

Math 140: Foundations of Real Analysis

Todd Kemp

Contents

Part 1. Math 140A	5
Chapter 1. Ordered Sets, Ordered Fields, and Completeness	7
1. Lecture 1: January 5, 2016	7
2. Lecture 2: January 7, 2016	9
3. Lecture 3: January 11, 2016	12
4. Lecture 4: January 14, 2014	14
Chapter 2. Sequences and Limits	19
1. Lecture 5: January 19, 2016	19
2. Lecture 6: January 21, 2016	21
3. Lecture 7: January 26, 2016	24
4. Lecture 8: January 28, 2016	27
Chapter 3. Extensions of \mathbb{R} ($\overline{\mathbb{R}}$ and \mathbb{C})	33
1. Lecture 9: February 2, 2016	33
2. Lecture 10: February 4, 2016	37
Chapter 4. Series	41
1. Lecture 11: February 9, 2016	41
2. Lecture 12: February 11, 2016	46
3. Lecture 13: February 16, 2016	51
Chapter 5. Metric Spaces	53
1. Lecture 14: February 18, 2016	56
2. Lecture 15: February 23, 2016	59
3. Lecture 16: February 25, 2016	61
4. Lecture 17: February 29, 2016	64
Chapter 6. Limits and Continuity	67
1. Lecture 18: March 3, 2016	67
2. Lecture 19: March 8, 2016	70
3. Lecture 20: March 10, 2016	73
Part 2. Math 140B	77
Chapter 7. More on Continuity	79
1. Lecture 1: March 29, 2016	79
2. Lecture 2: March 31, 2016	83
3. Lecture 3: April 5, 2016	85
Chapter 8. Differentiation of Functions of a Real Variable	91

1. Lecture 4: April 7, 2016	91
2. Lecture 5: April 12, 2016	95
3. Lecture 6: April 14, 2016	100
4. Lecture 7: April 19, 2016	103
Chapter 9. Integration	111
1. Lecture 8: April 21, 2016	111
2. Lecture 9: April 26, 2016	114
3. Lecture 10: April 28, 2016	118
4. Lecture 11: May 3, 2016	122
5. Lecture 12: May 5, 2016	127
Chapter 10. Sequences and Series of Functions	133
1. Lecture 13: May 10, 2016	133
2. Lecture 14: May 12, 2016	138
3. Lecture 15: May 17, 2016	142
4. Lecture 16: May 19, 2016	147
5. Lecture 17: May 24, 2016	151
6. Lecture 18: May 26, 2016	157
7. Lecture 19: May 31, 2016	161
8. Lecture 20: June 2, 2016	169

Part 1

Math 140A

CHAPTER 1

Ordered Sets, Ordered Fields, and Completeness

1. Lecture 1: January 5, 2016

- $\mathbb{N}, \mathbb{Z}, \mathbb{Q}, \mathbb{R}, \mathbb{C}$.
- \mathbb{R} is the “Réal numbers”. There is nothing *real* about them! That is the first, most important lesson to learn in this class. We will encounter many “obvious” statements that are, in fact, *false*. We will also see some counterintuitive statements that turn out to be true.
- Mathematicians roughly split into two groups: analysts and algebraists. (There’s lots of overlap, though.) Roughly speaking, algebraists are largely concerned about *equalities*, while analysts are largely concerned about *inequalities*.

DEFINITION 1.1. A **total order** is a binary relation $<$ on a set S which satisfies:

1. *transitive*: if $x, y, z \in S$, $x < y$, and $y < z$, then $x < z$.
2. *ordered*: given any $x, y \in S$, exactly one of the following is true: $x < y$, $x = y$, or $y < x$.

The usual order relation on \mathbb{Q} (and its subsets \mathbb{Z} and \mathbb{N}) is a total order. As usual, we write $x > y$ to mean $y < x$, and $x \leq y$ to mean “ $x < y$ or $x = y$ ”.

DEFINITION 1.2. Let $(S, <)$ be a totally ordered set. Let $E \subseteq S$. A **lower bound** for E is an element $\alpha \in S$ with the property that $\alpha \leq x$ for each $x \in E$. A **upper bound** for E is an element $\beta \in S$ with the property that $x \leq \beta$ for each $x \in E$. If E possesses an upper bound, we say E is **bounded above**; if it possesses a lower bound, it is **bounded below**.

For example, the set \mathbb{N} is bounded below in \mathbb{Z} , but it is not bounded above. Any set that has a maximal element is bounded above by its maximum; similarly, any set with a minimal element is bounded below by its minimum.

DEFINITION 1.3. Let $(S, <)$ be a totally ordered set, and let $E \subseteq S$ be bounded above. The **least upper bound** or **supremum** of E , should it exist, is

$$\sup E \equiv \min\{\beta \in S : \beta \text{ is an upper bound of } E\}.$$

Similarly, if F is bounded below, the **greatest lower bound** or **infimum** of F , should it exist, is

$$\inf F \subseteq S \equiv \max\{\alpha \in S : \alpha \text{ is a lower bound of } F\}.$$

To work with the definition (of \sup , say), we rewrite it slightly. A number $\sigma \in S$ is the supremum of E if the following two properties hold:

1. σ is an upper bound of E .
2. Given any $s \in S$ with $s < \sigma$, s is *not* an upper bound of E ; i.e. there exists some $x \in E$ with $s < x \leq \sigma$.

EXAMPLE 1.4. Consider the set $E = \{\frac{1}{n} : n \in \mathbb{N}\} \subset \mathbb{Q}$. This set has a maximal element:
1. So 1 is an upper bound. Moreover, if $s \in \mathbb{Q}$ is < 1 , then s is not an upper bound of E (since

$1 \in E$). Thus, $1 = \sup E$. (This argument shows in general that, if E has a maximal element, then $\max E = \sup E$.)

On the other hand, E has no minimal element. But note that all elements of E are positive, so 0 is a lower bound for E . If s is any rational number > 0 , there is certainly some $n \in \mathbb{N}$ with $0 < \frac{1}{n} < s$ (this is the Archimedean property of the rational field). Hence, no such s is a lower bound for E . This shows that 0 is the *greatest* lower bound: $0 = \inf E$.

EXAMPLE 1.5. It is well known that $\sqrt{2}$ is not rational: in other words, there is no rational number p satisfying $p^2 = 2$. You probably saw this proof in high school. Suppose, for a contradiction, that $p^2 = 2$. Since p is rational, we can write it in *lowest terms* as $p = m/n$ for $m, n \in \mathbb{Z}$. So we have $\frac{m^2}{n^2} = 2$, or $m^2 = 2n^2$. Thus m^2 is even, which means that m is even (since the square of an odd integer is odd). So $m = 2k$ for some $k \in \mathbb{Z}$, meaning $m^2 = 4k^2$, and so $4k^2 = 2n^2$, from which it follows that $n^2 = 2k^2$ is even. As before, this implies that n is even. But then both m and n are divisible by 2, which means they are not relatively prime. This contradicts the assumption that $p = m/n$ is in lowest terms.

A finer analysis of this situation shows that \mathbb{Q} has “holes”. Let

$$A = \{r \in \mathbb{Q} : r > 0, r^2 < 2\}, \quad \text{and} \quad B = \{r \in \mathbb{Q} : r > 0, r^2 > 2\}.$$

The set A is bounded above: if $q \geq \frac{3}{2}$ then $q^2 \geq \frac{9}{4} > 2$, meaning that $q \notin A$; the contrapositive is that if $q \in A$ then $q < \frac{3}{2}$, so $\frac{3}{2}$ is an upper bound for A . In fact, take any positive rational number r ; then $r^2 > 0$ is also rational. By the total order relation, exactly one of the following three statements is true: $r^2 < 2$, $r^2 = 2$, or $r^2 > 2$. In other words, $\mathbb{Q}_{>0} = A \sqcup \{r \in \mathbb{Q} : r > 0, r^2 = 2\} \sqcup B$. We just showed that the middle set is empty, so

$$\mathbb{Q}_{>0} = A \sqcup B.$$

- Every element $b \in B$ is an upper bound for A . Indeed, if $a \in A$ and $b \in B$, then $a^2 < 2 < b^2$ so $0 < b^2 - a^2 = (b-a)(b+a)$, and dividing through by the positive number $b+a$ shows $b-a > 0$ so $a < b$. (This also shows that every element $a \in A$ is a lower bound for B .)
- On the other hand, if $a \in A$, then a is *not* an upper bound for A ; i.e. given $a \in A$, there exists $a' \in A$ with $a < a'$. To see this, we can just take

$$a' = a + \frac{2 - a^2}{2 + a} = \frac{2a + 2}{a + 2}.$$

Since $a \in A$, we know $a^2 < 2$ so $2 - a^2 > 0$, and the denominator $2 + a > 2 > 0$, so $a' > a$. But we also have

$$2 - (a')^2 = \frac{2(a+2)^2 - (2a+2)^2}{(a+2)^2} = \frac{2a^2 + 8a + 8 - 4a^2 - 8a - 4}{(a+2)^2} = \frac{2(2 - a^2)}{(a+2)^2} > 0,$$

showing that $a' \in A$, as claimed.

Thus, B is equal to the set of upper bounds of A in $\mathbb{Q}_{>0}$, and similarly A is equal to the set of lower bounds of B in $\mathbb{Q}_{>0}$.

But then we have the following strange situation. The set A of lower bounds of B has no greatest element: we just showed that, given any $a \in A$, there is an $a' \in A$ with $a' > a$. Hence, B has no greatest lower bound: $\inf B$ does not exist in $\mathbb{Q}_{>0}$. Similarly, $\sup A$ does not exist in $\mathbb{Q}_{>0}$.

Example 1.5 viscerally demonstrates that there is a “hole” in \mathbb{Q} : the fact that $r^2 = 2$ has no solution in \mathbb{Q} forces the ordered set to be disconnected into two pieces, each of which is very incomplete: not only does each fail to possess a max/min, they also fail to possess a sup/inf.

2. Lecture 2: January 7, 2016

We now set the stage for the formal study of the real numbers: it is the (unique) *complete ordered field*. To understand these words, we begin with *fields*.

DEFINITION 1.6. A **field** is a set \mathbb{F} equipped with two binary operations $+, \cdot: \mathbb{F} \times \mathbb{F} \rightarrow \mathbb{F}$, called **addition and multiplication**, satisfying the following properties.

- (1) **Commutativity:** $\forall a, b \in \mathbb{F}, a + b = b + a$ and $a \cdot b = b \cdot a$.
- (2) **Associativity:** $\forall a, b, c \in \mathbb{F}, (a + b) + c = a + (b + c)$ and $(a \cdot b) \cdot c = a \cdot (b \cdot c)$.
- (3) **Identity:** there exists elements $0, 1 \in \mathbb{F}$ s.t. $\forall a \in \mathbb{F}, 0 + a = a = 1 \cdot a$.
- (4) **Inverse:** for any $a \in \mathbb{F}$, there is an element denoted $-a \in \mathbb{F}$ with the property that $a + (-a) = 0$. For any $a \in \mathbb{F} \setminus \{0\}$, there is an element denoted a^{-1} with the property that $a \cdot a^{-1} = 1$.
- (5) **Distributivity:** $\forall a, b, c \in \mathbb{F}, a \cdot (b + c) = (a \cdot b) + (a \cdot c)$.

EXAMPLE 1.7. Here are some examples of fields.

1. The field $\mathbb{Z}_p = \{[0], [1], \dots, [p-1]\}$ for any prime p , where the $+$ and \cdot are the usual ones inherited from the $+$ and \cdot on \mathbb{Z} (namely $[a] + [b] = [a + b]$ and $[a] \cdot [b] = [a \cdot b]$ – you studied this field in Math 109). All finite fields have this form.
2. \mathbb{Q} is a field.
3. \mathbb{Z} is *not* a field: it fails item (4), lacking multiplicative inverses of all elements other than ± 1 .
4. Let $\mathbb{Q}(t)$ denote the set of rational functions of a single variable t with coefficients in \mathbb{Q} :

$$\mathbb{Q}(t) = \left\{ \frac{p(t)}{q(t)} : p(t), q(t) \text{ are polynomials with coefficients in } \mathbb{Q} \text{ and } q(t) \text{ is not identically } 0 \right\}.$$

With the usual addition and multiplication of functions, $\mathbb{Q}(t)$ is a field. For example, $\left(\frac{p(t)}{q(t)}\right)^{-1} = \frac{q(t)}{p(t)}$, which exists so long as $p(t)$ is not identically 0 – i.e. as long as the original rational function $\frac{p(t)}{q(t)}$ is not the 0 function.

Fields are the kinds of number systems that behave the way you've grown up believing numbers behave, as summarized in the following lemma.

LEMMA 1.8. Let \mathbb{F} be a field. The following properties hold.

- (1) **Cancellation:** $\forall a, b, c \in \mathbb{F}$, if $a + b = a + c$ then $b = c$. If $a \neq 0$, if $a \cdot b = a \cdot c$ then $b = c$.
- (2) **Hungry Zero:** $\forall a \in \mathbb{F}, 0 \cdot a = 0$.
- (3) **No Zero Divisors:** $\forall a, b \in \mathbb{F}$, if $a \cdot b = 0$, then either $a = 0$ or $b = 0$.
- (4) **Negatives:** $\forall a, b \in \mathbb{F}, (-a)b = -(ab), -(-a) = a$, and $(-a)(-b) = ab$.

PROOF. We'll just prove (2), leaving the others to the reader. For any $a \in \mathbb{F}$, note that

$$0 \cdot a + a = 0 \cdot a + 1 \cdot a = (0 + 1) \cdot a = 1 \cdot a = a = 0 + a.$$

Hence, by (1) (cancellation), it follows that $0 \cdot a = 0$. □

EXAMPLE 1.9. As in Example 1.7.1, we can consider \mathbb{Z}_n for any positive integer n . This satisfies all of the properties of Definition 1.6 except (4): inverses don't always exist. For example, if n can be factored as $n = km$ for two positive integers $k, m > 1$, then we have two nonzero elements $[k], [m] \in \mathbb{Z}_n$ such that $[k] \cdot [m] = [km] = [n] = [0]$, which contradicts Lemma 1.8(3) – there are zero divisors. So \mathbb{Z}_n is not a field when n is composite.

Now, we combine fields with ordered sets.

DEFINITION 1.10. An **ordered field** is a field \mathbb{F} which is an ordered set $(\mathbb{F}, <)$, where the order relation also satisfies the following two properties:

- (1) $\forall a, b, c \in \mathbb{F}$, if $a < b$ then $a + c < b + c$.
- (2) $\forall a, b \in \mathbb{F}$, if $a > 0$ and $b > 0$, then $a \cdot b > 0$.

From here, all the usual properties mixing the order relation and the field operations follow. For example:

LEMMA 1.11. Let $(\mathbb{F}, <)$ be an ordered field. Then

- (1) $\forall a \in \mathbb{F}$, $a > 0$ iff $-a < 0$.
- (2) $\forall a \in \mathbb{F} \setminus \{0\}$, $a^2 > 0$. In particular, $1 = 1^2 > 0$.
- (3) $\forall a, b \in \mathbb{F}$, if $a > 0$ and $b < 0$, then $a \cdot b < 0$.
- (4) $\forall a \in \mathbb{F}$, if $a > 0$ then $a^{-1} > 0$.

PROOF. For (1), simply add $-a$ to both sides of the inequality. Note, by the properties of $<$, this means \mathbb{F} is the union of three disjoint subsets: the *positive* elements $a > 0$, the *negative* elements $a < 0$, and the zero element $a = 0$; and the operation of multiplication by -1 interchanges the positive and negative elements. So, for (2), we note that our given $a \neq 0$ must be either positive or negative; if $a > 0$ then $a^2 = a \cdot a > 0$ by Definition 1.10(2), while if $a < 0$ then $a^2 = (-a)^2 > 0$ by the same argument. For (3), we then have $a > 0$ and $-b > 0$, so $-(ab) = a \cdot (-b) > 0$, which means that $ab < 0$. Finally, for (4), suppose $a^{-1} < 0$. then by (3) we would have $1 = a \cdot a^{-1} < 0$; but by (2) we know $1 > 0$. This contradiction shows that $a^{-1} > 0$. \square

EXAMPLE 1.12. 1. \mathbb{Q} is an ordered field, with its usual order: $\frac{m_1}{n_1} < \frac{m_2}{n_2}$ iff $m_1 n_2 < m_2 n_1$. In fact, this is the *unique* total order on the set \mathbb{Q} which makes \mathbb{Q} into an ordered field.

2. \mathbb{Z}_p is not an ordered field for any prime p . For suppose it were; then by Lemma 1.11(2) we know that $[1] > [0]$. Then $[2] = [1] + [1] > [1] + [0] = [1]$, and so by transitivity $[2] > [0]$. Continuing this way by induction, we get to $[p-1] > [0]$. But we also have $[0] = [1] + [p-1] > [0] + [p-1] = [p-1]$. This is a contradiction.
3. Let \mathbb{F} be an ordered field. Denote by \mathbb{F}_c the following set of 2×2 matrices over \mathbb{F} :

$$\mathbb{F}_c = \left\{ \begin{bmatrix} a & -b \\ b & a \end{bmatrix} : a, b \in \mathbb{F} \right\}.$$

The determinant of such a matrix is $a^2 + b^2$. In an ordered field, we know that $a^2 > 0$ if $a \neq 0$, and thus we have the usual property that $a^2 + b^2 = 0$ iff $a = b = 0$. It follows that all nonzero matrices in \mathbb{F}_c are invertible: we can easily verify that

$$(a^2 + b^2)^{-1} \begin{bmatrix} a & b \\ -b & a \end{bmatrix} \begin{bmatrix} a & -b \\ b & a \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}.$$

If we define

$$I = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, \quad J = \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix}$$

then $\mathbb{F}_c = \{aI + bJ : a, b \in \mathbb{F}\}$. Note that $J^2 = -I$. It is now an easy exercise to show that \mathbb{F}_c is a field, with $+$, \cdot being given by matrix addition and multiplication, where I is the multiplicative identity and the additive identity is the 2×2 zero matrix. (Note: this is *not generally true* if \mathbb{F} is not an ordered field. For example, in \mathbb{Z}_2 we have $1^2 + (-1)^2 = 0$,

and as a result the matrix with $a = b = 1$ is not invertible in this case.) \mathbb{F}_c is the *complexification* of \mathbb{F} . We will later construct the complex numbers \mathbb{C} as $\mathbb{C} = \mathbb{R}_c$.

3.5 If \mathbb{F} is any ordered field, then \mathbb{F}_c cannot be ordered – there is no order relation that makes \mathbb{F}_c into an ordered field. This is actually what Problem 4 on HW1 asks you to prove.

Item 2 above noted that the finite fields \mathbb{Z}_p are not ordered fields. In fact, ordered fields must be infinite. The next results shows why this is true.

LEMMA 1.13. *Let $(\mathbb{F}, <)$ be an ordered field. Then, for any $n \in \mathbb{Z} \setminus \{0\}$, $n \cdot 1_{\mathbb{F}} \neq 0_{\mathbb{F}}$.*

Here $n \cdot 1_{\mathbb{F}} = 1_{\mathbb{F}} + 1_{\mathbb{F}} + \cdots + 1_{\mathbb{F}}$. Note that this property is *not* automatic for fields: for example, in \mathbb{Z}_p , $p \cdot [1] = [0]$.

PROOF. First, $1 \cdot 1_{\mathbb{F}} = 1_{\mathbb{F}} > 0_{\mathbb{F}}$ by Lemma 1.11(2). Proceeding by induction, suppose we've shown that $n \cdot 1_{\mathbb{F}} \neq 0_{\mathbb{F}}$. Then $(n+1) \cdot 1_{\mathbb{F}} = n \cdot 1_{\mathbb{F}} + 1_{\mathbb{F}} > 0 + 1_{\mathbb{F}} = 1_{\mathbb{F}} > 0_{\mathbb{F}}$. Thus, for every $n > 0$, $n \cdot 1_{\mathbb{F}} > 0_{\mathbb{F}}$, meaning it is $\neq 0$. If, on the other hand, $n < 0$ in \mathbb{Z} , then $n \cdot 1_{\mathbb{F}} = -(-n \cdot 1_{\mathbb{F}}) < 0_{\mathbb{F}}$, so also it is $\neq 0_{\mathbb{F}}$. \square

COROLLARY 1.14. *Let \mathbb{F} be an ordered field. The map $\varphi: \mathbb{Q} \rightarrow \mathbb{F}$ given by $\varphi(\frac{m}{n}) = (m \cdot 1_{\mathbb{F}}) \cdot (n \cdot 1_{\mathbb{F}})^{-1}$ is an injective ordered field homomorphism.*

An **ordered field homomorphism** is a function which preserves the field operations: $\varphi(a+b) = \varphi(a) + \varphi(b)$, $\varphi(a \cdot b) = \varphi(a) \cdot \varphi(b)$, and $\varphi(0) = 0$ and $\varphi(1) = 1$; and preserves the order relation: if $a < b$ then $\varphi(a) < \varphi(b)$. An injective ordered field homomorphism should be thought of as an **embedding**: we realize \mathbb{Q} as a subset of \mathbb{F} , in a way that respects all the ordered field structure.

