Introduction to Real Analysis I

Course notes for Math 425

Brandon Hanson

University of Maine, Orono

Spring 2022

Table of Contents

1 MOTIVATION

1.1 What is analysis and what are the reals?

Analysis, broadly speaking, is the branch of math that favours estimates rather than identities. Identities are algebraic. Estimates involve the use of inequalities.

The real numbers are a bit harder to define. We're all pretty comfortable (hopefully) with the rational numbers

$$\mathbb{Q} = \left\{\frac{a}{b} : a \in \mathbb{Z}, b > 0\right\}$$

consisting of ratios of integers. You might have seen these with the condition $b \neq 0$ instead of $b > 0$, but writing $b > 0$ doesn't really change anything, since we can always replace the numerator with $-a$ if need be. The rational numbers have a lot going for them: we can perform arithmetic – addition, multiplication, subtraction and division – and we can compare them – we can tell if one rational number is bigger than another. The problem with the rational numbers is that they have some holes (actually, a lot of holes). What this means is that there are bits missing from the "real line".

The real line is a tempting way to think of the real numbers. It is an infinitely long ruler, and a real number would then be a measurement taken on this rule. But if we only knew of rational numbers, we wouldn't be able to measure the hypotenuse of the right triangle with sidelengths 1. We know from the Pythagorean theorem that the hypotenuse should be $\sqrt{1^2+1^2}=\sqrt{2}$.

Theorem 1.1: An irrational number

There is no rational number x such that $x^2=2$.

Proof. If there were, say $x=a/b$, then we could keep dividing our factors of 2 from a and b until one of them was no longer even. We'd have

$$2=x^2=a^2/b^2$$

leading to

$$2b^2=a^2$$

so that a^2 has to be even, and in turn telling us that a is also even. Thus $a=2c$ and $a^2=4c^2$ so that

$$2b^2=4c^2$$

from which we conclude

$$b^2=2c^2.$$

Now the table has turned and we conclude that b is also even, but we already made clear that a and b should not both be even, a contradiction. □

So if there is no rational number that can be used to measure our hypotenuse, what do we do? Maybe decimals are the way to go? If you punch $\sqrt{2}$ into your calculator you'll see

$$\sqrt{2}=1.41421356237309504880168872\ldots$$

but then what do those … really mean? The digits can't terminate, since only rational numbers can have a terminating decimal expansion, and we already saw that $\sqrt{2}$ isn't rational. So this expansion goes on forever, which could mean adding up tenths and hundredths and thousandths and so on, but how do you add infinitely many numbers? Or maybe it means that you can get very close to $\sqrt{2}$ if you include enough digits. But how close is very close? All of these ideas are good ones, and our goal is to make them precise.

Back to the algebra vs. analysis question, a distinction is the following. Algebraists like to do algebra, so they would perhaps define 0 to be the thing that satisfies the rule

$$a+0=0+a=a,\ \forall a.$$

This is defining the oh so important number 0 by its algebraic property: it is the additive identity. And then one might go on to prove the following.

Lemma 1.1

There can only be one 0.

Proof. Suppose 0_1 and 0_2 are two additive identities. Then

$$0_1 = 0_1 + 0_2$$

since 0_2 is an additive identity. Meanwhile

$$0_2 = 0_1 + 0_2$$

because 0_1 is an additive identity. Thus $0_1 = 0_2$. □

This proof is completely algebraic. An analyst, on the other hand, might interpret 0 as the only number with no length. Before getting to this, we'll need the following essential property: the **Archimedean property** states that for any number x, there is a natural number N which is larger than x.

Lemma 1.2: The analyst's 0

Let x be any number. Then

$$x = 0 \iff |x| < \frac{1}{n}, \forall n \in \mathbb{N}.$$

Proof. If $x = 0$ then $|x| = 0$ too, while $1/n$ is positive for each $n \in \mathbb{N}$. Conversely, if $x \neq 0$ then $|x| > 0$ so $1/|x|$ is a well-defined number. Thus, by the Archimedean property,

$$\frac{1}{|x|} < N$$

for some natural number N which rearranges to

$$|x| > \frac{1}{N}.$$

□

This proof is worth thinking about a bit, because it is, first of all, a fundamental perspective of analysis, and second of all, it is a preview of the sorts of arguments we'll rely on.

2 INEQUALITIES

2.1 Basic properties of inequalities

The basic inequalities we need to get underway start from the order on the natural numbers. We enumerate

$$\mathbb{N} = \{1, 2, 3, \ldots\}$$

and then say $m < n$ if m comes before n in this enumeration. The first important property of this is **transitivity**: if $m < n$ and $l < m$ then $l < n$. Now we extend this to integers first by declaring that $0 < n$ and $-n < 0$ for any natural number n. By transitivity $-n < m$ whenever m and n are natural numbers. To compare two negative numbers, we say

$$-n < -m \iff m < n$$

whenever $m, n \in \mathbb{N}$. So negative numbers have the “reverse" order of their positive counterparts.

Next we want to see how these inequalities play with arithmetic. First we have **translation invariance**

$$m \leq n \iff m + a \leq n + a$$

whenever $a \in \mathbb{Z}$ and **dilation invariance**

$$m \leq n \iff ma \leq na$$

whenever $a \in \mathbb{Z}$ and $a > 0$. For negative dilation, we have

$$m \le n \iff am > an$$

whenever $a \in \mathbb{Z}$ and $a < 0$.

Dilation invariance then tells us how we should compare rationals:

$$\frac{a}{b} < \frac{c}{d} \iff ac < bd,$$

and the right hand side is now just an inequality involving integers. As always, we're assuming that $b, d > 0$. Let's take stock.

Lemma 2.1

We have the following properties for manipulation of inequalities involving numbers.

1. Transitivity: if $x < y$ and $y < z$ then $x < z$.
2. Translation invariance: if $x < y$ then $x + z < y + z$.
3. Dilation invariance: if $z > 0$ and $x < y$ then $xz < yz$, while if $z < 0$ and $x < y$ then $xz > yz$.

These basic properties will serve as the building blocks for more sophisticated inequalities.

Theorem 2.1: Inequalities with sums

Let $a_1, \ldots, a_N$ and $b_1, \ldots, b_N$ be numbers with $a_j \le b_j$ for each j. Then

$$a_1 + \ldots + a_N \le b_1 + \ldots + b_N.$$

Proof. This is by induction on N, and when $N = 1$ there is nothing to prove. Now assume that we know the theorem holds for $N - 1$. This means we can assert that

$$a_1 + \ldots + a_{N-1} \le b_1 + \ldots + b_{N-1}.$$

We add a_N to both sides, using translation invariance, to get

$$a_1 + \ldots + a_{N-1} + a_N \le b_1 + \ldots + b_{N-1} + a_N.$$

Now, by the hypothesis of the theorem,

$$a_N \le b_N$$

so that adding $b_1 + \ldots + b_{N-1}$ to either side of this inequality (by translation invariance again) we get

$$(b_1 + \ldots + b_{N-1}) + a_N \leq (b_1 + \ldots + b_{N-1}) + b_N.$$

So by transitivity,

$$a_1 + \ldots + a_{N-1} + a_N \leq b_1 + \ldots + b_N.$$

□

Another important piece that factors into our inequalities is the absolute value.

Lemma 2.2

We have

$$|x| \leq y \iff -y \leq x \leq y.$$

Proof. If $x \geq 0$ then

$$-y \leq 0 \leq x = |x| \leq y$$

while if $x \leq 0$ then

$$-x = |x| \leq y$$

so that

$$-y \leq x \leq 0 \leq y.$$

Conversely, if $-y \leq x \leq y$ then either $x \geq 0$ so that $|x| = x \leq y$ or $x \leq 0$ so that $y \geq -x = |x|$. □

Example: Solve $|x-7| \leq 4$. This is equivalent to $-4 \leq x-7 \leq 4$ and adding 7 throughout gives

$$3 \leq x \leq 11.$$

2.2 The fundamental inequalities

Probably the most important inequality in the whole course is the triangle inequality.

Theorem 2.2: Triangle Inequality

For any x and y,

$$|x+y| \leq |x| + |y|.$$

Furthermore, equality holds exactly when x and y have the same sign.

Proof. Suppose first the x and y have the same sign (or either of them is zero). If need be, we can multiply them both by -1 to make them both non-negative, so let's assume that too. Then

$$|x+y| = x+y = |x|+|y|.$$

Next, suppose that x and y have opposite signs and neither is zero. Furthermore, assume that $|y| \leq |x|$ (they are just letters after all, so we can just call the one with the larger absolute value x). Again, if we have to, we can multiply by -1 to make x positive and y negative. So $x > 0 > y$ and $x = |x| \geq |y| = -y$. Because $x \geq -y$ we have

$$x + y \geq -y + y = 0$$

so $|x+y| = x+y$, while

$$|x|+|y| = x-y$$

and so we are left to check

$$x-y > x+y$$

which (adding $-y-x$ to both sides) is the same as

$$-2y > 0$$

and this is true since $-y > 0$. □

Example: We check that if $|x-y|$ and $|y-z|$ are at most 1 then $|x-z|$ is at most 2. The trick is to "add 0 creatively". We have

$$|x-z| = |(x-y)+(y-z)| \leq |x-y|+|y-z| \leq 1+1 = 2.$$

Theorem 2.3: Reverse Triangle Inequality

For any x and y,

$$|x-y| \geq ||x|-|y||.$$

Proof. By the triangle inequality,

$$|x| = |x-y+y| \leq |x-y|+|y|$$

so

$$|x|-|y| \leq |x-y|.$$

Similarly

$$|y| = |y-x+x| \leq |y-x|+|x|$$

so

$$|y|-|x| \leq |y-x| = |x-y|.$$

Since $||x|-|y|| = |x|-|y|$ or $|y|-|x|$, and both are small than $|x-y|$, we're done. □

Theorem 2.4: The generalized triangle inequality

For numbers $a_1, \dots, a_N$ we have

$$|a_1 + \dots + a_N| \le |a_1| + \dots + |a_N|.$$

Proof. When $N = 2$ this is just the regular triangle inequality. By induction, it is enough to demonstrate how to proceed from N to $N+1$. We have

$$|(a_1 + \dots + a_N) + a_{N+1}| \le |a_1 + \dots + a_N| + |a_{N+1}| \le |a_1| + \dots + |a_N| + |a_{N+1}|$$

where the first inequality was the plain old triangle inequality and the second inequality was an application of the induction hypothesis. □

Example: Let $a, b \ge 50$ and $|c| \le 25$. Then $|a+b+c| \ge 75$. Indeed , by the reverse triangle inequality,

$$|a+b+c| \ge ||a+b| - |c|| \ge |a+b| - |c| \ge 100 - 25 = 75.$$

Example: Suppose $0 \le a, b, c, d \le 1$ and further that $|a-b| \le 1/10$ and $|c-d| \le 1/10$. How big can $ac - bd$ be? To answer this we need to investigate

$$|ac - bd|.$$

We somehow need to invoke what we know about $a-b$ and $c-d$ being small. To that end we introduce the intermediate term bc:

$$\begin{aligned} |ac - bd| &= |ac - bc + bc - bd| \\ &\le |ac - bc| + |bc - bd| \\ &= |c(a-b)| + |b(c-d)| \\ &= |c||a-b| + |b||c-d| \\ &\le 1/10 + 1/10 = 1/5. \end{aligned}$$

Theorem 2.5: The Cauchy-Schwarz inequality

Let $a_1, \dots, a_N$ and $b_1, \dots, b_N$ be numbers. Then

$$\left(\sum_{n=1}^{N} a_n b_n\right)^2 \le \left(\sum_{n=1}^{N} a_n^2\right)\left(\sum_{n=1}^{N} b_n^2\right).$$

Equality holds if and only if there is a constant c such that $a_n = cb_n$ for each n.

Proof. Let's first explore the equality case. If $a_n = cb_n$ for each n then the left hand side and the right hand side are both

$$c^2\left(\sum_{n=1}^{n} b_N^2\right)^2.$$

Now, for the rest of the proof, consider the quadratic polynomial (in x) given by

$$f(x) = \sum_{n=1}^{N}(a_n - xb_n)^2.$$

This quadratic function is non-negative because it is a sum of squares of numbers. It can only be zero if each of those numbers is itself 0, which is to say that for some x we have $a_n = xb_n$ for each n, and we already know that the Cauchy-Schwarz inequality is an equality in this case. Otherwise, $f(x) > 0$ for each x and we can expand f to get

$$f(x) = \sum_{n=1}^{N} a_n^2 - 2x\sum_{n=1}^{N} a_n b_n + x^2\sum_{n=1}^{N} b_n^2.$$

Now this is a quadratic function in x which is always positive, and so has no roots. This means we cannot solve the quadratic equation, so it must be that the discriminant $B^2 - 4AC$ is negative. In our case

$$A = \sum_{n=1}^{N} b_n^2,\ B = -2\sum_{n=1}^{N} a_n b_n,\ C = \sum_{n=1}^{N} a_n^2$$

and so $B^2 - 4AC < 0$ tells us that

$$4\left(\sum_{n=1}^{N} a_n b_n\right) < 4\left(\sum_{n=1}^{N} b_n^2\right)\left(\sum_{n=1}^{N} a_n^2\right)$$

which gives us the strict Cauchy-Schwarz inequality. □

Example: Suppose $a_1, \ldots, a_N$ are numbers which add to 1. Then

$$\sum_{n=1}^{N} \frac{a_n^2}{n} \geq \frac{2}{N(N+1)}.$$

This is an application of the Cauchy-Schwarz inequality, but seeing it might take some practice. There are no b_n's right off the bat, so we get to choose them, and then adjust the a_n's accordingly. In other words, we won't just blindly apply Cauchy-Schwarz, but instead we'll do a bit of setup first. We start with

$$1 = \sum_{n=1}^{N} a_n$$

which is given. We need a $1/n$ factor, but when we apply Cauchy-Schwarz, we end up squaring the terms, so instead we introduce a $1/\sqrt{n}$ factor. However, we can't just put this factor in for free, we also need to balance the books, so we re-write this as

$$1 = \sum_{n=1}^{N} a_n = \sum_{n=1}^{N} \frac{a_n}{\sqrt{n}} \sqrt{n}.$$

Thus

$$1 = 1^2 = \left(\sum_{n=1}^{N} \frac{a_n}{\sqrt{n}} \sqrt{n} \right)^2 \leq \left(\sum_{n=1}^{N} \frac{a_n^2}{n} \right) \left(\sum_{n=1}^{N} n \right).$$

The second sum on the right is

$$\sum_{n=1}^{N} n = \frac{N(N+1)}{2}$$

and if we divide through we get the conclusion we wanted.

