1. Introduction

1 Motivation and foreshadowing

Most of the fundamental concepts in this course are generalizations of concepts with which you are familiar from previous courses in analysis and/or linear algebra. For example, you have likely encountered of at least a few of the following concepts:

- The notions of open intervals of the form (a, b) and closed intervals of the form $[a, b]$ in $\mathbb{R}$. More generally, the notion of an open subset of $\mathbb{R}^2$ or $\mathbb{R}^n$, of whose definitions we will see in more detail later.

  The field of point-set topology (on which this course focuses) generalises this concept to a much broader context.

- The definition of a continuous function. For example you probably recall that a function $f : \mathbb{R}^2 \to \mathbb{R}$ is continuous at $\mathbf{a} \in \mathbb{R}^2$ if for every $\varepsilon > 0$ there is a $\delta > 0$ such that $\|\mathbf{x} - \mathbf{a}\| < \delta$ implies $|f(\mathbf{x}) - f(\mathbf{a})| < \varepsilon$.

  This definition is much more concisely and usefully stated in terms of open sets. We will see this early on in this course.

- The idea of "closeness". This is somewhat vague at the moment, but for example what does it mean to say that two points in $\mathbb{R}^n$ are close to one another? What about two subsets of $\mathbb{R}^n$ being close to another one? A point and a set?

  In the context of $\mathbb{R}^n$ we are used to using expressions like "$\|\mathbf{x} - \mathbf{a}\| < \varepsilon$" to denote something like "$\mathbf{x}$ is close to $\mathbf{a}$", but this requires some more context to make sense (namely, what is ε?).

  What if we want to discuss two points being close together in a context in which there is no notion of distance?

- The notion of convergence of sequences. For example you probably recall that a sequence $\{\mathbf{x}_n\}_{n=1}^{\infty}$ in $\mathbb{R}^n$ is said to converge to a point $\mathbf{x} \in \mathbb{R}^n$ if for every $\varepsilon > 0$ there is a $N \in \mathbb{N}$ such that $n > N$ implies $\|\mathbf{x}_n - \mathbf{x}\| < \varepsilon$.

  This is a concept we will be able to generalize significantly. Most notably we will see how to define sequence convergence *without* a notion of distance.

- What does it mean for a subset of $\mathbb{R}^n$, for example, to be "nice"? We recall a number of theorems from first and second year calculus and/or analysis that require assumptions about the domains of functions; for example the Extreme Value Theorem from single variable calculus says that a continuous function defined on a *closed and bounded interval* achieves a minimum and maximum value. Can we generalise this theorem after we generalise the definition of a continuous function, as we hinted at earlier? And if so, what is the analogue of a "closed and bounded" domain in this more general setting?

These are some of the questions we will seek to answer in this course.

2 Topologies in a familiar setting

Broadly speaking, a topology on a given set X is simply a specified collection of subsets of X, which we will call open sets, that satisfies certain properties.

Before we define what a topology is in general, we will define them in the familiar context of $\mathbb{R}^n$. Later in the course we will see that $\mathbb{R}^n$ with its "usual topology" is a *very* well-behaved topological space (a metric space). For the moment, we will just present the basic idea and some intuition behind it.

Throughout the following section, let $x = (x_1, \ldots, x_n)$ and $y = (y_1, \ldots, y_n)$ denote points in $\mathbb{R}^n$.

Recall the notion of an ε-ball around a point x in $\mathbb{R}^n$:

$$B_\varepsilon(x) = \{\, y \in \mathbb{R}^n \, : \, d(x, y) < \varepsilon \,\}$$

where $d(x, y) = \sqrt{(x_1 - y_1)^2 + \cdots + (x_n - y_n)^2}$ is the usual Euclidean distance between x and y. (At the beginning of this note we used the notation $\|x - y\|$ to refer to the same quantity. The notation $d(x, y)$ is preferable for us, because it will generalize later.) Using the idea of an ε-ball, we had a definition of what it meant for a subset of $\mathbb{R}^n$ to be open:

Definition 2.1. *A set $U \subseteq \mathbb{R}^n$ is called open if for every $x \in U$ there is an $\varepsilon > 0$ such that $B_\varepsilon(x) \subseteq U$.*

This is a fine definition, but it is very dependent on the specific structure of $\mathbb{R}^n$. Most notably it makes use of the fact that $\mathbb{R}^n$ has a notion of distance between points. However if we were to carefully go through some of the theorems that have to do with open subsets of $\mathbb{R}^n$, we would find that we do not have to busy ourselves with the distance function. In fact, we would find that much of what we need to show about open sets follows from these three properties:

Fact 2.2.