PROOF. First we must check that φ is well defined: if $\frac{m_1}{n_1} = \frac{m_2}{n_2}$, then $m_1 n_2 = m_2 n_1$. It then follows (by an easy induction) that $(m_1 \cdot 1_{\mathbb{F}}) \cdot (n_2 \cdot 1_{\mathbb{F}}) = (m_2 \cdot 1_{\mathbb{F}}) \cdot (n_1 \cdot 1_{\mathbb{F}})$. Dividing out on both sides then shows that $(m_1 \cdot 1_{\mathbb{F}})(n_1 \cdot 1_{\mathbb{F}})^{-1} = (m_2 \cdot 1_{\mathbb{F}}) \cdot (n_2 \cdot 1_{\mathbb{F}})^{-1}$. Thus, φ is well-defined. It is similar and routine to verify that it is an ordered field homomorphism. Finally, to show it is one-to-one, suppose that $\varphi(q_1) = \varphi(q_2)$ for $q_1, q_2 \in \mathbb{Q}$. Using the homomorphism property, this means $\varphi(q_1 - q_2) = \varphi(q_1) - \varphi(q_2) = 0$. Let $q_1 - q_2 = \frac{m}{n}$; thus, we have $\varphi(\frac{m}{n}) = (m \cdot 1_{\mathbb{F}}) \cdot (n \cdot 1_{\mathbb{F}})^{-1} = 0_{\mathbb{F}}$. But then, multiplying through by the non-zero (by Lemma 1.13) element $n \cdot 1_{\mathbb{F}}$, we have $m \cdot 1_{\mathbb{F}} = 0_{\mathbb{F}}$, and again by Lemma 1.13, it follows that $m = 0$. but this means $q_1 - q_2 = \frac{m}{n} = 0$, so $q_1 = q_2$. Thus, φ is injective. \square

Thus, we will from now on think if \mathbb{Q} as a subset of any ordered field.

In Lecture 1, we saw that \mathbb{Q} “has holes”. In example 1.5, we found two subsets $A, B \subset \mathbb{Q}$ with the property that $B =$ the set of upper bounds of A , $A =$ the set of lower bounds of B , and A has no maximal element, while B has no minimal element. Thus, $\sup A$ and $\inf B$ do not exist. This turns out to be a serious obstacle to doing the kind of analysis we're used to in calculus, so we'd like to fill in these holes. This motivates our next definition.

3. Lecture 3: January 11, 2016

DEFINITION 1.15. An ordered set $(S, <)$ is called **complete** if every nonempty subset $\emptyset \neq E \subseteq S$ that is bounded above possesses a supremum $\sup E \in S$. We also denote this by saying that $(S, <)$ has the **least upper bound property**.

We could also formulate things in terms of \inf , with the greatest lower bound property. Example 1.5 demonstrates how these two are typically related. In fact, they are equivalent.

PROPOSITION 1.16. An ordered set $(S, <)$ has the least upper bound property if and only if, for every nonempty subset $\emptyset \neq F \subseteq S$ that is bounded below, $\inf F \in S$ exists.

PROOF. We will argue the forward implication: the least upper bound property implies the greatest lower bound property. The converse is very similar.

Let $F \neq \emptyset$ be bounded below; then $L \equiv \{\text{lower bounds for } F\}$ is a nonempty subset of S . If $x \in L$ and $y \in F$, then $x \leq y$, which shows that every $y \in F$ is an upper bound for L . Thus, L is bounded above and nonempty; by the least upper bound property of S , $\sigma = \sup L \in S$ exists. By definition of supremum, if $x < \sigma$ then x is not an upper bound for L ; since every element of F is an upper bound for L , this means that such x is not in F . Taking contrapositives, this says that if $z \in F$ then $x \geq \sigma$. So σ is a lower bound for F – i.e. $\sigma \in L$. This shows that $\sigma = \max L$: i.e. σ is the greatest lower bound of F : $\sigma = \inf F$. So $\inf F$ exists, as claimed. \square

Let us now prove some important properties that complete ordered fields possess – properties that are critical for doing all of analysis.

THEOREM 1.17. Let \mathbb{F} be a complete ordered field.

- (1) (**Archimedean**) Let $x, y \in \mathbb{F}$ with $x > 0$. Then there exists $n \in \mathbb{N}$ so that $nx > y$.
- (2) (**Density of \mathbb{Q}**) Let $x, y \in \mathbb{F}$, with $x < y$. Then there exists $r \in \mathbb{Q}$ so that $x < r < y$.

A field with property (1) is called **Archimedean**. It tells us (by setting $x = 1$) that the set \mathbb{N} is *not bounded above* in the field: there is no $y \in \mathbb{F}$ that is \geq every integer. It also tells us (by setting $y = 1$) that there are no “infinitesimals” – that is, no matter how small a positive number x is, there is always a positive integer n such that $0 < \frac{1}{n} < x$. This is an absolutely crucial property for a field to have if we want to talk about limits. And it does *not* hold in every ordered field.

EXAMPLE 1.18. In the field $\mathbb{Q}(t)$ of rational functions with rational coefficients, it is always possible to uniquely express a function $f(t) \in \mathbb{Q}(t)$ in the form $f(t) = \lambda \cdot \frac{p(t)}{q(t)}$ where $\lambda \in \mathbb{Q}$ and $p(t), q(t)$ are *monic* polynomials: their highest order terms have coefficient 1. This allows us to define an order on $\mathbb{Q}(t)$: say $f(t) < g(t)$ iff $g(t) - f(t) = \lambda \frac{p(t)}{q(t)}$ where $p(t), q(t)$ are monic and $\lambda > 0$. (This is the same as insisting that the leading coefficients of the numerator and denominator of $f(t) - g(t)$ have the same sign.) For example $\frac{t^2 - 25t + 7}{t^4 - 10^{23}} > 0$ while $\frac{-t^2 - 25t + 7}{t^4 - 10^{23}} < 0$. Then it is easy but laborious to check that this makes $\mathbb{Q}(t)$ into an ordered field. Note: $t - n = 1 \cdot \frac{t-n}{1} > 0$ for any integer n ; this means that, in the ordered field $\mathbb{Q}(t)$, the element t is *greater than every integer*. I.e. the set $\mathbb{Z} \subset \mathbb{Q}(t)$ actually has an upper bound (e.g. t) in $\mathbb{Q}(t)$. This means $\mathbb{Q}(t)$ is a non-Archimedean field. In particular, by Theorem 1.17, $\mathbb{Q}(t)$ is not a complete ordered field.

PROOF OF THEOREM 1.17. (1) Suppose, for a contradiction, there there is no such n : that is, $nx \leq y$ for every $n \in \mathbb{N}$. Let $E = \{nx : n \in \mathbb{N}\}$. Then our assumption is that y is an upper bound for E , so E is bounded above. It is also non-empty (it contains x , for example). Thus, since \mathbb{F} is complete, it follows that $\alpha = \sup E$ exists. In particular, since $\alpha - x < \alpha$, this means that $\alpha - x$ is

not an upper bound for E , so there is some element $e \in E$ with $\alpha - x < e$. There is some integer $m \in \mathbb{N}$ so that $e = mx$, so we have $\alpha - x < mx$. But then $\alpha < (m + 1)x$, and $(m + 1)x \in E$. This contradicts $\alpha = \sup E$ being an upper bound. This contradiction proves the claim.

(2) Since $y - x > 0$, by (1) there is an $n \in \mathbb{N}$ so that $n(y - x) > 1$. Now, letting $y = \pm nx$ and applying (1) again, we can find two positive integers $m_1, m_2 \in \mathbb{N}$ so that $m_1 > nx$ and $m_2 > -nx$; in other words

$$-m_2 < nx < m_1.$$

This shows that the set $\{k \in \mathbb{Z} : nx < k \leq m_1\}$ is finite: it is contained in the finite set $\{-m_2 + 1, -m_2 + 2, \dots, m_1\}$. So, let $m = \min\{k \in \mathbb{Z} : nx < k\}$. Then since $m - 1 \in \mathbb{Z}$ and $m - 1 < m$, we must have $m - 1 \leq nx$.

Thus, we have two inequalities:

$$n(y - x) > 1, \quad m - 1 \leq nx < m.$$

Combining these gives us

$$nx < m \leq nx + 1 < ny.$$

Dividing through by (the positive) n shows that $x < \frac{m}{n} < y$, so setting $r = \frac{m}{n}$ completes the proof. \square

Here is another extremely important property that holds in ordered fields; this is crucial for doing calculus.

PROPOSITION 1.19. *Let \mathbb{F} be a complete ordered field. For each $n \in \mathbb{N}$, let $a_n, b_n \in \mathbb{F}$ satisfy*

$$a_1 \leq a_2 \leq \dots \leq a_n \leq \dots \leq b_n \leq \dots \leq b_2 \leq b_1.$$

Further, suppose that $b_n - a_n < \frac{1}{n}$. Then $\bigcap_{n \in \mathbb{N}} [a_n, b_n]$ is nonempty, and consists of exactly one point.

This is sometimes called the *nested intervals property*. It is actually equivalent to the least upper bound property. On HW2, you will prove the converse.

PROOF. By construction, b_1 is an upper bound for $\{a_n : n \in \mathbb{N}\}$, which is a nonempty set. Thus, by completeness, $\alpha = \sup a_n$ exists in \mathbb{F} . Since α is an upper bound for $\{a_n\}$, we have $a_n \leq \alpha$ for every n . On the other hand, since $b_m \geq a_n$ for every m, n , b_m is an upper bound for $\{a_n\}$, and since α is the *least* upper bound, it follows that $\alpha \leq b_m$ as well. Thus $\alpha \in [a_n, b_n]$ for every n , and so it is in the intersection.

Now, suppose $\beta \in \bigcap_n [a_n, b_n]$. Then either $\alpha < \beta$, $\alpha > \beta$, or $\alpha = \beta$. Suppose, for the moment, that $\alpha < \beta$. Then we have $a_n \leq \alpha < \beta \leq b_n$ for every n , and since $b_n - a_n < \frac{1}{n}$, it follows that $0 < \beta - \alpha < \frac{1}{n}$ for every n . But this violates the Archimedean property of \mathbb{F} . A similar contradiction arises if we assume $\alpha > \beta$. Thus $\alpha = \beta$, and so α is the unique element of the intersection. \square

Note: in the setup of the lemma, it is similar to see that the intersection consists of $\inf_n b_n$; so $\sup_n a_n = \inf_n b_n$.

4. Lecture 4: January 14, 2014

We have now seen several properties possessed by complete ordered fields. We would hope to find some examples as well. Here comes the big punchline.

THEOREM 1.20. *There exists exactly one complete ordered field. We call this field \mathbb{R} , the Real numbers.*

We will talk about the proof of Theorem 1.20 as we proceed in the course. The textbook relegates an existence proof to the end of Chapter 1, through **Dedekind cuts**. This is an old-fashioned proof, and not very intuitive. We are not going to discuss it presently. Once we have developed a little more technology, we will prove the existence claim of the theorem using Cauchy's construction of \mathbb{R} (through sequences).

We can, however, prove the uniqueness claim. To be precise, here is what uniqueness means in this case: suppose \mathbb{F} and \mathbb{G} are two complete ordered fields. Then there exists an ordered field isomorphism $\varphi: \mathbb{F} \rightarrow \mathbb{G}$. That means φ is an ordered field homomorphism that is also a bijection. So, from the point of view of ordered fields, \mathbb{F} and \mathbb{G} are indistinguishable.

The first question is: given two complete ordered fields \mathbb{F} and \mathbb{G} , how do we define $\varphi: \mathbb{F} \rightarrow \mathbb{G}$? By Corollary 1.14, \mathbb{Q} embeds in each of \mathbb{F} and \mathbb{G} via $\mathbb{Q} \cdot 1_{\mathbb{F}}$ and $\mathbb{Q} \cdot 1_{\mathbb{G}}$. So we can define φ as a *partial* function by its action on \mathbb{Q} :

$$\varphi(r1_{\mathbb{F}}) = r1_{\mathbb{G}}, \quad r \in \mathbb{Q}.$$

The question is: how should we define φ on elements of \mathbb{F} that are not necessarily in $\mathbb{Q} \cdot 1_{\mathbb{F}}$? Well, let $x \in \mathbb{F} \setminus \mathbb{Q}$. By Theorem 1.17(2), there are rationals $a_n, b_n \in \mathbb{Q}$ such that

$$x - \frac{1}{2n}1_{\mathbb{F}} < a_n1_{\mathbb{F}} < x < b_n1_{\mathbb{F}} < x + \frac{1}{2n}1_{\mathbb{F}}.$$

In particular, $b_n - a_n < \frac{1}{n}$. We should do this carefully and also make sure that $a_1 \leq a_2 \leq \dots \leq b_2 \leq b_1$ — this can be achieved by choosing the a_n and b_n successively, increasing the a_n or decreasing the b_n each step as needed. It follows from Proposition 1.19 that $\bigcap_n [a_n1_{\mathbb{G}}, b_n1_{\mathbb{G}}]$ contains exactly one point, $\alpha = \sup_n(a_n1_{\mathbb{G}}) = \inf_n(b_n1_{\mathbb{G}})$. So we define

$$\varphi(x) = \alpha.$$

Note: if $x \in \mathbb{Q}$, then $x \cdot 1_{\mathbb{G}}$ is the unique element in the intersection, meaning that we can take the above nested intervals definition as the formula for φ on all of \mathbb{F} , not just the irrational elements. This will be our starting point.

THEOREM 1.21. *If \mathbb{F} and \mathbb{G} are two complete ordered fields, then there exists an ordered field isomorphism $\varphi: \mathbb{F} \rightarrow \mathbb{G}$.*

PROOF. Following our outline from above, we define φ as follows. To begin, using the denseness of \mathbb{Q} in \mathbb{F} , select $a_1, b_1 \in \mathbb{Q}$ so that

$$x - \frac{1}{2}1_{\mathbb{F}} < a_11_{\mathbb{F}} < x < b_11_{\mathbb{F}} < x + \frac{1}{2}1_{\mathbb{F}}.$$

Now proceed inductively: once we've constructed a_1, \dots, a_{n-1} and b_1, \dots, b_{n-1} , choose a_n and b_n so that

$$\max \left\{ x - \frac{1}{2n}1_{\mathbb{F}}, a_{n-1} \right\} < a_n1_{\mathbb{F}} < x < b_n1_{\mathbb{F}} < \min \left\{ x + \frac{1}{2n}1_{\mathbb{F}}, b_{n-1} \right\}. \quad (1.1)$$

Then we have $a_1 < a_2 < \dots < a_n < \dots < b_n < \dots < b_2 < b_1$, and also

$$b_n - a_n < \left(x + \frac{1}{2n}1_{\mathbb{F}}\right) - \left(x - \frac{1}{2n}1_{\mathbb{F}}\right) = \frac{1}{n}.$$

So by the nested intervals property Proposition 1.19 applied in the field \mathbb{G} , we have

$$\bigcap_{n \in \mathbb{N}} [a_n 1_{\mathbb{G}}, b_n 1_{\mathbb{G}}] = \{\alpha\}$$

where $\alpha = \sup_n a_n 1_{\mathbb{G}} = \inf_n b_n 1_{\mathbb{G}}$. We thus define $\varphi(x) = \alpha$.

Now we must verify that:

- **φ is well-defined:** if a'_n, b'_n are some other rational elements satisfying (1.1) then $\sup_n a_n 1_{\mathbb{G}} = \sup_n a'_n 1_{\mathbb{G}}$. In fact, this follows because we also then have the mixed inequalities

$$x - \frac{1}{2n}1_{\mathbb{F}} < a'_n 1_{\mathbb{F}} < x < b_n 1_{\mathbb{F}} < x + \frac{1}{2n}1_{\mathbb{F}}$$

and, as above, we have $\sup_n a'_n 1_{\mathbb{G}} = \inf_n b_n 1_{\mathbb{G}} = \sup_n a_n 1_{\mathbb{G}}$.

- **φ is an ordered field homomorphism.** This is laborious. Let's check one of the field homomorphism properties: preservation of addition. Let $x, y \in \mathbb{F}$, and let $a_n < x < b_n$ and $c_n < y < d_n$ where $b_n - a_n < \frac{1}{2n} < \frac{1}{n}$ and $d_n - c_n < \frac{1}{2n} < \frac{1}{n}$. Then $\varphi(x) = \sup_n a_n$ and $\varphi(y) = \sup_n c_n$. Now, on the other hand, we have

$$a_n + c_n < x + y < b_n + d_n, \quad \text{and} \quad (b_n + d_n) - (a_n + c_n) = (b_n - a_n) + (d_n - c_n) < \frac{1}{2n} + \frac{1}{2n} = \frac{1}{n}.$$

It follows that $\varphi(x + y) = \sup(a_n + c_n)$. So, to see that $\varphi(x + y) = \varphi(x) + \varphi(y)$, it suffices to show that

$$\text{if } a_n \uparrow \ \& \ c_n \uparrow \quad \text{then} \quad \sup_n (a_n + c_n) = \sup_n a_n + \sup_n c_n.$$

This is also on HW2. The other ordered field homomorphism properties are verified similarly.

- **φ is a bijection.** First, suppose that $x \neq y \in \mathbb{F}$. Then either $x < y$ or $x > y$; wlog $x < y$. Since φ is an ordered field homomorphism, it follows that $\varphi(x) < \varphi(y)$. In particular, $\varphi(x) \neq \varphi(y)$. A similar argument in the case $x > y$ shows that φ is one-to-one.

Now, fix $y \in \mathbb{G}$. For each n , choose $a_n, b_n \in \mathbb{Q}$ nested so that $b_n - a_n < \frac{1}{n}$ and $a_n 1_{\mathbb{G}} < y < b_n 1_{\mathbb{G}}$. Mirroring the above arguments, we know that $a = \sup_n a_n 1_{\mathbb{F}} \in \bigcap_n [a_n 1_{\mathbb{F}}, b_n 1_{\mathbb{F}}]$. Since $a_n 1_{\mathbb{F}} < a < b_n 1_{\mathbb{F}}$, we have $a_n 1_{\mathbb{G}} = \varphi(a_n 1_{\mathbb{F}}) < \varphi(a) < \varphi(b_n 1_{\mathbb{F}}) = b_n 1_{\mathbb{G}}$. Thus $\varphi(a) \in \bigcap_n [a_n 1_{\mathbb{G}}, b_n 1_{\mathbb{G}}]$, and this intersection consists of the singleton element y , by Proposition 1.19. Hence, $\varphi(a) = y$, and so φ is onto. □

So, we see that there can be *only one* complete ordered field. (They're like Highlanders.) A priori, that doesn't preclude the possibility that there aren't any at all. To prove that \mathbb{R} exists, we need to first start talking about convergence properties of sequences. That will be our next task.

Before proceeding, let's return to our motivation for studying sup and inf and introducing completeness: we wanted to fill the "hole" in \mathbb{Q} where $\sqrt{2}$ should be. To see that we've filled at least that hole, the next result shows that \mathbb{R} (the complete ordered field) contains square roots, and in fact n th roots, of all positive numbers. First, let's state some standard results on "absolute value".

LEMMA 1.22. *Let \mathbb{F} be an ordered field. For $x \in \mathbb{F}$, define (as usual)*

$$|x| = \begin{cases} x, & \text{if } x \geq 0 \\ -x, & \text{if } x < 0 \end{cases}.$$

Then we have the following properties.

- (1) *For all $x \in \mathbb{F}$, $|x| \geq 0$, and $|x| = 0$ iff $x = 0$.*
- (2) *For all $x, y \in \mathbb{F}$, $|x + y| \leq |x| + |y|$.*
- (3) *For all $x, y \in \mathbb{F}$, $|xy| = |x||y|$.*

All of these properties are straightforward but annoying to prove in cases. We will use the absolute value frequently in all that follows.

THEOREM 1.23. *Let $n \in \mathbb{N}$, $n \geq 1$. For any $x \in \mathbb{R}$, $x > 0$, there is a unique $y \in \mathbb{R}$, $y > 0$, so that $y^n = x$. We denote it by $y = x^{1/n}$.*

PROOF OF THEOREM 1.23. First, for uniqueness: let $y_1 \neq y_2$ be two positive real numbers, wlog $y_1 < y_2$. Then $y_1^2 = y_1y_1 < y_1y_2 < y_2y_2 = y_2^2$; continuing by induction, we see that $y_1^n < y_2^n$. That is: the function $y \mapsto y^n$ is strictly increasing. In particular, it is one-to-one. It follows that there can be at most one y with $y^n = x$.

Now for existence. Let $E = \{y \in \mathbb{R} : y > 0, y^n < x\}$.

