function $\varphi : \mathcal{X} \to \mathcal{Y}$ is said to be an *isometry* if φ is linear and $\|\varphi(\vec{x})\|_{\mathcal{Y}} = \|\vec{x}\|_{\mathcal{X}}$ for all $\vec{x} \in \mathcal{X}$.

If in addition to being an isometry φ is a bijection, it is said that φ is an *isomorphism*. Finally, its is said that $(\mathcal{X}, \|\cdot\|_{\mathcal{X}})$ and $(\mathcal{Y}, \|\cdot\|_{\mathcal{Y}})$ are *isomorphic* if there exists a isomorphism from $\mathcal{X}$ to $\mathcal{Y}$.

Note $\|\varphi(\vec{x})\|_{\mathcal{Y}} = \|\vec{x}\|_{\mathcal{X}}$ for all $\vec{x} \in \mathcal{X}$ along with the fact that φ is linear implies

$$\|\varphi(\vec{x}_1) - \varphi(\vec{x}_2)\|_{\mathcal{Y}} = \|\vec{x}_1 - \vec{x}_2\|_{\mathcal{X}}.$$

In particular, isometries for normed linear spaces are isometries for metric spaces.

Of course, when dealing with normed linear spaces, we would like our completions to behave well with respect to the vector space structures. Thus we provide an alternate and improved definition for the completion of a normed linear space.

Definition 4.2.27. Let $(\mathcal{X}, \|\cdot\|_{\mathcal{X}})$ be a normed linear space. A *completion* of $(\mathcal{X}, \|\cdot\|_{\mathcal{X}})$ is a Banach space $(\mathcal{Y}, \|\cdot\|_{\mathcal{Y}})$ such that there exists an isometry $\varphi : \mathcal{X} \to \mathcal{Y}$ such that $\overline{\varphi(\mathcal{X})} = \mathcal{Y}$.

Of course Corollary 4.2.24 demonstrates that every normed linear space has a completion as a metric space whereas Proposition 4.2.25 shows that there is only one possible completion for each metric space. As a normed linear space completion is a metric space completion, the completion in Corollary 4.2.24 is the only candidate for a normed linear space completion. However, it is not clear whether the function $\vec{z} \mapsto f_{\vec{z}}$ (where $f_{\vec{z}}(\vec{x}) = \|\vec{z} - \vec{x}\| - \|\vec{x} - \vec{a}\|$ for all $\vec{x} \in \mathcal{X}$ and $\vec{a} \in \mathcal{X}$ is fixed) is linear. Thus it is unclear that every normed linear space has a normed linear space completion.

It turns out that every normed linear space has a completion as a normed linear space. There are two methods we could take to proving this. The first is to take the metric space completion of a normed linear space $(\mathcal{X}, \|\cdot\|)$ and define a vector space structure on the completion via the vector space structure on $\mathcal{X}$. The difficulty then comes in definition the norm and verifying the definition does produce a norm.

We will proceed with an alternative description of the completion for normed linear spaces. This description uses an equivalence relation on Cauchy sequences (and is one way of constructing $\mathbb{R}$ from $\mathbb{Q}$).

Theorem 4.2.28. *Every normed linear space has a completion.*

Proof. Let $(\mathcal{X}, \|\cdot\|_{\mathcal{X}})$ be a normed linear space. Let $\mathcal{V}$ denote the set of all Cauchy sequences in $\mathcal{X}$. Note that $\mathcal{V}$ is non-empty as every constant sequence is Cauchy. In fact, the constant sequences will give us the embedding of $\mathcal{X}$ into its completion. Furthermore, since given Cauchy sequences $(\vec{x}_n)_{n\geq 1}$ and $(\vec{y}_n)_{n\geq 1}$ and $\alpha \in \mathbb{K}$, the sequences

$$(\vec{x}_n + \vec{y}_n)_{n\geq 1} \qquad \text{and} \qquad (\alpha \vec{x}_n)_{n\geq 1}$$

are Cauchy by the properties of the norm, $\mathcal{V}$ is a vector space over $\mathbb{K}$. However, $\mathcal{V}$ is not the normed linear space we want. To construct a normed linear space, we require a quotient.

Let

$$\mathcal{W} = \left\{ (\vec{x}_n)_{n\geq 1} \in V \mid \lim_{n\to\infty} \vec{x}_n = \vec{0} \right\}.$$

Clearly $\mathcal{W}$ is a subspace of $\mathcal{V}$. Recall an equivalence relation $\sim$ may be placed on $\mathcal{V}$ via $\vec{v}_1 \sim \vec{v}_2$ if and only if $\vec{v}_1 - \vec{v}_2 \in \mathcal{W}$. Furthermore, recall if $[\vec{v}]$ denotes the equivalence class of $\vec{v}$ and $\mathcal{V}/\mathcal{W}$ is the set of all equivalence classes, then $\mathcal{V}/\mathcal{W}$ is a vector space with operations $[\vec{v}_1] + [\vec{v}_2] = [\vec{v}_1 + \vec{v}_2]$ and $\alpha[\vec{v}] = [\alpha\vec{v}]$. In particular, two elements $(\vec{x}_n)_{n\geq 1}, (\vec{y}_n)_{n\geq 1} \in V$ produce the same element in $\mathcal{V}/\mathcal{W}$ if and only if

$$\lim_{n\to\infty} \|\vec{x}_n - \vec{y}_n\|_{\mathcal{X}} = 0$$

Define $\|\cdot\| : \mathcal{V}/\mathcal{W} \to [0,\infty)$ by

$$\|[(\vec{x}_n)_{n\geq 1}]\| = \limsup_{n\to\infty} \|\vec{x}_n\|_{\mathcal{X}}$$

and note that since $(\vec{x}_n)_{n\geq 1}$ is Cauchy and thus bounded by Lemma 4.1.4, $\|\cdot\|$ does indeed map into $[0,\infty)$. However, since we are dealing with equivalence classes, we must check that $\|\cdot\|$ is well-defined. To see this, notice if $[(\vec{x}_n)_{n\geq 1}] = [(\vec{y}_n)_{n\geq 1}]$ then

$$\lim_{n\to\infty} \|\vec{x}_n - \vec{y}_n\|_{\mathcal{X}} = 0.$$

so

$$\limsup_{n\to\infty} \|\vec{x}_n\|_{\mathcal{X}} = \limsup_{n\to\infty} \|\vec{y}_n\|_{\mathcal{X}}$$

by the reverse triangle inequality. Hence $\|\cdot\|$ is well-defined. To see that $\|\cdot\|$ is indeed a norm, note that $\|[(\vec{x}_n)_{n\geq 1}]\| = 0$ if and only if $\limsup_{n\to\infty} \|\vec{x}_n\|_{\mathcal{X}} = 0$ if and only if $(\vec{x}_n)_{n\geq 1} \in \mathcal{W}$ if and only if $[(\vec{x}_n)_{n\geq 1}] = \vec{0}_{\mathcal{V}/\mathcal{W}}$. As the other properties from Definition 1.2.5 are trivial to verify, $(\mathcal{V}/\mathcal{W}, \|\cdot\|)$ is a normed linear space.

We will postpone the proof that $(\mathcal{V}/\mathcal{W}, \|\cdot\|)$ is complete momentarily in order to demonstrate some facts in relation to $\mathcal{X}$. Define $\varphi : \mathcal{X} \to \mathcal{V}/\mathcal{W}$ by $\varphi(\vec{x}) = [(\vec{x})_{n\geq 1}]$; that is, map each element of $\mathcal{X}$ to a constant sequence. Clearly φ is well-defined, linear, and an isometry. We claim that $\varphi(\mathcal{X})$ is dense in $\mathcal{V}/\mathcal{W}$.

To see that $\overline{\varphi(\mathcal{X})} = \mathcal{V}/\mathcal{W}$, let $[(\vec{x}_n)_{n\geq 1}] \in \mathcal{V}/\mathcal{W}$ be arbitrary and let $\epsilon > 0$ be arbitrary. Since $(\vec{x}_n)_{n\geq 1}$ is Cauchy in $\mathcal{X}$, there exists an $N \in \mathbb{N}$ such that $\|\vec{x}_n - \vec{x}_m\|_{\mathcal{X}} < \epsilon$ for all $n, m \geq N$. Hence

$$\|\varphi(\vec{x}_N) - [(\vec{x}_n)_{n\geq 1}]\| \leq \epsilon$$

by the definition of $\|\cdot\|$. Therefore, as $\epsilon > 0$ was arbitrary, $[(\vec{x}_n)_{n\geq 1}]$ is in the closure of $\varphi(\mathcal{X})$. Therefore, as $[(\vec{x}_n)_{n\geq 1}] \in \mathcal{V}/\mathcal{W}$ was arbitrary, $\overline{\varphi(\mathcal{X})} = \mathcal{V}/\mathcal{W}$.

To see that $(\mathcal{V}/\mathcal{W}, \|\cdot\|)$ is complete, let $(\vec{z}_n)_{n\geq 1}$ be an arbitrary Cauchy sequence in $(\mathcal{V}/\mathcal{W}, \|\cdot\|)$. Since $\overline{\varphi(\mathcal{X})} = \mathcal{V}/\mathcal{W}$, for each $n \in \mathbb{N}$ there exists an $\vec{x}_n \in \mathcal{X}$ such that

$$\|\varphi(\vec{x}_n) - \vec{z}_n\| < \frac{1}{n}.$$

We claim that $(\vec{x}_n)_{n\geq 1}$ is a Cauchy sequence of elements of $\mathcal{X}$ and thus is an element of $\mathcal{V}$. To see this, notice for all $n, m \in \mathbb{N}$ that

$$\begin{aligned}
\|\vec{x}_n - \vec{x}_m\|_{\mathcal{X}} &= \|\varphi(\vec{x}_n) - \varphi(\vec{x}_m)\| \\
&\leq \|\varphi(\vec{x}_n) - \vec{z}_n\| + \|\vec{z}_n - \vec{z}_m\| + \|\vec{z}_m - \varphi(\vec{x}_m)\| \\
&\leq \frac{1}{n} + \frac{1}{m} + \|\vec{z}_n - \vec{z}_m\|.
\end{aligned}$$

Therefore, as $(\vec{z}_n)_{n\geq 1}$ is Cauchy, it is elementary to verify the above inequality implies $(\vec{x}_n)_{n\geq 1}$ is Cauchy. Finally, to see that $(\vec{z}_n)_{n\geq 1}$ converges to $\vec{z} = [(\vec{x}_n)_{n\geq 1}]$, we notice that

$$\lim_{n\to\infty} \|\varphi(\vec{x}_n) - \vec{z}\| = 0$$

as $(\vec{x}_n)_{n\geq 1}$ is Cauchy. Hence as

$$\|\vec{z}_n - \vec{z}\| \leq \|\vec{z}_n - \varphi(\vec{x}_n)\| + \|\varphi(\vec{x}_n) - \vec{z}\| \leq \frac{1}{n} + \|\varphi(\vec{x}_n) - \vec{z}\|,$$

we obtain that $(\vec{z}_n)_{n\geq 1}$ converges to $\vec{z} = [(\vec{x}_n)_{n\geq 1}]$. Therefore, as $(\vec{z}_n)_{n\geq 1}$ was an arbitrary Cauchy sequence, $\mathcal{V}/\mathcal{W}$ is complete thereby completing the proof. ∎

4.3 Compact Metric Spaces

With the discussion of complete metric spaces complete, we turn our attention to our main goal: understanding compact metric spaces. Of course, we saw in Theorem 4.1.12 that every compact metric space must be complete. Thus we endeavour to see how close the converse is to being true. Of course, there are several examples of complete metric spaces that are not compact, such as $\mathbb{R}$. Consequently, there must be another property on metric spaces related to compactness. If one believes that being complete is an analogue of being closed for metric spaces, then, based on the Heine-Borel Theorem, we are looking for a property related to boundedness. However, this property is not just 'be bounded' as we know the discrete metric on an infinite set produces a complete (as every Cauchy sequence is eventually constant) bounded metric space that is not compact.

In order to arrive at the correct notion of boundedness required for a metric space to be compact, we turn our attention to another problem. Recall that Theorem 3.2.2 implies that any net in a compact topological space has a convergent subnet. Thus this must be true of any compact metric space. As it is enough to discuss convergent sequences in metric spaces by Theorem 1.5.28, we desire an analogue for sequences. As a subnet of a sequence need not be a sequence, such a result need not be immediately implied by Theorem 1.5.28. Consequently, we give a name to the property we desire to be equivalent to compactness for a metric space in order to study it.

Definition 4.3.1. A topological space $(X, \mathcal{T})$ is said to be *sequentially compact* if every sequence in $(X, \mathcal{T})$ has a convergent subsequence.

Recall in the realm of metric spaces that sequences suffice to determine the topology. Thus our goal based on Theorem 3.2.2 is to show the notions of compactness and sequential compactness agree. To begin to demonstrate this fact and to provide our first examples of sequentially compact spaces, we note that compact metric spaces are sequentially compact. To show this, we quickly adapt one portion of the proof of Theorem 3.2.2 to the metric setting.

Theorem 4.3.2. *Every compact metric space is sequentially compact.*

Proof. Let (X, d) be a compact metric space. To see that X is sequentially compact, let $(x_n)_{n\geq 1}$ be an arbitrary sequence of elements of X. For each $n \in \mathbb{N}$, let

$$F_n = \overline{\{x_k \mid k \geq n\}}.$$

Therefore $\{F_n\}_{n\in\mathbb{N}}$ is a collection of closed subsets of X which has the finite intersection property since if $n_1, \ldots, n_q \in \mathbb{N}$, then

$$\bigcap_{k=1}^{q} F_{n_q} = F_{\max\{n_1,\ldots,n_q\}}.$$

Hence, as (X, d) is compact, Theorem 1.5.28 implies $\bigcap_{n=1}^{\infty} F_n \neq \emptyset$.

Let $x_0 \in \bigcap_{n=1}^{\infty} F_n$ be arbitrary. We claim there exists a subsequence of $(x_n)_{n\geq 1}$ that converges to x_0. To see this, notice since $x \in F_1$ that there exists an $k_1 \in \mathbb{N}$ such that $d(x_0, x_{k_1}) \leq 1$ by Theorem 1.6.21 and the fact that the open balls centred at x_0 is a neighbourhood basis for x_0. Hence since $x_0 \in F_{k_1+1}$, there exists an $k_2 > k_1$ such that $d(x_0, x_{k_2}) \leq \frac{1}{2}$ by Theorem 1.6.21. By repeating this process ad infinitum, there exists an increasing sequence $(k_n)_{n\geq 1}$ of natural numbers such that $d(x_0, x_{k_n}) \leq \frac{1}{n}$. Hence $(x_{k_n})_{n\geq 1}$ is a subsequence of $(x_n)_{n\geq 1}$ that converges to x_0. Therefore, as $(x_n)_{n\geq 1}$ was arbitrary, (X, d) is sequentially compact. ∎

As we know that nets are required to determine the topology for a general topological space, it is unsurprising that the notions of compactness and sequential compactness differ for general topological spaces. The following two examples illustrate this fact.

Example 4.3.3. Let S denote the set of all increasing sequence of natural numbers, let $Y = \{0, 1\}$ equipped with the discrete topology, and let $X = \prod_{s\in S} Y$. As Y is compact, Tychonoff's Theorem (Theorem 3.3.4) implies X is compact when equipped with its product topology. However, X is not sequentially compact. To see this, consider the sequence $(x_n)_{n\geq 1}$ in X defined as follow: for a fixed $n \in \mathbb{N}$ and $s = (n_k)_{k\geq 1} \in S$, define

$$x_n(s) = \begin{cases} 0 & \text{if } n = n_k \text{ for some even } k \in \mathbb{N} \\ 1 & \text{otherwise} \end{cases}.$$

We claim that $(x_n)_{n\geq 1}$ has no convergent subsequences thereby showing that X is not sequentially compact. To see this, suppose $(x_{n_k})_{k\geq 1}$ is a subsequence of $(x_n)_{n\geq 1}$ that converges to some $x \in X$. Hence $s = (n_k)_{k\geq 1} \in S$. Therefore, by Theorem 1.5.25, it must be the case that $(x_{n_k}(s))_{k\geq 1}$ converges to $x(s)$ in Y. However

$$x_{n_k}(s) = \begin{cases} 0 & \text{if } k \text{ is even} \\ 1 & \text{if } k \text{ is odd} \end{cases}.$$

As a sequence in the discrete topology on Y converges if and only if the sequence is eventually constant, we see that $(x_{n_k}(s))_{k\geq 1}$ cannot possibly converge to $x(s)$ thereby yielding a contradiction. Hence X is compact but not sequentially compact.

For our second example, we encapsulate a useful fact via the following lemma.

Lemma 4.3.4. *Let $\{(X_n, d_n)\}_{n\in\mathbb{N}}$ be a countable set of metric spaces and let $X = \prod_{n\in\mathbb{N}} X_n$. Define $d_p : X \times X \to [0,\infty)$ by*

$$d_p\left((x_n)_{n\geq 1}, (y_n)_{n\geq 1}\right) = \sum_{n=1}^{\infty} \frac{d_n(x_n, y_n)}{2^n(1 + d_n(x_n, y_n))}$$

for all $(x_n)_{n\geq 1}, (y_n)_{n\geq 1} \in X$. Then d_p is a metric on X. Furthermore, a net $(x_\lambda)_{\lambda\in\Lambda}$ converges to a point $x \in X$ with respect to d_p if and only if $(x_\lambda)_{\lambda\in\Lambda}$ converges to x in the product topology. Hence the metric and product topologies on X agree by Theorem 1.5.27.

Proof. First, since $\frac{a}{1+a} \in [0,1)$ for all $a \in [0,\infty)$ and $\sum_{n=1}^{\infty} \frac{1}{2^n} = 1$, we see that d_p is a well-defined map from $X \times X$ to $[0,1)$. Furthermore, $d_p((x_n)_{n\geq 1}, (y_n)_{n\geq 1}) = 0$ if and only if

$$\frac{d_n(x_n, y_n)}{2^n(1 + d_n(x_n, y_n))} = 0$$

for all $n \in \mathbb{N}$ if and only if $d_n(x_n, y_n) = 0$ for all $n \in \mathbb{N}$ if and only if $x_n = y_n$ for all $n \in \mathbb{N}$ if and only if $(x_n)_{n\geq 1} = (y_n)_{n\geq 1}$. Hence d_p satisfies the first property of a metric. Similarly, since clearly

$$d_p((x_n)_{n\geq 1}, (y_n)_{n\geq 1}) = d_p((y_n)_{n\geq 1}, (x_n)_{n\geq 1})$$

for all $(x_n)_{n\geq 1}, (y_n)_{n\geq 1} \in X$, d_p satisfies the second property of a metric. Since

$$\frac{d_n(x_n, y_n)}{1 + d_n(x_n, y_n)} \leq \frac{d_n(x_n, z_n)}{1 + d_n(x_n, z_n)} + \frac{d_n(z_n, y_n)}{1 + d_n(z_n, y_n)}$$

for all $n \in \mathbb{N}$ and $(x_n)_{n\geq 1}, (y_n)_{n\geq 1}, (z_n)_{n\geq 1} \in X$ by the same arguments used in Example 3.1.28, we see that d_p is a metric on X as desired.

To see the second claim, it suffices by Theorem 1.5.25 to prove that a net $(x_\lambda)_{\lambda\in\Lambda}$ converges to a point $x \in X$ with respect to d_p if and only if $(x_\lambda(n))_{\lambda\in\Lambda}$ converges to $x(n)$ in (X_n, d_n) for all $n \in \mathbb{N}$.

To begin, suppose that $(x_\lambda)_{\lambda\in\Lambda}$ is an arbitrary net that converges to a point $x \in X$ with respect to d_p. For a fixed $n \in \mathbb{N}$, notice that

$$0 \leq \frac{d_n(x_\lambda(n), x(n))}{2^n(1 + d_n(x_\lambda(n), x(n)))} \leq d_p(x_\lambda, x).$$

Hence, as $(x_\lambda)_{\lambda\in\Lambda}$ converges to x with respect to d_p,

$$\lim_{\lambda\in\Lambda} d_p(x_\lambda, x) = 0$$

so

$$\lim_{\lambda\in\Lambda} \frac{d_n(x_\lambda(n), x(n))}{2^n(1 + d_n(x_\lambda(n), x(n)))} = 0$$

and thus

$$\lim_{\lambda\in\Lambda} d_n(x_\lambda(n), x(n)) = 0.$$

Hence, as $n \in \mathbb{N}$ was arbitrary, $(x_\lambda(n))_{\lambda\in\Lambda}$ converges to $x(n)$ in (X_n, d_n) for all $n \in \mathbb{N}$. Thus as $(x_\lambda)_{\lambda\in\Lambda}$ was arbitrary, the first direction is complete.

Conversely, suppose that $(x_\lambda)_{\lambda\in\Lambda}$ is a net in X and $x \in X$ are such that $(x_\lambda(n))_{\lambda\in\Lambda}$ converges to $x(n)$ in (X_n, d_n) for all $n \in \mathbb{N}$. To see that $(x_\lambda)_{\lambda\in\Lambda}$ converges to a point $x \in X$ with respect to d_p, let $\epsilon > 0$ be arbitrary. Since

$$\sum_{n=1}^{\infty} \frac{1}{2^n} < \infty,$$

there exists an $N \in \mathbb{N}$ such that

$$\sum_{n=N+1}^{\infty} \frac{1}{2^n} < \frac{\epsilon}{2}.$$

Furthermore, since $(x_\lambda(n))_{\lambda\in\Lambda}$ converges to $x(n)$ in (X_n, d_n) for each $n \in \mathbb{N}$ and since $x \mapsto \frac{x}{1+x}$ is a continuous function on $[0, \infty)$ that vanishes at 0, for each $n \in \{1, \ldots, N\}$ there exists a $\lambda_n \in \Lambda$ such that if $\lambda \geq \lambda_n$ then

$$\frac{d_n(x_\lambda(n), x(n))}{2^n(1 + d_n(x_\lambda(n), x(n)))} < \frac{\epsilon}{2N}.$$

Due to the existence of least upper bounds for nets, there exists a $\lambda_0 \in \Lambda$ such that $\lambda_0 \geq \lambda_n$ for all $n \in \{1, \ldots, N\}$. Hence for all $\lambda \geq \lambda_0$ we have that

$$\begin{aligned} d_p(x_\lambda, x) &= \sum_{n=1}^{N} \frac{d_n(x_\lambda(n), x(n))}{2^n(1 + d_n(x_\lambda(n), x(n)))} + \sum_{n=N+1}^{\infty} \frac{d_n(x_\lambda(n), x(n))}{2^n(1 + d_n(x_\lambda(n), x(n)))} \\ &\leq \sum_{n=1}^{N} \frac{\epsilon}{2N} + \sum_{n=1}^{\infty} \frac{1}{2^n} \\ &< \frac{\epsilon}{2} + \frac{\epsilon}{2} = \epsilon. \end{aligned}$$

Hence, as ϵ was arbitrary, $(x_\lambda)_{\lambda\in\Lambda}$ converges to x as desired. ∎

Example 4.3.5. Let $Y = [0, 1]$ equipped with the subspace topology inherited from the canonical topology on $\mathbb{R}$ and let $X = \prod_{x\in\mathbb{R}} Y = \mathcal{F}(\mathbb{R}, [0, 1])$. As Y is compact by the Heine-Borel Theorem (Theorem 3.1.25), Tychonoff's Theorem (Theorem 3.3.4) implies X is compact when equipped with its product topology. Consider the subspace

$$Z = \{f \in \mathcal{F}(\mathbb{R}, [0, 1]) \mid \{x \in \mathbb{R} \mid f(x) \neq 0\} \text{ is countable}\}.$$

Then we claim Z is a sequentially compact topological space that is not compact. To see that Z is not compact, suppose to the contrary that Z is compact. Therefore, since Y is Hausdorff so X is Hausdorff by Example 1.5.39, Z is a compact subspace of a Hausdorff space and thus closed in X by Theorem 3.1.13. To obtain a contradiction, we will use Theorem 1.6.14 to show that Z is not closed in X.

To begin, let

$$\Lambda = \{F \subseteq \mathbb{R} \mid F \text{ finite}\}.$$

For $F_1, F_2 \in \Lambda$, define $F_1 \leq F_2$ if and only if $F_1 \subseteq F_2$. Then $(\Lambda, \leq)$ is a directed set by Example 1.5.7. For each $F \in \Lambda$, let $f_F \in \mathcal{F}(\mathbb{R}, [0, 1])$ be defined by

$$f_F(x) = \begin{cases} 1 & \text{if } x \in F \\ 0 & \text{otherwise} \end{cases}.$$

Hence $(f_F)_{F\in\Lambda}$ is a net in Z by definition.

Let $f_0 = (1)_{x\in\mathbb{R}} \in X \setminus Z$. We claim that $(f_F)_{F\in\Lambda}$ converges f_0 in X thereby completing our contradiction. To see this, notice if $x \in X$ then if $F_x = \{x\}$ and $F \geq F$ then $f_F(x) = 1 = f_0(x)$. Hence $(f_F(x))_{F\in\Lambda}$ converges to $f_0(x)$ for all $x \in X$ so $(f_F)_{F\in\Lambda}$ converges f_0 in X by Theorem 1.5.25. Thus Z is not compact.

To see that Z is sequentially compact, let $(f_n)_{n\geq1}$ be an arbitrary sequence in Z. Let

$$G = \{x \in \mathbb{R} \mid f_n(x) \neq 0 \text{ for some } n \in \mathbb{N}\}.$$

Hence, as the countable union of countable sets is countable, G is countable by the definition of Z. Let

$$Z' = \{f \in \mathcal{F}(\mathbb{R}, [0, 1]) \mid \{x \in \mathbb{R} \mid f(x) \neq 0\} \subseteq G\}$$

so that $(f_n)_{n\geq1}$ is a sequence in the subspace Z' of Z.

Consider the map $\Phi : Z' \to \prod_{x\in G}[0, 1]$ defined by

$$\Phi(f) = f|_G.$$

Clearly Φ is a bijection. Furthermore, by Theorem 1.5.25, it is elementary to see that a net $(h_\lambda)_{\lambda\in\Lambda}$ in Z' converges to an element $h \in Z'$ if and only if $(\Phi(h_\lambda))_{\lambda\in\Lambda}$ in converges $\Phi(h)$ when $\prod_{x\in G}[0, 1]$ is equipped with the product topology. Hence Φ is a homeomorphism.

Consider the sequence $(\Phi(f_n))_{n\geq 1}$ in $\prod_{x\in G}[0,1]$. Since the product topology on $\prod_{x\in G}[0,1]$ is compact by Tychonoff's Theorem (Theorem 3.3.4) and induced by the metric d_p by Lemma 4.3.4 as G is countable, $\prod_{x\in G}[0,1]$ is sequentially compact. Thus there exists a subsequence $(\Phi(f_{n_k}))_{k\geq 1}$ of $(\Phi(f_n))_{n\geq 1}$ that converges to some element $y \in \prod_{x\in G}[0,1]$. Thus, as Φ is a homeomorphism, $(f_{n_k})_{k\geq 1}$ is a subsequence of $(f_n)_{n\geq 1}$ that converges to $\Phi^{-1}(y)$ in Z' and thus in Z. Thus, as $(f_n)_{n\geq 1}$ was an arbitrary sequence in Z, Z is a sequentially compact topological space that is not compact.

Of course, our goal now is to prove the converse of Theorem 4.3.2 holds for metric spaces. To do so leads us again towards developing the correct notion of boundedness to characterize compactness. In order to develop this correct notion, we note that we can always cover a metric space with open balls of a certain radius. Consequently, if a metric space (X, d) is compact there must be a finite cover of (X, d) using open balls of a specific radii. This causes us to define the following two terms.

Definition 4.3.6. Let (X, d) be a metric space and let $\epsilon > 0$. A subset $\{x_\alpha\}_{\alpha\in I} \subseteq X$ is said to be an *ϵ-net* of (X, d) if $X = \bigcup_{\alpha\in I} B_d(x_\alpha, \epsilon)$; that is, for all $x \in X$ there exists an $\alpha \in I$ such that $d(x, x_\alpha) < \epsilon$.

Definition 4.3.7. A metric space (X, d) is said to be *totally bounded* if (X, d) has a finite ϵ-net for all $\epsilon > 0$; that is, for each $\epsilon > 0$ there exists $x_1, \ldots, x_n \in X$ such that $X = \bigcup_{k=1}^n B_d(x_k, \epsilon)$.

Remark 4.3.8. It may be very tempting to claim that every totally bounded metric space (X, d) is automatically compact as if one has an open cover of (X, d) then we would hope that there is an ϵ-net of X where each ball is contained in a single element of the open cover. However, this argument clearly has the flaw in the statement 'where each ball is contained in a single element of the open cover'. Indeed consider $X = (0, 1)$ as a subspace of $\mathbb{R}$ equipped with its canonical metric. Then $(0, 1)$ is totally bounded. Indeed for every $\epsilon > 0$, choose $n \in \mathbb{N}$ with $\frac{1}{n} < \epsilon$ and consider the set $\{\frac{k}{n}\}_{k=1}^n$, which is clearly an ϵ-net of $(0, 1)$. Hence $(0, 1)$ is totally bounded but not compact by the Heine-Borel Theorem (Theorem 3.1.25) as $(0, 1)$ is not closed. Moreover, a similar argument can be used to show that any bounded subset of $\mathbb{K}^n$ is totally bounded.

In theory, checking a metric space (X, d) is totally bounded is easier than it is to check (X, d) is compact using the definitions of open covers. Indeed it quite difficult to describe all open covers of a metric space and determine whether each open cover has a finite subcover. However, checking a metric space has a finite ϵ-net for every $\epsilon > 0$ is often not too difficult as one need to simply find a correct set of points in the metric space for a given $\epsilon > 0$.

Thus our goal is to connect the notions of totally boundedness and compactness in metric spaces. To do so, we begin by developing the properties of totally bounded metric spaces as the most instructive examples will following once we have the connection between compactness and total boundedness in metric spaces. Note only some of these properties will be used in this section whereas others will be used in the next section.

Proposition 4.3.9. *Every sequentially compact metric space is totally bounded. Consequently compact metric spaces are totally bounded by Theorem 4.3.2.*

Proof. Let (X, d) be a sequentially compact metric space. To see that X is totally bounded, suppose to the contrary that there exists an $\epsilon > 0$ such that X does not have a finite ϵ-net. Let $x_1 \in X$ be arbitrary. Since $\{x_1\}$ is not an ϵ-net, there exists an $x_2 \in X \setminus B_d(x_1, \epsilon)$. Since $\{x_1, x_2\}$ is not an ϵ-net, there exists an $x_3 \in X \setminus (B_d(x_1, \epsilon) \cup B_d(x_2, \epsilon))$; that is, $d(x_3, x_j) \geq \epsilon$ for all $j \in \{1, 2\}$. By repeating this process ad infinitum, there exists a sequence $(x_n)_{n\geq 1}$ such that $d(x_n, x_m) \geq \epsilon$ for all $n, m \in \mathbb{N}$ with $n \neq m$. Clearly the sequence $(x_n)_{n\geq 1}$ does not have any convergent subsequences since it does not have any Cauchy subsequences. Hence X cannot be sequentially compact, which is a contradiction. Hence sequentially compact metric spaces are totally bounded. ∎

Proposition 4.3.10. *Every totally bounded metric space is bounded.*

Proof. Let (X, d) be a totally bounded metric space. Since (X, d) is totally bounded, there exists a finite 1-net $\{x_1, \ldots, x_n\}$ for (X, d). Let

$$M = \max(\{1 + d(x_k, x_1) \mid k, j \in \{1, \ldots, n\}\}).$$

We claim that $d(x, x_1) \leq M$ for all $x \in X$ which implies X is bounded by Lemma 3.1.21. To see this, let $x \in X$ be arbitrary. Since $\{x_1, \ldots, x_n\}$ is a 1-net for (X, d), there exists a $k \in \{1, \ldots, n\}$ such that $d(x, x_k) < 1$. Hence

$$d(x, x_1) \leq d(x, x_k) + d(x_k, x_1) < 1 + d(x_k, x_1) \leq M,$$

as desired. Therefore, since $x \in X$ was arbitrary, (X, d) is bounded as desired. ∎

Proposition 4.3.11. *Let (X, d) be a metric space and let $A \subseteq X$. If (X, d) is totally bounded, then $(A, d|_A)$ is totally bounded.*

Proof. The caveat of this proof is that the elements of each ϵ-net for A must come from A and, a priori, they only come from X.

To see that $(A, d|_A)$ is totally bounded, let $\epsilon > 0$ be arbitrary. Since (X, d) is totally bounded, there exists a finite $\frac{\epsilon}{2}$-net $\{x_k\}_{k=1}^n$ of (X, d). Hence

$$A \subseteq \bigcup_{k=1}^n B_d\left(x_k, \frac{\epsilon}{2}\right).$$

Let $I \subseteq \{1, \ldots, n\}$ consist of all indices k such that $A \cap B_d\left(x_k, \frac{\epsilon}{2}\right) \neq \emptyset$. For each $k \in I$, choose $a_k \in A \cap B_d\left(x_k, \frac{\epsilon}{2}\right)$. We claim that $\{a_k\}_{k\in I}$ is an ϵ-net for A. To see this, note the claim is trivial if $A = \emptyset$. Otherwise, let $a \in A$ be arbitrary. Therefore there exists a $k_0 \in \{1, \ldots, n\}$ such that $a \in B_d\left(x_{k_0}, \frac{\epsilon}{2}\right)$. Hence $k_0 \in I$ and

$$d(a, a_{k_0}) \leq d(a, x_{k_0}) + d(a_{k_0}, x_{k_0}) < \frac{\epsilon}{2} + \frac{\epsilon}{2} = \epsilon$$

as $a, a_{k_0} \in B_d\left(x_{k_0}, \frac{\epsilon}{2}\right)$, Therefore, as $a \in A$ was arbitrary, $\{a_k\}_{k\in I}$ is an ϵ-net for A by definition. Hence as $\epsilon > 0$ was arbitrary, $(A, d|_A)$ is totally bounded. ∎

Proposition 4.3.12. *Let (X, d) be a metric space and let $A \subseteq X$. If $(A, d|_A)$ is totally bounded, then $(\overline{A}, d|_{\overline{A}})$ is totally bounded.*

Proof. Let $\epsilon > 0$ be arbitrary. Since $(A, d|_A)$ is totally bounded, there exists a finite $\frac{\epsilon}{2}$-net $\{a_k\}_{k=1}^n$ for A. Hence $\{a_k\}_{k=1}^n \subseteq \overline{A}$. We claim that $\{a_k\}_{k=1}^n$ is an ϵ-net for $\overline{A}$. To see this, let $x \in \overline{A}$ be arbitrary. By the fact that open balls form a neighbourhood basis for each point in a metric space, Theorem 1.6.21 implies there exists an $a \in A$ such that $d(x, a) < \frac{\epsilon}{2}$. As $\{a_k\}_{k=1}^n$ is an $\frac{\epsilon}{2}$-net for A, there exists a $k \in \{1, \ldots, n\}$ such that $d(a, a_k) < \frac{\epsilon}{2}$. Hence $d(x, a_k) < \epsilon$ by the Triangle Inequality. Therefore, as $x \in \overline{A}$ was arbitrary, $\{a_k\}_{k=1}^n$ is an ϵ-net for $\overline{A}$. Since $\epsilon > 0$ was arbitrary, $\overline{A}$ is totally bounded by definition. ∎

If compactness and sequential compactness are the same notion in metric spaces, the following lemma should not be surprising.

Lemma 4.3.13. *Let (X, d_X) and (Y, d_Y) be metric spaces and let $f : X \to Y$ be continuous. If (X, d_X) is sequentially compact, then $f(X)$ is sequentially compact subspace of (Y, d_Y). Consequently, if $Y = \mathbb{R}$ with the canonical metric, there exists $x_1, x_2 \in X$ such that $f(x_1) \leq f(x) \leq f(x_2)$ for all $x \in X$.*

Proof. To see that $f(X)$ is sequentially compact subspace of (Y, d_Y), let $(y_n)_{n\geq1}$ be an arbitrary sequence of elements of $f(X)$. Hence there exists a sequence $(x_n)_{n\geq1}$ such that $y_n = f(x_n)$ for all $n \in \mathbb{N}$. Since (X, d_X) is sequentially compact, there exists a subsequence $(x_{k_n})_{n\geq1}$ that converges in (X, d_X) to some element $x \in X$. As f is continuous, $(y_{k_n})_{n\geq1}$ converges to $f(x)$ in $f(X)$. Therefore, as $(y_n)_{n\geq1}$ was arbitrary, $f(X)$ is sequentially compact by definition.

To see the later claim, suppose $Y = \mathbb{R}$. Since $f(X)$ is sequentially compact, $f(X)$ is totally bounded by Proposition 4.3.9 and thus bounded by Proposition 4.3.10. Hence $\inf(f(X))$ and $\sup(f(X))$ are finite. Since $f(X)$ is sequentially compact, the limits of any convergent sequences with elements in $f(X)$ must be elements of $f(X)$. As we may construct sequences of elements of $f(X)$ converging to $\inf(f(X))$ and $\sup(f(X))$ and respectively, we obtain that $\sup(f(X)), \inf(f(X)) \in f(X)$. Hence there exists $x_1, x_2 \in X$ such that $f(x_1) = \inf(f(X))$ and $f(x_2) = \sup(f(X))$ so $f(x_1) \leq f(x) \leq f(x_2)$ for all $x \in X$ as desired. ∎

Our next lemma is the last key ingredient we need in order to verify the equivalence of the notions of compactness and sequential compactness in metric spaces. In particular, this lemma allows us to choose 'large' open balls inside any open cover. Consequently, once we demonstrate the notions of compact and sequentially compact sets are the same, we may apply the following lemma to any open cover of a compact metric space.

Lemma 4.3.14. *Let (X, d) be a sequentially compact metric space. If $\{U_\alpha\}_{\alpha\in I}$ is an open cover of (X, d), then there exists an $\epsilon_0 > 0$ (called the* Lebesgue number *for $\{U_\alpha\}_{\alpha\in I}$) such that for any $0 < \delta < \epsilon_0$ and any $x \in X$ there exists an $\alpha_x \in I$ such that $B_d(x, \delta) \subseteq U_{\alpha_x}$.*

Proof. To begin, note (X, d) is totally bounded by Proposition 4.3.9 and thus bounded by Proposition 4.3.10. Hence there exists an $x_0 \in X$ and an $R > 0$ such that $B_d(x_0, R) = X$ by Lemma 3.1.21. Hence for any $x \in X$, $B_d(x, 2R) = X$ by the triangle inequality.

Fix an open cover $\{U_\alpha\}_{\alpha\in I}$ of X and consider a function $\varphi : X \to \mathbb{R}$ defined by

$$\varphi(x) = \sup\{r \in \mathbb{R} \mid r \leq 2R, B_d(x, r) \subseteq U_\alpha \text{ for some } \alpha \in I\}$$

for all $x \in X$. Clearly φ is well-defined. Furthermore, we claim that $\varphi(x) > 0$ for all $x \in X$. To see this, notice if $x \in X$ then $x \in \bigcup_{\alpha\in I} U_\alpha$. Hence there exists an $\alpha_x \in I$ such that $x \in U_{\alpha_x}$. Since U_{α_x} is open, there exists an $r > 0$ such that $B(x, r) \subseteq U_{\alpha_x}$ and thus $\varphi(x) > r$.

We claim that φ is continuous. To see this, let $x, y \in X$ be arbitrary. By definition of φ, for all $r < \varphi(x)$ there exists an $\alpha \in I$ such that $B_d(x, r) \subseteq U_\alpha$. If $r \in [0, 2R]$ is such that $r < \varphi(x)$, then if $r - d(x, y) > 0$ we must have $B_d(y, r - d(x, y)) \subseteq U_\alpha$ by the triangle inequality so $\varphi(y) \geq r - d(x, y)$. Otherwise, if $r - d(x, y) \leq 0$ then clearly $\varphi(y) \geq r - d(x, y)$. In either case, $\varphi(y) \geq r - d(x, y)$ for all $r < \varphi(x)$ so $\varphi(y) \geq \varphi(x) - d(x, y)$. By replacing the roles of x and y, we see that

$$|\varphi(x) - \varphi(y)| \leq d(x, y).$$

Therefore, as $x, y \in X$ were arbitrary, φ is clearly continuous.

Since (X, d) is sequentially compact, Lemma 4.3.13 implies there exists an $x_0 \in X$ such that $\varphi(x_0) \leq \varphi(x)$ for all $x \in X$. Hence if $\epsilon_0 = \varphi(x_0)$, then $\epsilon_0 > 0$. Furthermore, for all $0 < \delta < \epsilon_0$ and $x \in X$ we see that $\delta < \varphi(x)$ so by the definition of φ there exists an $\alpha_x \in I$ with $B_d(x, \delta) \subseteq U_{\alpha_x}$ as desired. ∎

Using Lemma 4.3.14, we obtain the equivalence of the compactness and sequential compactness in metric spaces.

Theorem 4.3.15 (Borel-Lebesgue Theorem). *A metric space is compact if and only if it is sequentially compact.*

Proof. As compact metric spaces are sequentially compact by Theorem 4.3.2, one direction is complete.

For the other direction, suppose (X, d) is a sequentially compact metric space. To see that (X, d) is compact, let $\{U_\alpha\}_{\alpha \in I}$ be an arbitrary open cover of (X, d). As (X, d) is sequentially compact, Lemma 4.3.14 implies there exists an $\epsilon_0 > 0$ such that for any $0 < \delta < \epsilon_0$ and any $x \in X$ there exists an $\alpha_x \in I$ such that $B_d(x, \delta) \subseteq U_{\alpha_x}$.

Since (X, d) is sequentially compact, (X, d) is totally bounded by Proposition 4.3.9. Hence there exists a finite $\frac{\epsilon_0}{2}$-net $\{x_k\}_{k=1}^n$ for (X, d). Hence

$$X = \bigcup_{k=1}^n B_d\left(x_k, \frac{\epsilon_0}{2}\right)$$

By the above paragraph there exists $\alpha_1, \ldots, \alpha_n \in I$ such that $B_d\left(x_k, \frac{\epsilon_0}{2}\right) \subseteq U_{\alpha_k}$ for all $k \in \{1, \ldots, n\}$. Hence

$$X = \bigcup_{k=1}^n U_{\alpha_k}$$

so $\{U_{\alpha_1}, \ldots, U_{\alpha_n}\}$ is a finite subcover of (X, d). Therefore, as $\{U_\alpha\}_{\alpha \in I}$ was arbitrary, (X, d) is compact. ∎

With the connection between compactness and sequential compactness completed by the Borel-Lebesgue Theorem (Theorem 4.3.15) we can obtain the full version of the Heine-Borel Theorem (Theorem 3.1.25) where 'closed' is replaced by 'complete' and 'bounded' is replaced by 'totally bounded'.

Theorem 4.3.16. *Let (X, d) be a metric space. Then the following are equivalent:*

(i) (X, d) is compact.

(ii) (X, d) is complete and totally bounded.

Proof. First, suppose (X, d) is compact. Hence (X, d) is complete by Theorem 4.1.12. Furthermore, since (X, d) is sequentially compact by Theorem 4.3.15, (X, d) is totally bounded by Proposition 4.3.9 as desired.

For the other direction, suppose (X, d) is complete and totally bounded. To show that (X, d) is compact, we will demonstrate that (X, d) is sequentially compact and apply the Borel-Lebesgue Theorem (Theorem 4.3.15).

To see that (X, d) is sequentially compact, let $(x_n)_{n\geq 1}$ be an arbitrary sequence of elements of X. Since (X, d) is totally bounded, $F_1 = \{x_n\}_{n\geq 1}$ is totally bounded by Proposition 4.3.11. Hence F_1 has a finite 1-net. This implies there exists an $n_1 \in \mathbb{N}$ such that $x_{n_1} \in F_1$ and

$$I_1 = \{n \in \mathbb{N} \mid n > n_1 \text{ and } x_n \in B_d(x_{n_1}, 1)\}$$

is infinite. Let $F_2 = \{x_n\}_{n\in I_1}$. Since (X, d) is totally bounded, F_2 is totally bounded by Proposition 4.3.11 and thus F_2 has finite $\frac{1}{2}$-net. As I_1 is infinite, there exists a $n_2 \in \mathbb{N}$ such that $n_2 \in I_1$ (so $n_2 > n_1$) such that $x_{n_2} \in F_2$ and

$$I_2 = \left\{ n \in \mathbb{N} \,\middle|\, n > n_2 \text{ and } x_n \in B_d\left(x_{n_2}, \frac{1}{2}\right) \right\}$$

is infinite. Let $F_3 = \{x_n\}_{n\in I_2}$. By repeating this process ad infinitum, there exists infinite subsets F_n of F_1 and an increasing sequence $(n_k)_{k\geq 1}$ of natural number such that $x_{n_k} \in F_m$ for all $k \geq m$ and $x_{n_k} \in B_d(x_{n_m}, \frac{1}{m})$ for all $k > m$. Hence $(x_{n_k})_{k\geq 1}$ is a Cauchy subsequence of $(x_n)_{n\geq 1}$. Since (X, d) is complete, $(x_{n_k})_{k\geq 1}$ is a convergent subsequence of $(x_n)_{n\geq 1}$. Therefore, as $(x_n)_{n\geq 1}$ was arbitrary, (X, d) is sequentially compact as desired. ∎

Theorem 4.3.16 is incredibly useful in verifying a metric space is compact. Indeed verifying that a metric space is totally bounded simply comes down to fixing an ϵ and picking a bunch of points that forms an ϵ-net. Moreover we have already seen several methods of verifying a metric spaces is complete in Section 4.1 and Section 4.2, such as verifying the metric space under consideration is a closed subset of a known complete metric space. Verifying a set is closed is simple by Theorem 4.1.13 as we need to check that the point of convergence of a convergent sequence is actually in the set. This is a far simpler task than verifying sequential compactness as we need not construct a convergent subsequence for every possible sequence in the space.

With Theorem 4.1.13, we give one example of how to apply Theorem 4.3.16.

Example 4.3.17. Let

$$K = \left\{ (x_n)_{n\geq 1} \in \ell_2(\mathbb{R}) \,\middle|\, \sum_{n=1}^{\infty} n^2 |x_n|^2 \leq 1 \right\}.$$

Then K is a compact subspace of $\ell_2(\mathbb{R})$. To see this, it suffices by Theorem 4.3.16 to show that K is complete and totally bounded. Since $\ell_2(\mathbb{R})$ is complete by Proposition 4.1.18, it suffices to show that K is a closed, totally bounded subspace of $\ell_2(\mathbb{R})$.

To see that K is closed in $(\ell_2(\mathbb{R}), \|\cdot\|_2)$, let $(\vec{v}_k)_{k\geq 1}$ be an arbitrary sequence of elements of K that converges to some $\vec{x} \in \ell_2(\mathbb{R})$. For each $k \in \mathbb{N}$ write

$$\vec{v}_k = (x_{k,n})_{n\geq 1} \qquad \text{and} \qquad \vec{x} = (x_n)_{n\geq 1}.$$

Since

$$\lim_{k\to\infty} \|\vec{v}_k - \vec{x}\|_2 = 0 \qquad \text{and} \qquad |x_{k,n} - x_n| \leq \|\vec{v}_k - \vec{x}\|_2 \text{ for all } k, n \in \mathbb{N},$$

we obtain that $\lim_{k\to\infty} |x_{k,n} - x_n| = 0$ for all $n \in \mathbb{N}$. Furthermore, $\vec{v}_k \in K$ for all $k \in \mathbb{N}$, we obtain by the definition of K that

$$\sum_{n=1}^{\infty} n^2 |x_{k,n}|^2 \leq 1$$

for all $k \in \mathbb{N}$. Hence for all $N \in \mathbb{N}$ we see that

$$\sum_{n=1}^{N} n^2 |x_n|^2 = \lim_{k\to\infty} \sum_{n=1}^{N} n^2 |x_{k,n}|^2 \leq \limsup_{k\to\infty} \sum_{n=1}^{\infty} n^2 |x_{k,n}|^2 \leq 1.$$

Therefore, since the above holds for all $N \in \mathbb{N}$, we obtain that $\sum_{n=1}^{\infty} n^2 |x_n|^2 \leq 1$ and thus $\vec{x} \in K$. Thus, as $(\vec{v}_k)_{k\geq 1}$ was arbitrary, we obtain that K is closed.

To see that K is totally bounded, let $\epsilon > 0$ be arbitrary. Without loss of generality, we may assume that $\epsilon < 1$. To see that K has an ϵ-net, first notice if $\vec{x} = (x_n)_{n\geq 1} \in K$ then $\sum_{n=1}^{\infty} n^2 |x_n|^2 \leq 1$ so $n^2 |x_n|^2 \leq 1$ for all $n \in \mathbb{N}$ and thus $|x_n| \leq \frac{1}{n}$ for all $n \in \mathbb{N}$. To begin to use this, we note since $\sum_{n=1}^{\infty} \frac{1}{n}^2 < \infty$ there exists and $N \in \mathbb{N}$ such that

$$\sum_{n=N+1}^{\infty} \frac{1}{n^2} < \frac{\epsilon^2}{2}.$$

Hence, if $\vec{x} = (x_n)_{n\geq 1} \in K$ then the above shows that

$$\sum_{n=N+1}^{\infty} |x_n|^2 < \frac{\epsilon^2}{2}.$$

Consider the set

$$K_0 = \left\{ (x_1, \ldots, x_N) \in \mathbb{R}^N \;\middle|\; \sum_{n=1}^{N} n^2 |x_n|^2 \leq 1 \right\} \subseteq \mathbb{R}^N.$$

By the same arguments used above, K_0 is a closed subset of $\mathbb{R}^n$. Furthermore, if $(x_1, \ldots, x_n) \in K_0$, then $|x_n| \leq \frac{1}{n}$ for all $n \in \mathbb{N}$ and thus

$$\|(x_1, \ldots, x_N)\|_2 \leq \left(\sum_{n=1}^{N} \frac{1}{n^2} \right)^{\frac{1}{2}}.$$

Therefore K_0 is bounded in $(\mathbb{R}^2, \|\cdot\|_2)$ and hence compact by the Heine-Borel Theorem (Theorem 3.1.25). Thus K_0 has a finite $\frac{\epsilon}{2}$-net.

Let $\vec{v}_1, \ldots, \vec{v}_m$ be a finite $\frac{\epsilon}{\sqrt{2}}$-net for K_0. Clearly each $\vec{v}_k$ defines an element of K by extending the N-tuple to a sequence by letting every term in the sequence with index greater than N be zero. We claim that $\vec{v}_1, \ldots, \vec{v}_m$ then forms an ϵ-net of K. To see this, let $\vec{x} = (x_n)_{n\geq 1} \in K$ be arbitrary. Then

$$(x_1, \ldots, x_N) \in K_0 \text{ by construction} \qquad \text{and} \qquad \sum_{n=N+1}^{\infty} |x_n|^2 < \frac{\epsilon^2}{2}.$$

Since $\vec{v}_1, \ldots, \vec{v}_m$ is a finite $\frac{\epsilon}{\sqrt{2}}$-net for K_0, there exists a $k \in \{1, \ldots, m\}$ such that

$$\|\vec{v}_k - (x_1, \ldots, x_N)\|_2 < \frac{\epsilon}{\sqrt{2}}.$$

Hence

$$\begin{aligned}\|\vec{v}_k - \vec{x}\|_2^2 &= \|\vec{v}_k - (x_1, \ldots, x_N)\|_2^2 + \sum_{n=N+1}^{\infty} |x_n|^2 \\ &< \frac{\epsilon^2}{2} + \frac{\epsilon^2}{2} = \epsilon^2.\end{aligned}$$

Therefore $\|\vec{v}_k - \vec{x}\|_2 < \epsilon$. Hence, as $\vec{x} \in K$ was arbitrary, $\vec{v}_1, \ldots, \vec{v}_m$ is an ϵ-net of K. Therefore, since $\epsilon > 0$ was arbitrary, K is totally bounded. Consequently, K is compact as desired.

4.4 Compact Function Spaces

Now that we have characterized the compact metric spaces as those that are complete and totally bounded in Theorem 4.3.16, we turn our attention back to function spaces. Indeed, recall from Theorem 4.2.22 that every metric space is isomorphic to a subset of a function space. Thus by studying compactness in function spaces, we are studying compactness for all metric spaces!

In this section, we endeavour to determine when specific collections of functions in a function space form compact subspaces. This is particularly useful in deriving properties of functions from other functions. For example, suppose we have a compact set of functions Φ with a specific property. Then, if we construct a net of functions from Φ in a specific way, we know by Theorem 3.2.2 that this net then has a subnet that converges to an element of Φ and thus must have the same properties. Of course, we will want to study closed set of complete function spaces because Theorem 4.1.14 implies closed sets are complete. As often one desires only to describes a collection of functions without their closure, we define the following.

Definition 4.4.1. Let $(X, \mathcal{T})$ be a topological space. A subset $A \subseteq X$ is said to be *relatively compact* if $\overline{A}$ is compact.

Remark 4.4.2. Notice that if (X, d) is a complete metric space and $A \subseteq X$, then A is relatively compact if and only if A is totally bounded by Theorem 4.1.14, Theorem 4.3.16, and Proposition 4.3.11.

Thus, if we want to study relatively compact subsets of functions space, we need only study which collections of functions are totally bounded. Of course, verifying totally boundedness from definition is easier than verifying compactness from definition, but it still is not simple. Thus we desire to find simpler conditions to verify a collection of functions is totally bounded.

Of course, if a collection of functions is totally bounded with respect to the sup metric, every function will be close to another function from a finite collection. Knowing how each element of this finite collection is continuous at a point then yields information about how the entire collection is continuous at a point. This leads us to the following notion of a collection of functions being 'equally continuous'.

Definition 4.4.3. Let $(X, \mathcal{T})$ be a topological space, let (Y, d) be a metric space, let $x_0 \in X$, and let $\mathcal{F} \subseteq \mathcal{C}(X, Y)$. It is said that $\mathcal{F}$ is *equicontinuous at* x_0 if for all $\epsilon > 0$ there exists a neighbourhood U of x_0 in $(X, \mathcal{T})$ such that $d(f(x), f(x_0)) < \epsilon$ for all $x \in U$ and $f \in \mathcal{F}$.

It is said that $\mathcal{F}$ is *equicontinuous* if $\mathcal{F}$ is equicontinuous at every point in X.

Of course, examples are easy to come by.

Example 4.4.4. For each $n \in \mathbb{N}$ let $f_n : [-1, 1] \to \mathbb{R}$ be defined by $f_n(x) = x^n$ for all $x \in [-1, 1]$. The collection $\mathcal{F} = \{f_n\}_{n\in\mathbb{N}}$ is equicontinuous at 0. Indeed if $\epsilon > 0$ let $\delta = \min\{\epsilon, 1\} > 0$. Then if $|x| < \delta$ so

$$|f_n(x)| = |x^n| \leq \delta^n \leq \epsilon.$$

Hence $\mathcal{F}$ is equicontinuous at 0. However, $\mathcal{F}$ is not equicontinuous at 1. To see this, notice for all $\delta > 0$,

$$\lim_{n\to\infty} |f_n(1) - f_n(1 - \delta)| = \lim_{n\to\infty} |1 - (1 - \delta)^n| = 1$$

so no δ-ball cented at 1 can work in Definition 4.4.3 for $\epsilon = \frac{1}{2}$.

To emphasize the idea that equicontinuity should stem from total boundedness, we note the following lemma.

Lemma 4.4.5. *Let $(X, \mathcal{T})$ be a compact topological space, let (Y, d_Y) be a metric space, let $\mathcal{F} \subseteq \mathcal{C}(X, Y)$ be totally bounded with respect to the sup metric. Then $\mathcal{F}$ is equicontinuous.*

Proof. To see that $\mathcal{F}$ is equicontinuous, let $\epsilon > 0$ and $x_0 \in X$ be arbitrary. Since $\mathcal{F}$ is totally bounded with respect to the sup metric on $\mathcal{C}(X, Y)$, there exists a finite $\frac{\epsilon}{3}$-net for $\mathcal{F}$. Hence there exists an $n \in \mathbb{N}$ and $f_1, f_2, \ldots, f_n \in \mathcal{F}$ such that

$$\mathcal{F} \subseteq \bigcup_{k=1}^{n} B_{d_{\sup}}\left(f_k, \frac{\epsilon}{3}\right).$$

Since f_k is continuous at x_0 for all $k \in \{1, \ldots, n\}$, there exists a $U_k \in \mathcal{T}$ such that

$$d_Y(f_k(x), f_k(x_0)) < \frac{\epsilon}{3}$$

for all $x \in U_k$. Let $U = \bigcap_{k=1}^{n} U_k \in \mathcal{T}$. We claim that U an open set that works for ϵ in Definition 4.4.3 to show that $\mathcal{F}$ is equicontinuous at x_0. To see this, let $f \in \mathcal{F}$ be arbitrary. Hence, as $\mathcal{F} \subseteq \bigcup_{k=1}^{n} B_{d_{\sup}}\left(f_k, \frac{\epsilon}{3}\right)$, there exists a $k \in \{1, \ldots, n\}$ such that

$$d_{\sup}(f, f_k) < \frac{\epsilon}{3}.$$

Thus $d_Y(f(x), f_k(x)) < \frac{\epsilon}{3}$ for all $x \in X$. Therefore, for all $x \in U \subseteq U_k$, we obtain from above that

$$\begin{aligned} d_Y(f(x), f(x_0)) &\leq d_Y(f(x), f_k(x)) + d_Y(f_k(x), f_k(x_0)) + d_Y(f_k(x_0), f(x_0)) \\ &< \frac{\epsilon}{3} + \frac{\epsilon}{3} + \frac{\epsilon}{3} = \epsilon. \end{aligned}$$

Therefore, as $f \in \mathcal{F}$, $\epsilon > 0$, and $x_0 \in X$ were arbitrary, $\mathcal{F}$ is equicontinuous as desired. ∎

Of course, equicontinuity is a nice property as it passes to closures of sets; something we expect as we are studying relative compactness of function spaces.

Proposition 4.4.6. *Let $(X, \mathcal{T})$ be a compact topological space, let (Y, d) be a metric space, let $\mathcal{F} \subseteq \mathcal{C}(X, Y)$ be equicontinuous. Then $\overline{\mathcal{F}}$ (computed with respect to the sup metric on $\mathcal{C}(X, Y)$ is equicontinuous.*

Proof. To see that $\overline{\mathcal{F}}$ is equicontinuous, fix an arbitrary element $x_0 \in X$ and let $\epsilon > 0$ be arbitrary. Since $\mathcal{F}$ is equicontinuous, there exists a neighbourhood U of x_0 in $(X, \mathcal{T})$ such that $d(f(x), f(x_0)) < \frac{\epsilon}{3}$ for all $x \in U$ and $f \in \mathcal{F}$. To see that U works for ϵ in the definition of equicontinuity for $\overline{\mathcal{F}}$, let $g \in \overline{\mathcal{F}}$ be arbitrary. By Theorem 1.6.21 and the definition of the sup metric, there exists an $f \in \mathcal{F}$ such that $d(g(x), f(x)) < \frac{\epsilon}{3}$ for all $x \in X$. Therefore, if $x \in U$, we obtain that

$$\begin{aligned} d(g(x), g(x_0)) &\leq d(g(x), f(x)) + d(f(x), f(x_0)) + d(f(x_0), g(x_0)) \\ &\leq \frac{\epsilon}{3} + \frac{\epsilon}{3} + \frac{\epsilon}{3} = \epsilon. \end{aligned}$$

Therefore, as $g \in \overline{\mathcal{F}}$, $\epsilon > 0$, and $x_0 \in X$ were arbitrary, we obtain that $\overline{\mathcal{F}}$ is equicontinuous. ∎

Unfortunately, equicontinuity does not immediately imply total boundedness.

Example 4.4.7. For each $a \in \mathbb{R}$, let $f_a : [0,1] \to \mathbb{R}$ be defined by $f(x) = x+a$ for all $x \in [0,1]$. Clearly $\{f_a\}_{a\in\mathbb{R}}$ is equicontinuous as

$$|f_a(x) - f_a(y)| = |f_0(x) - f_0(y)|$$

for all $x, y \in [0,1]$. However $\{f_a\}_{a\in\mathbb{R}}$ cannot be totally bounded with respect to the ∞-norm on $\mathcal{C}([0,1],\mathbb{R})$ since $\|f_a\|_\infty = a+1$ for all $a \geq 0$ so $\{f_a\}_{a\in\mathbb{R}}$ is not bounded with respect to $\|\cdot\|_\infty$ and thus cannot be totally bounded by Proposition 4.3.10.

Thus the problem is that equicontinuity does not yield any information about a collection of functions behaving like a bounded collection of functions. Of course we could just ask that the collection of functions is bounded with respect to the sup metric. However, there is also a much simpler notion of boundedness we can ask for.

Definition 4.4.8. Let $(X, \mathcal{T})$ be a topological space, let (Y, d) be a metric space, and let $\mathcal{F} \subseteq \mathcal{F}(X,Y)$. It is said that $\mathcal{F}$ is *pointwise bounded* if $\{f(x) \mid f \in \mathcal{F}\}$ is bounded in (Y, d) for all $x \in X$.

However, there no immediate connection between pointwise boundedness and boundedness (and hence total boundedness) of collections of functions.

Example 4.4.9. For each $n \in \mathbb{N}$ let $f_n : [0,1] \to \mathbb{R}$ be defined by

$$f_n(x) = \begin{cases} n^2 x & \text{if } x \in \left[0, \frac{1}{n}\right] \\ n^2\left(\frac{2}{n} - x\right) & \text{if } x \in \left[\frac{1}{n}, \frac{2}{n}\right] \\ 0 & \text{otherwise} \end{cases}$$

for all $x \in [0,1]$. We claim the collection $\mathcal{F} = \{f_n\}_{n\in\mathbb{N}}$ is pointwise bounded. To see this, notice $f_n(0) = 0$ for all $n \in \mathbb{N}$ so $\mathcal{F}$ is bounded at 0. Otherwise, if $x \in (0,1]$, choose $N \in \mathbb{N}$ such that $\frac{2}{N} < x$. Then it is easy to see that $\mathcal{F}$ is bounded by

$$\max(\{f_1(x), f_2(x), \ldots, f_{N-1}(x), f_N(x) = 0\}).$$

Hence $\mathcal{F}$ is pointwise bounded.

However, $\mathcal{F}$ is not bounded in $(\mathcal{C}(X,Y), \|\cdot\|_\infty)$ as $|f_n\left(\frac{1}{n}\right)| = n$ for all $n \in \mathbb{N}$ so $\|f_n\|_\infty \geq n$ for all $n \in \mathbb{N}$. Therefore, as $\mathcal{F}$ is not bounded, $\mathcal{F}$ is not totally bounded by Proposition 4.3.10.

Luckily, the reason the example in Example 4.4.9 is pointwise bounded but not imply total bounded is that the collection of functions was not equicontinuous at 0. The fact the collection of functions in Example 4.4.9 is not equicontinuous at 0 can be seen by similar arguments to those used in Example 4.4.4, or via the following lemma. In particular, the following lemma shows that the union of the ranges of a collection of functions that is pointwise bounded and equicontinuous is actually bounded in the co-domain. This is equivalent to the collection of functions being bounded with respect to the sup metric by consider Lemma 3.1.21 and any constant function (i.e. if a collection of functions is such that the union of the ranges is bounded, every function is a finite sup metric distance from any constant function, and if every function is a finite sup metric distance from a constant function, the range of each function is contained in a fixed open ball centred at the constant).

Lemma 4.4.10. *Let* $(X, \mathcal{T})$ *be a compact topological space, let* (Y, d) *be a metric space, and let* $\mathcal{F} \subseteq \mathcal{C}(X, Y)$*. If* $\mathcal{F}$ *is equicontinuous and pointwise bounded, then*

$$Y_0 = \bigcup_{f \in \mathcal{F}} f(X)$$

is a bounded subset of (Y, d)*.*

Proof. Let $\epsilon = 1$. Since $\mathcal{F}$ is equicontinuous, for each $x_0 \in X$ there exists a $U_{x_0} \in \mathcal{T}$ such that $d(f(x), f(x_0)) < 1$ for all $x \in U_{x_0}$ and $f \in \mathcal{F}$. Thus $\{U_x\}_{x \in X}$ is an open cover of $(X, \mathcal{T})$. Therefore, since $(X, \mathcal{T})$ is compact, there exists an $n \in \mathbb{N}$ and $x_1, x_2, \ldots, x_n \in X$ such that $X = \bigcup_{k=1}^n U_{x_k}$.

Let

$$B = \bigcup_{k=1}^{n} \{f(x_k) \mid f \in \mathcal{F}\} \subseteq Y.$$

Since $\mathcal{F}$ is pointwise bounded, B is a union of bounded subsets of Y and thus bounded by part (2) of Lemma 3.1.21 together with the triangle inequality. Hence there exists an $M \in \mathbb{R}$ such that $d(b_1, b_2) \leq M$ for all $b_1, b_2 \in B$.

To see that Y_0 is bounded in Y, let $y_1, y_2 \in Y_0$ be arbitrary. Hence there exists $z_1, z_2 \in X$ and $f_1, f_2 \in \mathcal{F}$ such that $y_1 = f_1(z_1)$ and $y_2 = f_2(z_2)$. Since $X = \bigcup_{k=1}^n U_{x_k}$ there exist $k_1, k_2 \in \{1, \ldots, n\}$ such that $z_1 \in U_{x_{k_1}}$ and $z_2 \in U_{x_{k_2}}$. Thus

$$d(f_1(z_1), f_1(x_{k_1})) < 1 \qquad \text{and} \qquad d(f_2(z_2), f_2(x_{k_2})) < 1$$

by the construction of $\{U_x\}_{x \in X}$. Therefore, since $f_1(x_{k_1}), f_2(x_{k_2}) \in B$, we obtain that

$$\begin{aligned} d(y_1, y_2) &\leq d(f_1(z_1), f_1(x_{k_1})) + d(f_1(x_{k_1}), f_2(x_{k_2})) + d(f_2(z_2), f_2(x_{k_2})) \\ &< 1 + M + 1 = M + 2. \end{aligned}$$

Therefore, since $y_1, y_2 \in Y_0$ were arbitrary, Y_0 is bounded as desired. ∎

As boundedness was our only previous obstruction to a set being totally bounded, Lemma 4.4.10 says, provided the collection of functions is equicontinuous and pointwise bounded, we no longer have a simple obstruction. In fact, the following important theorem says we never will provided we have a nice co-domain!

Theorem 4.4.11 (The Arzelà-Ascoli Theorem). *Let $(X, \mathcal{T})$ be a compact topological space and let $\mathcal{F} \subseteq \mathcal{C}(X, \mathbb{K}^n)$ for some $n \in \mathbb{N}$. The following are equivalent:*

(i) $\mathcal{F}$ is relatively compact in $(\mathcal{C}(X, \mathbb{K}^n), \|\cdot\|_\infty)$.

(ii) $\mathcal{F}$ is equicontinuous and pointwise bounded.

Proof. We will present the proof only in the case that $n = 1$ and $\mathbb{K} = \mathbb{R}$ to simplify notation and the arguments. However, we will emphasize the necessary changes to deal with the general case when they occur.

To begin, suppose $\mathcal{F}$ is relatively compact in $(\mathcal{C}(X, \mathbb{K}^n), \|\cdot\|_\infty)$. Hence $\overline{\mathcal{F}}$ is compact and thus complete and totally bounded by Theorem 4.3.16. Thus $\overline{\mathcal{F}}$ is bounded with respect to $\|\cdot\|_\infty$ by Proposition 4.3.10 so $\mathcal{F}$ is bounded with respect to $\|\cdot\|_\infty$ and hence pointwise bounded. To see that $\mathcal{F}$ is equicontinuous, we note that since $\overline{F}$ is totally bounded, $\mathcal{F}$ is totally bounded by Proposition 4.3.11. Hence $\mathcal{F}$ is equicontinuous by Lemma 4.4.5. Hence the first direction of the proof is complete.

For the other direction, suppose $\mathcal{F}$ is equicontinuous and pointwise bounded. Recall that it suffices to prove that $\mathcal{F}$ is totally bounded by Remark 4.4.2. To begin this process, recall Lemma 4.4.10 implies that

$$\bigcup_{f \in \mathcal{F}} f(X)$$

is a bounded subsets of $\mathbb{R}$. Hence there exists an $M \in \mathbb{R}$ such that $f(x) \in [-M, M]$ for all $x \in X$ and $f \in \mathcal{F}$ (clearly if we are using $\mathbb{R}^n$ we would use $[-M, M]^n$ here, and for $\mathbb{C}^n \equiv \mathbb{R}^{2n}$ we would use $[-M, M]^{2n}$). To extend this form of boundedness for $\mathcal{F}$ to total boundedness, our goal is to divide up X and $[-M, M]$ into suitably small pieces, take one function that maps each piece of X into a chosen piece of $[-M, M]$ that is actually obtained, and show this collection of functions is an ϵ-net for $\mathcal{F}$.

Thus, to see that $\mathcal{F}$ is totally bounded, let $\epsilon > 0$ be arbitrary. Since $\mathcal{F}$ is equicontinuous, for every $x_0 \in X$ there exists a $U_{x_0} \in \mathcal{T}$ such that

$$|f(x) - f(x_0)| < \frac{\epsilon}{3}$$

for all $x \in U_{x_0}$ and $f \in \mathcal{F}$ (for $\mathbb{K}^n$ we would use the $\|\cdot\|_\infty$ on $\mathbb{K}^n$). Thus $\{U_x\}_{x \in X}$ is an open cover of $(X, \mathcal{T})$. Therefore, since $(X, \mathcal{T})$ is compact, there exists an $n \in \mathbb{N}$ and $x_1, x_2, \ldots, x_n \in X$ such that $X = \bigcup_{j=1}^n U_{x_j}$.