1. *$\emptyset$ and $\mathbb{R}^n$ are both open.*
2. *The union of any collection of open sets (finite or infinite) is again open.*
3. *The intersection of finitely many open sets is open.*

It turns out that these three are the fundamental properties that open sets should satisfy to allow us to do useful things with them, like defining continuous functions, convergence of sequences, and so on. Everything we do in this course will flow from asking what happens when we fix a collection of sets satisfying these three properties.

3 Topologies and topological spaces

Definition 3.1. *Let X be a set. A collection $\mathcal{T} \subseteq \mathcal{P}(X)$ of subsets of X is called a topology on X provided that the following three properties are satisfied:*

1. *$\emptyset \in \mathcal{T}$ and $X \in \mathcal{T}$.*
2. *$\mathcal{T}$ is closed under finite intersections. That is, given any finite collection $U_1, \ldots, U_n$ of sets in $\mathcal{T}$, their common intersection $U_1 \cap \cdots \cap U_n$ is also an element of $\mathcal{T}$.*
3. *$\mathcal{T}$ is closed under arbitrary unions. That is, if $\{ U_\alpha : \alpha \in I \}$ is a family of sets in $\mathcal{T}$ (here I is some indexing set, which may be infinite), then their union $\bigcup_{\alpha \in I} U_\alpha$ is also an element of $\mathcal{T}$.*

Given a set X and a topology $\mathcal{T}$ on X, the pair $(X, \mathcal{T})$ is called a topological space. We will often conflate a topological space $(X, \mathcal{T})$ with its underlying set X if the topology in question is clear from context.

The elements $U \in \mathcal{T}$ of a topology on X are called open subsets of X, or simply open sets.

This is our fundamental definition. By only dealing with open sets in topological spaces, we will be able to solidify all of the concepts mentioned earlier in the introduction. Before going any further, however, let's see some examples.

Example 3.2. Let $X = \mathbb{R}^n$, and let

$$\mathcal{T}_{\text{usual}} := \{ U \subseteq \mathbb{R}^n : \forall x \in U, \exists \varepsilon > 0 \text{ such that } B_\varepsilon(x) \subseteq U \}.$$

As we already noted above, $\mathcal{T}_{\text{usual}}$ forms a topology on $\mathbb{R}^n$. We call this the usual topology on $\mathbb{R}^n$, and refer to the topological space $(\mathbb{R}^n, \mathcal{T}_{\text{usual}})$ as $\mathbb{R}^n$ with the usual topology. We will often simply write $\mathbb{R}^n_{\text{usual}}$ for convenience.

Of particular note is the case where $n = 1$. In this case our definition above takes the following form:

$$\mathcal{T}_{\text{usual}} = \{ U \subseteq \mathbb{R} : \forall x \in U, \exists \delta > 0 \text{ such that } (x - \delta, x + \delta) \subseteq U \}.$$

The nonempty open subsets of $\mathbb{R}_{\text{usual}}$ are precisely the open intervals you have been working with since first year calculus—those intervals of the form (a, b), (a, ∞), $(-\infty, b)$, and $(-\infty, \infty) = \mathbb{R}$ for any real numbers $a < b$—along with arbitrary unions of them.

Exercise 3.3. Fix real numbers $a < b$. Explicitly show that the interval (a, b) is open in $\mathbb{R}_{\text{usual}}$. Show that the interval $[a, b)$ is not open in $\mathbb{R}_{\text{usual}}$.