- $E \neq \emptyset$: note that $t = \frac{x}{x+1} \in (0, 1)$. This means that $0 < t^n < t$, and so since $\frac{x}{x+1} < x$, we have $0 < t^n < x$, meaning that $t \in E$.
- E is bounded above: let $s = 1 + x$. Then $s > 1$, and so $s^n > s > x$. Thus, if $y \in E$, then $y^n < x < s^n$, and so $0 < s^n - y^n = (s - y)(s^{n-1} + s^{n-2}y + \cdots + y^{n-1})$. The sum of terms is strictly positive, so we can divide out and find that $s - y > 0$. Thus s is an upper bound for E .

Hence, by completeness of \mathbb{R} , $\alpha = \sup E$ exists. Since α is the least upper bound, it follows that, for each k , there is an element $y_k \in E$ such that $y_k > \alpha - \frac{1}{k}$. Since $y_k^n < x$, we therefore have

$$\left(\alpha - \frac{1}{k}\right)^n < y_k^n < x, \quad \text{for all } k \in \mathbb{N}.$$

But we can expand

$$\left(\alpha - \frac{1}{k}\right)^n = \sum_{j=0}^n \binom{n}{j} \alpha^{n-j} \left(-\frac{1}{k}\right)^{-j} = \alpha^n - \frac{1}{k} \sum_{j=1}^n \binom{n}{j} \alpha^{n-j} \left(-\frac{1}{k}\right)^{j-1}.$$

Thus, we have

$$\alpha^n < x + \frac{1}{k} \sum_{j=1}^n \binom{n}{j} \alpha^{n-j} \left(-\frac{1}{k}\right)^{j-1}$$

and so, applying the triangle inequality – Lemma 1.22(2) – repeatedly, we have

$$\alpha^n < x + \frac{1}{k} \left| \sum_{j=1}^n \binom{n}{j} \alpha^{n-j} \left(-\frac{1}{k}\right)^{j-1} \right| \leq x + \frac{1}{k} \cdot \sum_{j=1}^n \binom{n}{j} \alpha^{n-k} \left(\frac{1}{k}\right)^{j-1}.$$

Note that n is fixed, and $\frac{1}{k} \leq 1$, so for $k \geq 1$ we have $\left(\frac{1}{k}\right)^{j-1} \leq 1$. Let $M = \sum_{k=1}^n \binom{n}{j} \alpha^{n-k}$; then we have

$$\forall k \in \mathbb{N} \quad \alpha^n < x + \frac{M}{k}; \quad \text{i.e.} \quad \alpha^n - x < \frac{M}{k}.$$

By the Archimedean property, it follows that $\alpha^n - x \leq 0$; thus, we have shown that $\alpha^n \leq x$.

On the other hand, let $y \in E$. Then for any $k \in \mathbb{N}$ we have, by similar calculations,

$$\left(y + \frac{1}{k}\right)^n = y^n + \frac{1}{k} \sum_{j=1}^n \binom{n}{j} y^{n-j} \left(\frac{1}{k}\right)^{j-1} \leq y^n + \frac{1}{k} \cdot \sum_{j=1}^n \binom{n}{j} y^{n-j}.$$

Since $y \in E$, we know $y^n < x$, so $\epsilon = x - y^n > 0$. Let $L = \sum_{j=1}^n \binom{n}{j} y^{n-j}$, which is a positive constant; by the Archimedean property, there is some $k \in \mathbb{N}$ so that $\frac{1}{k} \cdot L < \epsilon$. Thus, for such k ,

$$\left(y + \frac{1}{k}\right)^n \leq y^n + \frac{L}{k} < y^n + \epsilon = x.$$

That is: $y + \frac{1}{k} \in E$. But $y + \frac{1}{k} > y$. That is, for any $y \in E$, there is $y' > y$ with $y \in E$. So E has no maximal element. This shows that $\alpha \notin E$, and hence $\alpha^n \geq x$.

In conclusion: we've shown that $\alpha^n \leq x$ and $x \leq \alpha^n$. It follows that $\alpha^n = x$. \square

On Homework 2, you will flesh out extending this argument to defining x^r for $x > 0$ in \mathbb{R} and $r \in \mathbb{Q}$, and then extending this further to define x^y for $x > 0$ and $y \in \mathbb{R}$. One can use similar arguments to define $\log_b(x)$ for $x, b > 0$. We will wait a little while until we have a firm grounding in sequences and limits before rigorously developing the calculus of these well-known functions.

CHAPTER 2

Sequences and Limits

1. Lecture 5: January 19, 2016

DEFINITION 2.1. Let X be a set. A **sequence** in X is a function $a: \mathbb{N} \rightarrow X$. Instead of the usual notation $a(n)$ for the value of the function at $n \in \mathbb{N}$, we usually use the notation $a_n = a(n)$; accordingly, we often refer to the function as $(a_n)_{n \in \mathbb{N}}$ or $\{a_n\}_{n \in \mathbb{N}}$, or (when being sloppy) simply (a_n) or $\{a_n\}$.

In ordered fields, we can talk about limits of sequences. The following definition took half a century to finalize; its invention (by Weierstraß) is one of the greatest achievements of analysis.

DEFINITION 2.2. Let \mathbb{F} be an ordered field, and let (a_n) be a sequence in \mathbb{F} . Let $a \in \mathbb{F}$. Say that a_n **converges to** a , written $a_n \rightarrow a$ or $\lim_{n \rightarrow \infty} a_n = a$, if the following holds true:

$$\forall \epsilon > 0 \exists N \in \mathbb{N} \text{ s.t. } \forall n \geq N |a_n - a| < \epsilon.$$

Let's decode the three-quantifier sentence here. What this says is, no matter how small a tolerance $\epsilon > 0$ you want, there is some time N after which all the terms a_n (for $n \geq N$) are within ϵ of a . Some convenient language for this is:

Given any $\epsilon > 0$, we have $|a_n - a| < \epsilon$ for **almost all** n .

Here we colloquially say that a set $S \subseteq \mathbb{N}$ contains **almost all** positive integers if the complement $\mathbb{N} \setminus S$ is finite. This is equivalent to saying that, after some N , all $n \geq N$ are in S . So, the limit definition is that, for any positive tolerance, no matter how small, almost all of the terms are within that tolerance of the limit.

If (a_n) is a sequence and there exists a so that $a_n \rightarrow a$, we say that (a_n) **converges**; if there is no such a , we say that (a_n) **diverges**. Here are some examples.

EXAMPLE 2.3. Consider each of the following sequences in an Archimedean field.

- (1) $a_n = 1$ converges to 1. More generally, if (a_n) is equal to a constant a for almost all n , then $a_n \rightarrow a$.
- (2) $a_n = \frac{1}{n}$ converges to 0.
- (3) $a_n = n + \frac{1}{n}$ diverges.
- (4) $a_n = (-1)^n$ diverges.
- (5) $a_n = 1 + \frac{1}{n}(-1)^n$ converges to 1.
- (6) $a_n = \frac{4n+1}{7n-4}$ (defined for $n \geq 1$) converges to $\frac{4}{7}$.

In all these examples, we proved convergence (when the sequences converged) to a given value. However, a priori, it is not clear whether it might also have been possible to prove convergence to a *different* value as well. This is not the case: limits are unique.

LEMMA 2.4. Let \mathbb{F} be an ordered field, and let (a_n) be a sequence in \mathbb{F} . Suppose $a, b \in \mathbb{F}$ and $a_n \rightarrow a$ and $a_n \rightarrow b$. Then $a = b$.

PROOF. Fix $\epsilon > 0$. We know that there is N_1 so that $|a_n - a| < \frac{\epsilon}{2}$ for all $n > N_1$, and there is N_2 so that $|a_n - b| < \frac{\epsilon}{2}$ for all $n > N_2$. Thus, for any $n > \max\{N_1, N_2\}$, we have

$$|a - b| = |a - a_n + a_n - b| \leq |a - a_n| + |a_n - b| < \frac{\epsilon}{2} + \frac{\epsilon}{2} = \epsilon.$$

Now, suppose that $a \neq b$. Thus $a - b \neq 0$, which means that $|a - b| > 0$. So we can take $\epsilon = |a - b|$ above, and we find that $|a - b| < |a - b|$ – a contradiction. Hence, it must be true that $a = b$. \square

REMARK 2.5. Note, in an *Archimedean* field, we are free to restrict $\epsilon = \frac{1}{k}$ for some $k \in \mathbb{N}$; that is, an equivalent statement of $a_n \rightarrow a$ is

Given any $k \in \mathbb{N}$, we have $|a_n - a| < \frac{1}{k}$ for almost all n .

In non-Archimedean fields, this does not suffice. For example, in the field $\mathbb{Q}(t)$, to show $a_n(t) \rightarrow a(t)$ it does not suffice to show that, for any $k \in \mathbb{N}$, $|a_n(t) - a(t)| < \frac{1}{k}$ for all sufficiently large n . Indeed, what if $a_n(t) - a(t) = \frac{1}{t}$? This does not go to 0, but it is $< \frac{1}{k}$ for all $k \in \mathbb{N}$. Similarly, the sequence $a_n = \frac{1}{n}$ diverges in a non-Archimedean field.

2. Lecture 6: January 21, 2016

PROPOSITION 2.6. *Let \mathbb{F} be a complete ordered field. Let (a_n) be a sequence in \mathbb{F} , and suppose $a_n \uparrow$ (i.e. $a_n \leq a_{n+1}$ for all n) and bounded above. Let $\alpha = \sup\{a_n\}$. Then $a_n \rightarrow \alpha$. Similarly, if $b_n \downarrow$ and bounded below, then $\beta = \inf\{b_n\}$ exists and $b_n \rightarrow \beta$.*

PROOF. Since \mathbb{F} is a complete field, $\alpha = \sup\{a_n\}$ exists in \mathbb{F} . Let $\epsilon > 0$. Then $\alpha - \epsilon < \alpha$, and so by definition there exists some element $a_N \in \{a_n\}$ so that $\alpha - \epsilon < a_N \leq \alpha$. Now, suppose $n \geq N$; then $a_n \leq \alpha$ of course, but also since $a_n \uparrow$ we have $a_n \geq a_N > \alpha - \epsilon$. Thus, we have shown that $|a_n - \alpha| = \alpha - a_n < \epsilon$ for all $n \geq N$, which is to say that $a_n \rightarrow \alpha$.

The decreasing case is similar; alternatively, one can look at $a_n = -b_n$, which is increasing and bounded above; then we have by the first part that $-b_n = a_n \rightarrow \alpha$ where $\alpha = \sup\{-b_n\} = -\inf\{a_n\} = -\beta$. It follows that $b_n \rightarrow -\beta$, using the limit theorems below. □

In the proposition, we needed (a_n) to be bounded (above or below); indeed, the sequence $a_n = n$ is increasing, but not convergent. This is generally true: for any sequence to be convergent, it must be bounded (above and below). A sequence that is either increasing or decreasing is called **monotone**. So the proposition shows that monotone sequences either converge, or grow (in absolute value) without bound.

This gives us a new perspective on the motivating example that began our discussion of sup and inf. Consider, again, the sets $A = \{r \in \mathbb{Q} : r > 0, r^2 < 2\}$ and $B = \{r \in \mathbb{Q} : r > 0, r^2 > 2\}$. We saw that the set of positive rationals is equal to $A \sqcup B$, and therefore $\sup A$ and $\inf B$ do not exist in \mathbb{Q} . Note that the sequence 1, 1.4, 1.41, 1.414, 1.4142, 1.42431, ... is in the set A . We recognize the terms as the decimal approximations to $\sqrt{2}$. This sequence looks like it's *going somewhere*; but in fact the only place it can go is stuck in between A and B , which is not in \mathbb{Q} . The question is: *why* does it look like it's going somewhere?

DEFINITION 2.7. *A sequence (a_n) in an ordered set is called **Cauchy**, or is said to be a **Cauchy sequence**, if*

$$\forall \epsilon > 0 \exists N \in \mathbb{N} \text{ s.t. } \forall n, m \geq N |a_n - a_m| < \epsilon.$$

That is: a sequence is Cauchy if its terms get and stay close to *each other*. That is: for any given tolerance $\epsilon > 0$, there is some time N after which all the terms are within distance ϵ of a_N . This notion is very close to convergence. Indeed:

LEMMA 2.8. *Any convergent sequence is Cauchy.*

PROOF. Let (a_n) be a convergent sequence, with limit a . Fix $\epsilon > 0$, and choose N large enough so that $|a_n - a| < \frac{\epsilon}{2}$ for $n > N$. Then for any $n, m > N$,

$$|a_n - a_m| = |a_n - a + a - a_m| \leq |a_n - a| + |a_m - a| < \frac{\epsilon}{2} + \frac{\epsilon}{2} = \epsilon.$$

Hence, (a_n) is Cauchy. □

But the converse need not be true.

EXAMPLE 2.9. In \mathbb{Q} , the sequence 1, 1.4, 1.41, 1.414, 1.4142, 1.42431, ... is Cauchy. Indeed, by the definition of decimal expansion, if a_n is the n -decimal expansion of a number, then a_{n+1} and a_n agree on the first n digits. This means exactly that $|a_m - a_n| < \frac{1}{10^n}$ for any $m > n$. So, fix $\epsilon > 0$. We can certainly find N so that $\frac{1}{10^N} < \epsilon$ (since, for example, $\frac{1}{10^N} < \frac{1}{N}$). Thus, for $n, m > N$, we have $|a_n - a_m| < \frac{1}{10^{\min\{m,n\}}} < \frac{1}{10^N} < \epsilon$.

Here are some more important facts about Cauchy sequences. Note that, by Lemma 2.8, any fact about Cauchy sequences is also a fact about convergent sequences.

PROPOSITION 2.10. *Let (a_n) be a Cauchy sequences. Then (a_n) is bounded: there is a constant $M > 0$ so that $|a_n| \leq M$ for all n .*

PROOF. Taking $\epsilon = 1$, it follows from the definition of Cauchy that there is some $N \in \mathbb{N}$ so that $|a_n - a_m| < 1$ for all $n, m > N$. In particular, this shows that $|a_n - a_{N+1}| < 1$ for all $n > N$, which is to say that $a_{N+1} - 1 < a_n < a_{N+1} + 1$. Hence $|a_n| < \max\{|a_{N+1} - 1|, |a_{N+1} + 1|\}$ for $n > N$. So, define $M = \max\{|a_1|, \dots, |a_N|, |a_{N+1} - 1|, |a_{N+1} + 1|\}$. If $n \leq N$, then $|a_n| \leq M$ since $|a_n|$ appears in this list we maximize over; if $n > N$ then, as just shown, $|a_n| < \max\{|a_{N+1} - 1|, |a_{N+1} + 1|\} \leq M$. The result follows. \square

Another useful concept when working with sequences is **subsequences**.

DEFINITION 2.11. *Let $\{n_k : k \in \mathbb{N}\}$ be a set of positive integers with the property that $n_k < n_{k+1}$ for all k ; that is n_k is an increasing sequence in \mathbb{N} . Let (a_n) be a sequence. The function $k \mapsto a_{n_k}$ is called a **subsequence** of (a_n) , usually denoted (a_{n_k}) .*

EXAMPLE 2.12. (a) Let $a_n = \frac{1}{n}$. Then $a_{2n} = \frac{1}{2n}$ and $a_{2^n} = \frac{1}{2^n}$ are subsequences. However

$$b_n = \begin{cases} a_n & \text{if } n \text{ is odd} \\ a_{n/2} & \text{if } n \text{ is even} \end{cases}$$

is **not** a subsequence of (a_n) . Indeed, $b_k = a_{n_k}$ where $(n_k)_{k=1}^\infty = (1, 1, 3, 2, 5, 3, 7, 4, 9, 5, \dots)$, and this is not an increasing sequence of integers.

(b) Let $a_n = (-1)^n$. Then $a_{2n} = 1$ and $a_{2n+1} = -1$ are subsequences.

Here is an extremely useful fact about the indices of subsequences: if (n_k) is an increasing sequence in \mathbb{N} , then $n_k \geq k$ for every k . (This follows by a simple induction.)

PROPOSITION 2.13. *Let (a_n) be a sequence in an ordered set, and (a_{n_k}) a subsequence.*

- (1) *If (a_n) is Cauchy, then (a_{n_k}) is Cauchy.*
- (2) *If (a_n) is convergent with limit a , then (a_{n_k}) is convergent with limit a .*
- (3) *If (a_n) is Cauchy, and (a_{n_k}) is convergent with limit a , then (a_n) is convergent with limit a .*

PROOF. For (1): fix $\epsilon > 0$ and let $N \in \mathbb{N}$ be chosen so that $|a_n - a_m| < \epsilon$ for $n, m > N$. Then whenever $k, \ell > N$, we have $n_k \geq k > N$ and $n_\ell \geq \ell > N$, so by definition $|a_{n_k} - a_{n_\ell}| < \epsilon$. Thus (a_{n_k}) is Cauchy. The proof of (2) is very similar. Item (3) is on HW3. \square

Before proceeding with the theory of Cauchy sequences, here are some useful facts about convergent sequences sequences.

THEOREM 2.14. *Let (a_n) and (b_n) be convergent sequences in an ordered field \mathbb{F} .*

- (1) *If $a_n \leq b_n$ for all sufficiently large n , then $\lim_n a_n \leq \lim_n b_n$.*
- (2) *(Squeeze Theorem) Suppose also that $\lim_n a_n = \lim_n b_n$. If (c_n) is another sequence, and $a_n \leq c_n \leq b_n$ for all sufficiently large n , then (c_n) is convergent, and $\lim_n c_n = \lim_n a_n = \lim_n b_n$.*

PROOF. Let $a = \lim_n a_n$ and $b = \lim_n b_n$. For (1), fix $\epsilon > 0$. There is $N_a \in \mathbb{N}$ so that $|a_n - a| < \frac{\epsilon}{2}$ for $n > N_a$, and there is $N_b \in \mathbb{N}$ so that $|b_n - b| < \frac{\epsilon}{2}$ for $n > N_b$. Thus, letting $N = \max\{N_a, N_b\}$, we have $a_n - a > -\frac{\epsilon}{2}$ and $b_n - b < \frac{\epsilon}{2}$ for $n > N$. But then

$$a_n - b_n > a - \frac{\epsilon}{2} - b - \frac{\epsilon}{2} = a - b - \epsilon.$$

Since $a_n \leq b_n$ for all large n , we therefore have $0 \geq a_n - b_n > a - b - \epsilon$ for such n , and therefore $a - b - \epsilon < 0$. This is true for any $\epsilon > 0$, and therefore $a - b \leq 0$, as claimed.

For (2), we have $a = b$. Choosing N_a, N_b , and N as above, we have $-\frac{\epsilon}{2} < a_n - a \leq c_n - a \leq b_n - a < \frac{\epsilon}{2}$ for all $n \geq N$. That is: $|c_n - a| < \frac{\epsilon}{2} < \epsilon$ for all $n \geq N$. This shows $c_n \rightarrow a$, as claimed. \square

Cauchy sequences give us a way of talking about completeness that is not so wrapped up in the order properties. As discussed in Example 2.9 last lecture, the “hole” in \mathbb{Q} where $\sqrt{2}$ should be is the limit of a sequence in \mathbb{Q} which is Cauchy, but does not converge in \mathbb{Q} . Instead of filling in the holes by demanding bounded nonempty sets have suprema, we could instead demand that Cauchy sequences have limits.

DEFINITION 2.15. Let S be an ordered set. Call S **Cauchy complete** if every Cauchy sequence in S actually converges in S .

\mathbb{Q} is not Cauchy complete. But, as we will see, \mathbb{R} is. In fact, Cauchy completeness is *equivalent* to the least upper bound property in any Archimedean field. We can prove half of this assertion now.

3. Lecture 7: January 26, 2016

THEOREM 2.16. *Let \mathbb{F} be an Archimedean field. If \mathbb{F} is Cauchy complete, then \mathbb{F} has the nested intervals property and hence is complete in the sense of Definition 1.15.*

PROOF. That the nested intervals property implies the least upper bound property is the content HW2 Exercise 3; so it suffices to verify that \mathbb{F} has the nested intervals property. Let (a_n) and (b_n) be sequences in \mathbb{F} with $a_n \uparrow$, $b_n \downarrow$, $a_n \leq b_n$, and $b_n - a_n < \frac{1}{n}$. Fix $\epsilon > 0$, and let $N \in \mathbb{N}$ be large enough that $\frac{1}{N} < \epsilon$ (here is where the Archimedean property is needed). Thus, for $n \geq N$, we have $b_n - a_n < \frac{1}{n} \leq \frac{1}{N} < \epsilon$. Then for $m, n > N$, wlog $m \geq n$, we have

$$a_n \leq a_m \leq b_n$$

and so it follows that $|a_n - a_m| = a_m - a_n \leq b_n - a_n < \epsilon$. Thus (a_n) is a Cauchy sequence. By the Cauchy completeness assumption on \mathbb{F} , we conclude that $a = \lim_n a_n$ exists in \mathbb{F} .