Choose numbers $\{m_k\}_{k=1}^q$ such that

$$-M = m_1 < m_2 < \cdots < m_q = M$$

and $|m_{k+1} - m_k| < \frac{\epsilon}{3}$ for all $k \in \{1, \ldots, q-1\}$ (for $\mathbb{K}^n$, we would divided up $[-M, M]^p$ into boxes of $\|\cdot\|_\infty$ radius less than $\frac{\epsilon}{3}$). For each n-tuple $(k_1, \ldots, k_n) \in \{1, \ldots, q-1\}^n$, let

$$\mathcal{F}_{(k_1,\ldots,k_n)} = \{f \in \mathcal{F} \mid f(x_j) \in [m_{k_j}, m_{k_j+1}] \text{ for all } j \in \{1, \ldots, n\}\}$$

(for $\mathbb{K}^n$ we would divide $\mathcal{F}$ into collections based on the boxes constructed earlier). Clearly

$$\mathcal{F} = \bigcup_{(k_1,\ldots,k_n)\in\{1,\ldots,q-1\}^n} \mathcal{F}_{(k_1,\ldots,k_n)}$$

by construction and the fact that $\bigcup_{f\in\mathcal{F}} f(X) \subseteq [-M, M]$.

For each $(k_1, \ldots, k_n) \in \{1, \ldots, q-1\}^n$ for which $\mathcal{F}_{(k_1,\ldots,k_n)} \neq \emptyset$, choose a $f_{(k_1,\ldots,k_n)} \in \mathcal{F}_{(k_1,\ldots,k_n)}$. We claim the collection of all $f_{(k_1,\ldots,k_n)}$ (which is a finite set) is an ϵ-net for $\mathcal{F}$. To see this, let $f \in \mathcal{F}$ be arbitrary. Hence $f \in \mathcal{F}_{(k_1,\ldots,k_n)}$ for some $(k_1, \ldots, k_n) \in \{1, \ldots, q-1\}^n$. To see that

$$\left\|f - f_{(k_1,\ldots,k_n)}\right\|_\infty \leq \epsilon,$$

let $x_0 \in X$ be arbitrary. Hence, as $X = \bigcup_{j=1}^n U_{x_j}$ there exists a $j \in \{1, \ldots, n\}$ such that $x \in U_{x_j}$. Hence

$$|f(x_0) - f(x_j)| < \frac{\epsilon}{3} \qquad \text{and} \qquad |f_{(k_1,\ldots,k_n)}(x_0) - f_{(k_1,\ldots,k_n)}(x_j)| < \frac{\epsilon}{3}.$$

However, as $f \in \mathcal{F}_{(k_1,\ldots,k_n)}$, the fact that $|m_{k+1} - m_k| < \frac{\epsilon}{3}$ for all $k \in \{1, \ldots, q-1\}$ implies that

$$|f(x_j) - f_{(k_1,\ldots,k_n)}(x_j)| < \frac{\epsilon}{3}.$$

Hence the triangle inequality implies

$$|f(x_0) - f_{(k_1,\ldots,k_n)}(x_0)| < \epsilon.$$

Therefore, as $x_0 \in X$ was arbitrary, $\left\|f - f_{(k_1,\ldots,k_n)}\right\|_\infty \leq \epsilon$. Therefore, as $f \in \mathcal{F}$ was arbitrary, we have proven the existence of an ϵ-net for $\mathcal{F}$. Hence, as $\epsilon > 0$ was arbitrary, $\mathcal{F}$ is totally bounded as desired. ∎

Of course, the Arzelà-Ascoli Theorem (Theorem 4.4.11) as stated above only applies to functions on compact spaces into $\mathbb{K}^n$. One may think this is not too powerful until one realizes that most situations that one desires to study is for functions into $\mathbb{R}$ for which the Arzelà-Ascoli Theorem applies. Of course there are many generalizations of the Arzelà-Ascoli Theorem one may find in the literature. In fact, the proof of Arzelà-Ascoli Theorem given above immediately implies one of these theorems. Indeed note the only part of the proof of Theorem 4.4.11 where a property of $\mathbb{K}^n$ was applied was that we could take a bounded set containing the union of the ranges of the collection of functions and divide it up into a finite collection of arbitrary small portions. Thus, if the union of the ranges is contained in a compact subset of a metric space, then this can be done using a finite collection of $\frac{\epsilon}{6}$-balls instead of the box decomposition used above. Consequently, repeating the proof pretty much verbatim yields the following result.

Corollary 4.4.12. *Let $(X, \mathcal{T})$ be a compact topological space, let (Y, d_Y) be a compact metric space, and let $\mathcal{F} \subseteq \mathcal{C}(X, Y)$. The following are equivalent:*

(i) $\mathcal{F}$ is relatively compact in $(\mathcal{C}(X, Y), d_{\text{sup}})$.

(ii) $\mathcal{F}$ is equicontinuous and pointwise bounded.

To complete our discussion of the Arzelà-Ascoli Theorem (Theorem 4.4.11) we note it is a powerful tool to verify sets of functions are relatively compact. Indeed verifying a set of functions is pointwise bounded is generally trivial and verify a collection of functions is equicontinuous is no more difficult then verifying a single function is continuous using ϵ-δ. Consequently, if one desires to verify a collection of function is actually compact, one need only verify the collection is relatively compact and closed. One example of this is as follows.

Example 4.4.13. Let

$$K = \left\{ f \in \mathcal{C}[0,1] \,\middle|\, |f(x) - f(y)| \leq \sqrt{|x-y|}\, \forall\, x, y \in [0,1] \text{ and } f(0) = 0 \right\}.$$

Then one can use the Arzelà-Ascoli Theorem (Theorem 4.4.11) to show without much difficulty that K is a compact subset of $\mathcal{C}[0,1]$.

To prove that K is a compact subset of $(\mathcal{C}[0,1], \|\cdot\|_\infty)$, we will show that K is closed and relatively compact as this implies $K = \overline{K}$ is compact.

To see that K is a closed subset of $(\mathcal{C}[0,1], \|\cdot\|_\infty)$, let $(f_n)_{n\geq 1}$ be an arbitrary sequence in K that converges to some $f \in \mathcal{C}[0,1]$ with respect to $\|\cdot\|_\infty$. By the definition of the infinity norm, we see that $(f_n)_{n\geq 1}$ converges pointwise to f. Therefore, since

$$|f_n(x) - f_n(y)| \leq \sqrt{|x-y|} \text{ for all } x, y \in [0,1] \qquad \text{and} \qquad f_n(0) = 0$$

for all $n \in \mathbb{N}$ due to the defining properties of K, we obtain that

$$|f(x) - f(y)| \leq \sqrt{|x-y|} \text{ for all } x, y \in [0,1] \qquad \text{and} \qquad f(0) = 0$$

so $f \in K$ by definition. Therefore, since $(f_n)_{n\geq 1}$ was arbitrary, K is closed in $(\mathcal{C}[0,1], \|\cdot\|_\infty)$.

To see that K is relatively compact in $(\mathcal{C}[0,1], \|\cdot\|_\infty)$, it suffices to show by the Arzelà-Ascoli Theorem (Theorem 4.4.11) that K is equicontinuous and pointwise bounded. To see that K is pointwise bounded, notice for all $f \in K$ and $x \in [0,1]$ that

$$|f(x)| = |f(x) - f(0)| \leq \sqrt{x-0} = \sqrt{x}.$$

Consequently, K is clearly pointwise bounded. To see that K is equicontinuous, let $\epsilon > 0$ and $x \in [0,1]$ be arbitrary. Since the function $g : [0,1] \to \mathbb{R}$ defined by $g(y) = \sqrt{|x-y|}$ is continuous and vanishes at x, there exists a $\delta > 0$ such that if $y \in [0,1]$ and $|x-y| < \delta$ then $g(y) < \epsilon$. Hence for all $f \in K$ and $y \in [0,1]$ such that $|x-y| < \delta$, we obtain that

$$|f(x) - f(y)| \leq \sqrt{|x-y|} < \epsilon.$$

Therefore, since $f \in K$, $\epsilon > 0$, and $x \in [0,1]$ were arbitrary, K is equicontinuous as desired.

4.5 Weierstrass Approximation Theorem

The notion of relative compactness raises the question about how one goes about taking the closure of a set of functions with respect to the sup metric. In particular, as an element x is in the closure of a set if and only if there is a net from the set converging to x by Theorem 1.6.21, and as a net of functions converges with respect to the sup metric if and only if it converges uniformly, we are asking when one function can be uniformly approximated by other functions. This is often useful as there may be a nice collection of functions one understands that approximate all other functions. Hence one may use this nice collection to understand all functions.

In this section, we will delve into this question by proving the simplest case of such a theorem. Namely, we will demonstrate the Weierstrass Approximation Theorem (Theorem 4.5.7) which states every real-valued continuous function on a finite closed interval may be uniformly approximated by a polynomial. As we are often going to want to say a collection of functions approximate all functions, we given this concept a name.

Definition 4.5.1. Let $(X, \mathcal{T})$ be a topological space. A subset $A \subseteq X$ is said to be *dense* if $\overline{A} = X$.

Of course, to say a set of functions is dense with respect to the topology induced by the sup metric is to precisely say that the set of functions can uniformly approximate any other function. Thus this is precisely the notion we are after.

To prove the Weierstrass Approximation Theorem (Theorem 4.5.7) we need three ingredients plus a delicate proof. The first ingredient says we can study any particular finite closed interval we choose.

Lemma 4.5.2. *Consider the linear map $T : \mathcal{C}[a,b] \to \mathcal{C}[0,1]$ by*

$$T(f)(x) = f(a + (b-a)x)$$

for all $x \in [0,1]$ and $f \in \mathcal{C}[a,b]$. Then T is an isometric isomorphism such that $T(p)$ is a polynomial if and only if p is a polynomial.

Proof. Clearly $T(f)$ is well-defined and a continuous function on $[0,1]$ for all $f \in \mathcal{C}[a,b]$. It is elementary to see that T is linear and that $\|T(f)\|_\infty = \|f\|_\infty$ for all $f \in \mathcal{C}[a,b]$. Therefore, as $T^{-1} : \mathcal{C}[0,1] \to \mathcal{C}[a,b]$, defined by

$$T^{-1}(f)(x) = f\left(\frac{x-a}{b-a}\right)$$

for all $x \in [a,b]$ and $f \in \mathcal{C}[0,1]$, exists, we see that T is an isometric isomorphism. In addition, it is clear that if p is a polynomial then $T(p)$ is polynomial and $T^{-1}(p)$ is a polynomial. Hence the result follows. ∎

Our second ingredient is a technical result for a function we will encounter and is proved using elementary calculus.

Lemma 4.5.3. *If $x \in [-1,1]$ and $n \in \mathbb{N}$, then*

$$(1-x^2)^n \geq 1 - nx^2.$$

Proof. Clearly it suffices to consider $x \in [0,1]$ as $(1-(-x)^2)^n = (1-x^2)^n$ and $1 - n(-x)^2 = 1 - nx^2$ for all $x \in [-1,1]$.

Consider the functions $f, g : [0,1] \to \mathbb{R}$ defined by

$$f(x) = (1-x^2)^n \qquad \text{and} \qquad g(x) = 1 - nx^2$$

for all $x \in [0,1]$. Clearly $f(0) = 1 = g(0)$. Furthermore, f and g are differentiable with

$$f'(x) = n(1-x^2)(-2x) \qquad \text{and} \qquad g'(x) = -2nx.$$

As $-2nx \leq 0$ and $0 \leq 1 - x^2 \leq 1$ for all $x \in [0,1]$, we see that $f'(x) \geq g'(x)$ for all $x \in [0,1]$. Hence it follows that $f(x) \geq g(x)$ for all $x \in [0,1]$ as desired. ∎

Our third ingredient is a stronger notion of continuity for functions between metric spaces.

Definition 4.5.4. Let (X, d_X) and (Y, d_Y) be metric spaces. A function $f : X \to Y$ is said to be *uniformly continuous* if for all $\epsilon > 0$ there exists a $\delta > 0$ such that if $x_1, x_2 \in X$ are such that $d_X(x_1, x_2) < \delta$ then $d_Y(f(x_1), f(x_2)) < \epsilon$.

Of course, it is elementary to see that the notion of uniform continuity is stronger than the notion of continuity for metric spaces since the δ for a given ϵ works throughout X; that is, there is one δ to rule them all, one δ to find them, one δ to bring them all, and in the darkness bind them. To emphasize these notions are not the same, we consider the following example.

Example 4.5.5. Let $f : (0, 1) \to \mathbb{R}$ be defined by $f(x) = \frac{1}{x}$ for all $x \in (0, 1)$. Clearly f is continuous by properties of $\mathbb{R}$. However f is not uniformly continuous. To see this, for each $n \in \mathbb{N}$ let $x_n = \frac{1}{n}$ and $y_n = \frac{2}{n}$. Then $|x_n - y_n| < \frac{1}{n-1}$ yet

$$|f(x_n) - f(y_n)| = \left|\frac{1}{\frac{1}{n}} - \frac{1}{\frac{2}{n}}\right| = \left|n - \frac{n}{2}\right| = \frac{n}{2} \geq 1.$$

Hence, for $\epsilon < 1$, no $\delta > 0$ can work in Definition 4.5.4 for f. Thus f is not uniformly continuous as claimed.

However, we want to work with compact domains since that enables the existence of the sup metric. Consequently, we note the notions of continuity and uniform continuity are the same on compact spaces.

Theorem 4.5.6. *Let (X, d_X) and (Y, d_Y) be metric spaces and let $f : X \to Y$ be continuous. If X is compact, then f is uniformly continuous.*

Proof. To see that f is uniformly continuous, let $\epsilon > 0$ be arbitrary. Since f is continuous at each point in (X, d_X), for each $x_0 \in X$ there exists a $U_{x_0} \in \mathcal{T}$ such that if $x \in U_{x_0}$ then $d_Y(f(x), f(x_0)) < \frac{\epsilon}{2}$. Thus $\{U_x\}_{x \in X}$ is an open cover of (X, d). Therefore, since (X, d) is compact and thus sequentially compact by the Borel-Lebesgue Theorem (Theorem 4.3.15), Lemma 4.3.14 implies there exists a Lebesgue number $\epsilon_0 > 0$ for $\{U_x\}_{x \in X}$; that is, for any $0 < \delta < \epsilon_0$ we choose, then for any $x \in X$ there exists an $x_0 \in X$ such that $B_{d_X}(x, \delta) \subseteq U_{x_0}$.

We claim that δ works for ϵ in Definition 4.5.4 for f. To see this, let $x_1, x_2 \in X$ be arbitrary points such that $d_X(x_1, x_2) < \delta$. Hence, by the definition of δ, there exists an $x_0 \in X$ such that $B_{d_X}(x_1, \delta) \subseteq U_{x_0}$. Thus, as $d_X(x_1, x_2) < \delta$, $x_1, x_2 \in B_{d_X}(x_1, \delta) \subseteq U_{x_0}$ so the definition of U_{x_0} implies that

$$d_Y(f(x_1), f(x_2)) \leq d_Y(f(x_1), f(x_0)) + d_Y(f(x_0), f(x_2)) < \frac{\epsilon}{2} + \frac{\epsilon}{2} = \epsilon.$$

Hence, as $x_1, x_2 \in X$ were arbitrary, δ works for ϵ in Definition 4.5.4 for f. Therefore, as $\epsilon > 0$ was arbitrary, f is uniformly continuous as desired. ∎

With the above notions complete, we can now prove the main theorem of this section.

Theorem 4.5.7 (Weierstrass Approximation Theorem). *The set of polynomials is dense in* $(\mathcal{C}[a,b], \|\cdot\|_\infty)$.

Proof. By Lemma 4.5.2 we may assume without loss of generality that $a = 0$ and $b = 1$.

Let $g \in \mathcal{C}[0,1]$ be arbitrary. Define the function $f : [0,1] \to \mathbb{R}$ by

$$f(x) = g(x) - (g(0) + (g(1) - g(0))x)$$

for all $x \in [0,1]$. Clearly $f \in \mathcal{C}[0,1]$ and $f(0) = f(1) = 0$. We will demonstrate there exists a sequence $(p_n)_{n\geq 1}$ of polynomials such that

$$\lim_{n\to\infty} \|f - p_n\|_\infty = 0.$$

This will complete the proof as $r_n(x) = p_n(x) + (g(0) + (g(1) - g(0))x)$ are polynomials such that $\lim_{n\to\infty} \|g - r_n\|_\infty = 0$.

To see that f is a uniform limit of polynomials on $[0,1]$, let $\epsilon > 0$ be arbitrary. First note that as $f \in \mathcal{C}[0,1]$ and $f(0) = 0 = f(1)$, we can extend f to be a continuous function on $\mathbb{R}$ by defining $f(x) = 0$ for all $x \in (-\infty, 0) \cup (1, \infty)$. Since f is then continuous on $[-2,2]$, f is uniformly continuous on $[-2,2]$ by Theorem 4.5.6 so there exists a $0 < \delta < 1$ such that if $x \in [-1,1]$ and $|t| < \delta$ then

$$|f(x+t) - f(x)| < \frac{1}{2}\epsilon.$$

Notice for each $n \in \mathbb{N}$ that

$$\int_{-1}^{1} (1-x^2)^n \, dx > 0$$

as $(1-x^2)^n > 0$ for all $x \in (-1,1)$. Hence for each $n \in \mathbb{N}$ there exists a $c_n > 0$ such that

$$c_n \int_{-1}^{1} (1-x^2)^n \, dx = 1.$$

Therefore, by Lemma 4.5.3,

$$\begin{aligned}
\frac{1}{c_n} &= \int_{-1}^{1} (1-x^2)^n \, dx \\
&= 2\int_{0}^{1} (1-x^2)^n \, dx \\
&\geq 2\int_{0}^{\frac{1}{\sqrt{n}}} (1-x^2)^n \, dx \\
&\geq 2\int_{0}^{\frac{1}{\sqrt{n}}} 1 - nx^2 \, dx \\
&= 2\left(x - \frac{n}{3}x^3\right)\Big|_{x=0}^{\frac{1}{\sqrt{n}}} \\
&= \frac{4}{3\sqrt{n}} \geq \frac{1}{\sqrt{n}}.
\end{aligned}$$

Hence $0 < c_n \leq \sqrt{n}$ for all $n \in \mathbb{N}$.

For each $n \in \mathbb{N}$ define $q_n : \mathbb{R} \to \mathbb{R}$ by

$$q_n(x) = c_n(1 - x^2)^n$$

for all $x \in \mathbb{R}$. Thus $q_n(x) \geq 0$ for all $x \in [-1, 1]$ and

$$\int_{-1}^{1} q_n(x)\, dx = 1$$

by the definition of c_n. Notice by the definition of q_n that if $x \in [-1, -\delta] \cup [\delta, 1]$, then

$$q_n(x) = c_n(1 - x^2)^n \leq c_n(1 - \delta^2)^n \leq \sqrt{n}(1 - \delta^2)^n.$$

For each $n \in \mathbb{N}$, define the function $f * q_n : [0, 1] \to \mathbb{R}$ by

$$(f * q_n)(x) = \int_{-1}^{1} f(x + t) q_n(t)\, dt.$$

Due to the translation invariance of the Riemann (Lebesgue) integral, for all $x \in [0, 1]$ we see using the substitution $u = x + t$ that

$$\begin{aligned}
(f * q_n)(x) &= \int_{-1}^{1} f(x + t) q_n(t)\, dt \\
&= \int_{-x}^{1-x} f(x + t) q_n(t)\, dt \qquad f \text{ is 0 except on } [0, 1] \\
&= \int_{0}^{1} f(u) q_n(u - x)\, du.
\end{aligned}$$

Thus, as $q_n(u - x)$ is a polynomial in x with coefficients being continuous functions in u, $f(u)q_n(u-x)$ is a polynomial with coefficients being continuous functions in u. Hence integrating $f(u)q_n(u - x)$ is performed by integrating the coefficients thereby resulting in a polynomial. Hence $f * q_n$ is a polynomial on $[0, 1]$.

Finally, we claim that $\lim_{n\to\infty} \|(f * q_n) - f\|_\infty = 0$. To see this, note for each $x \in [0,1]$ that

$$\begin{aligned}
&|(f * q_n)(x) - f(x)| \\
&= \left| \int_{-1}^{1} f(x+t) q_n(t)\, dt - f(x) \right| \\
&= \left| \int_{-1}^{1} f(x+t) q_n(t)\, dt - f(x) \int_{-1}^{1} q_n(t)\, dt \right| \quad \text{as } \int_{-1}^{1} q_n(x)\, dx = 1 \\
&= \left| \int_{-1}^{1} (f(x+t) - f(x)) q_n(t)\, dt \right| \\
&\leq \int_{-1}^{1} |f(x+t) - f(x)| q_n(t)\, dt \quad \text{as } q_n(x) \geq 0 \text{ on } [-1,1] \\
&= \int_{[-1,-\delta]\cup[\delta,1]} |f(x+t) - f(x)| q_n(t)\, dt + \int_{-\delta}^{\delta} |f(x+t) - f(x)| q_n(t)\, dt \\
&\leq \int_{[-1,-\delta]\cup[\delta,1]} 2\left\|f\right\|_\infty \sqrt{n}(1-\delta^2)^n\, dt + \int_{-\delta}^{\delta} |f(x+t) - f(x)| q_n(t)\, dt \\
&= 4\sqrt{n}\left\|f\right\|_\infty (1-\delta^2)^n (1-\delta) + \int_{-\delta}^{\delta} |f(x+t) - f(x)| q_n(t)\, dt \\
&\leq 4\sqrt{n}\left\|f\right\|_\infty (1-\delta^2)^n (1-\delta) + \int_{-\delta}^{\delta} \frac{\epsilon}{2} q_n(t)\, dt \quad \text{by uniform continuity} \\
&\leq 4\sqrt{n}\left\|f\right\|_\infty (1-\delta^2)^n (1-\delta) + \frac{\epsilon}{2} \int_{-1}^{1} q_n(t)\, dt \\
&= 4\sqrt{n}\left\|f\right\|_\infty (1-\delta^2)^n (1-\delta) + \frac{\epsilon}{2}.
\end{aligned}$$

Therefore, as $0 < 1 - \delta^2 < 1$ so

$$\lim_{n\to\infty} 4\sqrt{n}\left\|f\right\|_\infty (1-\delta^2)^n (1-\delta) = 0,$$

we see that for sufficiently large n that $\left\|(f * q_n) - f\right\|_\infty < \epsilon$. Hence, as $\epsilon > 0$ was arbitrary, the result follows. ■

4.6 Stone-Weierstrass Theorem, Lattice Form

Although the Weierstrass Approximation Theorem (Theorem 4.5.7) is powerful, it is limited as we need not have the notion of polynomials on arbitrary compact topological spaces. In this section, we will develop one of two theorems which will produce dense subsets of $\mathcal{C}(X)$. The theorem of this section, the lattice form of the Stone-Weierstrass Theorem (Theorem 4.6.14), will be motivated by a poset structure on $\mathcal{C}(X, Y)$. As such, we will only consider $Y = \mathbb{R}$ as this guarantees the existence of the following poset structure.

Given two functions $f, g \in \mathcal{C}(X)$ for some compact topological space $(X, \mathcal{T})$, is it easy to define a poset structure on $\mathcal{C}(X)$ by defining $f \leq g$ if and only if $f(x) \leq g(x)$ for all $x \in X$. However this poset structure is something stronger in that maximums and minimums occur.

Definition 4.6.1. Let $(X, \mathcal{T})$ be a compact topological space and let $f, g \in \mathcal{C}(X)$. The functions $f \vee g, f \wedge g : X \to \mathbb{R}$ defined by

$$\begin{aligned}(f \vee g)(x) &= \max(\{f(x), g(x)\}) = \frac{1}{2}f(x) + \frac{1}{2}g(x) - \frac{1}{2}|f(x) - g(x)| \\ (f \wedge g)(x) &= \min(\{f(x), g(x)\}) = -((-f) \vee (-g))(x)\end{aligned}$$

for all $x \in X$ are continuous functions (as they are a combination of compositions, sums, and scalar multiples of continuous functions) called the *maximum* and *minimum functions* respectively.

It is elementary to see that $f \vee g$ is the smallest function in $\mathcal{C}(X)$ that is larger than both f and g and $f \wedge g$ is the largest function in $\mathcal{C}(X)$ that is smaller than both f and g. In particular, we will be interested in the following subspaces of $\mathcal{C}(X)$.

Definition 4.6.2. Let $(X, \mathcal{T})$ be a compact topological space. A vector subspace $\mathcal{F} \subseteq \mathcal{C}(X)$ is said to be a *lattice* if $f \vee g \in \mathcal{F}$ for all $f, g \in \mathcal{F}$.

Example 4.6.3. It is elementary to see that $\mathcal{C}(X)$ is lattice of $\mathcal{C}(X)$.

Example 4.6.4. Consider the vector subspace $\mathcal{F}$ of $\mathcal{C}[a, b]$ consisting of all piecewise linear functions; that is,

$$\mathcal{F} = \left\{ f : [a, b] \to \mathbb{R} \,\middle|\, \substack{f \in \mathcal{C}[a,b] \text{ and there exists a partition } \{t_k\}_{k=0}^n \\ \text{such that } f \text{ is linear on } [t_{k-1}, t_k] \text{ for all } k.} \right\}.$$

It is not difficult to check that $\mathcal{F}$ is lattice in $\mathcal{C}[a, b]$ (i.e. the max of two linear functions is piecewise linear and the union of two partitions is a partition). In fact, it is not difficult to check that $\mathcal{F}$ is the smallest lattice in $\mathcal{C}[a, b]$ that contains the functions $f(x) = x$ and $g(x) = 1$ for all $x \in [a, b]$.

Remark 4.6.5. The reason we require a lattice of $\mathcal{C}(X)$ to be a vector subspace is that we know $\mathcal{C}(X)$ is a vector space so it would be difficult to approximate a general function from a set that is not a vector subspace of $\mathcal{C}(X)$. This is not too much of a restriction since we may always take the span of a given a subset of $\mathcal{C}(X)$.

Remark 4.6.6. As it is not difficult to see that $(-f) \vee (-g) = -(f \wedge g)$, we have that any lattice in $\mathcal{C}(X)$ is closed under taking the maximum and minimum of the functions it contains (as lattices are subspaces and thus closed under scalar multiplication).

Of course, not every lattice can be dense in $\mathcal{C}(X)$ as the constant functions are clearly a lattice that is closed with respect to the sup metric. To avoid such lattices, we consider the following property.

Definition 4.6.7. Let $(X, \mathcal{T})$ be a compact topological and let (Y, d) be a metric space. A collection of continuous functions $\mathcal{F} \subseteq \mathcal{C}(X, Y)$ is said to *separate points* if for all $x_1, x_2 \in X$ with $x_1 \neq x_2$ there exists a function $f \in \mathcal{F}$ such that $f(x_1) \neq f(x_2)$.

Of course, many examples of set of functions that separate points.

Example 4.6.8. The piecewise linear functions on $C[a, b]$ separate points as the function $f(x) = x$ is (piecewise) linear and clearly separates points.

Example 4.6.9. If (X, d) is a compact metric space, then $\mathcal{C}(X)$ separate points. Indeed if $x_1, x_2 \in X$ are such that $x_1 \neq x_2$, then the function $f : X \to \mathbb{R}$ defined by $f(x) = d(x, x_1)$ for all $x \in X$ is continuous, $f(x_1) = 0$, and $f(x_2) = d(x_1, x_2) > 0$.

Using the above example, we may obtain the following.

Proposition 4.6.10. *If (X, d) be a compact metric space and $\mathcal{F}$ be a dense subset of $\mathcal{C}(X)$, then $\mathcal{F}$ separate points.*

Proof. Let $\mathcal{F}$ be a dense subset of $\mathcal{C}(X)$. To see that $\mathcal{F}$ separate points, let $x_1, x_2 \in X$ such that $x_1 \neq x_2$ be arbitrary. Define the function $f : X \to \mathbb{R}$ by $f(x) = d(x, x_1)$ for all $x \in X$. Clearly f is continuous, $f(x_1) = 0$ and $f(x_2) = d(x_1, x_2) > 0$.

Let $\epsilon = \frac{1}{3}d(x_1, x_2) > 0$. Since $\mathcal{F}$ is dense in $\mathcal{C}(X)$, there exists a $g \in \mathcal{F}$ such that $\|f - g\|_\infty < \epsilon$. Hence

$$\begin{aligned}
3\epsilon = d(x_2, x_1) &= |f(x_1) - f(x_2)| \\
&\leq |f(x_1) - g(x_1)| + |g(x_1) - g(x_2)| + |g(x_2) - f(x_2)| \\
&\leq 2\epsilon + |g(x_1) - g(x_2)|
\end{aligned}$$

Hence $|g(x_1) - g(x_2)| \geq \epsilon > 0$ so $g(x_1) \neq g(x_2)$. Hence, as $x_1, x_2 \in X$ were arbitrary, we obtain that $\mathcal{F}$ separate points. ∎

Remark 4.6.11. Of course, not every dense subset of $\mathcal{C}(X)$ separate points for a general compact topological space $(X, \mathcal{T})$. Indeed, $\mathcal{C}(X)$ need not separate points for a general compact topological space $(X, \mathcal{T})$. To see this, consider the trivial topology $\mathcal{T}$ on any set X with at least two points. Then $(X, \mathcal{T})$ is compact as $\mathcal{T}$ is finite so $\mathcal{C}(X)$ is precisely the constant functions and thus does not separate points.

Of course, the reason the $\mathcal{C}(X)$ above does not separate points is that the topology $\mathcal{T}$ does not distinguish points; that is, there are points $x_1, x_2 \in X$ such that the neighbourhoods of x_1 and x_2 are precisely the same. This raises the question, "For which topological space $(X, \mathcal{T})$ is $\mathcal{C}(X)$ point separating?" Unsurprisingly, we turn to the notion of a Hausdorff topological space.

Proposition 4.6.12. *Let $(X, \mathcal{T})$ be a compact topological space. If $\mathcal{C}(X)$ separates points, then $(X, \mathcal{T})$ is Hausdorff.*

Proof. To see that $(X, \mathcal{T})$ is Hausdorff, let $x_1, x_2 \in X$ be such that $x_1 \neq x_2$. Since $\mathcal{C}(X)$ is point separating, there exists a continuous function $f \in \mathcal{C}(X)$ such that $f(x_1) \neq f(x_2)$. Since $\mathbb{R}$ is Hausdorff, there exists neighbourhoods

U_1 and U_2 in $\mathbb{R}$ such that $f(x_1) \in U_1$, $f(x_2) \in U_2$, and $U_1 \cap U_2 = \emptyset$. Hence, as f is continuous, $f^{-1}(U_1)$ and $f^{-1}(U_2)$ are disjoint neighbourhoods of x_1 and x_2 respectively. Therefore, since $x_1, x_2 \in X$ were arbitrary, $(X, \mathcal{T})$ is Hausdorff. ∎

Remark 4.6.13. Of course Proposition 4.6.12 raises the question, "If $(X, \mathcal{T})$ is a compact Hausdorff topological space, does $\mathcal{C}(X)$ separate points?" This seems like a difficult question. Indeed given two points in X, we know $\mathcal{T}$ being a Hausdorff topology on X means we can find disjoint neighbourhoods of these points, but how would we construct a continuous function that separates these two points? This very difficult question will be postponed until Chapter 5 where many different notions of separability for topological spaces are discussed and many important results are obtained. In particular, Urysohn's Lemma (Theorem 5.2.1) implies that $\mathcal{C}(X)$ separates points provided $(X, \mathcal{T})$ is a compact Hausdorff topological space thereby completely resolving this question!

For now, even though we do not know which $(X, \mathcal{T})$ have the property that $\mathcal{C}(X)$ separate points, we can demonstrate the following collections of sets are dense in $\mathcal{C}(X)$. Of course, if $\mathcal{C}(X)$ does not separate points, we cannot find an $\mathcal{F}$ to apply the following theorem. As such, by Proposition 4.6.12, any version of the Stone-Weierstrass Theorem will only be applicable for compact Hausdorff topological spaces.

Theorem 4.6.14 (Stone-Weierstrass Theorem - Lattice Version). *Let $(X, \mathcal{T})$ be compact Hausdorff topological space and let $\mathcal{F} \subseteq \mathcal{C}(X)$ be a vector subspace such that*

(1) $1 \in \mathcal{F}$ (the constant function that is one everywhere),

(2) $\mathcal{F}$ separates points, and

(3) $\mathcal{F}$ is a lattice.

Then $\mathcal{F}$ is dense in $(\mathcal{C}(X), \|\cdot\|_\infty)$.

Proof. First we claim that for all $x_1, x_2 \in X$ with $x_1 \neq x_2$ and for all $\alpha, \beta \in \mathbb{R}$ there exists a function $h \in \mathcal{F}$ such that

$$h(x_1) = \alpha \qquad \text{and} \qquad h(x_2) = \beta.$$

To see this, let $x_1, x_2 \in X$ with $x_1 \neq x_2$ be arbitrary. Since $\mathcal{F}$ separates points, there exists a function $g \in \mathcal{F}$ such that $g(x_1) \neq g(x_2)$. Hence if we define $h : X \to \mathbb{R}$ by

$$h(x) = \alpha + \frac{\beta - \alpha}{g(x_2) - g(x_1)}(g(x) - g(x_1)),$$

for all $x \in X$, then clearly $h \in \mathcal{F}$ as $\mathcal{F}$ is a subspace and $1 \in \mathcal{F}$, and $h(x_1) = \alpha$ and $h(x_2) = \beta$ as desired. We will use these functions to uniformly approximate any function f in $\mathcal{C}(X)$.

To prove that $\mathcal{F}$ is dense in $(\mathcal{C}(X), \|\cdot\|_\infty)$, let $f \in \mathcal{C}(X)$ be arbitrary and let $\epsilon > 0$. To begin, we will demonstrate that for each $z \in X$ there exists a function $h_z \in \mathcal{F}$ such that $h_z(z) = f(z)$ and $h_z(x) < f(x) + \epsilon$ for all $x \in X$.

To see this, fix $z \in X$. By the above paragraph for each $y \in X$ there exists a $h_{z,y} \in \mathcal{F}$ such that $h_{z,y}(z) = f(z)$ and $h_{z,y}(y) = f(y)$. Since the function $h_{z,y} - f$ is continuous and $h_{z,y}(y) - f(y) = 0$, there exists a open set U_y containing y such that $h_{z,y}(x) - f(x) < \epsilon$ for all $x \in U_y$. However, since $\{U_y\}_{y \in X}$ is an open cover of $(X, \mathcal{T})$ and as $(X, \mathcal{T})$ is compact, there exists an $n \in \mathbb{N}$ and $y_1, \ldots, y_n \in X$ such that $X = \bigcup_{k=1}^n U_{y_k}$. Let

$$h_z = h_{z,y_1} \wedge h_{z,y_2} \wedge \cdots \wedge h_{z,y_n},$$

which is an element of $\mathcal{F}$ as $\mathcal{F}$ is a lattice. In addition, as $h_{z,y_k}(z) = f(z)$ for all $k \in \{1, \ldots, n\}$, we clearly see that $h_z(z) = f(z)$. Moreover if $x \in X$ then there exists a $k_0 \in \{1, \ldots, n\}$ such that $x \in U_{y_{k_0}}$ and thus

$$h_z(y) \leq h_{z,y_{k_0}}(x) < f(x) + \epsilon.$$

Hence, as $x \in X$ was arbitrary, h_z has the desired properties.

We may now use the $h_z \in \mathcal{F}$ along with a similar technique to obtain an $h \in \mathcal{F}$ such that $\|f - h\|_\infty \leq \epsilon$. To see this, notice for each $z \in X$ that $h_z - f$ is continuous and $h_z(z) - f(z) = 0$ so there exists an open set $V_z \in \mathcal{T}$ containing z such that $h_z(x) - f(x) > -\epsilon$ for all $x \in V_z$. However, since $\{V_z\}_{z \in X}$ is an open cover of $(X, \mathcal{T})$ and as $(X, \mathcal{T})$ is compact, there exists $z_1, \ldots, z_m \in X$ such that $X = \bigcup_{k=1}^m V_{z_k}$. Let

$$h = h_{z_1} \vee h_{z_2} \vee \cdots \vee h_{z_m},$$

which is an element of $\mathcal{F}$ as $\mathcal{F}$ is a lattice. Furthermore, as $h_{z_k}(x) < f(x) + \epsilon$ for all $x \in X$ and for all $k \in \{1, \ldots, m\}$, we see that $h(x) < f(x) + \epsilon$ for all $x \in X$ by the definition of the maximum. Furthermore, if $x \in X$ then there exists a $k_0 \in \{1, \ldots, m\}$ such that $x \in V_{z_{k_0}}$ and thus

$$h(x) \geq h_{z_k}(x) > f(x) - \epsilon.$$

Therefore, as $x \in X$ was arbitrary, we have that

$$f(x) - \epsilon < h(x) < f(x) + \epsilon$$

for all $x \in X$. Hence $\|h - f\|_\infty \leq \epsilon$. Therefore, as $\epsilon > 0$ and $f \in \mathcal{C}(X)$ were arbitrary, the result follows. ■

Note Theorem 4.6.14, Example 4.6.4, and Example 4.6.8 imply that the piecewise linear functions on $\mathcal{C}[a,b]$ are dense in $\mathcal{C}[a,b]$. Of course, one could verify the density of piecewise linear functions in $\mathcal{C}[a,b]$ directly using uniform continuity. Indeed given $f \in \mathcal{C}[a,b]$ and an $\epsilon > 0$, choose the δ from uniform continuity. Then choose a partition with intervals of length at most δ and define a piecewise linear function g that takes the values that f does at each end of each interval in the partition. Uniform continuity and piecewise linearity will then implies that $\|f - g\|_\infty \leq 2\epsilon$.

4.7 Stone-Weierstrass Theorem, Subalgebra Form

Of course constructing a subspace that is a lattice in $\mathcal{C}(X)$ may not be an easy task as making a subspace closed under maximum and minimum may not be an easy task. In this section, we will discuss another version of the Stone-Weierstrass Theorem (Theorem 4.7.5) that is very easy to verify the conditions for in general. Furthermore, note the lattice version of the Stone-Weierstrass Theorem (Theorem 4.6.14) cannot possibly extend to complex-valued functions as there is no natural ordering on $\mathbb{C}$. However, using our new version of the Stone-Weierstrass Theorem (Theorem 4.7.5) we will be able to develop a version of the Stone-Weierstrass Theorem for complex-valued functions (Theorem 4.7.7). Subsequently, we will also be able to extend these results to continuous functions that vanish at infinity on a locally compact topological space (Theorem 4.7.10).

To replace the lattice structure for dense subsets, we will consider the following structure.

Definition 4.7.1. Let $(X, \mathcal{T})$ be a compact topological space. A vector subspace $\mathcal{A} \subseteq \mathcal{C}(X, \mathbb{K})$ is said to be a *subalgebra of* $\mathcal{C}(X, \mathbb{K})$ if whenever $f, g \in \mathcal{A}$ it is the case that $fg \in \mathcal{A}$.

Example 4.7.2. Clearly $\mathcal{C}(X, \mathbb{K})$ is a subalgebra of $\mathcal{C}(X, \mathbb{K})$ and any ideal of $\mathcal{C}(X, \mathbb{K})$ is a subalgebra of $\mathcal{C}(X, \mathbb{K})$. Furthermore, it is clear that the polynomials are a subalgebra of $C[a,b]$.

Example 4.7.3. Let
$$\mathbb{T} = \{z \in \mathbb{C} \mid |z| = 1\}.$$
The *trigonometric polynomials* on $\mathbb{T}$ is the subset of $\mathcal{C}(\mathbb{T}, \mathbb{C})$ defined by
$$Trig(\mathbb{T}) = \mathrm{span}_{\mathbb{C}}(\{f_n : \mathbb{T} \to \mathbb{C} \mid n \in \mathbb{Z}, f_n(z) = z^n \text{ for all } z \in \mathbb{T}\}).$$
As $f_n f_m = f_{n+m}$ for all $n, m \in \mathbb{Z}$, clearly $Trig(\mathbb{T})$ is a subalgebra of $\mathcal{C}(\mathbb{T}, \mathbb{C})$.

To see why these are called the trigonometric polynomials, recall if $z \in \mathbb{T}$ then $z = e^{i\theta}$ for some $\theta \in [0, 2\pi)$. Hence
$$f_n(z) = e^{in\theta} = \cos(n\theta) + i\sin(n\theta)$$

for all $n \in \mathbb{Z}$. Therefore, as

$$\frac{1}{2}(f_n(z) + f_{-n}(z)) = \cos(n\theta) \qquad \text{and} \qquad \frac{1}{2i}(f_n(z) - f_{-n}(z)) = \sin(n\theta),$$

we see that

$$Trig(\mathbb{T}) = \operatorname{span}_{\mathbb{C}}(\{\cos(n\theta), \sin(n\theta) : [0, 2\pi] \to \mathbb{C} \mid n \in \mathbb{N} \cup \{0\}\}).$$

This is why $Trig(\mathbb{T})$ is called the trigonometric polynomials.

As the closure of a subspace is a subspace and as we appear to want to show that specific algebras are dense in $\mathcal{C}(X, \mathbb{K})$, which is an algebra, it is not difficult to believe the closure of a subalgebra is a subalgebra.

Lemma 4.7.4. *Let $(X, \mathcal{T})$ be a compact topological space and let $\mathcal{A}$ be a subalgebra of $\mathcal{C}(X, \mathbb{K})$. Then the closure of $\mathcal{A}$ in $(\mathcal{C}(X, \mathbb{K}), \|\cdot\|)_\infty)$ is a subalgebra of $\mathcal{C}(X, \mathbb{K})$*

Proof. To begin, notice for all functions $f, g \in \mathcal{C}(X, \mathbb{K})$ that

$$\|fg\|_\infty \leq \|f\|_\infty \|g\|_\infty$$

due to the definition of the infinity norm.

To see that $\overline{\mathcal{A}}$ is a subalgebra, let $f, g \in \overline{\mathcal{A}}$ be arbitrary. By Theorem 1.6.21 $n \in \mathbb{N}$ there exist $f_n, g_n \in \mathcal{A}$ such that

$$\|f - f_n\|_\infty < \frac{1}{n} \qquad \text{and} \qquad \|g - g_n\|_\infty < \frac{1}{n}.$$

Hence there exists sequences $(f_n)_{n\geq 1}$ and $(g_n)_{n\geq 1}$ of functions in $\mathcal{A}$ such that

$$\lim_{n\to\infty} \|f - f_n\|_\infty = 0 = \lim_{n\to\infty} \|g - g_n\|_\infty .$$

Clearly for all $\alpha \in \mathbb{K}$ the sequence $(\alpha f_n + g_n)_{n\geq 1}$ consists of elements of $\mathcal{A}$ as $\mathcal{A}$ is a subspace and

$$\lim_{n\to\infty} \|(\alpha f + g) - (\alpha f_n + g_n)\|_\infty \leq \limsup_{n\to\infty} |\alpha| \|f - f_n\|_\infty + \|g - g_n\|_\infty = 0.$$

Therefore $\alpha f + g \in \overline{\mathcal{A}}$ so $\overline{\mathcal{A}}$ is subspace. To see that $\overline{\mathcal{A}}$ is a subalgebra, notice the sequence $(f_n g_n)_{n\geq 1}$ consists of elements of $\mathcal{A}$ as $\mathcal{A}$ is subalgebra. Since $\sup_{n\geq 1} \|f_n\|_\infty < \infty$ as $0 = \lim_{n\to\infty} \|f - f_n\|_\infty$, we obtain that

$$\begin{aligned} \lim_{n\to\infty} \|fg - f_n g_n\|_\infty &\leq \limsup_{n\to\infty} \|fg - f_n g\|_\infty + \|f_n g - f_n g_n\|_\infty \\ &\leq \limsup_{n\to\infty} \|f - f_n\|_\infty \|g\|_\infty + \|f_n\|_\infty \|g - g_n\|_\infty \\ &= 0. \end{aligned}$$

Therefore $fg \in \overline{\mathcal{A}}$ so $\overline{\mathcal{A}}$ is a subalgebra. ∎

Using nothing but Lemma 4.7.4, the Weierstrass Approximation Theorem (Theorem 4.5.7), and the lattice version of the Stone-Weierstrass Theorem (Theorem 4.6.14), we obtain the following Stone-Weierstrass Theorem with next to no difficulty.

Theorem 4.7.5 (Stone-Weierstrass Theorem - Algebra Version). *Let $(X, \mathcal{T})$ be compact Hausdorff topological space and let $\mathcal{A} \subseteq \mathcal{C}(X)$ be a subalgebra such that*

(1) $1 \in \mathcal{A}$ *and*

(2) $\mathcal{A}$ *separates points.*

Then $\mathcal{A}$ *is dense in* $(\mathcal{C}(X), \|\cdot\|_\infty)$.

Proof. By Lemma 4.7.4 it is clear that $\overline{\mathcal{A}}$ is a closed subalgebra of $\mathcal{C}(X)$ that contains one and separates points. Our goal is to prove $\overline{\mathcal{A}} = \mathcal{C}(X)$.

First we claim that if $f \in \overline{\mathcal{A}}$ then $|f| \in \overline{\mathcal{A}}$. To see this, consider the function $a : [-\|f\|_\infty, \|f\|_\infty] \to \mathbb{R}$ defined by $a(x) = |x|$ for all $x \in [-\|f\|_\infty, \|f\|_\infty]$. As $a \in \mathcal{C}[-\|f\|_\infty, \|f\|_\infty]$, the Weierstrass Approximation Theorem (Theorem 4.5.7) implies there exists a sequence of polynomials p_n such that $\lim_{n\to\infty} \|p_n - a\|_\infty = 0$ as continuous functions on $[-\|f\|_\infty, \|f\|_\infty]$. Hence, as $f : X \to [-\|f\|_\infty, \|f\|_\infty]$, we see that

$$\lim_{n\to\infty} \|p_n \circ f - a \circ f\|_\infty$$

as continuous functions on X. Clearly $a \circ f = |f|$. Moreover, notice for any polynomial $p(x) = a_m x^m + a_{m-1}x^{m-1} + \cdots + a_1 x + a_0$ where $a_0, a_1, \ldots, a_m \in \mathbb{R}$ that

$$p \circ f = a_m f^m + a_{m-1} f^{m-1} + \cdots + a_1 f + a_0 1 \in \overline{\mathcal{A}}$$

as $1, f \in \overline{\mathcal{A}}$ and $\overline{\mathcal{A}}$ is a subalgebra. Hence we see that $p_n \circ f \in \mathcal{A}$ for all $n \in \mathbb{N}$ and hence $|f| \in \overline{\mathcal{A}}$.

Next let $f, g \in \overline{\mathcal{A}}$ be arbitrary. Then $f + g, f - g \in \overline{\mathcal{A}}$ as $\overline{\mathcal{A}}$ is subspace so $|f - g| \in \overline{\mathcal{A}}$ by the above paragraph. Hence, as $\overline{\mathcal{A}}$ is a subspace, we see that

$$f \vee g = \frac{1}{2}(f + g) + \frac{1}{2}|f - g| \in \overline{\mathcal{A}}.$$

Hence $\overline{\mathcal{A}}$ is a lattice that contains one and separates points. Therefore the lattice form of the Stone-Weierstrass Theorem (Theorem 4.6.14 implies that $\overline{\mathcal{A}}$ is dense in $\mathcal{C}(X)$. Therefore, as $\overline{\mathcal{A}}$ is closed, we obtain that $\overline{\mathcal{A}} = \mathcal{C}(X)$ so $\mathcal{A}$ is dense in $(\mathcal{C}(X), \|\cdot\|_\infty)$ as desired. ∎

To emphasize the simplicity of applying Stone-Weierstrass Theorem (Theorem 4.7.5), we note the following.

Example 4.7.6. As it is not difficult to verify that

$$\mathcal{A} = \operatorname{span}\{x^n \mid n \in \{0, 3, 6, 9, \ldots\}\}$$

is a subalgebra of $\mathcal{C}[0,1]$ that separates points and contains 1, the Stone-Weierstrass Theorem (Theorem 4.7.5) implies that $\mathcal{A}$ is dense in $\mathcal{C}[0,1]$.

It is not difficult to develop a version of the Stone-Weierstrass Theorem (Theorem 4.7.5) for complex-valued functions now. To do so, recall that if $f \in \mathcal{C}(X, \mathbb{C})$, then the function $\overline{f} : X \to \mathbb{C}$ defined by $\overline{f}(x) = \overline{f(x)}$ (the complex conjugate) is a continuous function being the composition of two continuous functions.

Theorem 4.7.7 (Stone-Weierstrass Theorem - Complex Version). *Let $(X, \mathcal{T})$ be compact Hausdorff topological space and let $\mathcal{A} \subseteq \mathcal{C}(X, \mathbb{C})$ be a subalgebra such that*

(1) $1 \in \mathcal{A}$,

(2) $\mathcal{A}$ *separates points, and*

(3) $\overline{f} \in \mathcal{A}$ *whenever* $f \in \mathcal{A}$.

Then $\mathcal{A}$ *is dense in* $(\mathcal{C}(X, \mathbb{C}), \|\cdot\|_\infty)$.

Proof. Consider the set

$$\mathcal{A}_0 = \{f \in \mathcal{A} \mid f(X) \subseteq \mathbb{R}\}.$$

Clearly $\mathcal{A}_0$ is a subalgebra of $\mathcal{C}(X, \mathbb{R})$ that contains the constant function 1. We claim that $\mathcal{A}_0$ separates points. To see this, let $x_1, x_2 \in X$ with $x_1 \neq x_2$ be arbitrary. Since $\mathcal{A}$ separates points, there exists an $f \in \mathcal{A}$ such that $f(x_1) \neq f(x_2)$. Hence it must be the case that either $\operatorname{Re}(f)(x_1) \neq \operatorname{Re}(f)(x_2)$ or $\operatorname{Im}(f)(x_1) \neq \operatorname{Im}(f)(x_2)$. Since $\operatorname{Re}(f), \operatorname{Im}(f) \in \mathcal{A}_0$ as $\mathcal{A}$ is a subspace closed under complex conjugates, we obtain that $\mathcal{A}_0$ separates points.

By the algebra version of the Stone-Weierstrass Theorem (Theorem 4.7.5), we obtain that $\mathcal{A}_0$ is dense in $(\mathcal{C}(X, \mathbb{R}), \|\cdot\|_\infty)$. To see that $\mathcal{A}$ is dense in $(\mathcal{C}(X, \mathbb{C}), \|\cdot\|_\infty)$, let $f \in \mathcal{C}(X, \mathbb{C})$ and let $\epsilon > 0$ be arbitrary. Since $\operatorname{Re}(f), \operatorname{Im}(f) \in \mathcal{C}(X, \mathbb{R})$ and since $\mathcal{A}_0$ is dense in $(\mathcal{C}(X, \mathbb{R}), \|\cdot\|_\infty)$, there exists $g_1, g_2 \in \mathcal{A}_0$ such that

$$\|\operatorname{Re}(f) - g_1\|_\infty < \frac{\epsilon}{2} \qquad \text{and} \qquad \|\operatorname{Im}(f) - g_2\|_\infty < \frac{\epsilon}{2}.$$

As $\mathcal{A}$ is a subalgebra over $\mathbb{C}$, we see that $g_1 + ig_2 \in \mathcal{A}$ and

$$\begin{aligned}\|f - (g_1 + ig_2)\|_\infty &= \|(\operatorname{Re}(f) + i\operatorname{Im}(f)) - (g_1 + ig_2)\|_\infty \\ &\leq \|\operatorname{Re}(f) - g_1\|_\infty + |i|\,\|\operatorname{Im}(f) - g_2\|_\infty < \frac{\epsilon}{2} + \frac{\epsilon}{2} = \epsilon.\end{aligned}$$

Hence, as $f \in \mathcal{C}(X, \mathbb{C})$ and $\epsilon > 0$ were arbitrary, we obtain that $\mathcal{A}$ is dense in $(\mathcal{C}(X, \mathbb{C}), \|\cdot\|_\infty)$ as desired. ∎

Again Theorem 4.7.7 is simple to use.

Example 4.7.8. The trigonometric polynomials are a dense subset of $(\mathcal{C}(\mathbb{T}, \mathbb{C}), \|\cdot\|_\infty)$. Indeed by Example 4.7.3 we know that $Trig(\mathbb{T})$ is a subalgebra of $\mathcal{C}(\mathbb{T}, \mathbb{C})$. Since $z^0 = 1$ for all $z \in \mathbb{T}$, clearly $1 \in Trig(\mathbb{T})$. Furthermore, as every $z \in \mathbb{T}$ can be written as $z = e^{i\theta}$ for some $\theta \in [0, 2\pi]$, we see that

$$\overline{z^n} = \overline{e^{in\theta}} = e^{-in\theta} = z^{-n}$$

for all $n \in \mathbb{Z}$ and $z \in \mathbb{T}$. Hence $Trig(\mathbb{T})$ is closed under complex conjugates. Finally, to see that $Trig(\mathbb{T})$ is point separating, we note that the function $f(z) = z$ for all $z \in \mathbb{T}$ is an element of $Trig(\mathbb{T})$ and clearly separate points. Hence, by the Stone-Weierstrass Theorem (Theorem 4.7.7), we obtain that $Trig(\mathbb{T})$ is dense in $\mathcal{C}(\mathbb{T}, \mathbb{C})$.

Finally, if we want to obtain a version of the Stone-Weierstrass Theorem for continuous functions that vanish at infinity for a locally compact Hausdorff topological space $(X, \mathcal{T})$, by Theorem 4.2.18 we should simply be able to study collections of functions on a compact topological spaces that vanish at a single point. Thus the following lemma should lead us to the answer.

Lemma 4.7.9. *Let $(X, \mathcal{T})$ be a compact Hausdorff topological space and let $\mathcal{A} \subseteq \mathcal{C}(X, \mathbb{K})$ be a subalgebra such that*

(1) there exists an $x_0 \in X$ such that $f(x_0) = 0$ for all $f \in \mathcal{A}$,

(2) $\mathcal{A}$ separates points, and

(3) $\overline{f} \in \mathcal{A}$ whenever $f \in \mathcal{A}$.

Then the closure of $\mathcal{A}$ in $(\mathcal{C}(X, \mathbb{K}), \|\cdot\|_\infty)$ is

$$\mathcal{A}_{x_0} = \{f \in \mathcal{C}(X, \mathbb{K}) \mid f(x_0) = 0\}.$$

Proof. Notice that $\mathcal{A}_{x_0}$ is closed in $(\mathcal{C}(X, \mathbb{K}), \|\cdot\|_\infty)$ as if a net $(f_\lambda)_{\lambda \in \Lambda}$ from $\mathcal{A}_{x_0}$ converges to a function $f \in \mathcal{C}(X, \mathbb{K})$ with respect to $\|\cdot\|_\infty$, then $(f_\lambda(x_0))_{\lambda \in \Lambda}$ converges to $f(x_0)$ in $\mathbb{K}$ and thus $f(x_0) = 0$ as $f_\lambda(x_0) = 0$ for all $\lambda \in \Lambda$ so $f \in \mathcal{A}_{x_0}$. As clearly $\mathcal{A} \subseteq \mathcal{A}_{x_0}$ by definition, the fact that $\mathcal{A}_{x_0}$ is closed implies that $\overline{\mathcal{A}} \subseteq \mathcal{A}_{x_0}$.

To see the other direction, first consider

$$\mathcal{B} = \{\alpha 1 + f \mid \alpha \in \mathbb{K} \text{ and } f \in \mathcal{A}\} \subseteq \mathcal{C}(X, \mathbb{K}).$$

Since $\mathcal{A}$ is a subalgebra of $\mathcal{C}(X, \mathbb{K})$ (and thus a vector subspace), it is clear that $\mathcal{B}$ is a subalgebra of $\mathcal{C}(X, \mathbb{K})$ that contains 1 (as $0 \in \mathcal{A}$ as $\mathcal{A}$ is a vector subspace). Furthermore, since $\mathcal{A}$ separates points and is closed under complex conjugation, $\mathcal{B}$ separates points and is closed under complete conjugation.

Hence the Stone-Weierstrass Theorem (Theorem 4.7.5 for $\mathbb{R}$ and Theorem 4.7.7 for $\mathbb{C}$) implies that $\mathcal{B}$ is dense in $(\mathcal{C}(X, \mathbb{K}), \|\cdot\|_\infty)$.

To see that $\mathcal{A}_{x_0} \subseteq \overline{\mathcal{A}}$, let $g \in \mathcal{A}_{x_0}$ and let $\epsilon > 0$ be arbitrary. Since $\mathcal{B}$ is dense in $(\mathcal{C}(X, \mathbb{K}), \|\cdot\|_\infty)$ there exists an $f \in \mathcal{A}$ and an $\alpha \in \mathbb{K}$ such that

$$\|g - (\alpha 1 + f)\|_\infty < \frac{\epsilon}{2}.$$

Since $g(x_0) = 0 = f(x_0)$, the definition of the infinity norm implies that

$$|\alpha| = |g(x_0) - (\alpha + f(x_0)| \leq \|g - (\alpha 1 + f)\|_\infty < \frac{\epsilon}{2}.$$

Hence, by the triangle inequality, we obtain that

$$\|g - f\|_\infty \leq \|g - (\alpha 1 + f)\|_\infty + \|\alpha 1\|_\infty < \frac{\epsilon}{2} + \frac{\epsilon}{2} = \epsilon.$$

Therefore, as $g \in \mathcal{A}_{x_0}$ and $\epsilon > 0$ were arbitrary, $\mathcal{A}_{x_0} \subseteq \overline{\mathcal{A}}$. Hence $\mathcal{A}_{x_0} = \overline{\mathcal{A}}$ as desired. ∎

Thus our knowledge of locally compact Hausdorff topological space immediately implies the following.

Theorem 4.7.10 (Stone-Weierstrass Theorem - Locally Compact Version). *Let $(X, \mathcal{T})$ be a locally compact Hausdorff topological space and let $\mathcal{A} \subseteq \mathcal{C}_0(X, \mathbb{K})$ be a subalgebra such that*

(1) for all $x \in X$ there exists a $f \in \mathcal{A}$ such that $f(x) \neq 0$,

(2) $\mathcal{A}$ separates points, and

(3) $\overline{f} \in \mathcal{A}$ whenever $f \in \mathcal{A}$.

Then $\mathcal{A}$ is dense in $(\mathcal{C}_0(X, \mathbb{K}), \|\cdot\|_\infty)$.

Proof. Let $(Y, \mathcal{T}_Y)$ be the one-point compactification of $(X, \mathcal{T})$ with $Y \setminus X = \{\infty\}$. By Theorem 4.2.18 we can view $(\mathcal{C}_0(X, \mathbb{K}), \|\cdot\|_\infty)$ as a subspace of $(\mathcal{C}(Y, \mathbb{K}), \|\cdot\|_\infty)$. However, inside of $(\mathcal{C}(Y, \mathbb{K}), \|\cdot\|_\infty)$, $\mathcal{A}$ is a subalgebra that separates points, is closed under complex conjugation, and has the property that $f(\infty) = 0$ for all $f \in \mathcal{A}$. Hence Lemma 4.7.9 implies that the closure of $\mathcal{A}$ in $(\mathcal{C}(Y, \mathbb{K}), \|\cdot\|_\infty)$ is precisely $\mathcal{C}_0(X, \mathbb{K})$. Hence, as $(\mathcal{C}_0(X, \mathbb{K}), \|\cdot\|_\infty)$ as a subspace of $(\mathcal{C}(Y, \mathbb{K}), \|\cdot\|_\infty)$, the result follows. ∎

Of course, the locally compact version of the Stone-Weierstrass Theorem (Theorem 4.7.10) has many potential applications. We note the following, which is actually easier to prove by hand using uniform continuity on compact sets but the following argument can be upgraded to other topological spaces with the material from Chapter 5.

Example 4.7.11. Let

$$\mathcal{C}_c(\mathbb{R}) = \left\{ f \in \mathcal{C}(\mathbb{R}) \mid \overline{\{x \in \mathbb{R} \mid f(x) \neq 0\}} \text{ is compact} \right\}.$$

The elements of $\mathcal{C}_c(\mathbb{R})$ are called the continuous, *compactly supported functions*. Clearly $\mathcal{C}_c(\mathbb{R}) \subseteq \mathcal{C}_0(\mathbb{R})$ by our knowledge of $\mathcal{C}_0(\mathbb{R})$. It is not difficult to see that $\mathcal{C}_c(\mathbb{R})$ is a subalgebra of $\mathcal{C}_0(\mathbb{R})$ as the union and intersection of two compact subsets of $\mathbb{R}$ are compact.

For each $x_0 \in \mathbb{R}$, let $f_{x_0} : \mathbb{R} \to \mathbb{R}$ be defined by

$$f_{x_0}(x) = \begin{cases} (x - x_0 + 1) & \text{if } x \in [x_0 - 1, x_0] \\ 1 - (x - x_0) & \text{if } x \in [x_0, x_0 + 1] \\ 0 & \text{otherwise} \end{cases}$$

for all $x \in \mathbb{R}$. Then clearly $\{f_{x_0}\}_{x_0 \in \mathbb{R}}$ have the properties that $f_{x_0}(x_0) = 1$, $f_{x_0}(x) \neq 1$ for all $x \neq x_0$ so $\{f_{x_0}\}_{x_0 \in \mathbb{R}}$ separates, and $\{f_{x_0}\}_{x_0 \in \mathbb{R}} \subseteq \mathcal{C}_c(\mathbb{R})$. Hence $\mathcal{C}_c(\mathbb{R})$ satisfies the assumptions of Stone-Weierstrass Theorem (Theorem 4.7.10) and thus is a dense subset of $\mathcal{C}_0(\mathbb{R})$.

Chapter 5

Separability Axioms and Theorems

The main question that remains from our study of the Stone-Weierstrass Theorems is whether or not $\mathcal{C}(X)$ separates points for an arbitrary compact Hausdorff topological space $(X, \mathcal{T})$. Of course the Stone-Weierstrass Theorems developed in Chapter 4 required a collection of functions to separate points in order for us to conclude they are dense in $(\mathcal{C}(X), \|\cdot\|_\infty)$ and knowing that $\mathcal{C}(X)$ separates points for an arbitrary compact Hausdorff topological space $(X, \mathcal{T})$ would make it necessary for a collection of functions to separate points in order to be dense in $(\mathcal{C}(X), \|\cdot\|_\infty)$. Thus, one of the main goals of this chapter is to show that $\mathcal{C}(X)$ separates points for an arbitrary compact Hausdorff topological space $(X, \mathcal{T})$.

To show that $\mathcal{C}(X)$ separates points causes us to delve into the world of constructing separation axioms on topological spaces. The idea of a separation axiom is to use various topologically motivated concepts, such as open sets or continuous functions, to separate out different points in a topological space. Each of these separation axioms has various strengths and thus raises the questions of which set of axioms is stronger and which topological spaces satisfy which axioms. This inevitably raises the questions of trying to 'classify' topological spaces by developing some sort of axioms and invariants that distinguish topological spaces upto isomorphism. However, it has been mathematically proven that the problem of classifying topological spaces with simple invariants is impossible. As such, our goal is not to try and classify topological spaces, but use the properties of specific topological spaces to obtain as much information and useful results as possible on these spaces.

After developing some basic separation axioms, we will demonstrate that $\mathcal{C}(X)$ separates points for an arbitrary compact Hausdorff topological space $(X, \mathcal{T})$ via specific properties of $(X, \mathcal{T})$. This will lead to a discussion of a more general class of topological spaces that are particularly nice in that they have well-behaved compactifications.

5.1 Some Separability Axioms

As mention in the introduction to this chapter, our first objective is to study various ways to separate out points in a topological space. In particular, our goal is to develop and study these property so we can use the correct separation axiom strength to obtain the most general results when possible. As such, we will often need to exhibit examples of topological spaces that satisfy one set of axioms and not the others in order to know we have the weakest assumptions as possible on our topological spaces in order for a result to hold.

In order for there to be some structure to the naming of our separations axioms, we will generally denote each separation axiom by 'T_n' for some natural number $n \in \mathbb{N}$. The 'T' comes from the German work "Trennungsaxiome" which translates to "separation axiom" and the $n \in \mathbb{N}$ is to denote the relative strength of the separation axiom. In particular, we endeavour to go from weakest separation axioms to the strongest while increasing n along the way. Of course, this clearly means that some topological spaces (e.g. metric spaces, compact Hausdorff topological spaces) will satisfy many of the first separation axioms although we will not mention them until we arrive at later separation axioms. In addition, we note that the study of separation axioms is notorious for conflicts with naming conventions used so the names used in these notes may or may not agree with the conventions other authors use in the literature (the axioms in this section are really standard; some in later sections not so much).

To being our study of separation axioms, we define the easiest way to separate two distinct points with an open set in a topological space.

Definition 5.1.1. A topological space $(X, \mathcal{T})$ is said to be a *T_0 space* if for all $x_1, x_2 \in X$ such that $x_1 \neq x_2$ there exists a $U \in \mathcal{T}$ such that either $x_1 \in U$ and $x_2 \notin U$ or $x_1 \notin U$ and $x_2 \in U$.

Remark 5.1.2. Clearly requiring a topological space $(X, \mathcal{T})$ to be T_0 is a very weak assumption thereby causing most spaces are T_0. The one main reason for requiring a topological space to be T_0 is that distinct points are topologically distinguishable in the sense that distinct points do not have the same neighbourhoods. This is a good property to have for otherwise there would exist two distinct points in X for which no aspect of the topology (e.g. continuous functions, cluster points, closed sets, etc.) distinguish.

Of course, not all topological spaces are T_0.

Example 5.1.3. Let X be a set with at least two points and let $\mathcal{T}$ denote the trivial topology on X. Then $(X, \mathcal{T})$ is not a T_0 space.

Instead of giving examples of a T_0 space (of which there are clearly numerous), we will instead strengthen the notion of separation used to define a T_0 space by making sure that we can separate both points using an open set and provide examples of this topological spaces with this stronger property.

Definition 5.1.4. A topological space $(X, \mathcal{T})$ is said to be a T_1 *space* if for all $x_1, x_2 \in X$ such that $x_1 \neq x_2$ there exists a $U \in \mathcal{T}$ such that $x_1 \in U$ and $x_2 \notin U$.

Example 5.1.5. It is trivial to see based on definitions that every T_1 space is automatically a T_0 space. However, the converse need not be true. Indeed consider the set $X = \{0, 1\}$ and the topology

$$\mathcal{T} = \{\emptyset, \{1\}, X\}$$

on X. Clearly $\mathcal{T}$ is indeed a topology on X. Furthermore, $(X, \mathcal{T})$ is a T_0 space since if we take two distinct points in X, the points must be 0 and 1 in which case the open set $U = \{1\} \in \mathcal{T}$ has the property that $1 \in U$ and $0 \notin U$. However, $(X, \mathcal{T})$ is not a T_1 space since there does not exists a $U \in \mathcal{T}$ such that $0 \in U$ and $1 \notin U$.

Of course, as we will later strengthen the notion of a T_1 space, we hold off giving examples of T_1 spaces until later. However, there is an alternate characterization of a T_1 space that gives particularly good insightful into why having a topological space be T_1 is desirable.

Lemma 5.1.6. *A topological space $(X, \mathcal{T})$ is a T_1 space if and only if every singleton in X is a closed set in $(X, \mathcal{T})$.*

Proof. To being, suppose $(X, \mathcal{T})$ is a T_1 space and let $x \in X$ be arbitrary. To see that $\{x\}$ is closed in $(X, \mathcal{T})$, note for every $y \in X \setminus \{x\}$ the assumption that $(X, \mathcal{T})$ is T_1 implies there exists a $U_y \in \mathcal{T}$ such that $y \in U_y$ but $x \notin U_y$. Hence

$$X \setminus \{x\} = \bigcup_{y \in X \setminus \{x\}} U_y \in \mathcal{T}$$

and thus $\{x\}$ is closed in $(X, \mathcal{T})$.

Conversely, suppose that every singleton in X is a closed set in $(X, \mathcal{T})$. To see that $(X, \mathcal{T})$ is a T_1 space, let $x, y \in X$ be arbitrary points such that $x \neq y$. Notice since $\{x\}$ is closed in $(X, \mathcal{T})$ that $U = X \setminus \{x\}$ is an open set in $(X, \mathcal{T})$ such that $y \in U$ and $x \notin U$. Therefore, since $x, y \in X$ were arbitrary, $(X, \mathcal{T})$ is a T_1 space. ∎

Of course, we have already seen that Hausdorff topological spaces are T_1 (so we have examples of T_1 spaces). In particular, if we re-examine the definition of a Hausdorff topological spaces, we can see the definition is very similar to the separation axioms we have seen above and seemingly one step above the notion of a T_1 space. Consequently, we define the following.

Definition 5.1.7. A topological space $(X, \mathcal{T})$ is said to be a T_2 *space* if $(X, \mathcal{T})$ is Hausdorff; that is, for all $x_1, x_2 \in X$ such that $x_1 \neq x_2$ there exist $U_1, U_2 \in \mathcal{T}$ such that $x_1 \in U_1$, $x_2 \in U_2$, and $U_1 \cap U_2 = \emptyset$.