The friendship paradox states that in most populations, the average person is less popular than their average friend. In fact the only situation in which this is not the case is when every person in said population has exactly the same number of friends – no one is more or less popular than anyone else.

To establish the friendship paradox, we need a bit of notation. Let $P = \{p_1, \ldots, p_n\}$ is a population of n people. We write

$$p_i \sim p_j$$

is p_i and p_j are friends and we write $d(p_i)$ for the popularity of person i, that is, the number of friends of person i. Let E denote the total number of relationships, which is to say, the number of pairs $\{p_i, p_j\}$ with $p_i \sim p_j$. Then

$$\sum_{i=1}^{n} d(p_i) = 2E.$$

This is because if $p_i \sim p_j$ is one of the E relationships, then both $d(p_i)$ and $d(p_j)$ count it. So the average person has

$$\frac{1}{n} \sum_{i=1}^{n} d(p_i) = \frac{2E}{n}$$

friends. Now how many friends might we expect the average friend of p_i to have? Well, if p_i has an average number of friends, then their average friend has

$$\frac{n}{2E} \sum_{p_j \sim p_i} d(p_j)$$

friends. Thus the average p_i has

$$\frac{1}{n} \sum_{i=1}^{n} \frac{n}{2E} \sum_{p_j \sim p_i} d(p_j) = \frac{1}{2E} \sum_{i=1}^{n} \sum_{p_j \sim p_i} d(p_j)$$

friends. But if we switch the order of the sums, we get

$$\frac{1}{2E}\sum_{j=1}^{n} d(p_j) \sum_{p_i \sim p_j} 1 = \frac{1}{2E}\sum_{j=1}^{n} d(p_j)^2.$$

So the claim of the friendship paradox is that the average person, whose popularity is $2E/n$ is less popular than their average friend, whose popularity is

$$\frac{1}{2E}\sum_{j=1}^{n} d(p_j)^2.$$

By Cauchy-Schwarz

$$\left(\frac{2E}{n}\right)^2 = \frac{1}{n^2}\left(\sum_{i=1}^{n} 1 \cdot d(p_i)\right)^2 \leq \frac{1}{n^2}\left(\sum_{i=1}^{n} 1^2\right)\left(\sum_{i=1}^{n} d(p_i)^2\right) = \frac{1}{n}\left(\sum_{i=1}^{n} d(p_i)^2\right)$$

and this rearranges to

$$\frac{2E}{n} \leq \frac{1}{2E}\sum_{i=1}^{n} d(p_i)^2.$$

Equality can only hold if it holds in the Cauchy-Schwarz inequality. That happens when the two sequences, all 1's (so constant) and $d(p_i)$, are proportional. This can only happen if $d(p_i)$ is constant, which means everyone has an equal number of friends.

3 REAL NUMBERS

3.1 Why real numbers?

A good way to motivate what we need from the real numbers is to reflect on what that rational numbers – something we can already get our hands on – already have, and what they are missing.

The first important property of the rationals is that they have all the necessary ingredients needed to perform arithmetic. The rationals form what is called a **field**, meaning they satisfy the following axioms.

Name	Formula	
associativity:	$(a+b)+c = a+(b+c)$	$(ab)c = a(bc)$
commutativity:	$a+b = b+a$	$ab = ba$
distributivity:	$a(b+c) = ab+ac$	$(a+b)c = ac+bc$
identities:	$a+0 = a = 0+a$	$a \cdot 1 = a = 1 \cdot a$
inverses:	$a+(-a) = 0 = (-a)+a$	$aa^{-1} = 1 = a^{-1}a$ if $a \neq 0$

Table 1: The field axioms

So we would like the reals to also satisfy these axioms. The rationals are ordered, meaning we can compare any two rational numbers and decide which of the two is larger. This is useful for measuring things, so we'd like the reals to be ordered as well.

It turns out that $\sqrt{2}$ is not rational (we've proved this) and π is not rational either (this is harder to prove!). So certain equations like

$$x^2 - 2 = 0$$

and

$$\frac{\text{sidelength of a square of perimeter 1}}{\text{radius of a square of circumference 1}} = x$$

cannot be solved. Well neither can the equation $x^2 + 1 = 0$, and the reals won't help with this. But there is a difference! We can *approximate* both $\sqrt{2}$ and π be rational numbers, we cannot approximate a solution to $x^2 + 1$ by rational numbers. For instance, if $a_1, a_2, \ldots$ were a sequence of better and better approximations to $\sqrt{2}$, we'd like to imagine $\sqrt{2}$ as the limit of these approximations, but from the point of view of the rational numbers, no such limit exists. The goal of the reals will be to add these missing limits in, a process called *completion*.

It turns out that a good way to describe the feature which we would like the reals to have is using **upper bounds** and **lower bounds**.

Definition 3.1: Bounded set

A set A is called **bounded above** if there is a number u such that

$$u \geq a, \quad \text{for each } a \in A.$$

Any number u with this property is called an upper bound for A. We say A is **bounded below** if there is a number l with

$$l \leq a, \quad \text{for each } a \in A.$$

Any number l with this property is called a lower bound for A. The set A is said to be bounded if it is both bounded above and bounded below.

Example: The set

$$A = \{a/b \in \mathbb{Q} : a^2/b^2 \leq 2\}$$

is a set with upper bound 2. Indeed, if $a/b \in A$ then

$$2^2 > 2 \geq a^2/b^2$$

so $2 > a/b$. We could just as well check that 1.5 is also an upper bound.

In the previous example, the most efficient upper bound we could find, if we new about numbers outside of $\mathbb{Q}$, would be $\sqrt{2}$ – it is an upper bound and we'll see that there is no number smaller that is still an upper bound. This turns out to be an important property.

Definition 3.2: Supremum and infimum

Let A be a set of numbers. A number u is called the supremum of A if it is an upper bound for A and if for any $\varepsilon > 0$, $u - \varepsilon$ is not an upper bound for A. A number l is called the infimum of A if it is a lower bound for A and if for any $\varepsilon > 0$, $l + \varepsilon$ is not a lower bound for A.

3.2 Defining the reals and using completeness

We can take $\mathbb{R}$, the set of real numbers, to be a field containing $\mathbb{Q}$, which is ordered in a manner consistent with the ordering of $\mathbb{Q}$, and with the **completeness property**: if A is any set which is bounded above, A has a supremum. This raises the question of whether such a field exists, but for now we just take it on faith. We now explore what we gain by using completeness by calculating suprema of various sets.

Example: If A is a finite, non-empty set of real numbers, then $\sup A = \max A$. By finiteness, we can just enumerate A as $A = \{a_1, \ldots, a_N\}$ in order, so that $a_1 < \ldots < a_N$. Then $a_N = \max A$. First, a_N is upper bound since it is the largest element of A. If $\varepsilon > 0$ the $a_N - \varepsilon < a_N$ so $a_N - \varepsilon$ is no longer an upper bound.

Example: If A is a the empty set, then $\sup A = -\infty$. To see this, notice that for any number r, $r \geq a$ for every $a \in A$, vacuously. This is true for every number r, and the smallest possible r would be $-\infty$.

The next example takes a bit of preparation, because we need to establish the following lemma.

Lemma 3.1: Denseness of $\mathbb{Q}$ in $\mathbb{R}$

If x is an real number then for any $\varepsilon > 0$, there are rational numbers a/b and c/d with

$$a/b - \varepsilon \leq x \leq a/b, \text{ and }, c/d \leq x \leq c/d + \varepsilon.$$

Proof. Let $\varepsilon > 0$ and let $N > 1/\varepsilon$ be a natural number (using the Archimedean principle). We can subdivide the real line into intervals $I_j = (j/N, (j+1)/N]$ of length $1/N$, where $j \in \mathbb{Z}$. Let's check this really is a partition of $\mathbb{R}$. This means we need to show that any x belongs to exactly one such interval. Well Nx is also a real number, so let

$$A = \{j \in \mathbb{Z} : j \geq Nx - 1\}.$$

Then A is a set of integers, and each integer in A is bounded below by $Nx - 1$. This means there is a smallest integer $j \in A$. Since $j \in A$,

$$j \geq Nx - 1 \implies j + 1 \geq Nx$$

and since $j - 1 < j$, $j - 1$ is not in A and so

$$j - 1 < Nx - 1 \implies j < Nx$$

and these inequalities combine to

$$j < Nx \leq j + 1.$$

Dividing by N we get

$$\frac{j}{N} < x \leq \frac{j+1}{N}$$

which means $x \in I_j$. These intervals are disjoint, so they have to partition $\mathbb{R}$. Now since x belongs to such an interval,

$$j/N \leq x \leq \frac{j+1}{N}$$

and so

$$\frac{j+1}{N} - \varepsilon \leq \frac{j+1}{N} - \frac{1}{N} \leq x \leq \frac{j+1}{N}$$

and

$$\frac{j}{N} \leq x \leq \frac{j}{N} + \frac{1}{N} \leq \frac{j}{N} + \varepsilon$$

So $a/b = (j+1)/N$ and $c/d = j/N$ are the fractions we're looking for. □

Example: If A is a the set

$$A = \{a/b \in \mathbb{Q} : a^2/b^2 \leq 2\}$$

then $\sup A = \sqrt{2}$. First $\sqrt{2}$ is an upper bound: if $a/b \in A$ then either $a/b < 0 < \sqrt{2}$ or else $a/b > 0$. But if $a/b > \sqrt{2}$ then

$$a^2/b^2 = a/b \cdot a/b > a/b \cdot \sqrt{2} > \sqrt{2} \cdot \sqrt{2} = 2$$

so $a/b \notin A$. Next, for $\varepsilon > 0$, we can find a fraction a/b with

$$a/b \leq \sqrt{2} \leq a/b + \varepsilon.$$

Then

$$a^2/b^2 \leq 2$$

so $a/b \in A$ but

$$\sqrt{2} \leq a/b + \varepsilon \implies \sqrt{2} - \varepsilon \leq a/b$$

so $\sqrt{2} - \varepsilon$ is not an upper bound for A.

Example: Let A and B be sets of real numbers each bounded above. Then

$$\sup(A \cup B) = \max\{\sup A, \sup B\}.$$

Indeed, if $M_A = \sup A$ and $M_B = \sup B$ then

$$M = \max\{M_A, M_B\} \geq M_A \geq a$$

for each $a \in A$ which shows M is an upper bound for A, and

$$M \geq M_B \geq b$$

showing M is an upper bound for B. Without loss of generality, $M = M_A$ and if $\varepsilon > 0$ there is an $a \in A$ (and hence in $A \cup B$) such that $M - \varepsilon < a$ which shows that $M - \varepsilon$ is not an upper bound of $A \cup B$.

Example: A tempting claim to make is that $m = \min\{\sup A, \sup B\} = \sup(A \cap B)$. It is true that m is an upper bound: if, without loss of generality, $m = \sup B$ then for each $c \in A \cap B$ we know $c \in B$ so $c \leq m$. However, if $A = [0,1] \cup \{3\}$ and $B = [0,2]$ then $\sup A = 3$, $\sup B = 2$ but $A \cap B = [0,1]$ and its supremum is 1.

Theorem 3.1: The Nested Interval Theorem

For each $n \in \mathbb{N}$ suppose $I_n = [a_n, b_n]$ is an interval and suppose further that these intervals are nested in the sense that $I_{n+1} \subseteq I_n$. Then $\bigcap_{n=1}^{\infty} I_n$ is not empty.