Note that while the usual topology on $\mathbb{R}$ contains, for example, sets like $(0, 1) \cup (7, 8) \cup (100, 127)$, we can describe all of these sets by specifying the usual open intervals, then allowing all unions of them. The idea that an entire topology can be specified by some smaller collection of special open sets (open intervals in this case) along with arbitrary unions of them is an important one. We will return to it in the next section.

The following two examples are somehow trivial sorts of topologies. In fact, the second of them is often called exactly that! They are the simplest topologies we can define, and we can define both of them on *any* nonempty set.

Example 3.4. Let X be any set. Define $\mathcal{T}_{\text{discrete}} := \mathcal{P}(X)$. That is, $\mathcal{T}_{\text{discrete}}$ is the collection of *all* subsets of X. Then $\mathcal{T}_{\text{discrete}}$ is called the discrete topology on X.

Example 3.5. Let X be any nonempty set. Define $\mathcal{T}_{\text{indiscrete}} := \{\emptyset, X\}$. Then $\mathcal{T}_{\text{indiscrete}}$ is called the indiscrete topology on X, or sometimes the trivial topology on X.

These two are the "extremal" topologies on a given set, in the sense that they have the most open sets and the fewest open sets, respectively; it is not possible to have more open sets than the discrete topology, or to have fewer open sets than the indiscrete topology. There is quite a substantial difference in their usefulness in mathematics, however.

The discrete topology comes up relatively frequently. It is even a metric space (which for now you should just read as "very nice space"). That said, it still has some weird properties that might make you uneasy. For example, every function whose domain is a discrete topological space is continuous. For another example, the only sequences that converge in a discrete topological space are the (eventually) constant sequences. Even though we have not yet formally defined sequence convergence or continuity in a general topological context, these should strike you as weird.

The indiscrete topology, on the other hand, practically never comes up while doing mathematics. It is a barren, sad place. We only mention it when we have to, and we never go there if we can help it. To give you a feeling for how awful indiscrete spaces are, we will later discover that in an indiscrete topological space, every sequence converges to every point in the space. For example in $(\mathbb{R}, \mathcal{T}_{\text{indiscrete}})$, the constant sequence $0, 0, 0, 0, \ldots$ converges to π. It also converges to 7. What a mess!

Exercise 3.6. Fix an arbitrary nonempty set X, and let $\mathcal{T}_{\text{discrete}}$ and $\mathcal{T}_{\text{indiscrete}}$ be as above. Show that these are both topologies on X.

We will see many different topologies on many different sets throughout the course, so our current catalogue of topological spaces is not very exhaustive. Here are just a few more examples, and any proofs that may be required will be left to the Big List.

Example 3.7. Let $X = \{\triangle, \square, \diamondsuit, \heartsuit\}$. Define

$$\mathcal{T} := \{\emptyset, X, \{\triangle, \square, \diamondsuit\}, \{\diamondsuit, \heartsuit\}, \{\diamondsuit\}\}.$$

Then $(X, \mathcal{T})$ is a topological space.

Example 3.8. Working with $\mathbb{R}$ as the underlying set, define

$$\mathcal{T}_{\text{ray}} := \{\, (a, \infty) \,:\, a \in \mathbb{R} \,\} \cup \{\emptyset, \mathbb{R}\}.$$

Then $\mathcal{T}_{\text{ray}}$ is a topology on $\mathbb{R}$ that we will call the "ray topology", for obvious reasons.

Example 3.9. Let X be any nonempty set. Define

$$\mathcal{T}_{\text{co-finite}} := \{\, U \subseteq X \,:\, X \setminus U \text{ is finite} \,\} \cup \{\emptyset\}.$$

Then $\mathcal{T}_{\text{co-finite}}$ is called the co-finite topology on X.

Example 3.10. Let X be any nonempty set. Define

$$\mathcal{T}_{\text{co-countable}} := \{\, U \subseteq X \,:\, X \setminus U \text{ is countable} \,\} \cup \{\emptyset\}.$$

Then $\mathcal{T}_{\text{co-countable}}$ is called the co-countable topology on X.