Remark 5.1.8. Of course, we have already seen the importance and useful properties of Hausdorff topological spaces in this course. For example, we have already seen a topological space is Hausdorff if and only if every convergent net has a unique point of convergence (Theorem 1.5.40 and Theorem 1.5.42). Moreover, we have seen that subspaces of Hausdorff spaces are Hausdorff, products of Hausdorff spaces are Hausdorff with respect to both the box and product topologies, and metric spaces are Hausdorff.

Of course, we now need to examine how the T_1 and T_2 separation axioms compare to each other.

Example 5.1.9. Clearly every Hausdorff topological space is T_1 by definition (or Example 1.6.8 and Lemma 5.1.6). However, not every T_1 topological space is Hausdorff. Indeed recall if X is an infinite set and $\mathcal{T}$ is the cofinite topology on X, then $(X, \mathcal{T})$ is not Hausdorff as the intersection of any two non-empty subsets from $\mathcal{T}$ will be non-empty. However $(X, \mathcal{T})$ is clearly T_1 by Lemma 5.1.6 as finite subsets of $(X, \mathcal{T})$ are closed being the complements of cofinite sets.

Of course, we can ask for stronger forms of separation. We know that a Hausdorff space allows us to separate points using disjoint neighbourhood of the points. Since Hausdorff spaces are T_1 spaces so points are closed, this means we are separating points from specific closed sets (namely points). This is very similar to what we saw in Lemma 3.1.12 in that we could separate a point from a compact set inside a Hausdorff space. However, as compact sets need not be closed inside Hausdorff spaces, we define the following in hopes that it is a stronger form of separation.

Definition 5.1.10. A topological space $(X, \mathcal{T})$ is said to be *regular* if whenever F is a closed subset of $(X, \mathcal{T})$ and $x \in X \setminus F$ there exist $U, V \in \mathcal{T}$ such that $x \in U$, $F \subseteq V$, and $U \cap V = \emptyset$.

Example 5.1.11. Unfortunately, a regular topological space need not be Hausdorff. Indeed if X is any non-empty set with at least two points and $\mathcal{T}$ is the trivial topology on X, then $(X, \mathcal{T})$ is clearly not Hausdorff but is regular since, in Definition 5.1.10, either $F = X$ so there is no $x \in X \setminus F$, or $F = \emptyset$ so one can take $U = X$ and $V = \emptyset$ for any $x \in X$.

As the above example is undesirable and as we would like a nice linear ordering on our separation axioms, we define the following to ensure that we have a stronger separation axiom than T_2.

Definition 5.1.12. A topological space $(X, \mathcal{T})$ is said to be a T_3 *space* (also called a *regular Hausdorff space*) if $(X, \mathcal{T})$ is regular and Hausdorff.

It is clear by construction that every T_3 topological space is Hausdorff by construction. In fact, one can weaken the Hausdorff requirement in the definition of a T_3 space.

Proposition 5.1.13. *A topological space $(X, \mathcal{T})$ is a T_3 space if and only if $(X, \mathcal{T})$ is a regular T_0 space.*

Proof. Let $(X, \mathcal{T})$ be a T_3 space. Hence $(X, \mathcal{T})$ is regular and Hausdorff. Since every Hausdorff space is a T_1 space and every T_1 space is a T_0 space, clearly $(X, \mathcal{T})$ is a regular T_0 space as desired.

Conversely, suppose that $(X, \mathcal{T})$ is a regular T_0 space. To see that $(X, \mathcal{T})$ is a T_3 space, it suffices to prove that $(X, \mathcal{T})$ is Hausdorff. To see that $(X, \mathcal{T})$ is Hausdorff, let $x_1, x_2 \in X$ be arbitrary points such that $x_1 \neq x_2$. Since $(X, \mathcal{T})$ is a T_0 space, there exists a $U \in \mathcal{T}$ such that $x_1 \in U$ and $x_2 \notin U$, or $x_1 \notin U$ and $x_2 \in U$. By exchanging the labels if necessary, we may assume that $x_1 \in U$ and $x_2 \notin U$. Hence $F = X \setminus U$ is a closed set in $(X, \mathcal{T})$ such that $x_1 \notin F$ and $x_2 \in F$. Since $(X, \mathcal{T})$ is regular, there exists $V_1, V_2 \in \mathcal{T}$ such that $x_1 \in V_1$, $x_2 \in F \subseteq V_2$, and $V_1 \cap V_2 = \emptyset$. Hence, as $x_1, x_2 \in X$ were arbitrary, $(X, \mathcal{T})$ is Hausdorff as desired. ∎

To indeed demonstrate that T_3 is a stronger axiom than T_2 and that we cannot replace 'compact' with 'closed' in Lemma 3.1.12, we examine the following example.

Example 5.1.14. Clearly every T_3 topological space is T_2 (i.e. Hausdorff) by definition. However, not every Hausdorff topological space is T_3. To see this, let $\mathcal{T}_\mathbb{R}$ be the canonical topology on $\mathbb{R}$, let $K = \{\frac{1}{n}\}_{n \geq 1} \subseteq \mathbb{R}$, let

$$\mathcal{K} = \{U \setminus C \mid U \in \mathcal{T}_\mathbb{R}, C \subseteq K\}$$

and let $\mathcal{T} = \mathcal{T}_\mathbb{R} \cup \mathcal{K}$. We claim that $\mathcal{T}$ is a topology on $\mathbb{R}$. Indeed clearly $\emptyset, \mathbb{R} \in \mathcal{T}_\mathbb{R} \subseteq \mathbb{R}$. Since both $\mathcal{T}_\mathbb{R}$ and $\mathcal{K}$ are closed under arbitrary unions and finite intersections as $\mathcal{T}_\mathbb{R}$ is a topology, since the union of an element of $\mathcal{T}_\mathbb{R}$ and $\mathcal{K}$ is an element of $\mathcal{K}$ or an element of $\mathcal{T}_\mathbb{R}$, and since the intersection of an element of $\mathcal{T}_\mathbb{R}$ and $\mathcal{K}$ is an element of $\mathcal{K}$, we obtain that $\mathcal{T}$ is a topology on $\mathbb{R}$.

We claim that $(\mathbb{R}, \mathcal{T})$ is a Hausdorff space that is not a T_3 space. To see that $(\mathbb{R}, \mathcal{T})$ is Hausdorff, we note that $\mathcal{T}_\mathbb{R} \subseteq \mathcal{T}$ so as $\mathcal{T}_\mathbb{R}$ is a Hausdorff topology and as topologies finer than Hausdorff topologies are Hausdorff by the definition of a Hausdorff topology, we obtain that $(\mathbb{R}, \mathcal{T})$ is Hausdorff. To see that $(\mathbb{R}, \mathcal{T})$ is not a T_3 space, we will show that $(\mathbb{R}, \mathcal{T})$ is not regular. To see this, suppose to the contrary that $(\mathbb{R}, \mathcal{T})$ is regular and consider the point $x = 0 \in \mathbb{R}$ and the set $K \subseteq \mathbb{R}$. Clearly $\mathbb{R} \setminus K \in \mathcal{T}_\mathbb{R} \subseteq \mathcal{T}$ so K is closed in $(\mathbb{R}, \mathcal{T})$. Thus, as we are assuming $(\mathbb{R}, \mathcal{T})$ is regular, there must exist $U, V \in \mathcal{T}$ such that $0 \in U$, $K \subseteq V$, and $U \cap V = \emptyset$. As this implies $K \cap U = \emptyset$, we must have that $U = U' \setminus K$ for some $U' \in \mathcal{T}_\mathbb{R}$. As $0 \in U = U' \setminus K$, there exists an $\epsilon > 0$ such that $(-\epsilon, \epsilon) \subseteq U'$. However, there exists an $n \in \mathbb{N}$ such that $\frac{1}{n} \in (-\epsilon, \epsilon) \cap K$. Moreover, since $K \subseteq V$ and $V \in \mathcal{T}$, the definition of $\mathcal{T}$ implies there exists a $\delta > 0$ such that $\left(\frac{1}{n} - \delta, \frac{1}{n} + \delta\right) \subseteq V$. However, as $\frac{1}{n} \in (-\epsilon, \epsilon) \subseteq U'$, we see that $U \cap V \neq \emptyset$ thereby yielding a contradiction. Hence $(\mathbb{R}, \mathcal{T})$ is not regular and thus is a Hausdorff space that is not a T_3 space.

One incredible use of knowing a topological space is regular is the ability to find neighbourhoods with closures inside other neighbourhoods.

Lemma 5.1.15. *Let $(X, \mathcal{T})$ be a topological space. Then $(X, \mathcal{T})$ is regular if and only if for every $x \in X$ and neighbourhood U of x there exists a neighbourhood V of x such that $\overline{V} \subseteq U$.*

Proof. First, suppose $(X, \mathcal{T})$ is a regular topological space. To see the desired properties, let $x \in X$ and let U be an arbitrary neighbourhood of x. Since $F = X \setminus U$ is closed in $(X, \mathcal{T})$ and since $x \notin F$, the fact that $(X, \mathcal{T})$ is regular implies there exists $V_1, V_2 \in \mathcal{T}$ such that $x \in V_1$, $F \subseteq V_2$, and $V_1 \cap V_2 = \emptyset$. Clearly V_1 is a neighbourhood of x. Furthermore, we claim that $\overline{V_1} \subseteq U$. To see this, we note since $F \subseteq V_2$ and $V_1 \cap V_2 = \emptyset$ that $F \cap \overline{V_1} = \emptyset$ by Theorem 1.6.21. Hence $\overline{V_1} \subseteq X \setminus F = U$. Therefore, since x and U were arbitrary, the result follows as desired.

Conversely, suppose $(X, \mathcal{T})$ is a topological space such that for every $x \in X$ and neighbourhood U of x there exists a neighbourhood V of x such that $\overline{V} \subseteq U$. To see that $(X, \mathcal{T})$ is regular, let F be an arbitrary closed subset of $(X, \mathcal{T})$ and let $x \in X \setminus F$ be arbitrary. Since $U = X \setminus F$ is then a neighbourhood of x, the assumptions imply there exists a neighbourhood V of x such that $\overline{V} \subseteq U$. Let $U_0 = X \setminus \overline{V}$ which is open in $(X, \mathcal{T})$ as $\overline{V}$ is closed. Therefore $V, U_0 \in \mathcal{T}$ are such that $x \in V$, $F \subseteq U_0$ since $\overline{V} \subseteq U = X \setminus F$, $V \cap U_0 = \emptyset$ since $V \subseteq \overline{V}$. Therefore, as x and F were arbitrary, $(X, \mathcal{T})$ is regular as desired. ∎

Lemma 5.1.15 along with definitions immediately enables us to show that T_3 is a separation axiom that behaves well with respect to our usual operations on topological spaces.

Theorem 5.1.16. *Any subspace of a regular space is regular. Consequently, a subspace of a T_3 space is a T_3 space.*

Proof. Since subspaces of Hausdorff spaces are Hausdorff, the second claim will follow from the first. To see the first claim, let $(X, \mathcal{T})$ be a regular topological space and let $Y \subseteq X$ be an arbitrary subspace. To see that Y is regular, let $F \subseteq Y$ be an arbitrary closed set and let $y \in Y \setminus F$ be arbitrary. By Lemma 1.6.12 there exists a closed subset C of $(X, \mathcal{T})$ such that $F = C \cap Y$. Since $y \in Y$, we know that $y \in X \setminus C$ so, since $(X, \mathcal{T})$ is regular, there exists $U_0, V_0 \in \mathcal{T}$ such that $y \in U_0$, $C \subseteq V_0$, and $U_0 \cap V_0 = \emptyset$.

Hence if $U = U_0 \cap Y$ and $V = V_0 \cap Y$, then U and V are open subset of Y such that $y \in U$, $F \subseteq V$, and $U \cap V = \emptyset$. Therefore, since F and y were arbitrary, Y is regular by definition. ∎

Theorem 5.1.17. *Let $\{(X_\alpha, \mathcal{T}_\alpha)\}_{\alpha \in I}$ be a non-empty collection of non-empty regular spaces. Then $\prod_{\alpha \in I} X_\alpha$ is regular with respect to both the product and box topologies. Hence if $\{(X_\alpha, \mathcal{T}_\alpha)\}_{\alpha \in I}$ is a non-empty collection of non-empty T_3 spaces, then $\prod_{\alpha \in I} X_\alpha$ is T_3 with respect to both the product and box topologies.*

Proof. Since products of Hausdorff spaces are Hausdorff with respect to both the box and product topologies, the second claim will follow from the first. To see the first claim, we will use Lemma 5.1.15 proceeding with the box and product topologies simultaneously and making the necessary refinements to deal with the coarser product topology.

Let $f \in \prod_{\alpha \in I} X_\alpha$ and U a neighbourhood of f be arbitrary. Due to the definition of the box and product topologies on $\prod_{\alpha \in I} X_\alpha$, for each $\alpha \in I$ there exists a $U_\alpha \in \mathcal{T}_\alpha$ such that if

$$U' = \prod_{\alpha \in I} U_\alpha$$

then U' is open and $f \in U' \subseteq U$ (where, in the case of the product topology, only a finite number of $\alpha \in I$ have the property that $U_\alpha \neq X_\alpha$). Since $(X_\alpha, \mathcal{T}_\alpha)$ is regular for all $\alpha \in I$ and since U_α is a neighbourhood of $f(\alpha)$, Lemma 5.1.15 implies that there exists a $V_\alpha \in \mathcal{T}_\alpha$ such that

$$f(\alpha) \in V_\alpha \subseteq \overline{V_\alpha} \subseteq U_\alpha.$$

Furthermore, in the case that $U_\alpha = X_\alpha$, we can clearly take $V_\alpha = X_\alpha$ and we do so to handle the proof for the product topology. Therefore, if we define

$$V = \prod_{\alpha \in I} V_\alpha$$

then V is open in the topology under consideration (the product topology is why we take $V_\alpha = X_\alpha$ so that only a finite number of $\alpha \in I$ have the property that $V_\alpha \neq X_\alpha$ by construction) and thus V is a neighbourhood of f. Furthermore, by Proposition 1.6.22, we see that

$$\overline{V} = \prod_{\alpha \in I} \overline{V_\alpha} \subseteq U' \subseteq U.$$

Therefore, as f and U were arbitrary, Lemma 5.1.15 implies that $\prod_{\alpha \in I} X_\alpha$ is regular as desired. ∎

As we could see, many of the spaces we have studied in this course, such as metric spaces and compact Hausdorff topological spaces, are T_3 spaces. However, we can go one small step further thereby strengthening the notion of a T_3 space. Indeed why should we only ask for points to be separated from closed sets? Why don't we ask for pairs of closed sets to be separated?

Definition 5.1.18. A topological space $(X, \mathcal{T})$ is said to be *normal* if whenever F_1 and F_2 are non-empty closed subsets of $(X, \mathcal{T})$ such that $F_1 \cap F_2 = \emptyset$ there exists $U_1, U_2 \in \mathcal{T}$ such that $F_1 \subseteq U_1$, $F_2 \subseteq U_2$, and $U_1 \cap U_2 = \emptyset$.

Remark 5.1.19. To aid in keeping track of the notions of Hausdorff, regular, and normal and how these notions compare, we provide the following diagrams:

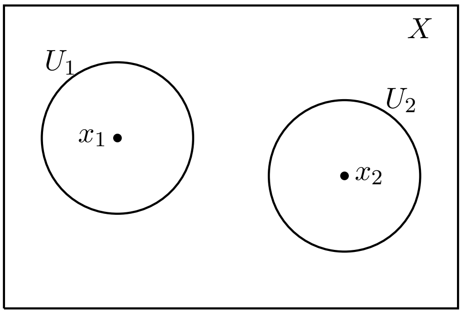

Hausdorff

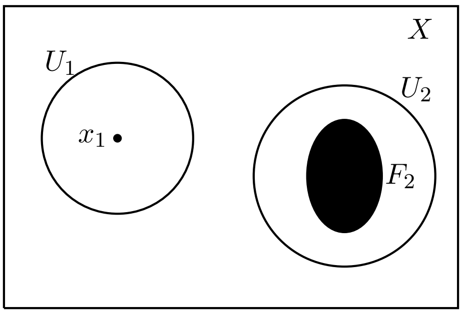

regular

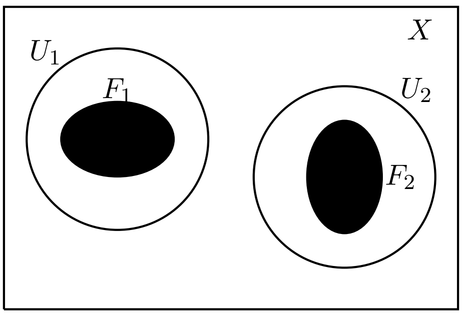

normal

Remark 5.1.20. Of course, it is easy to see by definitions that if points are closed, then every normal topological spaces is automatically regular and

Hausdorff. However, when points are not closed, a topological space can be normal but not regular nor Hausdorff. To see this, let $X = \{0, 1\}$ and let $\mathcal{T} = \{\emptyset, \{0, 1\}, \{0\}\}$ which is clearly a topology on X. To see that $(X, \mathcal{T})$ is normal, note the closed subsets of $(X, \mathcal{T})$ are $\emptyset$, X, and $\{1\}$. Hence as there does not exist non-empty closed sets F_1 and F_2 in $(X, \mathcal{T})$ such that $F_1 \cap F_2 = \emptyset$, $(X, \mathcal{T})$ is normal by definition. However $(X, \mathcal{T})$ is not Hausdorff as the point $\{0\}$ is not closed. Furthermore, $(X, \mathcal{T})$ is not regular since $F = \{1\}$ is a closed set, $0 \in X \setminus F$, and clearly there do not exists $U, V \in \mathcal{T}$ such that $0 \in U$ (so $U = \{0\}$ or $U = X$), $F \subseteq V$ (so $V = X$), and $U \cap V = \emptyset$ (as $0 \in U \cap V$ for all of the possible cases).

As we can see that a topological space being normal is a stronger property than a topological space being being regular provided points are closed, we can easily define a separation axiom that is stronger than being a T_3 space.

Definition 5.1.21. A topological space $(X, \mathcal{T})$ is said to be a *T_4 space* (or a *normal Hausdorff space*) if $(X, \mathcal{T})$ is normal and Hausdorff.

As discussed above, clearly every T_4 topological space is T_3. Furthermore, instead of requiring that a topological space is Hausdorff and normal in order to be T_4, we could have relaxed the conditions.

Proposition 5.1.22. *A topological space $(X, \mathcal{T})$ is T_4 if and only if $(X, \mathcal{T})$ is a normal T_1 space.*

Proof. First suppose $(X, \mathcal{T})$ is a T_4 topological space. Then clearly $(X, \mathcal{T})$ is normal and Hausdorff. Since every Hausdorff space is automatically T_1, $(X, \mathcal{T})$ is a normal T_1 space.

Conversely, suppose $(X, \mathcal{T})$ is a normal T_1 space. To show that $(X, \mathcal{T})$ is a T_4 space, it suffices to show that $(X, \mathcal{T})$ is Hausdorff. To see this, let $x_1, x_2 \in X$ such that $x_1 \neq x_2$. Since $(X, \mathcal{T})$ is T_1, Lemma 5.1.6 implies that $\{x_1\}$ and $\{x_2\}$ are closed in $(X, \mathcal{T})$. Therefore, since $(X, \mathcal{T})$ is normal and $\{x_1\}$ and $\{x_2\}$ are disjoint closed in $(X, \mathcal{T})$, there exists $U_1, U_2 \in \mathcal{T}$ such that $\{x_1\} \subseteq U_1$, $\{x_2\} \subseteq U_2$, and $U_1 \cap U_2 = \emptyset$. Thus, as $x_1, x_2 \in X$ were arbitrary, $(X, \mathcal{T})$ is Hausdorff and thus a T_4 space. ■

Of course, some of the nicest topological spaces we have studied in this course are T_4 spaces.

Theorem 5.1.23. *Every metric space is a T_4 space.*

Proof. First, as every metric spaces is Hausdorff by Example 1.5.36, it suffices to prove that every metric spaces is normal. To see this, let (X, d) be an arbitrary metric space and let A and B be arbitrary non-empty closed subsets of (X, d) such that $A \cap B = \emptyset$. We claim for all $a \in A$ that

$$\mathrm{dist}(a, B) = \inf(\{d(a, b) \mid b \in B\}) > 0.$$

To see this, suppose otherwise that $\text{dist}(a, B) = 0$. Then for all $n \in \mathbb{N}$ there exists a $b_n \in B$ such that $d(a, b_n) < \frac{1}{n}$. Therefore, $(b_n)_{n\geq 1}$ is a sequence in B that converges to a. Hence, as B is closed, $a \in B$ by Theorem 1.6.14 thereby implying $a \in A \cap B$ contradicting the fact that $A \cap B = \emptyset$. Hence $\text{dist}(a, B) > 0$.

For each $a \in A$, the above shows us that we may choose an $\epsilon_a \in \left(0, \frac{1}{2}\text{dist}(a, B)\right)$. Similarly, by reversing the roles of A and B, for each $b \in B$ we may choose an $\epsilon_b \in \left(0, \frac{1}{2}\text{dist}(b, A)\right)$. Let

$$U = \bigcup_{a\in A} B_d(a, \epsilon_a) \qquad \text{and} \qquad V = \bigcup_{b\in B} B_d(b, \epsilon_b).$$

Clearly U and V are open sets in (X, d) such that $A \subseteq U$ and $B \subseteq V$. Furthermore, we claim that $U \cap V = \emptyset$. To see this, suppose otherwise there exists a $x \in U \cap V$. By the definition of U and V, this implies there exists an $a \in A$ and $b \in B$ such that

$$d(a, x) < \epsilon_a \qquad \text{and} \qquad d(b, x) < \epsilon_b.$$

Hence

$$\begin{aligned} d(a, b) &\leq d(a, x) + d(x, b) \\ &< \frac{1}{2}\text{dist}(a, B) + \frac{1}{2}\text{dist}(b, A) \\ &\leq \frac{1}{2}d(a, b) + \frac{1}{2}d(b, a) = d(a, b) \end{aligned}$$

thereby yielding a contradiction. Hence $U \cap V = \emptyset$ as desired. Therefore, since A and B were arbitrary, (X, d) is normal as desired. ∎

Our next goal is to show that compact Hausdorff topological spaces are T_4 spaces. To show this, we need to improve upon Lemma 3.1.12.

Theorem 5.1.24. *Let $(X, \mathcal{T})$ be a Hausdorff topological space. If $K_1, K_2 \subseteq X$ are compact subspaces of $(X, \mathcal{T})$ such that $K_1 \cap K_2 = \emptyset$, then there exists $U, V \in \mathcal{T}$ such that $K_1 \subseteq U$, $K_2 \subseteq V$, and $U \cap V = \emptyset$.*

Proof. By Lemma 3.1.12 for all $y \in K_2$ there exists $U_y, V_y \in \mathcal{T}$ such that $K_1 \subseteq U_y$, $y \in V_y$, and $U_y \cap V_y = \emptyset$. Clearly

$$\{V_y \mid y \in K_2\}$$

is an open cover of K_2. Hence, as K_2 is a compact subspace of $(X, \mathcal{T})$ there exists an $m \in \mathbb{N}$ and $y_1, y_2, \ldots, y_m \in K_1$ such that $\{V_{y_j}\}_{j=1}^m$ is an open cover of K_2. Let

$$U = \bigcap_{j=1}^m U_{y_j} \qquad \text{and} \qquad V = \bigcup_{j=1}^m V_{y_j}.$$

Clearly $U, V \in \mathcal{T}$ as $\mathcal{T}$ is a topology, $K_1 \subseteq U$ by construction, and $K_2 \subseteq V$ as $\{V_{y_j}\}_{j=1}^m$ is an open cover of K_2. Furthermore, since $U_{y_j} \cap V_{y_j} = \emptyset$ for all $j \in \{1, \ldots, m\}$, we see that $U \cap V = \emptyset$. Hence, as K_1 and K_2 were arbitrary, the result holds. ∎

Corollary 5.1.25. *Every compact Hausdorff topological space is a T_4 space.*

Proof. Let $(X, \mathcal{T})$ be a compact Hausdorff topological space. As $(X, \mathcal{T})$ is Hausdorff, to show that $(X, \mathcal{T})$ is T_4 it suffices to prove that $(X, \mathcal{T})$ is normal. To see that $(X, \mathcal{T})$ is normal, let F_1 and F_2 be arbitrary non-empty closed subsets of $(X, \mathcal{T})$ such that $F_1 \cap F_2 = \emptyset$. Since $(X, \mathcal{T})$ is closed, Theorem 3.1.14 implies that F_1 and F_2 are compact subspaces of $(X, \mathcal{T})$. Hence since $F_1 \cap F_2 = \emptyset$ and $(X, \mathcal{T})$ is Hausdorff, Theorem 5.1.24 implies there exists $U, V \in \mathcal{T}$ such that $F_1 \subseteq U$, $F_2 \subseteq V$, and $U \cap V = \emptyset$. Therefore, since F_1 and F_2 were arbitrary, $(X, \mathcal{T})$ is normal as desired. ∎

Example 5.1.26. Let $\mathcal{T}_L$ be the lower limit topology on $\mathbb{R}$. Then $(\mathbb{R}, \mathcal{T}_L)$ is a T_4 space. As $(\mathbb{R}, \mathcal{T}_L)$ is Hausdorff by Example 1.5.37, it suffices to show that $(\mathbb{R}, \mathcal{T}_L)$ is normal.

To see that $(\mathbb{R}, \mathcal{T}_L)$ is normal, let A and B be arbitrary non-empty closed subsets of $(\mathbb{R}, \mathcal{T}_L)$ such that $A \cap B = \emptyset$. Hence $\mathbb{R} \setminus B$ and $\mathbb{R} \setminus A$ are open subsets of $(\mathbb{R}, \mathcal{T}_L)$ that contain A and B respectively. Since $A \subseteq \mathbb{R} \setminus B$, since $\mathbb{R} \setminus B$ is open in $(\mathbb{R}, \mathcal{T}_L)$, and since for all $x \in \mathbb{R}$,

$$\mathcal{B}_x = \left\{ \left[x, x + \frac{1}{n}\right) \,\middle|\, n \in \mathbb{N} \right\}$$

is a neighbourhood basis of x , for each $a \in A$ there exists an $\epsilon_a > 0$ such that if $U_a = [a, a + \epsilon_a)$, then $a \in U_a \subseteq \mathbb{R} \setminus B$. Similarly, for each $b \in B$ there exists an $\epsilon_b > 0$ such that if $V_b = [b, b + \epsilon_b)$, then $b \in V_b \subseteq \mathbb{R} \setminus A$

Let

$$U = \bigcup_{a \in A} U_a \qquad \text{and} \qquad V = \bigcup_{b \in B} V_b.$$

Clearly $U, V \in \mathcal{T}_L$, $A \subseteq U$, and $B \subseteq V$ by construction. We claim that $U \cap V = \emptyset$. To see this, suppose to the contrary that $U \cap V \neq \emptyset$. Therefore, by the definition of U and V there exist $a \in A$ and $b \in B$ such that $U_a \cap V_b \neq \emptyset$. Thus

$$[a, a + \epsilon_a) \cap [b, b + \epsilon_b) \neq \emptyset.$$

Thus it must be the case that $b < a + \epsilon_a$ and $a < b + \epsilon_b$. Since $A \cap B = \emptyset$, $a \neq b$. If $a < b$, then as $b < a + \epsilon_a$ we have that $b \in U_a \subseteq \mathbb{R} \setminus B$ contradicting the fact that $b \in B$. However, if $b < a$, then as $a < b + \epsilon_b$ we have that $a \in V_b \subseteq \mathbb{R} \setminus A$ contradicting the fact that $a \in A$. As we have a contradiction, it must be the case that $U \cap V = \emptyset$. Therefore, as A and B were arbitrary, $(\mathbb{R}, \mathcal{T}_L)$ is normal as desired.

Of course, unlike when discussing the other separation axioms, we have yet to given an example of a T_3 space that is not T_4. In particular, the following shows that a simple example on a finite set does not exist.

Proposition 5.1.27. *Let $(X, \mathcal{T})$ be a regular space with a countable basis. Then $(X, \mathcal{T})$ is normal.*

Proof. To see that $(X, \mathcal{T})$ is normal, let $\mathcal{B}$ be a countable basis for $(X, \mathcal{T})$ and let A and C be non-empty closed subsets of $(X, \mathcal{T})$ such that $A \cap C = \emptyset$. To construct our disjoint open sets containing A and C, first note since $(X, \mathcal{T})$ is regular and C is closed in $(X, \mathcal{T})$ that for all $a \in A$ there exists a neighbourhood U_a of a such that $U_a \cap C = \emptyset$. Furthermore, since $(X, \mathcal{T})$ is regular, Lemma 5.1.15 implies there exists $V_a \in \mathcal{T}$ such that

$$a \in V_a \subseteq \overline{V_a} \subseteq U_a.$$

Since $\mathcal{B}$ is a basis of $(X, \mathcal{T})$, for each $a \in A$ we can choose a $B_a \in \mathcal{B}$ such that $a \in B_a \subseteq V_a$. Hence, by Theorem 1.6.21 we see that

$$\overline{B_a} \subseteq \overline{V_a} \subseteq U_a$$

so that $\overline{B_a} \cap C = \emptyset$ as $U_a \cap C = \emptyset$. Hence, as $\mathcal{B}$ is countable, we can find a countable set $\{W_n\}_{n\geq 1}$ of elements of $\mathcal{B}$ such that

$$A \subseteq \bigcup_{n=1}^{\infty} W_n \qquad \text{and} \qquad \overline{W_n} \cap C = \emptyset \text{ for all } n \in \mathbb{N}.$$

By using identical arguments and by reversing the roles of of A and C, there exists a countable set $\{Z_n\}_{n\geq 1}$ of elements of $\mathcal{B}$ such that

$$C \subseteq \bigcup_{n=1}^{\infty} Z_n \qquad \text{and} \qquad \overline{Z_n} \cap A = \emptyset \text{ for all } n \in \mathbb{N}.$$

Of course, we would like to take the union of $\{W_n\}_{n\geq 1}$ and the union of $\{Z_n\}_{n\geq 1}$ in order to obtain open subsets of $(X, \mathcal{T})$ that contain A and C. However, this need not due the trick since we do not know that the unions will be disjoint. To solve this problem, we need to correct these sets.

For every $n \in \mathbb{N}$, let

$$W_n' = W_n \setminus \left(\bigcup_{k=1}^{n} \overline{Z}_k \right) \qquad \text{and} \qquad Z_n' = Z_n \setminus \left(\bigcup_{k=1}^{n} \overline{W}_k \right).$$

Clearly $\{\overline{W_n}\}_{n\geq 1}$ and $\{\overline{Z_n}\}_{n\geq 1}$ are collections of closed subsets of $(X, \mathcal{T})$ so $\{\bigcup_{k=1}^{n} \overline{W}_k\}_{n\geq 1}$ and $\{\bigcup_{k=1}^{n} \overline{Z}_k\}_{n\geq 1}$ are closed subsets of $(X, \mathcal{T})$. Therefore, since $\{W_n\}_{n\geq 1}$ and $\{Z_n\}_{n\geq 1}$ are collections of open subsets of $(X, \mathcal{T})$ and since $D \setminus E = D \cap (X \setminus E)$ for all $D, E \subseteq X$, we see that $\{W_n'\}_{n\geq 1}$ and $\{Z_n'\}_{n\geq 1}$ are collections of open subsets of $(X, \mathcal{T})$.

Let

$$W = \bigcup_{n=1}^{\infty} W'_n \qquad \text{and} \qquad Z = \bigcup_{n=1}^{\infty} Z'_n,$$

which are open subsets of $(X, \mathcal{T})$ by the above discussions. We claim that $A \subseteq W$, $C \subseteq Z$, and $W \cap Z = \emptyset$. To see this, first notice since $\overline{Z_n} \cap A = \emptyset$ for all $n \in \mathbb{N}$ that $W'_n \cap A = W_n \cap A$ for all $n \in \mathbb{N}$. Hence, since $A \subseteq \bigcup_{n=1}^{\infty} W_n$ we obtain that $A \subseteq W$. Furthermore, similar arguments show that $C \subseteq Z$. Finally, suppose to the contrary that $W \cap Z \neq \emptyset$ so that there exists an $x \in W \cap Z$. By the definition of W and Z, there must exists $n, m \in \mathbb{N}$ so that $x \in W'_n$ and $x \in Z'_m$. If $n \geq m$, then $x \in Z'_m$ implies that $x \in Z_m$ and $x \in W'_n$ implies that

$$x \in W_n \setminus \left(\bigcup_{k=1}^{n} \overline{Z}_k \right) \subseteq W_n \setminus Z_m,$$

which is an obvious contradiction. Similarly, if $m \geq n$ then $x \in W_n$ and $x \in Z_m \setminus W_n$ which is also a contradiction. Hence is must be the case that $W \cap Z = \emptyset$.

Therefore, since A and C were arbitrary, $(X, \mathcal{T})$ is normal as desired. ∎

So, if there is an example of a T_3 space that is not a T_4 space, it is not a straightforward example. Of course, some subspaces of T_4 spaces are clearly automatically T_4 and thus can aid us in our search.

Proposition 5.1.28. *Every closed subspace of a T_4 space is a T_4 space.*

Proof. Let Y be a closed subspace of a T_4 topological space $(X, \mathcal{T})$. To see that Y is T_4, we note that Y is Hausdorff being a subspace of a Hausdorff space. To see that Y is normal, let F_1 and F_2 be non-empty closed subsets of Y such that $F_1 \cap F_2 = \emptyset$. By the definition of the subspace topology, there exist closed subsets C_1 and C_2 in $(X, \mathcal{T})$ such that

$$F_1 = Y \cap C_1 \qquad \text{and} \qquad F_2 = Y \cap C_2.$$

Therefore, since Y is closed in $(X, \mathcal{T})$, F_1 and F_2 are non-empty closed subsets of $(X, \mathcal{T})$ such that $F_1 \cap F_2 = \emptyset$. Therefore, since $(X, \mathcal{T})$ is T_4, there exist open sets $U', V' \in \mathcal{T}$ such that $F_1 \subseteq U'$, $F_2 \subseteq V'$, and $U' \cap V' = \emptyset$. Hence if

$$U = Y \cap U' \qquad \text{and} \qquad V = Y \cap V',$$

then U and V are open subsets of Y such that $F_1 \subseteq U$, $F_2 \subseteq V$, and $U \cap V = \emptyset$. Therefore, as F_1 and F_2 were arbitrary, Y is a T_4 space as desired. ∎

We are now situated to show an example of a T_3 topological space that is not T_4.

Example 5.1.29. Let $\mathbb{R}$ be equipped with its canonical topology. Consider $\mathcal{F}(\mathbb{R}, \mathbb{R}) = \prod_{x\in\mathbb{R}} \mathbb{R}$ equipped with the product topology. Since $\mathbb{R}$ is a metric space, $\mathbb{R}$ is a T_4 space and thus a T_3 space. Hence, as the product topology on a product of T_3 spaces is a T_3 space by Theorem 5.1.17, $\mathcal{F}(\mathbb{R}, \mathbb{R})$ is a T_3 space.

However, $\mathcal{F}(\mathbb{R}, \mathbb{R})$ is not a T_4 space.

$$X = \mathcal{F}(\mathbb{R}, \mathbb{N}) = \prod_{\alpha\in\mathbb{R}} \mathbb{N}$$

equipped with the product topology. Clearly X is a closed subset of $\mathcal{F}(\mathbb{R}, \mathbb{R})$ being a product of closed sets. Therefore, if $\mathcal{F}(\mathbb{R}, \mathbb{R})$ is normal, then X would be normal by Proposition 5.1.28. Hence it suffices to show that X is not normal.

To see that X is not normal, we will require some structures. Namely we will require neighbourhood bases and some specific sets. Given $f \in X = \mathcal{F}(\mathbb{R}, \mathbb{N})$ and $I \subseteq \mathbb{R}$ finite, let

$$U(f, I) = \{g \in \mathcal{F}(\mathbb{R}, \mathbb{N}) \mid g(\alpha) = f(\alpha) \text{ for all } \alpha \in I\}.$$

Since $\{\{n\} \mid n \in \mathbb{N}\}$ is a basis for the topology on $\mathbb{N}$, we know that the collection of sets of the form

$$\prod_{\alpha\in\mathbb{R}} U_\alpha$$

where $U_\alpha = \mathbb{N}$ for all α except in a finite set $I \subseteq \mathbb{R}$ and $U_\alpha \in \{\{n\} \mid n \in \mathbb{N}\}$ for all $\alpha \in I$ is a basis for X. Therefore, as

$$\mathcal{B} = \{U(f, I) \mid f \in X, I \subseteq \mathbb{R} \text{ finite}\}$$

is precisely the collection of all of the above sets, $\mathcal{B}$ is a basis for X.

For each $n \in \mathbb{N}$, let

$$P_n = \{f \in X \mid f \text{ is injective on } \mathbb{R} \setminus f^{-1}(\{n\})\}.$$

The sets P_n will be instrumental in showing that X is not normal. To do so, we first need some properties of P_n.

First, we claim that P_n is closed for all $n \in \mathbb{N}$. To see this, let $(f_\lambda)_{\lambda\in\Lambda}$ be an arbitrary net in P_n that converges to some element $f \in X$. To complete the proof that P_n is closed, it suffices to show that $f \in P_n$ by Theorem 1.6.14. Therefore, by the definition of P_n, it remains only to show that f is injective on $\mathbb{R} \setminus f^{-1}(\{n\})$. To see this, let $x_1, x_2 \in \mathbb{R} \setminus f^{-1}(\{n\})$ such that $x_1 \neq x_2$ be arbitrary. Hence $f(x_1) \neq n$ and $f(x_2) \neq n$. Since $(f_\lambda)_{\lambda\in\Lambda}$ converges to f in X and the product topology has been placed on X, it must be the case that $(f_\lambda(x_1))_{\lambda\in\Lambda}$ converges to $f(x_1)$ in $\mathbb{N}$ and $(f_\lambda(x_2))_{\lambda\in\Lambda}$ converges to $f(x_2)$ in $\mathbb{N}$. Recall that since $\mathbb{N}$ is equipped with the discrete topology, a net converges in $\mathbb{N}$ if and only if it is eventually constant. Hence there exists $\lambda_1, \lambda_2 \in \mathbb{N}$ such that $f_\lambda(x_1) = f(x_1)$ for all $\lambda \geq \lambda_1$ and $f_\lambda(x_2) = f(x_2)$ for all $\lambda \geq \lambda_2$. By the properties of nets, there exists a $\lambda_0 \in \Lambda$ such that $\lambda_0 \geq \lambda_1$ and $\lambda_0 \geq \lambda_2$. Hence $f_{\lambda_0}(x_1) = f(x_1)$ and $f_{\lambda_0}(x_2) = f(x_2)$. Thus, as $f(x_1) \neq n$ and $f(x_2) \neq n$, $x_1, x_2 \in \mathbb{R} \setminus f_{\lambda_0}^{-1}(\{n\})$. Therefore, since $f_{\lambda_0} \in P_n$ so that f_{λ_0} is injective on $\mathbb{R} \setminus f^{-1}(\{n\})$, we obtain that

$$f(x_1) = f_{\lambda_0}(x_1) \neq f_{\lambda_0}(x_2) = f(x_2)$$

as desired. Hence P_n is closed.

Next we claim that $P_1 \cap P_2 = \emptyset$. To see this, let $f \in P_1$ be arbitrary. Hence f is injective on $\mathbb{R} \setminus f^{-1}(\{1\})$. However, as $f : \mathbb{R} \to \mathbb{N}$ and $\mathbb{N}$ is countable, f being injective on $\mathbb{R} \setminus f^{-1}(\{1\})$ implies that $\mathbb{R} \setminus f^{-1}(\{1\})$ is countable. Therefore, since $\mathbb{R}$ is not countable, $f^{-1}(\{1\})$ must be uncountable. Hence, as $f^{-1}(\{1\})$ is an uncountable subset of $\mathbb{R} \setminus f^{-1}(\{2\})$, we see that f is not injective on $\mathbb{R} \setminus f^{-1}(\{2\})$ as it maps an uncountable subset of $\mathbb{R} \setminus f^{-1}(\{2\})$ to 1. Hence $f \notin P_2$. Therefore, since f was arbitrary, $P_1 \cap P_2 = \emptyset$.

To proceed with showing that X is not normal, let U be an arbitrary open subset of X containing P_1. We claim that there exists distinct elements $(\alpha_n)_{n\geq 1}$ of $\mathbb{R}$ and an increasing sequence of natural numbers $(k_n)_{n\geq 1}$ such that if $k_0 = 0$, if for each $q \in \mathbb{N}$ we define $I_q = \{\alpha_n\}_{n=1}^{k_q}$, and if for each $q \in \mathbb{N}$ we define $f_q \in \mathcal{F}(\mathbb{R}, \mathbb{N})$ by

$$f_q(\alpha) = \begin{cases} j & \text{if } \alpha = \alpha_j \text{ for some } j \in \{1, \ldots, k_{q-1}\} \\ 1 & \text{otherwise} \end{cases},$$

then $U(f_q, I_q) \subseteq U$. To see this, let $f_1 \in \mathcal{F}(\mathbb{R}, \mathbb{N})$ be the constant function 1. Hence $f_1 \in P_1 \subseteq U$. Therefore, as $\{U(f, I) \mid f \in X, I \subseteq \mathbb{R} \text{ finite}\}$ is a basis for the topology on X, there exists a finite subset $I_1 \subseteq \mathbb{R}$ such that $U(f_1, I_1) \subseteq U$. Since I_1 is finite, we can write

$$I_1 = \{\alpha_1, \alpha_2, \ldots, \alpha_{k_1}\}$$

for some distinct $\alpha_1, \alpha_2, \ldots, \alpha_{k_1} \in \mathbb{R}$ and some $k_1 \in \mathbb{N}$.

Let $f_2 \in \mathcal{F}(\mathbb{R}, \mathbb{N})$ be defined by

$$f_2(\alpha) = \begin{cases} j & \text{if } \alpha = \alpha_j \text{ for some } j \in \{1, \ldots, k_1\} \\ 1 & \text{otherwise} \end{cases}.$$

Clearly f_2 is injective on $\mathbb{R} \setminus f_2^{-1}(\{1\})$ and hence $f_2 \in P_1 \subseteq U$. Therefore, as $\{U(f, I) \mid f \in X, I \subseteq \mathbb{R} \text{ finite}\}$ is a basis for the topology on X, there exists a finite subset $J_2 \subseteq \mathbb{R}$ such that $U(f_2, J_2) \subseteq U$. Choose an element $x_2 \in \mathbb{R}$ such that $x_2 \notin I_1$ and let $I_2 = I_1 \cup J_2 \cup \{x_2\}$ which is a finite subset of $\mathbb{R}$ such that $I_1 \subsetneq I_2$. Since $J_2 \subseteq I_2$, we see that $U(f_2, I_2) \subseteq U(f_2, J_2) \subseteq U$. Since I_2 is finite and since $I_1 \subsetneq I_2$, we can find distinct $\alpha_{k_1+1}, \alpha_{k_1+2}, \ldots, \alpha_{k_2} \in \mathbb{R}$ for some $k_2 \in \mathbb{N}$ with $k_2 > k_1$ such that

$$I_2 = \{\alpha_1, \alpha_2, \ldots, \alpha_{k_2}\}$$

and $\alpha_1, \alpha_2, \ldots, \alpha_{k_2}$ are distinct.

To proceed by recursion, suppose for some $q \in \mathbb{N}$ we have found natural numbers $\{k_n\}_{n=1}^{q}$ such that $k_n < k_{n+1}$ for all $n \in \{1, \ldots, q-1\}$ and distinct elements $I_q = \{\alpha_n\}_{n=1}^{k_q}$ in $\mathbb{R}$. Let $f_{q+1} \in \mathcal{F}(\mathbb{R}, \mathbb{N})$ be defined by

$$f_{q+1}(\alpha) = \begin{cases} j & \text{if } \alpha = \alpha_j \text{ for some } j \in \{1, \ldots, k_q\} \\ 1 & \text{otherwise} \end{cases}.$$

Clearly f_{q+1} is injective on $\mathbb{R} \setminus f_{q+1}^{-1}(\{1\})$ and hence $f_{q+1} \in P_1 \subseteq U$. Therefore, as $\{U(f, I) \mid f \in X, I \subseteq \mathbb{R} \text{ finite}\}$ is a basis for the topology on X, there exists a finite subset $J_{q+1} \subseteq \mathbb{R}$ such that $U(f_{q+1}, J_{q+1}) \subseteq U$. Choose an element $x_{q+1} \in \mathbb{R}$ such that $x_{q+1} \notin I_q$ and let $I_{q+1} = I_q \cup J_{q+1} \cup \{x_{q+1}\}$ which is a finite subset of $\mathbb{R}$ such that $I_q \subsetneq I_{q+1}$. Since $J_{q+1} \subseteq I_{q+1}$, we see that $U(f_{q+1}, I_{q+1}) \subseteq U(f_{q+1}, J_{q+1}) \subseteq U$. Since I_{q+1} is finite and since $I_q \subsetneq I_{q+1}$, we can find distinct $\alpha_{k_q+1}, \alpha_{k_q+2}, \ldots, \alpha_{k_{q+1}} \in \mathbb{R}$ for some $k_{q+1} \in \mathbb{N}$ with $k_{q+1} > k_q$ such that

$$I_{q+1} = \{\alpha_1, \alpha_2, \ldots, \alpha_{k_{q+1}}\}$$

and $\alpha_1, \alpha_2, \ldots, \alpha_{k_{q+1}}$ are distinct.

Hence, by the above recursive process, the desired objects have been constructed. Note that we indeed get an infinite collection $\{\alpha_n\}_{n=1}^{\infty}$ of distinct elements of $\mathbb{R}$ since $k_{q+1} > k_q$ for all $q \in \mathbb{N}$ so we are always adding at least on α_n at each step.

Finally, to see that X is not normal, let V be an arbitrary open subset of X containing P_2. Define $g \in \mathcal{F}(\mathbb{R}, \mathbb{N})$ by

$$g(\alpha) = \begin{cases} n & \text{if } \alpha = \alpha_n \text{ for some } n \in \mathbb{N} \\ 2 & \text{otherwise} \end{cases}$$

for all $\alpha \in \mathbb{R}$. Clearly $g \in P_2 \subseteq V$ so there exists a finite subset $J \subseteq \mathbb{R}$ such that $U(g, J) \subseteq V$. As J is finite, there exists a $q \in \mathbb{N}$ such that $J \cap \{\alpha_n\}_{n=1}^{\infty} \subseteq I_q$.

Define $h \in \mathcal{F}(\mathbb{R}, \mathbb{N})$ by

$$h(\alpha) = \begin{cases} n & \text{if } \alpha = \alpha_n \text{ for some } \alpha_n \in I_q \\ g(\alpha) & \text{if } \alpha \in J \setminus \{\alpha_n\}_{n=1}^{\infty} \\ f_{q+1}(\alpha) & \text{if } \alpha \in I_{q+1} \setminus I_q \\ 3 & \text{otherwise} \end{cases}$$

for all $\alpha \in \mathbb{R}$. Since $J \cap \{\alpha_n\}_{n=1}^{\infty} \subseteq I_q$ so

$$(J \setminus \{\alpha_n\}_{n=1}^{\infty}) \cap (I_{q+1} \setminus I_q) = \emptyset,$$

h is well-defined. We claim that $h \in U(f_{q+1}, I_{q+1}) \cap U(g, J)$. To see this, first we demonstrate that $h \in U(f_{q+1}, I_{q+1})$. To see this, let $\alpha \in I_{q+1}$ be arbitrary. Hence $\alpha = \alpha_n$ for some $n \in \{1, \ldots, k_q\}$. If $\alpha_n \in I_q \subseteq I_{q+1}$, then $n \in \{1, \ldots, k_{q-1}\}$ so $f_{q+1}(\alpha_n) = n = h(\alpha_n)$ as desired. Otherwise, if $\alpha_n \in I_{q+1} \setminus I_q$, then $h(\alpha_n) = f_{q+1}(\alpha_n)$ by construction. Hence, as $\alpha \in I_{q+1}$ was arbitrary, $h \in U(f_{q+1}, I_{q+1})$. Next, we demonstrate that $h \in U(g, J)$. To see this, let $\alpha \in J$ be arbitrary. If $\alpha \in J \cap \{\alpha_n\}_{n=1}^\infty \subseteq I_q$, then $\alpha = \alpha_n$ for some $\alpha_n \in I_q$ and thus $h(\alpha) = n = g(\alpha)$ by construction. Otherwise $\alpha \in J \setminus \{\alpha_n\}_{n=1}^\infty$ so $h(\alpha) = g(\alpha)$ by construction. Hence, as $\alpha \in J$ was arbitrary, $h \in U(g, J)$. Hence $h \in U(f_{q+1}, I_{q+1}) \cap U(g, J)$ as desired.

Thus $U(f_{q+1}, I_{q+1}) \cap U(g, J) \neq \emptyset$. Combining the above, we have obtained that it must be the case that $U \cap V \neq \emptyset$ for any open sets U and V of X containing P_1 and P_2 respectively. Hence $(X, \mathcal{T})$ is not normal as desired.

Remark 5.1.30. Of course, as $\mathbb{R}$ equipped with its canonical topology is a metric space and thus a T_4 space, Example 5.1.29 implies that the product topology on a product of T_4 spaces need not be T_4. Of course Example 5.1.29 requires an infinite product. It turns out that there does exist a pair of T_4 spaces with product being not T_4, but as describing all possible flaws for T_4 spaces is not our focus, and as a simple example exists with the material from Chapter 6, we postpone the example (see Example 6.2.24).

Example 5.1.29 also yield another fact related to T_4 spaces. In particular, since $\mathbb{R}$ is homeomorphic to $(0, 1)$, we know that $\prod_{x\in\mathbb{R}} \mathbb{R}$ is homeomorphic to $\prod_{x\in\mathbb{R}}(0, 1)$. Hence $\prod_{x\in\mathbb{R}} \mathbb{R}$ can be viewed as a subspace of $\prod_{x\in\mathbb{R}}[0, 1]$. Since $\prod_{x\in\mathbb{R}}[0, 1]$ is a product of compact Hausdorff topological spaces, $\prod_{x\in\mathbb{R}}[0, 1]$ is a compact Hausdorff topological space by Tychonoff's Theorem (Theorem 3.3.4). Hence, as every compact Hausdorff topological space is a T_4 space by Corollary 5.1.25 and since the notion of a T_4 space is invariant under homeomorphism (as homeomorphisms preserve closed and open sets), we have an example of a subspace of a T_4 space that is not T_4.

As metric spaces and compact Hausdorff topological spaces are T_4 spaces, we know we have not gone too far in creating separation axioms as these major spaces still satisfy these axioms. Of course, we do not want to go too much farther as we have already seen that T_4 spaces do not behave particularly well with respect to taking subspaces and products. Thus we complete this section by exhibiting an alternative characterization of a normal space in the same flavour as Lemma 5.1.15. In fact, the following result is of surprising vital importance for the subsequent section.

Lemma 5.1.31. *Let $(X, \mathcal{T})$ be a topological space. Then $(X, \mathcal{T})$ is normal if and only if for every closed subset F of (X, T) and every $U \in \mathcal{T}$ such that $F \subseteq U$ there exists a $V \in \mathcal{T}$ such that $F \subseteq V$ and $\overline{V} \subseteq U$.*

Proof. First, suppose $(X, \mathcal{T})$ is a normal topological space. To see the desired properties, let F be an arbitrary closed subset of $(X, \mathcal{T})$ and let $U \in \mathcal{T}$ be an arbitrary open set such that $F \subseteq U$. Since $C = X \setminus U$ is closed in $(X, \mathcal{T})$

and since $F \cap C = \emptyset$, the fact that $(X, \mathcal{T})$ is normal implies there exists $V_1, V_2 \in \mathcal{T}$ such that $F \subseteq V_1$, $C \subseteq V_2$, and $V_1 \cap V_2 = \emptyset$. We claim that $\overline{V_1} \subseteq U$. To see this, we note since $C \subseteq V_2$ and $V_1 \cap V_2 = \emptyset$ that $C \cap \overline{V_1} = \emptyset$ by Theorem 1.6.21. Hence $\overline{V_1} \subseteq X \setminus C = U$. Therefore, since x and U were arbitrary, the result follows as desired.

Conversely, suppose $(X, \mathcal{T})$ is a topological space such that for every closed subset F of (X, T) and every $U \in \mathcal{T}$ such that $F \subseteq U$ there exists a $V \in \mathcal{T}$ such that $F \subseteq V$ and $\overline{V} \subseteq U$. To see that $(X, \mathcal{T})$ is normal, let F_1 and F_2 be arbitrary closed subset of $(X, \mathcal{T})$. Since $U = X \setminus F_2$ is an open set in $(X, \mathcal{T})$ such that $F_1 \subseteq U$, the assumptions imply there exists a $V \in \mathcal{T}$ such that $F_1 \subseteq V$ and $\overline{V} \subseteq U$. Let $U_0 = X \setminus \overline{V}$ which is open in $(X, \mathcal{T})$ as $\overline{V}$ is closed. Therefore $V, U_0 \in \mathcal{T}$ are such that $F_1 \subseteq V$, $F_2 \subseteq U_0$ since $\overline{V} \subseteq U = X \setminus F_2$, $V \cap U_0 = \emptyset$ since $V \subseteq \overline{V}$. Therefore, as F_1 and F_2 were arbitrary, $(X, \mathcal{T})$ is normal as desired. ∎

5.2 Urysohn's Lemma

With our knowledge of the above separation axioms, we turn our attention back to showing that if $(X, \mathcal{T})$ is a compact Hausdorff topological space, then $\mathcal{C}(X)$ separates points. In particular, as we know points are closed subsets of Hausdorff spaces, the following theorem immediately implies we can separate points inside a compact Hausdorff topological space. In fact, the reason we had to wait until how to demonstrate this result in its full generality is that it is the property that compact Hausdorff topological spaces are normal along with Lemma 5.1.31 that make this result possible.

Theorem 5.2.1 (Urysohn's Lemma). *Let $(X, \mathcal{T})$ be a normal topological space (e.g. a T_4 space) and let A and B be disjoint closed subsets of $(X, \mathcal{T})$. Then for every finite closed interval $[a, b]$ in $\mathbb{R}$ there exists a continuous function $f : X \to [a, b]$ such that $f(x) = a$ for all $x \in A$ and $f(x) = b$ for all $x \in B$.*

Proof. Clearly as every finite closed interval $[a, b]$ in $\mathbb{R}$ is homeomorphic to $[0, 1]$, it suffices to prove the result for $[a, b] = [0, 1]$. As motivation for the proof, notice that for a function $f : X \to [0, 1]$ such that $f(x) = 0$ for all $x \in A$ and $f(x) = 1$ for all $x \in B$ to exist, we must have for all $t \in (0, 1)$ that $f^{-1}([0, t))$ is an open set in $(X, \mathcal{T})$ that contains A and is disjoint from B. Consequently, our goal is to build up a continuum of open subsets of $(X, \mathcal{T})$ containing A and culminating in $X \setminus B$.

To begin this construction, let $V = X \setminus B$. Since B is a closed subset of $(X, \mathcal{T})$ that is disjoint from A, V is an open subset of $(X, \mathcal{T})$ such that $A \subseteq V$. Consequently, as $(X, \mathcal{T})$ is normal, Lemma 5.1.31 implies there exists a $U_{\frac{1}{2}} \in \mathcal{T}$ such that

$$A \subseteq U_{\frac{1}{2}} \subseteq \overline{U_{\frac{1}{2}}} \subseteq V.$$

Moreover, since A is a closed subset of $(X, \mathcal{T})$ and $U_{\frac{1}{2}}$ is an open subset of $(X, \mathcal{T})$ such that $A \subseteq U_{\frac{1}{2}}$, and since $\overline{U_{\frac{1}{2}}}$ is a closed subset of $(X, \mathcal{T})$ and V is an open subset of $(X, \mathcal{T})$ such that $\overline{U_{\frac{1}{2}}} \subseteq \overline{U_{\frac{1}{2}}}$, Lemma 5.1.31 implies there exists $U_{\frac{1}{4}}, U_{\frac{3}{4}} \in \mathcal{T}$ such that

$$A \subseteq U_{\frac{1}{4}} \subseteq \overline{U_{\frac{1}{4}}} \subseteq U_{\frac{1}{2}} \subseteq \overline{U_{\frac{1}{2}}} \subseteq U_{\frac{3}{4}} \subseteq \overline{U_{\frac{3}{4}}} \subseteq V.$$

By repeating this process ad infinitum, if

$$I = \left\{ \frac{m}{2^n} \,\middle|\, n \in \mathbb{N}, m \in \{1, \ldots, 2^n\} \right\}$$

then there exists $\{U_q\}_{q \in I} \subseteq \mathcal{T}$ with $U_1 = X$ such that

$$A \subseteq U_{q_1} \subseteq \overline{U_{q_1}} \subseteq U_{q_2} \subseteq V$$

for all $q_1, q_2 \in I$ with $q_1 < q_2 < 1$.

Define $f : X \to [0, 1]$ by

$$f(x) = \inf(\{q \in I \mid x \in U_q\})$$

for all $x \in X$. Clearly f is a well-defined function into $[0, 1]$ by construction as $U_1 = X$ and $I \subseteq [0, 1]$. We claim that f is the function we are looking for. To see this, first we note since $A \subseteq U_q$ for all $q \in I$ that $f(x) < q$ for all $x \in A$ and $q \in I$ and thus $f(x) = 0$ for all $x \in A$. Furthermore, if $b \in B$ then $b \notin V$ so $b \notin U_q$ for all $q \in I \setminus \{1\}$ and thus $f(b) = 1$ by definition. Thus it remains only to show that f is continuous.

To see that $f : X \to [0, 1]$ is continuous, we must show that if U is an open subset of $[0, 1]$ then $f^{-1}(U)$ is open in $(X, \mathcal{T})$. To begin to show this, we consider a few subcases.

First, we claim that if $t \in (0, 1)$ then

$$f^{-1}([0, t)) = \bigcup_{q \in I, q < t} U_q$$

thereby showing that $f^{-1}([0, t)) \in \mathcal{T}$. To see the claim, we note that if $x \in \bigcup_{q \in I, q < t} U_q$ then $x \in U_{q_0}$ for some $q_0 \in I$ with $q_0 < t$ and thus the definition of f implies that $f(x) \leq q_0 < t$ so $x \in f^{-1}([0, t))$ as desired. Hence $f^{-1}([0, t)) \supseteq \bigcup_{q \in I, q < t} U_q$. To see the other inclusion, let $x \in f^{-1}([0, t))$ be arbitrary. Since $f(x) < t$, the definition of f implies there exists a $q_0 \in I$ such that $x \in U_{q_0}$ and $q_0 < t$. Hence $x \in \bigcup_{q \in I, q < t} U_q$ as desired. Thus $f^{-1}([0, t)) = \bigcup_{q \in I, q < t} U_q$ as claimed.

Secondly, we claim that if $t \in (0, 1)$ then

$$f^{-1}((t, 1]) = \bigcup_{q \in I, q > t} X \setminus \overline{U_q}$$

thereby showing that $f^{-1}((t,1]) \in \mathcal{T}$. To see this, first suppose that $x \in f^{-1}((t,1])$. Hence $f(x) > t$. By the density of I in $[0,1]$, there exist $q_1, q_2 \in I$ such that $t < q_1 < q_2 < f(x)$. As $q_2 < f(x)$ we know that $x \notin U_{q_2}$ by the definition of f. Furthermore, as $q_1 < q_2$ we know that $\overline{U_{q_1}} \subseteq U_{q_2}$ and thus $x \notin \overline{U_{q_1}}$. Therefore, as $t < q_1$, we obtain that $x \in \bigcup_{q \in I, q > t} X \setminus \overline{U_q}$. Hence $f^{-1}((t,1]) \subseteq \bigcup_{q \in I, q>t} X \setminus \overline{U_q}$. To see the other inclusion, let $x \in \bigcup_{q \in I, q>t} X \setminus \overline{U_q}$. Hence there exists a $q_0 \in I$ such that $q_0 > t$ and $x \in X \setminus \overline{U_{q_0}}$. Hence $x \notin \overline{U_{q_0}}$ so $x \notin U_{q_0}$. Furthermore, since $U_q \subseteq U_{q_0}$ for all $q \in I$ with $q < q_0$, we see by the definition of f that $f(x) \geq q_0 > t$. Hence $x \in f^{-1}((t,1])$. Thus $f^{-1}((t,1]) = \bigcup_{q \in I, q>t} X \setminus \overline{U_q}$ as claimed.

Combining the above, we see for all $t_1, t_2 \in (0,1)$ such that $t_1 < t_2$ that

$$f^{-1}((t_1,t_2)) = f^{-1}([0,t_2)) \cap f^{-1}((t_1,1]) \in \mathcal{T}.$$

However, since every open subset of $[0,1]$ is the intersection of an open subset of $\mathbb{R}$ with $[0,1]$ and as every open subset of $\mathbb{R}$ is a countable union of open intervals by Proposition A.4.4, every open subset of $[0,1]$ is a countable union of set of the form $[0,t)$ for some $t \in (0,1)$, $(t,1]$ for some $t \in (0,1)$, and (t_1,t_2) for some $t_1,t_2 \in (0,1)$ such that $t_1 < t_2$. Consequently if V is an open subset of $[0,1]$ then $f^{-1}(V)$ is a countable union of sets of the form $f^{-1}([0,t))$ for some $t \in (0,1)$, $f^{-1}((t,1])$ for some $t \in (0,1)$, and $f^{-1}((t_1,t_2))$ for some $t_1,t_2 \in (0,1)$ such that $t_1 < t_2$, and thus open in $(X,\mathcal{T})$ by the above cases. Hence f is continuous by definition thereby yielding the proof. ∎

Of course, one natural question is whether a topological space being normal is a requirement for the conclusion of Urysohn's Lemma (Theorem 5.2.1) to hold. Indeed it is.

Proposition 5.2.2. *Let $(X,\mathcal{T})$ be a topological space such that for any disjoint closed subsets A and B of $(X,\mathcal{T})$ there exists a continuous function $f : X \to [0,1]$ such that $f(a) = 0$ for all $a \in A$ and $f(b) = 1$ for all $b \in B$. Then $(X,\mathcal{T})$ is a normal topological space.*

Proof. To see that $(X,\mathcal{T})$ is normal, let A and B be non-empty disjoint closed subsets of $(X,\mathcal{T})$. By the assumptions, there exists a continuous function $f : X \to [0,1]$ such that $f(a) = 0$ for all $a \in A$ and $f(b) = 1$ for all $b \in B$. Let

$$U = f^{-1}\left(\left[0, \frac{1}{4}\right)\right) \qquad \text{and} \qquad V = f^{-1}\left(\left(\frac{3}{4}, 1\right]\right).$$

Since f is continuous, we easily see that U and V are disjoint open subset of $(X,\mathcal{T})$ such that $A \subseteq U$ and $B \subseteq V$. Therefore, since A and B were arbitrary, $(X,\mathcal{T})$ is normal as desired. ∎

Urysohn's Lemma (Theorem 5.2.1) immediately lets us conclude our investigations related to the Stone-Weierstrass Theorems.

Corollary 5.2.3. *Let $(X, \mathcal{T})$ be a compact Hausdorff topological space. Then $\mathcal{C}(X)$ separates points. Consequently, if $\mathcal{F} \subseteq \mathcal{C}(X)$ is dense in $(\mathcal{C}(X), \|\cdot\|_\infty)$, then $\mathcal{F}$ separates points.*

Proof. To see that $\mathcal{C}(X)$ separates points, let $x_1, x_2 \in X$ be such that $x_1 \neq x_2$ be arbitrary. Since $(X, \mathcal{T})$ is Hausdorff so $\{x_1\}$ and $\{x_2\}$ are closed subsets of $(X, \mathcal{T})$, and since $(X, \mathcal{T})$ is normal by Corollary 5.1.25, Urysohn's Lemma (Theorem 5.2.1) implies there exists a $f \in \mathcal{C}(X)$ such that $f(x_1) = 0$ and $f(x_2) = 1$. Hence, as x_1 and x_2 were arbitrary, $\mathcal{C}(X)$ separates points.

To see the second claim, let $\mathcal{F}$ be dense in $(\mathcal{C}(X), \|\cdot\|_\infty)$ and let $x_1, x_2 \in X$ be such that $x_1 \neq x_2$ be arbitrary. By the above arguments there exists a $f \in \mathcal{C}(X)$ such that $f(x_1) = 0$ and $f(x_2) = 1$. Since $\mathcal{F}$ is dense in $(\mathcal{C}(X), \|\cdot\|_\infty)$ there exists a $g \in \mathcal{F}$ such that $\|f - g\|_\infty < \frac{1}{4}$. Hence, since $f(x_1) = 0$ and $f(x_2) = 1$, the definition of the infinity norm implies that

$$g(x_1) \in \left(-\frac{1}{4}, \frac{1}{4}\right) \qquad \text{and} \qquad g(x_2) \in \left(\frac{3}{4}, \frac{5}{4}\right).$$

Hence $g(x_1) \neq g(x_2)$. Hence, as x_1 and x_2 were arbitrary, $\mathcal{F}$ separates points. ∎

Urysohn's Lemma (Theorem 5.2.1) has other uses in regards to function spaces. One example of this is the following.

Theorem 5.2.4. *Let $(X, \mathcal{T})$ be a locally compact Hausdorff space. If*

$$\mathcal{C}_c(X) = \{f \in \mathcal{C}_0(X) \mid \overline{\{x \in X \mid f(x) \neq 0\}} \textit{ is compact.}\},$$

then $\mathcal{C}_c(X)$ is dense in $(\mathcal{C}_0(X), \|\cdot\|_\infty)$.

Proof. To see that $\mathcal{C}_c(X)$ is dense in $(\mathcal{C}_0(X), \|\cdot\|_\infty)$, let $f \in \mathcal{C}_0(X)$ and $\epsilon > 0$ be arbitrary. Recall since $(X, \mathcal{T})$ is a locally compact Hausdorff space, $(X, \mathcal{T})$ has a one-point compactification $(Y, \mathcal{T}_Y)$ (Theorem 3.4.7) where $Y = X \cup \{\infty\}$,

$$\mathcal{T}_Y = \mathcal{T} \cup \{Y \setminus K \mid K \subseteq X \text{ compact}\},$$

$(Y, \mathcal{T}_Y)$ is a compact Hausdorff space and thus a normal topological space (Corollary 5.1.25), and

$$\mathcal{C}_0(X) = \{g \in \mathcal{C}(Y) \mid g(\infty) = 0\}$$

(Theorem 4.2.18). Thus f extends to a function $g \in \mathcal{C}(Y)$ by $g(x) = f(x)$ for all $x \in X$ and $g(\infty) = 0$.

Let

$$U = \{y \in Y \mid |g(y)| < \epsilon\}.$$

We claim that U is open in Y. To see this, we will show that

$$F = Y \setminus U = \{y \in Y \mid |g(y)| \geq \epsilon\}$$

is closed in Y. To see this, let $(y_\lambda)_{\lambda\in\Lambda}$ be a net in F that converges to some point $y \in Y$. Since g is continuous, we know that $(g(y_\lambda))_{\lambda\in\Lambda}$ converges to $g(y)$. Since $|g(y_\lambda)| \geq \epsilon$ for all $\lambda \in \Lambda$, we obtain that $|g(y)| \geq \epsilon$. Hence $y \in F$ so F is closed by Theorem 1.6.9 and thus U is open.

Since U is an open neighbourhood of ∞ and since $(Y, \mathcal{T}_Y)$ is normal, there exists a neighbourhood $V \subseteq \mathcal{T}_Y$ of ∞ such that

$$\infty \in V \subseteq \overline{V} \subseteq U$$

where the closure of V is taken in Y. Therefore, as $F = Y \setminus U$, F and $\overline{V}$ are disjoint closed subsets of the normal topological space $(Y, \mathcal{T}_Y)$. Hence Urysohn's Lemma (Theorem 5.2.1) implies there exists a function $h : Y \to [0, 1]$ such that $h(a) = 1$ for all $a \in F$ and $h(b) = 0$ for all $b \in \overline{V}$.

Consider the function $f_0 : X \to \mathbb{R}$ defined by

$$f_0(x) = f(x)h(x)$$

for all $x \in X$. We claim that $f_0 \in \mathcal{C}_c(X)$ and $\|f - f_0\|_\infty \leq \epsilon$. To see the former, first note that if $g_0 : Y \to \mathbb{R}$ is defined by

$$g_0(y) = g(y)h(y)$$

for all $y \in Y$, then $g_0 \in \mathcal{C}(Y)$ is such that $g_0(\infty) = 0$ and $g_0(x) = f_0(x)$ for all $x \in X$. Hence $f_0 \in \mathcal{C}_0(X)$. Furthermore, notice as $h(b) = 0$ for all $b \in \overline{V}$ that

$$\{x \in X \mid f_0(x) \neq 0\} \subseteq X \setminus V.$$

However, since V was a $\mathcal{T}_Y$-neighbourhood of ∞, by the definition of $\mathcal{T}_Y$ we see that $K = X \setminus V$ is a compact set. Therefore, as $(Y, \mathcal{T}_Y)$ is a compact Hausdorff space so that K is closed in $(Y, \mathcal{T}_Y)$, we see that

$$\overline{\{x \in X \mid f_0(x) \neq 0\}} \subseteq K$$

where the closure is computed in X. Therefore, as $\overline{\{x \in X \mid f_0(x) \neq 0\}}$ is a closed subset of a compact space and thus compact, we have that $f_0 \in \mathcal{C}_c(X)$.