Proof. The nesting condition says that

$$a_n \leq a_{n+1} \leq b_{n+1} \leq b_n$$

for each n. We claim $a_n \leq b_m$ for each n and m. Indeed, if $n \leq m$

$$a_n \leq a_{n-1} \leq \cdots \leq a_m \leq b_m$$

by nesting, while if $n > m$ then

$$a_n \leq b_n \leq \cdots \leq b_m$$

by nesting. This shows that if $A = \{a_n : n \in \mathbb{N}\}$ then for each m, b_m is an upper bound for A. This means A is bounded above, so $\sup A$ exists and moreover $b_m \geq \sup A$. We claim $\sup A \in I_n$ for each n and this will prove the theorem. To see why, notice that for each n, $b_n \geq \sup A \geq a_n$ (the first inequality comes from what we just worked on, the second because $\sup A$ is an upper bound for A). □

3.3 The real numbers as cuts

We can think of a real number x as a point on the real line. It divides the line into a left and right half

$$L_x = (-\infty, x),\ R_x = [x, \infty).$$

It's just a convention that we have included x in the right half. Anyways, for each x we get two sets

$$L_x = \{y \in \mathbb{R} : y < x\},\ R_x = \{y \in \mathbb{R} : y \geq x\}$$

and converse from a division of $\mathbb{R}$ into two halves like this, we could recover x:

$$x = \sup L_x.$$

We could just as well have used x to partition the rationals, instead of the reals:

$$L_x = \{y \in \mathbb{Q} : y < x\},\ R_x = \{y \in \mathbb{Q} : y \geq x\}.$$

We can still recover x from this partition since we still have $x = \sup L_x$, but the upshot is we have defined this partition without having to know what $\mathbb{R}$ is.

A partition of $\mathbb{Q}$ into left and right halves is called a Dedekind cut. Formally, a pair (L, R) is a cut if

1. $\mathbb{Q} = L \cup R$,
2. neither L nor R is empty, and the sets are disjoint,
3. L has no greatest element, and
4. for $l \in L$ and $r \in R$, we have $l < r$, or equivalently, if $l \in L$ and $a \in \mathbb{Q}$ with $a \leq l$, then $a \in L$.

In the last point, we have given two conditions, and claimed them to be equivalent. They really are: if $a \leq l$ and a is a rational number, then $a \in L$ or R in view of (1). But every element in R has to be larger than every element in L, and $a \leq l$, so a has to be in L. Conversely, if l and r are two numbers from L and R respectively, were it that $r \leq l$, we would necessarily have $r \in L$. But the fact that the sets are disjoint means that r cannot be both in R and L.

For any cut (L, R) we might ask where the cut occurs? That is, for which x do we have $L = L_x$ and $R = R_x$. The fantastic idea here is that we can define L and R without knowing x, but x is uniquely determined by L and R. So we can think of the cut (L, R) as defining the number x, rather than the other way around, and this is Dedekind's construction of the real numbers.

Definition 3.3: Real numbers as Dedekind cuts

The real numbers consists of the set

$$\mathbb{R} = \{(L,R) : (L,R) \text{ is a cut of } \mathbb{Q}\}.$$

This will feel a little awkward as we have defined a set of numbers by interpreting a number as a pair of sets, rather than something with digits. However, they behave in just the same ways. For instance, $\mathbb{R}$ contains $\mathbb{Q}$ in a natural way: if a is a rational number, then we can represent a as the cut

$$((-\infty, a), [a, \infty))$$

with the intervals consisting of rational numbers. We can add two cuts together. **Example**: Suppose (L_1, R_1) and (L_2, R_2) are two cuts of $\mathbb{Q}$. Then so is

$$(L_1 + L_2, R_1 + R_2)$$

where for two sets of rational numbers A and B we define

$$A + B = \{a + b : a \in A,\ b \in B\}.$$

We can also order cuts, and talk about suprema. We say $(L_1, R_1) \le (L_2, R_2)$ if $L_1 \subseteq L_2$. This makes perfect sense: if we had interval $(-\infty, x_1)$ and $(-\infty, x_2)$ for L_1 and L_2, then saying $L_1 \subseteq L_2$ is exactly saying that $x_1 \le x_2$. Now if A is a set of cuts, then we can define a new cut

$$\sup_A = \left(\bigcup_{(L,R)\in A} L, \bigcap_{(L,R)\in A} R \right).$$

This is also a cut. It's greater than every element of A since if $(L_0, R_0) \in A$ then

$$L_0 \subseteq \bigcup_{(L,R)\in A} L.$$

But if (L', R') is another upper bound for A then

$$L' \supseteq L$$

for each $(L,R) \in A$ so

$$L' \supseteq \bigcup_{(L,R)\in A} L$$

and this shows $L' \ge \sup A$.

There are more properties to check in order to convince yourself that these cuts really to have all the qualities we want the real numbers to have. We'll leave these qualities for the reader to investigate, and instead provide another definition of the reals in terms of sequences later on.

4 SEQUENCES

4.1 Sequences and their basic limit properties

A sequence, formally, is a function $a : \mathbb{N} \to \mathbb{R}$. But rather than write a sequence as an input and an output, we write it sequentially:

$$a_1, a_2, a_3, \ldots.$$

We also write $\{a_n\}_{n=1}^{\infty}$ or else just $\{a_n\}$ for a sequence. While we're at it, subsequence of a sequence is just a sequence obtained by omitting some of the terms from the original sequence:

$$a_{n_1}, a_{n_2}, a_{n_3}, \ldots$$

where now $n_1 < n_2 < n_3 < \ldots$ are just some indices, and because we may have skipped some terms, n_k (the index of k'th term of the subsequence) is at least k (the index of the k'th term from the original sequence). So this subsequence is the new infinite list of numbers $\{a_{n_k}\}_{k=1}^{\infty}$.

Example: The sequence $0, 1, 0, 1, \ldots$ of alternating 0's and 1's is an infinite sequence. Some subsequences are $\{0\}_{k=1}^{\infty}$, the sequence of all 0', or the sequence $0, 0, 1, 0, 0, 1, 0, 0, 1, 0, 0, 1, \ldots$, obtained be omitting every second 1.

Example: The sequence $1, 4, 9, 16, 25, \ldots$ of perfect squares is an infinite sequence. In this example $a_n = n^2$. The perfect fourth powers is a subsequence: $1, 16, 81, 256, \ldots$.

Here the terms of the sequence are $a_{k^2} = (k^2)^2$, so $n_k = k^2$.

Definition 4.1: Limit of a sequence

A sequence $\{a_n\}$ converges to a real number L if for any $\varepsilon > 0$, there is a threshold $N = N(\varepsilon)$ so that for any $n \geq N$,

$$|a_n - L| < \varepsilon.$$

We say $\{a_m\}$ diverges to ∞ if for any $M \in \mathbb{N}$ there is a threshold $N = N(M)$ so that for any $n \geq N$,

$$a_n > M.$$

Example: The sequence of perfect squares diverges to ∞. Indeed, if M is any number, we take $M = \sqrt{M} + 1$ and if $n \geq N$ then

$$a_n = n^2 \geq N^2 > (\sqrt{M})^2 = M.$$

Example: The decimal approximations $a_n = \lfloor 10^n \pi \rfloor / 10^n$ of π converge to π. To see why, we notice first that $\lfloor 10^n \pi \rfloor$ means to round down. Thus

$$10^n \pi \geq \lfloor 10^n \pi \rfloor \geq 10^n \pi - 1$$

and so dividing through by 10^n,

$$\pi \geq a_n \geq \pi - 10^{-n}.$$

From here we see that if $N = \log_{10}(1/\varepsilon) + 1$ then $n \geq N$ tells us that

$$10^{-n} \leq 10^{-N} < \varepsilon$$

and so

$$\pi \geq a_n > \pi - \varepsilon$$

which in particular means that

$$|a_n - \pi| < \varepsilon.$$

Definition 4.2: Cauchy sequence

A sequence $\{a_m\}$ is called a Cauchy sequence if for any $\varepsilon > 0$ there is a threshold $N = N(\varepsilon)$ so that if $n, m \geq N$ then

$$|a_n - a_m| < \varepsilon.$$

The Cauchy property turns out to be a nice way for testing convergence.

Lemma 4.1: Convergent sequences are Cauchy

If a sequence $\{a_n\}$ converges to a number L then it is a Cauchy sequence.

Proof. Let $\varepsilon > 0$. Since $\{a_n\}$ converges, there is an N so that if $n \geq N$ we know $|a_n - L| < \varepsilon/2$. This same threshold tells us that if $n, m \geq N$ then

$$|a_n - a_m| \leq |a_n - L| + |a_m - L| < \varepsilon/2 + \varepsilon/2 = \varepsilon$$

which verifies the Cauchy property. □

Example: The sequence $0, 1, 0, 1, 0, 1\ldots$ does not converge. Indeed, if it did, it would have to be Cauchy, in particular with $\varepsilon = 1/2$. But for any threshold N, two consecutive terms are always separated by 1, which is bigger than ε.

Lemma 4.2

If $\{a_n\}$ converges to L, then so does every subsequence of $\{a_n\}$.

Proof. If $\{a_{n_k}\}_{k=1}^{\infty}$ is a subsequence then in particular we know $n_k \geq k$. Let $\varepsilon > 0$ and suppose N is the threshold so that $|a_n - L| < \varepsilon$ for $n \geq N$. Then if $k \geq N$, $n_k \geq k \geq N$ and so

$$|a_{n_k} - L| < \varepsilon$$

as well. □

Lemma 4.3

If $a_n \to A$ and $b_n \to B$ for some numbers A and B then $a_n + b_n \to A + B$.

Proof. Let $\varepsilon > 0$. From the convergence of a_n and b_n we know there are thresholds N_1 and N_2 so that if $n \geq N_1$ we have $|a_n - A| < \varepsilon/2$ while if $n \geq N_2$ we have $|b_n - B| < \varepsilon/2$. Let $N = \max\{N_1, N_2\}$. Then $n \geq N$ tells us n is beyond both thresholds N_1 and N_2 so

$$|(a_n + b_n) - (A + B)| = |(a_n - A) + (b_n - B)| \leq |a_n - A| + |b_n - B| < \varepsilon/2 + \varepsilon/2 = \varepsilon.$$

□

Lemma 4.4

If $a_n \to A$ and $b_n \to B$ for some numbers A and B then $a_n b_n \to AB$.

Proof. Let $\varepsilon > 0$. From the convergence of a_n and b_n we know there are thresholds N_1 and N_2 so that if $n \geq N_1$ we have

$$|a_n - A| < \frac{\varepsilon}{3|B|}$$

while if $n \geq N_2$ we have

$$|b_n - B| < \min\left\{\frac{\varepsilon}{3|A|}, |B|\right\}.$$

Let $N = \max\{N_1, N_2\}$. Then $n \geq N$ tells us n is beyond both thresholds N_1 and N_2 so

$$|a_n b_n - AB| = |a_n b_n - Ab_n + Ab_n - AB| \leq |a_n b_n - Ab_n| + |Ab_n - AB| = |A - a_n||b_n| + |A||b_n - B|.$$

We are allowed to estimate $|A - a_n|$ and $|b_n - B|$ because n is beyond the threshold, so we get

$$|A - a_n||b_n| + |A||b_n - B| < \frac{\varepsilon}{3|B|}|b_n| + |A|\frac{\varepsilon}{3|A|} = \frac{\varepsilon}{3}\frac{|b_n|}{|B|} + \frac{\varepsilon}{3}.$$

But

$$|b_n| \leq |b_n - B| + |B| < 2|B|.$$

So we get

$$\frac{\varepsilon}{3}\frac{|b_n|}{|B|} + \frac{\varepsilon}{3} \leq \frac{\varepsilon}{3}\frac{2|B|}{|B|} + \frac{\varepsilon}{3} = \varepsilon.$$

□

Just how we chose the right thresholds for N_1 and N_2 in the last proof takes practice. You can figure them out by reversing the process: apply all the inequalities you can first, and then figure out how close you need a_n to be to A and b_n to be to B to make it all work out.

Lemma 4.5

If $b_n \to B$ and $B \neq 0$ then b_n is eventually non-zero and $\lim_{n\to\infty} 1/b_n = 1/B$.

Proof. Let $\varepsilon = |B|/2$ in the definition of convergence for b_n. Then for n sufficiently large, we have

$$|B| = |B - b_n + b_n| \leq |B - b_n| + |b_n| \leq |B|/2 + |b_n|$$

which rearranges to $|b_n| \geq |B|/2$ and in particular b_n cannot be zero. Next, if $\varepsilon > 0$ is arbitrary, we would like to estimate

$$\left|\frac{1}{b_n} - \frac{1}{B}\right| = \frac{|b_n - B|}{|b_n||B|}.$$

Since $|b_n| \geq |B|/2$ for n sufficiently large, we know the right hand side above is at most

$$\frac{|b_n - B|}{|B|/2 \cdot |B|} = \frac{2|b_n - B|}{|B|^2}.$$

But if n is sufficiently large, we can be sure (by convergence of b_n) that

$$|b_n - B| < \frac{\varepsilon |B|^2}{2}$$

and so

$$\left| \frac{1}{b_n} - \frac{1}{B} \right| < \varepsilon.$$

□

Lemma 4.6

If $a_n \to A$ and $b_n \to B \neq 0$ then a_n/b_n is eventually defined and converges to A/B.

Proof. We know $1/b_n$ is eventually defined and converges to B. So

$$\frac{a_n}{b_n} = \frac{1}{b_n} \cdot a_n \to \frac{1}{B} \cdot A$$

by the lemma about products of sequences. □

Lemma 4.7

Suppose $a_n \to A$ and $b_n \to B$ and further that $a_n \leq b_n$ for each n. Then $A \leq B$.