Now, fix n_0 , and note that since $a_n \geq a_{n_0}$ for $n \geq n_0$, Theorem 2.14(1) shows that $a = \lim_n a_n \geq a_{n_0}$ (thinking of a_{n_0} as the limit of the constant sequences $(a_{n_0}, a_{n_0}, \dots)$). Similarly, since $a_n \leq b_{n_0}$ for all n , it follows that $a \leq b_{n_0}$. Thus $a \in \bigcap_n [a_n, b_n]$, proving this intersection is nonempty. As usual, it follows that the intersection consists only of $\{a\}$. Indeed, if $x, y \in \bigcap_n [a_n, b_n]$, without loss of generality label them so that $x \leq y$. Thus $a_n \leq x \leq y \leq b_n$ for every n . For given $\epsilon > 0$, choose n so that $b_n - a_n < \epsilon$; then $y - x < \epsilon$. So $0 \leq y - x < \epsilon$ for all $\epsilon > 0$; it follows that $x = y$. This concludes the proof of the nested intervals property for S . \square

REMARK 2.17. The use of the Archimedean property is very subtle here. It is tempting to think that we can do without it. This is true if we replace the nested intervals property by a slightly weaker version: say an ordered S satisfies the *weak* nested intervals property if, given $a_n \uparrow$, $b_n \downarrow$, $a_n \leq b_n$, and $b_n - a_n \rightarrow 0$, then $\bigcap_n [a_n, b_n]$ contains exactly one point. (This is weaker than the nested intervals property, because the assumption is stronger: we're assuming $b_n - a_n \rightarrow 0$ here, while in the usual nested intervals property we assume that $b_n - a_n < \frac{1}{n}$, which does not imply $b_n - a_n \rightarrow 0$ in the non-Archimedean setting.) The trouble is: this weak nested intervals property *does not imply* the least upper bound property in the absence of the Archimedean property. In fact, there *do exist* non-Archimedean fields (which therefore do not have the least upper bound property), but *are* Cauchy complete. (We may explore this a little later.) This is a prime example of how counterintuitive analysis can be without the Archimedean property. Soon enough, we will once-and-for-all demand that it holds true (in the Real numbers), and dispense with these weird pathologies.

We would like to show the converse is true: that the least upper bound property implies Cauchy completeness. (This turns out to be true in any ordered set: after all, the least upper bound property implies the Archimedean property in an ordered field.) Then we could characterize the real numbers as the unique Archimedean field that is Cauchy complete. To do this, we need to dig a little deeper into the connection between limits and suprema / infima.

DEFINITION 2.18. *Let S be an ordered set with the least upper bound property. Let (a_n) be a bounded sequence in S . Define two new sequences from (a_n) :*

$$\bar{a}_k = \sup\{a_n : n \geq k\}, \quad \underline{a}_k = \inf\{a_n : n \geq k\}.$$

Since $\{a_n\}$ is bounded above (and nonempty), by the least upper bound property \bar{a}_k exists for each k . Similarly, by Proposition 1.16, \underline{a}_k exists for each k .

Note that $\{a_n : n \geq k+1\} \subseteq \{a_n : n \geq k\}$. Thus \bar{a}_k is an upper bound for $\{a_n : n \geq k+1\}$. It follows that \bar{a}_k is \geq the least upper bound of $\{a_n : n \geq k+1\}$, which is defined to be \bar{a}_{k+1} .

This means that $\bar{a}_k \geq \bar{a}_{k+1}$: the sequence \bar{a}_k is monotone decreasing. Similarly, the sequence \underline{a}_k is monotone increasing.

By assumption, $\{a_n\}$ is bounded. Thus there is a lower bound $a_n \geq L$ for all n . Since $\bar{a}_1 \geq \bar{a}_k \geq a_k \geq L$ for all k , the sequence \bar{a}_k is also bounded. Similarly, the sequence \underline{a}_k is bounded.

Thus, \bar{a}_k is a decreasing, bounded-below sequence. By Proposition 2.6, $\lim_{k \rightarrow \infty} \bar{a}_k$ exists, and is equal to $\inf\{\bar{a}_k\}$. Similarly, $\lim_{k \rightarrow \infty} \underline{a}_k$ exists, and is equal to $\sup\{\underline{a}_k\}$. We define

$$\begin{aligned} \limsup_{n \rightarrow \infty} a_n &= \lim_{n \rightarrow \infty} \bar{a}_n = \lim_{k \rightarrow \infty} \sup\{a_n : n \geq k\} = \inf_{k \in \mathbb{N}} \sup_{n \geq k} a_n \\ \liminf_{n \rightarrow \infty} a_n &= \lim_{n \rightarrow \infty} \underline{a}_n = \lim_{k \rightarrow \infty} \inf\{a_n : n \geq k\} = \sup_{k \in \mathbb{N}} \inf_{n \geq k} a_n. \end{aligned}$$

EXAMPLE 2.19. Let $a_n = (-1)^n$. Note that $-1 \leq a_n \leq 1$ for all n . Now, for any k , there is some $k' \geq k$ so that $b_{k'} = 1$. Thus $\bar{b}_k = \sup_{n \geq k} a_k = 1$. Similarly $\underline{b}_k = -1$ for all k . Thus $\limsup_n b_n = 1$ and $\liminf b_n = -1$.

Here are a few more examples computing \limsup and \liminf .

EXAMPLE 2.20. (1) Let $a_n = \frac{1}{n}$. Since $a_n \downarrow$, $\bar{a}_k = \sup_{n \geq k} a_n = a_k = \frac{1}{k}$. Thus $\limsup_n a_n = \lim_k a_k = 0$. On the other hand, for any k , $\inf_k \underline{a}_k = 0$ (by the Archimedean property), and so $\liminf_n a_n = \lim_k 0 = 0$. In this case, the \limsup and \liminf agree.
 (2) Let $b_n = \frac{(-1)^n}{n}$. Note that $-1 \leq b_n \leq 1$ for all n , and more generally $|b_n| \leq \frac{1}{n}$. For any k , we therefore have $\bar{b}_k = \sup\{b_n : n \geq k\} \leq \sup\{|b_n| : n \geq k\} = \frac{1}{k}$ and similarly $\underline{b}_k \geq -\frac{1}{k}$. Now, $\underline{b}_k \leq \bar{b}_k$ (the sup of any set is \geq its inf). Thus

$$-\frac{1}{k} \leq \underline{b}_k \leq \bar{b}_k \leq \frac{1}{k}.$$

Since $\pm\frac{1}{k} \rightarrow 0$, it follows from the Squeeze Theorem that $\lim_k \underline{b}_k = \lim_k \bar{b}_k = 0$. Thus $\limsup_n b_n = \liminf_n b_n = 0$.

(3) The sequence $(1, 2, 3, 1, 2, 3, 1, 2, 3, \dots)$ has $\limsup = 3$ and $\liminf = 1$.

(4) Let $c_n = n$. This is not a bounded sequence, so it doesn't fit the mold for \limsup and \liminf . Indeed, for any k , $\sup_{n \geq k} n$ does not exist for any k , and so $\limsup_n c_n$ does not exist. On the other hand, $\inf_{n \geq k} a_n = k$ does exist, but this sequence is unbounded and has no limit, so $\liminf c_n$ does not exist. This highlights the fact that we need *both* and upper *and* a lower bound in order for either \limsup or \liminf to exist.

In (1) and (2) in the example, \limsup and \liminf agree. This will always happen for a convergent sequence.

PROPOSITION 2.21. Let (a_n) be a bounded sequence. Then $\lim_n a_n$ exists iff $\limsup_n a_n = \liminf_n a_n$, in which case all three limits have the same value.

PROOF. Suppose that $\limsup_n a_n = \liminf_n a_n$. Thus \underline{a}_k and \bar{a}_k both converge to the same value. Since $\underline{a}_k \leq a_k \leq \bar{a}_k$ for each k , by the Squeeze Theorem, a_k also converges to this value, as claimed. Conversely, suppose that $\lim_n a_n = a$ exists. Let $\epsilon > 0$, and choose $N \in \mathbb{N}$ large enough that $|a_n - a| < \epsilon$ for all $n \geq N$. That is

$$a - \epsilon < a_n < a + \epsilon, \quad n \geq N.$$

It follows that

$$a - \epsilon \leq \inf_{n \geq k} a_n \leq \sup_{n \geq k} a_n \leq a + \epsilon, \quad k \geq N$$

which shows that both \bar{a}_k and \underline{a}_k are in $[a - \epsilon, a + \epsilon]$ for $k \geq N$. Thus they both converge to a , as claimed. \square

As with sup and inf, there is a useful trick for transforming statements about lim sup into statements about lim inf.

PROPOSITION 2.22. *Let (a_n) be a bounded sequence. Then $\liminf_n(-a_n) = -\limsup_n a_n$.*

PROOF. Recall that, for any bounded set A , if $-A = \{-a : a \in A\}$, then $\sup(-A) = -\inf A$ and $\inf(-A) = -\sup A$. Now, Let $b_n = -a_n$. Then $\underline{b}_k = \inf\{b_n : n \geq k\} = \inf\{-a_n : n \geq k\} = -\sup\{a_n : n \geq k\} = -\bar{a}_k$. Thus

$$\liminf_{n \rightarrow \infty} b_n = \sup\{\underline{b}_k : k \in \mathbb{N}\} = \sup\{-\bar{a}_k : k \in \mathbb{N}\} = -\inf\{\bar{a}_k : k \in \mathbb{N}\} = -\limsup_{n \rightarrow \infty} a_n.$$

\square

Here is a useful characterization of lim sup and lim inf.

PROPOSITION 2.23. *Let (a_n) be a bounded sequence in a complete ordered field. Denote $\bar{a} = \limsup_n a_n$ and $\underline{a} = \liminf_n a_n$. Then \bar{a} and \underline{a} are uniquely determined by the following properties: for all $\epsilon > 0$,*

$$a_n \leq \bar{a} + \epsilon \text{ for all sufficiently large } n, \text{ and}$$

$$a_n \geq \bar{a} - \epsilon \text{ for infinitely many } n,$$

and

$$a_n \leq \underline{a} + \epsilon \text{ for infinitely many } n, \text{ and}$$

$$a_n \geq \underline{a} - \epsilon \text{ for all sufficiently large } n.$$

PROOF. This is an exercise on HW4. \square

To put this into words: there are many “approximate eventual upper bounds” for the sequence: numbers a large enough that the sequence eventually never gets much bigger than a . The lim sup, \bar{a} , is the *smallest* approximate eventual upper bound: it is the unique number that the sequence eventually never strays far above, but also regularly gets close to from below. Similarly, the lim inf, \underline{a} , is the *largest* approximate eventual lower bound.

4. Lecture 8: January 28, 2016

This brings us to an important understanding of \limsup and \liminf : they are the *maximal and minimal subsequential limits*.

THEOREM 2.24. *Let (a_n) be a bounded sequence in a complete ordered field. There exists a subsequence of (a_n) that converges to $\limsup_n a_n$, and there exists a subsequence of (a_n) that converges to $\liminf_n a_n$. Moreover, if (b_k) is any convergent subsequence of (a_n) , then*

$$\liminf_{n \rightarrow \infty} a_n \leq \lim_{k \rightarrow \infty} b_k \leq \limsup_{n \rightarrow \infty} a_n.$$

PROOF. Let $\bar{a} = \limsup_n a_n$. By Proposition 2.23, for any $k \in \mathbb{N}$ there are infinitely many n so that $a_n \geq \bar{a} - \frac{1}{k}$. So, we proceed inductively: choose some n_1 so that $a_{n_1} \geq \bar{a} - 1$. Then, since there are infinitely many of them, we can find some $n_2 > n_1$ so that $a_{n_2} \geq \bar{a} - \frac{1}{2}$. Proceeding, we find an increasing sequence $n_1 < n_2 < \dots < n_k < \dots$ so that $a_{n_k} \geq \bar{a} - \frac{1}{k}$ for each $k \in \mathbb{N}$. We therefore have

$$\bar{a} - \frac{1}{k} \leq a_{n_k} \leq \sup_{m \geq n_k} a_m = \bar{a}_{n_k}. \quad (2.1)$$

Note that (\bar{a}_{n_k}) is a subsequence of (\bar{a}_n) which converges to \bar{a} ; thus, by Proposition 2.13, $\lim_k \bar{a}_{n_k} = \bar{a}$. Hence, by (2.1) and the Squeeze Theorem, it follows that $a_{n_k} \rightarrow \bar{a}$, and we have constructed the desired subsequence. The proof for \liminf is very similar; alternatively, it can be reasoned using Proposition 2.22.

Now to prove the inequalities. Let (b_k) be a subsequence, so $b_k = a_{m_k}$ for some $m_1 < m_2 < m_3 < \dots$. Then

$$\underline{a}_{m_k} = \inf_{n \geq m_k} a_n \leq b_k \leq \sup_{n \geq m_k} a_n = \bar{a}_{m_k}.$$

Thus, applying the Squeeze theorem, it follows that

$$\liminf_{n \rightarrow \infty} a_n = \lim_{k \rightarrow \infty} \underline{a}_{m_k} \leq \lim_{k \rightarrow \infty} b_k \leq \lim_{k \rightarrow \infty} \bar{a}_{m_k} = \limsup_{n \rightarrow \infty} a_n$$

as desired. □

This allows us to immediately prove our first “named theorem” in Real Analysis: the **Bolzano-Weierstrass Theorem**.

THEOREM 2.25 (Bolzano-Weierstrass). *Let (a_n) be a bounded sequence in a complete ordered field, with $a_n \in [\alpha, \beta]$ for all n . Then (a_n) possesses a convergent subsequence, with limit in $[\alpha, \beta]$.*

PROOF. Let $\bar{a} = \limsup_n a_n$. By Theorem 2.24, there is a subsequence (a_{n_k}) of (a_n) that converges to \bar{a} . Note, then, since $\alpha \leq a_{n_k} \leq \beta$ for all k , it follows from the Squeeze Theorem that $\alpha \leq \lim_k a_{n_k} = \bar{a} \leq \beta$, concluding the proof. □

This finally leads us to the converse of Theorem 2.16.

THEOREM 2.26. *Let \mathbb{F} be a complete ordered field (i.e. possessing the least upper bound property). Then \mathbb{F} is Cauchy complete.*

PROOF. Let (a_n) be a Cauchy sequence in \mathbb{F} . By Proposition 2.10, (a_n) is bounded. Thus, by the Bolzano-Weierstrass theorem, there is a subsequence a_{n_k} that converges. It then follows from Proposition 2.13 that (a_n) is convergent, concluding the proof. □

To summarize: we now have three equivalent characterizations of the notion of “completeness” in an Archimedean field:

least upper bound property \iff nested intervals property \iff Cauchy completeness.

We also know, by the half of Theorem 1.20 we’ve proved, that such a field is unique. So, to finally prove the existence of \mathbb{R} , it will suffice to give a construction of a Cauchy complete field that is Archimedean. The supplementary notes “Construction of \mathbb{R} ” describe how this is done in gory detail.

Henceforth, we will deal with the field \mathbb{R} , which satisfies all of the three equivalent completeness properties.

Now comfortably working in \mathbb{R} , let us state a few more (standard) limit theorems.

THEOREM 2.27 (Limit Theorems). *Let (a_n) and (b_n) be convergent sequences in \mathbb{R} , with $a_n \rightarrow a$ and $b_n \rightarrow b$.*

- (1) *The sequence $c_n = a_n + b_n$ converges to $a + b$.*
- (2) *The sequence $d_n = a_n b_n$ converges to ab .*
- (3) *If $b \neq 0$, then $b_n \neq 0$ for almost all n , and $e_n = \frac{a_n}{b_n}$ converges to $\frac{a}{b}$.*

PROOF. For (1), choose $N_a, N_b \in \mathbb{N}$ so that $|a_n - a| < \frac{\epsilon}{2}$ if $n \geq N_a$ and $|b_n - b| < \frac{\epsilon}{2}$ for $n \geq N_b$. For any $n \geq N = \max\{N_a, N_b\}$, we then have $|c_n - (a + b)| = |(a_n - a) + (b_n - b)| \leq |a_n - a| + |b_n - b| < \frac{\epsilon}{2} + \frac{\epsilon}{2} = \epsilon$, proving that $\lim_n c_n = a + b$.

For (2), we need to be slightly more clever. Note that

$$|d_n - ab| = |a_n b_n - ab| = |a_n b_n - a_n b + a_n b - ab| \leq |a_n| |b_n - b| + |a_n - a| |b|.$$

By Proposition 2.10, there is some constant $M > 0$ so that $|a_n| \leq M$ for all n . So, for $\epsilon > 0$, fix N_1 large enough that $|b_n - b| < \frac{\epsilon}{2M}$ for all $n \geq N_1$, and fix N_2 large enough that $|a_n - a| < \frac{\epsilon}{2|b|}$ for all $n \geq N_2$. (If $b = 0$, we can take N_2 to be any number we like.) Then for $N = \max\{N_1, N_2\}$, if $n \geq N$ we have

$$|d_n - ab| \leq |a_n| |b_n - b| + |a_n - a| |b| < M \cdot \frac{\epsilon}{2M} + \frac{\epsilon}{2|b|} \cdot |b| = \epsilon,$$

proving that $\lim_n d_n = ab$.

For (3), first we need to show that (e_n) even makes sense. Note that $e_n = \frac{a_n}{b_n}$ is not well-defined for any n for which $b_n = 0$. But we’re only concerned about tails of sequences for limit statements, so once we’ve proven that $b_n \neq 0$ for almost all n , we know that e_n is well-defined for all large n . For this, we use the assumption that $b \neq 0$, and so $|b| > 0$. Since $\lim_n b_n = b$, there is an $N_0 \in \mathbb{N}$ so that, for $n > N_0$, $|b_n - b| < \frac{|b|}{2}$; i.e. $-\frac{|b|}{2} < b_n - b < \frac{|b|}{2}$, and so $b_n < b + \frac{|b|}{2}$ and also $b_n > b - \frac{|b|}{2}$. Now, $b \neq 0$ so either $b < 0$ or $b > 0$. If $b < 0$, then $|b| = -b$ in which case $b_n < b + \frac{|b|}{2} = b - \frac{b}{2} = \frac{b}{2} < 0$; that is, for $n > N_0$, $b_n < 0$. If, on the other hand, $b > 0$, then $|b| = b$, and so $b_n > b - \frac{|b|}{2} = b - \frac{b}{2} = \frac{b}{2} > 0$; that is, for $n > N_0$, $b_n > 0$. Thus, in all cases, $b_n \neq 0$ for $n > N_0$, proving the first claim.

For the limit statement, note that $e_n = a_n \cdot \frac{1}{b_n}$. So, by (2), it suffices to show that $\frac{1}{b_n} \rightarrow \frac{1}{b}$. Compute that

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

As shown above, there is N_0 so that, for $n > N_0$, then $b_n > \frac{b}{2} = \frac{|b|}{2}$ if $b_n > 0$ and $b_n < \frac{b}{2} = -\frac{|b|}{2}$ if $b_n < 0$; i.e. this means that $|b_n| > \frac{|b|}{2}$ for $n > N_0$. Hence, we have

$$\left| \frac{1}{b_n} - \frac{1}{b} \right| = \frac{|b_n - b|}{|b_n||b|} < 2 \frac{|b_n - b|}{|b|^2}, \quad n > N_0.$$

By assumption, $b_n \rightarrow b$, and so we can choose N' large enough that $|b_n - b| < \frac{|b|^2}{2}\epsilon$ for $n > N'$. Thus, letting $N = \max\{N_0, N'\}$, we have

$$\left| \frac{1}{b_n} - \frac{1}{b} \right| < 2 \frac{|b_n - b|}{|b|^2} < \frac{2}{|b|^2} \cdot \frac{|b|^2}{2}\epsilon = \epsilon, \quad n > N.$$

This proves that $\frac{1}{b_n} \rightarrow \frac{1}{b}$ as claimed. \square

One might hope that Theorem 2.27 carries over to \limsup and \liminf ; but this is not the case.

EXAMPLE 2.28. Consider the sequences $a_n = (-1)^n$ and $b_n = -a_n = (-1)^{n+1}$. Then $\limsup_n a_n = \limsup_n b_n = 1$, $\liminf_n a_n = \liminf_n b_n = -1$, but $a_n + b_n = 0$ so $\limsup_n (a_n + b_n) = \liminf_n (a_n + b_n) = 0$. Hence, in this example we have

$$-2 = \liminf_{n \rightarrow \infty} a_n + \liminf_{n \rightarrow \infty} b_n < \liminf_{n \rightarrow \infty} (a_n + b_n) = 0 = \limsup_{n \rightarrow \infty} (a_n + b_n) < \limsup_{n \rightarrow \infty} a_n + \limsup_{n \rightarrow \infty} b_n = 2.$$

The *inequalities* in the example do turn out to be true in general.