To see that $\|f - f_0\|_\infty \leq \epsilon$, notice for all $x \in F$ that

$$|f(x) - f_0(x)| = |f(x) - f(x)h(x)| = |f(x) - f(x)| = 0 \leq \epsilon$$

as $h(x) = 1$ for all $x \in F$. Moreover, for all $x \in X \setminus F = U$, we see that

$$|f(x) - f_0(x)| = |1 - h(x)||f(x)| \leq |f(x)| = |g(x)| < \epsilon.$$

Hence $|f(x) - f_0(x)| < \epsilon$ for all $x \in X$ so $\|f - f_0\|_\infty \leq \epsilon$. Therefore, as $f \in \mathcal{C}_0(X)$ and $\epsilon > 0$ were arbitrary, the proof is complete. ∎

5.3 Tychonoff Spaces

Of course Corollary 5.2.3 was what we were after since it shows that $\mathcal{C}(X)$ separates points for any compact Hausdorff topological space $(X, \mathcal{T})$. However, the way we showed $\mathcal{C}(X)$ separates points was to use that $(X, \mathcal{T})$ was Hausdorff so points were closed and show that $(X, \mathcal{T})$ was normal and therefore separated closed subsets via Urysohn's Lemma (Theorem 5.2.1). This is a little bit of an overkill because Urysohn's Lemma implies we can separate closed subsets of $(X, \mathcal{T})$ with continuous functions whereas we only needed to separate points with continuous functions. Thus, it is natural to ask. "for which topological spaces $(X, \mathcal{T})$ does $\mathcal{C}(X)$ separate points?" This leads us to some new notions of separations axioms.

Before we begin, we remind the reader that the study of separation axioms is notorious for conflicts with naming conventions used. Consequently, one need always be clear in their definitions when reading or writing about these topics.

The first separation axiom we desire to look at is the one that precisely implies that $\mathcal{C}(X)$ separates points.

Definition 5.3.1. A topological space $(X, \mathcal{T})$ is said to be a *functionally Hausdorff space* if for all $x_1, x_2 \in X$ such that $x_1 \neq x_2$ there exists an $f \in \mathcal{C}(X)$ such that $f(x_1) \neq f(x_2)$.

Before investigating the relation of functionally Hausdorff spaces to other separation axioms developed earlier, we note functionally Hausdorff spaces are given their name as the same ideas as Proposition 5.2.2 show that any two points can be separated via disjoint open sets coming from a function (hence the name). Using that same idea, functionally Hausdorff spaces actually have a stronger form of Hausdorffness.

Definition 5.3.2. A topological space $(X, \mathcal{T})$ is said to be *Urysohn space* if for all $x_1, x_2 \in X$ such that $x_1 \neq x_2$ there exist $U_1, U_2 \in \mathcal{T}$ and closed subsets F_1, F_2 of $(X, \mathcal{T})$ such that $x_1 \in U_1 \subseteq F_1$, $x_2 \in U_2 \subseteq F_2$, and $F_1 \cap F_2 = \emptyset$.

We note a Urysohn space receives its name as the sets described in Definition 5.3.2 that can be used to separate points are very reminiscent of the sets used in the proof of Urysohn's Lemma (Theorem 5.2.1).

Remark 5.3.3. It is essential to note that some authors use 'Urysohn space' to mean what we have defined above to be a 'functionally Hausdorff space' and use 'functionally Hausdorff space' to mean what we have defined above to be a 'Urysohn space'. Thus using these terms in the literature can be quite problematic.

Remark 5.3.4. It is not difficult to see that every functionally Hausdorff space is automatically a Urysohn space. Indeed to see this, suppose $(X, \mathcal{T})$

is a functionally Hausdorff space and let $x_1, x_2 \in X$ such that $x_1 \neq x_2$ be arbitrary. As $(X, \mathcal{T})$ is a functionally Hausdorff space, there exists a continuous function $f : X \to [0,1]$ such that $f(a) = 0$ for all $a \in A$ and $f(b) = 1$ for all $b \in B$. Let

$$U_1 = f^{-1}\left(\left[0, \frac{1}{4}\right)\right) \qquad \text{and} \qquad U_2 = f^{-1}\left(\left(\frac{3}{4}, 1\right]\right)$$

and let

$$F_1 = f^{-1}\left(\left[0, \frac{1}{4}\right]\right) \qquad \text{and} \qquad F_2 = f^{-1}\left(\left[\frac{3}{4}, 1\right]\right)$$

Since f is continuous, we easily see that U_1 and U_2 are disjoint open subset of $(X, \mathcal{T})$ and F_1 and F_2 are disjoint closed subsets of $(X, \mathcal{T})$ such that $x_1 \in U_1 \subseteq F_1$ and $x_2 \in U_2 \subseteq F_2$. Therefore, as x_1 and x_2 were arbitrary, $(X, \mathcal{T})$ is a Urysohn space.

Remark 5.3.5. Of course, one question immediate from Remark 5.3.4 is whether there are topological spaces that are Urysohn spaces but are not functionally Hausdorff spaces. Indeed there are examples of such spaces. As we have not stumbled upon a particularly simple example, we omit this example at this time. Perhaps we will include this example into a future iteration of these notes. For those that are curious, the simplest example we have found is the 'corrected Arens square'. A curious reader may Google the Arens square to find details on said example, but should be mindful of `https://math.stackexchange.com/questions/1715435/is-arens-square-a-urysohn-space`.

Remark 5.3.6. Clearly every Urysohn space is Hausdorff. Furthermore, every T_3 topological space is automatically a Urysohn space. To see this, let $(X, \mathcal{T})$ be a T_3 space and let $x_1, x_2 \in X$ with $x_1 \neq x_2$ be arbitrary. Since $(X, \mathcal{T})$ is a T_3 space and thus Hausdorff, there exists $U_1, U_2 \in \mathcal{T}$ such that $x_1 \in U_1$, $x_2 \in U_2$, and $U_1 \cap U_2 = \emptyset$. By Lemma 5.1.15 there exists $V_1, V_2 \in \mathcal{T}$ such that

$$x_1 \in V_1 \subseteq \overline{V_1} \subseteq U_1 \qquad \text{and} \qquad x_2 \in V_2 \subseteq \overline{V_2} \subseteq U_2.$$

Therefore, U_1 and U_2 are disjoint, $\overline{V_1}$ and $\overline{V_2}$ are closed disjoint subsets of $(X, \mathcal{T})$ that contain the neighbourhoods V_1 of x_1 and V_2 of x_2 respectively. Therefore, since x_1 and x_2 were arbitrary, $(X, \mathcal{T})$ is a Urysohn space.

In addition, the following two examples will illustrate that there is a topological space that is Hausdorff but not a Urysohn space and there is a topological space that is a Urysohn space but not T_3. Hence Urysohn spaces are sometimes called *$T_{2\frac{1}{2}}$ spaces*.

Example 5.3.7. In this example, we will illustrate that there exists a Hausdorff topological space that is not a Urysohn space. To begin, for each $n \in \mathbb{N}$ let $I_n = (n, n+1) \subseteq \mathbb{R}$. Let $X = \mathbb{R} \cup \{\infty_e, \infty_o\}$.

To define a topology on X, let $\mathcal{T}$ denote the set of all subsets $U \subseteq X$ such that $U \cap \mathbb{R}$ is open in canonical topology on $\mathbb{R}$, if $\infty_e \in U$ then U contains all but finitely many of I_n with n even, and if $\infty_o \in U$ then U contains all but finitely many of I_n with n odd. Some moments thought shows that $\mathcal{T}$ is indeed a topology on X as clearly $\emptyset, X \in \mathcal{T}$, the arbitrary union of sets of this form also has this form, and the finite intersection of sets of this form also has this form (as intersecting a finite number of sets that are missing a finite number of I_n with n even or a finite number of I_n with n odd also has this property).

We claim that $(X, \mathcal{T})$ is Hausdorff. To see this, let $x_1, x_2 \in X$ with $x_1 \neq x_2$ be arbitrary. We claim there exists $U_1, U_2 \in \mathcal{T}$ such that $x_1 \in U_1$, $x_2 \in U_2$, and $U_1 \cap U_2 = \emptyset$. If $x_1, x_2 \in \mathbb{R}$ such U_1 and U_2 exist since $\mathbb{R}$ is Hausdorff with its canonical topology. If $x_1 = \infty_e$ and $x_2 = \infty_o$, then $U_1 = \bigcup_{n\in\mathbb{N}} I_{2n}$ and $U_2 = \bigcup_{n\in\mathbb{N}} I_{2n+1}$ are open sets in $(X, \mathcal{T})$ that have the desired properties. Hence, by relabelling x_1 and x_2 if necessary, it remains only to consider the case that $x_1 \in \mathbb{R}$ and $x_2 \in \{\infty_e, \infty_o\}$. Suppose $x_2 = \infty_e$. If $x_1 \leq 0$, we may take $U_1 = (x_1 - 1, x_1 + 1)$ and $U_2 = \bigcup_{n\in\mathbb{N}} I_{2n}$ which are open sets in $(X, \mathcal{T})$ that have the desired properties. Otherwise $x_1 \in (n-1, n+1)$ for some $n \in \mathbb{N}$. Hence taking $U_1 = (n-1, n+1)$ and $U_2 = \bigcup_{n\in\mathbb{N}} I_{2n} \setminus U_1$ are open sets in $(X, \mathcal{T})$ that have the desired properties. As the case $x_2 = \infty_o$ proceeds in a similar fashion, $(X, \mathcal{T})$ is Hausdorff as desired.

Finally, we claim that $(X, \mathcal{T})$ is not a Urysohn space. To see this, let U_1 be an arbitrary neighbourhood of ∞_e, let U_2 be an arbitrary neighbourhood of ∞_o, and let F_1 and F_2 be arbitrary closed subsets of $(X, \mathcal{T})$ such that $U_1 \subseteq F_1$ and $U_2 \subseteq F_2$. As clearly $[n, n+1] \subseteq \overline{I_n}$ when the closure of I_n is taking in $\mathbb{R}$ equipped with the canonical topology, $[n, n+1] \subseteq \overline{I_n}$ when the closure of I_n is taking in $(X, \mathcal{T})$ by Theorem 1.6.21. Hence by the defining properties of U_1 and U_2, $\overline{U_1} \subseteq F_1$ and $\overline{U_2} \subseteq F_2$ contain all but a finite number of the natural numbers. Hence $F_1 \cap F_2$ must contain an infinite number of the natural numbers and thus is not empty. Hence, as U_1, U_2, F_1, and F_2 were arbitrary, $(X, \mathcal{T})$ is not a Urysohn space.

Example 5.3.8. In this example, we will illustrate that there exists a Urysohn space that is not T_3. Indeed our example is the same as Example 5.1.14. In Example 5.1.14, a topology $\mathcal{T}$ was constructed on $\mathbb{R}$ that was finer than the canonical topology on $\mathbb{R}$ and it was shown that $(\mathbb{R}, \mathcal{T})$ was not T_3. However, the the canonical topology on $\mathbb{R}$ is a Urysohn space. Indeed for all $x_1, x_2 \in \mathbb{R}$ with $x_1 < x_2$ one can take

$$\begin{aligned} U_1 &= (x_1 - \delta, x_1 + \delta), & F_1 &= [x_1 - \delta, x_1 + \delta], \\ U_2 &= (x_2 - \delta, x_2 + \delta), & F_2 &= [x_2 - \delta, x_2 + \delta] \end{aligned}$$

where $0 < \delta < x_2 - x_1$ in Definition 5.3.2. Therefore, since $\mathcal{T}$ is finer than the canonical topology on $\mathbb{R}$, we immediately obtain that $(\mathbb{R}, \mathcal{T})$ is a Urysohn space by Definition 5.3.2.

We also the (corrected) Arens square we omitted earlier is a Urysohn space but not T_3.

Instead of just reducing the closed subsets in Urysohn's Lemma (Theorem 5.2.1) down to points in order to obtain a separation axiom, we can instead take a similar thought pattern to our notion of regularity by simply reducing one of the closed sets to a point. Thus we define the following. Please note that we have not adopted one of the many unstandardised names used in the literature for the following property and instead decided to adopt our own which seems by far more rational than any other name seen in the literature.

Definition 5.3.9. A topological space $(X, \mathcal{T})$ is said to be a *functionally regular space* if whenever F is a closed subset of $(X, \mathcal{T})$ and $x_0 \in X \setminus F$ there exist an $f \in \mathcal{C}(X)$ such that $f(x_0) = 1$ and $f(y) = 0$ for all $y \in F$.

Remark 5.3.10. Of course, a topological spaces $(X, \mathcal{T})$ is functionally regular if and only if for any closed subsets F of $(X, \mathcal{T})$, $x_0 \in X \setminus F$, and $a, b \in \mathbb{R}$ with $a \neq b$ there exist an $f \in \mathcal{C}(X)$ such that $f(x_0) = a$ and $f(y) = b$ for all $y \in F$. That is, we can replace 0 and 1 in Definition 5.3.9 with any distinct selection of real numbers by composing with certain functions on $\mathbb{R}$.

Remark 5.3.11. It is not difficult to see that every functionally regular space is automatically regular. Indeed to see this, suppose $(X, \mathcal{T})$ is a functionally regular space space. To see that $(X, \mathcal{T})$ is regular, let F be an arbitrary closed subset of $(X, \mathcal{T})$ and $x_0 \in X \setminus F$ be arbitrary. As $(X, \mathcal{T})$ is a functionally regular space, there exists a continuous function $f : X \to [0, 1]$ such that $f(x_0) = 0$ and $f(x) = 1$ for all $x \in F$. Let

$$U = f^{-1}\left(\left[0, \frac{1}{4}\right)\right) \qquad \text{and} \qquad V = f^{-1}\left(\left(\frac{3}{4}, 1\right]\right).$$

Since f is continuous, we easily see that U and V are disjoint open subset of $(X, \mathcal{T})$ such that $x_0 \in U$ and $F \subseteq V$. Therefore, as x_0 and F were arbitrary, $(X, \mathcal{T})$ is a regular topological space.

Of course the trivial topology on any set with at least two points is automatically functionally regular as one will not be able to find a non-empty closed set whose complement contains a point. Thus, as we are investigating for which topological spaces $(X, \mathcal{T})$ does $\mathcal{C}(X)$ separate points, we desire to add in the property we all know and love to functional regularity in order to deduce that points are closed.

Definition 5.3.12. A topological space $(X, \mathcal{T})$ is said to be a *Tychonoff space* if $(X, \mathcal{T})$ is a functionally regular, Hausdorff space.

Of course, we do not need to add the full power of Hausdorff to functionally regular spaces in order to obtain a Tychonoff space.

Proposition 5.3.13. *A topological space $(X, \mathcal{T})$ is Tychonoff if and only if $(X, \mathcal{T})$ is a functionally regular T_0 space.*

Proof. Clearly every Tychonoff space is a functionally regular T_0 space. To see the converse, suppose that $(X, \mathcal{T})$ is a functionally regular T_0 space. To show that $(X, \mathcal{T})$ is Tychonoff, it suffices to prove that $(X, \mathcal{T})$ is Hausdorff. To see that $(X, \mathcal{T})$ is Hausdorff, let $x_1, x_2 \in X$ with $x_1 \neq x_2$ be arbitrary. Since $(X, \mathcal{T})$ is T_0, by relabelling x_1 and x_2 if necessary, there exists a $U \in \mathcal{T}$ such that $x_1 \in U$ and $x_2 \notin U$. Let $F = X \setminus U$ which is a closed subset of $(X, \mathcal{T})$ such that $x_2 \in F$ and $x_1 \notin F$. Therefore, since $(X, \mathcal{T})$ is functionally regular, there exist an $f \in \mathcal{C}(X)$ such that $f(x_1) = 1$ and $f(y) = 0$ for all $y \in F$. In particular, $f(x_1) \neq f(x_2)$. Therefore, since x_1 and x_2 were arbitrary, $(X, \mathcal{T})$ is Urysohn and thus Hausdorff by Remark 5.3.4 and Remark 5.3.6 as desired. ∎

Remark 5.3.14. Of course, in order to get examples of Tychonoff spaces, Urysohn's Lemma (Theorem 5.2.1) automatically implies every T_4 topological space is Tychonoff. Hence metric spaces and complete Hausdorff spaces are Tychonoff.

In addition, clearly Remark 5.3.11 implies that every Tychonoff space is a T_3 space. In fact, as there are examples of T_3 spaces that are not Tychonoff and examples of Tychonoff spaces that are not T_4 spaces, Tychonoff spaces are often called *$T_{3\frac{1}{2}}$ spaces*.

In fact, Tychonoff spaces are automatically functionally Hausdorff spaces since every point is a closed set in a Tychonoff space.

As the relations between our separation axioms are probably getting more and more difficult for the reader to follow as we add more and more definitions, we note the following diagram where a $A \implies B$ means that every space that satisfies axiom A satisfies axiom B too.

$$\begin{tikzcd}
\mathrm{T}_4 \arrow[r, Rightarrow] & \mathrm{T}_{3\frac{1}{2}} \arrow[r, Rightarrow] \arrow[d, Rightarrow] & \text{Functionally Hausdorff} \arrow[d, Rightarrow] & & & \\
 & \mathrm{T}_3 \arrow[r, Rightarrow] & \text{Urysohn} \arrow[r, Rightarrow] & \mathrm{T}_2 \arrow[r, Rightarrow] & \mathrm{T}_1 \arrow[r, Rightarrow] & \mathrm{T}_0
\end{tikzcd}$$

In fact, all of the $\implies$ in the above diagram are strict. However, some of these implications have not been demonstrated via examples. In particular, we have not given an example of a topological space $(X, \mathcal{T})$ that is

(1) Urysohn but not functionally Hausdorff;

(2) T_3 but not Tychonoff;

(3) functionally Hausdorff but not Tychonoff;

(4) T_3 but not functionally Hausdorff;

(5) functionally Hausdorff but not T_3;

(6) Tychonoff but not T_4.

Of course, we could spend time to demonstrate six examples (or less if there are repeats) thereby completing our complete picture of these few separation axioms that we have listed. This of course is useful material for those that want a plethora of examples, but not so important for understanding the theory of topology.

Instead, we shall go and study Tychonoff spaces. One reason for this is that they are just a step weaker than T_4 spaces and thus particularly nice spaces. In fact, by studying Tychonoff spaces, we will quite quickly see that there must be a Tychonoff space that is not T_4 because Tychonoff spaces are far better behaved than T_4 spaces under the topological operations we have seen in this course as the following result shows.

Lemma 5.3.15. *Any subspace of a Tychonoff space is Tychonoff.*

Proof. Let $(X, \mathcal{T})$ be a Tychonoff space and let Y be a subspace of X. To see that Y is Tychonoff, first note that Y is Hausdorff since $(X, \mathcal{T})$ is Hausdorff. Thus it suffices to show that Y is functionally regular. To see this, let F be an arbitrary closed subset of Y and let $y \in Y \setminus F$ be arbitrary. Since F is closed in Y, there must exists a closed set C in $(X, \mathcal{T})$ such that $F = Y \cap C$. Since $y \in Y \setminus F$, we see that $y \in X \setminus C$. Therefore, since $(X, \mathcal{T})$ is Tychonoff, there exists an $f \in \mathcal{C}(X)$ such that $f(y) = 1$ and $f(x) = 0$ for all $x \in C$. Since $g = f|_Y \in \mathcal{C}(Y)$ by Lemma 2.1.18, since $g(y) = 1$, and since $g(x) = 0$ for all $x \in F = Y \cap C$, we have obtained the necessary continuous function on Y. Therefore, since F and y were arbitrary, Y is a Tychonoff space. ∎

Corollary 5.3.16. *There exists a topological space $(X, \mathcal{T})$ that is a Tychonoff space but not a T_4 space.*

Proof. By Remark 5.3.14 every T_4 space is a Tychonoff space. Hence Lemma 5.3.15 implies every subspace of a T_4 space is a Tychonoff space. Therefore, if every Tychonoff space were a T_4 space, then every subspace of a T_4 space would be a T_4 space. However, this contradicts Remark 5.1.30 thereby proving the result. ∎

Not only do Tychonoff spaces play well with respect to subspaces, they also play well with respect to products.

Lemma 5.3.17. *Let $\{(X_\alpha, \mathcal{T}_\alpha)\}_{\alpha \in I}$ be a non-empty collection of non-empty functionally regular spaces. If $\mathcal{T}$ denotes the product topology on $X = \prod_{\alpha \in I} X_\alpha$, then $(X, \mathcal{T})$ is a functionally regular space. Consequently, the product of Tychonoff spaces is a Tychonoff space with respect to the product topology.*

Proof. We first note that if $(X_\alpha, \mathcal{T}_\alpha)$ is Hausdorff then $(X, \mathcal{T})$ is Hausdorff is Hausdorff. Hence the second claim easily follows from the first. To see that $(X, \mathcal{T})$ is functionally regular, let F be a non-empty closed subset of $(X, \mathcal{T})$ and let $x_0 \in X \setminus F$ be arbitrary. If $F = \emptyset$, then there is nothing to prove since one may take a constant function $f \in \mathcal{C}(X)$ to satisfy the definition of a functionally regular topological space. Hence we will assume that $F \neq \emptyset$.

Consider $U = X \setminus F$ which is a neighbourhood of x_0 in $(X, \mathcal{T})$ such that $U \neq X$. Therefore, due to our knowledge of a neighbourhood basis of x_0 in the product topology, there exists a set $V = \prod_{\alpha \in I} V_\alpha \in \mathcal{T}$ such that $x_0 \in V \subseteq U$, $V_\alpha \in \mathcal{T}_\alpha$ for all $\alpha \in I$, and there exists a finite subset $J \subseteq I$ such that $V_\alpha = X_\alpha$ if and only if $\alpha \in I \setminus J$. Furthermore, since $V \subseteq U \neq X$, it must be the case that $J \neq \emptyset$.

For each $\alpha \in J$, let $F_\alpha = X_\alpha \setminus V_\alpha$. Hence F_α is a non-empty closed subset of $(X_\alpha, \mathcal{T}_\alpha)$ for all $\alpha \in J$. For each $\beta \in J$, let

$$C_\beta = \prod_{\alpha \in I} C_{\alpha,\beta}$$

where

$$C_{\alpha,\beta} = \begin{cases} X_\alpha & \text{if } \alpha \neq \beta \\ F_\alpha & \text{if } \alpha = \beta \end{cases}.$$

Thus C_β is closed in $(X, \mathcal{T})$ by Example 1.6.15 for all $\beta \in J$. Hence, as J is finite, $F' = \bigcup_{\beta \in J} C_\beta$ is closed on $(X, \mathcal{T})$. Furthermore, it is elementary to see that $F' = X \setminus V \supseteq F$ and thus $x_0 \notin F'$ as $x_0 \in V$. Thus it suffices to construct a function $f \in \mathcal{C}(X)$ such that $f(x_0) = 1$ and $f(x) = 0$ for all $x \in F'$.

For each $\alpha \in J$, note F_α is a closed subset of $(X_\alpha, \mathcal{T}_\alpha)$ and $x_0(\alpha) \in X_\alpha \setminus F_\alpha$. Therefore, since $(X_\alpha, \mathcal{T}_\alpha)$ is functionally regular, for each $\alpha \in J$ there exists an $f_\alpha \in \mathcal{C}(X_\alpha)$ such that $f_\alpha(x_0(\alpha)) = 1$ and $f_\alpha(y) = 0$ for all $y \in F_\alpha$. Define $f : X \to \mathbb{R}$ by

$$f(x) = \prod_{\alpha \in J} f_\alpha(x(\alpha))$$

for all $x \in X$; that is

$$f = \prod_{\alpha \in J} f_\alpha \circ \pi_\alpha$$

where $\pi_\alpha : X \to X_\alpha$ is the projection map from Example 2.1.8. Since the projections maps are continuous and the composition of continuous maps are continuous, f is a finite product (as J is finite) of continuous functions in $\mathbb{R}$ and thus a well-defined continuous function from X into $\mathbb{R}$. Furthermore, since $f_\alpha(x_0(\alpha)) = 1$ and $f_\alpha(y) = 0$ for all $y \in F_\alpha$ and $\alpha \in J$, we see that $f(x_0) = 1$ and $f(x) = 0$ for all $x \in F'$ (as each $x \in F'$ has an entry in F_α for some $\alpha \in J$) as desired. Hence, as F and x_0 were arbitrary, $(X, \mathcal{T})$ is a Tychonoff space. ∎

Of course, Lemma 5.3.17 immediately allows us to explicitly give an example of a Tychonoff space that is not a T_4 space thereby removing the abstraction of Corollary 5.3.16.

Example 5.3.18. Recall from Example 5.1.29 that $\mathcal{F}(\mathbb{R}, \mathbb{R}) = \prod_{x\in\mathbb{R}} \mathbb{R}$ equipped with the product topology is not a T_4 space. However, since $\mathbb{R}$ is a metric space, $\mathbb{R}$ is a T_4 space by Theorem 5.1.23. Hence $\mathbb{R}$ is a Tychonoff space by Remark 5.3.14 so Lemma 5.3.17 implies that $\mathcal{F}(\mathbb{R}, \mathbb{R})$ is a Tychonoff space.

Of course, the proof of Lemma 5.3.17 breaks down if the product topology is replaced with the box topology since we cannot easily take the infinite product of a collection of continuous functions and get a continuous function. This may seem surprising since the box topology is finer than the product topology. However, knowing the box topology is finer than the product topology not only causes each function constructed in Lemma 5.3.17 to be continuous from the box topology (which is good because we have functions that separate some closed sets from points), it causes there to be a greater number of closed sets that need to be separated from points. Unfortunately, the above is not a proof that Lemma 5.3.17 fails for the box topology. In fact, Lemma 5.3.17 holds for the box topology on our favourite product space $\prod_{x\in\mathbb{R}} \mathbb{R}$. Thus, as we have seen time and time again that we are interested in the product topology and not the box topology and as we are really interested in understanding Tychonoff spaces, we will ignore this question and get on with what we want to study.

In particular, an engrossed reader is probably wondering why Tychonoff spaces are called Tychonoff spaces. Perhaps they were just named after Tychonoff who happened to have be the first to study them? Well, a better reason relates to Tychonoff's Theorem (Theorem 3.3.4). In particular, we endeavour to show that Tychonoff spaces are precisely the topological spaces that are subspaces of topological spaces studied in Tychonoff's Theorem. To do so, we first prove the following theorem that will aid us in this and future goals.

Theorem 5.3.19 (Embedding Theorem). *Let $(X, \mathcal{T})$ be a T_0 space and let $\{f_\alpha\}_{\alpha\in I} \subseteq \mathcal{C}(X, \mathbb{R})$ be such that for each $x_0 \in X$ and neighbourhood U of x_0 there exists an $\alpha_0 \in I$ such that $f_{\alpha_0}(x_0) \neq 0$ and $f(x) = 0$ for all $x \in X \setminus U$ (this automatically implies that $(X, \mathcal{T})$ is Tychonoff). If $Y = \prod_{\alpha\in I} \mathbb{R}$ is equipped with the product topology, then function $F : X \to Y$ defined by*

$$F(x) = (f_\alpha(x))_{\alpha\in I}$$

is an embedding of $(X, \mathcal{T})$ into Y.

Proof. To begin to show that F is an embedding of $(X, \mathcal{T})$ into Y, we first show that F is injective. To see this, let $x_1, x_2 \in X$ with $x_1 \neq x_2$

be arbitrary. Since $(X, \mathcal{T})$ is T_0 space, up to relabelling x_1 and x_2 there exists a neighbourhood U of x_1 such that $x_2 \in X \setminus U$. Thus, by the above construction, there exists an $\alpha \in I$ such that $f_\alpha(x_1) \neq 0 = f_\alpha(x_2)$. Hence $F(x_1) \neq F(x_2)$. Therefore, as x_1 and x_2 were arbitrary, F is injective.

Next, notice that F is clearly continuous by Theorem 2.1.10 as each f_α is continuous. Thus to show that F is an embedding of X into Y, it suffices to show that if $Z = F(X)$ and if $U \in \mathcal{T}$ then $F(U)$ is open in Z.

To see the above, let $U \in \mathcal{T}$ be arbitrary. To see that $F(U)$ is open in Z, let $z_0 \in F(U)$ be arbitrary. Our goal is to show there is a neighbourhood of z_0 from Z contained in $F(U)$. Thus, as $z_0 \in F(U)$, choose $x_0 \in U$ such that $F(x_0) = z_0$. By the above construction, there exists an $\alpha_0 \in I$ such that $f_{\alpha_0}(x_0) \neq 0$ and $f_{\alpha_0}(x) = 0$ for all $x \in X \setminus U$. Let

$$V = \pi_{\alpha_0}^{-1}(\mathbb{R} \setminus \{0\}) \cap Z$$

where π_{α_0} is the projection map from Example 2.1.8. Clearly $z_0 \in V$ by construction and V is open in Z as Z is a subspace of Y and $\pi_{\alpha_0}^{-1}(\mathbb{R} \setminus \{0\})$ is open in Y as π_{α_0} is a continuous function from the product topology to $\mathbb{R}$. To see that $V \subseteq F(U)$, let $y \in V$ be arbitrary. Then $y \in Z$ by definition so $y = F(x)$ for some $x \in X$. Since $y = F(x) \in V$, it must be the case that $f_{\alpha_0}(x) \neq 0$. Hence the defining property of f_{α_0} implies that $x \notin X \setminus U$. Thus $x \in U$ so $y = f(x) \in F(U)$ as desired. Hence $z_0 \in V \subseteq F(U)$. Therefore, since $z_0 \in F(U)$ was arbitrary, $F(U)$ is open in Z. Hence, as $U \in \mathcal{T}$ was arbitrary, F is an embedding of X into Y as claimed. ∎

Theorem 5.3.20. *A topological space $(X, \mathcal{T})$ is Tychonoff if and only if there exists a non-empty collection of non-empty compact Hausdorff topological spaces $\{(X_\alpha, \mathcal{T}_\alpha)\}_{\alpha \in I}$ such that $(X, \mathcal{T})$ is homeomorphic to a subspace of $(\prod_{\alpha \in I} X_\alpha, \mathcal{T}_p)$ where $\mathcal{T}_p$ is the product topology. In fact, if $(X, \mathcal{T})$ is Tychonoff, one may take $X_\alpha = [0, 1]$ with its canonical topology for all $\alpha \in I$.*

Proof. First, suppose there exists a non-empty collection of non-empty compact Hausdorff topological spaces $\{(X_\alpha, \mathcal{T}_\alpha)\}_{\alpha \in I}$ such that $(X, \mathcal{T})$ is homeomorphic to a subspace of $(\prod_{\alpha \in I} X_\alpha, \mathcal{T}_p)$ where $\mathcal{T}_p$ is the product topology. By Tychonoff's Theorem (Theorem 3.3.4), $(\prod_{\alpha \in I} X_\alpha, \mathcal{T}_p)$ is a compact Hausdorff topological space. Hence $(\prod_{\alpha \in I} X_\alpha, \mathcal{T}_p)$ is normal by Corollary 5.1.25 and thus a Tychonoff space by Remark 5.3.14. Alternatively, one can instead claim that $(X_\alpha, \mathcal{T}_\alpha)$ is normal by Corollary 5.1.25 and thus a Tychonoff space by Remark 5.3.14 and appeal to Lemma 5.3.17 to bypass Tychonoff's Theorem (Theorem 3.3.4). Either way, $(\prod_{\alpha \in I} X_\alpha, \mathcal{T}_p)$ is a Tychonoff space so Lemma 5.3.15 implies that a subspace of $(\prod_{\alpha \in I} X_\alpha, \mathcal{T}_p)$ is a Tychonoff space. Hence $(X, \mathcal{T})$ is homeomorphic to a Tychonoff space and thus easily seen to be Tychonoff.

Conversely, suppose that $(X, \mathcal{T})$ is a Tychonoff space. By the defining properties of a Tychonoff space, there exists a set $\{f_\alpha\}_{\alpha \in I} \subseteq \mathcal{C}(X)$ such that for each $x_0 \in X$ and neighbourhood U of x_0 there exists an $\alpha_0 \in I$ such that $f_{\alpha_0}(x_0) \neq 0$ and $f(x) = 0$ for all $x \in X \setminus U$. Hence the Embedding Theorem (Theorem 5.3.19) implies there exists an embedding of X into $Y = \prod_{\alpha \in I} \mathbb{R}$. Since $\mathbb{R}$ is homeomorphic to $(0, 1)$, we know that Y is homeomorphic to $\prod_{\alpha \in I}(0, 1)$ equipped with the product topology. As clearly $\prod_{\alpha \in I}(0, 1)$ is a subspace of $\prod_{\alpha \in I}[0, 1]$ equipped with product topology, we see that $(X, \mathcal{T})$ is homeomorphic to a subspace of $\prod_{\alpha \in I}[0, 1]$ as the composition of embeddings is an embedding. Hence the proof is complete. ∎

5.4 Stone-Čech Compactification

The beauty of Tychonoff spaces is that they are are precisely the topological spaces that embed directly into a product of compact Hausdorff topological spaces. This is very reminiscent of the result that every metric space is isomorphic to a subspace of a complete metric space and thus has a completion. We desire to use Theorem 5.3.20 to show that every Tychonoff space has some form of compactification similar to how we showed that every locally compact Hausdorff topological space has a one-point compactification. In fact, there will be a very special form of compactification we desire to study. Before we get to that point, we should formally define what we mean by a compactification.

Definition 5.4.1. A *compactification of a topological space* $(X, \mathcal{T}_X)$ is a compact Hausdorff topological space $(Y, \mathcal{T}_Y)$ with the property that there exists an embedding $f : X \to Y$ such that $\overline{f(X)} = Y$.

Two compactifications $(Y_1, \mathcal{T}_1)$ and $(Y_2, \mathcal{T}_2)$ of $(X, \mathcal{T}_X)$ are said to be *equivalent* if there exists a homeomorphism $h : Y_1 \to Y_2$ such that $h(f_1(x)) = f_2(x)$ for all $x \in X$ where $f_1 : X \to Y_1$ and $f_2 : X \to Y_2$ are the embeddings of $(X, \mathcal{T})$ into $(Y_1, \mathcal{T}_1)$ and $(Y_2, \mathcal{T}_2)$ respectively such that $\overline{f_1(X)} = Y_1$ and $\overline{f_2(X)} = Y_2$.

Of course, by our above work, we already know exactly which topological spaces have compactifications.

Theorem 5.4.2. *A topological space $(X, \mathcal{T})$ has a compactification if and only if $(X, \mathcal{T})$ is a Tychonoff space.*

Proof. Suppose $(X, \mathcal{T})$ is a topological space with a compactification $(Y, \mathcal{T}_Y)$. Hence $(Y, \mathcal{T}_Y)$ is a compact Hausdorff topological space and thus a Tychonoff space by Corollary 5.1.25 and Remark 5.3.14. Since every subspace of a Tychonoff space is a Tychonoff space by Lemma 5.3.15, $(X, \mathcal{T})$ is homeomorphic to a Tychonoff space and thus is a Tychonoff space.

Conversely, suppose that $(X, \mathcal{T})$ is a Tychonoff space. By Theorem 5.3.20 there exists a non-empty index set I such that if $[0, 1]$ is equipped with the canonical subspace topology inherited from $\mathbb{R}$ and $\mathcal{T}_p$ is the product topology on $Y = \prod_{\alpha \in I}[0,1]$, then there exists an embedding $f : X \to Y$. Let $Y_0 = \overline{f(X)}$ equipped with the subspace topology inherited from $(Y, \mathcal{T}_p)$. Therefore, since $(Y, \mathcal{T}_p)$ is a compact Hausdorff topological space, Y_0 is a compact Hausdorff topological space by Theorem 3.1.14. Hence Y_0 is a compactification of $(X, \mathcal{T})$ by definition. ∎

Immediately we can increase our known collection of topological spaces that are Tychonoff.

Corollary 5.4.3. *Let $(X, \mathcal{T})$ be a locally compact Hausdorff topological space that is not compact. Then the one-point compactification of $(X, \mathcal{T})$ is a compactification of $(X, \mathcal{T})$. Consequently, every locally compact Hausdorff topological space is a Tychonoff space.*

Proof. First suppose $(X, \mathcal{T})$ is a locally compact Hausdorff topological space that is not compact and let $(Y, \mathcal{T}_Y)$ be the one-point compactification of $(X, \mathcal{T})$. We claim that $(Y, \mathcal{T}_Y)$ is indeed a compactification of $(X, \mathcal{T})$ in the sense of Definition 5.4.1. To begin, we note $(Y, \mathcal{T}_Y)$ is a compact Hausdorff topological space so it suffices to show that $\overline{X} = Y$. Let $Y \setminus X = \{\infty\}$ and recall that from Theorem 3.4.7 that

$$\mathcal{T}_Y = \mathcal{T} \cup \{Y \setminus K \mid K \text{ a compact subset of } (X, \mathcal{T})\}.$$

Hence, as X is not compact, every neighbourhood of ∞ in (Y, T_Y) contains an element of X. Thus Theorem 1.6.21 implies that $\overline{X} = Y$ as desired.

For the second claim, suppose $(X, \mathcal{T})$ is a locally compact Hausdorff topological space. Clearly if $(X, \mathcal{T})$ is a compact, then $(X, \mathcal{T})$ is a Tychonoff space by Corollary 5.1.25 and Remark 5.3.14. Otherwise, if $(X, \mathcal{T})$ is not compact, the above shows that the one-point compactification of $(X, \mathcal{T})$ is indeed a compactification of $(X, \mathcal{T})$ in the sense of Definition 5.4.1 thereby implying that $(X, \mathcal{T})$ is Tychonoff by Theorem 5.4.2. ∎

Of course, the above begs the question of whether the one-point compactification of a compact Hausdorff topological space is actually a compactification?

Proposition 5.4.4. *Let $(X, \mathcal{T}_X)$ be a compact Hausdorff topological space. Then the only compactification of $(X, \mathcal{T})$ up to homeomorphism is $(X, \mathcal{T})$.*

Proof. Let $(Y, \mathcal{T}_Y)$ be a compactification of a compact Hausdorff topological space $(X, \mathcal{T})$. Hence there exists an embedding $f : X \to Y$ such that $\overline{f(X)} = Y$. However, since $(X, \mathcal{T})$ is compact and f is continuous, Theorem 3.1.27 implies that $f(X)$ is a compact subset of $(Y, \mathcal{T}_Y)$. However, since $(Y, \mathcal{T}_Y)$ is Hausdorff, this implies that $f(X)$ is closed in $(Y, \mathcal{T}_Y)$ by Theorem 3.1.13. Hence $f(X) = \overline{f(X)} = Y$ so f is a homeomorphism as desired. ∎

Remark 5.4.5. Of course, Proposition 5.4.4 still doesn't immediately imply that the one-point compactification of a compact Hausdorff topological space is not a compactification because it may be possible that the one-point compactifications of compact Hausdorff topological spaces are always homeomorphic to their original spaces. Unfortunately this is not the case. To see this, suppose $(X, \mathcal{T})$ is a compact Hausdorff topological space and let $(Y, \mathcal{T}_Y)$ be the one-point compactification of $(X, \mathcal{T})$. Let $Y \setminus X = \{\infty\}$. Since X is compact and thus closed in $(Y, \mathcal{T}_Y)$ by Theorem 3.1.13 as $(Y, \mathcal{T}_Y)$ is Hausdorff, we see that $\{\infty\}$ is an open subset of $(Y, \mathcal{T}_Y)$. Hence, if $\{x\}$ is not open in $(X, \mathcal{T})$ for every $x \in X$, then it is impossible that $(X, \mathcal{T})$ and $(Y, \mathcal{T}_Y)$ are homeomorphic. Clearly this is the case when $X = [0, 1]$ equipped with the canonical subspace topology. For another example, $X = \{0\}$ then $Y = \{0, \infty\}$ which cannot be homeomorphic to X as homeomorphisms must be bijections and we have a cardinality restriction.

Clearly this is not an issue. Indeed if $(X, \mathcal{T})$ is a compact Hausdorff topological space, then it is already a compact Hausdorff topological space and its own unique completion by Proposition 5.4.4. That is, why would we ever want to consider the one-point compactification of a compact Hausdorff topological space anyways?

Of course, if we do not have a compact Hausdorff topological space, there are many possible compactifications that may occur. This is even true for locally compact Hausdorff topological spaces, such as $(0, 1)$.

Example 5.4.6. Let $(0, 1)$ be equipped with the subspace topology inherited from the canonical topology on $\mathbb{R}$. Then Example 3.4.9 implies the one-point compactification of $(0, 1)$ is

$$S^1 = \{(x, y) \mid x^2 + y^2 = 1\}.$$

Hence S^1 is a compactification of $(0, 1)$ by Corollary 5.4.3.

Example 5.4.7. Let $(0, 1)$ be equipped with the subspace topology inherited from the canonical topology on $\mathbb{R}$. If $Y = [0, 1]$ equipped with the subspace topology inherited from the canonical topology on $\mathbb{R}$, then clearly Y is a compact Hausdorff topological space that contains $(0, 1)$ such that $\overline{(0, 1)} = Y$. Hence $[0, 1]$ is a compactification of $(0, 1)$.

Example 5.4.8. Let $(0, 1)$ be equipped with the subspace topology inherited from the canonical topology on $\mathbb{R}$. Consider the topologist's sine curve

$$Z = \{(0, y) \mid -1 \leq y \leq 1\} \cup \left\{ \left(x, \sin\left(\frac{1}{x}\right)\right) \,\middle|\, 0 < x \leq \frac{1}{\pi} \right\}.$$

Clearly Z is a compact Hausdorff subspace of $\mathbb{R}^2$ by the Heine-Borel Theorem as Z is closed by Example 2.3.10. We claim Z is a compactification of $(0, 1)$.

To see this, define $f : (0,1) \to Z$ by

$$f(x) = \left(\frac{x}{\pi}, \sin\left(\frac{\pi}{x}\right)\right).$$

It is elementary to see that f is an embedding of X into

$$A = \left\{ \left(x, \sin\left(\frac{1}{x}\right)\right) \,\middle|\, 0 < x < \frac{1}{\pi} \right\}.$$

Since the closure of A in Z is Z by similar arguments to those used in Example 2.3.10, we obtain that Z is a compactification of $(0,1)$.

Remark 5.4.9. Let S^1, Y, and Z be the compactifications of $(0,1)$ from Examples 5.4.6, 5.4.7, and 5.4.8 respectively. We claim these compactifications of $(0,1)$ are pairwise non-homeomorphic. To see this, we note by Example 2.4.5 that Z is not path connected whereas it is trivial to see that S^1 and Y are path connected. Therefore, as the notion of path connectedness is invariant under homeomorphisms, Z is not homeomorphic to S^1 nor to Y. To see that Y and S^1 are not homeomorphic, suppose to the contrary that $f : Y \to S^1$ is a homeomorphism. Hence if $y = f\left(\frac{1}{2}\right)$, then $f|_{Y \setminus \left\{\frac{1}{2}\right\}}$ is a homeomorphism from $Y \setminus \left\{\frac{1}{2}\right\}$ to $S^1 \setminus \{y\}$. However $Y \setminus \left\{\frac{1}{2}\right\}$ is not a connected set whereas $S^1 \setminus \{y\}$ is connected. Hence the Intermediate Value Theorem (Theorem 2.3.5) applied to $f^{-1}|_{S \setminus \{y\}}$ yields a contradiction. Hence S^1 and Y are not homeomorphic.

Since there are multiple compactifications of a Tychonoff space, which one should we take? Well, one avenue towards this question is to consider the continuous function functions. In particular, given a continuous function f on a Tychonoff space (X, T), wouldn't it be awesome if we could extend f to the compactification? This would then mean we could comprehend all of the continuous functions on Tychonoff spaces as continuous functions on compact Hausdorff topological spaces which are incredibly well understood by our previous chapters! Of course, we would only be able to handle the bounded continuous functions on the Tychonoff space due to functions on compact Hausdorff topological spaces being bounded, but there should be lots of such functions by the definition of a Tychonoff space. In addition, we do not need to worry about having multiple ways of extending a continuous function to a compactification as the following lemma shows.

Lemma 5.4.10. *Let $(X, \mathcal{T}_X)$ and $(Y, \mathcal{T}_Y)$ be topological spaces with $(Y, \mathcal{T}_Y)$ Hausdorff. If $A \subseteq X$ and if $g_1, g_2 : \overline{A} \to Y$ are continuous functions such that $g_1|_A = g_2|_A$, then $g_1 = g_2$.*

Proof. To see that $g_1 = g_2$, let $a \in \overline{A}$ be arbitrary. By Theorem 1.6.21 there exists a net $(a_\lambda)_{\lambda \in \Lambda}$ of elements of a that converge to a in $(X, \mathcal{T}_X)$. Hence, since g_1 and g_2 are continuous functions on A, Theorem 2.1.9 implies that $(g_1(a_\lambda))_{\lambda\in\Lambda}$ converges to $g_1(a)$ in $(Y, \mathcal{T}_Y)$ and $(g_2(a_\lambda))_{\lambda\in\Lambda}$ converges to $g_2(a)$ in $(Y, \mathcal{T}_Y)$. Therefore, since $g_1(a_\lambda) = g_2(a_\lambda)$ for all $\lambda \in \Lambda$ since $a_\lambda \in A$ and $g_1|_A = g_2|_A$, we obtain that $g_1(a) = g_2(a)$ as $(Y, \mathcal{T}_Y)$ is Hausdorff so points of convergence are unique by Theorem 1.5.40. Thus, as $a \in \overline{A}$ was arbitrary, $g_1 = g_2$. ∎

Of course, for certain compactifications, only certain bounded continuous functions will extend.

Remark 5.4.11. Let S^1, Y, and Z be the compactifications of $(0, 1)$ from Examples 5.4.6, 5.4.7, and 5.4.8 respectively. Let $f : (0, 1) \to \mathbb{R}$ be a bounded continuous function. It is elementary using the embedding $x \mapsto (\cos(2\pi x), \sin(2\pi x))$ of $(0, 1)$ into S^1 to see that f extends to a continuous function on S^1 if and only if $\lim_{x\to 0} f(x)$ and $\lim_{x\to 1} f(x)$ exist and are equal. However, using the canonical embedding of $(0, 1)$ into $[0, 1]$, f extends to a continuous function on $[0, 1]$ if and only if $\lim_{x\to 0} f(x)$ and $\lim_{x\to 1} f(x)$ exist but need not be equal.

It is not difficult to see that if $\lim_{x\to 0} f(x)$ and $\lim_{x\to 1} f(x)$ exist, then f also extends to a continuous function on Z by setting f to be equal to $\lim_{x\to 0} f(x)$ on all points in Z that lie on the y-axis. However, if $f(x) = \sin\left(\frac{\pi}{x}\right)$, then f also extends to a continuous function on Z even though $\lim_{x\to 0} f(x)$ does not exist. To see this, simply note that Z is a subspace of $\mathbb{R}^2$ and the projection map $\pi_2 : \mathbb{R}^2 \to \mathbb{R}$ defined by $\pi_2((x, y)) = y$ is a continuous map so $\pi_2|_Z$ is continuous by Lemma 2.1.18. It is elementary to then see via the embedding of $(0, 1)$ into Z from Example 5.4.8 that $\pi_2|_Z$ extends f as desired.

The topologist's sine curve embedding above yielded a compactification of $(0, 1)$ so that $f(x) = \sin\left(\frac{\pi}{x}\right)$ extended to a continuous function. This was done by placing this function inside one part of a product space to which $(0, 1)$ embedded into. In fact, this is how we can construct a compactification of any Tychonoff space so that every continuous bounded function extends to the compactification.

Theorem 5.4.12. *Let $(X, \mathcal{T})$ be a Tychonoff space. For every $f \in \mathcal{C}_b(X, \mathbb{R})$ let*

$$I_f = [\inf(f(X)), \sup(f(X))],$$

which is a closed interval of $\mathbb{R}$ of finite length. Then, if $\prod_{f\in\mathcal{C}_b(X,\mathbb{R})} I_f$ is equipped with the product topology, there exists a compact Hausdorff subspace Y of $\prod_{f\in\mathcal{C}_b(X,\mathbb{R})} I_f$ that is a compactification of $(X, \mathcal{T})$. Moreover, Y has the property every element of $\mathcal{C}_b(X, \mathbb{R})$ has unique extension to a continuous bounded function on Y.

Proof. Let $\{f_\alpha\}_{\alpha\in I} = \mathcal{C}_b(X, \mathbb{R})$. We claim $\{f_\alpha\}_{\alpha\in I}$ satisfies the assumptions of the Embedding Theorem (Theorem 5.3.19). To see this, let $x_0 \in X$ and

U a neighbourhood of x_0 be arbitrary. Since $(X, \mathcal{T})$ is a Tychonoff space, there exists a function $g \in \mathcal{C}(X, \mathbb{R})$ such that $g(x_0) \neq 0$ and $g(x) = 0$ for all $x \in X \setminus U$. As the only problem is that g is not bounded, we will make g bounded. Let $[a, b]$ be any closed interval of finite length that contains $g(x_0)$ and 0. Define $\varphi : \mathbb{R} \to [a, b]$ by

$$\varphi(x) = \begin{cases} x & \text{if } x \in [a, b] \\ a & \text{if } x < a \\ b & \text{if } x > b \end{cases}.$$

Then is is clear that $f = \varphi \circ g$ is a continuous bounded function on $(X, \mathcal{T})$ such that $f(x_0) \neq 0$ and $f(x) = 0$ for all $x \in X \setminus U$. Hence, as x_0 and U were arbitrary, $\{f_\alpha\}_{\alpha \in I} = \mathcal{C}_b(X, \mathbb{R})$ satisfy the assumptions of the Embedding Theorem (Theorem 5.3.19).

By the Embedding Theorem (Theorem 5.3.19), if $\prod_{f \in \mathcal{C}_b(X,\mathbb{R})} \mathbb{R}$ is equipped with the product topology, the function $F : X \to \prod_{f \in \mathcal{C}_b(X,\mathbb{R})} \mathbb{R}$ defined by

$$F(x) = (f(x))_{f \in \mathcal{C}_b(X,\mathbb{R})}$$

is an embedding of $(X, \mathcal{T})$ into $\prod_{f \in \mathcal{C}_b(X,\mathbb{R})} \mathbb{R}$. As clearly the range of F is $\prod_{f \in \mathcal{C}_b(X,\mathbb{R})} I_f$, we can view F as an embedding of $(X, \mathcal{T})$ into $\prod_{f \in \mathcal{C}_b(X,\mathbb{R})} I_f$.

Since I_f is a compact subset of the Hausdorff space $\mathbb{R}$ by the Heine-Borel Theorem, Tychonoff's Theorem (Theorem 3.3.4) implies that $\prod_{f \in \mathcal{C}_b(X,\mathbb{R})} I_f$ is a compact Hausdorff topological space. Therefore, if $Y = \overline{F(X)}$, then Y is a closed subset of a compact Hausdorff topological space and thus a compact Hausdorff topological space by Theorem 3.1.14. Hence Y is a compactification of $(X, \mathcal{T})$ by construction.

To complete the proof, let $f_0 \in \mathcal{C}_b(X, \mathbb{R})$ be arbitrary. To see that f_0 extends to a continuous function on Y, let $\pi_{f_0} : \prod_{f \in \mathcal{C}_b(X,\mathbb{R})} I_f \to I_{f_0}$ be the projection map defined in Example 2.1.8. Clearly $g = \pi_{f_0}|_Y$ is a continuous bounded function on Y. Furthermore, for all $x \in X$ we see that

$$g(F(x)) = \pi_{f_0}\left((f(x))_{f \in \mathcal{C}_b(X,\mathbb{R})}\right) = f_0(x)$$

thereby showing that g is a continuous extension of f_0 up to identifying $(X, \mathcal{T})$ with its homeomorphic image $F(X)$ in Y. Hence the proof is complete as Lemma 5.4.10 shows that any continuous extension to a compactification must be unique. ∎

Due to the use of the compactification in Theorem 5.4.12, we give it a name.

Definition 5.4.13. Let $(X, \mathcal{T})$ be a Tychonoff space. The *Stone-Čech compactification of* $(X, \mathcal{T})$ is the compactification of $(X, \mathcal{T})$ produced in Theorem 5.4.12. The Stone-Čech compactification of $(X, \mathcal{T})$ is denoted by $\beta(X)$.

Of course, we only know that the Stone-Čech compactification of a Tychonoff space enables us to extend real-valued bounded continuous functions. In fact, we can extend far more.

Lemma 5.4.14. *Let $(X, \mathcal{T}_X)$ be a Tychonoff space and let $(Y, \mathcal{T}_Y)$ be a compactification of $(X, \mathcal{T}_X)$ such that every element of $\mathcal{C}_b(X, \mathbb{R})$ has unique extension to a continuous bounded function on $(Y, \mathcal{T}_Y)$ (e.g. $Y = \beta(X)$). If $(Z, \mathcal{T}_Z)$ is a compact Hausdorff topological space and $f : X \to Z$ is continuous, then f extends uniquely to a continuous function $g : Y \to Z$.*

Proof. Let $(Z, \mathcal{T}_Z)$ be a compact Hausdorff topological space. Therefore $(Z, \mathcal{T}_Z)$ is a Tychonoff space (by Theorem 5.4.2 for example). Hence Theorem 5.3.20 implies there exists a non-empty index set I and an embedding $\Phi : Z \to \prod_{\alpha \in I}[0,1]$ where $\prod_{\alpha \in I}[0,1]$ is equipped with the product topology. Consequently $\Phi(Z)$ is a compact subset of $\prod_{\alpha \in I}[0,1]$ by Theorem 3.1.27 and thus a closed subset of $\prod_{\alpha \in I}[0,1]$ by Theorem 3.1.13 as $\prod_{\alpha \in I}[0,1]$ is Hausdorff.

For each $\alpha \in I$, let $\pi_\alpha : \prod_{\alpha \in I}[0,1] \to [0,1]$ be the continuous projection map from Example 2.1.8. Therefore, for all $\alpha \in I$ the function $f_\alpha = \pi_\alpha \circ \Phi \circ f : X \to [0,1]$ is a composition of continuous functions and thus a continuous bounded function. Hence the assumptions on $(Y, \mathcal{T}_Y)$ implies there exists a unique function $g_\alpha \in \mathcal{C}_b(Y, \mathbb{R})$ such that $g_\alpha|_X = f_\alpha$.

Define $g_0 : Y \to \prod_{\alpha \in I} \mathbb{R}$ by

$$g_0(y) = (g_\alpha(y))_{\alpha \in I}$$

for all $y \in Y$. As $f_\alpha = \pi_\alpha \circ \Phi \circ f$ and as $g_\alpha|_X = f_\alpha$, we see that $g_0(x) = \Phi(f(x))$ for all $x \in X$. Moreover, since g_0 is continuous, we see that

$$g_0(Y) = g_0\left(\overline{X}\right) \subseteq \overline{g_0(X)} = \overline{\Phi(f(X))} \subseteq \overline{\Phi(Z)} = \Phi(Z).$$

Hence $g_0 : Y \to \Phi(Z)$. Therefore, if we define $g = \Phi^{-1} \circ g : Y \to Z$, then g is a continuous function such that $g(x) = f(x)$ for all $x \in X$. Hence the proof is complete as Lemma 5.4.10 shows that any continuous extension to a compactification of a function valued in a Hausdorff space must be unique. ∎

In fact, not only does Lemma 5.4.14 extend the known properties of the Stone-Čech compactification, it enables us to demonstrate the Stone-Čech compactification is the unique topological space with this property.

Theorem 5.4.15. *Let $(X, \mathcal{T}_X)$ be a Tychonoff space and let $(Y, \mathcal{T}_Y)$ be a compactification of $(X, \mathcal{T}_X)$ such that every element of $\mathcal{C}_b(X, \mathbb{R})$ has unique extension to a continuous bounded function on $(Y, \mathcal{T}_Y)$. Then $(Y, \mathcal{T}_Y)$ is homeomorphic to the Stone-Čech compactification of $(X, \mathcal{T})$.*

Proof. To simplify discussions and avoid a lot of homeomorphisms floating around, we may assume without loss of generality that $X \subseteq Y$ and $X \subseteq \beta(X)$.

Consider the inclusion map $i_1 : X \to Y$ and $i_2 : X \to \beta(X)$. Since $(Y, \mathcal{T}_Y)$ and $\beta(X)$ are compact Hausdorff topological spaces, Lemma 5.4.14 implies there exists unique continuous functions $f_1 : \beta(X) \to Y$ and $f_2 : Y \to \beta(X)$ such that $f_1(x) = x = f_2(x)$ for all $x \in X$. Consequently, if $g_1 = f_1 \circ f_2 : Y \to Y$ and $g_2 = f_2 \circ f_1 : \beta(X) \to \beta(X)$, then g_1 and g_2 are continuous functions such that $g_1(x) = x = g_2(x)$ for all $x \in X$. Therefore, since $Y = \overline{X}$ and $\beta(X) = \overline{X}$, continuity of g_1 and g_2 implied that $g_1(y) = y$ for all $y \in Y$ and $g_2(z) = z$ for all $z \in \beta(X)$ by Theorem 2.1.9. Hence f_1 and f_2 are continuous functions which are inverses of each other and thus homeomorphisms. Hence $(Y, \mathcal{T}_Y)$ is homeomorphic to $\beta(X)$ as desired. ∎

Unfortunately, even the simplest Stone-Čech compactifications are difficult (read as near impossible) to describe in a simpler way than what was done in Theorem 5.4.12. However, the case of $\mathbb{N}$ equipped with the discrete topology is of particular importance. The description of this space (with proofs) is given in Appendix C.

5.5 Tietz Extension Theorem

Of course, while we are on the topic of extending continuous functions to compactifications, it is also useful to discuss when we can extend a continuous function on a closed subset of a topological space to the entire space. In particular, the following show that if we are in a normal topological space, then we can always do such an extension for bounded continuous functions.

Theorem 5.5.1 (Tietz's Extension Theorem - Bounded Version). *Let $(X, \mathcal{T})$ be a normal topological space (e.g. a T_4 space), let F be a closed subspace of $(X, \mathcal{T})$, and let $f \in \mathcal{C}_b(F, \mathbb{R})$ be continuous. There exists a $g \in \mathcal{C}_b(X, \mathbb{R})$ such that $g|_F = f$ and $\|g\|_\infty = \|f\|_\infty$.*

Proof. Since $f \in \mathcal{C}_b(F, \mathbb{R})$, we know that

$$\|f\|_\infty = \sup(\{|f(x)| \mid x \in F\}) < \infty.$$

Clearly if $\|f\|_\infty = 0$, we can take g to be the zero function thereby completing the claim. Hence we may assume that $\|f\|_\infty > 0$. Therefore, by scaling f if necessary, we can assume without loss of generality that $\|f\|_\infty = 1$. Thus $|f(x)| \leq 1$ for all $x \in F$.

To proceed with the proof, our goal is to Urysohn's Lemma (Theorem 5.2.1) to get several elements of $\mathcal{C}_b(X, \mathbb{R})$. We will construct these functions in a specific way so that their sum closer and closer approximates f on F. We will then take a limit of these functions to obtain the desired extension of f.

To begin, let

$$A_1 = \left\{x \in F \,\middle|\, f(x) \in \left[-1, -\frac{1}{3}\right]\right\} = f^{-1}\left(\left[-1, -\frac{1}{3}\right]\right) \text{ and}$$
$$B_1 = \left\{x \in F \,\middle|\, f(x) \in \left[\frac{1}{3}, 1\right]\right\} = f^{-1}\left(\left[\frac{1}{3}, 1\right]\right).$$

Therefore, since f is continuous on F, A_1 and B_1 disjoint closed subsets of F. Therefore, since F is closed in $(X, \mathcal{T})$ and since closed subsets of F are the intersection of F with a closed subset of $(X, \mathcal{T})$ and therefore closed in $(X, \mathcal{T})$, we see that A_1 and B_1 disjoint closed subsets of $(X, \mathcal{T})$. Since $(X, \mathcal{T})$ is normal, Urysohn's Lemma implies there exists a continuous function $h_1 : X \to \left[-\frac{1}{3}, \frac{1}{3}\right]$ such that $h_1(a) = -\frac{1}{3}$ for all $a \in A_1$ and $h_1(b) = \frac{1}{3}$ for all $b \in B_1$.

We claim that $|f(x) - h_1(x)| \leq \frac{2}{3}$ for all $x \in F$. To see this, notice if $x \in A_1$ then $-1 \leq f(x) \leq -\frac{1}{3}$ so the fact that $h_1(x) = -\frac{1}{3}$ as $x \in A_1$ implies $|f(x) - h_1(x)| \leq \frac{2}{3}$. Similarly, if $x \in B_1$ then $\frac{1}{3} \leq f(x) \leq 1$ so the fact that $h_1(x) = \frac{1}{3}$ as $x \in B_1$ implies $|f(x) - h_1(x)| \leq \frac{2}{3}$. Finally, if $x \in F \setminus (A_1 \cup B_1)$, then the definitions of A_1 and B_1 imply that $-\frac{1}{3} < f(x) < \frac{1}{3}$ so, as $h_1 : X \to \left[-\frac{1}{3}, \frac{1}{3}\right]$, we obtain that $|f(x) - h_1(x)| \leq \frac{2}{3}$. Hence $|f(x) - h_1(x)| \leq \frac{2}{3}$ for all $x \in F$.

Let $\alpha = \frac{2}{3}$. We claim that there exists a sequence $(h_n)_{n\geq 1}$ in $\mathcal{C}_b(X, \mathbb{R})$ such that

$$\|h_n\|_\infty \leq \frac{1}{3}\alpha^{n-1} \qquad \text{and} \qquad \left|f(x) - \sum_{k=1}^{n} h_k(x)\right| \leq \alpha^n \text{ for all } x \in F$$

for all $n \in \mathbb{N}$. To see this, we proceed by induction on n with the base case $n = 1$ completed by the above arguments. Thus, to proceed with the inductive step, suppose there exist $(h_k)_{k=1}^n$ in $\mathcal{C}_b(X, \mathbb{R})$ such that

$$\|h_m\|_\infty \leq \frac{1}{3}\alpha^{m-1} \qquad \text{and} \qquad \left|f(x) - \sum_{k=1}^{m} h_k(x)\right| \leq \alpha^m \text{ for all } x \in F$$

for all $m \in \{1, \ldots, n\}$.

$$A_{n+1} = \left\{x \in F \,\middle|\, f(x) - \sum_{k=1}^{n} h_k(x) \in \left[-\alpha^n, -\frac{1}{3}\alpha^n\right]\right\} \text{ and}$$
$$B_{n+1} = \left\{x \in F \,\middle|\, f(x) - \sum_{k=1}^{n} h_k(x) \in \left[\frac{1}{3}\alpha^n, \alpha^n\right]\right\}.$$

Since $h_k \in \mathcal{C}_b(X, \mathbb{R})$ for all $k \in \{1, \ldots, n\}$ and since F is a closed subspace of $(X, \mathcal{T})$, we see that $x \mapsto f(x) - \sum_{k=1}^{n} h_k(x)$ is a continuous function on F and thus A_{n+1} and B_{n+1} are disjoint closed subsets of F. Therefore, since F is closed in $(X, \mathcal{T})$ and since closed subsets of F are the intersection of F with a closed subset of $(X, \mathcal{T})$ and therefore closed in $(X, \mathcal{T})$, we see that A_{n+1} and B_{n+1} disjoint closed subsets of $(X, \mathcal{T})$. Thus, since $(X, \mathcal{T})$ is normal, Urysohn's Lemma implies there exists a continuous function $h_{n+1} : X \to \left[-\frac{1}{3}\alpha^n, \frac{1}{3}\alpha^n\right]$ such that $h_{n+1}(a) = -\frac{1}{3}\alpha^n$ for all $a \in A_{n+1}$ and $h_{n+1}(b) = \frac{1}{3}\alpha^n$ for all $b \in B_{n+1}$.

Clearly $\|h_{n+1}\|_\infty \leq \frac{1}{3}\alpha^n$ by construction. To see that

$$\left| f(x) - \sum_{k=1}^{n+1} h_k(x) \right| \leq \alpha^{n+1}$$

for all $x \in F$, we will proceed as we did in the $n = 1$ case. Indeed, if $x \in A_{n+1}$ then

$$h_{n+1}(x) = -\frac{1}{3}\alpha^n \qquad \text{and} \qquad f(x) - \sum_{k=1}^{n} h_k(x) \in \left[-\alpha^n, -\frac{1}{3}\alpha^n\right]$$

so

$$\left| f(x) - \sum_{k=1}^{n+1} h_k(x) \right| \leq \alpha^n - \frac{1}{3}\alpha^n = \frac{2}{3}\alpha^n = \alpha^{n+1}.$$

Similarly, if $x \in B_{n+1}$ then

$$h_{n+1}(x) = \frac{1}{3}\alpha^n \qquad \text{and} \qquad f(x) - \sum_{k=1}^{n} h_k(x) \in \left[\frac{1}{3}\alpha^n, \alpha^n\right]$$

so

$$\left| f(x) - \sum_{k=1}^{n+1} h_k(x) \right| \leq \alpha^n - \frac{1}{3}\alpha^n = \frac{2}{3}\alpha^n = \alpha^{n+1}.$$

Finally, if $x \in F \setminus (A_{n+1} \cup B_{n+1})$, then the definitions of A_{n+1} and B_{n+1} imply that

$$\left| f(x) - \sum_{k=1}^{n} h_k(x) \right| < \frac{1}{3}\alpha^n \qquad \text{and} \qquad |h_{n+1}(x)| \leq \frac{1}{3}\alpha^n$$

so

$$\left| f(x) - \sum_{k=1}^{n+1} h_k(x) \right| \leq \frac{1}{3}\alpha^n + \frac{1}{3}\alpha^n = \frac{2}{3}\alpha^n = \alpha^{n+1}.$$

Hence

$$\left| f(x) - \sum_{k=1}^{n+1} h_k(x) \right| \leq \alpha^{n+1}$$

for all $x \in F$. Therefore, the inductive step is complete so there exist $(h_n)_{n\geq 1}$ in $\mathcal{C}_b(X, \mathbb{R})$ with the desired properties.

Of course, by construction we know that

$$\sum_{n=1}^{\infty} \|h_n\|_\infty \leq \sum_{n=1}^{\infty} \frac{1}{3}\alpha^{n-1} = \frac{1}{3}\frac{1}{1-\alpha} = 1 < \infty.$$

Hence $\sum_{n=1}^{\infty} h_n$ is an absolutely summable series in $(\mathcal{C}_b(X,\mathbb{R}), \|\cdot\|_\infty)$. Therefore, since $(\mathcal{C}_b(X,\mathbb{R}), \|\cdot\|_\infty)$ is a Banach space by Theorem 4.2.14, Theorem 4.1.17 implies that $\sum_{n=1}^{\infty} h_n$ is summable. Hence

$$g = \sum_{n=1}^{\infty} h_n$$

is a well-defined element of $\mathcal{C}_b(X,\mathbb{R})$. We claim that g is the function we seek.

To begin to see that g has the desired properties, we note since the norm is a continuous function on any normed linear space that

$$\begin{aligned} \|g\|_\infty &= \lim_{n\to\infty} \left\| \sum_{k=1}^{n} h_k \right\|_\infty \\ &\leq \limsup_{n\to\infty} \sum_{k=1}^{n} \|h_k\|_\infty \\ &\leq \limsup_{n\to\infty} \sum_{k=1}^{n} \frac{1}{3}\alpha^{k-1} \\ &= \frac{1}{3}\frac{1}{1-\alpha} = 1 \end{aligned}$$

so $\|g\|_\infty \leq 1$. Hence, provided we can show that $g|_F = f$, we will then be able to use the fact that $\|f\|_\infty = 1$ to obtain that $\|g\|_\infty = 1 = \|f\|_\infty$ as desired. Hence, all that remains to be shown is that $g|_F = f$.

To see that $g|_F = f$, let $x \in F$ be arbitrary. Then the definition of g implies that

$$|f(x) - g(x)| = \lim_{n\to\infty} \left| f(x) - \sum_{k=1}^{n} h_k(x) \right| \leq \limsup_{n\to\infty} \alpha^n = 0$$

as $\alpha = \frac{2}{3}$. Hence $g(x) = f(x)$. Therefore, since $x \in F$ was arbitrary, $g|_F = f$ as desired. ∎

With Theorem 5.5.1, it is not too difficult to extend these results to unbounded functions.

Theorem 5.5.2 (Tietz's Extension Theorem - Unbounded Version). *Let $(X,\mathcal{T})$ be a normal topological space (e.g. a T_4 space), let F be a closed subspace of $(X,\mathcal{T})$, and let $f : F \to \mathbb{R}$ be continuous. There exists a continuous function $g : X \to \mathbb{R}$ such that $g|_F = f$.*

Proof. Our goal in this proof is to use a homeomorphism to reduce the result to the bounded case studied in Theorem 5.5.1. Indeed consider the function $\varphi : \mathbb{R} \to (-1, 1)$ defined by

$$\varphi(x) = \frac{x}{1 + |x|}$$

for all $x \in \mathbb{R}$. It is elementary to see that φ is a homeomorphism with inverse $\varphi^{-1} : (-1, 1) \to \mathbb{R}$ defined by

$$\varphi^{-1}(y) = \frac{y}{1 - |y|}$$

for all $y \in (-1, 1)$. Hence, if $f_0 : F \to (-1, 1)$ is defined by $f_0 = \varphi \circ f$, then $f_0 \in \mathcal{C}_b(F, \mathbb{R})$ is such that $\|f_0\|_\infty \leq 1$. Hence the bounded version of Tietz Extension Theorem (Theorem 5.5.1) implies there exists an $h_0 \in \mathcal{C}_b(X, \mathbb{R})$ such that $\|h_0\|_\infty = \|f_0\|_\infty \leq 1$ and $h_0|_F = f_0$.