Proof. If not, $A - B > 0$ and set $\varepsilon = (A - B)/2$ in the definition of convergence for a_n and b_n. We get two thresholds, one (say N_1) such that

$$a_n > A - \varepsilon = \frac{A+B}{2}$$

if $n \geq N_1$, and another (say N_2) such that

$$b_n < B + \varepsilon = \frac{A+B}{2}$$

if $n \geq N_2$. These combine to say that if $n \geq \max\{N_1, N_2\}$ then

$$a_n > \frac{A+B}{2} > b_n$$

which is a contradiction. □

4.2 Monotone sequences

Definition 4.3: Monotone sequence

A sequence $\{a_n\}$ is called increasing (resp. strictly increasing) if $a_m \leq a_n$ when $m < n$ (resp. $a_m < a_n$ when $m < n$). A sequence $\{a_n\}$ is called decreasing (resp. strictly decreasing) if $a_m \geq a_n$ when $m < n$ (resp. $a_m > a_n$ when $m < n$). A sequence which is any of the above is called monotone.

Definition 4.4: Bounded sequence

A sequence $\{a_n\}$ is called bounded if there is a number M such that for each n,

$$-M \leq a_n \leq M.$$

Theorem 4.1: Monotone convergence theorem

Let $\{a_n\}$ be a bounded sequence. If $\{a_n\}$ is increasing then

$$\lim_{n\to\infty} a_n = \sup\{a_n : n \in \mathbb{N}\}.$$

If $\{a_n\}$ is decreasing then

$$\lim_{n\to\infty} a_n = \inf\{a_n : n \in \mathbb{N}\}.$$

Proof. We just prove the increasing case. We can apply it to the sequence $\{-a_n\}$ to handle the decreasing case.

Let $L = \sup\{a_n : n \in \mathbb{N}\}$. By definition $a_n \leq L$ for each n. Furthermore, if $\varepsilon > 0$, there is some a_N with

$$L - \varepsilon < a_N.$$

Since the sequence is increasing, if $n \geq N$ then

$$a_n \geq a_N > L - \varepsilon$$

so that in particular

$$|a_n - L| < \varepsilon.$$

This means that $a_n \to L$. □

Theorem 4.2: Monotone convergence theorem (divergence case)

Let $\{a_n\}$ be an unbounded sequence. If $\{a_n\}$ is increasing then

$$\lim_{n\to\infty} a_n = \infty.$$

If $\{a_n\}$ is decreasing then

$$\lim_{n\to\infty} a_n = -\infty.$$

Proof. We just prove the increasing case. We can apply it to the sequence $\{-a_n\}$ to handle the decreasing case.

Let M be any positive integer larger than $|a_1|$. Then all terms of the sequence satisfy

$$a_n \geq a_1 \geq -M.$$

But the sequence is not bounded, so it cannot be that $a_n \leq M$ for all n. Thus there is some N with $a_N > M$. But if $n \geq N$,

$$a_n \geq a_N > M$$

which precisely shows that $a_n \to \infty$. □

Theorem 4.3: The Monotone Subsequence Theorem

Let $\{a_n\}$ be any sequence. Then $\{a_n\}$ has a monotone subsequence.

We will prove the Monotone Subsequence Theorem by first proving Ramsey's Theorem.

Theorem 4.4: Ramsey's Theorem

Suppose we have a finite set of r colours, which we just label $\{1,\dots,r\}$. Suppose further that for each pair $\{i,j\}$ of two natural numbers, we colour that pair with one of the colours $\{1,\dots,r\}$. Then there is an infinite set $I \subseteq \mathbb{N}$ and a colour $c \in \{1,\dots,r\}$ such that if $i,j \in I$ then $\{i,j\}$ is coloured c.

Let's see how Ramsey's Theorem proves the Monotone Subsequence Theorem. Let $\{a_n\}$ be our given sequence. If $i,j \in \mathbb{N}$ with $i<j$ we colour $\{i,j\}$ with one of u or d (for up or down): if $a_i \leq a_j$ we colour $\{i,j\}$ with u and if $a_i > a_j$, we colour $\{i,j\}$ with d. This lets us colour all of the pairs of natural numbers. By Ramsey's Theorem, there is an infinite set

$$I = \{n_1 < n_2 < \dots\}$$

of natural numbers so that $\{n_i, n_j\}$ is always coloured the same, say u. This means that if $i < j$ then $a_{n_i} \le a_{n_j}$ and this means that the infinite sequence $\{a_{n_k}\}$ is increasing.

Now we present the proof of Ramsey's Theorem.

Proof. We are given that each pair of numbers $\{i, j\}$ is labeled with one of the colours $\{1, \ldots, r\}$.

- **Step 1**: Start with the number $n_1 = 1$ and set $A_1 = \{2, 3, 4, \ldots\}$. We divide A_1 into r sets
$$A_{1,c} = \{m \in A_1 : \{n_1, m\} \text{ is coloured } c\}.$$
One of the sets $A_{1,j}$ has to be infinite, since their union is A_1 which is infinite. Lets say A_{1,c_1} is infinite. Then we set $A_2 = A_{1,c_1}$ and move on to step 2.

- **Step 2**: Pick the smallest number $n_2 \in A_2$. We divide A_2 into r sets
$$A_{2,c} = \{m \in A_2 : \{n_2, m\} \text{ is coloured } c\}.$$
One of the sets $A_{2,c}$ has to be infinite, since their union is A_2 which is infinite. Lets say A_{2,c_2} is infinite. Then we set $A_3 = A_{2,c_2}$ and move on to step 3...

  Eventually we reach

- **Step k**: Pick the smallest number $n_k \in A_k$. We divide A_k into r sets
$$A_{k,c} = \{m \in A_k : \{n_k, m\} \text{ is coloured } c\}.$$
One of the sets $A_{k,j}$ has to be infinite, since their union is A_k which is infinite. Lets say A_{k,c_k} is infinite. Then we set $A_{k+1} = A_{k,c_k}$ and move on to step $k+1$...

We started step k with an infinite set A_k. We chose a colour c_k from $\{1, \ldots, r\}$ and a number n_k from A_k such that every number in the infinite set A_{k+1} is connected to n_k by the colour c_k. Also notice that A_{k+1} was always a subset of A_k and so in fact we have the inclusions
$$A_1 \supseteq A_2 \supseteq A_3 \supseteq \ldots$$

Now, we have a sequence $c_1, c_2, c_3, \ldots$ of colours from $\{1, \ldots, r\}$ and so there is an infinite set $K \subseteq \mathbb{N}$ and a colour c such that $c_k = c$ for each $k \in K$. In other words, K is the set of indices of steps where we chose colour c. If k_1 and k_2 are two numbers from K, then $A_{k_1} = A_{k_1,c}$ and $A_{k_2} = A_{k_2,c}$ because at steps k_1 and k_2 we chose the colour c. If $k_1 < k_2$ then $A_{k_2} \subseteq A_{k_1}$ which means n_{k_2} was connected to n_{k_1} with the colour c. In other words, the set $I = \{n_k : k \in K\}$ satisfies the conclusion of the theorem. $\square$

4.3 Other convergence results

From the Montone Convergence Theorem and the Monotone Subsequence Theorem we learn the following.

Theorem 4.5: Bolzano-Weierstrass

If $\{a_n\}$ is a bounded sequence then it has a convergent subsequence.

Proof. We know that $\{a_n\}$ has a monotone subsequence $\{a_{n_k}\}$ and all of the terms in the sequence are bounded. Thus they converge by Monotone Convergence Theorem. □

We all learn that Cauchy sequences converge. This is an incredibly useful fact since the Cauchy condition is often easier to check than actually finding a limit.

Lemma 4.8

Suppose $\{a_n\}$ is a Cauchy sequence. Then it is bounded.

Proof. Let $\varepsilon = 1$ in the definition of Cauchy sequence. Then there is some number N such that $|a_m - a_n| \le 1$ if $n, m \ge N$ and in particular, for all $n \ge N$, $a_n \in [a_N - 1, a_N + 1]$. Thus all the terms of the sequence are either one of $a_1, \ldots, a_{N-1}$ or else lie in a bounded interval, and hence are themselves bounded numbers. □

Lemma 4.9

Suppose $\{a_n\}$ is a Cauchy sequence and it has a subsequence which converges to a limit L. Then the whole sequence converges to L.

Proof. This is a homework exercise, but the gist is that some number a_{n_k} will be within $\varepsilon/2$ of L (any term far enough down the hypothetical subsequence) and so for any sufficiently large n,

$$|a_n - L| \le |a_n - a_{n_k}| + |a_{n_k} - L| \le \varepsilon/2 + \varepsilon/2,$$

using the Cauchy condition. □

Theorem 4.6: Cauchy's Criterion

Any Cauchy sequence converges.

Proof. The sequence is bounded. By Bolzano-Weierstrass, it has a subsequence converging to some limit L. The whole sequence must then converge to L. □

So far we have proved a number of theorems to get to a place where we know Cauchy sequences converge. These theorems can all be traced back to the assumption that a bounded set has a supremum, so we have thus shown “$\sup A$ exists if A is bounded" implies that “Cauchy sequences converge". The converse is also true. To prove this, a very handy lemma will come into play.

Lemma 4.10: The Squeeze Theorem

Suppose $\{a_n\}$, $\{m_n\}$, $\{b_n\}$ are three sequences with the properties

- for some number L, $\lim_{n\to\infty} a_n = \lim_{n\to\infty} b_n = L$, and
- for each n, $a_n \le m_n \le b_n$.

Then $\lim_{n\to\infty} m_n = L$.

Proof. Let $\varepsilon > 0$ and suppose N is so large that if $n \ge N$ we can be sure that $|a_n - L|, |b_n - L| \le \varepsilon$. Then

$$L - \varepsilon \le a_n \le m_n \le b_n \le L + \varepsilon$$

so $|m_n - L| \le \varepsilon$ too. □

Theorem 4.7: Completeness from Cauchy's Criterion

Suppose we know that any Cauchy sequence of numbers converges. Then any bounded set has a supremum.

Proof. Let A be a non-empty bounded set. We will construct a pair of sequences $\{a_n\}$ and $\{b_n\}$ with the following properties:

- $\{a_n\}$ is an increasing sequence of numbers from A,
- $\{b_n\}$ is a decreasing sequence of upper bounds for A, and
- for each n, $0 \le b_n - a_n \le \frac{1}{2^{n-1}}(b_1 - a_1)$.

Initially we choose a_1 to be any element from A, and let b_1 be any upper bound for A (which exists, since A is bounded).

For the second entries of the sequences, we examine the midpoint is $m_1 = (a_1 + b_1)/2$. There are two cases. If m_1 is an upper bound for A, then we set $b_2 = m_1$ and we set $a_2 = a_1$. In this way $b_2 \le b_1$ and $a_2 \ge a_1$; b_2 is still an upper bound for A, and a_2 is still an element of A; and since b_2 is the midpoint, it is half as far from a_1 as b_1 is:

$$b_2 - a_2 = m_1 - a_1 = (b_1 - a_1)/2.$$

If m_1 is not an upper bound for A then there is some number $a \in A$ which is larger than m_1. We set $a_2 = a$ and $b_2 = b_1$. In this way $b_2 \leq b_1$ and $a_2 \geq a_1$; b_2 is still an upper bound for A, and a_2 is still an element of A; and since a_2 is larger than the midpoint, it's less than half as far from a_1 as b_1 is:

$$b_2 - a_2 \leq b_1 - m_1 = (b_1 - a_1)/2.$$

If m_1 is not an upper bound for A then there is some number $a \in A$ which is larger than m_1. We set $a_2 = a$ and $b_2 = b_1$. For subsequent entries in the sequences we proceed in the same fashion. Having constructed b_n and a_n, we look at the midpoint m_n. If it is an upper bound we set $b_{n+1} = m_n$ and $a_{n+1} = a_n$. If not, we find an element from $a_{n+1} \in A$ which is larger than m_n and set $b_{n+1} = b_n$. In any case,

$$b_{n+1} - a_{n+1} \leq \frac{1}{2}(b_n - a_n) \leq \frac{1}{2}\frac{1}{2^n}(b_1 - a_1) = \frac{1}{2^{n+1}}(b_1 - a_1)$$

so that (3) is always true.

Now, the sequence $\{a_n\}$ is Cauchy: if $n > m > N$ then

$$a_N \leq a_m \leq a_n \leq b_N$$

since $\{a_n\}$ is increasing and b_N is always an upper bound for A, so that

$$|a_n - a_m| \leq b_N - a_N \leq \frac{b_1 - a_1}{2^{N-1}}$$

which can be made as small as we like by choosing N large. In the same way, $\{b_n\}$ is Cauchy: if $n > m > N$

$$a_N \leq b_n \leq b_m \leq b_N.$$

This is enough to guaranteed that $a_n \to L_1$ and $b_n \to L_2$ for some limits L_1 and L_2. By the Squeeze Theorem, though, we have that

$$0 \leq b_n - a_n \leq \frac{b_1 - a_1}{2^{n-1}}$$

and both $\{0\}$ (as a constant sequence) and $\{\frac{b_1 - a_1}{2^{n-1}}\}$ converge to 0. So

$$L_2 - L_1 = \lim_{n\to\infty} b_n - a_n = 0$$

and thus $L_1 = L_2$. By the order rule for sequences, if $a \in A$ then $a \leq b_n$ for each n and so

$$L = \lim_{n\to\infty} b_n \geq a$$

which shows L is an upper bound for A. But $a_n \to L$ shows that in fact L is the supremum of A. □

4.4 The Erdős-Szekeres Theorem

The Monotone Subsequence Theorem is a purely combinatorial theorem that, when combined with the Monotone Convergence Theorem, unlocked a number of analytic results. The Erdős-Szekeres Theorem provides a quantitatively strong version of the Monotone Subsequence Theorem which lets us pass to a reasonably large subsequence.