PROPOSITION 2.29. *Let (a_n) and (b_n) be bounded sequences in \mathbb{R} . The following always hold true.*

$$\begin{aligned} \liminf_{n \rightarrow \infty} (a_n + b_n) &\geq \liminf_{n \rightarrow \infty} a_n + \liminf_{n \rightarrow \infty} b_n, \quad \text{and} \\ \limsup_{n \rightarrow \infty} (a_n + b_n) &\leq \limsup_{n \rightarrow \infty} a_n + \limsup_{n \rightarrow \infty} b_n. \end{aligned}$$

If $a_n \geq 0$ and $b_n \geq 0$ for all sufficiently large n , we also have the following.

$$\begin{aligned} \liminf_{n \rightarrow \infty} (a_n \cdot b_n) &\geq \liminf_{n \rightarrow \infty} a_n \cdot \liminf_{n \rightarrow \infty} b_n, \quad \text{and} \\ \limsup_{n \rightarrow \infty} (a_n \cdot b_n) &\leq \limsup_{n \rightarrow \infty} a_n \cdot \limsup_{n \rightarrow \infty} b_n. \end{aligned}$$

PROOF. The proofs of the \limsup inequalities are exercises on HW4. Assuming these, the \liminf statements follow from Proposition 2.22. For example, we have

$$\liminf_{n \rightarrow \infty} (a_n + b_n) = \liminf_{n \rightarrow \infty} [-(-a_n - b_n)] = - \limsup_{n \rightarrow \infty} [(-a_n) + (-b_n)].$$

Since $\limsup_n [(-a_n) + (-b_n)] \leq \limsup_n (-a_n) + \limsup_n (-b_n)$ by HW4, taking negatives reverses the inequality, giving

$$- \limsup_{n \rightarrow \infty} [(-a_n) + (-b_n)] \geq - \limsup_{n \rightarrow \infty} (-a_n) - \limsup_{n \rightarrow \infty} (-b_n).$$

Now using Proposition 2.22 again on each term, we then have

$$\liminf_{n \rightarrow \infty} (a_n + b_n) \geq - \limsup_{n \rightarrow \infty} (-a_n) - \limsup_{n \rightarrow \infty} (-b_n) = \liminf_{n \rightarrow \infty} a_n + \liminf_{n \rightarrow \infty} b_n$$

as claimed. The proof of the inequality for products is very similar. \square

Let us close out our discussion (for now) of limits of real sequences with a rigorous treatment of the following special kinds of sequences.

PROPOSITION 2.30. Let $p > 0$ and $\alpha \in \mathbb{R}$.

- (1) $\lim_{n \rightarrow \infty} \frac{1}{n^p} = 0$.
- (2) $\lim_{n \rightarrow \infty} p^{1/n} = 1$.
- (3) $\lim_{n \rightarrow \infty} n^{1/n} = 1$.
- (4) If $p > 1$ and $\alpha \in \mathbb{R}$, then $\lim_{n \rightarrow \infty} \frac{n^\alpha}{p^n} = 0$.
- (5) If $|p| < 1$, then $\lim_{n \rightarrow \infty} p^n = 0$.

PROOF. For (1): fix $\epsilon > 0$, and choose $N \in \mathbb{N}$ large enough that $\frac{1}{N} < \epsilon^{1/p}$. Then for $n \geq N$, $\frac{1}{n} \leq \frac{1}{N} < \epsilon^{1/p}$, and so $0 < \frac{1}{n^p} = \left(\frac{1}{n}\right)^p < \epsilon$. This shows that $\frac{1}{n^p} \rightarrow 0$ as claimed.

For (2): as above, in the case $p = 1$ the sequence is constant $1^{1/n} = 1$ with limit 1. If $p > 1$, put $x_n = p^{1/n} - 1$. Since $p > 1$ we have $p^{1/n} > 1$ and so $x_n > 0$. From the binomial theorem, then,

$$(1 + x_n)^n = \sum_{k=0}^n \binom{n}{k} x_n^k \geq 1 + nx_n.$$

By definition $(1 + x_n)^n = p$, and so

$$0 < x_n < \frac{(1 + x_n)^n - 1}{n} = \frac{p - 1}{n}.$$

Knowing that $\frac{p-1}{n} \rightarrow 0$, it now follows from the Squeeze Theorem that $x_n \rightarrow 0$. This proves the limit in the case $p > 1$. If, on the other hand, $0 < p < 1$, then $r = \frac{1}{p} > 1$, and $p^{1/n} = \left(\frac{1}{r}\right)^{1/n} = \frac{1}{r^{1/n}}$. We have just proved that $r^{1/n} \rightarrow 1$, and so it follows from Theorem 2.27(3) that $p^{1/n} \rightarrow \frac{1}{1} = 1$.

For (3): we follow a similar outline. Let $x_n = n^{1/n} - 1$, which is ≥ 0 (and > 0 for $n > 1$). We use the binomial theorem again, this time estimating with the quadratic term:

$$n = (1 + x_n)^n = \sum_{k=1}^n \binom{n}{k} x_n^k \geq \binom{n}{2} x_n^2 = \frac{n(n-1)}{2} x_n^2.$$

Thus, we have (for $n \geq 2$)

$$0 \leq x_n \leq \sqrt{\frac{2}{n-1}}$$

and by the Squeeze Theorem $x_n \rightarrow 0$.

For (4): Choose a positive integer $\ell > \alpha$. Let $p = 1 + r$, so $r > 0$. Applying the binomial theorem again, when $n > \ell$ we have

$$p^n = (1 + r)^n = \sum_{k=0}^n \binom{n}{k} r^k > \binom{n}{\ell} r^\ell = \frac{n(n-1) \cdots (n-\ell+1)}{\ell!} r^\ell.$$

Now, if we choose $n \geq 2\ell$, each term $n - \ell + j \geq \frac{n}{2}$ for $1 \leq j \leq \ell$, and so in this range

$$p^n > \frac{1}{\ell!} \left(\frac{n}{2}\right)^\ell r^\ell.$$

Hence, for $n \geq 2\ell$, we have

$$\frac{n^\alpha}{p^n} < n^\alpha \cdot \frac{\ell! 2^\ell}{n^\ell r^\ell} = \frac{\ell! 2^\ell}{r^\ell} \cdot n^{\alpha-\ell}.$$

This is a constant $\frac{\ell!2^\ell}{r^\ell}$ times $n^{\alpha-\ell}$, where $\alpha - \ell < 0$; applying part (1) with $p = \alpha - \ell$ proves the result.

Finally, for (5): the special case of (4) with $\alpha = 0$ yields $\frac{1}{r^n} \rightarrow 0$ when $r > 1$. Thus, with $|p| < 1$, setting $r = \frac{1}{|p|}$ gives us $|p^n| = |p|^n \rightarrow 0$. The reader should prove (if they haven't already) that $|a_n| \rightarrow 0$ iff $a_n \rightarrow 0$, so it follows that $p^n \rightarrow 0$ as claimed. \square

CHAPTER 3

Extensions of \mathbb{R} ($\overline{\mathbb{R}}$ and \mathbb{C})

1. Lecture 9: February 2, 2016

Now that we have a good understanding of real numbers, it is convenient to extend them a little bit to give us language about certain kinds of divergent sequences.

DEFINITION 3.1. Let (a_n) be a sequence in \mathbb{R} . Say that a_n **diverges to** $+\infty$ or $a_n \rightarrow +\infty$ if

$$\forall M > 0 \exists N \in \mathbb{N} \text{ s.t. } \forall n \geq N \ a_n > M.$$

That is: no matter how large a bound M we choose, it is a lower bound for a_n for all sufficiently large n . Similarly, we say that a_n **diverges to** $-\infty$ if $-a_n \rightarrow +\infty$; this is equivalent to

$$\forall M > 0 \exists N \in \mathbb{N} \text{ s.t. } \forall n \geq N \ a_n < -M.$$

The expressions $a_n \rightarrow \pm\infty$ are also sometimes written as $\lim_{n \rightarrow \infty} a_n = \pm\infty$, and accordingly it is sometimes pronounced as a_n **converges to** $\pm\infty$.

EXAMPLE 3.2. The sequence $a_n = n^p$ diverges to $+\infty$ for any $p > 0$. Indeed, fix a large $M > 0$. Then $M^{1/p} > 0$, and by the Archimedean property there is an $N \in \mathbb{N}$ with $N > M^{1/p}$. Thus, for $n \geq N$, $n > M^{1/p}$, and so $a_n = n^p > (M^{1/p})^p = M$, as desired.

On the other hand, the sequence $(a_n) = (1, 0, 2, 0, 3, 0, 4, 0, \dots)$ does not diverge to $+\infty$: no matter how large N is, there is some integer $n \geq N$ with $a_n = 0$. (Indeed, we can either take $n = N$ or $n = N + 1$.) This sequence diverges, but it does not diverge to $+\infty$.

This suggests that we include the symbols $\pm\infty$ in the field \mathbb{R} . We must be careful how to do this, however. We have already proved that \mathbb{R} is the *unique* complete ordered field, so no matter how we add $\pm\infty$, the resulting object cannot be a complete ordered field. In fact, it won't be a field at all, for we won't always be able to do algebraic operations.

DEFINITION 3.3. Let $\overline{\mathbb{R}} = \mathbb{R} \cup \{-\infty, +\infty\}$. We make $\overline{\mathbb{R}}$ into an ordered set as follows: given $x, y \in \overline{\mathbb{R}}$, if in fact $x, y \in \mathbb{R}$ then we use the order relation from \mathbb{R} to compare x, y . If one of the two (say x) is in \mathbb{R} , then we declare $-\infty < x < +\infty$. Finally, we declare $-\infty < +\infty$.

We make the following conventions. If $a \in \overline{\mathbb{R}}$ with $a > 0$, then $\pm\infty \cdot a = a \cdot \pm\infty = \pm\infty$; if $a \in \overline{\mathbb{R}}$ with $a < 0$ then $\pm\infty \cdot a = a \cdot \pm\infty = \mp\infty$. We also declare that $a + (\pm\infty) = \pm\infty$ for any $a \in \mathbb{R}$, and that $(+\infty) + (+\infty) = +\infty$ while $(-\infty) + (-\infty) = -\infty$. We leave all the following expressions **undefined**:

$$(+\infty) + (-\infty), (-\infty) + (+\infty), \frac{\infty}{\infty}, 0 \cdot (\pm\infty), \text{ and } (\pm\infty) \cdot 0.$$

EXAMPLE 3.4. Let $\alpha, \beta \in \mathbb{R}$ with $\alpha > 0$, and let $a_n = n$ while $b_n = -\alpha n + \beta$. Then

$$\lim_{n \rightarrow \infty} (a_n + b_n) = \begin{cases} +\infty, & \text{if } \alpha < 1 \\ \beta, & \text{if } \alpha = 1 \\ -\infty, & \text{if } \alpha > 1 \end{cases}.$$

Hence the value of the limit of the sum depends on the value of α . However, Example 3.2 shows that $a_n \rightarrow +\infty$ while a similar argument shows that $b_n \rightarrow -\infty$ for any α, β . So we ought to have

$$\lim_{n \rightarrow \infty} a_n + \lim_{n \rightarrow \infty} b_n \text{ “} = \text{”} (+\infty) + (-\infty).$$

This highlights why it is important to leave such expressions undefined: there is no way to consistently define them that respects the limit theorems.

We can also use these conventions to extend the notions of sup and inf to unbounded sets, and the notions of lim sup and lim inf to unbounded sequences.

DEFINITION 3.5. *Let $E \subseteq \mathbb{R}$ be any nonempty subset. If E is not bounded above, declare $\sup E = +\infty$; if E is not bounded below, declare $\inf E = -\infty$. We also make the convention that $\sup(\emptyset) = -\infty$ while $\inf(\emptyset) = +\infty$. (Note: this means that, in the one special case $E = \emptyset$, it is not true that $\inf E \leq \sup E$.)*

Similarly, let (a_n) be any sequence in \mathbb{R} . If (a_n) is not bounded above, declare $\limsup_n a_n = +\infty$; if (a_n) is not bounded below, declare $\liminf_n a_n = -\infty$.

With these conventions, essentially all of the theorems involving limits extend to unbounded sequences.

PROPOSITION 3.6. *Using the preceding conventions, Lemma 2.4, Proposition 2.6, Proposition 2.13(2), Squeeze Theorem 2.14, Proposition 2.21, Proposition 2.22, and Theorem 2.24 all generalize to the cases where the limits in the statements are allowed to be in $\overline{\mathbb{R}}$ rather than just \mathbb{R} . Moreover, Theorem 2.27 and Proposition 2.29 also hold in this more general setting whenever the statements make sense: i.e. excluding the cases when the involved expressions are undefined (like $(+\infty) + (-\infty)$).*

PROOF. It would take many pages to prove all of the special cases of all of these results remain valid in the extended reals. Let us choose just one to illustrate: Theorem 2.27(1): if $\lim_n a_n = a$ and $\lim_n b_n = b$, then $\lim_n (a_n + b_n) = a + b$. We already know this holds true when $a, b \in \mathbb{R}$. If $a = +\infty$ and $b = -\infty$, or $a = -\infty$ and $b = +\infty$, the sum $a + b$ is undefined, and so we exclude these cases from the statement of the theorem. So we only need to consider the cases that $a \in \mathbb{R}$ and $b = \pm\infty$, $a = \pm\infty$ and $b \in \mathbb{R}$, or $a = b = \pm\infty$.

- $a \in \mathbb{R}$ and $b = +\infty$: Since (a_n) is convergent in \mathbb{R} , it is bounded; thus say $|a_n| \leq L$. Then fix $M > 0$ and choose N so that $b_n > M + L$ for $n \geq N$. Thus $a_n + b_n > -L + (M + L) = M$ for $n \geq N$, and so $a_n + b_n \rightarrow +\infty$. The argument is similar when $b = -\infty$.
- $a = \pm\infty$ and $b \in \mathbb{R}$: this is the same as the previous case, just reverse the roles of a_n and b_n and a and b .
- $a = b = +\infty$: let $M > 0$, and choose N_1 so that $a_n > M/2$ for $n \geq N_1$; choose N_2 so that $b_n > M/2$ for $n \geq N_2$. Thus, for $n \geq N = \max\{N_1, N_2\}$, it follows that $a_n + b_n \geq M/2 + M/2 = M$, proving that $a_n + b_n \rightarrow +\infty$ as required. The argument when $a = b = -\infty$ is very similar.

□

Now we turn to a very different extension of \mathbb{R} : the Complex Numbers. We've already discussed them a little bit, in Example 1.12(3-3.5) and HW1.4, so we'll start by reiterating that discussion. We will rely on our knowledge of linear algebra.

DEFINITION 3.7. Let \mathbb{C} denote the following set of 2×2 matrices over \mathbb{R} :

$$\mathbb{C} = \left\{ \begin{bmatrix} a & -b \\ b & a \end{bmatrix} : a, b \in \mathbb{R} \right\}.$$

Then $\mathbb{C} = \text{span}_{\mathbb{R}}\{I, J\}$, where

$$I = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, \quad J = \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix}.$$

As is customary, we denote $I = 1$ and $J = i$. We can compute that $J^2 = -I$, so $i^2 = -1$. Every complex number has the form $a1 + bi$ for unique $a, b \in \mathbb{R}$; we often suppress the 1 and write this as $a + ib$. We think of $\mathbb{R} \subset \mathbb{C}$ via the identification $a \leftrightarrow a + i0$ (so a is the matrix aI).

It is convenient to construct \mathbb{C} this way, since, as a collection of matrices, we already have addition and multiplication built in; and we have all the tools of linear algebra to prove properties of \mathbb{C} .

PROPOSITION 3.8. Denote $1_{\mathbb{C}} = I$ and $0_{\mathbb{C}}$ the 2×2 zero matrix. Define $+$ and \cdot on \mathbb{C} by their usual matrix definitions. Then \mathbb{C} is a field.

PROOF. Most of the work is done for us, since $+$ and \cdot of matrices are associative and distributive, and $+$ is commutative, and $1_{\mathbb{C}}$ and $0_{\mathbb{C}}$ are multiplicative and additive identities. All that we are left to verify are the following three properties:

- \mathbb{C} is closed under $+$ and \cdot , i.e. we need to check that if $z, w \in \mathbb{C}$ then $z + w \in \mathbb{C}$ and $z \cdot w \in \mathbb{C}$. Setting $z = a + ib$ and $w = c + id$, we simply compute

$$z + w = \begin{bmatrix} a & -b \\ b & a \end{bmatrix} + \begin{bmatrix} c & -d \\ d & c \end{bmatrix} = \begin{bmatrix} a+c & -b-d \\ b+d & a+c \end{bmatrix} = (a+c) + (b+d)i \in \mathbb{C},$$

$$z \cdot w = \begin{bmatrix} a & -b \\ b & a \end{bmatrix} \begin{bmatrix} c & -d \\ d & c \end{bmatrix} = \begin{bmatrix} ac-bd & -ad-bc \\ ad+bc & ac-bd \end{bmatrix} = (ac-bd) + (ad+bc)i \in \mathbb{C}.$$

- \cdot is commutative: this follows from the computation above: if we exchange $z \leftrightarrow w$, meaning $a \leftrightarrow c$ and $b \leftrightarrow d$, the value of the product $z \cdot w$ is unaffected, so $z \cdot w = w \cdot z$.
- If $z \in \mathbb{C} \setminus \{0_{\mathbb{C}}\}$ then z^{-1} exists: here we use the criterion that a matrix z is invertible iff $\det(z) \neq 0$. We can readily compute that, for $z \in \mathbb{C}$,

$$\det(z) = \det \begin{bmatrix} a & -b \\ b & a \end{bmatrix} = a^2 + b^2$$

and this = 0 iff $a = b = 0$ meaning $a + ib = 0_{\mathbb{C}}$.

□

Now more notation.

DEFINITION 3.9. Let $z = a + ib \in \mathbb{C}$. We denote $a = \Re(z)$ and $b = \Im(z)$, the **Real and Imaginary parts** of z . Define the **modulus or absolute value** of z to be

$$|z| = \sqrt{\det(z)} = \sqrt{a^2 + b^2} = \sqrt{\Re(z)^2 + \Im(z)^2}.$$

For $z \in \mathbb{C}$, its **complex conjugate** \bar{z} is the complex number $\bar{z} = \Re(z) - i\Im(z)$; in terms of matrices, this is just the transpose $\bar{z} = z^{\top}$.

Note that $z + \bar{z} = 2\Re(z)$ and $z - \bar{z} = 2i\Im(z)$. Since i is invertible (indeed $i^{-1} = -i$), it follows that

$$\Re(z) = \frac{z + \bar{z}}{2}, \quad \Im(z) = \frac{z - \bar{z}}{2i}. \quad (3.1)$$

Note that, if $z \in \mathbb{C}$ happens to be in \mathbb{R} (meaning that $\Im z = 0$ so $z = \Re(z)$), then $|z| = \sqrt{\Re(z)^2 + 0} = \sqrt{z^2} = |z|$ corresponds to the absolute value in \mathbb{R} ; so the complex modulus generalizes the familiar absolute value.

2. Lecture 10: February 4, 2016

Here are some important properties of modulus and complex conjugate.

LEMMA 3.10. *Let $z, w \in \mathbb{C}$. Then we have the following.*

- (1) $\overline{\overline{z}} = z$.
- (2) $\overline{z + w} = \overline{z} + \overline{w}$ and $\overline{zw} = \overline{z} \cdot \overline{w}$.
- (3) $z\overline{z} = |z|^2$.
- (4) $|\overline{z}| = |z|$.
- (5) $|zw| = |z||w|$, and so $|z^n| = |z|^n$ for all $n \in \mathbb{N}$.
- (6) $|\Re(z)| \leq |z|$ and $|\Im(z)| \leq |z|$.
- (7) $|z + w| \leq |z| + |w|$.
- (8) $|z| = 0$ iff $z = 0$.
- (9) If $z \neq 0$ then z^{-1} (which we also write as $\frac{1}{z}$) is given by

$$z^{-1} = \frac{\overline{z}}{|z|^2}.$$

- (10) If $z \neq 0$ then $|z^{-1}| = |z|^{-1}$, and so $|z^n| = |z|^n$ for all $n \in \mathbb{Z}$.

PROOF. (1) is the familiar linear algebra fact that $(z^\top)^\top = z$, and (2) follows similarly from linear algebra (and the commutativity of \cdot in \mathbb{C}): $\overline{z + w} = (z + w)^\top = z^\top + w^\top = \overline{z} + \overline{w}$, and $\overline{zw} = (zw)^\top = w^\top z^\top = z^\top w^\top = \overline{z} \cdot \overline{w}$. For (3), writing $z = a + ib$ we have

$$z\overline{z} = (a + ib)(a - ib) = a^2 + b^2 + (ab - ab)i = a^2 + b^2 = |z|^2.$$

(4) then follows that from this and (1), and commutativity of complex multiplication: $|\overline{z}|^2 = \overline{z}z = z\overline{z} = |z|^2$; taking square roots (using the fact that $|z| \geq 0$) shows that $|\overline{z}| = |z|$. (Alternatively, for (4), we simply note that $|\overline{z}| = \det(z^\top) = \det(z) = |z|$.)