Of course, if $h_0(x) \neq \pm 1$ for all $x \in X$, then one can immediately take $g = \varphi^{-1} \circ h_0$ thereby completing the proof. Therefore, as we only know that $\|h_0\|_\infty \leq 1$ so it is possible that $h_0(x) = \pm 1$ for some $x \in X$, we must correct h_0.

Let $C = h_0^{-1}(\{-1, 1\})$. Since $h_0 \in \mathcal{C}_b(X, \mathbb{R})$, C is a closed (possibly empty subset) of X. We claim that $C \cap F = \emptyset$. To see this, notice if $x \in F$ then $h_0(x) = f_0(x) \in (-1, 1)$ so $x \notin C$ by definition. Hence $C \cap F = \emptyset$.

Since $(X, \mathcal{T})$ is normal and C and F are pairwise disjoint closed subsets of $(X, \mathcal{T})$, Urysohn's Lemma (Theorem 5.2.1) implies there exists a continuous function $h : X \to [0, 1]$ such that $h(x) = 0$ for all $x \in C$ and $h(x) = 1$ for all $x \in \mathcal{F}$. Define $g_0 : X \to \mathbb{R}$ by

$$g_0(x) = h_0(x)h(x)$$

for all $x \in X$. Since g_0 is a product of elements of $\mathcal{C}_b(X, \mathbb{R})$, it is elementary to see that $g_0 \in \mathcal{C}_b(X, \mathbb{R})$. Furthermore, we claim that $g(X) \subseteq (-1, 1)$. To see this, notice if $x \in C$ then $|h_0(x)| = 1$ and $h(x) = 0$ so $g_0(x) = 0 \in (-1, 1)$. Furthermore, if $x \in X \setminus C$ then $|h_0(x)| < 1$ and $h(x) \in [0, 1]$ so $g_0(x) \in (-1, 1)$. Hence $g(X) \subseteq (-1, 1)$ as claimed.

Define $g : X \to \mathbb{R}$ by

$$g(x) = \varphi^{-1}(g_0(x))$$

for all $x \in X$, which is well-defined as $g(X) \subseteq (-1, 1)$. Furthermore $g \in \mathcal{C}(X, \mathbb{R})$ as g is the composition of two continuous functions and thus continuous. Finally, to see that $g|_F = f$, let $x \in F$ be arbitrary. Then

$$g(x) = \varphi^{-1}(g_0(x)) = \varphi^{-1}(h_0(x)h(x)) = \varphi^{-1}(f_0(x)1) = \varphi^{-1}(f_0(x)) = f(x)$$

as desired. ∎

However, in some cases, one wants to have complex value functions and be sure that the same bound is retained. The following corollary does the trick.

Corollary 5.5.3. *Let $(X, \mathcal{T})$ be a normal topological space (e.g. a T_4 space), let F be a closed subspace of $(X, \mathcal{T})$, and let $f : F \to B$ be continuous where B is a closed ball centred at the origin in $(\mathbb{K}^n, \|\cdot\|_2)$. There exists a continuous function $g : X \to B$ such that $g|_F = f$.*

Proof. Since $(\mathbb{C}^n, \|\cdot\|_2)$ is isomorphic as a normed linear space to $(\mathbb{R}^{2n}, \|\cdot\|_2)$ (i.e. there is a linear bijection from $\mathbb{C}^n$ to $\mathbb{R}^{2n}$ that preserves the norm), it suffices to consider the case $\mathbb{K} = \mathbb{R}$. Furthermore, by scaling, we may assume without loss of generality that $B = B_2[\vec{0}, 1]$.

Since $f : F \to B$ is continuous, there exists $f_1, f_2, \ldots, f_n \in \mathcal{C}_b(F, \mathbb{R})$ such that

$$f(x) = (f_1(x), f_2(x), \ldots, f_n(x))$$

for all $x \in F$. Hence Tietz Extension Theorem (Theorem 5.5.2) implies there $g_1, g_2, \ldots, g_n \in \mathcal{C}_b(F, \mathbb{R})$ such that $g_k|_F = f_k$ for all $k \in \{1, \ldots, n\}$. Define $h : X \to \mathbb{R}^n$ by

$$h(x) = (g_1(x), g_2(x), \ldots, g_n(x))$$

for all $x \in X$. Clearly h is continuous by Theorem 2.1.10 and $h|_F = f$ by construction. However, it by the above construction, it is not necessarily the case that $h(x) \in B$ for all $x \in X$.

Define $\varphi : \mathbb{R}^n \to B$ by

$$\varphi(\vec{v}) = \begin{cases} \vec{v} & \text{if } \|\vec{v}\|_2 \leq 1 \\ \frac{1}{\|\vec{v}\|_2}\vec{v} & \text{if } \|\vec{v}\|_2 > 1 \end{cases}$$

for all $\vec{v} \in \mathbb{R}^n$. As $x \mapsto \frac{1}{x}$ is a continuous function on $[1, \infty)$ and since the norm in any normed linear space is continuous, it is elementary to verify that φ is a continuous function that maps into B. Furthermore, it is clear that if $\vec{v} \in B$, then $\varphi(\vec{v}) = \vec{v}$.

Define $g : X \to B$ by $g = \varphi \circ h$. Hence g is a continuous function that maps into B. Thus, it remains only to show that $g|_F = f$. To see this, notice for all $x \in F$ that

$$g(x) = \varphi(h(x)) = \varphi(f(x)) = f(x)$$

as $f(x) \in B$ for all $x \in F$. Hence the claim and proof is complete. ∎

Of course, at this stage it is natural to ask whether the assumption that our topological spaces are normal in the Tietz Extension Theorem is actually required or whether we can weaken the assumption. It turns out we can not.

Proposition 5.5.4. *Let $(X, \mathcal{T})$ be a topological space such that whenever F is a closed subspace of $(X, \mathcal{T})$ and $f \in \mathcal{C}_b(F, \mathbb{R})$ there exists a $g \in \mathcal{C}(X, \mathbb{R})$ such that $g|_F = f$. Then $(X, \mathcal{T})$ is normal.*

Proof. To see that $(X, \mathcal{T})$ is normal, let A and B be non-empty closed subsets of $(X, \mathcal{T})$ such that $A \cap B = \emptyset$. Let $F = A \cup B$ which is a closed subset of $(X, \mathcal{T})$ since $X \setminus F = (X \setminus A) \cap (X \setminus B)$ which is open in $(X, \mathcal{T})$. Define $f : F \to [0, 1]$ by

$$f(x) = \begin{cases} 0 & \text{if } x \in A \\ 1 & \text{if } x \in B \end{cases}.$$

Hence f is clearly continuous since the inverse image of any closed subset of $[0, 1]$ is either $\emptyset$, F, A, or B, all of which are closed subsets of F. Hence the assumptions imply there exists a $g \in \mathcal{C}(X, \mathbb{R})$ such that $g|_F = f$. Let

$$U = g^{-1}\left(\left(-\infty, \frac{1}{4}\right)\right) \qquad \text{and} \qquad U = g^{-1}\left(\left(\frac{3}{4}, \infty\right)\right).$$

Hence, by construction U and V are open subsets of $(X, \mathcal{T})$ such that $A \subseteq U$, $B \subseteq V$, and $U \cap V = \emptyset$. Therefore, as A and B were arbitrary, $(X, \mathcal{T})$ is normal as desired. ∎

Chapter 6

Metrizations

As Chapter 5 demonstrates, normal topological spaces are among the nicest topological spaces due to the plethora of results that hold. In particular, Urysohn's Lemma (Theorem 5.2.1) and Tietz's Extension Theorem (Theorem 5.5.1) hold in every normal topological spaces and T_4 topological spaces embed into the product of compact Hausdorff topological spaces and have Stone-Čech Compactifications (Theorem 5.4.12). Therefore, as every metric spaces is a T_4 space, all of these results hold for metric spaces.

However, metric spaces are particularly nice. For example, Theorem 1.5.28 implies sequences are enough to determine the topology and thus many topological properties in metric spaces. Moreover, it is very easy to prove Urysohn's Lemma (Theorem 5.2.1) for metric spaces and, in fact, one has an explicit formula for the function f in Theorem 5.2.1, that is

$$f(x) = \frac{\mathrm{dist}(x, A)}{\mathrm{dist}(x, A) + \mathrm{dist}(x, B)}.$$

Consequently, as metric spaces are some of the nicest T_4 spaces, it is natural to ask, "Which topologies are induced by metrics?" Thus the focus of this chapter is to answer this question in some general settings.

To begin this chapter, we will focus on the simplest requirement for a topology to be induced by a metric; having a countable neighbourhood basis of each point. Countability also plays a major role in the Baire Category Theorem (Theorem 6.2.10 and Theorem 6.2.12) which holds for compact Hausdorff topological spaces and complete metric spaces. From there we will generalize the idea of the Embedding Theorem (Theorem 5.3.19) to develop Urysohn's Metrization Theorem (Theorem 6.3.1) which provides the simplest method of checking certain topologies are induced by metrics. Finally, we will complete this chapter with two characterizations of when a topology is induced by a metric.

6.1 The Countability Axioms

To begin our study of determining when a topology is induced by a metric, we begin with the simplest requirement; that every point has a countable neighbourhood basis. Of course this is a requirement of a topology being induced by a metric since $\left\{B_d\left(x, \frac{1}{n}\right)\right\}_{n=1}^{\infty}$ is clearly a countable neighbourhood basis of a point x in a metric space (X, d). Before we begin this study, we first desire to simplify our terminology in saying that "$\mathcal{T}$ is induced by a metric".

Definition 6.1.1. A topological space $(X, \mathcal{T})$ is said to be *metrizable* if there exists a metric $d : X \times X \to [0, \infty)$ such that $\mathcal{T}$ is the topology induced by d.

Remark 6.1.2. It is necessary to point out that there is a difference between a metrizable topological spaces and metric spaces; that is, a metric space is a topological space with a fixed metric whereas a metrizable topological space where there is a metric that induces the topology. This distinction is not when it comes to topological properties, but when it comes to properties that are characterized by the metric. For example, if $\mathcal{T}$ is the canonical topology on $\mathbb{R}$, then $\mathcal{T}$ is induced by the metrics $d_1(x, y) = |x - y|$ and $d_2(x, y) = |e^{-x} - e^{-y}|$. Thus $(\mathbb{R}, \mathcal{T})$ is metrizable. However $(\mathbb{R}, d_1)$ and $(\mathbb{R}, d_2)$ are very different metric spaces as, for example, we know that $(\mathbb{R}, d_1)$ is complete but $(\mathbb{R}, d_2)$ is not complete by Example 4.1.10.

Remark 6.1.3. As every metric space is a T_4 space, every metrizable topological space must also be a T_4 space. Hence we immediately obtain that $\prod_{\alpha \in \mathbb{R}} \mathbb{R}$ is not metrizable as it is not normal by Example 5.1.29.

Clearly any subspace of a metrizable spaces is metrizable by the definition of the subspace topology and by restricting the metric. Thus, for a normal topological space to be metrizable, every subspace must also be normal. Furthermore, some products of metrizable spaces are metrizable.

Lemma 6.1.4. *The product topology on a countable product of metrizable spaces is metrizable.*

Proof. Let I be a countable set and for each $\alpha \in I$, let $(X_\alpha, \mathcal{T}_\alpha)$ be a metrizable topological space. Hence for each $\alpha \in I$ there exists a metric $d_\alpha : X_\alpha \times X_\alpha \to \mathbb{R}$ such that the metric topology on (X_α, d_α) is exactly T_α. Since Lemma 4.3.4 implies there exists a metric on $\prod_{\alpha \in I} X_\alpha$ that induces the product topology, the proof is complete. ∎

To be sure our terminology from the introduction of this chapter and section are clear, we define the following.

Definition 6.1.5. A topological space $(X, \mathcal{T})$ is said to have a *countable neighbourhood basis at a point* $x \in X$ if there exists a collection $\{B_n\}_{n=1}^\infty$ of neighbourhoods of x such that if U is a neighbourhood of x, then there exists an $N \in \mathbb{N}$ such that $B_N \subseteq U$.

Definition 6.1.6. A topological space $(X, \mathcal{T})$ is said to be *first countable* if every point in X has a countable neighbourhood basis.

Example 6.1.7. Clearly every metric space (X, d) is first countable since given $x \in X$, $\left\{B_d\left(x, \frac{1}{n}\right)\right\}_{n=1}^\infty$ is clearly a countable neighbourhood basis of x. Consequently, every metrizable topological space must be first countable whereas every topological space that is not first countable cannot be metrizable.

Example 6.1.8. Let $\mathcal{T}_L$ be the lower limit topology on $\mathbb{R}$. Then $(\mathbb{R}, \mathcal{T}_L)$ is first countable. To see this, let $x \in \mathbb{R}$ be arbitrary. Consider

$$\mathcal{B}_x = \left\{ \left[x, x + \frac{1}{n}\right) \,\middle|\, n \in \mathbb{N}\right\}.$$

Clearly $\mathcal{B}_x \subseteq \mathcal{T}_L$ and $\mathcal{B}_x$ is countable. Moreover, if $B = [a, b) \in \mathcal{B}$ is such that $x \in [a, b)$, then $a \leq x$ and $x < b$ so there exists an $n \in \mathbb{N}$ such that $x + \frac{1}{n} < b$. Hence

$$x \in \left[x, x + \frac{1}{n}\right) \subseteq [a, b).$$

Therefore, as $\left[x, x + \frac{1}{n}\right) \in \mathcal{B}_x$ and $B \in \mathcal{B}$ was arbitrary, $\mathcal{B}_x$ is a countable neighbourhood basis for x. Therefore, as $x \in \mathbb{R}$ was arbitrary, $(\mathbb{R}, \mathcal{T}_L)$ is first countable.

Of course, metrizable spaces are particularly nice because, as with metric spaces, sequences characterize the topology (see Theorem 1.5.28) and other nice topological properties. To see these nice properties, we first establish the following useful lemma.

Lemma 6.1.9. *Let $(X, \mathcal{T})$ be a topological space, let $x_0 \in X$, and let $\{U_n\}_{n=1}^\infty$ be a countable neighbourhood basis of x_0. Then there exists a countable neighbourhood basis $\{V_n\}_{n=1}^\infty$ of x_0 such that $V_{n+1} \subseteq V_n$ for all $n \in \mathbb{N}$.*

Proof. For each $n \in \mathbb{N}$, let

$$V_n = \bigcap_{k=1}^n U_k.$$

Clearly V_n is a neighbourhood of x_0 and $V_{n+1} \subseteq V_n$ for all $n \in \mathbb{N}$. To see that $\{V_n\}_{n=1}^\infty$ is a neighbourhood basis for x_0, let U be an arbitrary neighbourhood of x_0. Since $\{U_n\}_{n=1}^\infty$ is a neighbourhood basis of x_0, there exists an $m \in \mathbb{N}$ such that $x_0 \in U_m \subseteq U$. Hence $x_0 \in V_m \subseteq U_m \subseteq U$. Therefore, as U was arbitrary, $\{V_n\}_{n=1}^\infty$ has the desired properties. ∎

Theorem 6.1.10. *Let $(X, \mathcal{T})$ be a first countable topological space. Then the following hold:*

(1) If $A \subseteq X$, then $x_0 \in \overline{A}$ if and only if there exists a sequence $(a_n)_{n\geq 1}$ of elements of A that converge to x.

(2) If $(Y, \mathcal{T}_Y)$ is a topological space, then $f : X \to Y$ is continuous if and only if for every sequence $(x_n)_{n\geq 1}$ in $(X, \mathcal{T})$ that converges to a point $x_0 \in X$, the sequence $(f(x_n))_{n\geq 1}$ in $(Y, \mathcal{T}_Y)$ converges to $f(x_0)$.

Proof. To see (1), let $A \subseteq X$ be arbitrary. If there exists a sequence $(a_n)_{n\geq 1}$ of elements of A that converge to x_0, then $x_0 \in \overline{A}$ by Theorem 1.6.21 as sequences are nets. Conversely, let $x_0 \in \overline{A}$ be arbitrary. As $(X, \mathcal{T})$ is first countable, Lemma 6.1.9 implies there exists a countable neighbourhood basis $\{U_n\}_{n=1}^{\infty}$ of x_0 such that $U_{n+1} \subseteq U_n$ for all $n \in \mathbb{N}$. As $x_0 \in \overline{A}$, Theorem 1.6.21 implies there exists an $a_n \in U_n \cap A$ for all $n \in \mathbb{N}$. We claim that $(a_n)_{n\geq 1}$ converges to x_0 thereby completing the proof. To see this, let U be an arbitrary neighbourhood of x_0. Thus as $\{U_n\}_{n=1}^{\infty}$ is a neighbourhood basis of x_0, there exists an $N \in \mathbb{N}$ such that $U_N \subseteq U$. Hence for all $n \geq N$ we have that

$$a_n \in U_n \subseteq U_N \subseteq U.$$

Therefore, as U was arbitrary, $(a_n)_{n\geq 1}$ converges to x_0 as desired.

To see (2), let $f : X \to Y$ be arbitrary. If f is continuous and $(x_n)_{n\geq 1}$ is a sequence in $(X, \mathcal{T})$ that converges to a point $x_0 \in X$, then $(f(x_n))_{n\geq 1}$ converges to $f(x_0)$ in $(Y, \mathcal{T}_Y)$ by Theorem 2.1.9 as sequences are nets. Conversely, suppose that whenever $(x_n)_{n\geq 1}$ is a sequence in $(X, \mathcal{T})$ that converges to $x_0 \in X$, the sequence $(f(x_n))_{n\geq 1}$ in $(Y, \mathcal{T}_Y)$ converges to $f(x_0)$. To verify that f is continuous, it suffices by Theorem 2.1.9 to show that if $A \subseteq X$ then $f\left(\overline{A}\right) \subseteq \overline{f(A)}$. To proceed in this direction, let $A \subseteq X$ and let $x_0 \in \overline{A}$ be arbitrary. As $x_0 \in \overline{A}$ part (1) of this proof implies there exists a sequence $(a_n)_{n\geq 1}$ of points in A that converges to x_0. Therefore, by the assumptions of this direction of the proof, $(f(x_n))_{n\geq 1}$ is a sequence of points in $f(A)$ that converges to $f(x_0)$ in $(Y, \mathcal{T}_Y)$. Hence as sequences are nets, Theorem 1.6.21 implies that $f(x_0) \in \overline{f(A)}$. Therefore, as $x_0 \in \overline{A}$ was arbitrary, $f\left(\overline{A}\right) \subseteq \overline{f(A)}$ as desired. ∎

Often it is much easier to use Theorem 6.1.10 to demonstrate that a topological space is not first countable (and thus not metrizable) than it is to explicitly demonstrate that a point does not have a countable neighbourhood basis.

Example 6.1.11. Let $\mathbb{R}$ be equipped with its canonical topology, let $X = \prod_{\alpha\in\mathbb{N}} \mathbb{R}$ equipped with the box topology, let

$$A = \{(x_n)_{n\geq 1} \mid x_n > 0\} \subseteq X,$$

and let $\vec{0} = (0)_{n\geq 1}$. We claim that $\vec{0} \in \overline{A}$ but no sequence of elements of A converges to $\vec{0}$. This then shows that X is not first countable Theorem 6.1.10 and thus not metrizable.

To see that $\vec{0} \in \overline{A}$, let U be an arbitrary neighbourhood of $\vec{0}$. By the definition of the box topology there exists a sequence $(\epsilon_n)_{n\geq 1}$ of real numbers with $\epsilon_n > 0$ for all $n \in \mathbb{N}$ such that

$$\prod_{n\in\mathbb{N}} (-\epsilon_n, \epsilon_n) \subseteq U.$$

Hence, if $\vec{a} = \left(\frac{1}{2}\epsilon_n\right)_{n\geq 1}$, then $\vec{a} \in A$ and

$$\vec{a} \in \prod_{n\in\mathbb{N}} (-\epsilon_n, \epsilon_n) \subseteq U.$$

Therefore, as U was arbitrary, $\vec{0} \in \overline{A}$ by Theorem 1.6.21.

To see that no sequence of elements of A converges to $\vec{0}$, suppose to the contrary that there exists a sequence $(\vec{a}_m)_{m\geq 1}$ of elements of A that converge to $\vec{0}$ in the box topology. For each $m \in \mathbb{N}$, write

$$\vec{a}_m = (a_{m,n})_{n\geq 1}$$

where $a_{m,n} > 0$ for all $n, m \in \mathbb{N}$ by the definition of A. Consider the set

$$U = \prod_{n\in\mathbb{N}} (-a_{n,n}, a_{n,n}),$$

which clearly is a neighbourhood of $\vec{0}$ in the box topology as $a_{n,n} > 0$ for all $n \in \mathbb{N}$. However, since $a_{m,m} \notin (-a_{m,m}, a_{m,m})$ for all $m \in \mathbb{N}$, we have that $\vec{a}_m \notin U$ for all $m \in \mathbb{N}$ thereby contradicting the fact that $(\vec{a}_m)_{m\geq 1}$ converges to $\vec{0}$ in the box topology. Hence no sequence of elements of A converges to $\vec{0}$.

Of course the box topology is known for producing examples. Here is an example on the product topology that fails due to cardinality issues.

Example 6.1.12. Let $\mathbb{R}$ be equipped with its canonical topology, let $X = \mathcal{F}(\mathbb{R}, \mathbb{R}) = \prod_{\alpha\in\mathbb{R}} \mathbb{R}$ equipped with the product topology, let

$$A = \left\{ x \in X \,\middle|\, \substack{x(\alpha)\in\{0,1\} \text{ for all } \alpha\in\mathbb{R}, \text{ and} \\ \{\alpha\in\mathbb{R} \,|\, x(\alpha)=0\} \text{ is finite}} \right\} \subseteq X,$$

and let $\vec{0} \in X$ be the unique element such that $\vec{0}(\alpha) = 0$ for all $\alpha \in \mathbb{R}$. We claim that $\vec{0} \in \overline{A}$ but no sequence of elements of A converges to $\vec{0}$. This then shows that X is not first countable Theorem 6.1.10 and thus not metrizable. Alternatively X is not metrizable by Theorem 6.1.3.

To see that $\vec{0} \in \overline{A}$, let U be an arbitrary neighbourhood of $\vec{0}$. By the definition of the product topology there exists a finite set $J \subseteq \mathbb{R}$ and neighbourhoods V_α of 0 in $\mathbb{R}$ for all $\alpha \in J$ such that if $V_\alpha = \mathbb{R}$ for all $\alpha \in \mathbb{R}\setminus J$ and if

$$V = \prod_{\alpha\in\mathbb{R}} V_\alpha,$$

then V is a neighbourhood of $\vec{0}$ contained in U. Let $\vec{a} \in X$ be the unique element such that

$$\vec{a}(\alpha) = \begin{cases} 0 & \text{if } \alpha \in J \\ 1 & \text{if } \alpha \in \mathbb{R}\setminus J \end{cases}.$$

Clearly $\vec{a} \in V \subseteq U$ by construction. Furthermore, since J is finite, we see that $\vec{a} \in A$ so that $\vec{a} \in A \cap U$. Therefore, as U was arbitrary, $\vec{0} \in \overline{A}$ by Theorem 1.6.21.

To see that no sequence of elements of A converges to $\vec{0}$, suppose to the contrary that there exists a sequence $(\vec{a}_n)_{n\geq 1}$ of elements of A that converge to $\vec{0}$ in the product topology. Let

$$I = \bigcup_{n\geq 1} \{\alpha \in \mathbb{R} \mid a_n(\alpha) = 0\} \subseteq \mathbb{R}.$$

By the definition of A we see that I is a countable union of finite subsets of $\mathbb{R}$ and thus finite. Hence there exists an $\alpha_0 \in \mathbb{R}\setminus I$. For each $\alpha \in \mathbb{R}$ let

$$V_\alpha = \begin{cases} \mathbb{R} & \text{if } \alpha \neq \alpha_0 \\ \left(-\frac{1}{2}, \frac{1}{2}\right) & \text{if } \alpha = \alpha_0 \end{cases}$$

and let $V = \prod_{\alpha\in\mathbb{R}} V_\alpha$. Hence V is a neighbourhood of $\vec{0}$. However, as $a_n(\alpha_0) = 1$ for all $n \in \mathbb{N}$ as $\alpha_0 \notin I$, we see that $a_n \notin V$ for all $n \in \mathbb{N}$ thereby contradicting the fact that $(\vec{a}_n)_{n\geq 1}$ converges to $\vec{0}$ in the product topology. Hence no sequence of elements of A converges to $\vec{0}$.

The above demonstrates that topological spaces that are Hausdorff and 'too large' are generally not going to be metrizable because having a plethora of points and being Hausdorff means there are going to be a ton of open sets thereby making it difficult for the topological space to be first countable. As such, one may often want to restrict to topological spaces that are closer to being countable. In particular, we can also consider the following strengthening of first countability. Note this strengthening is particularly nice as it is the condition used in Proposition 5.1.27 that enabled us to show that a certain regular topological spaces are normal. Thus, as a topological space must be normal to be metrizable by Remark 6.1.3, the following strengthening of first countability is often desired.

Definition 6.1.13. A topological space $(X, \mathcal{T})$ is said to be *second countable* if $(X, \mathcal{T})$ has a countable basis.

Example 6.1.14. Let X be a non-empty set. The discrete topology on X is second countable if and only if X is countable. Indeed if X is countable, $\{\{x\} \mid x \in X\}$ is clearly a countable basis for X. However, if X is uncountable, any basis for the discrete topology on X must be uncountable as there needs to be an element of the basis that is contained in $\{x\}$ for all $x \in X$.

Example 6.1.15. The canonical topology on $\mathbb{R}$ is second countable. To see this, consider the set

$$\mathcal{B} = \{(p, q) \mid p, q \in \mathbb{Q}, p < q\}$$

which is clearly a countable set of open subsets of $\mathbb{R}$ as $\mathbb{Q}$ is countable. To see that $\mathcal{B}$ is a basis for the topology on $\mathbb{R}$, let $x \in \mathbb{R}$ and U be an arbitrary neighbourhood of x. Hence there exists $a, b \in \mathbb{R}$ such that $x \in (a, b) \subseteq U$. The density of $\mathbb{Q}$ in $\mathbb{R}$ (which is part of the definition of $\mathbb{R}$) implies that there exists $p, q \in \mathbb{Q}$ such that

$$a < p < x < q < b$$

so that $(p, q) \in \mathcal{B}$ and $x \in (p, q) \subseteq U$. Hence, as x and U were arbitrary, $\mathcal{B}$ is a basis for $\mathbb{R}$ by Proposition 1.3.12.

Example 6.1.16. Let $\mathcal{T}_L$ be the lower limit topology on $\mathbb{R}$. Then $(\mathbb{R}, \mathcal{T}_L)$ is not second countable. To see this, suppose that $\mathcal{B}_0$ is a basis of $(\mathbb{R}, \mathcal{T}_L)$. For each $x \in \mathbb{R}$, consider the neighbourhood $[x, x+1)$ of x. Since $\mathcal{B}_0$ is a basis of $(\mathbb{R}, \mathcal{T}_L)$, there exists a $B_x \in \mathcal{B}_0$ such that $x \in B_x \subseteq [x, x+1)$. Clearly this implies that $B_x = [x, x+\epsilon_x)$ for some $\epsilon_x > 0$. Hence

$$X = \{[x, x+\epsilon_x) \mid x \in \mathbb{R}\}$$

is an uncountable collection of elements of $\mathcal{B}_0$ as $\mathbb{R}$ is uncountable and no two elements in X are equal. Thus $\mathcal{B}_0$ must be uncountable. Therefore $(\mathbb{R}, \mathcal{T}_L)$ cannot have a countable basis and thus is not second countable.

Proposition 6.1.17. *Any subspace of a first (second) countable topological space is first (second) countable, and the product topology on a countable product of first (second) countable topological spaces is first (second) countable.*

Proof. The fact that a subspace of a first countable topological space is first countable follows from the definition of the subspace topology; Definition 1.4.2. The fact that a subspace of a second countable topological space is second countable follows from the description of a basis for the subspace topology; Proposition 1.4.4.

The fact that the product topology on a countable product of first (second) countable topological space is first (second) countable follows from the description of a basis for the product topology (Corollary 1.4.15), the fact that a finite product of countable sets is countable, and the fact that the countable union of countable sets is countable. ∎

Of course, not every metrizable space need be second countable.

Example 6.1.18. Let $X = \mathcal{F}(\mathbb{N}, \mathbb{R}) = \prod_{\alpha \in \mathbb{N}} \mathbb{R}$ equipped with the uniform topology. Recall from Example 1.4.19 that the topology on X is strictly finer than the product topology but strictly coarser than the box topology. Clearly X is first countable as X is metrizable. However X is not second countable. To see this, we note that $\mathcal{F}(\mathbb{N}, \{0,1\})$ is a subspace of X and the subspace topology on $\mathcal{F}(\mathbb{N}, \{0,1\})$ is the discrete topology (i.e. all of the open balls of radius $\frac{1}{2}$ contain a single point). Therefore, as $\mathcal{F}(\mathbb{N}, \{0,1\})$ is uncountable, $\mathcal{F}(\mathbb{N}, \{0,1\})$ is not second countable by Example 6.1.14. Hence Proposition 6.1.17 implies that X cannot be second countable.

Of course, the problem with the topological space in Example 6.1.18 being second countable is that the topological space was 'too large'.

Proposition 6.1.19. *Let $(X, \mathcal{T})$ be a second countable topological space. Then:*

(1) There exists a countable dense subset of $(X, \mathcal{T})$.

(2) Every open cover of $(X, \mathcal{T})$ has a countable subcover.

Proof. To see (1), as $(X, \mathcal{T})$ is second countable, there exists a basis $\mathcal{B} = \{B_n\}_{n=1}^\infty$ of $(X, \mathcal{T})$ with $B_n \neq \emptyset$ for all $n \in \mathbb{N}$ (unless $X = \emptyset$). For each $B_n \neq \emptyset$, choose a point $x_n \in B_n$. We claim that $A = \{x_n\}_{n=1}^\infty$ is a countable dense subset of $(X, \mathcal{T})$. To see this, first we note that clearly A is countable. To see that A is dense in X, let $x \in X$ and U a neighbourhood of x be arbitrary. As $\mathcal{B}$ is a basis for $(X, \mathcal{T})$, there exists an $N \in \mathbb{N}$ such that $x \in B_N \subseteq U$. Hence $x_N \in U$. Therefore, as x and U were arbitrary, $x \in \overline{A}$ for all $x \in X$ by Theorem 1.6.21. Hence A is dense in $(X, \mathcal{T})$ as desired.

To see (2), as $(X, \mathcal{T})$ is second countable, there exists a basis $\mathcal{B} = \{B_n\}_{n=1}^\infty$ of $(X, \mathcal{T})$. Let $\{U_\alpha\}_{\alpha \in I}$ be an arbitrary open cover of $(X, \mathcal{T})$. To see that $\{U_\alpha\}_{\alpha \in I}$ has a countable subcover, first note we may assume without loss of generality that $U_\alpha \neq \emptyset$ for all $\alpha \in I$ and $B_n \neq \emptyset$ for all $n \in \mathbb{N}$ (i.e. remove every such set that is empty - this is no problem unless $X = \emptyset$ in which case the result is trivial). Next for each $n \in \mathbb{N}$ for which $B_n \subseteq U_\alpha$ for some $\alpha \in I$, choose an $\alpha_n \in I$ such that $B_n \subseteq U_{\alpha_n}$. Let

$$J = \{\alpha_n \mid n \in \mathbb{N}, \alpha_n \text{ exists}\} \subseteq I.$$

Clearly J is countable.

We claim that $\{U_\alpha\}_{\alpha \in J}$ is a subcover of $(X, \mathcal{T})$. To see this, let $x \in X$ be arbitrary. As $\{U_\alpha\}_{\alpha \in I}$ is an open cover of $(X, \mathcal{T})$, there exists an $\alpha_x \in I$ such that $x \in U_{\alpha_x}$. Moreover, since $\mathcal{B}$ is a basis of $(X, \mathcal{T})$, there exists an $n_x \in \mathbb{N}$ such that $x \in B_{n_x} \subseteq U_{\alpha_x}$. Hence $x \in B_{n_x} \subseteq U_{\alpha_{n_x}}$ by the definition of α_{n_x}. Hence, as $\alpha_{n_x} \in J$, we have that $x \in \bigcup_{\alpha \in J} U_\alpha$. Therefore, as $x \in X$ was arbitrary, $\{U_\alpha\}_{\alpha \in J}$ is a countable subcover of $(X, \mathcal{T})$. ∎

As 'having a countable dense subset' is a very useful property in analysis, a name should be given.

Definition 6.1.20. A topological space $(X, \mathcal{T})$ is said to be *separable* if $(X, \mathcal{T})$ has as countable dense subset.

Example 6.1.21. It is well known that $\mathbb{R}$ equipped with its canonical topology is separable as $\mathbb{Q}$ is dense in $\mathbb{R}$. Furthermore, the lower limit topology on $\mathbb{R}$ is also separable as $\mathbb{Q}$ is also dense in $\mathbb{R}$ with respect to the lower limit topology. In fact, $\mathbb{Q}$ is dense in $(\mathbb{R}, \mathcal{T}_L)$. To see this, note for all $[a, b) \in \mathcal{B}$ that there exists a $q \in \mathbb{Q}$ such that $q \in [a, b)$ by the properties of the real numbers. This immediately implies that $\overline{\mathbb{Q}} = \mathbb{R}$ so that $\mathbb{Q}$ is dense in $(\mathbb{R}, \mathcal{T}_L)$. Therefore, since $\mathbb{Q}$ is countable, $(\mathbb{R}, \mathcal{T}_L)$ is separable by definition.

In fact, in metrizable topological spaces, there is a converse to Proposition 6.1.19

Proposition 6.1.22. *Let $(X, \mathcal{T})$ be a separable, metrizable topological space. Then $(X, \mathcal{T})$ is second countable.*

Proof. By assumption there exists a metric $d : X \times X \to [0, \infty)$ that induces $\mathcal{T}$ and there exists a countable dense subset $\{x_n\}_{n=1}^\infty$ of $(X, \mathcal{T})$. To see that $(X, \mathcal{T})$ is second countable, consider the set

$$\mathcal{B} = \left\{ B_d\left(x_n, \frac{1}{m}\right) \,\middle|\, n, m \in \mathbb{N}\right\}.$$

We claim that $\mathcal{B}$ is a countable basis for $(X, \mathcal{T})$.

Clearly $\mathcal{B}$ is a countable set of open subsets of $(X, \mathcal{T})$. To see that $\mathcal{B}$ is a basis for $(X, \mathcal{T})$, let $x \in X$ and U a neighbourhood of x be arbitrary. Since $\mathcal{T}$ is the topology induced by d, there exists an $\epsilon > 0$ such that $B_d(x, \epsilon) \subseteq U$. Choose $N \in \mathbb{N}$ such that $\frac{1}{N} < \epsilon$. Since $\{x_n\}_{n=1}^\infty$ is dense in $(X, \mathcal{T})$, there exists an $m \in \mathbb{N}$ such that

$$x_m \in B_d\left(x, \frac{1}{2N}\right).$$

Thus $d(x, x_m) < \frac{1}{2N}$ so $d(y, x_m) < \frac{1}{N}$ for all $y \in B_d\left(x, \frac{1}{2N}\right)$ by the triangle inequality. Hence if

$$V = B_d\left(x_m, \frac{1}{2N}\right),$$

then

$$x \in V \subseteq B_d\left(x, \frac{1}{N}\right) \subseteq B_d(x, \epsilon) \subseteq U.$$

Therefore, since x and U were arbitrary, $\mathcal{B}$ is a basis for $(X, \mathcal{T})$ as desired. ∎

As a result of Proposition 6.1.22, we have the following.

Corollary 6.1.23. *The lower limit topology on $\mathbb{R}$ is not metrizable.*

Proof. Recall that $(\mathbb{R}, \mathcal{T}_L)$ is not second countable by Example 6.1.16. However, $(\mathbb{R}, \mathcal{T}_L)$ is separable by Example 6.1.21. Hence Proposition 6.1.22 implies $(\mathbb{R}, \mathcal{T}_L)$ is not metrizable. ∎

6.2 The Baire Category Theorem

As Proposition 6.1.19 just showed, second countable topological space are separable and thus have nice dense subsets. It turns out that complete metric spaces and compact Hausdorff topological spaces are also well-behaved with respect to dense sets. In particular, this section is devoted to demonstrating the amazing Baire Category Theorems for compact Hausdorff topological spaces (Theorem 6.2.10) and for complete metric spaces (Theorem 6.2.12) which describe what happens when we take the countable (after all, countability is nice) intersection of open dense subsets. These results are particularly useful when it comes to discussing discontinuities of functions.

Before we get to our main results, we require some background terminology.

Definition 6.2.1. Let $(X, \mathcal{T})$ be a topological space. A subset $A \subseteq X$ is said to be

- *nowhere dense* if $\mathrm{int}(\overline{A}) = \emptyset$.
- *first category in* $(X, \mathcal{T})$ if $A = \bigcup_{n=1}^{\infty} A_n$ where each $A_n \subseteq X$ is nowhere dense.
- *second category in* $(X, \mathcal{T})$ if A is not first category.
- *residual* if $X \setminus A$ is first category

Remark 6.2.2. It is important that A is considered as a subset in each part of Definition 6.2.1 instead of considering the subspace topology on A as the notion of the interior of a set can change. For example, equip $\mathbb{R}^2$ with the usual Euclidean topology and consider the set

$$A = \{(x, 0) \mid x \in [0, 1]\}.$$

Clearly A is closed in $\mathbb{R}^2$ and it is not difficult to see that $\mathrm{int}(A) = \emptyset$ with the interior computed in $\mathbb{R}^2$. However the interior of A as a subset of A equipped with the subspace topology is A. Hence A is nowhere dense in $\mathbb{R}^2$ but is dense in A when equipped with the subspace topology.

Remark 6.2.3. In the literature, first category subsets of a topological space are often called *meagre sets* as they are thought of as sets that do not have much size. Consequently, second category subsets of a topological space are often called *non-meagre sets*. We will stick with the first and second category definitions for historical reasons and as the word 'category' is used in the name of the main theorem we desire to prove in this section.

Onto some examples!

Example 6.2.4. Consider $X = \mathbb{R}$ equipped with the canonical topology. Clearly for each $x \in \mathbb{R}$ the set $\{x\}$ is nowhere dense. Furthermore, from this it is clear that $\mathbb{Q}$ is of first category in $\mathbb{R}$ and their complements are residual in $\mathbb{R}$.

Example 6.2.5. The Cantor set Example 1.6.11 is nowhere dense. Indeed the Cantor set is closed being the intersection of closed sets. Furthermore, if x is a point in the Cantor set and I is an open interval centred at x, then I cannot be contained in P_n (as defined in Example 1.6.11) for sufficiently large n as P_n does not contain an interval of length exceeding $\frac{1}{3^n}$ for all $n \in \mathbb{N}$. Hence the Cantor set has no interior and thus is of first category in $\mathbb{R}$.

Remark 6.2.6. Of course, one natural question is, "Is $\mathbb{R}$ of first or of second category in of itself?" More generally, given a topological space $(X, \mathcal{T})$, we are often interested in when $(X, \mathcal{T})$ is of first or second category in itself. Notice that if

$$X = \bigcup_{n=1}^{\infty} A_n \quad \text{then} \quad X = \bigcup_{n=1}^{\infty} \overline{A_n}.$$

Therefore, as the closure of a nowhere dense set is clearly nowhere dense, $(X, \mathcal{T})$ is of first category in itself if and only if X is a countable union of closed nowhere dense sets. Of course, a closed nowhere dense subset of X is the same as a closed subset of X with empty interior. The simplification of this later condition is that one no longer needs to take the closure of the closed sets and need only consider their interiors.

Instead of considering closed sets with empty interior, the following enables us to flip things around.

Lemma 6.2.7. *Let $(X, \mathcal{T})$ be a topological space and let $A \subseteq X$. Then A has empty interior in $(X, \mathcal{T})$ if and only if $X \setminus A$ is dense in $(X, \mathcal{T})$.*

Proof. Suppose A has empty interior. Hence $\emptyset = \mathrm{int}(A)$. To see that $X \setminus A$ is dense in $(X, \mathcal{T})$, let $x \in X$ and U a neighbourhood of x be arbitrary. Since $\mathrm{int}(A) = \emptyset$, we know that U is not a subset of A so there exists a $y \in X \setminus A$ such that $y \in U$. Therefore, as U was arbitrary, $x \in \overline{X \setminus A}$. Hence, as $x \in X$ was arbitrary, $X \setminus A$ is dense in X.

Conversely, suppose $X \setminus A$ is dense in $(X, \mathcal{T})$. To see that A has empty interior, let $a \in A$ and U a neighbourhood of a be arbitrary. Since $X \setminus A$ is dense in X, there exists an $x \in X \setminus A$ such that $x \in U$. Hence U is not a subset of A. Hence, as U was arbitrary, $a \notin \mathrm{int}(A)$. Therefore, as $a \in A$ was arbitrary, $\mathrm{int}(A) = \emptyset$ as desired. ∎

As the complement of closed sets are open sets, we have the following.

Lemma 6.2.8. *Let $(X, \mathcal{T})$ be a topological space. Then $(X, \mathcal{T})$ is of first category in itself if and only if there exists a sequence $(U_n)_{n\geq 1}$ of open dense subsets of $(X, \mathcal{T})$ with $\bigcap_{n=1}^{\infty} U_n = \emptyset$.*

Proof. The result follows directly from Lemma 6.2.7 and Remark 6.2.6. ∎

Of course Remark 6.2.6 and Lemma 6.2.8 give nice characterizations of when a topological space is of first category in itself and thus when a topological space is of second category in itself; that is, when the countable union of closed sets with empty interior is not all of X and when the intersection of open dense subsets is non-empty. Of course, we can take both of these concepts to the extreme.

Definition 6.2.9. A topological space $(X, \mathcal{T})$ is said to be a *Baire space* if one of the following equivalent (by Lemma 6.2.7) conditions holds:

- If $\{F_n\}_{n=1}^{\infty}$ are closed subsets of $(X, \mathcal{T})$ with empty interior, then $\bigcup_{n=1}^{\infty} F_n$ has empty interior.
- If $\{U_n\}_{n=1}^{\infty}$ are open dense subsets of $(X, \mathcal{T})$, then $\bigcap_{n=1}^{\infty} U_n$ is dense in $(X, \mathcal{T})$.

Clearly every Baire space is of second category in itself by Remark 6.2.6 or Lemma 6.2.8. In fact, being a Baire space is a priori much stronger than being second category in itself (it is possible to construct a topological space that is second category in itself that is not a Baire space).

It turns out that some of the nicest spaces we have studied are Baire spaces.

Theorem 6.2.10 (Baire's Category Theorem - Compact Hausdorff Spaces). *Every compact Hausdorff topological space is a Baire space.*

Proof. Let $(X, \mathcal{T})$ be a compact Hausdorff topological space. To see that $(X, \mathcal{T})$ is a Baire space, let $\{U_n\}_{n=1}^{\infty}$ be a countable set of open dense subsets of $(X, \mathcal{T})$. To see that $\bigcap_{n=1}^{\infty} U_n$ is dense in $(X, \mathcal{T})$, let $x_0 \in X$ and U a neighbourhood of x_0 be arbitrary. We desire to show that there $U \cap (\bigcap_{n=1}^{\infty} U_n) \neq \emptyset$. To do this, we will use regularity and the finite intersection property.

Since compact Hausdorff topological spaces are T_4 and thus T_3, Lemma 5.1.15 implies there exists a neighbourhood V_0 of x_0 such that

$$x_0 \in V_0 \subseteq \overline{V_0} \subseteq U.$$

Since U_1 is dense in $(X, \mathcal{T})$, there exists an element $a_1 \in U_1 \cap V_0$. Since compact Hausdorff topological spaces are T_4 and thus T_3, and since $U_1 \cap V_0$

is a neighbourhood of a_1, Lemma 5.1.15 implies there exists a neighbourhood V_1 of a_1 such that

$$a_1 \in V_1 \subseteq \overline{V_1} \subseteq U_1 \cap V_0.$$

Since U_2 is dense in $(X, \mathcal{T})$, there exists an element $a_2 \in U_2 \cap V_1$. By repeating this process ad infinitum, there exists a sequence of points $(a_n)_{n\geq 1}$ in X and open sets $\{V_n\}_{n=1}^{\infty} \subseteq \mathcal{T}$ such that $a_{n+1} \in U_{n+1} \cap V_n$ for all $n \in \mathbb{N}$ and

$$a_{n+1} \in V_{n+1} \subseteq \overline{V_{n+1}} \subseteq U_{n+1} \cap V_n$$

for all $n \in \mathbb{N}$.

From the above construction, we see that

$$a_{n+1} \in V_{n+1} \subseteq \overline{V_{n+1}} \subseteq U_{n+1} \cap V_n \subseteq V_n \subseteq \overline{V_n}$$

for all $n \in \mathbb{N}$. Hence it is trivial to see that $\{\overline{V_n}\}_{n=1}^{\infty}$ has the finite intersection property. Therefore, since $(X, \mathcal{T})$ is compact, Theorem 3.2.2 implies that

$$\bigcap_{n=1}^{\infty} \overline{V_n} \neq \emptyset.$$

Let $y \in \bigcap_{n=1}^{\infty} \overline{V_n}$. We claim that $y \in U \cap (\bigcap_{n=1}^{\infty} U_n)$. To see this, notice that $\overline{V_n} \subseteq U_n$ for all $n \in \mathbb{N}$. Hence as $y \in \overline{V_n}$ for all $n \in \mathbb{N}$, $y \in U_n$ for all $n \in \mathbb{N}$ so $y \in \bigcap_{n=1}^{\infty} U_n$. Similarly, as $\overline{V_{n+1}} \subseteq \overline{V_n}$ for all $n \in \mathbb{N}$ and as

$$\overline{V_1} \subseteq U_1 \cap V_0 \subseteq V_0 \subseteq \overline{V_0} \subseteq U,$$

we see that $y \in \bigcap_{n=1}^{\infty} \overline{V_n} \subseteq U$ so $y \in U$ as desired. Hence $U \cap (\bigcap_{n=1}^{\infty} U_n) \neq \emptyset$. Therefore, as U was arbitrary, $x_0 \in \overline{\bigcap_{n=1}^{\infty} U_n}$. Hence, as $x_0 \in X$ was arbitrary, $\bigcap_{n=1}^{\infty} U_n$ is dense in $(X, \mathcal{T})$ as desired. ∎

Remark 6.2.11. By Theorem 6.2.10, the Cantor set is a Baire space even though Example 6.2.5 showed that the Cantor set is of first category in $\mathbb{R}$.

By using the proof of Theorem, 6.2.10 as a roadmap, we can obtain that every complete metric space is a Baire space where compactness is replaced with dwindling diameters.

Theorem 6.2.12 (Baire's Category Theorem - Metric Spaces). *Every complete metric space is a Baire space.*

Proof. Let (X, d) be a complete metric space. To see that (X, d) is a Baire space, let $\{U_n\}_{n=1}^{\infty}$ be a countable set of open dense subsets of (X, d). To see that $\bigcap_{n=1}^{\infty} U_n$ is dense in (X, d), let $x_0 \in X$ and U a neighbourhood of x_0 be arbitrary. We desire to show that there $U \cap (\bigcap_{n=1}^{\infty} U_n) \neq \emptyset$. Due to the fact that (X, d) is a metric space, we may assume without loss of generality that $U = B_d(x_0, \epsilon)$.

To show that $U \cap (\bigcap_{n=1}^{\infty} U_n) \neq \emptyset$, we will use an analogue of regularity that enables us to invoke Cantor's Theorem (Theorem 4.1.11). For our analogue of regularity, we note that if $y \in X$ and $r > 0$ then for any $0 < r' < r$ we have that

$$B_d[y, r'] \subseteq B_d(y, r).$$

Let $r_1 = \frac{1}{2}\epsilon$. Since U_1 is dense in (X, d), there exists an element $a_1 \in U_1$ such that $d(a_1, x_0) < r_1$. Since U_1 is open, by using the above comment there exists an $0 < r_2 < \frac{1}{4}\epsilon$ such that $B_d[a_1, r_2] \subseteq U_1$ (i.e. choose an open ball around a_1 contained in U_1 and then decrease the radius of the ball).

Since U_2 is dense in (X, d), there exists an element $a_2 \in U_2$ such that $d(a_2, a_1) < r_2$. Hence $a_2 \in B_d(a_1, r_2)$ so $a_2 \in U_2 \cap B_d(a_1, r_2)$. Hence, since $U_2 \cap B_d(a_1, r_2)$ is open, there exists an $0 < r_2 < \frac{1}{2^3}\epsilon$ such that

$$B_d[a_2, r_3] \subseteq U_2 \cap B_d(a_1, r_2).$$

By recursion, for each $n \in \mathbb{N}$ there exists an $a_n \in U_n \cap B_d(a_{n-1}, r_n)$ and an $0 < r_{n+1} < \frac{1}{2^{n+1}}\epsilon$ such that $d(a_n, a_{n-1}) < r_n$ and

$$B_d[a_n, r_{n+1}] \subseteq U_n \cap B_d(a_{n-1}, r_n).$$

For each $n \in \mathbb{N}$, let $F_n = B_d[a_n, r_{n+1}]$. Clearly $(F_n)_{n \geq 1}$ is a sequence of non-empty closed subsets of X such that $F_{n+1} \subseteq F_n$ and $\lim_{n\to\infty} \operatorname{diam}(F_n) = 0$ (as $\operatorname{diam}(F_n) \leq 2r_{n+1}$). Hence Cantor's Theorem (Theorem 4.1.11) implies that $\bigcap_{n=1}^{\infty} F_n \neq \emptyset$.

Let $y \in \bigcap_{n=1}^{\infty} F_n$. We claim that $y \in \bigcap_{n=1}^{\infty} U_n$ and $d(x_0, y) < \epsilon$. To see this, notice that $F_n \subseteq U_n$ for all $n \in \mathbb{N}$. Hence as $y \in F_n$ for all $n \in \mathbb{N}$, $y \in U_n$ for all $n \in \mathbb{N}$ so $y \in \bigcap_{n=1}^{\infty} U_n$. To see that $d(x_0, y) < \epsilon$, we note that $y \in F_1 = B_d[a_1, r_2]$ so $d(y, a_1) \leq r_2$. Hence

$$d(x_0, y) \leq d(x, a_1) + d(a_1, y) \leq r_1 + r_2 < \epsilon.$$

Thus $y \in B_d(x, \epsilon) = U$ so $U \cap (\bigcap_{n=1}^{\infty} U_n) \neq \emptyset$. Therefore, as U was arbitrary, $x_0 \in \overline{\bigcap_{n=1}^{\infty} U_n}$. Hence, as $x_0 \in X$ was arbitrary, $\bigcap_{n=1}^{\infty} U_n$ is dense in (X, d) as desired. ■

Of course for non-complete metric spaces may or may not be Baire spaces.

Example 6.2.13. Let $\mathbb{Z}$ be equipped with the discrete topology. Since the only open dense subset of $\mathbb{Z}$ is $\mathbb{Z}$ as every singleton is open, $\mathbb{Z}$ is clearly a Baire space.

Example 6.2.14. Let $\mathbb{Q}$ be equipped with the subspace topology inherited from the canonical topology on $\mathbb{R}$. Then $\mathbb{Q}$ is not a Baire space. Indeed for each $q \in \mathbb{Q}$ let $U_q = \mathbb{Q} \setminus \{q\}$. Clearly U_q is an open subset of $\mathbb{Q}$ which is dense in $\mathbb{Q}$. However, as $\mathbb{Q}$ is countable and $\bigcap_{q\in\mathbb{Q}} U_q = \emptyset$ by construction, we see that $\mathbb{Q}$ is not a Baire space and not second countable in itself.

As $\mathbb{R}$ is a Baire space, Example 6.2.14 demonstrates a subspace of a Baire space need not be a Baire space. However, certain subspace of Baire spaces are Baire spaces thereby increasing our repertoire of Baire spaces.

Proposition 6.2.15. *Every open subspace of a Baire space is a Baire space.*

Proof. Let (X, T) be a Baire space and let Y be an open subset of $(X, \mathcal{T})$. To see that Y is a Baire space when equipped with the subspace topology, let $\{F_n\}_{n=1}^\infty$ be arbitrary closed subsets of Y with empty interior in Y. We desire to show that $\bigcup_{n=1}^\infty F_n$ has empty interior in Y.

Suppose to the contrary that $\bigcup_{n=1}^\infty F_n$ does not have empty interior in Y. Hence there exists a non-empty open subset V of Y such that $V \subseteq \bigcup_{n=1}^\infty F_n$. By the definition of the subspace topology, there exists a $W \in \mathcal{T}$ such that $V = W \cap Y$. Hence $V \in \mathcal{T}$ as $Y \in T$.

For each $n \in \mathbb{N}$, let C_n be the closure of F_n in X. Hence, as F_n is closed in Y, Lemma 1.6.20 implies that $C_n \cap Y = F_n$. We claim that C_n has empty interior in $(X, \mathcal{T})$. To see this, suppose to the contrary that there exists a non-empty open set $U \in \mathcal{T}$ such that $U \subseteq C_n$. Since $U \subseteq C_n$ and C_n is the closure of F_n in X, it must be the case that $U \cap F_n \neq \emptyset$. Thus, as $F_n \subseteq Y$, $U \cap Y \neq \emptyset$. Hence $U \cap Y$ is a non-empty open subset of Y such that

$$U \cap Y \subseteq C_n \cap Y = F_n.$$

As this contradicts the fact that F_n has empty interior in Y, C_n has empty interior in $(X, \mathcal{T})$ for all $n \in \mathbb{N}$.

Since $(X, \mathcal{T})$ is a Baire space, $\bigcup_{n=1}^\infty C_n$ has empty interior in $(X, \mathcal{T})$. However $V \in \mathcal{T}$ is non-empty and

$$V \subseteq \bigcup_{n=1}^\infty F_n \subseteq \bigcup_{n=1}^\infty C_n$$

thereby contradicting the fact that $\bigcup_{n=1}^\infty C_n$ has empty interior in $(X, \mathcal{T})$. Hence, as we have obtained a contradiction, $\bigcup_{n=1}^\infty F_n$ has empty interior in Y. Thus, as $\{F_n\}_{n=1}^\infty$ were arbitrary, Y is a Baire space. ∎

Corollary 6.2.16. *Every locally compact Hausdorff topological space is a Baire space.*

Proof. As every compact Hausdorff topological space is a Baire space by the Baire Category Theorem (Theorem 6.2.10) and as every locally compact Hausdorff topological space is an open subspace its one-point compactification, which is a compact Hausdorff topological space, the result follows from Proposition 6.2.15. ∎

For closed subspaces, we are not so lucky.

Example 6.2.17. Let $\mathcal{T}$ be the subspace topology on $\mathbb{Q}$ inherited from the canonical topology on $\mathbb{R}$. Choose $x_0 \in \mathbb{R} \setminus \mathbb{Q}$, let $X = \mathbb{Q} \cup \{x_0\}$, and let

$$\mathcal{T}_0 = \{\emptyset\} \cup \{U \cup \{x_0\} \mid U \in \mathcal{T}\}.$$

Clearly $\mathcal{T}_0$ is a topology on X as $\mathcal{T}$ is a topology on $\mathbb{Q}$. Furthermore, by construction $(\mathbb{Q}, \mathcal{T})$ is a subspace of $(X, \mathcal{T}_0)$.

We claim that $(X, \mathcal{T}_0)$ is a Baire space. To see this, first note that since $\{x_0\}$ is open in $(X, \mathcal{T}_0)$ that every dense subset of $(X, \mathcal{T}_0)$ must contain x_0. Furthermore, by the definition of $\mathcal{T}_0$ we see that $\{x_0\}$ is dense in $(X, \mathcal{T}_0)$ as x_0 is contained in every non-empty open subset of $(X, \mathcal{T}_0)$. Thus as every dense subset of $(X, \mathcal{T}_0)$ contains x_0, the intersection of any collection of dense subsets of $(X, \mathcal{T}_0)$ contains x_0 and thus is dense in $(X, \mathcal{T}_0)$ as $\{x_0\}$ is dense in $(X, \mathcal{T}_0)$. Hence $(X, \mathcal{T}_0)$ is a Baire space as desired.

Since $\mathbb{Q} = X \setminus \{x_0\}$ is closed in $(X, \mathcal{T}_0)$, $(\mathbb{Q}, \mathcal{T})$ is a closed subspace of $(X, \mathcal{T}_0)$. However $(\mathbb{Q}, \mathcal{T})$ is not a Baire space by Example 6.2.14. Hence a closed subspace of a Baire space need not be Baire.

To conclude this section, we note there are numerous applications of the Baire Category Theorem. To illustrate one such example related to continuous functions, we first develop some technology that will be useful in subsequent sections.

Definition 6.2.18. Let $(X, \mathcal{T})$ be a topological space. A subset $A \subseteq X$ is said to be G_δ if there exists a countable set of open subsets $\{U_n\}_{n=1}^\infty$ of $(X, \mathcal{T})$ such that $A = \bigcap_{n=1}^\infty U_n$.

Similarly, a subset $B \subseteq X$ is said to be F_σ if there exists a countable set of closed subsets $\{F_n\}_{n=1}^\infty$ of $(X, \mathcal{T})$ such that $A = \bigcup_{n=1}^\infty F_n$.

Remark 6.2.19. It is not difficult to see using De Morgan's Laws that A is G_δ if and only if $X \setminus A$ is F_σ.

Of course we do not really need to give examples of G_δ and F_σ sets as clearly every open set is G_δ, every closed set is F_σ, and the notions are not that complicated. However, the following is quite useful to note.

Lemma 6.2.20. *Every closed subset of a metric space is G_δ.*

Proof. Let F be a closed subset of a metric space (X, d). If $F = \emptyset$ then, as $\emptyset$ is open and as $\bigcap_{n=1}^\infty \emptyset = \emptyset$, we obtain that F is G_δ.

Otherwise, suppose F is not empty. For each $n \in \mathbb{N}$, let

$$U_n = \bigcup_{x \in F} B_d\left(x, \frac{1}{n}\right).$$

Clearly each U_n is an open subset such that $F \subseteq U_n$. Hence

$$F \subseteq \bigcap_{n=1}^\infty U_n.$$

For the other inclusion, let $x \in X \setminus F$ be arbitrary. Therefore, as F is closed, $\mathrm{dist}(\{x\}, F) > 0$ (for otherwise there would exist a sequence of elements of F that converge to x thereby implying $x \in F$). Choose $n \in \mathbb{N}$ such that

$$\mathrm{dist}(\{x\}, F) \geq \frac{1}{n} > 0.$$

Hence $d(x, y) \geq \frac{1}{n}$ for all $y \in F$ so $x \notin U_n$ and thus $x \notin \bigcap_{n=1}^{\infty} U_n$. Therefore, as $x \in X \setminus F$ was arbitrary,

$$F = \bigcap_{n=1}^{\infty} U_n.$$

Hence F is G_δ as desired. ∎

One useful example of a set that is not G_δ is as follows.

Proposition 6.2.21. *Let $\mathbb{R}$ be equipped with its canonical topology. The rational numbers are not a G_δ subset of $\mathbb{R}$ and the irrational numbers are not an F_σ subset of $\mathbb{R}$.*

Proof. As $\mathbb{Q}$ is a G_δ subset of $\mathbb{R}$ if and only if $\mathbb{R} \setminus \mathbb{Q}$ is an F_σ subset of $\mathbb{R}$, it suffices to prove the former.

Suppose to the contrary that $\mathbb{Q}$ is a G_δ subset of $\mathbb{R}$. Hence there exists a countable set $\{U_n\}_{n=1}^{\infty}$ of open subsets of $\mathbb{R}$ such that $\mathbb{Q} = \bigcap_{n=1}^{\infty} U_n$. Therefore $\mathbb{Q} \subseteq U_n$ for all n so each U_n is dense in $\mathbb{R}$. Hence each $\mathbb{R} \setminus U_n$ is closed and nowhere dense by Lemma 6.2.7. Thus $\mathrm{int}(\mathbb{R} \setminus U_n) = \emptyset$ so $\mathbb{R} \setminus U_n$ contains no open intervals for all $n \in \mathbb{N}$.

Since $\mathbb{Q}$ is countable, we may write $\mathbb{Q} = \{r_n\}_{n=1}^{\infty}$. For each $n \in \mathbb{N}$, let

$$F_n = (\mathbb{R} \setminus U_n) \cup \{r_n\}.$$

Clearly each F_n is closed being the union of two closed sets and F_n does not contain an open interval since $\mathbb{R} \setminus U_n$ does not contain an open interval. Hence $\mathrm{int}(F_n) = \emptyset$ as F was closed and thus F_n is nowhere dense.

Since

$$\mathbb{R} \setminus \mathbb{Q} = \mathbb{R} \setminus \left(\bigcap_{n=1}^{\infty} U_n \right) = \bigcup_{n=1}^{\infty} \mathbb{R} \setminus U_n \subseteq \bigcup_{n=1}^{\infty} F_n$$

and since

$$r_m \in F_m \subseteq \bigcup_{n=1}^{\infty} F_n$$

for all $m \in \mathbb{N}$, we obtain that

$$\mathbb{R} = \bigcup_{n=1}^{\infty} F_n.$$

Therefore $\mathbb{R}$ is a countable union of nowhere dense sets and thus $\mathbb{R}$ is of first category. However, as $\mathbb{R}$ is complete, the Baire Category Theorem (Theorem 6.2.12) implies that $\mathbb{R}$ is not of first category thereby providing a contradiction. Hence $\mathbb{Q}$ is not a G_δ set. ∎

Using Proposition 6.2.21 and the following characterization of the discontinuities of a function between metric spaces, we can demonstrate that certain sets cannot be the discontinuities of a real-valued function.

Lemma 6.2.22. *Let (X, d_X) and (Y, d_Y) be metric spaces, let $f : X \to Y$, and let*

$$D(f) = \{x \in X \mid f \text{ is not continuous at } x\}.$$

For each $n \in \mathbb{N}$ let

$$D_n(f) = \left\{ x \in X \;\middle|\; \begin{array}{c} \text{for every } \delta > 0 \text{ there exists } x_1, x_2 \in X \text{ such that} \\ d_X(x, x_1) < \delta, d_X(x, x_2) < \delta, \text{ and} \\ d_Y(f(x_1), f(x_2)) \geq \frac{1}{n} \end{array} \right\}.$$

Then $D_n(f)$ is closed for all $n \in \mathbb{N}$ and $D(f) = \bigcup_{n=1}^\infty D_n(f)$. Hence $D_n(f)$ is an F_σ subset of (X, d_x).

Proof. Fix $n \in \mathbb{N}$. To see that $D_n(f)$ is closed, let $(x_\lambda)_{\lambda\in\Lambda}$ be an arbitrary net of elements of $D_n(f)$ that converges to some $x \in X$. To see that $x \in D_n(f)$, let $\delta > 0$ be arbitrary. Since $(x_\lambda)_{\lambda\in\Lambda}$ converges to x, there exists a $\lambda_0 \in \Lambda$ such that $d_X(x, x_{\lambda_0}) < \frac{1}{2}\delta$. Furthermore, since $x_{\lambda_0} \in D_n(f)$, there exists $x_1, x_2 \in X$ such that $d_X(x_{\lambda_0}, x_1) < \frac{1}{2}\delta$, $d_X(x_{\lambda_0}, x_2) < \frac{1}{2}\delta$, and $d_Y(f(x_1), f(x_2)) \geq \frac{1}{n}$. As $d_X(x, x_1) < \delta$ and $d_X(x, x_2) < \delta$ by the Triangle Inequality, and as $d_Y(f(x_1), f(x_2)) \geq \frac{1}{n}$, we obtain that $x \in D_n(f)$ as $\delta > 0$ was arbitrary. Hence as $(x_\lambda)_{\lambda\in\Lambda}$ was arbitrary, $D_n(f)$ is closed.

To see that $D(f) = \bigcup_{n=1}^\infty D_n(f)$, first suppose $x \in \bigcup_{n=1}^\infty D_n(f)$. Hence $x \in D_n(f)$ for some $n \in \mathbb{N}$. To see that f is discontinuous at x, suppose to the contrary that f is continuous at x. Notice by the definition of $D_n(f)$ that for each $m \in \mathbb{N}$ there exists points $x_{1,m}, x_{2,m} \in X$ such that $d_X(x, x_{1,m}) < \frac{1}{m}$, $d_X(x, x_{2,m}) < \frac{1}{m}$, and $d_Y(f(x_{1,m}, f(x_{2,m})) \geq \frac{1}{n}$. Since $(x_{1,m})_{m\geq1}$ and $(x_{2,m})_{m\geq1}$ converge to x, the continuity of f implies $\lim_{m\to\infty} d_Y(f(x), f(x_{1,m})) = 0 = \lim_{m\to\infty} d_Y(f(x), f(x_{1,m}))$, which, together with the Triangle Inequality, contradicts the fact that

$$d_Y(f(x_{1,m}), f(x_{2,m})) \geq \frac{1}{n}$$

for all $m \geq 1$. Hence we have obtained a contradiction so $x \in D(f)$. Hence $\bigcup_{n=1}^\infty D_n(f) \subseteq D(f)$.

For the other inclusion, notice if $x \in D(f)$ then f is discontinuous at x. Therefore there exists an $\epsilon > 0$ such that for all $\delta > 0$ there exists a $x_1 \in X$ such that $d_X(x, x_1) < \delta$ yet $d_Y(f(x), f(x_1)) \geq \epsilon$. Choose $n \in \mathbb{N}$ such that $\frac{1}{n} < \epsilon$. By taking $x_2 = x$ in the definition of $D_n(f)$, we see that $x \in D_n(f)$. Hence, as x was arbitrary, $D(f) \subseteq \bigcup_{n=1}^\infty D_n(f)$ as desired. ∎

Theorem 6.2.23. *There does not exists a function $f : \mathbb{R} \to \mathbb{R}$ that is continuous at each point in $\mathbb{Q}$ yet discontinuous at each point in $\mathbb{R} \setminus \mathbb{Q}$.*

Proof. Let $f : \mathbb{R} \to \mathbb{R}$. By Lemma 6.2.22 the set of discontinuities of f are F_σ. Thus the points where f is continuous must be a G_δ set. As $\mathbb{Q}$ is not G_δ by Proposition 6.2.21, f cannot be continuous at each point in $\mathbb{Q}$ yet discontinuous at each point in $\mathbb{R} \setminus \mathbb{Q}$. ∎

Furthermore, we can use the fact that the irrational numbers are not a F_σ subset of $\mathbb{R}$ to finally demonstrate an example of a product of two normal topological spaces that is not normal.

Example 6.2.24. Let $\mathcal{T}_L$ be the lower limit topology on $\mathbb{R}$, let $X = \mathbb{R} \times \mathbb{R}$ equipped with the product topology where both copies of $\mathbb{R}$ are equipped with $\mathcal{T}_L$, and let

$$A = \{(x, -x) \mid x \in \mathbb{Q}\} \qquad \text{and} \qquad B = \{(x, -x) \mid x \in \mathbb{R} \setminus \mathbb{Q}\}.$$

Our goal is to show that A and B are closed subsets of X such that whenever U and V are open subsets of X such that $A \subseteq U$ and $B \subseteq V$, then $U \cap V \neq \emptyset$. This then shows us that X is not normal. Therefore, as the lower limit topology is T_4 by Example 5.1.26, we have an example of a product of two T_4 topological spaces that is not normal.

To see that A and B are closed subsets of X, first suppose $(y_\lambda)_{\lambda \in \Lambda}$ is a net of elements of $A \cup B$ that converge to some element $y \in X$. By the definition of $A \cup B$, for each $\lambda \in \Lambda$ we can write $y_\lambda = (x_\lambda, -x_\lambda)$ for some $x_\lambda \in \mathbb{R}$. Furthermore, write $y = (y_1, y_2) \in \mathbb{R}^2$.

Since $(y_\lambda)_{\lambda \in \Lambda}$ converges to y, we know Theorem 1.5.25 that $(x_\lambda)_{\lambda \in \Lambda}$ converges to y_1 in $(\mathbb{R}, \mathcal{T}_L)$ and $(-x_\lambda)_{\lambda \in \Lambda}$ converges to y_2 in $(\mathbb{R}, \mathcal{T}_L)$. By Proposition 1.5.23, we know that this implies for every $\epsilon > 0$ there exists $\lambda_1, \lambda_2 \in \Lambda$ such that $x_\lambda \in [y_1, y_1 + \epsilon)$ for all $\lambda \geq \lambda_1$ and $-x_\lambda \in [y_2, y_2 + \epsilon)$ for every $\lambda \geq \lambda_2$. As there exists a $\lambda_0 \in \Lambda$ such that $\lambda_0 \geq \lambda_1$ and $\lambda_0 \geq \lambda_2$ by the properties of a directed set, we have that $x_{\lambda_0} \in [y_1, y_1 + \epsilon)$ and $-x_{\lambda_0} \in [y_2, y_2 + \epsilon)$. Hence

$$[y_1, y_1 + \epsilon) \cap (-y_2 - \epsilon - y_2] \neq \emptyset$$

for every $\epsilon > 0$. This is only possible if $y_2 = -y_1$. Hence $y = (y_1, -y_1)$ for some $y_1 \in \mathbb{R}$.