Theorem 4.8: Erdős-Szekeres Theorem

Let $a_1, \ldots, a_{N^2}$ be a sequence of N^2 real numbers. It contains a monotone subsequence of length at least N.

Proof. Let M denote the length of the longest monotone subsequence of the sequence in question. We define a function

$$\phi : \{1, \ldots, N^2\} \to \{1, \ldots, M\} \times \{1, \ldots, M\}$$

defined as follows. We let I_j denote the length of the longest increasing subsequence which begins at a_j and D_j the length of the longest decreasing subsequence which begins at a_j. Then, since a_j is always a member of these subsequences, $1 \leq I_j, D_j \leq M$. The map ϕ is defined by the rule $\phi(j) = (I_j, D_j)$. If $M < N$ then the map ϕ cannot be injective, since the domain has N^2 elements but the co-domain has only M^2. This means that for some $j < k$ we have $\phi(j) = \phi(k)$. By definition, there is an increasing sequence of length I_k, say (a_{m_l}) starting at a_k and a decreasing sequence of length D_k, say (a_{n_l}) starting at a_k. But if $a_j \leq a_k$ then we can create an increasing sequence of length $I_k + 1$ starting at a_j by appending a_j to the beginning of the sequence (a_{m_l}), which would tell us that $I_j > I_k$; if $a_j > a_k$ then we can append a_j to the beginning of (a_{n_k}) and create a decreasing sequence of length $D_k + 1$ starting at a_j, so that $D_j > D_k$. In either case, we cannot have $\phi(j) = \phi(k)$, so $M \geq N$. □

5 SERIES

5.1 Convergent series

And infinite series is the result of trying to add infinitely many numbers together. Say we want to add all of the terms of the sequence $\{a_n\}$. Well, adding N terms results in the **partial sum**

$$S_N = a_1 + \ldots, a_N.$$

If we want to add all the terms of this sequence, we would let $N \to \infty$, which leads to the following.

Definition 5.1: Convergent series

The infinite series

$$S = \sum_{n=1}^{\infty} a_n$$

is said to converge (to S) if the sequence S_N of partial sums converges to S.

Right away, this rules out the convergence for many infinite series.

Lemma 5.1: Divergence Criterion

If

$$\sum_{n=1}^{\infty} a_n$$

converges then $a_n \to 0$.

Proof. The sequences $\{S_N\}_{N=1}^{\infty}$ and $\{S_{N+1}\}_{N=1}^{\infty}$ both converge to S (the second sequence is the same the first, just shifted by one). Thus

$$a_{N+1} = S_{N+1} - S_N \to S - S = 0.$$

□

The Divergence Criterion is enough to guarantee that a sequence like $1+1+1+1+\dots$ diverges, but not enough to guarantee *convergence*, as in the following example.

Theorem 5.1: Divergence of the harmonic series

The series

$$\sum_{n=1}^{\infty} \frac{1}{n}$$

diverges.

Actually this fact follows from a slightly more general fact.

Lemma 5.2: Cauchy Condensation Test

If $\{a_n\}$ is a decreasing sequence of non-negative numbers and if

$$\sum_{n=1}^{\infty} a_n$$

converges then

$$\sum_{j=0}^{\infty} 2^j a_{2^j}$$

also converges. In particular, $2^k a_{2^k} \to 0$.

Proof. Observe that since the terms of the sequence are non-negative, the partial sums S_N are increasing and hence their limit is $S = \sup\{S_N\}$. Now

$$\begin{aligned}
S_{2^k-1} &= (a_1) + (a_2 + a_3) + (a_4 + a_5 + a_6 + a_7) + \dots \\
&= \sum_{j=0}^{k-1} \sum_{n=2^j}^{2^{j+1}-1} a_n \\
&\geq \sum_{j=0}^{k-1} \sum_{n=2^j}^{2^{j+1}-1} a_{2^{j+1}} \\
&\geq \sum_{j=0}^{k-1} 2^j a_{2^{j+1}}.
\end{aligned}$$

So the partial sums of the series $\sum_{j=0}^{k-1} 2^j a_{2^{j+1}}$ bounded by $S_{2^k} \leq S$. Thus these partial sums also increase to their supremum and the series converges. □

In fact, Cauchy Condensation tells us that $\sum_{n=1}^{N} 1/n$ diverges like $\log N$. So the Divergence Test is not an if and only if, and in fact neither is the Cauchy Condensation Test. However, if the terms of the sequence happen to cancel out a bit, we can obtain convergence.

Theorem 5.2: Alternating Series Test

Suppose a_n is a decreasing sequence of positive numbers with $a_n \to 0$. Then

$$\sum_{n=1}^{\infty} (-1)^{n+1} a_n$$

converges.

Proof. We have

$$S_{2N} = (a_1 - a_2) + (a_3 - a_4) + \dots + (a_{2N} - a_{2N-1})$$

and this is a sum of non-negative terms, so S_{2N} is non-negative and increasing with N. Meanwhile

$$S_{2N+1} = a_1 - (a_2 - a_3) - (a_4 - a_5) + \dots + (a_{2N+1} - a_{2N}) = S_{2N} + a_{N+1}$$

and this sequence is decreasing with N, but still larger than S_{2N}, and in particular, larger than 0. This means S_{2N+1} converges by the Monotone Convergence Theorem. On the other hand,

$$S_{2N} \leq S_{2N+1} \leq S_1$$

so the sequence S_{2N} also converges by Monotone Convergence. Finally, $S_{2N+1} - S_{2N} = a_{2N+1} \to 0$, so in fact S_{2N+1} and S_{2N} converge to the same limit, S say, which means S_N also converges to S. □

So sometimes introducing cancellation among the terms of a series can make it converge. But cancellation can never make a series diverge.

Lemma 5.3

If $\sum_n |a_n|$ converges then so does $\sum_n a_n$.

Proof. We show that

$$S_N = \sum_{n=1}^{N} a_n$$

is Cauchy. Write

$$T_N = \sum_{n=1}^{N} |a_n|.$$

Then if $N > M$

$$S_N - S_M = \sum_{n=M+1}^{N} a_n$$

so

$$|S_N - S_M| = \sum_{n=M+1}^{N} |a_n| = T_N - T_M$$

and because T_N converges, it is a Cauchy sequence, and the right hand side is arbitrarily small provided M and N are sufficiently large. □

Definition 5.2: Absolutely convergent series

A series

$$\sum_{n=1}^{\infty} a_n$$

is called absolutely convergent if

$$\sum_{n=1} |a_n|$$

converges.

The above lemma says that absolutely convergent series are convergent series. In fact they converge even if you rearrange the terms.

Theorem 5.3: Rearrangements of absolutely convergent series converge

Let $\sum_n a_n$ be an absolutely convergent series converging to S and let $\sigma : \mathbb{N} \to \mathbb{N}$ be a bijection. Then

$$\sum_{n=1}^{\infty} a_{\sigma(n)}$$

is also absolutely convergent, and converges to S.

Proof. First we show that for any bijection $\sigma : \mathbb{N} \to \mathbb{N}$, the series

$$\sum_{n=1}^{\infty} |a_{\sigma(n)}|$$

converges, and the limit does not depend on σ. Indeed, write

$$T_N = \sum_{n=1}^{N} |a_N|, \ T_N^{\sigma} = \sum_{n=1}^{N} |a_{\sigma(n)}|.$$

Then, if $M_N = \max\{\sigma(1), \ldots, \sigma(N)\}$, we have

$$T_N^{\sigma} = \sum_{n=1}^{N} |a_{\sigma(n)}| \leq \sum_{n=1}^{M_N} |a_n| = T_{M_N}$$

because the summands in $|a_n|$ are all non-negative, and T_{M_N} has all the $|a_{\sigma(n)}|$ as summands. This shows that

$$T_N^{\sigma} \leq T_{M_N} \leq \sup T_M$$

and Since T_M is increasing,

$$\sup T_M = \lim_{M \to \infty} T_M = \sum_{n=1}^{\infty} |a_n|.$$

Thus

$$T_N^{\sigma} \leq \sum_{n=1}^{\infty} |a_n|,$$

and since T_N^{σ} is increasing,

$$\sum_{n=1}^{\infty} |a_{\sigma(n)}| = \sup_N T_N^{\sigma} \leq \sum_{n=1}^{\infty} |a_n|.$$

This shows that

$$\sum_{n=1}^{\infty} a_{\sigma(n)}$$

is absolutely convergent. Now we need only to show that $\sum a_n$ and $\sum a_{\sigma(n)}$ converge to the same thing. But, denoting by S_N and S_N^{σ} their respective partial sums, we get

$$|S_N - S_N^{\sigma}| = \left| \sum_{n=1}^{N} a_n - \sum_{n=1}^{N} a_{\sigma}(n) \right|.$$

But S_N and S_N^σ will both include $a_1, \ldots, a_M$ as summands if N is sufficiently large, since there are numbers $k_1, \ldots, k_M$ such that $\sigma(k_m) = m$ (by surjectivity), and all the summands common to S_N and S_N^σ are cancelled out. This means

$$|S_N - S_N^\sigma| = \left| \sum_{n=1}^{N} a_n - \sum_{n=1}^{N} a_\sigma(n) \right| \leq \sum_{m=M+1}^{\infty} |a_n| + \sum_{m=M+1}^{\infty} |a_{\sigma(n)}|$$

and both of these are the tails of convergent series, and so are smaller than ε if M (and in turn N) is large enough. So

$$\lim_{N\to\infty} S_N - S_N^\sigma = 0$$

and both series are in fact the same. □

The previous theorem shows that an absolutely convergent series can be added in any order and yield the same result. If the series is merely convergent (but not absolutely convergent) then this fails in the most spectacular way.

Theorem 5.4: Riemann's Rearrangement Theorem

Suppose $\sum_n a_n$ converges but not absolutely. Then for any $x \in \mathbb{R}$ there is a bijection $\sigma : \mathbb{N} \to \mathbb{N}$ such that $\sum_n a_{\sigma(n)} = x$.

Proof. Set

$$a_n^+ = \begin{cases} a_n & \text{if } a_n \geq 0, \\ 0 & \text{if } a_n < 0, \end{cases} \quad a_n^- = \begin{cases} -a_n & \text{if } a_n < 0, \\ 0 & \text{if } a_n \geq 0, \end{cases}$$

so that

$$a_n = a_n^+ - a_n^-.$$

Set

$$S_N^+ = \sum_{n=1}^{N} a_n^+, \; S_N^- = \sum_{n=1}^{N} a_n^-.$$

Because a_n^+ and a_n^- are both non-negative (from the way we defined them), S_N^+ and S_N^- are both increasing with N.

Now by definition,

$$S_N = \sum_{n=1}^{N} a_n = S_N^+ - S_N^+$$

while

$$\sum_{n=1}^{N} |a_n| = S_N^+ + S_N^-.$$

The second identity shows that one of S_N^+ or S_N^- must increase to infinity, since the left hand side does. But in fact the other must too, since if only S_N^+ diverged, then from the first identity

$$S_N + S_N^- = S_N^+$$

and the left hand side would converge because S_N converges by hypothesis, and S_N^- would be increasing and bounded. A similar argument shows that S_N^+ must diverge if S_N^- does. The important point is this: for any integer m,

$$\sum_{n=m}^{\infty} a_n^+ = \sum_{n=m}^{\infty} a_n^- = \infty.$$

This is because each series in its entirety diverges, and the first $m-1$ terms of that series have a finite sum.

We now proceed as follows. We will alternate adding some terms from S_N^+ in order and then subtracting some terms from S_N^- in order. In this way, all of the a_n's will be eventually added, and only once. The order in which we add these terms determines σ.

First, we set m_1 to be the minimum integer such that

$$S_{m_1}^+ \geq x.$$

This integer could be 0. So we have the first few non-negative a_n's until we surpassed x. Next let n_1 denote the smallest integer such that

$$S_{m_1}^+ - S_{n_1}^- < x.$$

Thus we have added in the first few negative a_n's to move to the other side of x. Next let m_2 denote the smallest integer such that

$$S_{m_2}^+ - S_{n_1}^- \geq x,$$

so we not continue adding positive terms from the series until we move to the other side of x again. And we continue this process indefinitely, adding non-negative terms, and then negative terms, so that we switch every time we move from one side of x to the other. We can always do this: because the tails of the non-negative and negative series are both infinite, we can always move left or right as far as we need to by adding terms from either series. Moreover, the moment we cross x, it is because we added a number $a_{m_k}^+$ or else $a_{n_k}^-$ which pushed us past x. So

$$x - a_{n_k}^- \leq S_{m_k}^+ - S_{n_k}^- < x, \; x + a_{m_{k+1}}^+ \geq S_{m_{k+1}}^+ - S_{n_k}^- \geq x$$

at each step of the way. This shows that at step k, we are never more than $a_{n_k}^-$ or $a_{m_{k+1}}^+$ away from x, and both of these quantities tend to 0 by the Divergence Criterion. Thus

$$S_{m_{k+1}}^+ - S_{n_k}^-, S_{m_k}^+ - S_{n_k}^- \to x.$$

All other partial sums lie between these, and hence the sequence of all partial sums converge to x. □

6 TOPOLOGY

6.1 Basic topology on $\mathbb{R}$

Definition 6.1: Open set

e call a subset U of $\mathbb{R}$ open if for each $x \in U$, there is some $\varepsilon > 0$ such that $(x-\varepsilon, x+\varepsilon) \subseteq U$.