(5) is a well-known property of determinants: $|zw| = \det(zw) = \det(z)\det(w) = |z||w|$; taking $z = w$ and doing induction shows that $|z^n| = |z|^n$. (6) follows easily from the fact that $|z| = \sqrt{|\Re(z)|^2 + |\Im(z)|^2}$. For (7), we have

$$|z + w|^2 = (z + w)(\overline{z + w}) = (z + w)(\overline{z} + \overline{w}) = z\overline{z} + z\overline{w} + w\overline{z} + w\overline{w} = |z|^2 + z\overline{w} + w\overline{z} + |w|^2.$$

The two middle terms can be written as $z\overline{w} + w\overline{z} = z\overline{w} + \overline{(z\overline{w})}$ and, by (3.1), this equals $2\Re(z\overline{w})$. Now, any real number x is $\leq |x|$, and so

$$|z + w|^2 = |z|^2 + 2\Re(z\overline{w}) + |w|^2 \leq |z|^2 + 2|\Re(z\overline{w})| + |w|^2 \leq |z|^2 + 2|z\overline{w}| + |w|^2$$

where we have used (6). From (4) and (5), $|z\overline{w}| = |z||\overline{w}| = |z||w|$, and so finally we have

$$|z + w|^2 \leq |z|^2 + 2|z||w| + |w|^2 = (|z| + |w|)^2.$$

Taking square roots proves the result.

For (8), it is immediate that $|0| = 0$; the converse was shown in the proof of Proposition 3.8: $|z| = \det(z) = 0$ iff $(\Re(z))^2 + (\Im(z))^2 = 0$ which happens only when $\Re(z) = \Im(z) = 0$, so $z = 0$. Part (9) follows similarly from the matrix representation; alternatively we can simply check from (3) that

$$z \cdot \frac{\overline{z}}{|z|^2} = \frac{z\overline{z}}{|z|^2} = 1$$

showing that $z^{-1} = \frac{\overline{z}}{|z|^2}$. Finally, for (10), using (5) we have

$$|z^{-1}||z| = |z^{-1}z| = |1| = 1$$

so $|z|^{-1} = |z^{-1}|$. An induction argument combining this with (5) shows that $|z|^{-n} = |z^{-n}|$ for $n \in \mathbb{N}$, and coupling this with the second statement of (5) concludes the proof. \square

Items (7) and (8) of Lemma 3.10 show that the complex modulus behaves just like the real absolute value: it satisfies the triangle inequality, and is only 0 at 0. These properties are all that were necessary to make most of the technology of limits of sequences in \mathbb{R} work, and so we can now use the complex modulus to extend these notions to \mathbb{C} .

DEFINITION 3.11. *Let (z_n) be a sequence in \mathbb{C} . Given $z \in \mathbb{C}$, say that $\lim_{n \rightarrow \infty} z_n = z$ iff*

$$\forall \epsilon > 0 \exists N \in \mathbb{N} \text{ s.t. } \forall n \geq N \ |z_n - z| < \epsilon.$$

Say that (z_n) is a Cauchy sequence if

$$\forall \epsilon > 0 \exists N \in \mathbb{N} \text{ s.t. } \forall n, m \geq N \ |z_n - z_m| < \epsilon.$$

Note that these are, symbolically, *exactly the same* as the definitions (6.1 and 2.7) of limits and Cauchy sequences of real numbers; the only difference is, we now interpret $|z|$ to mean the modulus of the complex number z rather than the absolute value of a real number.

The properties of complex modulus mirroring those of real absolute value allow us to prove the results of Lemmas 2.4 and 2.8, Propositions 2.10 and 2.13, and the Limit Theorems (Theorem 2.27) with nearly identical proofs. To summarize:

- THEOREM 3.12. (1) *Limits are unique: if $z_n \rightarrow z$ and $z_n \rightarrow w$, then $z = w$.*
 (2) *Every convergent sequence in \mathbb{C} is Cauchy.*
 (3) *Every Cauchy sequence in \mathbb{C} is bounded.*
 (4) (a) *If $z_n \rightarrow z$ then any subsequence of (z_n) converges to z .*
 (b) *If (z_n) is Cauchy then any subsequence of (z_n) is Cauchy.*
 (c) *If (z_n) is Cauchy and has a convergent subsequence with limit z , then $z_n \rightarrow z$.*
 (5) *If $z_n \rightarrow z$ and $w_n \rightarrow w$, then $z_n + w_n \rightarrow z + w$, $z_n w_n \rightarrow zw$, and if $z \neq 0$ then $z_n \neq 0$ for sufficiently large n and $\frac{1}{z_n} \rightarrow \frac{1}{z}$.*

To illustrate how to handle complex modulus in these proofs, let us look at the analog of Proposition 2.10: that Cauchy sequences are bounded. As before, we set $\epsilon = 1$ and let N be large enough that $|z_n - z_m| < 1$ whenever $n, m > N$. Thus, taking $m = N + 1$, for any $n > N$ we have $|z_n - z_{N+1}| < 1$. Now, $z_n = (z_n - z_{N+1}) + z_{N+1}$, and so by the triangle inequality

$$|z_n| = |(z_n - z_{N+1}) + z_{N+1}| \leq |z_n - z_{N+1}| + |z_{N+1}| < 1 + |z_{N+1}|, \quad \forall n > N.$$

So, as in the previous proof, if we set $M = \max\{|z_1|, |z_2|, \dots, |z_N|, 1 + |z_{N+1}|\}$ then $|z_n| \leq M$ for all n .

In fact, convergence and Cauchy-ness of complex sequences boils down to convergence and Cauchy-ness of the real and imaginary parts separately.

PROPOSITION 3.13. *Let (z_n) be a sequence in \mathbb{C} . Then (z_n) is Cauchy iff the two real sequences $(\Re(a_n))$ and $(\Im(b_n))$ are Cauchy, and $z_n \rightarrow z$ iff $\Re(z_n) \rightarrow \Re(z)$ and $\Im(z_n) \rightarrow \Im(z)$.*

PROOF. Let $z_n = a_n + ib_n$. Suppose (a_n) and (b_n) are Cauchy. Fix $\epsilon > 0$ and choose N_1 large enough that $|a_n - a_m| < \frac{\epsilon}{2}$ for $n, m > N_1$, and choose N_2 large enough that $|b_n - b_m| < \frac{\epsilon}{2}$ for

$n, m > N_2$. Then for $n, m > N = \max\{N_1, N_2\}$, we have

$$\begin{aligned} |z_n - z_m| &= |(a_n + ib_n) - (a_m + ib_m)| = |(a_n - a_m) + i(b_n - b_m)| \leq |a_n - a_m| + |i(b_n - b_m)| \\ &= |a_n - a_m| + |b_n - b_m| \\ &< \frac{\epsilon}{2} + \frac{\epsilon}{2} = \epsilon \end{aligned}$$

where, in the second last step, we used the fact that $|i(b_n - b_m)| = |i||b_n - b_m|$ and $|i| = 1$. Thus, (z_n) is Cauchy. For the converse, suppose that (z_n) is Cauchy. Fix $\epsilon > 0$, and choose N large enough that $|z_n - z_m| < \epsilon$ for $n, m > N$. Then we also have

$$\begin{aligned} |\Re(z_n) - \Re(z_m)| &= |\Re(z_n - z_m)| \leq |z_n - z_m| < \epsilon, \quad \text{and} \\ |\Im(z_n) - \Im(z_m)| &= |\Im(z_n - z_m)| \leq |z_n - z_m| < \epsilon \end{aligned}$$

for $n, m > N$. Thus, both $(\Re(z_n))$ and $(\Im(z_n))$ are Cauchy, as claimed.

The proof of the limit statements is very similar, and is left as a homework exercise (on HW6). \square

Now, \mathbb{C} is not an ordered field (you proved this on HW1), so it does not even make sense to ask if it has the least upper bound property (and likewise we cannot talk about a Squeeze Theorem, or \limsup and \liminf). This is one of the main reasons we gave an equivalent characterization of the least upper bound property – Cauchy completeness – that does not explicitly require an order relation.

THEOREM 3.14. *The field \mathbb{C} is Cauchy complete: any Cauchy sequence is convergent.*

PROOF. Let (z_n) be a Cauchy sequence in \mathbb{C} . By Proposition 3.13, the two real sequences $(\Re(z_n))$ and $(\Im(z_n))$ are both Cauchy. Since \mathbb{R} is Cauchy complete, it follows that there are real numbers $a, b \in \mathbb{R}$ so that $\Re(z_n) \rightarrow a$ and $\Im(z_n) \rightarrow b$. It then follows, again by Proposition 3.13, that $z_n \rightarrow a + ib$. \square

In \mathbb{R} , we proved the Bolzano-Weierstrass theorem (that bounded sequences have convergent subsequences) using the technology of \limsup and \liminf . As noted, since \mathbb{C} is not ordered, there is no way to talk about \limsup and \liminf for a complex sequence. Nevertheless, the Bolzano-Weierstrass theorem holds true in \mathbb{C} . We conclude our discussion of \mathbb{C} (for now) with its proof.

THEOREM 3.15 (Bolzano-Weierstrass). *Every bounded sequence in \mathbb{C} contains a convergent subsequence.*

PROOF. Let (z_n) be a bounded sequence. Letting $z_n = a_n + ib_n$, since $|a_n| \leq |z_n|$ and $|b_n| \leq |z_n|$, it follows that (a_n) and (b_n) are bounded sequences in \mathbb{R} . Now, by the Bolzano-Weierstrass theorem for \mathbb{R} , there is a subsequence a_{n_k} of (a_n) that converges to some real number a . Consider now the subsequence b_{n_k} of (b_n) . Since (b_n) is bounded, so is (b_{n_k}) , and so again applying the Bolzano-Weierstrass theorem for \mathbb{R} , there is a *further* subsequence $(b_{n_{k_\ell}})$ that converges to some $b \in \mathbb{R}$. The subsequence $(a_{n_{k_\ell}})$ is a subsequence of the convergent sequence a_{n_k} and hence also converges to a . Thus, by Proposition 3.13, the subsequence $z_{n_{k_\ell}}$ converges to $a + ib$ as $\ell \rightarrow \infty$. \square

REMARK 3.16. This proof highlights an important technique with subsequences in higher dimensional spaces. We chose the second subsequence as a *subsubsequence*, not only a subsequence. Had we tried to select the subsequences of the real and imaginary parts independently, we could not have concluded anything about the two together. Indeed, the Bolzano-Weierstrass theorem gives us a convergent subsequence a_{n_k} and also gives us a convergent subsequence b_{m_k} . But we need to

use the same index n for both a_n and b_n , which might not be possible with independent choices like this. A priori, the chosen convergent subsequence of a_n might have been (a_1, a_3, a_5, \dots) , while from b_n we might have chosen (b_2, b_4, b_6, \dots) , ne'er the 'tween shall meet.

CHAPTER 4

Series

1. Lecture 11: February 9, 2016

We now turn to a special class of sequences called **series**.

DEFINITION 4.1. Let (a_n) be a sequence in \mathbb{R} or \mathbb{C} . The **series** associated to (a_n) is the new sequence (s_n) given by

$$s_n = \sum_{k=1}^n a_k = a_1 + a_2 + \cdots + a_n.$$

It is a bit of a misnomer to refer to series as special kinds of sequences; indeed, any sequence is the series associated to some other sequence. For let (a_n) be a sequence. Define a new sequence (b_n) by

$$b_1 = a_1, \quad b_n = a_n - a_{n-1} \text{ for } n > 1.$$

Then $a_1 = b_1 = \sum_{k=1}^1 b_k$, and for $n > 1$ we compute that

$$\sum_{k=1}^n b_k = b_1 + b_2 + \cdots + b_n = a_1 + (a_2 - a_1) + (a_3 - a_2) + \cdots + (a_n - a_{n-1}) = a_n.$$

Thus, (a_n) (the arbitrary sequence we started with) is the series associated to the sequence (b_n) .

We will see, however, that the concept of convergence is quite different when applied to the series associated to a sequence rather than the sequence itself.

DEFINITION 4.2. Let (a_n) be a sequence in \mathbb{R} or \mathbb{C} , and let $s_n = \sum_{k=1}^n a_k$ be its series. We say that the **series converges** if the sequence (s_n) converges. If the limit is $s = \lim_{n \rightarrow \infty} s_n$, we denote it by

$$s = \sum_{n=1}^{\infty} a_n = \lim_{n \rightarrow \infty} \sum_{k=1}^n a_k.$$

In this case, we will often use the cumbersome-but-standard notation “the series $\sum_{n=1}^{\infty} a_n$ converges”.

EXAMPLE 4.3 (Geometric Series). Let $r \in \mathbb{C}$, and consider the sequence $a_n = r^n$ (in this case it is customary to start at $n = 0$). We can compute the terms in the series exactly, following a trick purportedly invented by Gauss at age 10.

$$\begin{aligned} s_n &= \sum_{k=0}^n a_k = 1 + r + r^2 + \cdots + r^n. \\ \therefore r s_n &= \quad \quad \quad r + r^2 + \cdots + r^n + r^{n+1}. \end{aligned}$$

So, subtracting the two lines, we have

$$(1 - r)s_n = s_n - r s_n = 1 - r^{n+1}.$$

Now, if $r = 1$, this gives no information. In that degenerate case, we simply have $s_n = 1 + 1 + \cdots + 1 = n$, and this series does not converge. In all other cases, we have the explicit formula

$$s_n = \frac{1 - r^{n+1}}{1 - r}.$$

Using the limit theorems, we can decide whether this converges, and to what, just looking at the shifted sequence $a_{n+1} = r^{n+1}$. If $|r| < 1$, then this converges to 0. If $|r| \geq 1$, this sequence does not converge. (This is something you should work out.) Thus, we have

$$\sum_{n=0}^{\infty} r^n = \frac{1}{1 - r}, \quad \text{if } |r| < 1$$

while the series diverges if $|r| \geq 1$.

EXAMPLE 4.4. Consider the sequence $a_n = \frac{1}{n(n+1)}$. We can employ a trick here: the partial fractions decomposition:

$$\frac{1}{n(n+1)} = \frac{1}{n} - \frac{1}{n+1}.$$

Thus, taking the series $s_n = \sum_{k=1}^n a_k$, we have

$$s_n = \sum_{k=1}^n \frac{1}{k(k+1)} = \sum_{k=1}^n \left(\frac{1}{k} - \frac{1}{k+1} \right) = \left(1 - \frac{1}{2} \right) + \left(\frac{1}{2} - \frac{1}{3} \right) + \cdots + \left(\frac{1}{n} - \frac{1}{n+1} \right).$$

This is a *telescoping sum*: all terms except for the first and the last cancel in pairs. Thus, we have a closed formula

$$s_n = 1 - \frac{1}{n+1}.$$

Hence, the series converges, and we have explicitly

$$\sum_{n=1}^{\infty} \frac{1}{n(n+1)} = \lim_{n \rightarrow \infty} s_n = \lim_{n \rightarrow \infty} \left(1 - \frac{1}{n+1} \right) = 1.$$

EXAMPLE 4.5 (Harmonic Series). Consider the series $s_n = \sum_{k=1}^n \frac{1}{k}$. That is

$$s_n = 1 + \frac{1}{2} + \frac{1}{3} + \cdots + \frac{1}{n}.$$

If you add up the first *billion* terms (i.e. s_{10^9}) you get about 21.3. This seems to suggest convergence; after all, the terms are getting arbitrarily small. However, this series *does not converge*. To see why, look at terms s_N with $N = 2^m + 2^{m-1} + \cdots + 2 + 1$ for some positive integer m . (By the way, from the previous example, this could be written explicitly as $N = 2^{m+1} - 1$.) Then we can group terms as

$$s_N = (1) + \left(\frac{1}{2} + \frac{1}{3} \right) + \left(\frac{1}{4} + \frac{1}{5} + \frac{1}{6} + \frac{1}{7} \right) + \cdots + \left(\frac{1}{2^m + 1} + \frac{1}{2^m + 2} + \cdots + \frac{1}{2^{m+1} - 1} \right).$$

That is: we break up the sum into $m + 1$ groups, the first group with 1 term, the second with 2, the third with 4, up to the last with 2^m terms. Now, $1 > \frac{1}{2}$. In the second group of terms, both $\frac{1}{2}$ and $\frac{1}{3}$

are $> \frac{1}{4}$. In the next group, each of the four terms is $> \frac{1}{8}$. That is, we have

$$\begin{aligned} s_N &> \left(\frac{1}{2}\right) + \left(\frac{1}{4} + \frac{1}{4}\right) + \left(\frac{1}{8} + \frac{1}{8} + \frac{1}{8} + \frac{1}{8}\right) + \cdots + \left(\frac{1}{2^{m+1}} + \frac{1}{2^{m+1}} + \cdots + \frac{1}{2^{m+1}}\right) \\ &= \frac{1}{2} + \frac{1}{2} + \frac{1}{2} + \cdots + \frac{1}{2} = \frac{m+1}{2}. \end{aligned}$$

Now, $s_{n+1} = s_n + \frac{1}{n+1} \geq s_n$, so (s_n) is an increasing sequence. We've just shown that, for any integer m , we can find some time N so that $s_N \geq \frac{m+1}{2}$, and so it follows that for all larger $n \geq N$, $s_n \geq s_N \geq \frac{m+1}{2}$. Since $\frac{m+1}{2}$ is arbitrarily larger, we've just proved that $s_n \rightarrow +\infty$ as $n \rightarrow \infty$. So the series diverges.

In Example 4.3, we were able to compute the n th term in the series as a closed formula, and compute the limit directly. It is rare that we can do this explicitly; more often, we will need to make estimates like we did in Example 4.5. So we now begin to discuss some general tools for attacking such limits.

PROPOSITION 4.6 (Cauchy Criterion). *Let a_n be a sequence in \mathbb{R} or \mathbb{C} . Then the series $\sum_{n=1}^{\infty} a_n$ converges if and only if: for every $\epsilon > 0$, there is a natural number $N \in \mathbb{N}$ so that, for all $m \geq n \geq N$,*

$$\left| \sum_{k=n+1}^m a_k \right| < \epsilon.$$

PROOF. This is just a restatement of the Cauchy completeness of \mathbb{R} and \mathbb{C} . Let $s_n = \sum_{k=1}^n a_k$. Then

$$\sum_{k=n+1}^m = a_{n+1} + \cdots + a_m = s_m - s_n.$$

Thus, having decided to always use m to denote the larger of m, n , the statement is that, for every $\epsilon > 0$ there is $N \in \mathbb{N}$ such that, for $m \geq n \geq N$, $|s_m - s_n| < \epsilon$; this is precisely the statement that (s_n) is a Cauchy sequence. In \mathbb{R} or \mathbb{C} , this is equivalent to (s_n) being convergent, as desired. \square

COROLLARY 4.7. *Let (a_n) be a sequence such that the series $\sum_{n=1}^{\infty} a_n$ converges. Then $a_n \rightarrow 0$.*

PROOF. By the Cauchy Criterion (Proposition 4.6), given $\epsilon > 0$ we may find $N \in \mathbb{N}$ so that (letting $n = m - 1$) for $m > N$,

$$\epsilon > \left| \sum_{k=(m-1)+1}^m a_k \right| = |a_m|.$$

This is precisely the statement that $a_m \rightarrow 0$ as $m \rightarrow \infty$. \square

As Example 4.5 points out, the converse to Corollary 4.7 is false: there are sequences, such as $a_n = \frac{1}{n}$, that tend to 0, but for which the series $\sum_{n=1}^{\infty} a_n$ diverges.

It is often impossible to compute the exact value of the sum $\sum_{n=1}^{\infty} a_n$ of a convergent series. More often, we use estimates to approximate the value. More basically, we use estimates to determine whether the series converges or not, without any direct knowledge of the value if it does converge. The most basic test for convergence is the **comparison theorem**.

THEOREM 4.8 (Comparison). *Let (a_n) and (b_n) be sequences in \mathbb{C} .*

- (1) If $b_n \geq 0$ and $\sum_n b_n$ converges, and if $|a_n| \leq b_n$ for all sufficiently large n , then $\sum_n a_n$ converges, and $|\sum_n a_n| \leq \sum_n b_n$.
- (2) If $a_n \geq b_n \geq 0$ for all sufficiently large n and $\sum_n b_n$ diverges, then $\sum_n a_n$ diverges.

PROOF. For item 1: by assumption $\sum_n b_n$ converges, and so by the Cauchy criterion, for given $\epsilon > 0$ we can choose $N_0 \in \mathbb{N}$ so that, for $m \geq n \geq N_0$,

$$\sum_{k=n+1}^m b_k < \epsilon$$

(here we have used the fact that $b_n \geq 0$ to drop the modulus). Now, let N_1 be large enough that $|a_n| \leq b_n$ for $n \geq N_1$. Then for $m \geq n \geq \max\{N_0, N_1\}$, we have

$$\left| \sum_{k=n+1}^m a_k \right| \leq \sum_{k=n+1}^m |a_k| \leq \sum_{k=n+1}^m b_k < \epsilon.$$

So the series $\sum_n a_n$ satisfies the Cauchy criterion, and therefore is convergent.