We now claim that $(y_\lambda)_{\lambda \in \Lambda}$ converges to y only if there exists a $\lambda_3 \in \Lambda$ such that $y_\lambda = y$ for all $\lambda \geq \lambda_3$. Indeed consider the set

$$U = [y_1, y_1 + 1) \cup [-y_1, -y_1 + 1).$$

Clearly $y \in U$ and U is open in X by the definition of the product topology. However, we notice that $(x, -x) \in U$ if and only if $x = y_1$. Hence, by the definition of a convergent net, $(y_\lambda)_{\lambda\in\Lambda}$ converges to y only if there exists a $\lambda_3 \in \Lambda$ such that $y_\lambda = y$ for all $\lambda \geq \lambda_3$.

By the above, we see that if a net in A converges to an element $y \in X$, then $y \in A$ and if a net in B converges to an element $y \in X$, then $y \in B$. Hence A and B are closed by Theorem 1.6.14.

Next, let V be an arbitrary open subset of X such that $B \subseteq V$. For each $x \in \mathbb{R} \setminus \mathbb{Q}$, we notice since $(x, -x) \in B \subseteq V$, since V is open in X, and since a neighbourhood basis for the product topology is the Cartesian product of bases for the respective topologies, there exists $a, b \in \mathbb{R}$ with $a, b > 0$ such that $[x, x+a) \times [-x, -x+b) \subseteq V$. Hence for each $x \in \mathbb{R} \setminus \mathbb{Q}$ we may choose a $\delta_x > 0$ such that $[x, x+\delta_x) \times [-x, -x+\delta_x) \subseteq V$.

For each $n \in \mathbb{N}$ let $X_n = \{x \in \mathbb{R} \setminus \mathbb{Q} \mid \delta_x > \frac{1}{n}\}$. We claim there exists a $z \in \mathbb{Q}$ and an $n \in \mathbb{N}$ such that $z \in \overline{X_n}$ where the closure is taken in the canonical topology on $\mathbb{R}$. To see this, notice by construction, we clearly have $\mathbb{R} \setminus \mathbb{Q} = \bigcup_{n=1}^\infty X_n$. Suppose for all $n \in \mathbb{N}$ that $\overline{X_n} \cap \mathbb{Q} = \emptyset$ where the closures are taken in the canonical topology on $\mathbb{R}$. Then we have that $\mathbb{R} \setminus \mathbb{Q} = \bigcup_{n=1}^\infty \overline{X_n}$ where $\overline{X_n}$ are closed subsets of the canonical topology on $\mathbb{R}$. This implies that $\mathbb{R} \setminus \mathbb{Q}$ is an F_σ subset of $\mathbb{R}$. As this contradicts Proposition 6.2.21, the proof is complete.

Now, we claim that $(z, -z) \in \overline{V}$ where the closure is in X. To see this, let U_0 be an arbitrary neighbourhood of $(z, -z)$ in X. By the same argument as above, there exists an $\epsilon > 0$ such that $[z, z+\epsilon) \times [-z, -z+\epsilon) \subseteq U_0$. Furthermore, we may assume without loss of generality that $\epsilon < \frac{1}{n}$.

Since $z \in \overline{X_n}$ with the closure being with respect to the canonical topology, there exists an $x \in X_n$ such that $|z - x| < \frac{\epsilon}{2}$. Therefore

$$z < x + \frac{\epsilon}{2} < z + \epsilon \qquad \text{and} \qquad -z < -x + \frac{\epsilon}{2} < -z + \epsilon.$$

Hence if $c = x + \frac{\epsilon}{2}$ and $d = -x + \frac{\epsilon}{2}$, then

$$c \in [z, z+\epsilon) \cap [x, x+\epsilon) \qquad \text{and} \qquad d \in [-z, -z+\epsilon) \cap [-x, -x+\epsilon).$$

Clearly this implies that $(c, d) \in U_0$. Moreover, as

$$(c, d) \in [x, x+\epsilon) \times [-x, -x+\epsilon) \subseteq \left[x, x + \frac{1}{n}\right) \times \left[-x, -x + \frac{1}{n}\right),$$

we see that $(c, d) \in V$ by the definition of X_n. Therefore, as U_0 was arbitrary, we see that V intersects every neighbourhood of $(z, -z)$ and thus $(z, -z) \in \overline{V}$.

To conclude the proof that X is not normal, suppose to the contrary that X is normal. Therefore, as A and B are closed subsets of X, there exists open subsets U and V of X such that $A \subseteq U$, $B \subseteq V$, and $U \cap V = \emptyset$. By the above we know that there exists an element $(z, -z) \in A$ such that $(z, -z) \in \overline{V}$. Therefore, since $(z, -z) \in A \subseteq U$ and U is open, the properties of the closure implies that $U \cap V \neq \emptyset$. As this is a contradiction, the proof that X is not normal is complete.

6.3 Urysohn's Metrization Theorem

As the Baire Category Theorem provides additional proof that metric spaces are awesome (and an additional property topological spaces must have in order to be metrizable), it is about time we have a way of verifying that certain topological spaces are metrizable. Our main result of this section, Urysohn's Metrization Theorem (Theorem 6.3.1) does exactly that. In particular, Urysohn's Metrization Theorem can be thought of as the continuation of Proposition 5.1.27 as, after all, every metrizable topological space must be Hausdorff, first countable, and normal by Remark 6.1.3 so strengthening first countable to second countable allows us to weaken normal to regular. Furthermore, Urysohn's Metrization Theorem is exactly why Urysohn's Lemma (Theorem 5.2.1) is called a lemma.

Theorem 6.3.1 (Urysohn's Metrization Theorem). *If $(X, \mathcal{T})$ is a second countable, T_3 topological space, then $(X, \mathcal{T})$ is metrizable.*

Proof. By Proposition 5.1.27, $(X, \mathcal{T})$ is normal. Hence, as $(X, \mathcal{T})$ is T_3 and thus Hausdorff, $(X, \mathcal{T})$ is a Tychonoff space. Thus $(X, \mathcal{T})$ embeds into $\prod_{\alpha \in I}[0,1]$ equipped with the product topology by Theorem 5.3.20. However, if I is uncountable, $\prod_{\alpha \in I}[0,1]$ is probably not metrizable as, for example, $\prod_{\alpha \in \mathbb{R}} \mathbb{R}$ is not normal and thus not metrizable. However, if I is countable, Lemma 4.3.4 implies that the product topology on $\prod_{\alpha \in I}[0,1]$ (or even $\prod_{\alpha \in I} \mathbb{R}$) is metrizable. This would then complete the proof as every subspace of a metrizable topological space is metrizable. Hence we need only analyze the proof of Theorem 5.3.20 to see if we can take I to be countable.

Reviewing the part of the proof of Theorem 5.3.20 that we need, only the Embedding Theorem (Theorem 5.3.19) was used to construct the embedding. As the Embedding Theorem produces an embedding of $(X, \mathcal{T})$ into a product of copies of $\mathbb{R}$ indexed by a collection a functions $\{f_\alpha\}_{\alpha \in I} \subseteq \mathcal{C}(X, \mathbb{R})$ with the properties that for each $x_0 \in X$ and neighbourhood U of x_0 there exists an $\alpha_0 \in I$ such that $f_{\alpha_0}(x_0) \neq 0$ and $f(x) = 0$ for all $x \in X \setminus U$, we need only verify that we can take I to be countable in this context. This is where $(X, \mathcal{T})$ being second countable and Urysohn's Lemma (Theorem 5.2.1) comes into play.

To construct the desired countable collection of functions, recall that since $(X, \mathcal{T})$ is second countable that there exist a countable basis $\mathcal{B} = \{B_n\}_{n=1}^\infty$ of $(X, \mathcal{T})$. Let

$$I = \{(n, m) \in \mathbb{N}^2 \mid \overline{B_m} \subseteq B_n\} \subseteq \mathbb{N}^2.$$

Clearly I is a countable set. Moreover for each $(n, m) \in I$ we have by definition that $\overline{B_m}$ and $X \setminus B_n$ are disjoint closed subsets of $(X, \mathcal{T})$ by construction. Therefore, since $(X, \mathcal{T})$ is normal, Urysohn's Lemma implies that for each $(n, m) \in I$ there exists an $f_{(n,m)} \in \mathcal{C}(X, [0,1])$ such that $f_{(n,m)}(x) = 1$ for all $x \in \overline{B_m}$ and $f_{(n,m)}(x) = 0$ for all $x \in X \setminus B_n$.

We claim that $\{f_{(n,m)}\}_{(n,m)\in I}$ has the desired properties to invoke the Embedding Theorem to embed $(X, \mathcal{T})$ into a countable product of copies of $\mathbb{R}$ which is then metrizable by Lemma 4.3.4. To see this, let $x_0 \in X$ and U a neighbourhood of x_0 be arbitrary. Since $\mathcal{B}$ is a basis for $(X, \mathcal{T})$, there exists an $n \in \mathbb{N}$ such that

$$x_0 \in B_n \subseteq U.$$

Since $(X, \mathcal{T})$ is regular, Lemma 5.1.15 implies there exists a $V \in \mathcal{T}$ such that

$$x_0 \in V \subseteq \overline{V} \subseteq B_n \subseteq U.$$

Since $\mathcal{B}$ is a basis for $(X, \mathcal{T})$, there exists an $m \in \mathbb{N}$ such that

$$x_0 \in B_m \subseteq V.$$

Hence by Theorem 1.6.21 we have that

$$x_0 \in B_m \subseteq \overline{B_m} \subseteq \overline{V} \subseteq B_n \subseteq U.$$

Hence $(n, m) \in I$. Since $x_0 \in B_m$ we see by the definition of $f_{(n,m)}$ that $f_{(n,m)}(x_0) = 1$. Furthermore, since $B_n \subseteq U$ so that $X \setminus U \subseteq X \setminus B_n$, we see by the definition of $f_{(n,m)}$ that $f_{(n,m)}(x) = 0$ for all $x \in X \setminus U$. Hence, as x and U were arbitrary, we have verified the necessary assumptions for the Embedding Theorem hold for $\{f_{(n,m)}\}_{(n,m)\in I}$ thereby completing the proof. ∎

Remark 6.3.2. Of course, it is natural to ask, "How useful is Urysohn's Metrization Theorem (Theorem 6.3.1)?" After all, how often can one see a topological space $(X, \mathcal{T})$ is second countable, regular, and Hausdorff without easily seeing that there is a metric that induces $\mathcal{T}$? In terms of analysis, the answer is Urysohn's Metrization Theorem is not that useful. This is due to the fact that in analysis topologies are often defined via norms, seminorms, or in other natural ways (such as weak and weak* topologies) which either are obviously metrizable topologies or are not second countable.

However, one important use of Urysohn's Metrization Theorem comes from differential topology. One main focus of differential topology is topological manifolds, which are Hausdorff topological spaces $(X, \mathcal{T})$ for which there exists an $n \in \mathbb{N}$ such that every point $x \in X$ has a neighbourhood that is homeomorphic to $\mathbb{R}^n$ equipped with the Euclidean topology. One can easily see that topological manifolds are locally compact Hausdorff topological spaces by definitions. Hence as Corollary 5.4.3 implies that every locally compact Hausdorff topological space is Tychonoff and thus regular by Remark 5.3.11, every topological manifold is regular. Thus, provided a topological manifold is not too large and thus second countable, Urysohn's Metrization Theorem automatically implies that a metric exists.

Of course, it is natural to ask whether provided $(X, \mathcal{T})$ is normal, can the assumption that $(X, \mathcal{T})$ is second countrable in Urysohn's Metrization Theorem (Theorem 6.3.1) be reduced to the assumption that $(X, \mathcal{T})$ is first countable? After all, it is a lot easier to demonstrate a topological space is first countable than it is to show a topological space is second countable. Unfortunately, this is not the case. In particular, the lower limit topology is such an example. Indeed, to summarize, we have the following.

Corollary 6.3.3. *Let $\mathcal{T}_L$ be the lower limit topology on $\mathbb{R}$. Then $(\mathbb{R}, \mathcal{T}_L)$ is a first countable, separable T_4 topological space that is not second countable and thus not metrizable.*

Proof. The fact that $(\mathbb{R}, \mathcal{T}_L)$ is a first countable, separable T_4 topological space that is not second countable follows from combining Examples 1.5.37, 5.1.26, 6.1.8, 6.1.16, and 6.1.21. Moreover, $(\mathbb{R}, \mathcal{T}_L)$ is not metrizable by 6.1.23. ∎

6.4 Local Finiteness

Although Urysohn's Metrization Theorem (Theorem 6.3.1) is useful to some in order to show that certain topological spaces are metrizable, we have seen that we cannot simply weaken the assumption of second countability to first countability in the Urysohn's Metrization Theorem due to the lower limit topology on $\mathbb{R}$. Thus, as we know from Example 6.1.18 that the uniform metric topology on $\mathcal{F}(\mathbb{N}, \mathbb{R})$ is not second countable, we have no hope in directly improving Urysohn's Metrization Theorem to obtain necessary and sufficient conditions for a topological space to be metrizable. Thus, a new concept is needed in order to progress in this direction.

As second countability is too strong an assumption on a topological space in order to prove metrizability as there are metrizable topological spaces that are not second countable, we need a weaker assumption. Of course first countable is a possibility, but we want better control over the neighbourhoods of a point in a metrizable topological space. Thus we take some motivation from compact topological spaces and try to invoke some notion of finiteness of the neighbourhoods near a point. In particular, we want to make sure that we can find neighbourhoods that do not have too much overlap.

Definition 6.4.1. Let $(X, \mathcal{T})$ be a topological space. A set $\mathcal{A} \subseteq \mathcal{P}(X)$ is said to be *locally finite in* $(X, \mathcal{T})$ if every point in X has a neighbourhood that intersects only finitely many elements of $\mathcal{A}$.

Example 6.4.2. Given a topological space $(X, \mathcal{T})$, clearly every finite subset of $\mathcal{P}(X)$ is locally finite. Consequently, every topology on a finite number of points has a locally finite basis.

Example 6.4.3. Consider the canonical topology on $\mathbb{R}$. The set

$$\mathcal{A} = \{(n, n+2) \mid n \in \mathbb{Z}\}$$

is locally finite open cover of $\mathbb{R}$. However, the collection

$$\mathcal{B} = \left\{ \left(-\frac{1}{n}, \frac{1}{n}\right) \,\middle|\, n \in \mathbb{N} \right\}$$

is not locally finite in $\mathbb{R}$ as every neighbourhood of 0 intersects infinitely many elements of $\mathcal{B}$.

Locally finite sets of a topological space clearly behave well under certain operations. Furthermore, they are particularly well-behaved when it comes to taking closures of their unions as the following result demonstrates.

Lemma 6.4.4. *Let $\mathcal{A}$ be a locally finite set of subsets of a topological space $(X, \mathcal{T})$. Then the following hold:*

(1) Any subset $\mathcal{A}_0 \subseteq \mathcal{A}$ is locally finite in $(X, \mathcal{T})$.

(2) The collection $\mathcal{A}_c = \left\{ \overline{A} \mid A \in \mathcal{A} \right\}$ is locally finite in $(X, \mathcal{T})$.

(3) $\bigcup_{A \in \mathcal{A}} \overline{A} = \overline{\bigcup_{A \in \mathcal{A}} A}$.

Proof. Clearly (1) holds due to the definition of a locally finite set.

To see that (2) holds, let $x \in X$ be arbitrary. Since $\mathcal{A}$ is locally finite, there exists a neighbourhood U of x such that U intersects only finitely many elements of $\mathcal{A}$. However, if $A \in \mathcal{A}$ and $U \cap \overline{A} \neq \emptyset$, then $U \cap A \neq \emptyset$ by Theorem 1.6.21. Hence, as U intersects only finitely many elements of $\mathcal{A}$, U is a neighbourhood of x that intersects only finitely many elements of $\mathcal{A}_c$. Thus, as x was arbitrary, $\mathcal{A}_c$ is locally finite in $(X, \mathcal{T})$.

To see that (3) holds, first note that for all $A_1 \in \mathcal{A}$ that $A_1 \subseteq \bigcup_{A \in \mathcal{A}} A \subseteq \overline{\bigcup_{A \in \mathcal{A}} A}$. Therefore, since $\overline{\bigcup_{A \in \mathcal{A}} A}$ is closed, the definition of the closure of a set implies that $\overline{A_1} \subseteq \overline{\bigcup_{A \in \mathcal{A}} A}$ for all $A_1 \in \mathcal{A}$. Hence

$$\bigcup_{A \in \mathcal{A}} \overline{A} \subseteq \overline{\bigcup_{A \in \mathcal{A}} A}.$$

To see the reverse inclusion, suppose to the contrary that there exists an $a \in \overline{\bigcup_{A \in \mathcal{A}} A}$ such that $a \notin \bigcup_{A \in \mathcal{A}} \overline{A}$. Since $\mathcal{A}$ is locally finite in $(X, \mathcal{T})$, there exists a neighbourhood U of a such that U intersects only finitely many elements of $\mathcal{A}$. Let $A_1, A_2, \ldots, A_n \in \mathcal{A}$ be precisely the elements of $\mathcal{A}$ that have non-empty intersection with U. As $a \notin \bigcup_{A \in \mathcal{A}} \overline{A}$, for each $k \in \{1, \ldots, n\}$ there exists a $U_k \in \mathcal{T}$ such that $a \in U_k$ yet $U_k \cap A_k = \emptyset$. Let

$$V = U \cap \left(\bigcap_{k=1}^{n} U_k \right).$$

Clearly V is a neighbourhood of a such that $V \cap A_k = \emptyset$ for all $k \in \{1, \ldots, n\}$. However, since $U \cap A = \emptyset$ for all $A \in \mathcal{A} \setminus \{A_k\}_{k=1}^n$, $V \cap A = \emptyset$ for all $A \in \mathcal{A} \setminus \{A_k\}_{k=1}^n$. Hence $V \cap (\bigcup_{A \in \mathcal{A}} A) = \emptyset$ thereby contradicting the fact that V is a neighbourhood of a and $a \in \overline{\bigcup_{A \in \mathcal{A}} A}$. Thus, as we have obtained a contradiction, the reverse inclusion holds as desired. ∎

If our goal is to use locally finite sets to help describe a given topology, obtain a property that is weaker than second countable but stronger than first countable, and prove to metrizability of a topological space, then we would likely want locally finite collections of open sets that describe a basis; that is, we would like a locally finite basis of a topological space. However, given a metrizable topological space, it is unlikely that we will be able to find a locally finite basis since, as Example 6.4.3 shows, the requisite of having arbitrary small neighbourhoods around each point is an immediate obstacle to having a locally finite basis. However, as we can consider balls of a fixed radius at a given time and as we only need to consider rational radii, the following is not out of reach.

Definition 6.4.5. A set $\mathcal{A}$ of subsets of a topological space $(X, \mathcal{T})$ is said to be *σ-locally finite* if $\mathcal{A}$ is a countable union of locally finite subsets of $(X, \mathcal{T})$.

Remark 6.4.6. The term 'σ-locally finite' comes from the common notation in mathematics that 'σ' refers to 'countable sums'. Thus σ-locally finite is often called 'countably locally finite', but we prefer the σ-notation.

Example 6.4.7. Consider the canonical topology on $\mathbb{R}$. For each $n \in \mathbb{N}$, let

$$\mathcal{A}_n = \left\{ \left(\frac{m}{n}, \frac{m+2}{n} \right) \,\middle|\, m \in \mathbb{Z} \right\}.$$

It is elementary to see that $\mathcal{A}_n$ is a locally finite set of open subsets of $\mathbb{R}$ for every $n \in \mathbb{N}$. Hence $\mathcal{A} = \bigcup_{n=1}^{\infty} \mathcal{A}_n$ is a σ-locally finite subset of $\mathbb{R}$ consisting of open sets. Moreover, it is not difficult to see that $\mathcal{A}$ is a basis for the canonical topology on $\mathbb{R}$ due to the density of $\mathbb{Q}$ in $\mathbb{R}$ and the description of $\mathcal{A}$.

If a topological space having a σ-locally finite basis is the property we are searching for in order to obtain a better a necessary requirement for a topological space to be metrizable, we better show that every metrizable topological space as a σ-locally finite basis. However, as we will see through several proofs, it is important to be able to extract a σ-locally finite set from a collection of sets. This notion of refinement is formalized below.

Definition 6.4.8. Let $\mathcal{A}_1$ and $\mathcal{A}_2$ be sets of subsets of a topological space $(X, \mathcal{T})$. It is said that $\mathcal{A}_1$ is a *refinement* of $\mathcal{A}_2$ if for each $A_1 \in \mathcal{A}_1$ there exists an $A_2 \in \mathcal{A}_2$ such that $A_1 \subseteq A_2$. If every set in $\mathcal{A}_1$ is open in $(X, \mathcal{T})$, we call $\mathcal{A}_1$ an *open refinement* of $\mathcal{A}_2$. Similarly, if every set in $\mathcal{A}_1$ is closed in $(X, \mathcal{T})$, we call $\mathcal{A}_1$ a *closed refinement* of $\mathcal{A}_2$.

As an example of obtaining a σ-finite refinement, we demonstrate the following lemma. Note this is the best analogue of 'every open cover of a compact topological space has a finite subcover' that we can possibly obtain for a metrizable topological space. Therefore, as compactness is such a nice property, we are perhaps on the right track to study metrizable topological spaces.

Lemma 6.4.9. *Let $(X,\mathcal{T})$ be a metrizable space and let $\mathcal{A}$ be an open cover of $(X,\mathcal{T})$. Then there exists an open refinement $\mathcal{A}'$ of $\mathcal{A}$ that is σ-locally finite and covers $(X,\mathcal{T})$.*

Proof. To see the result, let $\mathcal{A} = \{U_\alpha\}_{\alpha\in I}$ be an arbitrary open cover of $(X,\mathcal{T})$. By the Well-Ordering Theorem (Theorem A.6.3), there exists a well-ordering $\leq$ on I.

Since $(X,\mathcal{T})$ is metrizable, there exists a metric $d : X \times X \to [0,\infty)$ that induces $\mathcal{T}$. To construct one portion of the σ-locally finite refinement that covers $(X,\mathcal{T})$, for each $n \in \mathbb{N}$ we will use d to first shrink each element of $\mathcal{A}$ by $\frac{1}{n}$, then we will disjointify the resulting sets resulting in sets with positive separation, and then we will expand these sets slightly in order to obtain a locally finite open refinement of $\mathcal{U}$.

Fix a natural number $n \in \mathbb{N}$. For each $\alpha \in I$, let

$$S_n(\alpha) = \left\{ x \in X \,\middle|\, B_d\left(x, \frac{1}{n}\right) \subseteq U_\alpha \right\}$$

and let

$$D_n(\alpha) = S_n(\alpha) \setminus \bigcup_{\substack{\beta \in I\setminus\{\alpha\} \\ \beta \leq \alpha}} U_\beta.$$

We claim that if $\alpha, \beta \in I$ are such that $\alpha \neq \beta$, then

$$\mathrm{dist}(D_n(\alpha), D_n(\beta)) = \inf(\{d(a,b) \mid a \in D_n(\alpha), b \in D_n(\beta)\}) \geq \frac{1}{n}.$$

To see this, let $\alpha, \beta \in I$ such that $\alpha \neq \beta$ be arbitrary. As $\leq$ is a well-ordering on I, by interchanging α and β if necessary, we may assume that $\beta \leq \alpha$. Let $a \in D_n(\alpha)$ and $b \in D_n(\beta)$ be arbitrary. Hence $b \in S_n(\beta)$ so $B_d\left(b, \frac{1}{n}\right) \subseteq U_\beta$. Since $D_n(\alpha) \subseteq X \setminus U_\beta$ by construction, $a \notin B_d\left(b, \frac{1}{n}\right)$ so $d(a,b) \geq \frac{1}{n}$. Since a and b were arbitrary, $\mathrm{dist}(D_n(\alpha), D_n(\beta)) \geq \frac{1}{n}$ as desired.

Unfortunately, $\{D_n(\alpha)\}_{\alpha\in I}$ are not the droids... I mean sets we are looking for (as it is possible to check that $S_n(\alpha)$ is closed and thus $D_n(\alpha)$ is also closed). To rectify the situation, for all $\alpha \in I$ let

$$E_n(\alpha) = \left\{ x \in X \,\middle|\, \mathrm{dist}(x, D_n(\alpha)) < \frac{1}{3n} \right\} = \bigcup_{a \in D_n(\alpha)} B_d\left(a, \frac{1}{3n}\right) \in \mathcal{T}.$$

Thus clearly $E_n(\alpha)$ is open for all $\alpha \in I$. We claim that if $\alpha, \beta \in I$ are such that $\alpha \neq \beta$, then

$$\mathrm{dist}(E_n(\alpha), E_n(\beta)) \geq \frac{1}{3n}.$$

To see this, let $\alpha, \beta \in I$ such that $\alpha \neq \beta$ be arbitrary. Let $a \in E_n(\alpha)$ and $b \in E_n(\beta)$ be arbitrary. By the definition of $E_n(\alpha)$ and $E_n(\beta)$ there exists $a' \in D_n(\alpha)$ and $b' \in D_n(\beta)$ such that

$$d(a, a') < \frac{1}{3n} \qquad \text{and} \qquad d(b, b') < \frac{1}{3n}.$$

Hence, as $\mathrm{dist}(D_n(\alpha), D_n(\beta)) \geq \frac{1}{n}$, we obtain that

$$\frac{1}{n} \leq d(a', b') \leq d(a', a) + d(a, b) + d(b, b') < \frac{2}{3n} + d(a, b)$$

so that $d(a, b) \geq \frac{1}{3n}$. Hence, as a and b were arbitrary, the claim holds.

For each $n \in \mathbb{N}$, let $\mathcal{A}_n = \{E_n(\alpha) \mid \alpha \in I\}$ and let $\mathcal{A}' = \bigcup_{n\in\mathbb{N}} \mathcal{A}_n$. We claim that $\mathcal{A}'$ is the set we are looking for. To begin to see this, first notice that $E_n(\alpha) \in \mathcal{T}$ for all $\alpha \in I$ and $n \in \mathbb{N}$ so $\mathcal{A}$ consists of open sets. To see that $\mathcal{A}'$ is a refinement of $\mathcal{A}$, consider an arbitrary $E_n(\alpha) \in \mathcal{A}'$ for some $n \in \mathbb{N}$ and $\alpha \in I$. By construction

$$E_n(\alpha) \subseteq \bigcup_{a \in S_n(\alpha)} B_d\left(a, \frac{1}{3n}\right) \subseteq U_\alpha$$

by the definition of $E_n(\alpha)$, $D_n(\alpha)$, and $S_n(\alpha)$ (i.e. $E_n(\alpha)$ is a union of open balls of radius $\frac{1}{3n}$ centred at the elements of $D_n(\alpha)$ which are elements of $S_n(\alpha)$ and thus have the property that the open ball of radius $\frac{1}{n}$ centred at them is contained in U_α). Hence $\mathcal{A}'$ is an open refinement of $\mathcal{A}$.

To see that $\mathcal{A}'$ is σ-locally finite, it suffices to show that $\mathcal{A}_n$ is locally finite for all $n \in \mathbb{N}$. To see that $\mathcal{A}_n$ is locally finite, let $x \in X$ be arbitrary. Since the open ball of radius $\frac{1}{6n}$ can intersect at most one $E_n(\alpha)$ as $\mathrm{dist}(E_n(\alpha), E_n(\beta)) \geq \frac{1}{3n}$ for all $\alpha, \beta \in I$ with $\alpha \neq \beta$, x has a neighbourhood that intersects a finite number of elements of $\mathcal{A}_n$. Hence, as x was arbitrary, $\mathcal{A}_n$ is locally finite for all $n \in \mathbb{N}$ and thus $\mathcal{A}'$ is σ-locally finite.

Finally, it remains only to show that $\mathcal{A}'$ is a cover of $(X, \mathcal{T})$. To see this, let $x \in X$ be arbitrary. Since $\mathcal{A}$ is a cover of $(X, \mathcal{T})$, there exists an $\alpha \in I$ such that $x \in U_\alpha$. Since $\leq$ is a well-ordering on I and as

$$J = \{\alpha \in I \mid x \in U_\alpha\} \neq \emptyset,$$

there exists a least element of J. Hence there exists an $\alpha_x \in I$ such that $x \in U_{\alpha_x}$ and $x \notin U_\beta$ for all $\beta < \alpha_x$.

Since $x \in U_{\alpha_x}$ and U_{α_x} is open, there exists an $n_x \in \mathbb{N}$ such that $B_d\left(x, \frac{1}{n_x}\right) \subseteq U_{\alpha_x}$. Hence $x \in S_{n_x}(\alpha_x)$ Therefore, as $x \notin U_\beta$ for all $\beta < \alpha_x$, we have that $x \in D_{n_x}(\alpha_x)$ by definition. Thus $x \in E_{n_x}(\alpha_x)$ by definition. Therefore, as $E_{n_x}(\alpha_x) \in \mathcal{A}'$, we have that

$$x \in \bigcup_{A \in \mathcal{A}'} A.$$

Therefore, as $x \in X$ was arbitrary, $\mathcal{A}'$ is a cover of $(X, \mathcal{T})$ thereby completing the proof. ∎

Using Lemma 6.4.9, we can actually prove that metrizable spaces have nice bases thereby showing that having a σ-locally finite basis is a requirement of being metrizable.

Corollary 6.4.10. *Every metrizable topological space has a σ-locally finite basis.*

Proof. Let $(X, \mathcal{T})$ be a metrizable topological space and let d be a metric that induces $\mathcal{T}$. For every $n \in \mathbb{N}$, let

$$\mathcal{A}_n = \left\{ B_d\left(x, \frac{1}{n}\right) \,\middle|\, x \in X \right\}.$$

Since $\mathcal{A}_n$ is clearly an open cover of $(X, \mathcal{T})$, Lemma 6.4.9 implies that there exists an open refinement $\mathcal{B}_n$ of $\mathcal{A}_n$ that is σ-locally finite and covers $(X, \mathcal{T})$. Since $\mathcal{B}_n$ is a refinement of $\mathcal{A}_n$, notice if $B \in \mathcal{B}_n$ then $B \subseteq B_d\left(x, \frac{1}{n}\right)$ for some $x \in X$ and thus $\mathrm{diam}(B) \leq \frac{2}{n}$.

Let $\mathcal{B} = \bigcup_{n=1}^{\infty} \mathcal{B}_n$. We claim that $\mathcal{B}$ is a σ-locally finite basis of $(X, \mathcal{T})$. To see this, note $\mathcal{B}$ is clearly σ-locally finite being the countable union of σ-locally finite subset of $(X, \mathcal{T})$. To see that $\mathcal{B}$ is a basis for $(X, \mathcal{T})$, let $x \in X$ and $\epsilon > 0$ be arbitrary. Choose $n \in \mathbb{N}$ such that $\frac{1}{n} < \frac{\epsilon}{2}$. Since $\mathcal{B}_n$ covers $(X, \mathcal{T})$, there exists a $B \in \mathcal{B}_n \subseteq \mathcal{B}$ such that $x \in B$. Therefore, since $\mathrm{diam}(B) \leq \frac{2}{n}$, it must be the case that $x \in B \subseteq B_d\left[x, \frac{2}{n}\right] \subseteq B_d(x, \epsilon)$. Therefore, since $x \in X$ and $\epsilon > 0$ were arbitrary and since d induces $\mathcal{T}$, $\mathcal{B}$ is a σ-locally finite basis of $(X, \mathcal{T})$ as desired. ∎

6.5 The Nagata-Smirnov Metrization Theorem

As Corollary 6.4.10 shows that every metrizable topological space must have a σ-locally finite basis, it is natural to ask whether we can obtain a converse to Corollary 6.4.10. Of course we must add the condition that the topological space under investigation is normal as every metrizable topological space is normal by Remark 6.1.3. However, as verifying a topological space is normal is often difficult, we desire to replace the assumption of being normal with being regular as Lemma 5.1.15 provides a simpler method for verifying a topological space is regular.

Thus the main goal of this section is to prove the Naga-Smirnov Metrization Theorem (Theorem 6.5.5) showing that every regular topological space with a σ-locally finite basis is metrizable. To proceed, we begin by developing additional properties of regular topological spaces with σ-locally finite bases.

Lemma 6.5.1. *Let $(X, \mathcal{T})$ be a regular topological space with a σ-locally finite basis. If $V \in \mathcal{T}$, then there exists $\{U_n\}_{n=1}^\infty \subseteq \mathcal{T}$ such that*

$$V = \bigcup_{n=1}^\infty U_n = \bigcup_{n=1}^\infty \overline{U_n}.$$

Proof. By assumption there exists a basis $\mathcal{B}$ of $(X, \mathcal{T})$ such that $\mathcal{B} = \bigcup_{n=1}^\infty \mathcal{B}_n$ where each $\mathcal{B}_n$ is a locally finite subset of $(X, \mathcal{T})$. For each $n \in \mathbb{N}$, let

$$\mathcal{A}_n = \{B \in \mathcal{B}_n \mid \overline{B} \subseteq V\}.$$

Since clearly $\mathcal{A}_n \subseteq \mathcal{B}_n$, $\mathcal{A}_n$ is a locally finite subset of $(X, \mathcal{T})$ for all $n \in \mathbb{N}$.

For each $n \in \mathbb{N}$, let

$$U_n = \bigcup_{B \in \mathcal{A}_n} B.$$

Clearly $U_n \in \mathcal{T}$ for all $n \in \mathbb{N}$. Furthermore, by Lemma 6.4.4 and the definition of $\mathcal{A}_n$, we see that

$$U_n \subseteq \overline{U_n} = \bigcup_{B \in \mathcal{A}_n} \overline{B} \subseteq V$$

for all $n \in \mathbb{N}$. Hence

$$\bigcup_{n=1}^\infty U_n \subseteq \bigcup_{n=1}^\infty \overline{U_n} \subseteq V.$$

To see the reverse inclusion, let $x \in V$ be arbitrary. Since $(X, \mathcal{T})$ is regular and $\mathcal{B}$ is a basis for $(X, \mathcal{T})$, Lemma 5.1.15 implies that there exists a $B \in \mathcal{B}$ such that

$$x \in B \subseteq \overline{B} \subseteq V.$$

Hence $B \in \mathcal{A}_n$ for some $n \in \mathbb{N}$, so $x \in U_n$ for some $n \in \mathbb{N}$, and thus $x \in \bigcup_{n=1}^\infty U_n$. Therefore, as $x \in V$ was arbitrary, the result follows. ∎

If our goal is to prove that every regular topological space with a σ-locally finite basis is metrizable, we must be able to prove such topological spaces are normal. The following lemma does just this. Note the idea of the proof of this lemma is to follow the ideas of Proposition 5.1.27. Indeed we can repeat the proof once we bypass the second countability assumption using our σ-locally finite basis.

Lemma 6.5.2. *Let $(X, \mathcal{T})$ be a regular topological space with a σ-locally finite basis. Then $(X, \mathcal{T})$ is normal.*

Proof. Let A and B be arbitrary closed subsets of (X, T) such that $A \cap B = \emptyset$. Since $X \setminus B$ and $X \setminus A$ are open sets in $(X, \mathcal{T})$, Lemma 6.5.1 implies that there exists $\{U_n\}_{n=1}^{\infty}, \{V_n\}_{n=1}^{\infty} \subseteq \mathcal{T}$ such that

$$X \setminus B = \bigcup_{n=1}^{\infty} U_n = \bigcup_{n=1}^{\infty} \overline{U_n} \qquad \text{and} \qquad X \setminus A = \bigcup_{n=1}^{\infty} V_n = \bigcup_{n=1}^{\infty} \overline{V_n}.$$

By taking complements, we see that $\{U_n\}_{n=1}^{\infty}$ is an open cover of A such that $B \cap \overline{U_n} = \emptyset$ for all $n \in \mathbb{N}$, and $\{V_n\}_{n=1}^{\infty}$ is an open cover of B such that $A \cap \overline{V_n} = \emptyset$ for all $n \in \mathbb{N}$. As this is precisely what was originally constructed in Proposition 5.1.27 to obtain the desired open subsets of $(X, \mathcal{T})$ in that context, we need only repeat the proof to ensure that $\{U_n\}_{n=1}^{\infty}$ and $\{V_n\}_{n=1}^{\infty}$ are pairwise disjoint.

For every $n \in \mathbb{N}$, let

$$U_n' = U_n \setminus \left(\bigcup_{k=1}^{n} \overline{V}_k \right) \qquad \text{and} \qquad V_n' = V_n \setminus \left(\bigcup_{k=1}^{n} \overline{U}_k \right).$$

Clearly $\{\overline{U_n}\}_{n\geq 1}$ and $\{\overline{V_n}\}_{n\geq 1}$ are collections of closed subsets of $(X, \mathcal{T})$ so $\{\bigcup_{k=1}^{n} \overline{U}_k\}_{n\geq 1}$ and $\{\bigcup_{k=1}^{n} \overline{V}_k\}_{n\geq 1}$ are closed subsets of $(X, \mathcal{T})$. Therefore, since $\{U_n\}_{n\geq 1}$ and $\{V_n\}_{n\geq 1}$ are collections of open subsets of $(X, \mathcal{T})$ and since $D \setminus E = D \cap (X \setminus E)$ for all $D, E \subseteq X$, we see that $\{U_n'\}_{n\geq 1}$ and $\{V_n'\}_{n\geq 1}$ are collections of open subsets of $(X, \mathcal{T})$.

Let

$$U = \bigcup_{n=1}^{\infty} U_n' \qquad \text{and} \qquad V = \bigcup_{n=1}^{\infty} V_n',$$

which are open subsets of $(X, \mathcal{T})$ by the above discussions. We claim that $A \subseteq U$, $B \subseteq V$, and $U \cap V = \emptyset$. To see this, first notice since $\overline{V_n} \cap A = \emptyset$ for all $n \in \mathbb{N}$ that $U_n' \cap A = U_n \cap A$ for all $n \in \mathbb{N}$. Hence, since $A \subseteq \bigcup_{n=1}^{\infty} U_n$ we obtain that $A \subseteq U$. Furthermore, similar arguments show that $B \subseteq V$. Finally, suppose to the contrary that $U \cap V \neq \emptyset$ so that there exists an $x \in U \cap V$. By the definition of U and V, there must exists $n, m \in \mathbb{N}$ so that $x \in U_n'$ and $x \in V_m'$. If $n \geq m$, then $x \in V_m'$ implies that $x \in V_m$ and $x \in U_n'$ implies that

$$x \in U_n \setminus \left(\bigcup_{k=1}^{n} \overline{V}_k \right) \subseteq U_n \setminus V_m,$$

which is an obvious contradiction. Similarly, if $m \geq n$ then $x \in U_n$ and $x \in V_m \setminus U_n$ which is also a contradiction. Hence is must be the case that $U \cap V = \emptyset$.

Therefore, since A and B were arbitrary, $(X, \mathcal{T})$ is normal as desired. ∎

In fact, regular topological spaces with σ-locally finite bases share another topological property with metric spaces. Indeed recall Lemma 6.2.20 showed that closed subset of a metric spaces is G_δ and the following shows the same is true for regular topological spaces with σ-locally finite bases.

Lemma 6.5.3. *Let $(X, \mathcal{T})$ be a regular topological space with a σ-locally finite basis. Then every closed subset of $(X, \mathcal{T})$ is a G_δ subset of $(X, \mathcal{T})$.*

Proof. Let F be an arbitrary closed subset of $(X, \mathcal{T})$. Since $X \setminus F$ is open, Lemma 6.5.1 implies there $\{V_n\}_{n=1}^\infty \subseteq \mathcal{T}$ such that

$$X \setminus F = \bigcup_{n=1}^\infty V_n = \bigcup_{n=1}^\infty \overline{V_n}.$$

For each $n \in \mathbb{N}$, let $U_n = X \setminus \overline{V_n} \in \mathcal{T}$. We claim that $F = \bigcap_{n=1}^\infty U_n$ thereby showing that F is a G_δ set. Indeed this follows directly from the above set equality due to De Morgan's laws. Therefore, as F was arbitrary, every closed subset of $(X, \mathcal{T})$ is G_δ. ∎

Going back to our main goal of proving metrizability, in order to prove that every regular topological space with a σ-locally finite basis is metrizable, we either need to explicitly construct a metric (which is likely a daunting task) or we must embed our topological space into a metrizable space. Using Lemma 6.5.3 we can obtain functions that look very similar to those required in the assumptions of the Embedding Theorem (Theorem 5.3.19).

Lemma 6.5.4. *Let $(X, \mathcal{T})$ be a normal topological space and let A be a closed G_δ subset of $(X, \mathcal{T})$. There exists an $f \in \mathcal{C}(X, [0, 1])$ such that $f(a) = 0$ for all $a \in A$ and $f(x) > 0$ for all $x \in X \setminus A$.*

Proof. Let A be an arbitrary closed G_δ subset of $(X, \mathcal{T})$. Since A is G_δ, there exists a countable set $\{U_n\}_{n=1}^\infty$ of open subsets of $(X, \mathcal{T})$ such that $A = \bigcap_{n=1}^\infty U_n$. Since $A \subseteq U_n$ for all $n \in \mathbb{N}$, A and $X \setminus U_n$ are disjoint closed subsets of $(X, \mathcal{T})$ for all $n \in \mathbb{N}$. Therefore, since $(X, \mathcal{T})$ is normal, Urysohn's Lemma (Theorem 5.2.1) implies that for all $n \in \mathbb{N}$ there exists an $f_n \in \mathcal{C}(X, [0, 1])$ such that $f_n(a) = 0$ for all $a \in A$ and $f_n(x) = 1$ for all $x \in X \setminus U_n$.

Since $(\mathcal{C}_b(X, \mathbb{R}), \|\cdot\|_\infty)$ is a Banach space by Theorem 4.2.14 and since

$$\sum_{n=1}^\infty \left\| \frac{1}{2^n} f_n \right\|_\infty \leq \sum_{n=1}^\infty \frac{1}{2^n} = 1,$$

Theorem 4.1.17 implies the function $f : X \to \mathbb{R}$ defined by

$$f(x) = \sum_{n=1}^\infty \frac{1}{2^n} f_n(x)$$

for all $x \in X$ is an element of $\mathcal{C}_b(X, \mathbb{R})$. Furthermore, the above norm estimate implies that $f(x) \leq 1$ for all $x \in X$ and clearly $f(x) \geq 0$ for all $x \in X$ by construction.

To see that f is the function we are looking for, first notice for all $a \in A$ that

$$f(a) = \sum_{n=1}^{\infty} \frac{1}{2^n} f_n(a) = 0$$

by construction. To see that f has the second desired property, let $x \in X \setminus A$ be arbitrary. Since $A = \bigcap_{n=1}^{\infty} U_n$, $x \in \bigcup_{n=1}^{\infty} X \setminus U_n$. Thus there exists an $n_x \in \mathbb{N}$ such that $x \in X \setminus U_{n_x}$. Therefore $f_{n_x}(x) = 1$ by construction so that

$$f(x) = \sum_{n=1}^{\infty} \frac{1}{2^n} f_n(x) \geq \frac{1}{2^{n_x}} f_{n_x}(x) > 0$$

as desired. Hence, as $x \in X \setminus A$ was arbitrary, f has the desired properties. ∎

Now we have all the tools to prove our main result. The motivation and idea of the proof are described as we begin the proof below.

Theorem 6.5.5 (Nagata-Smirnov Metrization Theorem). *A topological space $(X, \mathcal{T})$ is metrizable if and only if $(X, \mathcal{T})$ is T_0, regular, and has a σ-locally finite basis.*

Proof. First, suppose that $(X, \mathcal{T})$ is metrizable. Hence $(X, \mathcal{T})$ is a T_4 space by Theorem 5.1.23 and thus regular. Moreover, $(X, \mathcal{T})$ has a σ-locally finite basis by Corollary 6.4.10. Hence one direction of the proof is complete.

Conversely, suppose that $(X, \mathcal{T})$ is regular and has a σ-locally finite basis. Our goal is to use the same idea as the Embedding Theorem (Theorem 5.3.19) to embed $(X, \mathcal{T})$ into $(\prod_{\alpha \in I}[0, 1], \mathcal{T}_m)$ where $\mathcal{T}_m$ is the uniform metric topology. This will complete the proof as a subspace of a metric space is a metric space.

By Lemma 6.5.2 and Lemma 6.5.3 we know that $(X, \mathcal{T})$ is normal and that every closed subset of $(X, \mathcal{T})$ is G_δ. Hence Lemma 6.5.4 implies that for each closed subset A of $(X, \mathcal{T})$ that there exists an $f \in \mathcal{C}(X, [0, 1])$ such that $f(a) = 0$ for all $a \in A$ and $f(x) > 0$ for all $x \in X \setminus A$. These are precisely functions we desire to use to build our embedding into a product of metric spaces equipped with the uniform topology. Of course, we desire to use our σ-locally finite basis in order to have greater control over the functions we are using.

Since $(X, \mathcal{T})$ has a σ-locally finite basis there exists a basis $\mathcal{B}$ of $(X, \mathcal{T})$ such that $\mathcal{B} = \bigcup_{n=1}^{\infty} \mathcal{B}_n$ where each $\mathcal{B}_n$ is a locally finite subset of $(X, \mathcal{T})$. By the previous paragraph, for each $n \in \mathbb{N}$ and each $B \in \mathcal{B}_n$ we obtain as $X \setminus B$ is closed that there exists an $f_{(n,B)} \in \mathcal{C}(X, [0, 1])$ such that $f_{(n,B)}(x) > 0$ for all $x \in B$ and $f_{(n,B)}(x) = 0$ for all $x \in X \setminus B$. By scaling $f_{(n,B)}$ by $\frac{1}{n}$ if necessary, we may assume that

$$f_{(n,B)} \in \mathcal{C}\left(X, \left[0, \frac{1}{n}\right]\right)$$

for all $B \in \mathcal{B}_n$ and $n \in \mathbb{N}$.

Let

$$I = \{(n, B) \mid n \in \mathbb{N}, B \in \mathcal{B}_n\} \quad \text{and} \quad Y = \prod_{(n,B)\in I} \left[0, \frac{1}{n}\right]$$

and define $F : X \to Y$ by

$$F(x) = (f_{n,B}(x))_{(n,B)\in I}$$

for all $x \in X$. We claim that $\{f_{(n,B)}\}_{(n,B)\in I}$ satisfy the assumptions of the Embedding Theorem (Theorem 5.3.19) so that F is an embedding from $(X, \mathcal{T})$ into Y equipped with the product topology. Of course, this is not the desired topology we want on Y as we wanted the uniform metric topology. However, since the uniform metric topology is finer than the product topology by Lemma 4.2.4, this will immediately imply that F is injective and if $U \in \mathcal{T}$ then $F(U)$ is open with respect to the product and thus with respect to the uniform metric topology.

To see that $\{f_{(n,B)}\}_{(n,B)\in I}$ satisfy the assumptions of the Embedding Theorem, let $x_0 \in X$ and U a neighbourhood of x_0 be arbitrary. Since $\mathcal{B}$ is a basis of $(X, \mathcal{T})$, there exists an $n \in \mathbb{N}$ and a $B \in \mathcal{B}_n$ such that $x_0 \in B \subseteq U$. Hence by the definition of $f_{(n,B)}$, $f_{(n,B)}(x_0) > 0$ and $f_{(n,B)}(x) = 0$ for all $x \in X \setminus U$ as $X \setminus U \subseteq X \setminus B$. Therefore, as x_0 and U were arbitrary, the claim is complete. Hence, by previous discussions, the proof will be complete once it has been demonstrated that F is continuous from $(X, \mathcal{T})$ into Y equipped with the uniform metric topology.

Let $d : Y \times Y \to [0, \infty)$ denote the uniform metric on Y. Therefore, due to the explicit descriptions of Y and the uniform metric on Y, we know that

$$d(g, h) = \sup_{(n,B)\in I} |g((n, B)) - h((n, B))| \leq 1$$

for all $g, h \in Y$ (i.e. taking the minimum with 1 is not required as every term in the sup is at most 1).

To see that F is continuous from $(X, \mathcal{T})$ into (Y, d), let $x_0 \in X$ and $\epsilon > 0$ be arbitrary. To complete the proof, it suffices to construct a neighbourhood V of x_0 such that if $x \in V$, then $d(F(x), F(x_0)) \leq \epsilon$.

Choose $N \in \mathbb{N}$ such that $\frac{1}{N} \leq \epsilon$. Since $0 \leq f_{(n,B)}(x) \leq \frac{1}{n}$ for all $n \in \mathbb{N}$, we see that for all $n \geq N$ and for all $B \in \mathcal{B}_n$ that

$$|f_{(n,B)}(x) - f_{(n,B)}(x_0)| \leq \frac{1}{N} \leq \epsilon$$

for all $x \in X$. Thus we will need only consider $(n, B) \in I$ such that $n < N$ when analyzing the values of $d(F(x), F(x_0))$ for $x \in X$.

For a fixed $n \in \mathbb{N}$, we recall that $\mathcal{B}_n$ is locally finite. Hence there exists a neighbourhood U_n of x_0 such that U_n intersects a finite number of elements of $\mathcal{B}_n$. Let

$$\{B_{n,k}\}_{k=1}^m = \{B \in \mathcal{B}_n \mid B \cap U_n \neq \emptyset\}.$$

Since $f_{(n,B_{n,k})} \in \mathcal{C}(X)$ for all $k \in \{1, \ldots, m\}$, there exists a neighbourhood $V_{n,k}$ of x_0 such that

$$|f_{(n,B_{n,k})}(x) - f_{(n,B_{n,k})}(x_0)| \leq \epsilon$$

for all $x \in V_{n,k}$. Let

$$V_n = U_n \cap \left(\bigcap_{k=1}^m V_{n,k} \right)$$

so that V_n is a neighbourhood of x_0. We claim that

$$|f_{(n,B)}(x) - f_{(n,B)}(x_0)| \leq \epsilon$$

for all $x \in V_n$ and $B \in \mathcal{B}_n$. Indeed if $B \in \{B_{n,k}\}_{k=1}^m$ then the claim follow from the definitions of V_n and $V_{n,k}$. Otherwise $B \notin \{B_{n,k}\}_{k=1}^m$. Hence $B \cap U_n = \emptyset$ so that $U_n \subseteq X \setminus B$ and thus

$$f_{(n,B)}(x) = f_{(n,B)}(x_0) = 0$$

for all $x \in V_n$ as $x_0 \in V_n$. Hence the claim follows.

Let

$$V = \bigcap_{n=1}^N V_n,$$

which is an open neighbourhood by construction. Furthermore, by construction, if $n \in \mathbb{N}$ is such that $n \leq N$, if $B \in \mathcal{B}_n$, and if $x \in V$ then

$$|f_{(n,B)}(x) - f_{(n,B)}(x_0)| \leq \epsilon.$$

Hence as

$$|f_{(n,B)}(x) - f_{(n,B)}(x_0)| \leq \frac{1}{N} \leq \epsilon$$

for all $n \geq N$, $B \in \mathcal{B}_n$, and $x \in X$, we obtain by the definition of the uniform metric that

$$d(F(x), F(x_0)) \leq \epsilon$$

for all $x \in V$ thereby completing the proof. ∎

6.6 Paracompactness

Of course the Nagata-Smirnov Metrization Theorem (Theorem 6.5.5) has one limitation in that one needs to verify that a topological space has a σ-locally finite basis, which is often not an easy task. As the idea of a σ-locally finite basis was motivated by trying to weaken second countability via an idea similar to compactness, in this section we will introduce a generalization of compactness called paracompactness. It turns out that paracompactness is particularly useful for applications in topology and differential geometry. However, our only goal will be to relate paracompactness to the existence of σ-locally finite bases.

Onto the definition.

Definition 6.6.1. A topological space $(X, \mathcal{T})$ is said to be *paracompact* if every open cover of $(X, \mathcal{T})$ has a locally finite open refinement that covers $(X, \mathcal{T})$.

Clearly every compact topological space is paracompact as every finite subcover of an open cover $\mathcal{U}$ is clearly a locally finite open refinement of $\mathcal{U}$ that covers the space. However, additional topological spaces are paracompact.

Example 6.6.2. Let $(X, \mathcal{T})$ be a Hausdorff topological space such that there exists a sequence $(V_n)_{n\geq 1}$ of open subsets of $(X, \mathcal{T})$ such that $K_n = \overline{V_n}$ is compact for all $n \in \mathbb{N}$, $K_n \subseteq V_{n+1}$ for all $n \in \mathbb{N}$, and $X = \bigcup_{n=1}^{\infty} K_n$ (e.g. $X = \mathbb{R}^n$ equipped with the Euclidean topology). Then $(X, \mathcal{T})$ will be paracompact. To see this, let $\mathcal{U}$ be an arbitrary open cover of X. For each $n \in \mathbb{N}$, we note since K_n is compact and since $\mathcal{U}$ is an open cover of $(X, \mathcal{T})$ and therefore K_n that there exists a finite subset $\mathcal{U}_n \subseteq \mathcal{U}$ such that $\mathcal{U}_n$ covers K_n. Let

$$\mathcal{A}_n = \{U \cap (V_n \setminus K_{n-1}) \mid U \in \mathcal{U}_n\}$$

where $K_0 = \emptyset$. Since K_n is compact in $(X, \mathcal{T})$ and thus closed by Theorem 3.1.13, $\mathcal{A}_n \subseteq \mathcal{T}$.

Let $\mathcal{A} = \bigcup_{n=1}^{\infty} \mathcal{A}_n$. We claim that $\mathcal{A}$ does the trick. To see this, note since $\mathcal{A}_n$ is an open refinement of $\mathcal{U}_n \subseteq \mathcal{U}$ for all $n \in \mathbb{N}$ that $\mathcal{A}$ is an open refinement of $\mathcal{U}$. Moreover, since $\mathcal{U}_n$ covers K_n, we see that $\mathcal{A}_n$ covers $K_n \setminus K_{n-1}$ for all $n \in \mathbb{N}$. Therefore, since $X = \bigcup_{n=1}^{\infty} K_n$, we obtain that $\mathcal{A}$ covers $(X, \mathcal{T})$.

Finally, to see that $\mathcal{A}$ is locally finite, let $x \in X$ be arbitrary. Hence $x \in V_n$ for some $n \in \mathbb{N}$. Since no element of $\mathcal{A}_m$ intersects V_n for all $m > n$, we see that the set of elements of $\mathcal{A}$ that intersect V_n is contained in $\bigcup_{k=1}^{n} \mathcal{A}_k$ which is finite since $\mathcal{U}_k$ and thus $\mathcal{A}_k$ is finite for all $k \in \mathbb{N}$. Therefore, as x was arbitrary, $\mathcal{A}$ is locally finite as desired.

Since paracompactness is very similar to compactness in that locally finite replaces finite, we can prove some results for paracompact topological spaces that we had only for compact topological spaces. For example, the following result is similar to Lemma 3.1.12 and is proved in a similar manner.

Lemma 6.6.3. *Every paracompact Hausdorff topological space is regular.*

Proof. Let $(X, \mathcal{T})$ be a paracompact Hausdorff space. To see that $(X, \mathcal{T})$ is regular, suppose A is a non-empty closed subset of $(X, \mathcal{T})$ and $x \in X \setminus A$. Since $(X, \mathcal{T})$ is Hausdorff, for each $a \in A$ there exists $U_a, V_a \in \mathcal{T}$ such that $a \in U_a$, $x \in V_a$, and $U_a \cap V_a = \emptyset$. Thus $\{X \setminus A\} \cup \{U_a\}_{a \in A}$ is an open cover of $(X, \mathcal{T})$. Therefore, since $(X, \mathcal{T})$ is paracompact, there exists a locally finite open refinement $\mathcal{U}$ of $\{X \setminus A\} \cup \{U_a\}_{a \in A}$ that covers $(X, \mathcal{T})$.

Let
$$\mathcal{C} = \{C \in \mathcal{U} \mid C \cap A \neq \emptyset\}$$
and let
$$U = \bigcup_{C \in \mathcal{C}} C.$$
Clearly U is an open subset of $(X, \mathcal{T})$. Moreover, since $\mathcal{U}$ was a cover of $(X, \mathcal{T})$, $\mathcal{C}$ is an open cover of A and thus $A \subseteq U$. Furthermore, since $\mathcal{U}$ is locally finite, $\mathcal{C}$ is locally finite and thus
$$\overline{U} = \bigcup_{C \in \mathcal{C}} \overline{C}$$
by Lemma 6.4.4.

We claim that $x \in X \setminus \overline{U}$. To see this, suppose to the contrary that $x \in \overline{U}$. Hence there exists an $C_x \in \mathcal{C}$ such that $x \in \overline{C_x}$. Since $C_x \in \mathcal{C}$, $C_x \cap A \neq \emptyset$ so $C_x \nsubseteq X \setminus A$. However, since $\mathcal{U}$ was a refinement of $\{X \setminus A\} \cup \{U_a\}_{a \in A}$, it must be the case that there exists an $a \in A$ such that $C_x \subseteq U_a$. Hence $C_x \cap V_a \subseteq U_a \cap V_a = \emptyset$ thereby contradicting the fact that $x \in \overline{C_x}$ by Theorem 1.6.21 as V_a is a neighbourhood of x. Thus $x \in X \setminus \overline{U}$.

Since U and $X \setminus \overline{U}$ are disjoint open subset of $(X, \mathcal{T})$ such that $A \subseteq U$ and $x \in X \setminus \overline{U}$, we have demonstrated the desired separation. ∎

Thus, as Theorem 5.1.24 used Lemma 3.1.12 to show that compact Hausdorff topological spaces are normal, so too may we use Lemma 6.6.3 in a similar way to show the following.

Theorem 6.6.4. *Every paracompact Hausdorff topological space is normal.*

Proof. Let $(X, \mathcal{T})$ be a paracompact Hausdorff space. To see that $(X, \mathcal{T})$ is normal, let A and B be arbitrary disjoint non-empty closed subset of $(X, \mathcal{T})$. Since $(X, \mathcal{T})$ is regular by Lemma 6.6.3, for each $a \in A$ there exists $U_a, V_a \in \mathcal{T}$ such that $a \in U_a$, $B \subseteq V_a$, and $U_a \cap V_a = \emptyset$. Thus $\{X \setminus A\} \cup \{U_a\}_{a \in A}$ is an open cover of $(X, \mathcal{T})$. Therefore, since $(X, \mathcal{T})$ is paracompact, there exists a locally finite open refinement $\mathcal{U}$ of $\{X \setminus A\} \cup \{U_a\}_{a \in A}$ that covers $(X, \mathcal{T})$.

Let
$$\mathcal{C} = \{C \in \mathcal{U} \mid C \cap A \neq \emptyset\}$$

and let

$$U = \bigcup_{C \in \mathcal{C}} C.$$

Clearly U is an open subset of $(X, \mathcal{T})$. Moreover, since $\mathcal{U}$ was a cover of $(X, \mathcal{T})$, $\mathcal{C}$ is an open cover of A and thus $A \subseteq U$. Furthermore, since $\mathcal{U}$ is locally finite, $\mathcal{C}$ is locally finite and thus

$$\overline{U} = \bigcup_{C \in \mathcal{C}} \overline{C}$$

by Lemma 6.4.4.

We claim that $B \subseteq X \setminus \overline{U}$. To see this, suppose to the contrary that there exists a $b \in B$ such that $b \in \overline{U}$. Hence there exists an $C_b \in \mathcal{C}$ such that $b \in \overline{C_x}$. Since $C_b \in \mathcal{C}$, $C_b \cap A \neq \emptyset$ so $C_b \nsubseteq X \setminus A$. However, since $\mathcal{U}$ was a refinement of $\{X \setminus A\} \cup \{U_a\}_{a \in A}$, it must be the case that there exists an $a \in A$ such that $C_b \subseteq U_a$. Hence $C_b \cap V_a \subseteq U_a \cap V_a = \emptyset$ thereby contradicting the fact that $b \in \overline{C_b}$ by Theorem 1.6.21 as V_a is a neighbourhood of b. Thus $B \subseteq X \setminus \overline{U}$.

Since U and $X \setminus \overline{U}$ are disjoint open subset of $(X, \mathcal{T})$ such that $A \subseteq U$ and $B \subseteq X \setminus \overline{U}$, we have demonstrated the desired separation. ∎

Moreover, in a similar fashion to how Theorem 3.1.14 shows that a closed subspace of a compact space is compact, we can prove the following.

Theorem 6.6.5. *Every closed subspace of a paracompact topological space is paracompact.*

Proof. Let $(X, \mathcal{T})$ be a paracompact topological space and let A be a closed subspace of $(X, \mathcal{T})$. To see that A is paracompact, let $\{U_\alpha\}_{\alpha \in I}$ be an arbitrary open cover of A. By the definition of the subspace topology, for each $\alpha \in I$ there exists a $V_\alpha \in \mathcal{T}$ such that

$$U_\alpha = A \cap V_\alpha.$$

Since A is a closed subset of $(X, \mathcal{T})$ and $\{U_\alpha\}_{\alpha \in I}$ is an open cover of A,

$$\mathcal{V} = \{X \setminus A\} \cup \{V_\alpha\}_{\alpha \in I}$$

is an open cover of $(X, \mathcal{T})$. Therefore, as $(X, \mathcal{T})$ is paracompact, there exists a locally finite open refinement $\mathcal{V}_0$ of $\mathcal{V}$ that covers $(X, \mathcal{T})$. It is then elementary to see that

$$\mathcal{U}_0 = \{A \cap V \mid V \in \mathcal{V}_0\}$$

is a locally finite (by Lemma 6.4.4 and as intersecting with a fixed set preserves being locally finite) open (because $V \in \mathcal{V}_0 \subseteq \mathcal{T}$ so $A \cap V$ is open in A) refinement (reductions of refinements are refinements) of $\mathcal{U}$ that covers A (as $\mathcal{V}_0$ covered $(X, \mathcal{T})$) as desired. ∎

Let's return to our motivation of using the idea of paracompactness to simplifying the task of finding σ-locally finite bases. In particular, our goal in a regular topological space is to relate paracompactness and the existence of a σ-locally finite basis. Thus the equivalence of (i) and (ii) in the following lemma is of particular interest. Although we are not interested in (iii) and (iv) in these notes, we include them out of interest and the fact that the proof that (ii) implies (i) must go through (iii) and (iv) anyways, so we might as well break down the proof as much as possible.

Lemma 6.6.6. *Let $(X, \mathcal{T})$ be a regular topological space. The following are equivalent:*

(i) Every open cover of $(X, \mathcal{T})$ has an open refinement that is locally finite and covers X (i.e. $(X, \mathcal{T})$ is paracompact).

(ii) Every open cover of $(X, \mathcal{T})$ has an open refinement that is σ-locally finite and covers X.

(iii) Every open cover of $(X, \mathcal{T})$ has a refinement that is locally finite and covers X.

(iv) Every open cover of $(X, \mathcal{T})$ has a closed refinement that is locally finite and covers X.

Proof. To being, clearly (i) implies (ii).

To see that (ii) implies (iii), let $\mathcal{U}$ be an arbitrary open cover of $(X, \mathcal{T})$. By the assumption of (ii), there exists an open refinement $\mathcal{A}$ of $\mathcal{U}$ that σ-locally finite and covers X. Since $\mathcal{A}$ is σ-locally finite, there exists locally finite sets $\{\mathcal{A}_n\}_{n=1}^\infty \subseteq \mathcal{P}(\mathcal{T})$ such that $\mathcal{A} = \bigcup_{n=1}^\infty \mathcal{A}_n$.

In order to obtain a refinement of $\mathcal{U}$ that is locally finite and covers X, we will systematically modify each $\mathcal{A}_n$ to make sets from different $\mathcal{A}_n$ disjoint in order to improve σ-locally finite to locally finite at the cost of no longer having open sets. To begin this process, for each $m \in \mathbb{N}$ let

$$V_m = \bigcup_{A \in \mathcal{A}_m} A,$$

for each $n \in \mathbb{N}$ and $A \in \mathcal{A}_n$ let

$$S_n(A) = A \setminus \left(\bigcup_{m<n} V_m \right),$$

for each $n \in \mathbb{N}$ let

$$\mathcal{B}_n = \{S_n(A) \mid A \in \mathcal{A}_n\},$$

and let

$$\mathcal{B} = \bigcup_{n=1}^\infty \mathcal{B}_n.$$

Since $\mathcal{B}_n$ is a refinement of $\mathcal{A}_n$ for all $n \in \mathbb{N}$ as $S_n(A) \subseteq A$ for all $A \in \mathcal{A}_n$, clearly $\mathcal{B}$ is a refinement of $\mathcal{A}$ and thus a refinement of $\mathcal{U}$ by definition. Hence it remains only to show that $\mathcal{B}$ is locally finite and covers X.

To see that $\mathcal{B}$ is locally finite and covers X, let $x \in X$ be arbitrary. Our goal is to show that $x \in B$ for some $B \in \mathcal{B}$ and that there exists a neighbourhood of x that intersects only finitely many elements of $\mathcal{B}$. To begin, recall since $\mathcal{A}$ covers X, there exists an $n \in \mathbb{N}$ and an $A \in \mathcal{A}_n$ such that $x \in A$. Choose $n_x \in \mathbb{N}$ such that there exists an $A_x \in \mathcal{A}_{n_x}$ such that $x \in A_x$ but $x \notin A$ for all $A \in \mathcal{A}_n$ with $n < n_x$. Thus $x \notin V_m$ for all $m < n_x$. Hence $x \in A_x \in \mathcal{A}_{n_x}$ and $x \notin V_m$ for all $m < n_x$ implies that $x \in S_{n_x}(A_x) \in \mathcal{B}_{n_x} \subseteq \mathcal{B}$. Therefore, as x was arbitrary, $\mathcal{B}$ covers X.

To construct the desired neighbourhood of x, note for each $n \leq n_x$ that $\mathcal{A}_n$ is locally finite so there exists a neighbourhood W_n of x that intersects only finitely name elements of $\mathcal{A}_n$. Therefore, since $S_n(A) \subseteq A$ for all $A \in \mathcal{A}_n$, we see that W_n intersects only a finite number of elements of $\mathcal{B}_n$ for all $n \leq n_x$. Let

$$V = A_{n_x} \cap \left(\bigcap_{n=1}^{n_x} W_n \right),$$

which is a neighbourhood of x being a finite intersection of neighbourhoods of x. As $V \subseteq W_n$ for all $n \leq n_x$, we see that V intersects only a finite number of elements of $\mathcal{B}_n$ for all $n \leq n_x$. However, as $A_{n_x} \subseteq V_{n_x}$ so that $S_m(A) \cap A_{n_x} = \emptyset$ for all $A \in \mathcal{A}_m$ and $m > n_x$, we obtain that V does not intersect any element of $\mathcal{B}_n$ for all $n > n_x$. Hence, as a finite union of finite sets is finite, V intersects only as finite number of elements of $\mathcal{B}$. Therefore, as x was arbitrary, $\mathcal{B}$ is locally finite as desired. Hence (ii) implies (iii).

To see that (iii) implies (iv), let $\mathcal{U}$ be an arbitrary open cover of $(X, \mathcal{T})$. Let

$$\mathcal{A} = \{V \in \mathcal{T} \mid \overline{V} \subseteq U \text{ for some } U \in \mathcal{U}\}.$$

Clearly $\mathcal{A}$ is an open refinement of $\mathcal{U}$. Furthermore, since $(X, \mathcal{T})$ is regular, Lemma 5.1.15 implies that $\mathcal{A}$ covers X. Hence by the assumption of (iii) there exists a refinement $\mathcal{B}$ of $\mathcal{A}$ that is locally finite and covers X. Let

$$\mathcal{C} = \{\overline{B} \mid B \in \mathcal{B}\}.$$

Clearly $\mathcal{C}$ is a collection of closed subset of $(X, \mathcal{T})$ that cover $(X, \mathcal{T})$ since $\mathcal{B}$ covers X. As $\mathcal{C}$ is locally finite by Lemma 6.4.4, to complete the proof, we claim that $\mathcal{C}$ is a refinement of $\mathcal{U}$. To see this, let $C \in \mathcal{C}$ be arbitrary. Hence there exists a $B \in \mathcal{B}$ such that $C = \overline{B}$. However, since $\mathcal{B}$ is a refinement of $\mathcal{A}$, there exists a $V \in \mathcal{A}$ such that $B \subseteq V$. Hence, by the definition of A, there exists a $U \in \mathcal{U}$ such that

$$C = \overline{B} \subseteq \overline{V} \subseteq U.$$

Therefore, as $C \in \mathcal{C}$ was arbitrary, $\mathcal{C}$ is a refinement of $\mathcal{U}$. Hence (iii) implies (iv).

To see that (iv) implies (i), let $\mathcal{U}$ be an arbitrary open cover of $(X, \mathcal{T})$. By the assumption of (iv), there exists a (closed) refinement $\mathcal{A}$ of $\mathcal{U}$ that is locally finite and covers X. To obtain our open refinement of $\mathcal{U}$ that is locally finite and covers X by using (iv), we will need some nice closed sets to take complements of.

Let

$$\mathcal{B} = \{V \in \mathcal{T} \mid V \text{ intersects at most a finite number of elements of } \mathcal{A}\}.$$

Since $\mathcal{A}$ is locally finite, for each $x \in X$ there exists a $B \in \mathcal{B}$ such that $x \in B$. Hence $\mathcal{B}$ is an open cover of $(X, \mathcal{T})$. By the assumption of (iv), there exists a closed refinement $\mathcal{C}$ of $\mathcal{B}$ that is locally finite and covers X. Since $\mathcal{C}$ is a refinement of $\mathcal{B}$, the definition of $\mathcal{B}$ implies that each element of $\mathcal{C}$ intersects at most a finite number of elements of $\mathcal{A}$.