In other words, U is open if every point in U is surrounded by a neighbourhood of points from U.

Lemma 6.1

Both $\mathbb{R}$ and $\{\}$ are open.

Proof. For each $x \in \mathbb{R}$, we have $(x-1, x+1) \subseteq \mathbb{R}$. For each $x \in \{\}$ we also have $(x-1, x+1) \subseteq \{\}$, vacuously. □

Lemma 6.2

Open intervals are open.

Proof. Let (a,b) be an open interval. We will assume $a,b<\infty$, that case can be handled similarly. If $a<x<b$ then $x-a$ and $b-x$ are both positive. Let

$$\varepsilon=\min\{x-a,b-x\}.$$

Then $(x-\varepsilon,x+\varepsilon)\subseteq(a,b)$ since for $t\in(x-\varepsilon,x+\varepsilon)$ we have

$$t>x-\varepsilon\geq x-(x-a)=a,\ t<x+\varepsilon\leq x+(b-x)=b.$$

□

Theorem 6.1

If $\mathscr{U}$ is any collection of open sets then $\bigcup_{U\in\mathscr{U}}U$. If $U_1,\ldots,U_N$ are open sets, then so is $U_1\cap\cdots\cap U_N$. Thus an arbitrary union of open sets is open, and a finite intersection of open sets is open.

Proof. If $x\in\bigcup_{U\in\mathscr{U}}U$ then $x\in U$ for some $U\in\mathscr{U}$, and since U is open, it must be that $(x-\varepsilon,x+\varepsilon)\subseteq U$ for some $\varepsilon>0$. But thus means $(x-\varepsilon,x+\varepsilon)\subseteq\bigcup_{U\in\mathscr{U}}U$.

If $x\in U_1\cap\ldots\cap U_N$ then for $n=1,\ldots,N$, there is some $\varepsilon_n>0$ such that $(x-\varepsilon_, x+\varepsilon_n)\subseteq U_n$. But if $\varepsilon=\min\{\varepsilon_1,\ldots,\varepsilon_N\}$ then $\varepsilon>0$ and

$$(x-\varepsilon,x+\varepsilon)\subseteq(x-\varepsilon_n,x+\varepsilon_n)\subseteq U_n$$

so that in fact $(x-\varepsilon,e+\varepsilon)\subseteq U_1\cap\ldots\cap U_N$. □

Example: It is **not** the case that arbitrary intersections of open sets are still open. For instance $(-1/n,1/n)$ is open for each $n\in\mathbb{N}$ but the intersection

$$\bigcap_{n\in\mathbb{N}}(-1/n,1/n)=\{0\}$$

is not.

Definition 6.2: Topological space

A set of points X along with a collection $\mathscr{U}$ consisting of subsets of X is said to form a topological space if the following conditions hold:

1. $X,\emptyset\in\mathscr{U}$,
2. if $\mathscr{U}'\subseteq\mathscr{U}$ is any collection of sets from $\mathscr{U}$ then $\bigcup_{U\in\mathscr{U}'}U$ also belongs to $\mathscr{U}$, and
3. if $U_1,\ldots,U_N\in\mathscr{U}$ then so is $U_1\cap\ldots\cap U_N$.

The sets $\mathscr{U}$ are called the open subsets of X.

What we have just shown is that $\mathbb{R}$ forms a topological space where $\mathscr{U}$ is the collection of open subsets U of $\mathbb{R}$ defined by the neighbourhood condition $x \in U \implies (x-\varepsilon, x+\varepsilon) \subseteq U$ for some $\varepsilon > 0$. A closed set in a topological space is one whose complement is open.

6.2 Open and closed sets

In a general topological space, the open sets are declared: U is open if and only if $U \in \mathscr{U}$. Similarly, the closed sets are the complements of the open sets. But these rules don't give a good sense of what the open or closed sets *are*. In the case of $\mathbb{R}$, we can characterize the open and closed sets.

Theorem 6.2: Closed using sequences

A set C is closed if and only if whenever $\{c_n\}$ is a sequence with $c_n \in C$ for each $n \in \mathbb{N}$ and $c_n \to L$, then $L \in C$ as well.

Proof. First we prove that if C is closed then if $\{c_n\}$ is a sequence with $c_n \in C$ for each $n \in \mathbb{N}$ and $c_n \to L$, we know $L \in C$. If $L \notin C$, then $L \in C^c$ which is open because C is closed. So $(L-\varepsilon, L+\varepsilon) \in C^c$ for some $\varepsilon > 0$. However, since $c_n \to L$, this means that $c_n \in (L-\varepsilon, L+\varepsilon)$ for all sufficiently large n. Then $c_n \in C$ and $c_n \in C^c$, a contradiction.

Conversely, suppose we know that every convergent sequence of numbers from C has its limit in C too. Let $x \in C^c$. If, for each n, $(x-1/n, x+1/n)$ is not a subset of C^c, then we can find an element, c_n, from this interval which belongs to C. By design, $c_n \to x$, but $c_n \in C$ for each n. Thus $x \in C$ also, but $x \in C^c$, again a contradiction. □

To characterize open sets we need the following lemma.

Lemma 6.3

Let U be an open set. Then for x in U, there is a largest interval $I_x = (l_x, r_x)$ such that $x \in I_x \subseteq U$, in the sense that $l_x, r_x \notin U$.

Proof. Let $r_x = \sup\{r \in \mathbb{R} : (x, r) \subseteq U\}$. Then we claim that $r_x > x$, $r_x \notin X$,and $(x, r_x) \subseteq X$. First, since $x \in U$ and U is open, there is an $\varepsilon > 0$ so that $(x, x+\varepsilon) \subseteq (x-\varepsilon, x+\varepsilon) \subseteq U$. Thus $r_x \geq x+\varepsilon > x$. Also, if $r_x \in U$ then we would also know that for some $\varepsilon > 0$, $(r_x-\varepsilon, r_x+\varepsilon) \subseteq U$. However, $r_x - \varepsilon/2 < r_x$, so $(x, r_x-\varepsilon/2) \subseteq U$ and hence

$$(x, r_x+\varepsilon) = (x, r_x-\varepsilon/2) \cup (r_x-\varepsilon, r_x+\varepsilon) \subseteq U$$

which is impossible since $r_x + \varepsilon > r_x$. Finally, since $r_x > x$ we have $x < r_x - 1/n < r_x$

for all n larger than $(r_x - x)^{-1}$. Thus $(x, r_x - 1/n) \subseteq U$ and so is

$$\bigcup_{n > (r_x - x)^{-1}} (x, r_x - 1/n) = (x, r_x).$$

In a similar way we show $l_x = \inf\{l : (l, x) \subseteq U\}$ satisfies $l_x < x$, $l_x \notin U$ and $(l_x, x) \subseteq U$. Thus $(l_x, r_x) \subseteq U$ while $l_x, r_x \notin U$. □

Theorem 6.3: The structure of open sets

Let U be an open set. Then U is the disjoint union of at most countably many open intervals.

Proof. For each $x \in U$ let I_x be the largest interval in U containing x. We claim that for $x \neq y$, either $I_x = I_y$ or else $I_x \cap I_y = \emptyset$. Indeed, if $z \in I_x \cap I_y$ then $m = \min\{a_x, a_y\} < z < \max\{b_x, b_y\} = M$ so (m, M) is an interval in U containing both I_x and I_y, and hence $(m, M) = I_x = I_y$ since the intervals I_x and I_y were defined to be as large as possible.

This shows that the set

$$\mathscr{I} = \{I_x : x \in U\}$$

is a collection of disjoint intervals (provided we this of I as a set, so no intervals are repeated in it). Since $x \in I_x \in \mathscr{I}$,

$$U \subseteq \bigcup_{I \in \mathscr{I}} I,$$

and since each $I \in \mathscr{I}$ is a subset of U, we also have

$$\bigcup_{I \in \mathscr{I}} I \subseteq U.$$

So U is a union of the disjoint intervals from $\mathscr{I}$. To see that $\mathscr{I}$ is countable, observe that since each $I \in \mathscr{I}$ is open and non-empty, it contains a rational number. By choosing one rational number from each I in $\mathscr{I}$, and observing that these rational numbers must be distinct for each I by disjointness, we create an injection from $\mathscr{I}$ to $\mathbb{Q}$, so that $\mathscr{I}$ must be countable or else finite. □

6.3 Connected sets

Definition 6.3: Connected set

In a topological space, a set D is said to be **disconnected** if one can find open sets U_1 and U_2 such that

1. $U_1 \cap U_2 = \emptyset$,
2. $U_1 \cap D \neq \emptyset$ and $U_2 \cap D \neq \emptyset$, and
3. $D \subseteq U_1 \cup U_2$.

A set which is not disconnected is called connected.

This seems to be a very abstract definition but it turns out to be very good at capturing our intuition of connectedness all the while being flexible enough to work with.

Theorem 6.4: Connectedness Criterion

A set $A \subseteq \mathbb{R}$ is connected if and only if it is an interval (i.e. for each $a, b \in A$ with $a < b$, we have $(a, b) \subseteq A$).

Proof. Suppose A is connected. If $a, b \in A$ with $a < b$, let $c \in (a, b)$. If $c \notin A$ then $(-\infty, c) \cup (c, \infty)$ disconnects A since each set is open, they are disjoint, $a \in (-\infty, c)$, $b \in (c, \infty)$, and $(-\infty, c) \cup (c, \infty) = \mathbb{R} \setminus \{c\} \supset A$.

Conversely suppose we have an interval and U_1 and U_2 are open sets which together contain A, each intersecting A in a non-empty set. We will show that $U_1 \cap U_2$ is non-empty, so that A cannot be disconnected. Suppose $a, b \in A$ be such that $a \in U_1$ and $b \in U_2$, and without loss of generality $a < b$. Then $[a, b] \subseteq A$ since A is an interval. Since U_1 is open, there is a maximal interval $I_a = (l_a, r_a) \subseteq U_1$ containing a. If $r_a \geq b$ then $b \in U_1$ as well and we're done. Otherwise $r_a < b$. Since I_a is maximal, $r_a \notin U_1$ but $r_a \in (a, b)$ and so $r_a \in U_2$. Since U_2 is open, it contains an interval $(r_a - \varepsilon, r_a + \varepsilon)$. But then $(a, r_a) \cap (r_a - \varepsilon, r_a + \varepsilon)$ is non-empty and hence so is $U_1 \cap U_2$. □

6.4 Compactness

We are going to introduce compactness, which is a tool for passing from the infinite to the finite. There are a few definitions of compactness, and they are different in some topological spaces. In the context of $\mathbb{R}$, however, they turn out the be the same thing, which will be the main theorem we will prove.

Definition 6.4: Sequential compactness

A set A is said to be sequentially compact if whenever $\{a_n\}$ is a sequence with $a_n \in A$ for each $n \in \mathbb{N}$, there is a subsequence $\{a_{n_k}\}$ which converges to a limit in A.

Example: Any finite set is sequentially compact. Indeed if $A = \{x_1, \ldots, x_n\}$ and $\{a_n\}$ is a sequence of numbers from A then, by the pigeonhole principle, one of the numbers x_i appears infinitely often as some a_n. In other words $\{a_n\}$ has a subsequence which is just the constant sequence $\{x_i\}$, which converges to $x_i \in A$.

Example: The set $[0,1)$ is not sequentially compact. Indeed, the sequence $\{1-1/n\}$ converges to 1, and hence so too does any subsequence. So no subsequence of $\{1-1/n\}$ can converge to a number in $[0,1)$.

The failure of sequential compactness here comes from the fact that $(0,1]$ is not closed. This is a general phenomenon.

Lemma 6.4

If A is sequentially compact, then A is closed.

Proof. Our goal is to show that A^c is open. If $A^c = \emptyset$, then it is open by definition. If not, let $x \in A^c$. Consider the intervals $(x-1/n, x+1/n)$ with $n \in \mathbb{N}$. If one of these intervals is contained in A^c, then because x was arbitrary, we will have shown A^c is open. If not, then each such interval intersects A at some point, say a_n. By design, $a_n \to x$. But any subsequence of $\{a_n\}$ also converges to x. Since $\{a_n\}$ has a subsequence converging to a limit from A, we would deduce that $x \in A$, contradicting the fact that $x \in A^c$. □

We could have used the alternative characterization of closed sets in the above lemma, which would have been faster. In the proof provided, we constructed a sequence, which is more or less the same is the proof of the alternative characterization.

Example: The set $\mathbb{Z}$ is not sequentially compact. This is because it contains the sequence $\{n\}$ which diverges to ∞.

Now the issue is that our sequence is not bounded, but rather diverges to infinity. Again, this is a general thing.

Lemma 6.5

If A is sequentially compact then A is bounded.

Proof. If A is not bounded, then $A \not\subseteq [-N, N]$ for any N. So we can construct a sequence of numbers a_N with $|a_N| \to \infty$. Any subsequence of this will fail to converge, since it has to be unbounded. □

Theorem 6.5

A set A is sequentially compact if and only if it is both closed and bounded.