Item 2 follows from item 1 by contrapositive: if $\sum_n a_n$ converges, then since $b_n = |a_n| \leq a_n$ for all large n , we have just proven that $\sum_n b_n$ converges. Thus, if $\sum_n b_n$ diverges, so must $\sum_n a_n$. \square

EXAMPLE 4.9. The series $\sum_n \frac{1}{\sqrt{n}}$ diverges, since $\frac{1}{\sqrt{n}} \geq \frac{1}{n}$ and, by Example 4.5, $\sum_n \frac{1}{n}$ diverges. On the other hand, note that

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

for $n \geq 1$. Thus $\frac{1}{n^2} \leq \frac{2}{n(n+1)}$. As we computed in Example 4.4, $\sum_{n=1}^{\infty} \frac{1}{n(n+1)} = 1$, so by the limit theorems, $\sum_{n=1}^{\infty} \frac{2}{n(n+1)} = 2$. That is: this series converges. It follows from the comparison test that $\sum_{n=1}^{\infty} \frac{1}{n^2}$ converges.

In showing that the harmonic series diverges, we broke the terms up into groups of exponentially increasing size. This is an important trick known as the **lacunary technique**, and it works well when the sequence of terms is positive and decreasing.

PROPOSITION 4.10 (Lacunary Series). *Suppose (a_n) is a sequence of non-negative numbers that is decreasing: $a_n \geq a_{n+1} \geq 0$ for all n . Then $\sum_{n=1}^{\infty} a_n$ converges if and only if the series*

$$\sum_{k=0}^{\infty} 2^k a_{2^k} = a_1 + 2a_2 + 4a_4 + 8a_8 + \cdots$$

converges.

PROOF. Since $a_k \geq 0$ for all k , the series of partial sums $s_n = \sum_{k=1}^n a_k$ is monotone increasing. Hence, convergence of s_n is equivalent to the *boundedness* of (s_n) . Let $t_k = a_1 + 2a_2 + 4a_4 + \cdots + 2^k a_{2^k}$. We will show that (s_n) is bounded iff (t_k) is bounded.

Note that $2^k \leq 2^{k+1} - 1$, and so if $n < 2^k$ then $n < 2^{k+1} - 1$. Then we have for such n

$$\begin{aligned} s_n &= a_1 + a_2 + a_3 + \cdots + a_n \leq a_1 + a_2 + a_3 + \cdots + a_{2^{k+1}-1} \\ &= (a_1) + (a_2 + a_3) + \cdots + (a_{2^k} + \cdots + a_{2^{k+1}-1}) \\ &\leq a_1 + 2a_2 + \cdots + 2^k a_{2^k} = t_k. \end{aligned}$$

(In the last inequality, we used the fact that a_n is decreasing.) This shows that if (t_k) is bounded, then so is (s_n) . For the converse, we just group terms slightly differently (exactly as we did in the proof of the divergence of the harmonic series): if $n > 2^k$, then

$$\begin{aligned} s_n &= a_1 + a_2 + a_3 + \cdots + a_n \geq a_1 + a_2 + a_3 + \cdots + a_{2^k} \\ &= (a_1) + (a_2) + (a_3 + a_4) + \cdots + (a_{2^{k-1}+1} + \cdots + a_{2^k}) \\ &\geq \frac{1}{2}a_1 + a_2 + 2a_4 \cdots + 2^{k-1}a_{2^k} = \frac{1}{2}t_k. \end{aligned}$$

Thus $t_k \leq 2s_n$ whenever $n > 2^k$. This shows that if (s_n) is bounded then so is (t_k) , concluding the proof. \square

EXAMPLE 4.11. Let $p \in \mathbb{R}$, and consider the series $\sum_{n=1}^{\infty} \frac{1}{n^p}$. We've already seen that this series diverges when $p = 1$. If $p < 1$, then $\frac{1}{n^p} \geq \frac{1}{n}$; it follows by the comparison theorem that the series $\sum_{n=1}^{\infty} \frac{1}{n^p}$ diverges for $p \leq 1$.

On the other hand, consider $p > 1$. Here the sequence of terms $a_n = \frac{1}{n^p}$ is positive and decreasing, so we may use the lacunary series test to determine whether the series converges. Compute that

$$\sum_{k=0}^{\infty} 2^k a_{2^k} = \sum_{k=0}^{\infty} 2^k \frac{1}{2^{kp}} = \sum_{k=0}^{\infty} 2^{(1-p)k}.$$

This is a geometric series, with base $r = 2^{1-p}$. So $0 < r < 1$ provided that $p > 1$, in which case the series converges. Hence, by Proposition 4.10, the series $\sum_{n=1}^{\infty} \frac{1}{n^p}$ converges if and only if $p > 1$.

2. Lecture 12: February 11, 2016

Two generally effective tools for deciding convergence, that you already saw in your calculus class, are the **Root Test** and the **Ratio Test**. Both of them are predicated on rough comparison with a geometric series, cf. Example 4.3. If $a_n = r^n$, then $\sum_n a_n$ converges iff $|r| < 1$. Now, note for this series that this important constant $|r|$ can be computed either as $|a_n|^{1/n}$ or as $|\frac{a_{n+1}}{a_n}|$. Even when these quantities are not constant, they still can give a lot of information about the convergence of the series.

THEOREM 4.12 (Root Test). *Let (a_n) be a sequence in \mathbb{C} . Define*

$$\alpha = \limsup_{n \rightarrow \infty} |a_n|^{1/n}.$$

If $\alpha < 1$, then $\sum_n a_n$ converges. If $\alpha > 1$, then $\sum_n a_n$ diverges.

REMARK 4.13. It is important to note that the theorem gives no information when $\alpha = 1$. Indeed, consider Examples 4.5 and 4.9, showing that $\sum_{n=1}^{\infty} \frac{1}{n}$ diverges, while $\sum_{n=1}^{\infty} \frac{1}{n^2}$ converges. But, in both cases, we have

$$\lim_{n \rightarrow \infty} \left(\frac{1}{n}\right)^{1/n} = \lim_{n \rightarrow \infty} \left(\frac{1}{n^2}\right)^{1/n} = 1$$

(see Theorem 3.20(c) in Rudin). Thus, $\limsup_n |a_n|^{1/n} = 1$ can happen whether $\sum_n a_n$ converges or diverges.

PROOF. Suppose $\alpha < 1$. Then choose any $r \in \mathbb{R}$ with $\alpha < r < 1$. That is, we have $\limsup_n |a_n|^{1/n} < r$. Let $b_n = |a_n|^{1/n}$; then the statement is that $\limsup_n b_n = \lim_n \bar{b}_n < r$. That means that, for all sufficiently large n , $\bar{b}_n < r$, and so since $b_n \leq \bar{b}_n$, we have $|a_n|^{1/n} < r$ for all sufficiently large n . That is: there is some N so that $|a_n| < r^n$ for $n \geq N$. Since the series $\sum_n r^n$ converges (as $0 < r < 1$), it now follows that $\sum_n a_n$ converges by the comparison theorem.

Now, suppose $\alpha > 1$. As above, let $b_n = |a_n|^{1/n}$. Since $\alpha = \limsup_n b_n$, from Theorem 2.24 there is a subsequence b_{n_k} that converges to α . (This is even true of $\alpha = +\infty$; in this case, it is quite easy to see that the series diverges.) In particular, this means that $b_{n_k} > 1$ for all k , and so $|a_{n_k}| = b_{n_k}^{n_k} > 1$ as well. It follows that a_n does not converge to 0, and so by Corollary 4.7, $\sum_n a_n$ diverges. \square

The Ratio Test, which we state and prove below, is actually weaker than the Root Test. Its proof is based on comparison with the Root Test, using the following result.

LEMMA 4.14. *Let c_n be a sequence of positive real numbers. Then*

$$\begin{aligned} \limsup_{n \rightarrow \infty} c_n^{1/n} &\leq \limsup_{n \rightarrow \infty} \frac{c_{n+1}}{c_n}, \quad \text{and} \\ \liminf_{n \rightarrow \infty} c_n^{1/n} &\geq \liminf_{n \rightarrow \infty} \frac{c_{n+1}}{c_n}. \end{aligned}$$

PROOF. We prove the lim sup inequality, and leave the similar lim inf case as an exercise. Let $\gamma = \limsup_n \frac{c_{n+1}}{c_n}$. If $\gamma = +\infty$, there is nothing to prove, since every extended real number x satisfies $x \leq +\infty$. So, assume $\gamma \in \mathbb{R}$. Then we can choose some $\beta > \gamma$, and as in the proof of the Root Test above, it follows that $\frac{c_{n+1}}{c_n} < \beta$ for all sufficiently large n , say $n \geq N$. But then, by induction, we have

$$\frac{c_{N+k}}{c_N} = \frac{c_{N+k}}{c_{N+k-1}} \frac{c_{N+k-1}}{c_{N+k-2}} \cdots \frac{c_{N+1}}{c_N} < \beta^k.$$

Thus, for $n \geq N$, letting $k = n - N$, we have

$$c_n = c_{N+k} < C_N \beta^k = c_N \beta^{n-N} = (c_N \beta^{-N}) \cdot \beta^n$$

and so

$$c_n^{1/n} < (c_N \beta^{-N})^{1/n} \cdot \beta.$$

From the Squeeze Theorem, it follows that

$$\limsup_{n \rightarrow \infty} c_n^{1/n} \leq \limsup_{n \rightarrow \infty} (c_N \beta^{-N})^{1/n} \cdot \beta = \beta \cdot \lim_{n \rightarrow \infty} (c_N \beta^{-N})^{1/n} = \beta.$$

(Here we have used the fact that $p = c_N \beta^{-N}$ is a positive constant, and $\lim_n p^{1/n} = 1$ for any $p > 0$; this last well-known limit can be found as Theorem 3.20(b) in Rudin.) Thus, for any $\beta > \gamma$, we have $\limsup_n c_n^{1/n} \leq \beta$. It follows that $\limsup_n c_n^{1/n} \leq \gamma$, as claimed. \square

THEOREM 4.15 (Ratio Test). *Let (a_n) be a sequence in \mathbb{C} .*

- (1) *If $\limsup_{n \rightarrow \infty} \left| \frac{a_{n+1}}{a_n} \right| < 1$, then $\sum_n a_n$ converges.*
- (2) *If $\liminf_{n \rightarrow \infty} \left| \frac{a_{n+1}}{a_n} \right| > 1$, then $\sum_n a_n$ diverges.*

PROOF. For (1): from Lemma 4.14, $\limsup_n |a_n|^{1/n} \leq \limsup_n \frac{|a_{n+1}|}{|a_n|} < 1$, and so by the Root Test, $\sum_n a_n$ converges. For (2): from Lemma 4.14, $\limsup_n |a_n|^{1/n} \geq \liminf_n |a_n|^{1/n} \geq \liminf_n \frac{|a_{n+1}|}{|a_n|} > 1$, and so by the Root Test, $\sum_n a_n$ diverges. \square

REMARK 4.16. Once again, if the \limsup or \liminf of the ratio of successive terms = 1, the test cannot give any information: letting $a_n = \frac{1}{n}$ and $b_n = \frac{1}{n^2}$, in both cases we have $\lim_n \left| \frac{a_{n+1}}{a_n} \right| = \lim_n \left| \frac{b_{n+1}}{b_n} \right| = 1$, and yet $\sum_n a_n$ diverges while $\sum_n b_n$ converges.

EXAMPLE 4.17. Consider the sequence $(a_n) = (\frac{1}{2}, \frac{1}{3}, \frac{1}{2^2}, \frac{1}{3^2}, \dots)$. That is: $a_{2n-1} = \frac{1}{2^n}$ and $a_{2n} = \frac{1}{3^n}$ for $n \geq 1$. Thus $|a_{2n-1}|^{1/(2n-1)} = (\frac{1}{2})^{n/(2n-1)} \rightarrow \frac{1}{\sqrt{2}}$ while $|a_{2n}|^{1/2n} = (\frac{1}{3})^{n/2n} = \frac{1}{\sqrt{3}}$. Thus $\limsup_n |a_n|^{1/n} = \frac{1}{\sqrt{2}}$, and so by the Root Test, the series $\sum_n a_n$ converges. But the Ratio Test is no use here. Note that

$$\begin{aligned} \frac{a_{(2n-1)+1}}{a_{2n-1}} &= \frac{1/3^n}{1/2^n} = \left(\frac{2}{3}\right)^n \rightarrow 0, \\ \frac{a_{2n+1}}{a_{2n}} &= \frac{1/2^{n+1}}{1/3^n} = \frac{1}{2} \left(\frac{3}{2}\right)^n \rightarrow +\infty. \end{aligned}$$

Thus $\limsup_n \frac{a_{n+1}}{a_n} = +\infty > 1$ while $\liminf_n \frac{a_{n+1}}{a_n} = 0 < 1$; so the Ratio Test gives no information.

REMARK 4.18. You may remember the Ratio and Root Tests as being described as equivalent in your calculus class. This is only true if you restrict to the case when $\lim_n \left| \frac{a_{n+1}}{a_n} \right|$ exists. In this case, the limit is equal to both the \liminf and the \limsup , and then Lemma 4.14 shows that $\lim_n |a_n|^{1/n}$ also exists. But this rules out series like the one above, that somehow “alternate” between different kinds of terms, all of which are shrinking fast enough for the series to converge.

EXAMPLE 4.19 (The number e). Consider the series

$$\sum_{n=0}^{\infty} \frac{1}{n!}.$$

Note that the sequence of terms a_n satisfies $\frac{a_{n+1}}{a_n} = \frac{1/(n+1)!}{1/n!} = \frac{1}{n+1} \rightarrow 0$ as $n \rightarrow \infty$, and so by the ratio test the series converges. Its exact value is called e . It is sometimes called *Napier's constant*, since it was first alluded to in a table of logarithms in an appendix of a book written by the Scottish mathematician / natural philosopher John Napier, circa 1618. It was first directly studied by Jacob Bernoulli, who used the letter b to denote it. But, like everything else from that era, it was eventually Euler who proved much of what we know about it, and Euler called it e .

The approximate value is

$$e \approx 2.71828182845904523536028747135266249775724709369995.$$

Fun fact: when Google went public in 2004, their IPO (initial public offering) was \$2,718,281,828. *Nerrrrrrds*.

LEMMA 4.20. *The number e is given by $e = \lim_{n \rightarrow \infty} \left(1 + \frac{1}{n}\right)^n$.*

PROOF. Let $s_n = \sum_{k=0}^n \frac{1}{k!}$ be the n th partial sum of the series defining e . Let $t_n = \left(1 + \frac{1}{n}\right)^n$. Now, for fixed m , we can use the binomial theorem to expand

$$\left(1 + \frac{1}{n}\right)^m = \sum_{k=0}^m \binom{m}{k} \frac{1}{n^k} = \sum_{k=0}^m \frac{m(m-1)\cdots(m-k+1)}{k!} \cdot \frac{1}{n^k}.$$

Write the k th term as

$$\frac{1}{k!} \cdot \frac{m}{n} \cdot \frac{m-1}{n} \cdots \frac{m-k+1}{n}.$$

Thus, specializing to the case $m = n$, we have

$$t_n = \left(1 + \frac{1}{n}\right)^n = \sum_{k=0}^n \frac{1}{k!} \cdot 1 \cdot \frac{n-1}{n} \cdots \frac{n-k+1}{n} < \sum_{k=0}^n \frac{1}{k!} = s_n.$$

We have yet to prove that $\lim_{n \rightarrow \infty} t_n$ exists, but since s_n converges to the finite number e , it follows from $t_n < s_n$ that t_n is bounded above, and so $\limsup_n t_n$ exists, and (by HW5)

$$\limsup_{n \rightarrow \infty} t_n \leq \limsup_{n \rightarrow \infty} s_n = e.$$

Now, on the other hand, let m be fixed. Then for $n \geq m$,

$$t_n = \sum_{k=0}^n \frac{1}{k!} \cdot 1 \cdot \frac{n-1}{n} \cdots \frac{n-k+1}{n} \geq \sum_{k=0}^m \frac{1}{k!} \cdot 1 \cdot \frac{n-1}{n} \cdots \frac{n-k+1}{n} \equiv t_n^m.$$

So, for fixed m , the two sequences (t_n) and (t_n^m) are comparable: $t_n \geq t_n^m$. Again by HW5, and using the limit theorems (for the finite sum with m terms), we have

$$\liminf_{n \rightarrow \infty} t_n \geq \liminf_{n \rightarrow \infty} t_n^m = \sum_{k=0}^m \frac{1}{k!} \lim_{n \rightarrow \infty} \frac{n-1}{n} \cdots \frac{n-k+1}{n} = \sum_{k=0}^m \frac{1}{m!} = s_m.$$

As this holds true for every m , it follows from the squeeze theorem that

$$\liminf_{n \rightarrow \infty} t_n \geq \lim_{m \rightarrow \infty} s_m = e.$$

Thus

$$e \leq \liminf_{n \rightarrow \infty} t_n \leq \limsup_{n \rightarrow \infty} t_n \leq e$$

which implies that the lim sup and lim inf are both equal to e , as claimed. \square

The second form of e , as a limit, is (one of the) reason(s) it is so important: this shows that e shows up in many problems related to compound interest or exponential decay. However, as a means of approximating e , this limit is *very* slow: for example

$$t_{10} \approx 2.5937 \text{ (4.6\% error)}, \quad t_{100} \approx 2.7048 \text{ (0.50\% error)}.$$

That is: you need 100 terms in order to get 2 digits of accuracy. On the other hand, s_{10} is within 10^{-8} of e , and s_{100} is so close to e my computer cannot compute the difference. But we can give a bound on this tiny error as follows. First note that s_n is increasing, so $|e - s_n| = e - s_n$. Now

$$e - s_n = \sum_{k=0}^{\infty} \frac{1}{k!} - \sum_{k=0}^n \frac{1}{k!} = \sum_{k=n+1}^{\infty} \frac{1}{k!}.$$

Now, we factor the terms (all of which have $k \geq n+1$) as

$$\frac{1}{k!} = \frac{1}{(n+1)!(n+2)(n+3)\cdots k} = \frac{1}{(n+1)!} \cdot \frac{1}{(n+2)(n+3)\cdots k} < \frac{1}{(n+1)!} \cdot \frac{1}{(n+1)^{k-n-1}}.$$

Thus

$$e - s_n < \frac{1}{(n+1)!} \sum_{k=n+1}^{\infty} \frac{1}{(n+1)^{k-n-1}} = \frac{1}{(n+1)!} \sum_{j=0}^{\infty} \frac{1}{(n+1)^j}.$$

This is a geometric series, and $0 < \frac{1}{n+1} < 1$, so we know the sum is

$$\sum_{j=0}^{\infty} \frac{1}{(n+1)^j} = \frac{1}{1 - \frac{1}{n+1}} = \frac{n+1}{n}.$$

Thus, we have our estimate:

$$e - s_n < \frac{1}{(n+1)!} \frac{n+1}{n} = \frac{1}{n! \cdot n}.$$

This is a *tiny* number. Since $10! = 3,628,800$, this shows that $e - s_{10} < \frac{1}{3 \times 10^7}$ (and in fact it's 3 times smaller than this). For $n = 100$, we have $100! \cdot 100 \approx 10^{160}$, so s_{100} differs from e only after the 160th decimal digit!

This is one of the rare occasions where a perfectly practical question of error approximation actually allows us to prove something entirely theoretical.

PROPOSITION 4.21. *The number e is irrational.*

PROOF. For a contradiction, let us suppose $e \in \mathbb{Q}$. Since $e > 0$, this means there are positive integers m, n so that $e = \frac{m}{n}$. Now, from the above estimate, we have

$$0 < e - s_n < \frac{1}{n! \cdot n}, \quad \therefore 0 < n!e - n!s_n < \frac{1}{n}.$$

Now,

$$n!s_n = n! \sum_{k=0}^n \frac{1}{k!} = \sum_{k=0}^n \frac{n!}{k!} = \sum_{k=0}^n n(n-1)\cdots(n-k+1) \in \mathbb{N}.$$

Also, by assumption $e = \frac{m}{n}$, and so $n!e = m \cdot (n-1)! \in \mathbb{N}$. Thus $\ell = n!e - n!s_n \in \mathbb{N}$. But this means $0 < \ell < \frac{1}{n}$ for some $n \in \mathbb{N}$, and that is a contradiction (there are no integers between 0 and $\frac{1}{n}$). \square

Moving to our final topic on the subject of series, let us consider *absolute convergence*.