For each $A \in \mathcal{A}$, let

$$C_A = \bigcup_{\substack{C \in \mathcal{C} \\ C \subseteq X \setminus A}} C \subseteq X \setminus A$$

and let

$$U(A) = X \setminus C_A.$$

Since $\mathcal{C}$ is locally finite so any subset of $\mathcal{C}$ is locally finite, and since every element of $\mathcal{C}$ is closed, Lemma 6.4.4 implies that C_A is a closed subset of $(X, \mathcal{T})$ for all $A \in \mathcal{A}$. Hence $U(A)$ is open in $(X, \mathcal{T})$ for all $A \in \mathcal{A}$. Furthermore, since $A \subseteq U(A)$ for all $A \in \mathcal{A}$ by construction, and since $\mathcal{A}$ covers X,

$$\mathcal{U}_0 = \{U(A) \mid A \in \mathcal{A}\}$$

is an open cover of X.

Unfortunately $\mathcal{U}_0$ need not be a refinement of $\mathcal{U}$. To rectify the situation, note since $\mathcal{A}$ is a refinement of $\mathcal{U}$, for each $A \in \mathcal{A}$ there exists a $V(A) \in \mathcal{U}$ such that $A \subseteq V(A)$. Let

$$\mathcal{D} = \{U(A) \cap V(A) \mid A \in \mathcal{A}\}.$$

Clearly $\mathcal{D}$ is a refinement of $\mathcal{A}$ by construction and thus a refinement of $\mathcal{U}$. Moreover, since $\mathcal{U}$ and $\mathcal{U}_0$ are open covers of X, it is elementary to see that $\mathcal{D}$ is an open cover of X.

We claim that $\mathcal{D}$ is locally finite thereby completing the proof. To see this, let $x \in X$ be arbitrary. Since $\mathcal{C}$ is locally finite, there exists a neighbourhood W of x such that W intersects only a finite number of elements of $\mathcal{C}$. Therefore, since $\mathcal{C}$ covers X and thus W, there exists a finite set $\{C_k\}_{k=1}^n \subseteq \mathcal{C}$ such that

$$W \subseteq \bigcup_{k=1}^n C_k.$$

We claim that each C_k for $k \in \{1, \ldots, n\}$ intersects only a finite number of elements of $\mathcal{D}$ thereby completing the proof. To see this, let $C \in \mathcal{C}$ be arbitrary. If C intersects $U(A) \cap V(A)$ for some $A \in \mathcal{A}$, then $C \cap U(A) \neq \emptyset$ which implies C is not a subset of C_A and thus implies C is not contained in $X \setminus A$ thereby yielding that C intersects A. Since $\mathcal{C}$ is a refinement of $\mathcal{B}$ and since each element of $\mathcal{B}$ intersects only a finite number of elements of $\mathcal{A}$, C can only intersect a finite number of elements of $\mathcal{A}$. Hence as C intersects A whenever C intersects $U(A) \cap V(A)$, C can only intersect a finite number of elements of $\mathcal{D}$. Hence $\mathcal{D}$ is a locally finite open refinement of $\mathcal{U}$ as desired. ∎

Lemma 6.6.6 now easily enables us to show the following using σ-locally finiteness as a stopping point.

Lemma 6.6.7. *Every metrizable topological space is paracompact.*

Proof. Let $(X, \mathcal{T})$ be a metrizable topological space. Hence $(X, \mathcal{T})$ is regular. Moreover, Lemma 6.4.9 implies that ever open cover of $(X, \mathcal{T})$ has a σ-locally finite open refinement that covers $(X, \mathcal{T})$. Hence Lemma 6.6.6 implies that every open cover of $(X, \mathcal{T})$ has an open refinement that is locally finite and covers X. Thus $(X, \mathcal{T})$ is paracompact as desired. ∎

6.7 Smirnov Metrization Theorem

To conclude our discussion of metrizability, we improve on the Nagata-Smirnov Theorem (Theorem 6.5.5) by replacing having a σ-finite basis with paracompactness provided our topological space also has the following property.

Definition 6.7.1. A topological space $(X, \mathcal{T})$ is said to be *locally metrizable* if every point in X has a neighbourhood that is metrizable in the subspace topology.

Recall from Remark 6.3.2 that one use of Urysohn's Metrization Theorem (Theorem 6.3.1) was to imply the existence of metrics on topological manifolds. By definition topological manifolds are locally Euclidean and thus locally metrizable. Thus the following theorem is by far an improvement as paracompactness is more general than second countability (i.e. a topological manifold need not be second countable, although some authors make this assumption) and often paracompactness is not difficult to check.

Theorem 6.7.2 (Smirnov Metrization Theorem). *A topological space $(X, \mathcal{T})$ is metrizable if and only if it is a locally metrizable, paracompact, Hausdorff topological space.*

Proof. To begin, suppose that $(X, \mathcal{T})$ is a metrizable topological space. Then clearly $(X, \mathcal{T})$ is locally metrizable and Hausdorff. Furthermore, $(X, \mathcal{T})$ is paracompact by Lemma 6.6.7. Hence one direction of the proof is complete.

Conversely, suppose that $(X, \mathcal{T})$ is a locally metrizable, paracompact, Hausdorff topological space. To see that $(X, \mathcal{T})$ is metrizable, it suffices by the Nagata-Smirnov Metrization Theorem (Theorem 6.5.5) to show that $(X, \mathcal{T})$ is regular and has a σ-locally finite basis. Notice $(X, \mathcal{T})$ is regular by Lemma 6.6.3.

To see that $(X, \mathcal{T})$ has a σ-locally finite basis, recall since $(X, \mathcal{T})$ is locally metrizable that for each $x \in X$ there exists a $U_x \in \mathcal{T}$ such that U_x is metrizable. Hence $\mathcal{U} = \{U_x\}_{x \in X}$ is an open cover of $(X, \mathcal{T})$. Therefore, since $(X, \mathcal{T})$ is paracompact and regular, Lemma 6.6.6 implies that there exists an open refinement $\mathcal{A}$ of $\mathcal{U}$ that is locally finite and covers $(X, \mathcal{T})$. Since $\mathcal{A}$ is a refinement of $\mathcal{U}$, for each $A \in \mathcal{A}$ there exists an $x \in X$ such that $A \subseteq U_x$. Thus, as a subspace of a metrizable space is metrizable, every element of $\mathcal{A}$ is metrizable. Hence for each $A \in \mathcal{A}$ there exists a metric $d_A : A \times A \to [0, \infty)$.

For each $n \in \mathbb{N}$ let

$$\mathcal{A}_n = \left\{ B_{d_A}\left(a, \frac{1}{n}\right) \,\middle|\, a \in A \text{ and } A \in \mathcal{A} \right\}.$$

Clearly $\mathcal{A}_n$ is clearly an open cover of $(X, \mathcal{T})$ as $\mathcal{A}$ is a cover of $(X, \mathcal{T})$. Therefore, since $(X, \mathcal{T})$ is paracompact and regular, Lemma 6.6.6 implies that there exists an open refinement $\mathcal{B}_n$ of $\mathcal{A}_n$ that is σ-locally finite and covers $(X, \mathcal{T})$.

Let $\mathcal{B} = \bigcup_{n=1}^\infty \mathcal{B}_n$. We claim that $\mathcal{B}$ is a σ-locally finite basis of $(X, \mathcal{T})$. To see this, note $\mathcal{B}$ is clearly σ-locally finite being the countable union of σ-locally finite subset of $(X, \mathcal{T})$.

To see that $\mathcal{B}$ is a basis for $(X, \mathcal{T})$, let $x \in X$ and V a neighbourhood of x be arbitrary. Since $\mathcal{A}$ is locally finite, there exists a finite subset $\{A_k\}_{k=1}^m \subseteq \mathcal{A}$ such that $A_k \cap V \neq \emptyset$ for all $k \in \{1, \ldots, m\}$ and $A \cap V = \emptyset$ for all $A \in \mathcal{A} \setminus \{A_k\}_{k=1}^m$. Since $A_k \cap V \subseteq A_k$ is a neighbourhood of x for all $k \in \{1, \ldots, m\}$, there exists an $\epsilon_k > 0$ such that

$$B_{d_{A_k}}(x, \epsilon_k) \subseteq A_k \cap V.$$

Let $\epsilon = \min(\{\epsilon_1, \epsilon_2, \ldots, \epsilon_m\})$ and choose $n \in \mathbb{N}$ such that $\frac{1}{n} < \frac{\epsilon}{2}$.

Since $\mathcal{B}_n$ covers $(X, \mathcal{T})$, there exists a $B \in \mathcal{B}_n \subseteq \mathcal{B}$ such that $x \in B$. Since $\mathcal{B}_n$ is a refinement of $\mathcal{A}_n$, there must exist an $A \in \mathcal{A}_n$ and an $a \in A$ such that $B \subseteq B_{d_A}\left(a, \frac{1}{n}\right)$. Therefore $x \in V$ and

$$x \in B \subseteq B_{d_A}\left(a, \frac{1}{n}\right) \subseteq A,$$

so $A = A_{k_0}$ for some $k_0 \in \{1, \ldots, m\}$. Moreover, since $B \subseteq B_{d_{A_{k_0}}}\left(a, \frac{1}{n}\right)$, we have that $\operatorname{diam}_{d_{A_{k_0}}}(B) \leq \frac{2}{n}$. Therefore, since $x \in B$, we have that

$$x \in B \subseteq B_{A_{k_0}}\left[x, \frac{2}{n}\right] \subseteq B_{A_{k_0}}(x, \epsilon) \subseteq A_{k_0} \cap V \subseteq V.$$

Therefore, as x and V were arbitrary, $\mathcal{B}$ is a basis for $(X, \mathcal{T})$ thereby completing the proof. ∎

Appendix A

Basic Set Theory

All mathematics must contain some notation in order for one to adequately describe the objects of study. As such, we begin the notation for the most basic structures in this course.

A.1 Sets

One of the most natural mathematical objects is the following:

Heuristic Definition. A *set* is a collection of distinct objects.

The following table list several sets, the symbol used to represent the set, and a set notational way to describe the set.

Set	**Symbol**	**Set Notation**
natural numbers	$\mathbb{N}$	$\{1, 2, 3, 4, \ldots\}$
integers	$\mathbb{Z}$	$\{0, 1, -1, 2, -2, 3, -3, \ldots\}$
rational numbers	$\mathbb{Q}$	$\{\frac{a}{b} \mid a, b \in \mathbb{Z}, b \neq 0\}$
real numbers	$\mathbb{R}$	$\{\text{real numbers}\}$
complex numbers	$\mathbb{C}$	$\{a + bi \mid a, b \in \mathbb{R}\}$

Notice two different types of notation are used in the above table to describe sets: namely $\{\text{objects}\}$ and $\{\text{objects} \mid \text{conditions on the objects}\}$. Furthermore, the symbol $\emptyset$ will denote the *empty set*; that is, the set with no elements.

Given a set X and an object x, we say that x is an *element* of X, denoted $x \in X$, when x is one of the objects that make up X. Furthermore, we will use $x \notin X$ when x is not an element of X. For example, $\sqrt{2} \in \mathbb{R}$ yet $\sqrt{2} \notin \mathbb{Q}$ and $0 \in \mathbb{Z}$ but $0 \notin \mathbb{N}$. Furthermore, given two sets X and Y, we say that Y is a *subset* of X, denoted $Y \subseteq X$, if each element of Y is an element of X; that is, if $a \in Y$ then $a \in X$. For example $\mathbb{N} \subseteq \mathbb{Z} \subseteq \mathbb{Q} \subseteq \mathbb{R}$. Furthermore, note the empty set is a subset of all sets, and if $X \subseteq Y$ and $Y \subseteq X$ then $X = Y$.

One important question that has not been addressed is, "What exactly is a set?" This questions must be asked as we have not provided a rigorous definition of a set. This leads to some interesting questions, such as, "Does the collection of all sets form a set?"

Let us suppose that there is a set of all sets; that is

$$Z = \{X \mid X \text{ is a set}\}$$

makes sense. Note Z has the interesting property that $Z \in Z$. Furthermore, if Z exists, then

$$Y = \{X \mid X \text{ is a set and } X \notin X\}$$

would be a valid subset of Z. However, we clearly have two disjoint cases: either $Y \in Y$ or $Y \notin Y$ (that is, either Y is an element of Y or Y is not an element of Y).

If $Y \in Y$, then the definition of Y implies $Y \notin Y$ which is a contradiction since we cannot have both $Y \in Y$ and $Y \notin Y$. Thus, if $Y \in Y$ is false, then it must be the case that $Y \notin Y$.

However, $Y \notin Y$ implies by the definition of Y that $Y \in Y$. Again this is a contradiction since we cannot have both $Y \notin Y$ and $Y \in Y$. This argument is known as Russell's Paradox and demonstrates that there cannot be a set of all sets.

The above paradox illustrates the necessity of a rigorous definition of a set. However, said definition takes us in a different direction than desired in this course. That being said, a rigorous definition of a set would provide us with the ability to take subsets of a given set and would permit the following operations on sets.

Definition A.1.1. Let X be a set. The *power set of* X, denote $\mathcal{P}(X)$, is

$$\mathcal{P}(X) = \{A \mid A \subseteq X\}.$$

Note $\emptyset \in \mathcal{P}(X)$ and $X \in \mathcal{P}(X)$.

Definition A.1.2. Let $\{X_\alpha\}_{\alpha\in I}$ denote a collection of subsets of a set X indexed by a set I.

- The *union* of $\{X_\alpha\}_{\alpha\in I}$, denoted $\bigcup_{\alpha\in I} X_\alpha$, is the set

$$\bigcup_{\alpha\in I} X_\alpha = \{a \mid a \in X_\alpha \text{ for at least one } \alpha \in I\}.$$

- The *intersection* of $\{X_\alpha\}_{\alpha\in I}$, denoted $\bigcap_{\alpha\in I} X_\alpha$, is the set

$$\bigcap_{\alpha\in I} X_\alpha = \{a \mid a \in X_\alpha \text{ for all } \alpha \in I\}.$$

Definition A.1.3. Given two sets X and Y, the *set difference* of X and Y, denoted $X \setminus Y$, is the set

$$X \setminus Y = \{a \mid a \in X \text{ and } a \notin Y\}.$$

In this course, we will often have a set X and will be considering subsets of X. Consequently, given a subset Y of X, the set difference $X \setminus Y$ will be called the *complement* of Y (in X) and will be denoted Y^c for convenience.

To conclude this section, we note the following set inequalities that will be used surprisingly often in this course.

Theorem A.1.4 (De Morgan's Laws). *Let X and I be non-empty sets and for each $\alpha \in I$ let X_α be a subset of X. Then*

$$X \setminus \left(\bigcup_{\alpha \in I} X_\alpha \right) = \bigcap_{\alpha \in I} (X \setminus X_\alpha) \qquad \textit{and} \qquad X \setminus \left(\bigcap_{\alpha \in I} X_\alpha \right) = \bigcup_{\alpha \in I} (X \setminus X_\alpha).$$

Proof. Notice that

$$\begin{aligned} x \in \left(\bigcup_{i \in I} EXi \right)^c &\iff x \notin \bigcup_{i \in I} X_i \\ &\iff x \notin X_i \text{ for all } i \in I \\ &\iff x \in X_i^c \text{ for all } i \in I \\ &\iff x \in \bigcap_{i \in I} X_i^c \end{aligned}$$

which completes the proof since we have shown that $x \in (\bigcup_{i \in I} X_i)^c$ if and only if $x \in \bigcap_{i \in I} X_i^c$ (which implies the sets are the same).

We can play a similar game to prove that

$$\left(\bigcap_{i \in I} X_i \right)^c = \bigcup_{i \in I} X_i^c.$$

Alternatively, we can use the first result to prove the second. To do this, we must first show that if $E \subseteq X$ and $F = E^c$, then $F^c = E$. Indeed notice $x \in F^c$ if and only if $x \notin F$ if and only if $x \notin E^c$ if and only if $x \in E$. Hence $F^c = E$.

To prove this new equality using the old, for each $i \in I$ let $F_i = X_i^c$. By applying the first equation using the F_i's instead of the X_i's, we obtain that

$$\left(\bigcup_{i \in I} F_i \right)^c = \bigcap_{i \in I} F_i^c.$$

Since $F_i = X_i^c$ so $F_i^c = X_i$ by the above proof, we have that

$$\left(\bigcup_{i \in I} X_i^c \right)^c = \bigcap_{i \in I} X_i.$$

Hence

$$\bigcup_{i \in I} X_i^c = \left(\bigcap_{i \in I} X_i \right)^c$$

by taking the complement of both sides and using the proof in the above paragraph. ■

A.2 Functions

In any analysis course, functions will play a fundamental role. The most useful and accurate method for defining functions is to use the following operation on sets (which is also valid by the actual definition of what a set is).

Definition A.2.1. Given two non-empty sets X and Y, the *Cartesian product* of X and Y, denoted $X \times Y$, is the set

$$X \times Y = \{(x, y) \mid x \in X, y \in Y\}.$$

Definition A.2.2. Given two non-empty sets X and Y, a *function* f from X to Y, denoted $f : X \to Y$, is a subset S of $X \times Y$ such that for each $x \in X$ there is an unique element denoted $f(x) \in Y$ such that $(x, f(x)) \in S$ (that is, a function is defined by its graph).

Example A.2.3. Given two non-empty sets X and Y, there is a natural way to view

$$X \times Y = \{f : \{1, 2\} \to X \cup Y \mid f(1) \in X, f(2) \in Y\}.$$

Indeed, a function $f : \{1, 2\} \to X \cup Y$ is uniquely determined by the values $f(1)$ and $f(2)$. Consequently, an $f : \{1, 2\} \to X \cup Y$ as defined in the above set can be viewed as the pair $(f(1), f(2))$. Conversely a pair $(x, y) \in X \times Y$ can be represented by the function $f : \{1, 2\} \to X \cup Y$ defined by $f(1) = x$ and $f(2) = y$.

The above example can be extended from a pair of sets to a finite number of sets. Let $X_1, \ldots, X_n$ be non-empty sets. We define the *product* of these sets to be

$$X_1 \times \cdots \times X_n = \{(x_1, \ldots, x_n) \mid x_j \in X_j \text{ for all } j \in \{1, \ldots, n\}\}.$$

If $X = X_1 = \cdots = X_n$, we will write X^n for $X_1 \times \cdots \times X_n$.

Notice we can view $X_1 \times \cdots \times X_n$ as a set of functions in a similar manner to Example A.2.3. Indeed

$$X_1 \times \cdots \times X_n = \left\{ f : \{1, \ldots, n\} \to \bigcup_{k=1}^{n} X_k \;\middle|\; f(j) \in X_j \;\forall\, j \in \{1, \ldots, n\} \right\}.$$

But what happens if we want to take a product of an infinite number of sets?

Given a non-empty set I and a collection of non-empty sets $\{X_\alpha\}_{\alpha\in I}$, we define the product

$$\prod_{\alpha\in I} X_\alpha = \left\{ f : I \to \bigcup_{\alpha\in I} X_\alpha \,\middle|\, f(i) \in X_i \text{ for all } i \in I \right\}.$$

However, we must ask, "Is the above set non-empty?" That is, how do we know there is always such a function? The answer is, because we add an axiom to make it so.

Axiom A.2.4 (The Axiom of Choice). *Given a non-empty set I and a collection of non-empty sets $\{X_\alpha \mid \alpha \in I\}$, the product $\prod_{\alpha\in I} X_\alpha$ is non-empty. Any function $f \in \prod_{\alpha\in I} X_\alpha$ is called a choice function.*

One may ask, "Why Mr. Anderson? Why? Why do we include the Axiom of Choice?" The short answer is, of course, "Because I choose to."

It turns out that the Axiom of Choice is independent from the axioms of (Zermelo–Fraenkel) set theory. This means that if one starts with the standard axioms of set theory, one can neither prove nor disprove the Axiom of Choice. Thus we have the option on whether to include or exclude the Axiom of Choice from our theory. We will allow the use of the Axiom of Choice (and almost surely you have made use of it in a previous analysis course and didn't even know it!).

A.3 Bijections

As we will be using functions throughout the remainder of the course, we will need some notation and definitions.

Given a function $f : X \to Y$ and $A \subseteq X$, we define

$$f(A) = \{f(x) \mid x \in A\} \subseteq Y.$$

Definition A.3.1. Given a function $f : X \to Y$, the *range* of f is $f(X)$.

Using the notion of the range, we can define an important property we may desire our functions to have.

Definition A.3.2. A function $f : X \to Y$ is said to be *surjective* (or *onto*) if $f(X) = Y$; that is, for each $y \in Y$ there exists an $x \in X$ such that $f(x) = y$.

To illustrate when a function is surjective or not, consider the following diagrams.

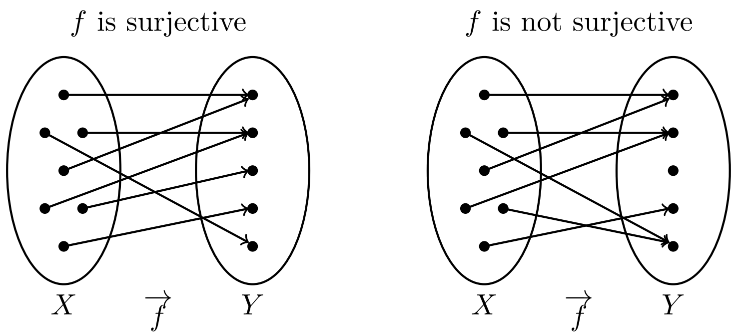

Example A.3.3. Consider the function $f : [0, 1] \to [0, 2]$ defined by $f(x) = x^2$. Notice f is not surjective since $f(x) \neq 2$ for all $x \in [0, 1]$. However, the function $g : [0, 1] \to [0, 1]$ defined by $g(x) = x^2$ is surjective. Consequently, the target set (known as the *co-domain*) matters.

One useful tool when dealing with functions is to be able to describe all points in the initial space that map into a predetermined set. Thus we make the following definition.

Definition A.3.4. Given a function $f : X \to Y$ and a $B \subseteq Y$, the *preimage* of B under f is the set

$$f^{-1}(B) = \{x \in X \mid f(x) \in B\} \subseteq X.$$

Note the notation used for the preimage does not assume the existence of an inverse of f (see Theorem A.3.8). Using preimages, we can define an important property we may desire our functions to have.

Definition A.3.5. A function $f : X \to Y$ is said to be *injective* (or *one-to-one*) if for all $y \in Y$, the preimage $f^{-1}(\{y\})$ has at most one element; that is, if $x_1, x_2 \in X$ are such that $f(x_1) = f(x_2)$, then $x_1 = x_2$.

To illustrate when a function is injective or not, consider the following diagrams.

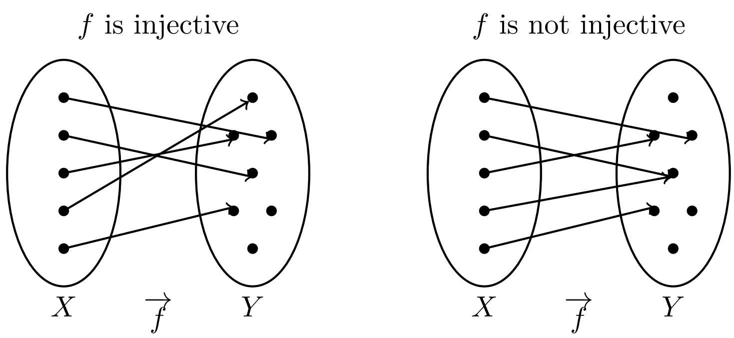

Example A.3.6. Consider the function $f : [-1, 1] \to [0, 1]$ defined by $f(x) = x^2$. Notice f is not injective since $f(-1) = f(1)$. However, the function $g : [0, 1] \to [0, 1]$ defined by $g(x) = x^2$ is injective. Consequently, the initial set (known as the *domain*) matters.

We desire to combine the notions of injective and surjective.

Definition A.3.7. A function $f : X \to Y$ is said to be a *bijection* if f is injective and surjective.

Using the above examples, we have seen several functions that are not bijective. Furthermore, we have seen that $f : [0, 1] \to [0, 1]$ defined by $f(x) = x^2$ is bijective. One way to observe that f is injective is to consider the function $g : [0, 1] \to [0, 1]$ defined by $g(x) = \sqrt{x}$. Notice that f and g 'undo' what the other function does. In fact, this is true of all bijections.

Theorem A.3.8. *A function $f : X \to Y$ is a bijection if and only if there exists a function $g : Y \to X$ such that*

- *$g(f(x)) = x$ for all $x \in X$, and*
- *$f(g(y)) = y$ for all $y \in Y$.*

Furthermore, if f is a bijection, there is exactly one function $g : Y \to X$ that satisfies these properties, which is called the inverse of f and is denoted by $f^{-1} : Y \to X$. Notice this implies f^{-1} is also a bijection with $(f^{-1})^{-1} = f$.

Proof. Suppose that f is a bijection. Since f is surjective, for each $y \in Y$ there exists an $z_y \in X$ such that $f(z_y) = y$. Furthermore, note z_y is the unique element of X that f maps to y since f is injective.

Define $g : Y \to X$ by $g(y) = z_y$. Clearly g is a well-defined function.

To see that g satisfies the two properties, first let $x \in X$ be arbitrary. Then $y = f(x) \in Y$. However, since $f(z_y) = y = f(x)$, it must be the case that $z_y = x$ as f is injective. Therefore

$$g(f(x)) = g(y) = z_y = x$$

as desired. For the second property, let $y \in Y$ be arbitrary. Then

$$f(g(y)) = f(z_y) = y$$

by the definition of z_y. Hence g satisfies the desired properties.

Conversely, suppose $g : Y \to X$ satisfies the two properties. To see that f is injective, suppose $x_1, x_2 \in X$ are such that $f(x_1) = f(x_2)$. Then

$$x_1 = g(f(x_1)) = g(f(x_2)) = x_2$$

as desired. To see that f is surjective, let $y \in Y$ be arbitrary. Then $g(y) \in X$ so

$$y = f(g(y)) \in f(X).$$

Since $y \in Y$ is arbitrary, we have $Y \subseteq f(X)$. Hence $f(X) = Y$ so f is surjective. Therefore, as f is both injective and surjective, f is bijective by definition.

Finally, suppose f is bijective and $g : Y \to X$ satisfies the above properties. Suppose $h : Y \to X$ is another function such that $h(f(x)) = x$ for all $x \in X$, and $f(h(y)) = y$ for all $y \in Y$. Then for all $y \in Y$,

$$h(y) = g(f(h(y))) = g(y)$$

(where we have used $g(f(x_1)) = x_1$ when $x_1 = h(y)$ and $f(h(y)) = y$). Therefore $g = h$ as desired. ∎

Remark A.3.9. If $f : X \to Y$ is injective, consider the function $g : X \to f(X)$ defined by $g(x) = f(x)$ for all $x \in X$. Clearly g is injective since f is, and, by construction, g is surjective. Hence g is bijective and thus has an inverse $g^{-1} : f(X) \to X$. The function g^{-1} is called the *inverse of f on its image*.

A.4 Equivalence Relations

Using the same idea as we used for defining functions (i.e. as subsets of a Cartesian product), we can define another useful notion in mathematics.

Definition A.4.1. Given two non-empty sets X and Y, a *relation* is a subset of the product $X \times Y$. Given a relation R, we write xRy if $(x, y) \in R$.

Given a non-empty set X, by a relation on X we will mean a relation on $X \times X$.

Using a specific type of relation, we can generalize the notion of equality.

Definition A.4.2. Let X be a set. A relation $\sim$ on the elements of X is said to be an *equivalence relation* if:

(1) (reflexive) $x \sim x$ for all $x \in X$,

(2) (symmetric) if $x \sim y$, then $y \sim x$ for all $x, y \in X$, and

(3) (transitive) if $x \sim y$ and $y \sim z$, then $x \sim z$ for all $x, y, z \in X$.

Given an $x \in X$, the set $\{y \in X \mid y \sim x\}$ is called the *equivalence class* of x and is denoted $[x]$.

Notice that $[x] \cap [y] \neq \emptyset$ if and only if $x \sim y$. Thus by taking an index set consisting of one element from each equivalence class, the set X can be written as the disjoint union of its equivalence classes.

Example A.4.3. Let V be a vector space and let W be a subspace of V. It is elementary to check that if we define $\vec{x} \sim \vec{y}$ if and only if $\vec{x} - \vec{y} \in W$, then $\sim$ is an equivalence relation on V. Note that the equivalence classes of V then become a vector space, denoted V/W, with the operations $[\vec{x}] + [\vec{y}] = [\vec{x} + \vec{y}]$ and $\alpha[\vec{x}] = [\alpha\vec{x}]$. Note the necessity of checking that these operations are well-defined; that is, for addition to make sense, one must show that if $\vec{x}_1 \sim \vec{x}_2$ and $\vec{y}_1 \sim \vec{y}_2$ then $\vec{x}_1 + \vec{y}_1 \sim \vec{x}_2 + \vec{y}_2$.

One of the most useful examples of equivalence relations is to characterize the open subsets of $\mathbb{R}$. Recall a subset $U \subseteq \mathbb{R}$ is said to be open if for every $x \in U$ there exists an $\epsilon > 0$ such that $(x - \epsilon, x + \epsilon) \subseteq U$. We refer the reader to Section 1.2 more on the definition of an open subset of $\mathbb{R}$ based on the metric induced by the absolute value.

Proposition A.4.4. *Every open subset of* $\mathbb{R}$ *is a countable union of open intervals*

Proof. Let U be an arbitrary non-empty open subset of $\mathbb{R}$. Define a relation $\sim$ on U by $x \sim y$ if and only if whenever $x < z < y$ or $y < z < x$ then $z \in U$. We claim that $\sim$ is an equivalence relation on U.

To see this first notice that if $x \in U$, then $x \sim x$ trivially. Furthermore, clearly if $x \sim y$ then $z \in U$ whenever $x < z < y$ or $y < z < x$, and thus $y \sim x$. Finally, suppose $x, y, w \in U$ are such that $x \sim y$ and $y \sim w$. To see that $x \sim w$, we divide the discussion into five cases:

Case 1: $x \leq y \leq w$. In this case, we have $x < z < y$ implies $z \in U$ and $y < z < w$ implies $z \in U$. If z is such that $x < z < w$, then either $x < z < y$, $y < z < w$, or $y = z$. As all of these imply $z \in U$, we have $z \sim w$ in this case.

Case 2: $w \leq y \leq x$. This case follows from Case 1 by interchanging x and w.

Case 3: $y \leq x \leq w$. In this case, we have $y < z < w$ implies $z \in U$. Thus if $x < z < w$ then $y < z < w$ so $z \in U$. Hence $z \sim x$ in this case.

Case 4: $y \leq w \leq x$. This case follows from Case 3 by interchanging x and w.

Case 5: $x \leq w \leq y$ or $w \leq x \leq y$. This case follows from Cases 3 and 4 by reversing the inequalities.

Thus, in any case $x \sim w$. Thus $\sim$ is an equivalence relation.

Next we claim that each equivalence class is an open interval. To see this let $x \in U$ be arbitrary and let E_x denote the equivalence class of x with respect to $\sim$. To see that E_x is an open interval, let

$$\alpha_x = \inf(E_x) \qquad \text{and} \qquad \beta_x = \sup(E_x).$$

We claim that $E_x = (\alpha_x, \beta_x)$.

First, we claim that $\alpha_x < \beta_x$. To see this, notice that $x \in E_x \subseteq U$. Hence, as U is open, there exists an $\epsilon > 0$ such that $(x - \epsilon, x + \epsilon) \subseteq U$. Clearly $y \sim x$ for all $y \in (x - \epsilon, x + \epsilon)$ so

$$\alpha_x \leq x - \epsilon < x + \epsilon \leq \beta_x.$$

To see that $(\alpha_x, \beta_x) \subseteq E_x$, let $y \in (\alpha_x, \beta_x)$ be arbitrary. Since $\alpha_x < y < \beta_x$, by the definition of inf and sup there exists $z_1, z_2 \in E_x$ such that

$$\alpha_x \leq z_1 < y < z_2 \leq \beta_x.$$

Since $z_1, z_2 \in E_x$, we have $z_1 \sim x$ and $z_2 \sim x$. Thus $z_1 \sim z_2$ so $[z_1, z_2] \subseteq U$. Hence $y \in [z_1, z_2] \subseteq U$. Therefore, as $y \in (\alpha_x, \beta_x)$ was arbitrary, $(\alpha_x, \beta_x) \subseteq E_x$.

To see that $E_x \subseteq (\alpha_x, \beta_x)$, note that $E_x \subseteq (\alpha_x, \beta_x) \cup \{\alpha_x, \beta_x\}$ by the definition of α_x and β_x. Thus it suffices to show that $\alpha_x, \beta_x \notin E_x$. Suppose $\beta_x \in E_x$ (this implies $\beta_x \neq \infty$). Then $\beta_x \in U$ so there exists an $\epsilon > 0$ so that $(\beta_x - \epsilon, \beta_x + \epsilon) \subseteq U$. Hence $\beta_x + \frac{1}{2}\epsilon \sim \beta_x \sim x$ (as $\beta_x \in E_x$). Hence $\beta_x + \frac{1}{2}\epsilon \in E_x$. However $\beta_x + \frac{1}{2}\epsilon > \beta_x$ so $\beta_x + \frac{1}{2}\epsilon \in E_x$ contradicts the fact that $\beta_x = \sup(E_x)$. Hence we have obtained a contradiction so $\beta_x \notin E_x$. Similar arguments show that $\alpha_x \notin E_x$. Hence $E_x = (\alpha_x, \beta_x)$ as desired.

To complete the proof, first notice that clearly

$$U = \bigcup_{x \in U} E_x$$

so U is a union of open intervals. It remains to be verified that the above union can be made countable. Since each E_x is an open interval, $E_x \cap \mathbb{Q} \neq \emptyset$. Hence, as each $E_x \cap \mathbb{Q}$ is non-empty, by the Axiom of Choice there exists a function $f : \{E_x \mid x \in U\} \to \mathbb{Q}$ such that $f(E_x) \in E_x$ for all $x \in U$. Hence, as $E_x \cap E_y = \emptyset$ if $E_x \neq E_y$, f is an injective function. Hence $\{E_x \mid x \in U\}$ is countable. Thus the union $U = \bigcup_{x \in U} E_x$ can be made into a countable union by choosing one representative from each equivalence class (or, alternatively, $U = \bigcup_{q \in \mathbb{Q}} f^{-1}(\{q\})$). ∎

A.5 Zorn's Lemma

In this section, we will briefly review Zorn's Lemma, which is a necessary step in order for us to prove Tychonoff's Theorem (Theorem 3.3.4). We begin with the basics.

Definition A.5.1. Let X be a set. A relation $\preceq$ on the elements of X is called a *partial ordering* if:

(1) (reflexivity) $a \preceq a$ for all $a \in X$,

(2) (antisymmetry) if $a \preceq b$ and $b \preceq a$, then $a = b$ for all $a, b \in X$, and

(3) (transitivity) if $a, b, c \in X$ are such that $a \preceq b$ and $b \preceq c$, then $a \preceq c$.

Clearly $\leq$ is a partial ordering on $\mathbb{R}$. Here is another example:

Example A.5.2. Given a set X, the relation $\preceq$ on $\mathcal{P}(X)$ defined by

$$Z \preceq Y \qquad \text{if and only if} \qquad Z \subseteq Y$$

is a partial ordering on $\mathcal{P}(X)$.

The partial ordering in the previous example is not as nice as our ordering on $\mathbb{R}$. To see this, consider the sets $Z = \{1\}$ and $Y = \{2\}$. Then $Z \not\preceq Y$ and $Y \not\preceq Z$; that is, we cannot use the partial ordering to compare X and Z. However, if $x, y \in \mathbb{R}$, then either $x \leq y$ or $y \leq x$. Consequently, a partial ordering is nicer if it has the following property:

Definition A.5.3. Let X be a set. A partial ordering $\preceq$ on X is called a *total ordering* if for all $x, y \in X$, either $x \preceq y$ or $y \preceq x$ (or both).

Of course, we desire to equip a set with a partial ordering. Thus we give the following name to such an object.

Definition A.5.4. A *partially ordered set* (or *poset*) is a pair $(X, \preceq)$ where X is a non-empty set and $\preceq$ is a partial ordering on X.

Our main focus is a 'result' about totally ordered subsets of partially ordered sets:

Definition A.5.5. Let $(X, \preceq)$ be a partially ordered set. A non-empty subset $Y \subseteq X$ is said to be a *chain* if Y is totally ordered with respect to $\preceq$; that is, if $a, b \in Y$, then either $a \preceq b$ or $b \preceq a$.

Clearly any non-empty subset of a totally ordered set is a chain. Here is a less obvious example.

Example A.5.6. Recall that the power set $\mathcal{P}(\mathbb{R})$ of $\mathbb{R}$ has a partial ordering $\preceq$ where

$$A \preceq B \quad \Longleftrightarrow \quad A \subseteq B.$$

If $Y = \{A_n\}_{n=1}^{\infty} \subseteq \mathcal{P}(\mathbb{R})$ are such that $A_n \subseteq A_{n+1}$ for all $n \in \mathbb{N}$, then Y is a chain.

Like with the real numbers, upper bounds play an important role with respect to chains.

Definition A.5.7. Let $(X, \preceq)$ be a partially ordered set. A non-empty subset $Y \subseteq X$ is said to be a *bounded above* if there exists a $z \in X$ such that $y \leq z$ for all $y \in Y$. Such an element z is said to be an *upper bound* for Y.

Example A.5.8. Recall from Example A.5.6 that if $Y = \{A_n\}_{n=1}^\infty \subseteq \mathcal{P}(\mathbb{R})$ are such that $A_n \subseteq A_{n+1}$ for all $n \in \mathbb{N}$, then Y is a chain with respect to the partial ordering defined by inclusion. If

$$A = \bigcup_{n=1}^\infty A_n$$

then clearly $A \in \mathcal{P}(\mathbb{R})$ and $A_n \subseteq A$ for all $n \in \mathbb{N}$. Hence A is an upper bound for Y.

Recall there are optimal upper bounds of subsets of $\mathbb{R}$ called least upper bounds which need not be in the subset. We desire a slightly different object when it comes to partially ordered sets as the lack of a total ordering means there may not be a unique 'optimal' upper bound.

Definition A.5.9. Let X be a non-empty set and let $\preceq$ be a partial ordering on X. An element $x \in X$ is said to be *maximal* if there does not exist a $y \in X \setminus \{x\}$ such that $x \preceq y$; that is, there is no element of X that is larger than x with respect to $\preceq$.

Notice that $\mathbb{R}$ together with its usual ordering $\leq$ does not have a maximal element (by, for example, the Archimedean Property). However, many partially ordered sets do have maximal elements. For example $([0,1], \leq)$ has 1 as a maximal element (although $((0,1), \leq)$ does not).

For an example involving a partial ordering that is not a total ordering, suppose $X = \{x, y, z, w\}$ and $\preceq$ is defined such that $a \preceq a$ for all $a \in X$, $a \preceq b$ for all $a \in \{x, y\}$ and $b \in \{z, w\}$, and $a \not\preceq b$ for all other pairs $(a, b) \in X \times X$. It is not difficult to see that z and w are maximal elements and x and y are not maximal elements. Thus it is possible, when dealing with a partial ordering that is not a total ordering, to have multiple maximal elements.

The result we require for the next subsection may now be stated using the above notions.

Axiom A.5.10 (Zorn's Lemma). *Let $(X, \preceq)$ be a non-empty partially ordered set. If every chain in X has an upper bound, then X has a maximal element.*

We will not prove Zorn's Lemma. To do so, we would need to use the Axiom of Choice. In fact, Zorn's Lemma and the Axiom of Choice are logically equivalent; that is, assuming the axioms of (Zermelo–Fraenkel) set theory, one may use the Axiom of Choice to prove Zorn's Lemma, and one may use Zorn's Lemma to prove the Axiom of Choice.

As a simple example of the use of Zorn's Lemma, we present the following.

Example A.5.11. Let V be a (non-zero) vector space. We claim that V has a basis; that is, a linearly independent spanning set. To see this, let

$\mathcal{L}$ denote the collection of all linearly independent subsets of V (which is clearly non-empty) and define a partial ordering on $\mathcal{L}$ by $A \preceq B$ if and only if $A \subseteq B$ (clearly this is a partial ordering on $\mathcal{L}$).

To invoke Zorn's Lemma (Axiom A.5.10), we need to demonstrate that every chain in $\mathcal{L}$ has an upper bound. Let $\{A_\alpha\}_{\alpha\in I}$ be an arbitrary chain in $\mathcal{L}$ and let

$$A = \bigcup_{\alpha\in I} A_\alpha.$$

We claim that $A \in \mathcal{L}$. To see this, suppose $\vec{v}_1, \ldots, \vec{v}_n \in A$ and $a_1\vec{v}_1 + \cdots a_n\vec{v}_n = 0$ for some scalars a_k. By the definition of A and the fact that $\{A_\alpha\}_{\alpha\in I}$ is a chain, there exists an $i \in I$ such that $\vec{v}_1, \ldots, \vec{v}_n \in A_i$ (that is, each $\vec{v}_k$ is in some A_α and as the A_α are totally ordered, take the largest). Hence, as A_i is a linearly independent set, $a_1\vec{v}_1 + \cdots a_n\vec{v}_n = 0$ implies $a_1 = \cdots = a_n = 0$. Hence $A \in \mathcal{L}$. As A is clearly an upper bound for $\{A_\alpha\}_{\alpha\in I}$, ever chain in $\mathcal{L}$ has an upper bound.

By Zorn's Lemma there exists a maximal element $B \in \mathcal{L}$. We claim that B is a basis for V. To see this, suppose to the contrary that $\mathrm{span}(B) \neq V$. Thus there exists a non-zero vector $\vec{v} \in V \setminus \mathrm{span}(B)$. This implies that $B \cup \{\vec{v}\}$ is linearly independent. However, as $B \preceq B \cup \{\vec{v}\}$ and $B \neq B \cup \{\vec{v}\}$, we have a contradiction to the fact that B is a maximal element in $\mathcal{L}$. Hence it must have been the case that $\mathrm{span}(B) = V$ and thus B is a basis for V.

A.6 The Well-Ordering Theorem

Zorn's Lemma (Axiom A.5.10) has another application that is useful in topology; namely the existence of 'nice' orderings of arbitrary sets. To describe these orderings and the desired result, we require the following strengthening of total orderings.

Definition A.6.1. A total ordering $\preceq$ on a set X is said to be a *well-ordering* if for all non-empty $Y \subseteq X$ there exists an $y_0 \in Y$ such that $y_0 \preceq y$ for all $y \in Y$; that is, every subset of X has a least element.

Example A.6.2. The canonical ordering $\leq$ on $\mathbb{R}$ is not a well-ordering since $\mathbb{R}$ has no least element with respect to $\leq$. However, $\leq$ is a well-ordering on $\mathbb{N}$ by Peano's Axioms.

Although the canonical ordering on $\mathbb{R}$ is not a well-ordering, $\mathbb{R}$ does have a well-ordering as the following result demonstrates.

Theorem A.6.3 (Well-Ordering Theorem). *A well-ordering can be placed on every non-empty set.*

Proof. Let X be a non-empty set. Let

$$\mathcal{I} = \{(A, \leq_A) \mid A \subseteq X, \leq_A \text{ is a well-ordering of } A\}.$$

Note that $\mathcal{I} \neq \emptyset$ as $X \neq \emptyset$ and every singleton in X has a trivial well-ordering to be equipped with.

Given $(A, \leq_A), (B, \leq_B) \in \mathcal{I}$, define $(A, \leq_A) \preceq (B, \leq_B)$ if and only if $A \subseteq B$, $\leq_B|_{A\times A} = \leq_A$, and if $a \in A$ and $b \in B$ are such that $b \leq_B a$ then $b \in A$. We claim that $\preceq$ is a partial ordering on $\mathcal{I}$. To begin to see this, we note clearly if $(A, \leq_A), (B, \leq_B) \in \mathcal{I}$, then $(A, \leq_A) \preceq (A, \leq_A)$ and if $(A, \leq_A) \preceq (B, \leq_B)$ and $(B, \leq_B) \preceq (A, \leq_A)$ then $(A, \leq_A) = (B, \leq_B)$. Lastly, suppose $(A, \leq_A), (B, \leq_B), (C, \leq_C) \in \mathcal{I}$ are such that $(A, \leq_A) \preceq (B, \leq_B)$ and $(B, \leq_B) \preceq (C, \leq_C)$. Clearly this implies that $A \subseteq C$ and $\leq_C|_A = \leq_A$. Finally, suppose $a \in A$ and $c \in C$ are such that $c \leq_C a$. As $a \in A \subseteq B$ and $(B, \leq_B) \preceq (C, \leq_C)$, we have that $c \in B$. Thus $c \leq_C a$ implies $c \leq_B a$. Hence, as $(A, \leq_A) \preceq (B, \leq_B)$, we have that $c \in A$ as desired. Hence $(\mathcal{I}, \preceq)$ is a partially ordered set.

To invoke Zorn's Lemma (Axiom A.5.10) we need to demonstrate that every chain in $\mathcal{I}$ has an upper bound. Let $\mathcal{C} = \{(A_\alpha, \leq_\alpha)\}_{\alpha\in J}$ be an arbitrary chain in $\mathcal{I}$. Let

$$A = \bigcup_{\alpha\in J} A_\alpha \subseteq X$$

For $a, b \in A$, define $a \leq b$ if and only if there exists an $\alpha_0 \in J$ such that $a, b \in A_{\alpha_0}$ and $a \leq_{\alpha_0} b$. We claim that $\leq$ is a well-ordering on A thereby showing that $(A, \leq) \in \mathcal{I}$.

To see this, we first note that $\leq$ is trivially reflexive as each $(A_\alpha, \leq_\alpha)$ is reflexive. To see that $\leq$ is antisymmetric, let $a, b \in A$ be arbitrary elements such that $a \leq b$ and $b \leq a$. By the definition of $\leq$, there exists $\alpha_1, \alpha_2 \in J$ such that $a, b \in A_{\alpha_1}$, $a \leq_{\alpha_1} b$, $a, b \in A_{\alpha_2}$, and $b \leq_{\alpha_2} a$. As $\mathcal{C}$ is a chain, by reversing the roles of a and b if necessary, we may assume that $(A_{\alpha_1}, \leq_{\alpha_1}) \preceq (A_{\alpha_2}, \leq_{\alpha_2})$. Hence, by the definition of $\preceq$ we have that $a, b \in A_{\alpha_2}$, $a \leq_{\alpha_1} b$, and $b \leq_{\alpha_2} a$. Therefore, as $(A_{\alpha_2}, \leq_{\alpha_2})$ is a total ordering, $a = b$. Hence, as a and b were arbitrary, $\leq$ is antisymmetric.

To see that $\leq$ is transitive, let $a, b, c \in A$ be arbitrary elements such that $a \leq b$ and $b \leq c$. By the definition of $\leq$, there exists $\alpha_1, \alpha_2 \in J$ such that $a, b \in A_{\alpha_1}$, $a \leq_{\alpha_1} b$, $a, c \in A_{\alpha_2}$, and $b \leq_{\alpha_2} c$. As $\mathcal{C}$ is a chain, if $\alpha = \max(\{\alpha_1, \alpha_2\})$ (which is either α_1 or α_2, we have by the definition of $\preceq$ that $a, b, c \in A_\alpha$, $a \leq_\alpha b$, and $b \leq_\alpha c$. Therefore, as $(A_\alpha, \leq_\alpha)$ is a partial ordering, $a \leq_\alpha c$ thereby showing that $a \leq c$. Hence, as $a, b, c \in A$ were arbitrary, $\leq$ is transitive.

To see that $\leq$ is a total ordering, let $a, b \in A$ be arbitrary elements. Since $\mathcal{C}$ is a chain, there exists an $\alpha \in J$ such that $a, b \in A_\alpha$. Hence, as $(A_\alpha, \leq_\alpha)$ is a total ordering, either $a \leq_\alpha b$ and thus $a \leq b$, or $b \leq_\alpha a$ and thus $b \leq a$. Hence, as $a, b \in A$ where arbitrary, $\leq$ is a total ordering.

To see that $\leq$ is a well-ordering, let $B \subseteq A$ be an arbitrary non-empty set. Due to the definition of A, there exists an $\alpha_0 \in J$ such that $B \cap A_{\alpha_0} \neq \emptyset$. Hence, as $(A_{\alpha_0}, \leq_{\alpha_0})$ is a well-ordering, there exists an element $x \in B \cap A_\alpha \subseteq B$ such that $x \leq_{\alpha_0} y$ for all $y \in B \cap A_\alpha$. We claim that x is a least element of B with respect to $\leq$. To see this, let $b \in B$ be arbitrary. By the definition of A, there exists an $\alpha_b \in J$ such that $b \in B \cap A_{\alpha_b}$. Since $\mathcal{C}$ is a chain, either $(A_{\alpha_0}, \leq_{\alpha_0}) \preceq (A_{\alpha_b}, \leq_{\alpha_b})$ or $(A_{\alpha_b}, \leq_{\alpha_b}) \preceq (A_{\alpha_0}, \leq_{\alpha_0})$. If $(A_{\alpha_0}, \leq_{\alpha_0}) \preceq (A_{\alpha_b}, \leq_{\alpha_b})$, then if $b \leq_{\alpha_b} x$, the definition of $\preceq$ implies that $b \in A_{\alpha_0}$ and thus $x \leq_{\alpha_b} b$ so $x \leq b$. Otherwise, if $(A_{\alpha_b}, \leq_{\alpha_b}) \preceq (A_{\alpha_0}, \leq_{\alpha_0})$, then $b \in B \cap A_{\alpha_b} \subseteq B \cap A_{\alpha_0}$ so $x \leq_{\alpha_0} b$ by the definition of x and thus $x \leq b$. Hence, as $b \in B$ was arbitrary, x is a least element of B. Therefore, as $B \subseteq A$ was arbitrary, $\leq$ is a well-ordering on A. Hence $(A, \leq) \in \mathcal{I}$ as desired.

By the definition of $\preceq$, is it is clear that $(A, \leq)$ is an upper bound for $\mathcal{C}$. Therefore, as $\mathcal{C}$ was arbitrary, every chain in $\mathcal{I}$ has an upper bound. Hence Zorn's Lemma implies there exists a maximal element $(M, \leq_M)$ of $\mathcal{I}$.

We claim that $M = X$. To see this, suppose to the contrary that there exists an $x \in X \setminus M$. Let $M' = M \cup \{x\}$ and for $a, b \in M'$, define $a \leq_0 b$ if and only if $a, b \in M$ and $a \leq_M b$ or if $a \in M$ and $b = x$. Clearly $\leq_0$ is a well-ordering on M' as $\leq_M$ is a well-ordering and as $a \leq_0 x$ for all $a \in M$. Thus $(M', \leq_0) \in \mathcal{I}$. Moreover, it is clear by construction that $(M, \leq_M) \preceq (M', \leq_0)$ by the definitions of $\preceq$ and $\leq_0$. As $(M, \leq_M) \neq (M', \leq_0)$, it contradicts the fact that $(M, \leq_M)$ is a maximal element of $\mathcal{I}$. Hence it must have been the case that $M = X$.

Since $M = X$, $\leq_M$ is a well-ordering on X by the definition of $\mathcal{I}$. Thus the result holds. ∎

Although we will not demonstrate it here, the Well-Ordering Theorem (Theorem A.6.3) is logically equivalent to Zorn's Lemma and thus the Axiom of Choice; that is, assuming the axioms of (Zermelo–Fraenkel) set theory, one may use the Well-Ordering Theorem to prove Zorn's Lemma, and one may use Zorn's Lemma to prove the Well-Ordering Theorem as we did above.

Appendix B

ℓ_p Spaces

In this chapter, we will look at some topological properties of ℓ_p spaces.

B.1 p-Norms

To see that $\ell_p(\mathbb{K})$ is a vector space on $\mathbb{K}$ and the p-norms as described in Examples 1.2.12 and 1.2.13 are indeed norms, it suffices to verify the Triangle Inequality for these p-norms. This will be done by verifying Minkowski's Inequality (Theorem B.1.3).

To do so, we will need to develop some additional inequalities. First fix $p \in (1, \infty)$ and consider the function $f : (1, \infty) \to (1, \infty)$ defined by $f(x) = \frac{x}{x-1}$. Using elementary calculus, f is a bijection between these two sets. In particular, for each $p \in (1, \infty)$ there exists a $q \in (1, \infty)$ such that $p = \frac{q}{q-1}$. Thus

$$\frac{1}{p} = \frac{q-1}{q} = 1 - \frac{1}{q}.$$

Hence for each $p \in (1, \infty)$ there exists a $q \in (1, \infty)$ such that $\frac{1}{p} + \frac{1}{q} = 1$.

Lemma B.1.1 (Young's Inequality). *Let $a, b \geq 0$ and let $p, q \in (1, \infty)$ be such that $\frac{1}{p} + \frac{1}{q} = 1$. Then $ab \leq \frac{1}{p}a^p + \frac{1}{q}b^q$.*

Proof. Notice $1 = \frac{1}{p} + \frac{1}{q} = \frac{p+q}{pq}$ implies $p + q - pq = 0$. Hence $q = \frac{p}{p-1}$.

Fix $b \geq 0$. Notice if $b = 0$, the inequality easily holds. Thus we will assume $b > 0$.

Define $f : [0, \infty) \to \mathbb{R}$ by $f(x) = \frac{1}{p}x^p + \frac{1}{q}b^q - bx$. Clearly $f(0) > 0$ and $\lim_{x\to\infty} f(x) = \infty$ as $p > 1$ so x^p grows faster than x. We claim that $f(x) \geq 0$ for all $x \in [0, \infty)$ thereby proving the inequality. Notice f is differentiable on $[0, \infty)$ with

$$f'(x) = x^{p-1} - b.$$

Therefore $f'(x) = 0$ if and only if $x = b^{\frac{1}{p-1}}$. Moreover, it is elementary to see from the derivative that f has a local minimum at $b^{\frac{1}{p-1}}$ and thus f has a global minimum at $b^{\frac{1}{p-1}}$ due to the boundary conditions. Therefore, since

$$f\left(b^{\frac{1}{p-1}}\right) = \frac{1}{p}b^{\frac{p}{p-1}} + \frac{1}{q}b^q - b^{1+\frac{1}{p-1}} = \frac{1}{p}b^q + \frac{1}{q}b^q - b^q = 0,$$

we obtain that $f(x) \geq 0$ for all $x \in [0, \infty)$ as desired. ∎

Using Young's Inequality, we have a stepping stone towards the Triangle Inequality.

Theorem B.1.2 (Hölder's Inequality). *Let $p, q \in (1, \infty)$ be such that $\frac{1}{p} + \frac{1}{q} = 1$. For any sequences $(x_n)_{n\geq 1}$ and $(y_n)_{n\geq 1}$ with complex entries,*

$$\sum_{k=1}^{\infty} |x_k y_k| \leq \left(\sum_{k=1}^{\infty} |x_k|^p\right)^{\frac{1}{p}} \left(\sum_{k=1}^{\infty} |y_k|^q\right)^{\frac{1}{q}}.$$

Proof. Let $\alpha = \left(\sum_{k=1}^{\infty} |x_k|^p\right)^{\frac{1}{p}}$ and let $\beta = \left(\sum_{k=1}^{\infty} |y_k|^q\right)^{\frac{1}{q}}$. It is clear that $\alpha = 0$ implies $x_k = 0$ for all k which implies $\sum_{k=1}^{\infty} |x_k y_k| = 0$ and thus the inequality will hold in this case. Similarly if $\beta = 0$, then the inequality holds. Subsequently, if $\alpha = \infty$ or $\beta = \infty$ the result holds. Hence we may assume that $0 < \alpha, \beta < \infty$.

Since $0 < \alpha, \beta < \infty$, we obtain that

$$\begin{aligned}
\sum_{k=1}^{\infty} |x_k y_k| &= \alpha\beta \sum_{k=1}^{\infty} \left|\frac{x_k}{\alpha}\right| \left|\frac{y_k}{\beta}\right| \\
&\leq \alpha\beta \left(\sum_{k=1}^{\infty} \frac{1}{p}\left|\frac{x_k}{\alpha}\right|^p + \frac{1}{q}\left|\frac{y_k}{\beta}\right|^q\right) \qquad \text{by Lemma B.1.1} \\
&= \alpha\beta \left(\frac{1}{p\alpha^p}\sum_{k=1}^{\infty} |x_k|^p + \frac{1}{q\beta^q}\sum_{k=1}^{\infty} |y_k|^q\right) \\
&= \alpha\beta \left(\frac{1}{p} + \frac{1}{q}\right) \\
&= \alpha\beta
\end{aligned}$$

as desired. ∎

Theorem B.1.3 (Minkowski's Inequality). *Let $p \in (1, \infty)$. For any sequences $(x_n)_{n\geq 1}$ and $(y_n)_{n\geq 1}$ with complex entries,*

$$\left(\sum_{k=1}^{\infty} |x_k + y_k|^p\right)^{\frac{1}{p}} \leq \left(\sum_{i=1}^{\infty} |x_k|^p\right)^{\frac{1}{p}} + \left(\sum_{k=1}^{\infty} |y_k|^p\right)^{\frac{1}{p}}.$$

Proof. Choose $q \in (1, \infty)$ so that $\frac{1}{p} + \frac{1}{q} = 1$. Thus $q = \frac{p}{p-1}$. Since $p \in (1, \infty)$, notice by Hölder's Inequality that

$$\begin{aligned}
&\sum_{k=1}^{\infty} |x_k + y_k|^p \\
&= \sum_{k=1}^{\infty} (|x_k + y_k|)(|x_k + y_k|)^{p-1} \\
&\leq \sum_{k=1}^{\infty} (|x_k| + |y_k|)(|x_k + y_k|)^{p-1} \\
&= \sum_{k=1}^{\infty} |x_k|(|x_k + y_k|)^{p-1} + \sum_{k=1}^{\infty} |y_k|(|x_k + y_k|)^{p-1} \\
&\leq \left(\sum_{i=1}^{\infty} |x_k|^p\right)^{\frac{1}{p}} \left(\sum_{k=1}^{\infty} (|x_k + y_k|^{p-1})^q\right)^{\frac{1}{q}} + \left(\sum_{k=1}^{\infty} |y_k|^p\right)^{\frac{1}{p}} \left(\sum_{k=1}^{\infty} (|x_k + y_k|^{p-1})^q\right)^{\frac{1}{q}} \\
&= \left(\left(\sum_{k=1}^{\infty} |x_k|^p\right)^{\frac{1}{p}} + \left(\sum_{i=1}^{\infty} |y_k|^p\right)^{\frac{1}{p}}\right) \left(\sum_{k=1}^{\infty} |x_k + y_k|^p\right)^{\frac{1}{q}}.
\end{aligned}$$

If $\sum_{k=1}^{\infty} |x_k + y_k|^p = 0$, the result follows trivially. Otherwise, we may divide both sides of the equation by $(\sum_{i=1}^{\infty} |x_k + y_k|^p)^{\frac{1}{q}}$ to obtain that

$$\left(\sum_{k=1}^{\infty} |x_k + y_k|^p\right)^{\frac{1}{p}} = \left(\sum_{k=1}^{\infty} |x_k + y_k|^p\right)^{1-\frac{1}{q}} \leq \left(\sum_{k=1}^{\infty} |x_k|^p\right)^{\frac{1}{p}} + \left(\sum_{k=1}^{\infty} |y_k|^p\right)^{\frac{1}{p}}$$

as desired. ∎

B.2 The Weak Topology on $\ell_1(\mathbb{R})$

In this section, we will demonstrate that sequences are not enough to distinguish topologies as there are distinct topologies on $\ell_1(\mathbb{R})$ that have the same convergence sequences.

Consider $\ell_1(\mathbb{R})$ and recall from the previous section that $\|\cdot\|_1$ is a norm on $\ell_1(\mathbb{R})$ and thus induces a topology, which will be denoted $\mathcal{T}_n$. To construct another topology on $\ell_1(\mathbb{R})$, for each $\vec{x} = (x_n)_{n\geq 1} \in \ell_1(\mathbb{R})$, for each $\vec{z} = (z_n)_{n\geq 1} \in \ell_\infty(\mathbb{R})$, and each $\epsilon > 0$, define

$$S_{\vec{x},\vec{z},\epsilon} = \left\{ (y_n)_{n\geq 1} \in \ell_1(\mathbb{R}) \,\middle|\, \left|\sum_{n=1}^{\infty} (x_n - y_n) z_n\right| < \epsilon \right\}.$$

The collection of all $S_{\vec{x},\vec{z},\epsilon}$ form a subbasis for a topology $\mathcal{T}_w$ on $\ell_1(\mathbb{R})$ known as the *weak topology*.

Of course, it is more interesting to know the neighbourhood basis surrounding each point in a topological space.

Lemma B.2.1. *For every $\vec{x} = (x_n)_{n\geq 1} \in \ell_1(\mathbb{R})$, for every $m \in \mathbb{N}$, for every $\vec{z}_k = (z_{k,n})_{n\geq 1} \in \ell_\infty(\mathbb{R})$ for $k \in \{1, \ldots, m\}$, and for every $\epsilon > 0$, define $B_{\vec{x},\{\vec{z}_k\}_{k=1}^m,\epsilon}$ to be the set*

$$\left\{ (y_n)_{n\geq 1} \in \ell_1(\mathbb{R}) \;\middle|\; \left| \sum_{n=1}^\infty (x_n - y_n) z_{k,n} \right| < \epsilon \text{ for all } k \in \{1, \ldots, m\} \right\}.$$

The collection over m, $\{\vec{z}_k\}_{k=1}^m$, and ϵ of all $B_{\vec{x},\{\vec{z}_k\}_{k=1}^m,\epsilon}$ is a $\mathcal{T}_w$-neighbourhood basis for $\vec{x}$.

Proof. Fix $\vec{x} = (x_n)_{n\geq 1} \in \ell_1(\mathbb{R})$. To see that the set over m, $\{\vec{z}_k\}_{k=1}^m$, and ϵ of all $B_{\vec{x},\{\vec{z}_k\}_{k=1}^m,\epsilon}$ is a $\mathcal{T}_w$-neighbourhood basis for $\vec{x}$, we must show that every $\mathcal{T}_w$-neighbourhood of $\vec{x}$ contains one of these sets. Thus let U be an arbitrary $\mathcal{T}_w$-neighbourhood of x. Since the collection of all $S_{\vec{x}',\vec{z}',\epsilon'}$ forms a subbasis for $\mathcal{T}_w$, there must exist $\{\vec{x}_k\}_{k=1}^m \subseteq \ell_1(\mathbb{R})$, $\{\vec{z}_k\}_{k=1}^m \subseteq \ell_\infty(\mathbb{R})$, and $\epsilon_1, \ldots, \epsilon_m > 0$ such that

$$\vec{x} \in \bigcap_{k=1}^m S_{\vec{x}_k,\vec{z}_k,\epsilon_k} \subseteq U.$$

Consider a fixed $k \in \{1, \ldots, m\}$. As $\vec{x} \in S_{\vec{x}_k,\vec{z}_k,\epsilon_k}$, if $\vec{x}_k = (x_{k,n})_{n\geq 1}$ and $\vec{z}_k = (z_{k,n})_{n\geq 1}$, then

$$r_k := \left| \sum_{n=1}^\infty (x_n - x_{k,n}) z_{k,n} \right| < \epsilon_k.$$

We claim that $B_{\vec{x},\vec{z}_k,\epsilon_k - r_k} \subseteq S_{\vec{x}_k,\vec{z}_k,\epsilon_k}$. Indeed, let $\vec{y} = (y_n)_{n\geq 1} \in B_{\vec{x},\vec{z}_k,\epsilon_k - r_k}$ be arbitrary. Then

$$\left| \sum_{n=1}^\infty (x_n - y_n) z_{k,n} \right| < \epsilon_k - r_k.$$

Therefore, since every sequence in $\ell_1(\mathbb{R})$ is absolutely summable and as scaling each entry in an absolutely summable series by a bounded sequence preserves the absolute summability, we are dealing with absolutely summable series so we can rearrange the series as desired to obtain that

$$\begin{aligned} \left| \sum_{n=1}^\infty (y_n - x_{k,n}) z_{k,n} \right| &= \left| \sum_{n=1}^\infty (y_n - x_n) z_{k,n} + (x_n - x_{k,n}) z_{k,n} \right| \\ &\leq \left| \sum_{n=1}^\infty (y_n - x_n) z_{k,n} \right| + \left| \sum_{n=1}^\infty (x_n - x_{k,n}) z_{k,n} \right| \\ &< \epsilon_k - r_k + r_k = \epsilon_k. \end{aligned}$$

Hence, as $\vec{y} \in B_{\vec{x},\vec{z}_k,\epsilon_k - r_k}$ was arbitrary, $B_{\vec{x},\vec{z}_k,\epsilon_k - r_k} \subseteq S_{\vec{x}_k,\vec{z}_k,\epsilon_k}$ for every $k \in \{1, \ldots, m\}$.

Finally, let $\delta = \min\{\epsilon_k - r_k \mid k \in \{1, \ldots, m\}\} > 0$. Then, as

$$B_{\vec{x},\vec{z}_k,\delta} \subseteq B_{\vec{x},\vec{z}_k,\epsilon_k - r_k}$$

for all k, we see that

$$\vec{x} \in B_{\vec{x},\{\vec{z}_k\}_{k=1}^m,\delta} = \bigcap_{k=1}^m B_{\vec{x},\vec{z}_k,\delta} \subseteq \bigcap_{k=1}^m B_{\vec{x},\vec{z}_k,\epsilon_k - r_k} \subseteq \bigcap_{k=1}^m S_{\vec{x}_k,\vec{z}_k,\epsilon_k} \subseteq U.$$

Therefore, as U was arbitrary, the proof is complete. ∎

Using Lemma B.2.1 along with a version of Hölder's inequality (Theorem B.1.2) for $\ell_1(\mathbb{R})$ and $\ell_\infty(\mathbb{R})$, we obtain the following.

Proposition B.2.2. *The norm topology on $\ell_1(\mathbb{R})$ is finer than the weak topology on $\ell_1(\mathbb{R})$.*

Proof. To see that $\mathcal{T}_w \subseteq \mathcal{T}_n$, let $U \in \mathcal{T}_w$ and $\vec{x} = (x_n)_{n\geq 1} \in U$ be arbitrary. By Lemma B.2.1 there exist an $m \in \mathbb{N}$, $\vec{z}_k = (z_{k,n})_{n\geq 1} \in \ell_\infty(\mathbb{R})$ for $k \in \{1, \ldots, m\}$, and $\epsilon > 0$ such that $\vec{x} \in B_{\vec{x},\{\vec{z}_k\}_{k=1}^m,\epsilon} \subseteq U$. Let

$$M = \max\{\|\vec{z}_k\| \mid k \in \{1, \ldots, m\}\} > 0.$$

We claim that the $\|\cdot\|_1$-ball V of radius $\frac{\epsilon}{M+1}$ centred at $\vec{x}$ is contained in $B_{\vec{x},\{\vec{z}_k\}_{k=1}^m,\epsilon}$. Indeed if $y = (y_n)_{n\geq 1} \in V$, then

$$\sum_{n=1}^\infty |x_n - y_n| < \frac{\epsilon}{M}.$$

Hence for all $k \in \{1, \ldots, m\}$ we see that

$$\left|\sum_{n=1}^\infty (x_n - y_n) z_{k,n}\right| \leq \sum_{n=1}^\infty |x_n - y_n||z_{k,n}| \leq \sum_{n=1}^\infty |x_n - y_n| M < \frac{\epsilon}{M+1} M < \epsilon.$$

Thus $\vec{y} \in B_{\vec{x},\{\vec{z}_k\}_{k=1}^m,\epsilon}$. Therefore, as $\vec{y}$ is arbitrary, we obtain that $V \subseteq B_{\vec{x},\{\vec{z}_k\}_{k=1}^m,\epsilon}$. Hence

$$\vec{x} \in V \subseteq B_{\vec{x},\{\vec{z}_k\}_{k=1}^m,\epsilon} \subseteq U.$$

Therefore, as $\vec{x} \in U$ was arbitrary, $U \in \mathcal{T}_n$ and thus, as $U \in \mathcal{T}_w$ was arbitrary, $\mathcal{T}_w \subseteq \mathcal{T}_n$. ∎

To emphasize that the weak topology is not the norm topology, we note the following.

Proposition B.2.3. *The weak topology on $\ell_1(\mathbb{R})$ is not a topology induced by a norm. Hence the weak and norm topologies on $\ell_1(\mathbb{R})$ differ.*

Proof. Suppose $\mathcal{T}_w$ was a topology induced by a norm $\|\cdot\|_w$. Let U be the $\|\cdot\|_w$-ball of radius 1 centred at the zero vector $\vec{0}$. By Lemma B.2.1 there must exist an $m \in \mathbb{N}$, $\vec{z}_k = (z_{k,n})_{n\geq 1} \in \ell_\infty(\mathbb{R})$ for $k \in \{1, \ldots, m\}$, and $\epsilon > 0$ such that $B_{\vec{0},\{\vec{z}_k\}_{k=1}^m,\epsilon} \subseteq U$.

For each $k \in \{1, \ldots, m\}$ let

$$\vec{v}_k = (z_{k,1}, z_{k,2}, \ldots, z_{k,m+1}) \in \mathbb{R}^{m+1}.$$

Then the set $\{\vec{v}_k \mid k \in \{1, \ldots, m\}\}$ is a set with m vectors in $\mathbb{R}^{m+1}$. Hence there exists a non-zero vector $\vec{v} = (x_1, x_2, \ldots, x_{m+1}) \in \mathbb{R}^{m+1}$ such that

$$0 = \vec{v} \cdot \vec{v}_k = \sum_{j=1}^{m+1} x_j z_{k,j}$$

for all $k \in \{1, \ldots, m\}$.