Proof. We have already shown A needs to be both closed and bounded to be sequentially compact. Now lets show the converse. Suppose that $\{a_n\}$ is a sequence from A. Since A is bounded, so is $\{a_n\}$, and hence $\{a_n\}$ has a convergent subsequence $\{a_{n_k}\}$ by the Bolzano-Weierstrass theorem. Since A is closed, we must also have that $\lim_{k\to\infty} a_{n_k} \in A$, showing that $\{a_n\}$ has a subsequence converging in A. Thus A is sequentially compact. □

Definition 6.5: Open cover

If X is a topological space and $Y \subseteq X$, then an open cover of Y is any collection $\mathscr{U}$ consisting of open sets and such that $Y \subseteq \bigcup_{U\in\mathscr{U}} U$.

Example: The sets $(n, n+2)$ with $n \in \mathbb{Z}$ are an open cover of $\mathbb{R}$. The sets $(n, n+1)$ with $n \in \mathbb{Z}$ are not an open cover of $\mathbb{R}$ because they fail to cover $\mathbb{Z}$.

Example: The sets $(-1/n, 1+1/n)$ with $n \in \mathbb{N}$ forms an open cover of $[0,1]$. It is pretty redundant, as any one of these sets will do the job.

Example: The sets $(1/n - 1/n^2, 1/n + 1/n^2)$ with $n \in \mathbb{N}$ forms an open cover of the set $\{1/n : n \in \mathbb{N}\}$.

Definition 6.6: Compactness

A set A is said to be compact if whenever $\mathscr{U}$ is an open cover of A, there is a finite subset $\mathscr{U}' \subseteq \mathscr{U}$ which is also an open cover of A.

Example: The set $(0,1]$ is not compact. Indeed $\mathscr{U} = \{(1/n, 2) : n \in \mathbb{N}\}$ forms an open cover, since $x \in (0,1]$ means $1 \geq x > 1/n$ for some $n \in \mathbb{N}$ and so $x \in (1/n, 2)$. However, if $\mathscr{U}'$ is finite subset of $\mathscr{U}$ then there is some largest m such that $(1/m, 2) \in \mathscr{U}'$ and this contains all the other sets from $\mathscr{U}'$. But $(1/m, 2)$ does not contain $(0,1]$.

Example: The set $[0,1]$ is compact. But rather than prove this explicitly, we'll just prove the following theorem.

Theorem 6.6: Heine-Borel

The following are equivalent for a set $A \subseteq \mathbb{R}$:

1. A is closed and bounded,
2. S is sequentially compact, and
3. A is compact.

Proof. We've already established the equivalence of (1) and (2). So now let's prove that they are equivalent to (3).

First, we'll show that (3) implies (1). So suppose A is compact. Let $x \in A^c$, and consider the open sets $(-\infty, x-1/n) \cup (x+1/n, \infty)$ with $n \in \mathbb{N}$. The union of these open sets is $\mathbb{R} \setminus \{x\}$ which contains A. So these sets form an open cover of A and hence have a finite subcover. But these sets are growing in size, and so the finite subcover can be reduced to the largest set in it, which means there is some n such that $A \subseteq (-\infty, x-1/n) \cup (x+1/n, \infty)$. Thus $(x-1/n, x+1/n) \subseteq A^c$, and since x was arbitrary, this shows that A^c is open, and A is closed. In addition, the sets $(-n, n)$ with $n \in \mathbb{N}$ form an open cover of $\mathbb{R}$ and hence of A. A finite subcover must contain some largest $(-n, n)$ which in turn contains A, which shows that A is bounded.

Now let's show that (1) implies (3). So suppose that $\mathscr{U}$ is an open cover of a closed and bounded set $A \subseteq [-M, M]$. We define the following process. Initially, we set $A_1 = A$ and $I_1 = [-M, M]$. We assume that A_1 cannot be covered by a finite subset of $\mathscr{U}$, or else we'd be done. At stage j, we have an interval $I_j = [l_j, r_j]$ and a set $A_j \subseteq A_{j-1}$ such that A_j cannot be covered by a finite subset of $\mathscr{U}$. We split the interval I_j in half at the midpoint to get two intervals I_j^l and I_j^r, each half as long as I_j. Then A_j gets split in half as $A_j^l = A_j \cap I_j^l$ and $A_j^r = A_j \cap I_j^r$. Since A_j cannot be covered by a finite subset of $\mathscr{U}$, the same must be true of either A_j^l or A_j^r (if each could be covered by a finite subset of $\mathscr{U}$, combining these two finite subsets would produce a, possibly larger, finite subset of $\mathscr{U}$ which covered all of A_j). We set A_{j+1} to be whichever of A_j^l or A_j^r cannot be covered, and I_{j+1} to be the corresponding half of I_j. In this way, we produce a decreasing sequence of sets $A_{j+1} \subseteq A_j \subseteq A$ and a decreasing sequence of intervals $I_{j+1} \subseteq I_j$, such that the intervals I_{j+1} are half as long as their predecessor. Now, each A_{j+1} has to be non-empty (or else it would certainly be covered by a finite subset of $\mathscr{U}$), so choose some element of A_j for each j to produce a sequence $\{a_j\}$. Because A is closed and bounded, we can pass to a subsequence $\{a_{j_k}\}$ which converges to some $a \in A$ (because A is closed). Now $a \in A$ means that there is some $U \in \mathscr{U}$ which contains it – after all, $\mathscr{U}$ covers all of A. Since U is open, there is some $\varepsilon > 0$ such that $(a-\varepsilon, a+\varepsilon) \subseteq U$. Let j be so large that $|a_j - a| < \varepsilon/2$ and such that I_j has length at most $\varepsilon/4$ (this is possible since the length of I_j is at most $2M/2^j$, having been halved at each step). Let's now take stock: $a_j \in A_j \subseteq I_j$ and $|a_j - a| < \varepsilon/2$. But the endpoints of $I_j = [l_j, r_j]$ are at most $\varepsilon/4$ away from a_j. So

$$|l_j - a| \leq |l_j - a_j| + |a_j - a| < 3\varepsilon/4 < \varepsilon$$

and similarly $|r_j - a| < \varepsilon$. This means that

$$A_j \subseteq I_j \subseteq (a-\varepsilon, a+\varepsilon) \subseteq U,$$

and so A_j can be covered by a *single* subset of $\mathscr{U}$, a contradiction. □

7 CONTINUITY

7.1 Continuous Limits

Recall that a sequence is really just a function $a : \mathbb{N} \to \mathbb{R}$ which we usually write as $a(n) = a_n$, and we write $\lim_{n\to\infty} a_n = L$ if $a_n = a(n)$ is within ε of L for all n sufficiently “close" to ∞. This is exactly the same for continuous limits.

Definition 7.1: Limit

If f is a function defined on an interval surrounding a then we write

$$\lim_{x\to a} f(x) = L$$

if for each $\varepsilon > 0$, there is a $\delta > 0$ such that we have $|f(x) - L| < \varepsilon$ for all x with $|x - a| < \delta$, save for possibly $x = a$

Definition 7.2: One sided limit

If f is a function defined on an interval with left endpoint at a then we write

$$\lim_{x\to a^+} f(x) = L$$

if for each $\varepsilon > 0$, there is a $\delta > 0$ such that we have $|f(x) - L| < \varepsilon$ for all x with $a < x < a + \delta$. Similarly, if f is a function defined on an interval with right endpoint at a then we write

$$\lim_{x\to a^-} f(x) = L$$

if for each $\varepsilon > 0$, there is a $\delta > 0$ such that we have $|f(x) - L| < \varepsilon$ for all x with $a - \delta < x < a$.

Definition 7.3: Continuity

We say f is continuous at a if $\lim_{x\to a} f(x) = f(a)$. We say f is continuous on A if for each $a \in A$, f is continuous at a.

Unravelling the definition, we see that f is continuous at a if for each $\varepsilon > 0$ there is some $\delta > 0$ such that $|x - a| < \delta$ tells us that $|f(x) - f(a)| < \varepsilon$. Notice that in this definition, the parameter δ depends implicitly on ε and on a. This is necessarily the case.

Example: The function $x \mapsto x^2$ is continuous on all of $\mathbb{R}$. Indeed, for $a \in \mathbb{R}$ and $|x - a| < \delta$ we have

$$|x^2 - a^2| = |x - a||x + a| < \delta(|x| + |a|)$$

and

$$|x| \le |x - a| + |a| \le \delta + |a|$$

so that

$$|x^2 - a^2| \le \delta(\delta + 2|a|)$$

and this can be made at most ε by taking δ sufficiently small. Indeed, if $\delta < \min\{\sqrt{\varepsilon}/2, \varepsilon/4|a|\}$ then

$$\delta(\delta + 2|a|) = \delta^2 + 2\delta|a| < \varepsilon/2 + \varepsilon/2 < \varepsilon.$$

Example: We have $\lim_{x\to 0} \sin(x)/x = 1$. Indeed if $\varepsilon > 0$ we shall show that $1 - \varepsilon < \sin(x)/x < 1$ as $x \to 0$ from the right, the left hand following from the fact that $\sin(x)/x$ is even. Now $\sin(x) < x$ for all positive x, while for x sufficiently small, we have $\cos(x) \le \sin(x)/x$. However,

$$\cos(x) = \sqrt{1 - \sin(x)^2} \ge \sqrt{1 - x^2} > \sqrt{1 - x} \ge \sqrt{1 - \delta},$$

if $0 < x < \delta$. Thus

$$1 - \frac{\sin(x)}{x} \leq 1 - \sqrt{1-\delta} = \frac{1-(1-\delta)}{1+\sqrt{1-\delta}} \leq \delta.$$

So if $\delta = \varepsilon$, we have $1 - \varepsilon < \sin(x)/x < 1$ as required.

7.2 Other characterizations of continuity

Definition 7.4: Image and preimage

If $f : X \to Y$ is a function and $A \subseteq X$ then

$$f(A) = \{f(a) : a \in A\}$$

is the image of A under f. If $B \subseteq Y$ then

$$f^{-1}(B) = \{x \in X : f(x) \in B\}$$

is the preimage (or inverse image) of B.

Theorem 7.1: Continuity from topology

A function f is continuous if and only if for any open set U,

Proof. Suppose f is continuous. If U is open, let $x \in f^{-1}(U)$, which is to say $f(x) \in U$. Then there is some $\varepsilon > 0$ such that $(f(x)-\varepsilon, f(x)+\varepsilon) \subseteq U$ as well, since U is open. There is some $\delta > 0$ such that $|f(x) - f(y)| < \varepsilon$ when $|x - y| < \delta$. That is to say, if $y \in (x-\delta, x+\delta)$ then $f(y) \in (f(x)-\varepsilon, f(x)+\varepsilon) \subseteq U$, which means that $(x-\delta, x+\delta) \subseteq f^{-1}(U)$.

Conversely, suppose the inverse image of any open set is open. Then if $\varepsilon > 0$, the inverse image of $(f(x)-\varepsilon, f(x)+\varepsilon)$ is open, and it contains x. So there is some $\delta > 0$ such that $(x-\delta, x+\delta) \subseteq f^{-1}((f(x)-\varepsilon, f(x)+\varepsilon))$, which means that $|x-y| < \delta$ implies $f(y) \in (f(x)-\varepsilon, f(x)+\varepsilon)$ which is the definition of continuity. □

Another characterization of continuity concerns sequences.

Theorem 7.2

The function f defined on an interval around a is continuous at a if and only if whenever $\{a_n\}$ is a sequence converging to a (but distinct from a) and contained in the domain of f, we have $f(a_n) \to f(a)$

Proof. Suppose f is continuous at a and let $\varepsilon > 0$. Let $\delta > 0$ be such that $|x - a| < \delta$

implies $|f(x)-f(a)|<\varepsilon$, and let N be so large that $|a_n-a|<\delta$ for $n>N$. Then $|f(a_n)-f(a)|<\varepsilon$ and this shows that $f(a_n)\to f(a)$.

Conversely, if f is not continuous at a then there is some $\varepsilon>0$ such that for *no* $\delta>0$ do we have $|x-a|<\delta$ implies $|f(x)-f(a)|<\varepsilon$. This means that if we attempt $\delta=1/n$, for each $n\in\mathbb{N}$, there will always be some number a_n with $|a_n-a|<1/n$ but $|f(a_n)-f(a)|>\varepsilon$. By construction $a_n\to a$ but $f(a_n)\not\to f(a)$, so the sequence condition also fails. □

The previous characterization of continuity is the exact same as the following.

Theorem 7.3: Composition and continuity

Suppose f is continuous at a and g is continuous at $f(a)$. Then $g\circ f$ is continuous at a.

Proof. Let $\varepsilon>0$ and let $\delta_g>0$ be such that $|g(y)-g(f(a))|<\varepsilon$ whenever $|y-f(a)|<\delta_g$. Next, let δ_f be such that $|f(x)-f(a)|<\delta_g$ whenever $|x-a|<\delta_f$. Then if $|x-a|<\delta_f$ we have $|g(f(x))-g(f(a))|<\varepsilon$, showing continuity of $g\circ f$ at a. □

7.3 Algebraic properties of continuity

Theorem 7.4

Suppose that f and g are two functions defined in an interval around a, and each is continuous at a. Then so is $f+g$, fg and f/g provided $g(a)\neq 0$.