Example 3.11. Let X be a nonempty set, and fix an element $p \in X$. Define

$$\mathcal{T}_p := \{\, U \subseteq X \,:\, p \in U \,\} \cup \{\emptyset\}.$$

Then $\mathcal{T}_p$ is called the particular point topology at p on X.

These last three examples are not quite as useless in "real" mathematics as the indiscrete topology, but they are close. However, they will serve as testing grounds for some of the properties of topological spaces we will soon define, and may serve to illustrate what a topology can do for us.

For example, we will learn that in the topological space $(\mathbb{R}, \mathcal{T}_0)$ (that is, the reals with the particular point topology at 0), the sequence $\{\frac{1}{n}\}_{n=1}^{\infty}$ does not converge to 0. In fact the only sequences that *do* converge to 0 in this topological space are eventually equal to the constant sequence $0, 0, 0, 0, \ldots$. Try to guess what point or points, if any, the sequence $1, 1, 1, 1, \ldots$ might converge to in this topology.

Exercise 3.12. Under what assumption(s) on the set X does $\mathcal{T}_{\text{co-finite}} = \mathcal{T}_{\text{co-countable}}$?

4 Comparing topologies

It is often useful to be able to compare two topologies on the same set. To this end, we have the following terminology.

Definition 4.1. *Let X be a set, and suppose $\mathcal{T}_1$ and $\mathcal{T}_2$ are topologies on X. We say that $\mathcal{T}_1$ refines $\mathcal{T}_2$, or that $\mathcal{T}_1$ is finer than $\mathcal{T}_2$, if $\mathcal{T}_1 \supseteq \mathcal{T}_2$. In other words, a topology with more open sets is finer than a topology with fewer open sets.*

We can (and will) also express this same idea by saying $\mathcal{T}_2$ is refined by $\mathcal{T}_1$, or by saying $\mathcal{T}_2$ is coarser than $\mathcal{T}_1$.

You may find this definition somewhat confusing at first, but it will become very natural as you strengthen your intuition about topological spaces and what they can do. For now, imagine that a finer topology on a set X can make finer (in the colloquial sense) distinctions between points in X than a coarser topology.

Example 4.2. Thus far we have discussed two topologies specific to the set of real numbers: the usual topology and the ray topology. We have also discussed the discrete and indiscrete topologies that any nonempty set can have. For the reals, we have:

$$\mathcal{T}_{\text{discrete}} \text{ refines } \mathcal{T}_{\text{usual}}, \text{ which refines } \mathcal{T}_{\text{ray}}, \text{ which refines } \mathcal{T}_{\text{indiscrete}}.$$

Exercise 4.3. Fix a nonempty set X. Show that $\mathcal{T}_{\text{co-finite}}$ is coarser than $\mathcal{T}_{\text{co-countable}}$. This exercise may be tricky at first.

Remark 4.4. Given two topologies $\mathcal{T}_1$ and $\mathcal{T}_2$ on a set X, it may be that neither one refines the other. In this case, we sometimes say $\mathcal{T}_1$ and $\mathcal{T}_2$ are incomparable topologies.

Example 4.5. $\mathcal{T}_{\text{usual}}$ and $\mathcal{T}_7$ (the particular point topology at 7) are incomparable topologies on $\mathbb{R}$. To see this, observe that the interval $(1,2)$ is open in the usual topology, but it does not contain 7 and so is not open in $\mathcal{T}_7$. This shows that $\mathcal{T}_{\text{usual}} \not\subseteq \mathcal{T}_7$, or in other words that the particular point topology does not refine the usual topology.

On the other hand, the set $\{\pi, 7\}$ is open in $\mathcal{T}_7$ since it contains 7, but is not open in $\mathcal{T}_{\text{usual}}$. This shows that $\mathcal{T}_7 \not\subseteq \mathcal{T}_{\text{usual}}$, or in other words that the usual topology does not refine the particular point topology.

(The number 7 is, of course, not really important in this example. $\mathcal{T}_{\text{usual}}$ and $\mathcal{T}_p$ are incomparable for any $p \in \mathbb{R}$.)