Since $\vec{v} \in \mathbb{R}^{m+1}$ is non-zero, if we define $x_j = 0$ if $j > m+1$, then the sequence $\vec{x} = (x_n)_{n\geq 1}$ is a non-zero element of $\ell_\infty(\mathbb{R})$ such that for all $t \in \mathbb{R}$

$$\sum_{n=1}^{\infty} t x_n z_{k,n} = 0$$

for all $k \in \{1, \ldots, m\}$. Hence $t\vec{x} \in B_{\vec{0},\{\vec{z}_k\}_{k=1}^m,\epsilon} \subseteq U$. Hence, as U is the $\|\cdot\|_w$-ball of radius 1 centred at the zero vector $\vec{0}$, $t\vec{x} \in U$ for all $t \in \mathbb{R}$ implies that $|t|\,\|\vec{x}\| = \|t\vec{x}\| < 1$ for all $t \in \mathbb{R}$. However, this is impossible as $\vec{x} \neq \vec{0}$ so $\|\vec{x}\| > 0$. Thus we have a contradiction so $\mathcal{T}_w$ cannot be induced by a norm. ∎

Although the weak and norm topologies on $\ell_1(\mathbb{R})$ are different topologies, they do have the same convergent sequences. Hence sequences are not enough to discuss convergence in topological spaces!

Theorem B.2.4. *A sequence $(\vec{x}_n)_{n\geq 1}$ converges to a point $\vec{x}$ in $(\ell_1(\mathbb{R}), \mathcal{T}_n)$ if and only if $(\vec{x}_n)_{n\geq 1}$ converges to $\vec{x}$ in $(\ell_1(\mathbb{R}), \mathcal{T}_w)$.*

Proof. Let $(\vec{x}_k)_{k\geq 1}$ be a sequence in $\ell_1(\mathbb{R})$. By Proposition B.2.2, we know that if $(\vec{x}_n)_{n\geq 1}$ converges to a point $\vec{x}$ in $(\ell_1(\mathbb{R}), \mathcal{T}_n)$ then $(\vec{x}_n)_{n\geq 1}$ converges to $\vec{x}$ in $(\ell_1(\mathbb{R}), \mathcal{T}_w)$.

To see the converse, notice that $(\vec{x}_k)_{k\geq 1}$ converges to a vector $\vec{x} \in \ell_1(\mathbb{R})$ in the norm topology if and only if $\lim_{k\to\infty} \|\vec{x}_k - \vec{x}\|_1 = 0$ if and only if the sequence $(\vec{x}_k - \vec{x})_{k\geq 1}$ converges to $\vec{0}$ in the norm topology. Similarly, due to the basis of the weak topology considered in Lemma B.2.1, $(\vec{x}_k)_{k\geq 1}$ converges to a vector $\vec{x} \in \ell_1(\mathbb{R})$ in the weak topology if and only if the sequence $(\vec{x}_k - \vec{x})_{k\geq 1}$ converges to $\vec{0}$ in the weak topology. Thus, for the purposes of this question, we need only consider sequences that converge to zero in the weak topology. Therefore, to proceed, suppose $(\vec{x}_k)_{k\geq 1}$ is a sequence in $\ell_1(\mathbb{R})$ that converges to $\vec{0}$ in the weak topology.

To see that $(\vec{x}_k)_{k\geq 1}$ converges to $\vec{0}$ in the norm topology, suppose to the contrary that $(\vec{x}_k)_{k\geq 1}$ does not converge to $\vec{0}$ in the norm topology. Thus there exists an $\epsilon > 0$ and a subsequence $(\vec{x}_{k_j})_{j\geq 1}$ such that $\left\|\vec{x}_{k_j}\right\|_{j\geq 1} \geq \epsilon$ for all $j \in \mathbb{N}$. Therefore, by replacing $(\vec{x}_k)_{k\geq 1}$ with $(\vec{x}_{k_j})_{j\geq 1}$ if necessary, we may assume that $(\vec{x}_k)_{k\geq 1}$ converges to $\vec{0}$ in the weak topology (i.e. a subsequence of a convergent sequence still converges) and that there exists an $\delta > 0$ such that $\|\vec{x}_k\| \geq \delta$ for all $k \in \mathbb{N}$.

Write $\vec{x}_k = (x_{k,n})_{n\geq 1}$ for all $k \in \mathbb{N}$. We claim for each $m \in \mathbb{N}$ that $\lim_{n\to\infty} x_{m,n} = 0$. Indeed fix $m \in \mathbb{N}$ and let $\vec{e}_m = (e_{m,n})_{n\geq 1}$ where

$$e_{m,n} = \begin{cases} 1 & \text{if } n = m \\ 0 & \text{otherwise} \end{cases}.$$

Notice for any $\epsilon > 0$ that $B_{\vec{0},\vec{e}_m,\epsilon}$ is a neighbourhood of $\vec{0}$. Thus, as $(\vec{x}_k)_{k\geq 1}$ converges to $\vec{0}$ in the weak topology, there exists an $N \in \mathbb{N}$ such that $\vec{x}_k \in B_{\vec{0},\vec{e}_m,\epsilon}$ for all $k \geq N$. Hence

$$|x_{k,m}| = \left|\sum_{n=1}^{\infty} x_{k,n}e_{m,n}\right| < \epsilon$$

for all $k \geq N$. Therefore, as $\epsilon > 0$ was arbitrary, $\lim_{n\to\infty} x_{m,n} = 0$ as desired.

Using the facts that $\|\vec{x}_k\| \geq \delta$ for all $k \in \mathbb{N}$ and that $\lim_{n\to\infty} x_{m,n} = 0$ for all $m \in \mathbb{N}$, we will obtain a contradiction to the fact that $(\vec{x}_k)_{k\geq 1}$ converges weakly to $\vec{0}$ by constructing a subsequence that does not converge weakly to $\vec{0}$. Let $k_1 = 1$ and let $n_1 \in \mathbb{N}$ be such that $\sum_{j=n_1+1}^{\infty} |x_{k_1,j}| < \frac{\delta}{6}$; which is possible since $\vec{x}_{k_1} \in \ell_1(\mathbb{N})$. As $\lim_{n\to\infty} x_{m,n} = 0$ for all $m \in \mathbb{N}$, there exists a $k_2 > k_1$ such that $\sum_{j=1}^{n_1} |x_{k,j}| < \frac{\delta}{6}$ for all $k \geq k_2$. Thus, as $\vec{x}_{k_2} \in \ell_1(\mathbb{R})$, there exists an $n_2 > n_1$ such that $\sum_{j=n_2+1}^{\infty} |x_{k_2,j}| < \frac{\delta}{6}$.

By repeating the above construction inductively, we obtain increasing sequences $(k_m)_{m\geq 1}$ and $(n_m)_{m\geq 1}$ such that $\sum_{j=1}^{n_{m-1}} |x_{k,j}| < \frac{\delta}{6}$ for all $k \geq k_m$ and $\sum_{j=n_m+1}^{\infty} |x_{k_m,j}| < \frac{\delta}{6}$.

Consider the subsequence $(\vec{x}_{k_m})_{m\geq 1}$. As $(\vec{x}_k)_{k\geq 1}$ converges weakly to 0, $(\vec{x}_{k_m})_{m\geq 1}$ converges weakly to 0. However, consider $y = (y_n)_{n\geq 1} \in \ell_\infty(\mathbb{R})$ defined by

$$y_n = \begin{cases} \mathrm{sgn}(x_{1,n}) & \text{if } n \leq n_1 \\ \mathrm{sgn}(x_{k_m,n}) & \text{whenever} \quad n_{m-1} + 1 \leq n \leq n_m \end{cases}$$

where

$$\mathrm{sgn}(x) = \begin{cases} 1 & \text{if } x \geq 0 \\ -1 & \text{if } x < 0 \end{cases}.$$

We claim that $\vec{x}_{k_m} \notin B_{\vec{0},\vec{y},\frac{\delta}{3}}$ for all $m \in \mathbb{N}$ thereby contradicting the fact that $(\vec{x}_{k_m})_{m\geq 1}$ converges weakly to 0. Indeed notice by construction that

$$\begin{aligned}\delta \leq \|\vec{x}_{k_m}\|_1 &= \sum_{j=1}^{n_{m-1}} |x_{k_m,j}| + \sum_{j=n_{m-1}+1}^{n_m} |x_{k_m,j}| + \sum_{j=n_m+1}^{\infty} |x_{k_m,j}| \\ &\leq \frac{\delta}{6} + \sum_{j=n_{m-1}+1}^{n_m} |x_{k_m,j}| + \frac{\delta}{6}\end{aligned}$$

so

$$\sum_{j=n_{m-1}+1}^{n_m} |x_{k_m,j}| \geq \frac{2\delta}{3}.$$

However, notice that

$$\begin{aligned}\left|\sum_{j=1}^{\infty} x_{k_m,j}y_j\right| &= \left|\sum_{j=1}^{n_{m-1}} x_{k_m,j}y_j + \sum_{j=n_{m-1}+1}^{n_m} |x_{k_m,j}| + \sum_{j=n_m+1}^{\infty} x_{k_m,j}y_j\right| \\ &\geq \sum_{j=n_{m-1}+1}^{n_m} |x_{k_m,j}| - \sum_{j=1}^{n_{m-1}} |x_{k_m,j}| - \sum_{j=n_m+1}^{\infty} |x_{k_m,j}| \\ &\geq \frac{2\delta}{3} - \frac{\delta}{6} - \frac{\delta}{6} = \frac{\delta}{3}.\end{aligned}$$

Hence $\vec{x}_{k_m} \notin B_{\vec{0},\vec{y},\frac{\delta}{3}}$ for all $m \in \mathbb{N}$ thereby yielding a contradiction and the proof. ∎

As an immediate corollary, we have the following.

Corollary B.2.5. *The weak topology on $\ell_1(\mathbb{R})$ is not metrizable.*

Proof. Suppose the weak topology on $\ell_1(\mathbb{R})$ is induced by a metric. Then Theorem B.2.4 along with Theorem 1.5.28 implies that $\mathcal{T}_n = \mathcal{T}_w$. However this contradicts Proposition B.2.3. Hence the weak topology on $\ell_1(\mathbb{R})$ is not induced by a metric. ∎

Appendix C

The Stone-Čech Compactification of $\mathbb{N}$

In this chapter, an explicit description of the Stone-Čech Compactification of $\mathbb{N}$ will be constructed which is useful for many applications. This description requires us to delve into an alternative approach to topology where the concepts of nets is replaced with the concepts of filters and ultrafilters. As there exists a notion of taking a limits along ultrafilters, it is possible to obtain analogous results to those obtained with nets. The reason nets were used throughout these notes instead of ultrafilters is that we feel nets behave in a far more similar fashion to sequences than ultrafilters do thereby enabling an easier progression from undergraduate real analysis to this graduate level topology course. However, ultrafilters do have their uses.

C.1 Ultrafilters

To begin, we define the notions of a filter and an ultrafilter in the most general setting possible.

Definition C.1.1. Let $(X, \preceq)$ be a partially ordered set (see Definition A.5.1). A subset $\mathcal{F} \subseteq X$ is said to be a *filter on* $(X, \preceq)$ if the following conditions hold:

- $\mathcal{F} \neq \emptyset$,
- (existence of lower bounds) for every $x, y \in \mathcal{F}$ there exists a $z \in \mathcal{F}$ such that $z \preceq x$ and $z \preceq y$, and
- (upper set) if $y \in \mathcal{F}$ and $x \in X$ are such that $y \preceq x$, then $x \in \mathcal{F}$.

A *proper filter* is a filter $\mathcal{F}$ on $(X, \preceq)$ such that $\mathcal{F} \neq X$.

An *ultrafilter* is a maximal proper filter; that is, a proper filter $\mathcal{F}_0$ on $(X, \preceq)$ is said to be an ultrafilter if whenever $\mathcal{F}$ is a proper filter on $(X, \preceq)$ such that $\mathcal{F}_0 \subseteq \mathcal{F}$, then $\mathcal{F} = \mathcal{F}_0$.

For topological applications, only certain filters and ultrafilters are required. The collection of such objects is given by the following definition.

Definition C.1.2. Let X be a non-empty set. For $A, B \subseteq X$, define $A \preceq B$ if and only if $A \subseteq B$ so that $(\mathcal{P}(X), \preceq)$ is a partially ordered set. A *filter on* X is a filter on $(\mathcal{P}(X), \subseteq)$. An *ultrafilter on* X is a ultrafilter on $(\mathcal{P}(X), \subseteq)$.

In the specific setting of filters and ultrafilters on a set, we have the following equivalent formulations.

Lemma C.1.3. *Let X be a non-empty set. A set $\mathcal{F} \subseteq \mathcal{P}(X)$ is a filter on X if and only if*

- $\mathcal{F} \neq \emptyset$,
- *if* $A \subseteq B \subseteq X$ *and* $A \in \mathcal{F}$ *then* $B \in \mathcal{F}$, *and*
- *if* $A, B \in \mathcal{F}$ *then* $A \cap B \in \mathcal{F}$,

Proof. This trivially follows from Definition C.1.1. ∎

Lemma C.1.4. *Let X be a non-empty set. A set $\mathcal{F} \subseteq \mathcal{P}(X)$ is an ultrafilter on X if and only if*

(A) if $A \subseteq B \subseteq X$ *and* $A \in \mathcal{F}$ *then* $B \in \mathcal{F}$,

(B) if $A, B \in \mathcal{F}$ *then* $A \cap B \in \mathcal{F}$,

(C) if $A \subseteq X$ *then either* $A \in \mathcal{F}$ *or* $X \setminus A \in \mathcal{F}$, *and*

(D) $\emptyset \notin \mathcal{F}$.

Note properties (B), (C), and (D) together immediate imply that if $A \subseteq X$ then either $A \in \mathcal{F}$ or $X \setminus A \in \mathcal{F}$ but not both (for otherwise $\emptyset = A \cap (X \setminus A) \in \mathcal{F}$).

Proof. First suppose that $\mathcal{F}$ has the four properties on set inclusions listed. We desire to prove that $\mathcal{F}$ is an ultrafilter. To begin this process, we first note that property (C) implies that $\mathcal{F}$ is non-empty. This fact together with properties (A) and (B) immediately imply that $\mathcal{F}$ is a filter on X by Lemma C.1.3. Moreover, by property (D), $\emptyset \notin \mathcal{F}$ so $\mathcal{F} \neq \mathcal{P}(X)$ and thus $\mathcal{F}$ is a proper filter on X.

To see that $\mathcal{F}$ is an ultrafilter on X, suppose to the contrary that $\mathcal{F}_0$ is a proper filter on X such that $\mathcal{F} \subsetneq \mathcal{F}_0$. Hence there exists an $A \in \mathcal{F}_0 \setminus \mathcal{F}$. As $A \notin \mathcal{F}$, property (C) on $\mathcal{F}$ implies that $X \setminus A \in \mathcal{F} \subseteq \mathcal{F}_0$. Hence $X \setminus A \in \mathcal{F}_0$ and $A \in \mathcal{F}_0$. Since $\mathcal{F}_0$ is a filter on X, $\emptyset = A \cap (X \setminus A) \in \mathcal{F}_0$. Therefore, since $\mathcal{F}_0$ is an upper set and $\emptyset \subseteq A$ for all $A \in \mathcal{P}(X)$, $\mathcal{F}_0 = \mathcal{P}(X)$ thereby contradicting the fact that $\mathcal{F}_0$ was a proper filter of X. Hence is an ultrafilter on X as desired.

Conversely, suppose that $\mathcal{F}$ is an ultrafilter on X. We desire to show that $\mathcal{F}$ has the four properties. To begin, Lemma C.1.3 immediately implies properties (A) and (B). Next, note that if $\emptyset \in \mathcal{F}$, then the fact that $\mathcal{F}$ is an upper set implies that $\mathcal{F} = \mathcal{P}(X)$ thereby contradicting the fact that $\mathcal{F}$ is an ultrafilter and thus a proper filter. Hence $\emptyset \notin \mathcal{F}$ so $\mathcal{F}$ satisfies property (D). Thus all that remains is to show property (C).

To see property (C), let $A \subseteq X$ be arbitrary. We claim that either every element of $\mathcal{F}$ intersects A or every element of $\mathcal{F}$ intersects $X \setminus A$. To see this, suppose to the contrary that there exists $C, D \in \mathcal{F}$ such that $C \cap A = \emptyset$ and $D \cap (X \setminus A) = \emptyset$. As $\mathcal{F}$ has the lower bound property, there exists a $Z \in \mathcal{F}$ such that $Z \subseteq C \subseteq X \setminus A$ and $Z \subseteq D \subseteq A$. Hence $Z = \emptyset$ thereby contradicting the fact that $\emptyset \notin \mathcal{F}$. Therefore, the claim must be true. Therefore, by interchanging A with $X \setminus A$ if necessary, we may assume that every element of $\mathcal{F}$ intersects A.

Let

$$\mathcal{I} = \{A \cap F \mid F \in \mathcal{F}\} \subseteq \mathcal{P}(X) \quad \text{and}$$
$$\mathcal{J} = \{G \subseteq X \mid \text{there exists a } I \in \mathcal{I} \text{ such that } I \subseteq G\} \subseteq \mathcal{P}(X).$$

We claim that $\mathcal{F}_0 = \mathcal{F} \cup \mathcal{I} \cup \mathcal{J}$ is a filter on X. To see this, we note that $\mathcal{F} \neq \emptyset$ as $\mathcal{F}$ is a filter so $\mathcal{F}_0 \neq \emptyset$. Next, to see that $\mathcal{F}_0$ has lower bounds let $B, C \in \mathcal{F}_0$ be arbitrary. Notice since $\mathcal{F}$ is closed under finite intersections by property (B), $\mathcal{I}$ is closed under finite intersections and the intersection of an element of $\mathcal{I}$ with an element of $\mathcal{F}$ is an element of $\mathcal{I}$. Hence if $B, C \in \mathcal{F} \cup \mathcal{I}$, then $B \cap C \in \mathcal{F} \cup I \subseteq \mathcal{F}_0$ is such that $B \cap C \subseteq B$ and $B \cap C \subseteq C$. Otherwise, if $B \in \mathcal{J}$ or $C \in \mathcal{J}$, we can replace B and C respectively with an element of $\mathcal{I}$ in order to prove the existence of lower bounds with respect to $\subseteq$. Hence $\mathcal{F}_0$ has lower bounds. Finally, to see that $\mathcal{F}_0$ is an upper set, let $B \in \mathcal{F}_0$ and $C \subseteq X$ be such that $B \subseteq C$. If $B \in \mathcal{F}$, then $C \in \mathcal{F} \subseteq \mathcal{F}_0$ as $\mathcal{F}$ is a filter. Next, if $B \in \mathcal{I}$, then $C \in \mathcal{J} \subseteq \mathcal{F}_0$ by definition. Finally, if $B \in \mathcal{J}$, then there exists an $I \in \mathcal{I}$ such that $I \subseteq B \subseteq C$ thereby showing that $C \in \mathcal{J} \subseteq \mathcal{F}_0$ by definition. Hence $\mathcal{F}_0$ is a filter on X.

Since $\emptyset \notin \mathcal{F}$ and since $A \cap F \neq \emptyset$ for all $F \in \mathcal{F}$, we see that $\emptyset \notin \mathcal{F} \cup \mathcal{I}$. Hence $\emptyset \notin \mathcal{J}$ by definition. Thus $\mathcal{F}_0$ is a proper filter on X. However, as $\mathcal{F} \subseteq \mathcal{F}_0$ and $\mathcal{F}$ is an ultrafilter, it must be the case that $\mathcal{F} = \mathcal{F}_0$. Therefore, since $X \in \mathcal{F}$ as $\mathcal{F} \neq \emptyset$ and $\mathcal{F}$ is an upper set, we see that $A = A \cap X \in \mathcal{I} \subseteq \mathcal{F}_0 = \mathcal{F}$ as desired. ∎

Of course, we have not explicitly given examples of ultrafilters on a set. Before doing so, it is useful to divide the examples of ultrafilters into two classes.

Definition C.1.5. An ultrafilter $\mathcal{F}$ on X is said to be a *principal ultrafilter* if $\mathcal{F}$ has a least element; that is, there exists an $A \in \mathcal{F}$ such that $A \subseteq B$ for all $B \in \mathcal{F}$ (and thus $\mathcal{F} = \{B \subseteq X \mid A \subseteq B\}$ by the upper set property of filters). Any ultrafilter that is not principal is said to be a *free ultrafilter*.

Describing and constructing examples of principal ultrafilters is easy.

Proposition C.1.6. *Let X be a non-empty set. For each $x \in X$, let*

$$\mathcal{F}_x = \{A \subseteq X \mid x \in A\}.$$

Then $\{\mathcal{F}_x\}_{x\in X}$ is precisely the collection of all principal ultrafilters on X.

Proof. Fix $x \in X$. We claim that $\mathcal{F}_x$ is a principal ultrafilter on X. To see this, first it must be demonstrated that $\mathcal{F}_x$ is an ultrafilter. To begin this process, first note that if $A \subseteq B \subseteq C$ are such that $A \in \mathcal{F}_X$, then $x \in A \subseteq B$ so $B \in \mathcal{F}_x$ by definition. Similarly, if $A, B \in \mathcal{F}_x$, then $x \in A$ and $x \in B$ so $x \in A \cap B$ and thus $A \cap B \in \mathcal{F}_x$ by definition. Next, if $A \subseteq X$, then either $x \in A$ or $x \in X \setminus A$ and thus either $A \in \mathcal{F}_x$ or $X \setminus A \in \mathcal{F}_x$ by definition. Finally, as $x \notin \emptyset$, $\emptyset \notin \mathcal{F}_x$. Hence $\mathcal{F}_x$ is an ultrafilter on X. To see that $\mathcal{F}_x$ is a principal ultrafilter, we note that

$$\mathcal{F}_x = \{A \subseteq X \mid \{x\} \subseteq A\}$$

thereby showing that $\mathcal{F}_x$ is a principal ultrafilter.

To see that $\{\mathcal{F}_x\}_{x\in X}$ is exactly the set of principal ultrafilters on X, let $\mathcal{F}$ be an arbitrary principal ultrafilter of X. Hence there exists an $A \subseteq X$ such that

$$\mathcal{F} = \{B \subseteq X \mid A \subseteq B\}.$$

If $A = \emptyset$ then $\emptyset = A \in \mathcal{F}$ thereby contradicting the fact that $\mathcal{F}$ is an ultrafilter. Next, suppose that there exist $x, y \in A$ such that $x \neq y$. Thus as A is not a subset of $\{x\}$ and not a subset of $X \setminus \{x\}$, we obtain that $\{x\} \notin \mathcal{F}$ and $X \setminus \{x\} \notin \mathcal{F}$. However, as this contradicts the third property of $\mathcal{F}$ being an ultrafilter on $(\mathcal{P}(X), \subseteq)$, it must be case that A is a singleton. Thus $A = \{x\}$ for some $x \in X$ so $\mathcal{F} = \mathcal{F}_x$. Therefore, as $\mathcal{F}$ was arbitrary, $\{\mathcal{F}_x\}_{x\in X}$ is exactly the set of all principal ultrafilters on X. ∎

Free ultrafilters are a more complicated beast then principal ultrafilters as they will not be so easy to describe. In fact, it is impossible to given an explicit description of a free ultrafilter on the natural numbers! However, we can exhibit the existence of free ultrafilters by constructing a free filter and using maximality.

Theorem C.1.7. *Let X be a non-empty set and let $\mathcal{F}_0$ be a proper filter on X. Then there exists an ultrafilter $\mathcal{U}$ on X such that $\mathcal{F}_0 \subseteq \mathcal{U}$.*

Proof. Note since $\mathcal{F}_0$ is a proper filter, Lemma C.1.3 implies that $\emptyset \notin \mathcal{F}_0$ for otherwise $\mathcal{F}_0$ being an upper set would imply $\mathcal{F}_0 = \mathcal{P}(X)$ contradicting the fact that $\mathcal{F}_0$ was a proper filter on X.

Let

$$\mathcal{Y} = \{\mathcal{F} \subseteq \mathcal{P}(X) \mid \mathcal{F} \text{ is a filter on } X \text{ such that } \mathcal{F}_0 \subseteq \mathcal{F} \text{ and } \emptyset \notin \mathcal{F}\}.$$

Notice that $(\mathcal{Y}, \subseteq)$ is a non-empty poset since $\mathcal{F}_0 \in \mathcal{Y}$. Thus we desire to apply Zorn's Lemma (Axiom A.5.10) to obtain a maximal element of $\mathcal{Y}$.

To see that the assumptions of Zorn's Lemma hold for $(\mathcal{Y}, \subseteq)$, let $\mathcal{C}$ be a non-empty chain in $\mathcal{Y}$. To see that $\mathcal{C}$ has an upper bound in $(\mathcal{Y}, \subseteq)$, let

$$\mathcal{B} = \bigcup_{C \in \mathcal{C}} C.$$

Clearly $C \subseteq \mathcal{B}$ for all $C \in \mathcal{C}$. Hence, provided $\mathcal{B} \in \mathcal{Y}$, $\mathcal{C}$ has an upper bound in $\mathcal{Y}$.

To see that $\mathcal{B} \in Y$, first note since $\mathcal{F}_0 \subseteq C$ for all $C \in \mathcal{C}$ that $\mathcal{F}_0 \subseteq \mathcal{B}$. Furthermore, since $\emptyset \notin C$ for all $C \in \mathcal{C}$, $\emptyset \notin \mathcal{B}$. To see that $\mathcal{B}$ is a filter on X, we will verify that $\mathcal{B}$ satisfies the conditions of Lemma C.1.3. To begin, we note since $\mathcal{F}_0$ is a filter on X and thus non-empty, $\mathcal{B}$ is non-empty as $\mathcal{F}_0 \subseteq \mathcal{B}$. Next, let $A \subseteq B \subseteq X$ be arbitrary elements such that $A \in \mathcal{B}$. By the definition of $\mathcal{B}$, there exists a $C \in \mathcal{C}$ such that $A \in C$. Hence, since C is a filter, $A \in C$ and $A \subseteq B \subseteq X$ implies that $B \in C \subseteq \mathcal{B}$ as desired. Finally, let $A, B \in \mathcal{B}$ be arbitrary elements. By the definition of $\mathcal{B}$ and since $\mathcal{C}$ is a chain, there exists a $C \in \mathcal{C}$ such that $A, B \in C$. Therefore, since C is a filter, $A \cap B \in C \subseteq \mathcal{B}$ as desired. Hence $\mathcal{B}$ is a filter on X by Lemma C.1.3.

Therefore, Zorn's Lemma implies that there is a maximal element $\mathcal{U}$ of $\mathcal{Y}$. Clearly this implies $\mathcal{F}_0 \subseteq \mathcal{U}$. Thus, to complete the proof, it suffices to show that $\mathcal{U}$ is an ultrafilter on X. To see this, we first note by the definition of $\mathcal{Y}$ that $\emptyset \notin \mathcal{U}$. Hence $\mathcal{U}$ is a proper filter on X. To see that $\mathcal{U}$ is an ultrafilter, let $\mathcal{F}$ be an arbitrary proper filter of $(\mathcal{P}(X), \subseteq)$ such that $\mathcal{U} \subseteq \mathcal{F}$. Clearly as $\mathcal{F}_0 \subseteq \mathcal{U}$, this implies $\mathcal{F}_0 \subseteq \mathcal{F}$. Moreover, since $\mathcal{F}$ is a proper filter, Lemma C.1.3 implies that $\emptyset \notin \mathcal{F}$ for otherwise $\mathcal{F}$ being an upper set would imply $\mathcal{F} = \mathcal{P}(X)$ contradicting the fact that $\mathcal{F}$ was a proper filter on X. Hence $\mathcal{F} \in \mathcal{Y}$ by the definition of $\mathcal{Y}$. Therefore, since $\mathcal{U}$ is a maximal element of $\mathcal{Y}$ and since $\mathcal{U} \subseteq \mathcal{F} \in \mathcal{Y}$, it must be the case that $\mathcal{U} = \mathcal{F}$. Hence $\mathcal{U}$ is an ultrafilter on X containing $\mathcal{F}_0$ as desired. ∎

Using Theorem C.1.7, we can demonstrate the existence of free ultrafilters.

Example C.1.8. Let X be an infinite set and let

$$\mathcal{F} = \{A \subseteq X \mid X \setminus A \text{ is finite}\}.$$

We claim that $\mathcal{F}$ is a proper filter on $(\mathcal{P}(X), \subseteq)$. Indeed clearly $\mathcal{F} \neq \emptyset$ and $\mathcal{F} \neq \mathcal{P}(X)$ by construction. Next let $A \subseteq B \subseteq X$ be arbitrary sets such that $A \in \mathcal{F}$. Thus $X \setminus A$ is finite so $A \subseteq B$ implies that $X \setminus B$ is finite and hence $B \in \mathcal{F}$ as desired. Finally, let $A, B \in \mathcal{F}$ be arbitrary. Thus $X \setminus A$ and $X \setminus B$ are finite sets. Since

$$X \setminus (A \cap B) = (X \setminus A) \cup (X \setminus B)$$

is finite, $A \cap B \in \mathcal{F}$ as desired. Therefore, $\mathcal{F}$ is a proper filter on $(\mathcal{P}(X), \subseteq)$.

By Theorem C.1.7, there exists an ultrafilter $\mathcal{U}$ on X containing $\mathcal{F}$. We claim that $\mathcal{U}$ is a free ultrafilter. To see this, we note by Proposition C.1.6 that if for each $x \in X$ we let

$$\mathcal{F}_x = \{A \subseteq X \mid x \in A\},$$

then $\{\mathcal{F}_x\}_{x\in X}$ are the set of all principal ultrafilters on X. We claim that $\mathcal{U} \neq \mathcal{F}_x$ for all $x \in X$ thereby proving that $\mathcal{U}$ is a free ultrafilter on X. To see this, suppose to the contrary that $\mathcal{U} = \mathcal{F}_x$ for some $x \in X$. Let $A = X \setminus \{x\}$. As $X \setminus A = \{x\}$ is finite, $A \in \mathcal{F} \subseteq \mathcal{U} = \mathcal{F}_x$ thereby contradicting the fact that $A \notin \mathcal{F}_x$ by the definition of $\mathcal{F}_x$. Hence $\mathcal{U}$ is a free ultrafilter as desired.

To complete this section, we note the useful fact that property (C) in Lemma C.1.4 can be extended further.

Lemma C.1.9. *If X is a non-empty set, if $\mathcal{F}$ is an ultrafilter on $(\mathcal{P}(X), \subseteq)$, and if $\{A_k\}_{k=1}^n \subseteq \mathcal{P}(X)$ are subsets such that $\bigcup_{k=1}^n A_k \in \mathcal{F}$, then $A_k \in \mathcal{F}$ for some $k \in \{1, \ldots, n\}$.*

Proof. To proceed, suppose to the contrary that $A_k \notin \mathcal{F}$ for all $k \in \{1, \ldots, n\}$. Hence property (C) of Lemma C.1.4 implies that $X \setminus A_k \in \mathcal{F}$ for all $k \in \{1, \ldots, n\}$. Thus, by using property (B) $n-1$ times, we see that

$$X \setminus \left(\bigcup_{k=1}^n A_k \right) = \bigcap_{k=1}^n (X \setminus A_k) \in \mathcal{F}.$$

Hence $\bigcup_{k=1}^n A_k \in \mathcal{F}$ and $X \setminus (\bigcup_{k=1}^n A_k) \in \mathcal{F}$ thereby contradicting the remark in C.1.4. Hence we have a contradiction so it must be the case that $A_k \in \mathcal{F}$ for some $k \in \{1, \ldots, n\}$. ∎

C.2 Limits Along Ultrafilters

As alluded to in the introduction to this chapter, ultrafilters give us an alternate approach to limits that we now describe.

Definition C.2.1. Let $(X, \mathcal{T})$ be a topological space, let I be a non-empty set, and let $\mathcal{F}$ be a proper filter on I. For each $\alpha \in I$, chose a point $x_\alpha \in X$. It is said that the collection $(x_\alpha)_{\alpha \in I}$ *converges along $\mathcal{F}$ to a point $x \in X$* if for every neighbourhood U of x in $(X, \mathcal{T})$ the set $\{\alpha \in I \mid x_\alpha \in U\}$ is an element of $\mathcal{F}$.

As Definition C.2.1 may seem unnatural to any analyst, we provide the simplest example.

Example C.2.2. Let $(X, \mathcal{T})$ be a topological space and let $(x_n)_{n\geq 1}$ be a sequence of points in X. For an $m \in \mathbb{N}$, let

$$\mathcal{F}_m = \{A \subseteq \mathbb{N} \mid m \in A\}$$

which is a principal ultrafilter on $\mathbb{N}$. We claim that $(x_n)_{n\geq 1}$ converges to x_m along $\mathcal{F}_m$. To see this, we note if $U \in \mathcal{T}$ is an arbitrary element such that $x_m \in U$, then

$$m \in \{n \in \mathbb{N} \mid x_n \in U\}$$

so

$$\{n \in \mathbb{N} \mid x_n \in U\} \in \mathcal{F}_m.$$

Hence as U was arbitrary, $(x_n)_{n\geq 1}$ converges to x_m along $\mathcal{F}_m$.

Based on the above example, it is natural to ask whether 'limits' along ultrafilters are unique. Of course, the answer is yes if we are in a Hausdorff space just as in Theorem 1.5.40. In fact, unique limits along ultrafilters is equivalent to Hausdorffness just like the result for nets in Theorem 1.5.42. To show this, we will show a little of the connection between nets and filters. We will not go too far into the depth of the connection as this is an appendix section after all!

The following shows the very natural way to get a filter from a net.

Lemma C.2.3. *Let $(x_\lambda)_{\lambda\in\Lambda}$ be a net in a set X and let*

$$\mathcal{F} = \{F \subseteq \Lambda \mid \text{there exists a } \lambda_0 \in \Lambda \text{ such that } \{\lambda \in \Lambda \mid \lambda \geq \lambda_0\} \subseteq F\}.$$

Then $\mathcal{F}$ is a proper filter on Λ.

Proof. To begin to see this, we note that for any $\lambda_0 \in \Lambda$ that

$$F = \{\lambda \mid \lambda \geq \lambda_0\} \in \mathcal{F}$$

by definition. Hence $\mathcal{F} \neq \emptyset$. Next let $A \subseteq B \subseteq X$ be arbitrary sets such that $A \in \mathcal{F}$. Since $A \in \mathcal{F}$, there exists a $\lambda_0 \in \Lambda$ such that $\lambda \in A$ for all $\lambda \geq \lambda_0$. Thus, as $A \subseteq B$, we see that $\lambda \in B$ for all $\lambda \geq \lambda_0$. Hence $B \in \mathcal{F}$ by definition as desired. Finally, let $F_1, F_2 \in \mathcal{F}$ be arbitrary. Hence there exists $\lambda_1, \lambda_2 \in \Lambda$ such that $\lambda \in F_1$ for all $\lambda \geq \lambda_1$ and $\lambda \in F_2$ for all $\lambda \geq \lambda_2$. As Λ is a directed set, there exists a $\lambda_0 \in \Lambda$ such that $\lambda_0 \geq \lambda_1$ and $\lambda_0 \geq \lambda_2$. Hence the above implies for all $\lambda \geq \lambda_0$ that $\lambda \in F_1 \cap F_2$. Thus $F_1 \cap F_2 \in \mathcal{F}$ by definition as desired. Hence $\mathcal{F}$ is a filter by Lemma C.1.3. Furthermore, as $\emptyset \notin \mathcal{F}$, $\mathcal{F}$ is a proper filter. ∎

It is useful to give a name to the above net.

Definition C.2.4. Let $(x_\lambda)_{\lambda\in\Lambda}$ be a net in a set X. The filter

$$\mathcal{F} = \{F \subseteq \Lambda \mid \text{there exists a } \lambda_0 \in \Lambda \text{ such that } \{\lambda \in \Lambda \mid \lambda \geq \lambda_0\} \subseteq F\}.$$

from Lemma C.2.3 is called the *filter derived from* $(x_\lambda)_{\lambda\in\Lambda}$.

With the above definitions complete, we can connect the convergence of nets and the converges along the filter derived from a net easily.

Lemma C.2.5. *Let $(X, \mathcal{T})$ be a topological space and let $x_0 \in X$. If $(x_\lambda)_{\lambda\in\Lambda}$ is a net in $(X, \mathcal{T})$ and $\mathcal{F}$ is the filter derived from $(x_\lambda)_{\lambda\in\Lambda}$, then $(x_\lambda)_{\lambda\in\Lambda}$ converges to x_0 as a net if and only if $(x_\lambda)_{\lambda\in\Lambda}$ converges along $\mathcal{F}$ to x_0.*

Proof. This easily follows from the definition of a convergent net, from the definition of convergences along a filter, and the definition of the filter derived from a net. ∎

Using the above, we easily obtain the filter version of Theorem 1.5.40 and Theorem 1.5.42.

Theorem C.2.6. *A topological space $(X, \mathcal{T})$ is Hausdorff if and only if for every proper filter $\mathcal{F}$ on a non-empty set I and every collection $(x_\alpha)_{\alpha\in I}$ in X, $(x_\alpha)_{\alpha\in I}$ converges to at most one point along $\mathcal{F}$.*

Proof. First suppose that $(X, \mathcal{T})$ is Hausdorff. To see the result, suppose to the contrary that there exists a proper filter $\mathcal{F}$ on a non-empty set I and a collection $(x_\alpha)_{\alpha\in I}$ in X such that there exists two points $x, y \in X$ with $x \neq y$ such that $(x_\alpha)_{\alpha\in I}$ converges to both x and y along $\mathcal{F}$. As $(X, \mathcal{T})$ is Hausdorff and $x \neq y$, there exists $U, V \in \mathcal{T}$ such that $x \in U$, $y \in V$, and $U \cap V = \emptyset$. By the definition of convergence along a filter, we obtain that

$$\{\alpha \in I \mid x_\alpha \in U\} \in \mathcal{F} \qquad \text{and} \qquad \{\alpha \in I \mid x_\alpha \in V\} \in \mathcal{F}.$$

However, as $\mathcal{F}$ is a filter, this implies that

$$\emptyset = \{\alpha \in I \mid x_\alpha \in U \cap V\} = \{\alpha \in I \mid x_\alpha \in U\} \cap \{\alpha \in I \mid x_\alpha \in V\} \in \mathcal{F}.$$

Hence, as $\mathcal{F}$ is filter and thus an upper set, $\mathcal{F} = \mathcal{P}(I)$ thereby contradicting the fact that $\mathcal{F}$ was a proper filter. Hence we have a contradiction so the result holds.

To see the converse, suppose for every proper filter $\mathcal{F}$ on a non-empty set I and every collection $(x_\alpha)_{\alpha\in I}$ in X, $(x_\alpha)_{\alpha\in I}$ converges to at most one point along $\mathcal{F}$. As this implies every net in $(X, \mathcal{T})$ has at most one point of convergence by Lemma C.2.5, Theorem 1.5.42 implies that $(X, \mathcal{T})$ is Hausdorff. ∎

In fact, we can add to our characterization of compact topological spaces seen in Theorem 3.2.2 with limits along ultrafilters.

Theorem C.2.7. *Let $(X, \mathcal{T})$ be a topological space. The following are equivalent:*

(i) $(X, \mathcal{T})$ *is compact.*

(ii) For every non-empty set I, every ultrafilter $\mathcal{F}$ on I, and every collection of points $(x_\alpha)_{\alpha\in I} \subseteq X$, the collection $(x_\alpha)_{\alpha\in I}$ converges along $\mathcal{F}$ in $(X, \mathcal{T})$.

Proof. To begin, let $(X, \mathcal{T})$ be a compact topological space. Suppose to the contrary that there exists a non-empty set I, an ultrafilter $\mathcal{F}$ on I, and a collection of points $(x_\alpha)_{\alpha\in I} \subseteq X$ such that $(x_\alpha)_{\alpha\in I}$ does not converges along $\mathcal{F}$. Hence for every $x \in X$ there exists a $U_x \in \mathcal{T}$ such that $x \in U_x$ and

$$\{\alpha \in I \mid x_\alpha \in U_x\} \notin \mathcal{F}.$$

Clearly $\{U_x\}_{x\in X}$ is an open cover of $(X, \mathcal{T})$. Therefore, since $(X, \mathcal{T})$ is compact, there exists an $n \in \mathbb{N}$ and $x_1, x_2, \ldots, x_n \in X$ such that $X = \bigcup_{k=1}^n U_{x_k}$. Hence

$$\bigcup_{k=1}^n \{\alpha \in I \mid x_\alpha \in U_{x_k}\} = I \in \mathcal{F}.$$

Therefore, by Lemma C.1.9 there exists a $k_0 \in \{1, \ldots, n\}$ such that

$$\{\alpha \in I \mid x_\alpha \in U_{x_{k_0}}\} \in \mathcal{F}.$$

Since this contradicts the defining property of $U_{x_{k_0}}$, we have a contradiction. Thus one direction of the proof is complete.

To prove the converse, suppose for every non-empty set I, every ultrafilter $\mathcal{F}$ on I, and every collection of points $(x_\alpha)_{\alpha\in I} \subseteq X$, the collection $(x_\alpha)_{\alpha\in I}$ converges along $\mathcal{F}$ in $(X, \mathcal{T})$. To see that $(X, \mathcal{T})$ is compact, suppose to the contrary that $(X, \mathcal{T})$ is not compact. Thus by Theorem 3.2.2 there exists a collection $\{F_\alpha\}_{\alpha\in I}$ of closed subsets of $(X, \mathcal{T})$ with the finite intersection property such that $\bigcap_{\alpha\in I} F_\alpha = \emptyset$.

Let

$$\mathcal{J} = \left\{ \bigcap_{k=1}^n F_{\alpha_k} \,\middle|\, n \in \mathbb{N}, \alpha_1, \ldots, \alpha_n \in I \right\} \subseteq \mathcal{P}(X) \quad \text{and}$$
$$\mathcal{F} = \{G \subseteq X \mid \text{there exists a } J \in \mathcal{J} \text{ such that } J \subseteq G\} \subseteq \mathcal{P}(X).$$

We claim that $\mathcal{F}$ is a proper filter on X. To see this, we first note that $\mathcal{J} \neq \emptyset$ so $X \in \mathcal{F}$ by construction. Next, let $A \subseteq B \subseteq X$ be arbitrary sets such that $A \in \mathcal{F}$. Hence there exists an $J \in \mathcal{J}$ such that $J \subseteq A$. Hence $J \subseteq B$ so $B \in \mathcal{F}$ by definition as desired. Finally, let $A, B \in \mathcal{F}$ be arbitrary. Hence there exists $J_1, J_2 \in \mathcal{J}$ such that $J_1 \subseteq A$ and $J_2 \subseteq B$. However, since $\mathcal{J}$ is closed under finite intersections since $\{F_\alpha\}_{\alpha\in I}$ has the finite intersection property, we see that $J_1 \cap J_2 \in \mathcal{J}$ has the property that $J_1 \cap J_2 \subseteq J_1 \subseteq A$ and $J_1 \cap J_2 \subseteq J_2 \subseteq B$. Hence $A \cap B \in \mathcal{F}$ as desired. Hence $\mathcal{F}$ is a filter on X. Finally, we note that $\emptyset \notin \mathcal{J}$ by the previous paragraph and thus $\emptyset \notin \mathcal{F}$. Hence $\mathcal{F}$ is a proper filter on X.

By Theorem C.1.7 there exists an ultrafilter $\mathcal{U}$ on X. We claim that the collection of points X in X does not converge along $\mathcal{U}$ to any point in $(X, \mathcal{T})$. To see this, suppose to the contrary that the collection of points X in X converges along $\mathcal{U}$ to a point $x_0 \in X$. Since

$$\bigcap_{\alpha\in I} F_\alpha = \emptyset,$$

there exists an $\alpha_0 \in I$ such that $x_0 \in X \setminus F_{\alpha_0}$. Let $U_{\alpha_0} = X \setminus F_{\alpha_0}$, which is an open subset of $(X, \mathcal{T})$ containing x_0. Thus, as X converges along $\mathcal{U}$ to x_0, the set

$$U_{\alpha_0} = \{x \in X \mid x \in U_{\alpha_0}\} \in \mathcal{U}.$$

However, as $F_{\alpha_0} \in \mathcal{J} \subseteq \mathcal{F} \subseteq \mathcal{U}$, we have that $F_{\alpha_0} \in \mathcal{U}$ and $X \setminus F_{\alpha_0} \in \mathcal{U}$ thereby contradicting the fact that $\mathcal{U}$ is an ultrafilter on X by Lemma C.1.4. Hence we have a contradiction thereby showing that the collection of points X in X does not converge along $\mathcal{U}$ to any point in $(X, \mathcal{T})$. As this contradicts the assumptions of this direction of the proof, the result is complete. ∎

C.3 The Stone-Čech Compactification of $\mathbb{N}$

With the above complete, we can now finally develop the descriptions of the Stone-Čech Compactification of $\mathbb{N}$ (equipped with the discrete topology of course). This is due to the fact that the Stone-Čech Compactification of $\mathbb{N}$ can be identified by placing a topology on all ultrafilters on $\mathbb{N}$. As such, we will let $\mathcal{U}(\mathbb{N})$ denote the set of all ultrafilters on $(\mathcal{P}(\mathbb{N}), \subseteq)$. To described the topology on $\mathcal{U}(\mathbb{N})$ in order to make $\mathcal{U}(\mathbb{N})$ the Stone-Čech Compactification of $\mathbb{N}$, we have the following.

Theorem C.3.1. *Let*

$$\mathcal{B} = \{\{\mathcal{F} \in \mathcal{U}(\mathbb{N}) \mid A \in \mathcal{F}\} \mid A \subseteq \mathbb{N}\}.$$

Then $\mathcal{B}$ is a basis for a topology $\mathcal{T}$ on $\mathcal{U}(\mathbb{N})$.

Proof. To see that $\mathcal{B}$ is a basis for a topology on $\mathcal{U}(\mathbb{N})$, we need only verify the two defining properties for a basis for a topology from Definition 1.3.2. For the first, property, let $\mathcal{F}_0 \in \mathcal{U}(\mathbb{N})$ be arbitrary. As $\mathcal{F}_0$ is an ultrafilter, we easily see that $\mathbb{N} \in \mathcal{F}_0$ so

$$\mathcal{F}_0 \in \{\mathcal{F} \in \mathcal{U}(\mathbb{N}) \mid \mathbb{N} \in \mathcal{F}\}.$$

Hence the first property has been verified. To see the second property, let $B_1, B_2 \in \mathcal{B}$ and $\mathcal{F}_0 \in B_1 \cap B_2$ be arbitrary. By the definition of $\mathcal{B}$, there exists $A_1, A_2 \subseteq \mathbb{N}$ such that

$$B_1 = \{\mathcal{F} \in \mathcal{U}(\mathbb{N}) \mid A_1 \in \mathcal{F}\} \qquad \text{and} \qquad B_2 = \{\mathcal{F} \in \mathcal{U}(\mathbb{N}) \mid A_2 \in \mathcal{F}\}.$$

Let $A_3 = A_1 \cap A_2$ and let

$$B_3 = \{\mathcal{F} \in \mathcal{U}(\mathbb{N}) \mid A_3 \in \mathcal{F}\} \in \mathcal{B}.$$

To see that B_3 is the desired element of $\mathcal{B}$, first note since $\mathcal{F}_0 \in B_1$ and $\mathcal{F}_0 \in B_2$ that $A_1 \in \mathcal{F}_0$ and $A_2 \in \mathcal{F}_0$. Therefore, since $\mathcal{F}_0$ is an ultrafilter, $A_1 \cap A_2 \in \mathcal{F}_0$ and thus $\mathcal{F}_0 \in B_3$. To see that $B_3 \subseteq B_1 \cap B_2$, let $\mathcal{F} \in B_3$

be an arbitrary. By the definition of B_3 we know that $A_3 \in \mathcal{F}$. Therefore, since $\mathcal{F}$ is an ultrafilter, since $A_3 \subseteq A_1$, and since $A_3 \subseteq A_2$, we have that $\mathcal{F} \in B_1$ and $\mathcal{F} \in B_2$ so $\mathcal{F} \in B_1 \cap B_2$. Therefore, since $\mathcal{F} \in B_3$ was arbitrary, $\mathcal{F}_0 \subseteq B_3 \subseteq B_1 \cap B_2$ as desired. Hence $\mathcal{B}$ is a basis for a topology on $\mathcal{U}(\mathbb{N})$. ∎

Now that we have a topology on $\mathcal{U}(\mathbb{N})$, in order to show that the Stone-Čech Compactification of $\mathbb{N}$ can be identified with $\mathcal{U}(\mathbb{N})$ we have to show that the topology on $\mathcal{U}(\mathbb{N})$ as the desired properties. In particular, we must show (among other things) that $\mathcal{U}(\mathbb{N})$ is Hausdorff, compact, and a compactification of $\mathbb{N}$. Thus we begin with the following.

Theorem C.3.2. *Let $\mathcal{T}$ be the topology on $\mathcal{U}(\mathbb{N})$ from Theorem C.3.1. Then $(\mathcal{U}(\mathbb{N}), \mathcal{T})$ is a Hausdorff space.*

Proof. To see that $(\mathcal{U}(\mathbb{N}), \mathcal{T})$ is a Hausdorff space, let $\mathcal{F}_1, \mathcal{F}_2 \in \mathcal{U}(\mathbb{N})$ be arbitrary elements such that $\mathcal{F}_1 \neq \mathcal{F}_2$. Hence there exists an element $A_0 \subseteq \mathbb{N}$ such that A_0 is in one of $\mathcal{F}_1$ or $\mathcal{F}_2$ but not the other. Therefore, since $\mathcal{F}_1$ and $\mathcal{F}_2$ are ultrafilters (so Lemma C.1.4 implies exactly one of A and $\mathbb{N} \setminus A$ is in a given ultrafilter), there exists an element $A \subseteq \mathbb{N}$ such that $A \in \mathcal{F}_1$, $\mathbb{N} \setminus A \notin \mathcal{F}_1$, $A \notin \mathcal{F}_2$, and $\mathbb{N} \setminus A \in \mathcal{F}_2$. Let

$$U = \{\mathcal{F} \in \mathcal{U}(\mathbb{N}) \mid A \in \mathcal{F}\} \qquad \text{and} \qquad V = \{\mathcal{F} \in \mathcal{U}(\mathbb{N}) \mid \mathbb{N} \setminus A \in \mathcal{F}\}.$$

As U and V are elements of $\mathcal{B}$, U and V are open subsets of $(\mathcal{U}(\mathbb{N}), \mathcal{T})$ that contain $\mathcal{F}_1$ and $\mathcal{F}_2$ respectively. Moreover $U \cap V = \emptyset$ as the properties of an ultrafilter imply exactly one of A and $\mathbb{N} \setminus A$ is in a given ultrafilter. Therefore, since $\mathcal{F}_1, \mathcal{F}_2 \in \mathcal{U}(\mathbb{N})$ were arbitrary, $(\mathcal{U}(\mathbb{N}), \mathcal{T})$ is Hausdorff. ∎

Theorem C.3.3. *Let $\mathcal{T}$ be the topology on $\mathcal{U}(\mathbb{N})$ from Theorem C.3.1. Then $(\mathcal{U}(\mathbb{N}), \mathcal{T})$ is compact.*

Proof. To see that $(\mathcal{U}(\mathbb{N}), \mathcal{T})$ is compact, let $\{U_\alpha\}_{\alpha \in I}$ be an arbitrary open cover of $(\mathcal{U}(\mathbb{N}), \mathcal{T})$. Since $\mathcal{B}$ is a basis for $(\mathcal{U}(\mathbb{N}), \mathcal{T})$, each U_α is a union of elements of $\mathcal{B}$ by Theorem 1.3.4. Let

$$J = \{B \in \mathcal{B} \mid B \subseteq U_\alpha \text{ for some } \alpha \in I\} \subseteq \mathcal{B}.$$

Therefore, as each element U_α is a union of elements of $\{B\}_{B \in J}$, since $\{B\}_{B \in J} \subseteq \mathcal{B}$ are open sets, and since $\{U_\alpha\}_{\alpha \in I}$ is an open cover of $(\mathcal{U}(\mathbb{N}), \mathcal{T})$, $\{B\}_{B \in J}$ is an open cover of $(\mathcal{U}(\mathbb{N}), T)$. Furthermore, notice if we can show that $\{B\}_{B \in J}$ has a finite subcover of $(\mathcal{U}(\mathbb{N}), \mathcal{T})$, then so too will $\{U_\alpha\}_{\alpha \in I}$ as each element of the finite subcover of $\{B\}_{B \in J}$ is contained in one of the elements of $\{U_\alpha\}_{\alpha \in I}$ and combining these finite number of elements yields a finite subcover of $(\mathcal{U}(\mathbb{N}), \mathcal{T})$ from $\{U_\alpha\}_{\alpha \in I}$. Hence it suffices to show that $\{B\}_{B \in J}$ has a finite subcover of $(\mathcal{U}(\mathbb{N}), \mathcal{T})$.

Suppose to the contrary that $\{B\}_{B\in J}$ has no finite subcovers of $(\mathcal{U}(\mathbb{N}), \mathcal{T})$. Note by the definition of $\mathcal{B}$, for each $B \in J$ there exists an $A_B \subseteq \mathbb{N}$ such that

$$B = \{\mathcal{F} \in \mathcal{U}(\mathbb{N}) \mid A_B \in \mathcal{F}\}.$$

Thus, as $\{B\}_{B\in J}$ has no finite subcovers of $(\mathcal{U}(\mathbb{N}), \mathcal{T})$, for every $n \in \mathbb{N}$ and every $B_1, B_2, \ldots, B_n \in J$ there exists an $\mathcal{F} \in \mathcal{U}(\mathbb{N})$ such that $A_{B_k} \notin \mathcal{F}$ for all $k \in \{1, \ldots, n\}$.

Let

$$\begin{aligned} \mathcal{C} &= \{\mathbb{N} \setminus A_B \mid B \in J\} \subseteq \mathcal{P}(\mathbb{N}) \\ \mathcal{I} &= \left\{ \bigcap_{k=1}^n C_k \,\middle|\, n \in \mathbb{N}, \{C_k\}_{k=1}^n \subseteq \mathcal{C} \right\} \subseteq \mathcal{P}(\mathbb{N}) \\ \mathcal{F}_0 &= \{D \subseteq \mathbb{N} \mid \text{there exists an } I \in \mathcal{I} \text{ such that } I \subseteq D\} \subseteq \mathcal{P}(\mathbb{N}). \end{aligned}$$

We claim that $\mathcal{F}_0$ is a filter on $\mathbb{N}$. To begin, we note that $\mathcal{C} \neq \emptyset$ for otherwise $J = \emptyset$ thereby contradicting the fact that $\{B\}_{B\in J}$ is an open cover of $\mathcal{U}(\mathbb{N})$. Hence, as $\mathcal{C} \neq \emptyset$, we obtain that $\mathcal{I} \neq \emptyset$ and thus $\mathcal{F}_0 \neq \emptyset$. Next, suppose $A \subseteq D \subseteq \mathbb{N}$ are arbitrary sets such that $A \in \mathcal{F}_0$. Hence there exists an $I \in \mathcal{I}$ such that $I \subseteq A$. Thus $I \subseteq D$ so $D \in \mathcal{F}_0$ as desired. Finally let $A, D \in \mathcal{F}$ be arbitrary. Hence there exists $I_1, I_2 \in \mathcal{I}$ such that $I_1 \subseteq A$ and $I_2 \subseteq D$. However, since $\mathcal{I}$ is closed under finite intersections by construction, we see that $I_1 \cap I_2 \in \mathcal{I}$ has the property that $I_1 \cap I_2 \subseteq I_1 \subseteq A$ and $I_1 \cap I_2 \subseteq I_2 \subseteq D$. Hence $A \cap D \in \mathcal{F}_0$ as desired. Hence $\mathcal{F}_0$ is a filter on $\mathbb{N}$.

Next we claim that $\mathcal{F}_0$ is a proper filter on $\mathbb{N}$. To see this, it suffices to show that $\emptyset \notin \mathcal{I}$ by the definition of $\mathcal{F}_0$. Thus, suppose to the contrary that $\emptyset \in \mathcal{I}$. Hence there exists an $n \in \mathbb{N}$ and $\{B_k\}_{k=1}^n \subseteq J$ such that

$$\bigcap_{k=1}^n \mathbb{N} \setminus A_{B_k} = \emptyset.$$

Hence

$$\mathbb{N} = \mathbb{N} \setminus \left(\bigcap_{k=1}^n \mathbb{N} \setminus A_{B_k} \right) = \bigcup_{k=1}^n A_{B_k}.$$

However the above equation, the fact that every ultrafilter contains $\mathbb{N}$ by Lemma C.1.4, and Lemma C.1.9 imply that for every $\mathcal{F} \in \mathcal{U}(\mathbb{N})$ that there exists a $k \in \{1, \ldots, n\}$ such that $A_{B_k} \in \mathcal{F}$. However, this contradicts the previous claim that we obtained from the assumption that $\{B\}_{B\in J}$ has no finite subcovers of $(\mathcal{U}(\mathbb{N}), \mathcal{T})$. Hence we have obtained a contradiction so $\emptyset \notin \mathcal{I}$ and thus $\mathcal{F}_0$ is a proper filter on $\mathbb{N}$.

Since $\mathcal{F}_0$ is a proper filter on $\mathbb{N}$, there exists an ultrafilter $\mathcal{F}_u \in \mathcal{U}(\mathbb{N})$ such that $\mathcal{F}_0 \subseteq \mathcal{F}_u$ by Theorem C.1.7. However, as $\mathcal{C} \subseteq \mathcal{I} \subseteq \mathcal{F}_0 \subseteq \mathcal{F}_u$ by construction, we see that $\mathbb{N} \setminus A_B \in \mathcal{F}_u$ for all $B \in J$. Hence Lemma C.1.4 implies that $A_B \notin \mathcal{F}_u$ for all $B \in J$. Hence $\mathcal{F}_u \notin B$ for all $B \in J$ thereby contradicting the fact that $\{B\}_{B\in J}$ is an open cover of $(\mathcal{U}(\mathbb{N}), \mathcal{T})$. Hence we have obtained a contradiction to the assumption that $\{B\}_{B\in J}$ has no finite subcovers of $(\mathcal{U}(\mathbb{N}), \mathcal{T})$. Hence $(\mathcal{U}(\mathbb{N}), \mathcal{T})$ is compact as desired. ∎

With the following, we are getting close.

Theorem C.3.4. *Let* $\mathcal{T}$ *be the topology on* $\mathcal{U}(\mathbb{N})$ *from Theorem C.3.1. Then* $(\mathcal{U}(\mathbb{N}), \mathcal{T})$ *is a compactification of* $\mathbb{N}$.

Proof. First, by Theorem C.3.2 and Theorem C.3.3, $(\mathcal{U}(\mathbb{N}), \mathcal{T})$ is a compact Hausdorff space and thus potentially a compactification of $\mathbb{N}$. Thus, to show that $(\mathcal{U}(\mathbb{N}), \mathcal{T})$ is a compactification of $\mathbb{N}$, we must find an embedding of $\mathbb{N}$ equipped with the discrete topology into $(\mathcal{U}(\mathbb{N}), \mathcal{T})$ with dense range.

For each $n \in \mathbb{N}$ let

$$\mathcal{F}_n = \{A \subseteq \mathbb{N} \mid n \in A\}$$

which is a principal ultrafilter on $\mathbb{N}$ by Proposition C.1.6. Define $\Phi : \mathbb{N} \to \mathcal{U}(\mathbb{N})$ by $\Phi(n) = \mathcal{F}_n$ for all $n \in \mathbb{N}$. As $\mathcal{F}_n \neq \mathcal{F}_m$ for all $n, m \in \mathbb{N}$ with $n \neq m$, we see that Φ is an injective function. Furthermore, as for all $n \in \mathbb{N}$ we know that

$$\mathcal{B}_n = \{\mathcal{F} \in \mathcal{U}(\mathbb{N}) \mid \{n\} \in \mathcal{F}\}$$

is a basis element in $(\mathcal{U}(\mathbb{N}), \mathcal{T})$ that contains $\mathcal{F}_n$ but not $\mathcal{F}_m$ for all $m \in \mathbb{N}\setminus\{n\}$, we see that each singleton in $\Phi(\mathbb{N})$ is open in the subspace topology inherited from $(\mathcal{U}(\mathbb{N}), \mathcal{T})$. Hence Φ is an embedding as desired.

Finally, it suffices to show that $\Phi(\mathbb{N})$ is dense in $(\mathcal{U}(\mathbb{N}), \mathcal{T})$. To see this, let $\mathcal{F}_0 \in \mathcal{U}(\mathbb{N})$ and $U \in \mathcal{T}$ such that $\mathcal{F}_0 \in U$ be arbitrary. By Theorem C.3.1 there exists an $A \subseteq \mathbb{N}$ such that if

$$B = \{\mathcal{F} \in \mathcal{U}(\mathbb{N}) \mid A \in \mathcal{F}\},$$

then $B \in \mathcal{T}$ and $\mathcal{F}_0 \in B \subseteq U$. As $\emptyset$ is not an element of any ultrafilter on $\mathbb{N}$ by Lemma C.1.4, we obtain that $A \neq \emptyset$ for otherwise $B = \emptyset$ thereby contradicting the fact that $\mathcal{F}_0 \in B$. Thus there exists an $n \in A$. Hence, as $\{n\} \subseteq A \subseteq \mathbb{N}$, as $\{n\} \in \mathcal{F}_n$, and as $\mathcal{F}_n$ is an ultrafilter, $A \in \mathcal{F}_n$ by Lemma C.1.4. Hence $\mathcal{F}_n \in \Phi(\mathbb{N}) \cap B \subseteq \Phi(\mathbb{N}) \cap U$ so $\Phi(\mathbb{N}) \cap U \neq \emptyset$. Therefore, as $U \in \mathcal{T}$ was arbitrary, we obtain that $\mathcal{F}_0 \in \overline{\Phi(\mathbb{N})}$ for all $\mathcal{F}_0 \in \mathcal{U}(\mathbb{N})$ by Theorem 1.6.21. Hence $\overline{\Phi(\mathbb{N})} = \mathcal{U}(\mathbb{N})$ thereby completing the proof that $(\mathcal{U}(\mathbb{N}), \mathcal{T})$ is a compactification of $\mathbb{N}$. ∎

Finally, we have arrived at the following.

Theorem C.3.5. *Let* $\mathcal{T}$ *be the topology on* $\mathcal{U}(\mathbb{N})$ *from Theorem C.3.1. Then* $(\mathcal{U}(\mathbb{N}), \mathcal{T})$ *is the Stone-Čech compactification of* $\mathbb{N}$.

Proof. Recall from Theorem C.3.4 that if for each $n \in \mathbb{N}$ we let

$$\mathcal{F}_n = \{A \subseteq \mathbb{N} \mid n \in A\}$$

and we define $\Phi : \mathbb{N} \to \mathcal{U}(\mathbb{N})$ by $\Phi(n) = \mathcal{F}_n$, then Φ is an embedding of $\mathbb{N}$ into $(\mathcal{U}(\mathbb{N}), \mathcal{T})$ such that $\overline{\Phi(\mathbb{N})} = \mathcal{U}(\mathbb{N})$. Therefore, by Theorem 5.4.15, it suffices to prove that if $(X, \mathcal{T}_X)$ is a compact Hausdorff space and if $f \in \mathcal{C}(\mathbb{N}, X)$, then there exists a $g \in \mathcal{C}(\mathcal{U}(\mathbb{N}), X)$ such that $g \circ \Phi = f$ (as every element of $\mathcal{C}_b(\mathbb{N}, \mathbb{R})$ has range contained in a compact subspace of $\mathbb{R}$).

Let $f \in \mathcal{C}(\mathbb{N}, X)$ be arbitrary. To extend f to an element of $\mathcal{C}(\mathcal{U}(\mathbb{N}), X)$, let $\mathcal{F} \in \mathcal{U}(\mathbb{N})$ be arbitrary. Therefore, since $\mathcal{F}$ is an ultrafilter on $\mathbb{N}$ and since $(f(n))_{n\in\mathbb{N}}$ is a collection of points in $(X, \mathcal{T}_X)$, Theorem C.2.7 and Theorem C.2.6 implies there exists a unique point $g(\mathcal{F}) \in X$ such that $(f(n))_{n\in\mathbb{N}}$ converges to $g(\mathcal{F})$ along $\mathcal{F}$. Furthermore, Example C.2.2 implies that $g(\mathcal{F}_n) = f(n)$ for all $n \in \mathbb{N}$. Hence $g \in \mathcal{F}(\mathcal{U}(\mathbb{N}), X)$ is such that $g \circ \Phi = f$. Hence it remains only to show that g is continuous.

To see that g is continuous, let $\mathcal{F}_0 \in \mathcal{U}(\mathbb{N})$ and let $U \in \mathcal{T}_X$ be such that $g(\mathcal{F}_0) \subseteq U$ be arbitrary. Since $(X, \mathcal{T}_X)$ is a compact Hausdorff space, $(X, \mathcal{T}_X)$ is a normal space by Corollary 5.1.25 and thus a regular space. Hence Lemma 5.1.15 impli there exists a $V \in \mathcal{T}_X$ such that

$$g(\mathcal{F}_0) \in V \subseteq \overline{V} \subseteq U.$$

Let

$$A = \{n \in \mathbb{N} \mid f(n) \in V\} \subseteq \mathbb{N}$$

so that

$$B = \{\mathcal{F} \in \mathcal{U}(\mathbb{N}) \mid A \in \mathcal{F}\} \in \mathcal{T}.$$

We claim that $\mathcal{F}_0 \in g^{-1}(V) \subseteq B \subseteq g^{-1}(U)$ thereby showing that B is a neighbourhood of $\mathcal{F}_0$ in $(\mathcal{U}(\mathbb{N}), \mathcal{T})$ contained in $g^{-1}(U)$. Therefore, as $\mathcal{F}_0$ and U were arbitrary, this claim will complete the proof that g is continuous thereby completing the proof.

Clearly $\mathcal{F}_0 \in g^{-1}(V)$ by construction. To see that $g^{-1}(V) \subseteq B$, let $\mathcal{F} \in \mathcal{U}(\mathbb{N})$ be such that $\mathcal{F} \in g^{-1}(V)$. Since $(f(n))_{n\in\mathbb{N}}$ converges to $g(\mathcal{F})$ along $\mathcal{F}$ and since $g(\mathcal{F}) \in V$, we obtain that

$$A = \{n \in \mathbb{N} \mid f(n) \in V\} \in \mathcal{F}$$

by definition of convergence along an ultrafilter. Hence, as $\mathcal{F}$ was arbitrary, $g^{-1}(V) \subseteq B$ as desired.

Finally, to see that $B \subseteq g^{-1}(U)$, suppose to the contrary that there exists an $\mathcal{F}' \in B$ such that $\mathcal{F}' \notin g^{-1}(U)$. Hence $g(\mathcal{F}') \in X \setminus \overline{V}$. Let

$$C = \{n \in \mathbb{N} \mid f(n) \in X \setminus \overline{V}\} \subseteq \mathbb{N}$$

and let

$$B' = \{\mathcal{F} \in \mathcal{U}(\mathbb{N}) \mid C \in \mathcal{F}\} \in \mathcal{T}.$$

As $X \setminus \overline{V}$ is open, the same arguments used above imply that

$$\mathcal{F}' \in g^{-1}(X \setminus \overline{V}) \subseteq B'.$$

Hence $A \in \mathcal{F}'$ and $C \in \mathcal{F}'$ as $\mathcal{F}' \in B$ and $\mathcal{F}' \in B'$ respectively. However, as $V \cap (X \setminus \overline{V}) = \emptyset$, we see that $\emptyset = A \cap C$ thereby contradicting the fact that $A \in \mathcal{F}'$, $C \in \mathcal{F}'$, and $\mathcal{F}'$ is an ultrafilter. Hence, as we have obtained a contradiction, we have that $B \subseteq g^{-1}(U)$ and thus the proof. ∎

To conclude this appendix, we note two things. First we notice that we did not use any properties of $\mathbb{N}$ in the above results; that is, the above easily generalizes to show that the Stone-Čech compactification of any infinite set X equipped with the discrete topology is the set of ultrafilters on X equipped an analogous version of the 'contains a set' topology from Theorem C.3.1.

Secondly, we end with one application of the Stone-Čech compactification of $\mathbb{N}$. First we claim that $\ell_\infty(\mathbb{N}) = C(\beta\mathbb{N})$. To see this, first notice if $f \in C(\beta\mathbb{N})$, then the range of f is bounded and thus the values of f on the principle ultrafilters yield an element of $\ell_\infty(\mathbb{N})$. Conversely, if $g \in \ell_\infty(\mathbb{N})$, then g is a continuous bounded function on $\mathbb{N}$ and thus extends to a unique element of $C(\beta\mathbb{N})$. As this bijection between $\ell_\infty(\mathbb{N})$ and $C(\beta\mathbb{N})$ is linear and an isometry, $\ell_\infty(\mathbb{N})$ and $C(\beta\mathbb{N})$ are equal as Banach spaces. This then means that $\ell_\infty(\mathbb{N})$ and $C(\beta\mathbb{N})$ have the same dual space. As the Riesz Representation Theorem immediately tells us that dual space of $C(\beta\mathbb{N})$ is isomorphic to the finite regular Borel measures on $\beta\mathbb{N}$, we have now determined the dual space of $\ell_\infty(\mathbb{N})$.

Index