Proof. We just show the product is continuous, the others are left as an exercise. So let $\varepsilon>0$ and suppose we know $|x-a|<\delta$. Then

$$|f(x)g(x)-f(a)g(a)|=|f(x)g(x)-f(a)g(x)+f(a)g(x)-f(a)g(a)|$$
$$\leq|f(x)g(x)-f(a)g(x)|+|f(a)g(x)-f(a)g(a)|\leq|g(x)||f(x)-f(a)|+|f(a)||g(x)-g(a)|.$$

If δ is sufficiently small, we can guarantee that $|g(x)-g(a)|<\varepsilon/(2|f(a)|)$, that $|g(x)|<|g(a)|+1$ and that $|f(x)-f(a)|<\varepsilon/2(|g(a)+1)$, all of which in turn guarantees that

$$|f(x)g(x)-f(a)g(a)|<\varepsilon.$$

□

7.4 Continuity and compactness

Theorem 7.5

Let $f : \mathbb{R} \to \mathbb{R}$ be a continuous function and let C be a compact set. Then on C, f is uniformly continuous.

Proof. Let $\varepsilon > 0$. By continuity, for $a \in C$, there is some $\delta_a > 0$ such that whenever $b \in C$ and $|b - a| < \delta_a$, we have $|f(b) - f(a)| < \varepsilon/2$. Let $I_a = (a - \delta_a/2, a + \delta_a/2)$. Then I_a is an open interval and $a \in I_a$. Let $\mathscr{U} = \{I_a : a \in C\}$, which is an open cover of C. By compactness, there is a finite subcover $\mathscr{U}' = \{I_{a_1}, \ldots, I_{a_n}\}$. Now suppose $a, b \in C$ are arbitrary, and $|a - b| < \delta$ where

$$\delta = \min\{\delta_{a_1}/2, \ldots, \delta_{a_n}/2\} > 0.$$

Then, since $\mathscr{U}'$ is a cover, $a \in I_{a_j}$ for some j, which we may assume to be 1, which means

$$|a - a_1| < \delta_{a_1}/2 < \delta_{a_1}.$$

Since $|b - a| < \delta_{a_1}/2$, we have

$$|b - a_1| \le |b - a| + |a - a_1| \le 2\delta_{a_1}/2 = \delta_{a_1}$$

and, since both a and b are within δ_{a_1} of a_1, we have

$$|f(b) - f(a)| \le |f(b) - f(a_1)| + |f(a_1) - f(a)| < 2\varepsilon/2 = \varepsilon.$$

□

Theorem 7.6

Let $f : \mathbb{R} \to \mathbb{R}$ be continuous and let C be a compact set. Then $f(C)$ is also compact.

Proof. Let $\mathscr{U}$ be an open cover of $f(C)$. Then $\mathscr{V} = \{f^{-1}(U) : U \in \mathscr{U}\}$ consists of open sets by continuity, thus forming an open cover of C, by the definition of inverse image. Since C is compact, $\mathscr{V}$ has a finite subcover $\mathscr{V}' = \{f^{-1}(U_1), \ldots, f^{-1}(U_n)\}$. The sets $\{U_1, \ldots, U_n\}$ form an open cover of $f(C)$, which is a finite subcover of $\mathscr{U}$. □

Lemma 7.1

Let C be compact, then C contains a maximal and minimal element.

Proof. We just prove the result for a maximal element, the minimal element is similar. By compactness, C is bounded, so $M = \sup(C)$ exists. If $M \notin C$ then the sets

$(-\infty, M - 1/n)$ with $n \in \mathbb{N}$ form an open cover of C which cannot have a finite subcover – the sets are increasing, so a finite subcover would mean that $(-\infty, M - 1/n)$ contains C for some n. But that would mean $M - 1/n$ is an upper bound for C, a contradiction. □

Corollary 7.1: Extreme Value Theorem

Let $f : C \to \mathbb{R}$ be a continuous function, where C is compact. Then f attains a maximum and minimum value.

Proof. We know that $f(C)$ is compact, and so contains maximal and minimal elements. □

Theorem 7.7: Intermediate Value Theorem

Suppose A is a connected set and $f : A \to \mathbb{R}$ is continuous. Then $f(A)$ is also connected. Hence $f([a, b])$ is an interval, and contains all numbers between $f(a)$ and $f(b)$.

Proof. If U, V form a disconnecting pair of sets for $f(A)$ then $f^{-1}(U)$ and $f^{-1}(V)$ form a disconnecting pair of sets for A, which was assumed to be connected. □

In the previous theorems, we insisted on a function f which is continuous on all of $\mathbb{R}$, not just C. This is only to avoid technicalities in the proofs in dealing with boundary points – if we have open sets U containing C, then $f^{-1}(U)$ may not be open as a subset of $\mathbb{R}$, and we need to use something called the subspace topology, which is beyond the scope of this course. For the most part, one can use sequences instead for easy alternative proofs, which avoid these technicalities. These are more real analysis than topology, which is appropriate for this course, but I wanted to emphasize the role played by topology here.

8

DIFFERENTIABILITY

8.1 Basics of differentiability

Definition 8.1: Differentiability at a point

A function f, defined on an interval around a number a, is said to be differentiable at a if

$$\lim_{h\to 0} \frac{f(a+h)-f(a)}{h}$$

exists. In that case, the limit is called the derivative of f at a and denoted $f'(a)$.

Example: The function x^2 is everywhere differentiable with derivative $2x$. Indeed, at the point a,

$$\frac{(a+h)^2-a^2}{h} = 2a+h$$

which plainly converges to $2a$ as $h \to 0$.

An alternative characterization of differentiability is this: the function f is differentiable at a if and only if there is a number $f'(a)$ and a function $e_a(x)$ such that

$$f(x) = f(a) + f'(a)(x-a) + e_a(x)$$

and $e_a(x)/(x-a) \to 0$ as $x \to a$. Indeed, we merely define

$$e_a(x) = f(x) - f(a) - f'(a)(x-a)$$

and observe that

$$\frac{e_a(x)}{x-a} = \frac{f(x)-f(a)}{x-a} - f'(a)$$

and the right hand side tends to zero as $x \to a$.

This alternate characterization is particularly useful in applications.

Lemma 8.1

If f is differentiable at a then it is continuous at a.

Proof. From our alternate characterization

$$f(x) = f(a) + (x-a)\left(f'(a) + \frac{e_a(x)}{x-a}\right)$$

and the second term on the right tends to 0 as $x \to a$. □

Continuity alone is, however, far from sufficient.

Example: The function $|x|$ is not differentiable at 0.

Proof. We have

$$\frac{f(0+h)-f(0)}{h} = \begin{cases} 1 & \text{if } h > 0 \\ -1 & \text{if } h < 0. \end{cases}$$

Thus the left and right hand limits as $h \to 0$ are distinct and so the limit does not exist. □

Here are a few more familiar rules from calculus.

Theorem 8.1: Product Rule

If f and g are each differentiable at a then so is fg and its derivative is $f'(a)g(a) + g'(a)f(a)$.

Proof. We use the alternate characterization,

$$\begin{aligned} f(x)g(x) &= \left(f(a) + f'(a)(x-a) + e_{f,a}(x)\right)\left(g(a) + g'(a)(x-a) + e_{g,a}(x)\right) \\ &= f(a)g(a) + (f'(a)g(a) + g'(a)f(a))(x-a) + \\ &\quad + \left(e_{f,a}(x)\left(g(a) + g'(a)(x-a) + e_{g,a}(x)\right) + e_{g,a}(x)\left(f(a) + f'(a)(x-a)\right)\right) \end{aligned}$$

and the final expression in brackets when divided by $x-a$ tends to 0 as $x \to a$. □

Theorem 8.2: Chain Rule

If f is differentiable at a and g is differentiable at $f(a)$ then $g \circ f$ is differentiable at a and its derivative is $g'(f(a))f'(a)$.

Proof. Again we use the alternate characterization,

$$f(x) = f(a) + f'(a)(x-a) + e_{f,a}(x), \; g(y) = g(f(a)) + g'(f(a))(y - f(a)) + e_{g,f(a)}(y).$$

Set $y = f(x)$, then since f is, in particular, continuous at a, $y \to f(a)$ as $x \to a$. Thus,

$$g(f(x)) = g(f(a)) + g'(f(a))(f(x) - f(a)) + e_{g,f(a)}(f(x))$$

and writing $f(x) - f(a) = f'(a)(x-a) + e_{f,a}(x)$, we get

$$g(f(x)) = g(f(a)) + g'(f(a))f'(a)(x-a) + \left(g'(f(a))e_{f,a}(x) + e_{g,f(a)}(f(x))\right).$$

The error term in brackets, when divided by $x - a$ is

$$\frac{e_{f,a}(x)}{x-a} g'(f(a)) + \frac{f(x)-f(a)}{x-a} \frac{e_{g,f(a)}(f(x))}{f(x)-f(a)}$$

which tends to 0 as $x \to a$. □

8.2 Derivatives and local behaviour

Lemma 8.2

Suppose f is differentiable at a. If $f'(a) > 0$ then f is increasing in an interval around a, while if $f'(a) < 0$ then f is decreasing in an interval around a.

Proof. Suppose $f'(a) > 0$ and x is sufficiently close to a. Since

$$f(x) = f(a) + f'(a)(x-a) + e_a(x)$$

we can choose x so close to a that $e_a(x) < |x-a| f'(a)/2$. Then

$$f(x) - f(a) = f'(a)(x-a)\left(1 + \frac{e_a(x)}{(x-a)f'(a)}\right)$$

and the quantity in brackets has to be positive, since it's at least $1/2$. From this we see that the sign of $f(x) - f(a)$ and the sign of $x - a$ are the same, which means f is increasing. A similar proof works when $f'(a) < 0$. □

Corollary 8.1

Suppose f is differentiable on the interval (a, b) and there is a local maximum or minimum at some $c \in (a, b)$. Then $f'(c) = 0$.

Proof. Local extrema occur when f changes from increasing to decreasing or vice versa. Since c is a local extremum, f' cannot be increasing at c, so $f'(c) > 0$ is impossible. Similarly f cannot be decreasing at c, so $f'(c) < 0$ is impossible. □

Theorem 8.3: Rolle's Theorem

Suppose f is differentiable inside the interval $[a, b]$, continuous at the endpoints, and such that $f(a) = f(b)$. Then there is a number $c \in (a, b)$ for which $f'(c) = 0$.

Proof. The function f is continuous on the compact interval $[a, b]$ and so achieves a maximum and minimum value. If both of these occur at the endpoints then f has to be constant, and so $f' = 0$ everywhere. So we can assume that f has a local extremum c inside (a, b) and there we must have that $f'(c) = 0$. □

Corollary 8.2: The Generalized Mean Value Theorem

Suppose f and g are differentiable functions on the open interval (a, b) which are continuous at the endpoints a and b. Then there is a point $c \in (a, b)$ for which

$$f'(c)(g(b) - g(a)) = g'(c)(f(b) - f(a)).$$

Proof. Consider the differentiable function

$$h(x) = f(x)(g(b) - g(a)) - g(x)(f(b) - f(a)).$$

Then

$$h(a) = f(a)g(b) - g(a)f(b) = h(b)$$

and so by Rolle's Theorem, we find a c for which

$$h'(c) = f'(c)(g(b) - g(a)) - g'(c)(f(b) - f(a)) = 0$$

and the theorem follows. □

8.3 Taylor's Theorem

We close the course with a very useful approximation theorem which lets us replace any sufficiently nice function with a polynomial, at least when close to a given point.

Theorem 8.4: Taylor's Theorem with remainder

Suppose f is a function which is $N+1$-times differentiable on $(a-R, a+R)$. Then for any t in this interval, there is a number c between a and t for which

$$f(t) = \sum_{n=0}^{N} \frac{f^{(n)}(a)}{n!}(t-a)^n + \frac{f^{(N+1)}(c)}{(N+1)!}(t-a)^{N+1}.$$

Proof. We set

$$E(x) = f(x) - \sum_{n=0}^{N} \frac{f^{(n)}(a)}{n!}(x-a)^n$$

and check that

$$E^n(a) = 0,\ 0 \le n \le N$$

and

$$E^{N+1}(x) = f^{(N+1)}(x).$$

Iteratively, we apply the Generalized Mean Value Theorem as follows. Start with the functions $E(x)$ and $(x-a)^{N+1}$ at the points a and t. This tells us that

$$\frac{E(t)}{(t-a)^{N+1}} = \frac{E(t)-E(a)}{(t-a)^{N+1}-(a-a)^{N+1}} = \frac{E'(t_1)}{(N+1)(t_1-a)^N}.$$

At stage n we apply this same strategy to functions $E^{(n-1)}(x)$ and $(x-a)^{N+2-n}$ the points a and t_{n-1} and get a point t_n between t_{n-1} and a such that

$$\frac{E^{(n)}(t_{n-1})}{(t_n-a)^{N+2-n}} = \frac{E^{(n+1)}(t_n)}{(N+2-n)(t_n-a)^{N+1-n}}.$$

We do this until $n = N+1$ at which point we get a number t_{N+1} between x and a for which

$$\frac{E^{(N)}(t_n)}{t_N - a} = E^{(N+1)}(t_{N+1}) = f^{(N+1)}(t_{N+1}).$$

We iteratively substitute back to get

$$\frac{E(t)}{(t-a)^{N+1}} = \frac{f^{(N+1)}(t_{N+1})}{(N+1)!}$$

and then

$$f(t) = \sum_{n=0}^{N} \frac{f^{(n)}(a)}{n!}(t-a)^n + E(t) = \sum_{n=0}^{N} \frac{f^{(n)}(a)}{n!}(t-a)^n + \frac{f^{(N+1)}(t_{N+1})}{(N+1)!}(t-a)^{N+1}$$

□