DEFINITION 4.22. Let (a_n) be a sequence in \mathbb{C} . We say that the series $\sum_{n=1}^{\infty} a_n$ **converges absolutely** if, in fact, $\sum_{n=1}^{\infty} |a_n|$ converges.

LEMMA 4.23. If $\sum_{n=1}^{\infty} a_n$ converges absolutely, then it converges.

PROOF. This follows immediately from the Cauchy criterion. Fix $\epsilon > 0$ and choose $N \in \mathbb{N}$ large enough that, for $m > n \geq N$, $\sum_{k=n+1}^m |a_k| < \epsilon$. Then by the triangle inequality

$$\left| \sum_{k=n+1}^m a_k \right| \leq \sum_{k=n+1}^m |a_k| < \epsilon$$

and so $\sum_{n=1}^{\infty} a_n$ converges. □

3. Lecture 13: February 16, 2016

The converse of Lemma 4.23 is quite false. To see why, let us study one particular class of real series known as *alternating series*.

PROPOSITION 4.24 (Alternating Series). *Let $a_n \geq 0$ be a monotone decreasing sequence with limit $a_n \rightarrow 0$. Then $\sum_{n=1}^{\infty} (-1)^{n-1} a_n = a_1 - a_2 + a_3 - a_4 + \cdots$ converges.*

PROOF. Fix $m > n \in \mathbb{N}$ and consider the tail sum

$$|(-1)^n a_{n+1} + (-1)^{n+1} a_{n+2} + \cdots + (-1)^{m-1} a_m| = |a_{n+1} - a_{n+2} + \cdots \pm a_m|.$$

We consider two cases: either $m - n$ is even or odd. If it is even, then we can group the terms as

$$|(a_{n+1} - a_{n+2}) + \cdots + (a_{m-1} - a_m)| = (a_{n+1} - a_{n+2}) + \cdots + (a_{m-1} - a_m),$$

where we have used the fact that $a_n \downarrow$. On the other hand, we may group terms as

$$= a_{n+1} - (a_{n+2} - a_{n+3}) - (a_{n+4} - a_{n+5}) - \cdots - a_m \leq a_{n+1}.$$

On the other hand, if $n - m$ is odd, then by similar reasoning

$$\left| \sum_{k=n+1}^m (-1)^{k-1} a_k \right| = (a_{n+1} - a_{n+2}) + \cdots + (a_{m-2} - a_{m-1}) + a_m$$

and we may group this as

$$= a_{n+1} - (a_{n+2} - a_{n+3}) - \cdots - (a_{m-1} - a_m) \leq a_{n+1}.$$

Hence, in all cases, we have

$$\left| \sum_{k=n+1}^m (-1)^{k-1} a_k \right| \leq a_{n+1}.$$

Thus, fix $\epsilon > 0$. Since $a_n \rightarrow 0$, we may choose $N \in \mathbb{N}$ so that, for $n \geq N$, $a_n = |a_n| < \epsilon$. Since $a_{n+1} \leq a_n$, we therefore have $|\sum_{k=n+1}^m (-1)^{k-1} a_k| \leq a_{n+1} < \epsilon$ whenever $m > n \geq N$, which verifies the Cauchy criterion showing that $\sum_{n=1}^{\infty} (-1)^{n-1} a_n$ converges. \square

EXAMPLE 4.25. The sequence $a_n = \frac{1}{n}$ is positive, decreasing, and satisfies $a_n \rightarrow 0$. Therefore, by Proposition 4.24,

$$\sum_{n=1}^{\infty} a_n = 1 - \frac{1}{2} + \frac{1}{3} - \frac{1}{4} + \cdots$$

converges. (Remembering your calculus, it in fact converges to $\ln 2$.) This is known as the **alternating harmonic series**. Note that the absolute series $\sum_{n=1}^{\infty} \frac{1}{n}$ diverges. So this is an example of a series that is convergent but not absolutely convergent. These are sometimes called **conditionally convergent** series.

Conditionally convergent series have strange properties, particularly with regard to *rearrangements*. That is: suppose we reorder the terms. Continuing Example 4.25, rearrange the terms as follows.

$$1 - \frac{1}{2} - \frac{1}{4} + \frac{1}{3} - \frac{1}{6} - \frac{1}{8} + \frac{1}{5} - \frac{1}{10} - \frac{1}{12} + \frac{1}{7} - \cdots + \frac{1}{2n-1} - \frac{1}{4n-2} - \frac{1}{4n} + \cdots$$

Each of the terms in the alternating harmonic series appears exactly once in this sum. It is no longer alternating, so we cannot apply a theorem to tell whether it converges; but we can in fact sum it as follows: in each three-term group, simplify

$$\frac{1}{2n-1} - \frac{1}{4n-2} - \frac{1}{4n} = \left(\frac{1}{2n-1} - \frac{1}{4n-2} \right) - \frac{1}{4n} = \frac{1}{4n-2} - \frac{1}{4n} = \frac{1}{2} \left(\frac{1}{2n-1} - \frac{1}{2n} \right).$$

So, the sum of the whole rearranged series is

$$\frac{1}{2} \left(1 - \frac{1}{2} \right) + \frac{1}{2} \left(\frac{1}{3} - \frac{1}{4} \right) + \cdots + \frac{1}{2} \left(\frac{1}{2n-1} - \frac{1}{2n} \right) + \cdots = \frac{1}{2} \sum_{n=1}^{\infty} \frac{(-1)^{n-1}}{n} = \frac{1}{2} \ln 2.$$

That is: this rearrangement produces *half* the value of the original series!

This is always possible for a conditionally convergent series of real numbers. Riemann proved this: if $\sum_{n=1}^{\infty} a_n$ is conditionally convergent and $a_n \in \mathbb{R}$, then there is a rearrangement a'_n of the terms so that the sequence $s'_n = \sum_{k=1}^n a'_k$ has *any tail behavior possible*: given any α, β with $-\infty \leq \alpha \leq \beta \leq +\infty$, one can find a rearrangement so that $\limsup_n s'_n = \beta$ and $\liminf_n s'_n = \alpha$. (This is proved as Theorem 3.54 in Rudin.) Fortunately, this kind of craziness is not possible for absolutely convergent series, as our final theorem in this section attests to.

THEOREM 4.26. *Let (a_n) be a complex sequence such that $\sum_{n=1}^{\infty} |a_n|$ converges. Then for any rearrangement a'_n of a_n , $\sum_{n=1}^{\infty} a'_n = \sum_{n=1}^{\infty} a_n$.*

PROOF. Fix ϵ , and choose $N \in \mathbb{N}$ so that $\sum_{k=n+1}^m |a_k| < \epsilon$ for $m > n \geq N$. Let $s_n = \sum_{k=1}^n a_k$ and $s'_n = \sum_{k=1}^n a'_k$. The numbers $1, 2, \dots, N$ appear as indices in the rearranged sequence (a'_n) each exactly once, so there must be some finite p so that they all appear by time p in (a'_n) . Thus, for $m > n \geq p$, in the difference $s_m - s'_m$ the N terms a_1, \dots, a_N cancel leaving only with (original) indices $> N$. Thus, by the choice of N , this difference is $\leq |\sum_{k=N+1}^m a_k| \leq \sum_{k=N+1}^m |a_k| < \epsilon$. This shows that the sequence $s_n - s'_n$ converges to 0, and it follows, since we know s_n converges to $\sum_{n=1}^{\infty} a_n$, that s'_n also converges to the sum. \square

CHAPTER 5

Metric Spaces

For the remainder of this course, we are going to generalize the concepts we've worked with (notably *convergence*) beyond the case of \mathbb{R} or \mathbb{C} . The key to this generalization was already discussed in the generalization from \mathbb{R} to \mathbb{C} : we replace the absolute value in \mathbb{R} (defined in terms of the order relation) with the complex modulus in \mathbb{C} . For all of the same technology to work, only a few basic properties of the absolute value / modulus were needed: that $|x| \geq 0$, that $|x| = 0$ only when $x = 0$, and finally the triangle inequality $|x + y| \leq |x| + |y|$.

This last property requires a notion of addition, and we'd like to move beyond vector spaces. The trick is that, in the notion of convergence, the absolute value / modulus only ever comes up as a means of measuring *distance* between two elements: $|x - y|$. Thinking of it this way, what do the three key properties say?

- For any two elements x, y , the distance $|x - y|$ is ≥ 0 .
- If the distance $|x - y|$ is 0, then actually $x = y$.
- The triangle inequality: for any *three* elements x, y, z , the distance $|x - z|$ is bounded above by $|x - y| + |y - z|$. (This is really why it's called the triangle inequality: draw the associated picture.) Indeed, we have

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

Interpreted in this light, we don't need a notion of addition: everything can be stated purely in terms of the notion of distance (in this case given by $(x, y) \mapsto |x - y|$). We generalize thus.

DEFINITION 5.1. *Let X be a nonempty set. A function $d: X \times X \rightarrow \mathbb{R}$ is called a **metric** if it satisfies the following three properties.*

- (1) For any $x, y \in X$, $d(x, y) = d(y, x) \geq 0$.
- (2) For any $x, y \in X$, if $d(x, y) = 0$, then $x = y$.
- (3) For any $x, y, z \in X$, $d(x, z) \leq d(x, y) + d(y, z)$.

The pair (X, d) is called a **metric space**.

EXAMPLE 5.2. (1) As above, if we let $d_{\mathbb{C}}(x, y) = |x - y|$, then $(\mathbb{C}, d_{\mathbb{C}})$ is a metric space. Same goes for \mathbb{R} equipped with the restriction of $d_{\mathbb{C}}$ to \mathbb{R} .

- (2) More generally, fix n , and consider the set \mathbb{C}^n of n -tuples of real numbers. Define the **Euclidean norm** on \mathbb{C}^n as follows:

$$\|(x_1, \dots, x_n)\|_2 = (|x_1|^2 + \dots + |x_n|^2)^{1/2}.$$

It is a simple but laborious exercise to verify that the **Euclidean metric** $d_2(\mathbf{x}, \mathbf{y}) = \|\mathbf{x} - \mathbf{y}\|_2$ is a metric on \mathbb{C}^n . As above, the restriction to \mathbb{R}^n is also a metric.

- (3) There are many other, *different* metrics on \mathbb{R}^n . The best known are the p -metrics: for $1 \leq p < \infty$,

$$\|(x_1, \dots, x_n)\|_p = (|x_1|^p + \dots + |x_n|^p)^{1/p}.$$

There is also the ∞ -norm, aka the sup norm

$$\|(x_1, \dots, x_n)\|_\infty = \max\{|x_1|, \dots, |x_n|\}.$$

As above, all of these norms yield metrics in the usual way, $d_p(\mathbf{x}, \mathbf{y}) = \|\mathbf{x} - \mathbf{y}\|_p$. Note: the definition still makes sense when $p < 1$, but it no longer gives a metric: the triangle inequality is violated. For example, taking $p = \frac{1}{2}$, we have

$$\begin{aligned} \|(9, 1) + (16, 0)\|_{1/2} &= \|(25, 1)\|_{1/2} = (25^{1/2} + 1^{1/2})^2 = 36 \\ \|(9, 1)\|_{1/2} + \|(16, 0)\|_{1/2} &= (9^{1/2} + 1^{1/2})^2 + (16^{1/2} + 0^{1/2})^2 = 32 < 36. \end{aligned}$$

- (4) Let $B[0, 1]$ consist of all bounded functions $[0, 1] \rightarrow \mathbb{R}$. Then define a function $d_u: B[0, 1] \times B[0, 1] \rightarrow \mathbb{R}$ by

$$d_u(f, g) = \sup_{x \in [0, 1]} |f(x) - g(x)|.$$

This is well-defined: since f and g are bounded, the set $\{f(x) - g(x) : x \in [0, 1]\}$ is a bounded, nonempty set, so it has a sup. It is ≥ 0 , and moreover if $d_u(f, g) = 0$, then for every $x_0 \in [0, 1]$, $|f(x_0) - g(x_0)| \leq \sup_{x \in [0, 1]} |f(x) - g(x)| = 0$, which implies that $f(x_0) - g(x_0) = 0$ — i.e. $f = g$. This verifies the first two properties of Definition 5.1. For the triangle inequality, we have

$$\begin{aligned} d_u(f, h) &= \sup_{x \in [0, 1]} |f(x) - h(x)| = \sup_{x \in [0, 1]} |f(x) - g(x) + g(x) - h(x)| \\ &\leq \sup_{x \in [0, 1]} (|f(x) - g(x)| + |g(x) - h(x)|) \\ &\leq \sup_{x \in [0, 1]} |f(x) - g(x)| + \sup_{x \in [0, 1]} |g(x) - h(x)| \\ &= d_u(f, g) + d_u(g, h) \end{aligned}$$

using the properties of sup we now know well. Thus the triangle inequality holds for d_u as well, and so it is a metric. Note, like the above examples, it has the form $d_u(f, g) = \|f - g\|_u$ for a *norm* $\|\cdot\|_u$: a function on $B[0, 1]$ which has the properties $\|f\|_u \geq 0$ and $= 0$ only if $f = 0$, and satisfies the triangle inequality $\|f + g\|_u \leq \|f\|_u + \|g\|_u$. Whenever we have a function like this defined on a vector space, it gives rise to a metric by subtraction.

- (5) Not every metric is given in terms of a norm like this. For example, consider on \mathbb{R} the function

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

It is easy to verify that this satisfies properties (1) and (2) in Definition 5.1. The triangle inequality is also easy to see, by breaking into eight cases (depending whether $|x - y|$, $|x - z|$, and $|y - z|$ are ≤ 1 or > 1); this boring proof is left to the reader.

- (6) Given any nonempty set X , one can define a metric on X by the silly rule

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

This is known as the **discrete metric**. It says two points are close only if they are equal; otherwise they are far apart. It is again simple to verify this is a metric.

One important observation was made at several points in the examples: if (X, d) is a metric space, and $Y \subseteq X$, then $(Y, d|_Y)$ is a metric space – that is, the metric d defined on all pairs $(x, y) \in X \times X$, also defines a metric when restricted only to pairs in $Y \times Y$, as is straightforward to verify. Thus, the Euclidean metric on \mathbb{C}^n automatically gives us a metric (also called the Euclidean metric) on \mathbb{R}^n . Similarly, the usual metric on \mathbb{R} restricts to a metric on $[0, 1]$.

Usually thinking of metric spaces using our intuition from \mathbb{R}^2 and \mathbb{R}^3 , we introduce the following notation.

DEFINITION 5.3. Let (X, d) be a metric space, and let $x_0 \in X$. For a fixed $r > 0$, the **ball** of radius r centered at x_0 , denoted $B_r(x_0)$, is the set

$$B_r(x_0) = \{x \in X : d(x_0, x) < r\}.$$


(Rudin calls this a **neighborhood** $N_r(x_0)$.) With $r = 1$, we refer to this as the **unit ball** centered at x_0 .

EXAMPLE 5.4. (1) In \mathbb{R}^n , using the definition of the Euclidean metric (and choosing the base point $\mathbf{0}$ to simplify things), we have

$$B_r(\mathbf{0}) = \{(x_1, \dots, x_n) : x_1^2 + \dots + x_n^2 < r^2\}$$

which is what we usually know as a ball (in n -dimensions).

(2) Consider (\mathbb{R}^2, d_p) , with the p -metric of Example 5.2(3). Here are some pictures of the unit ball:



(3) In a discrete metric space (X, d) as in Example 5.2(6), $B_r(x_0) = X$ if $r > 1$, and $B_r(x_0) = \{x_0\}$ if $r \leq 1$.

1. Lecture 14: February 18, 2016

Now, once more, we define **convergence** and **Cauchy** in the wider world of metric spaces.

DEFINITION 5.5. Let (X, d) be a metric space, and let (x_n) be a sequence in X , and let $x \in X$. Say that (x_n) **converges to** x , or $x_n \rightarrow x$, or $\lim_{n \rightarrow \infty} x_n = x$, if

$$\forall \epsilon > 0 \exists N \in \mathbb{N} \text{ s.t. } \forall n \geq N \ d(x_n, x) < \epsilon.$$

In other words, $x_n \rightarrow x$ means that the real sequence $d(x_n, x) \rightarrow 0$. Alternatively, we could state this as: for all sufficiently large n , $x_n \in B_\epsilon(x)$.

Similarly, say that (x_n) is a **Cauchy sequence** in X if

$$\forall \epsilon > 0 \exists N \in \mathbb{N} \text{ s.t. } \forall n, m \geq N \ d(x_n, x_m) < \epsilon.$$

As discussed in the generalization from \mathbb{R} to \mathbb{C} , limits are unique: if $x_n \rightarrow x$ and $x_n \rightarrow y$, then $x = y$ (this follows from the fact that $d(x, y) = 0$ implies that $x = y$; it is primarily for this reason that this *non-degeneracy* property is required in the definition of a metric).

EXAMPLE 5.6. Consider a discrete metric space: (X, d) where d is given as in Example 5.2(6). Let $x_n \rightarrow x$. In particular, this means that there is some time N such that, for $n \geq N$, $d(x_n, x) < \frac{1}{2}$. But by the definition of d , either $d(x_n, x) = 0$ or $d(x_n, x) = 1$; so if $d(x_n, x) < \frac{1}{2}$ then $d(x_n, x) = 0$, and so $x_n = x$. Thus, if $x_n \rightarrow x$, then $x_n = x$ for all large n . In a discrete metric space, *convergence* is the same thing as *eventually constant*. The same holds true for Cauchy.

In general, there is a fundamental difference between convergent sequences that are eventually constant and convergent sequences that are *not* eventually constant. We use this difference to define one of the most important topological concepts.

DEFINITION 5.7. Let (X, d) be a metric space, and let $E \subseteq X$ be a subset. A point $x \in X$ (not necessarily in E) is called a **limit point** of E if there is a sequence $x_n \in E \setminus \{x\}$ that converges to x , $x_n \rightarrow x$. That is: a limit point of E is a limit of some not eventually constant sequence in E . A point $e \in E$ that is not a limit point of E is called an **isolated point** of E .

EXAMPLE 5.8. In \mathbb{R} with the usual metric, take $E = (-1, 0] \cup \mathbb{N}$. Then -1 is a limit point of E : for example, $-1 = \lim(-1 + \frac{1}{n})$ and $-1 + \frac{1}{n} \in E$ for each n . Also, any point $x \in E$ is a limit point of E : take $x_n = x - \frac{1+x}{n}$ as the sequence. This is in E since $1+x > 0$ and so $x - \frac{1+x}{n} < x \leq 0$, but also $x - \frac{1+x}{n} > x - (1+x) = -1$.

On the other hand, the positive integers \mathbb{N} are isolated points of E . For example, consider 1. If y_n is any sequence in \mathbb{R} that converges to 1, then we must have $y_n \in (0.9, 1.1)$ for all large n ; but then if $y_n \in E$ it follows that $y_n = 1$ for all large n , which isn't allowed. Thus, no sequence in $E \setminus \{1\}$ converges to 1, showing that the point $1 \in E$ is not a limit point of E – it is an isolated point.

The set of all limit points of a set E is denoted E' . So E is closed iff $E' \subseteq E$.

DEFINITION 5.9. A subset E of a metric space is called **closed** if it contains all of its limit points.

EXAMPLE 5.10. (1) The set $E = (-1, 0] \cup \mathbb{N}$ from Example 5.8 is not closed: -1 is a limit point of E , but $-1 \notin E$.
 (2) The set $F = [-1, 0] \cup \mathbb{N}$ is closed. The argument in Example 5.8 shows that each of the points in $[-1, 0]$ is a limit point of F , while each of the points $n \in \mathbb{N}$ is an isolated point

of F . On the other hand, if x is a real number not in F , then either $x < -1$ or $x > 0$ and $x \notin \mathbb{N}$. In the former case, this means that no sequence in F can come within distance $1 + x > 0$ of x , and so cannot converge to x ; a similar argument with $x > 0$ shows that x is not a limit point of F . Thus, the set of limit points of F consists exactly of the set $[-1, 0]$, and this set is contained in F . So F is closed.

Definition 5.9 is stated in terms of limit points to make it clear that there are two kinds of points to consider in deciding whether a set is closed: isolated points and non-isolated points. For example, if one has a closed set, then adding to it a finite collection of isolated points will preserve closedness. But for the purposes of a concise definition, one need not be concerned about the distinction.

PROPOSITION 5.11. *A subset E of a metric space (X, d) is closed if and only if, for any sequence (x_n) in E that converges in X , the limit $\lim_{n \rightarrow \infty} x_n$ is actually in E .*

That is: *closed* means *closed under limits of sequences*.

PROOF. Suppose E is closed, so $E' \subseteq E$. Now, let (x_n) be any sequence in E that converges to some point x . If $x_n \neq x$ for any n , then by definition $x \in E'$, and therefore by assumption $x \in E$. If, on the other hand, there exists n with $x_n = x$, then since $x_n \in E$ for each n , we have $x \in E$. Thus, the E is closed under limits of sequences.

Conversely, suppose E is closed under limits of sequences. Let $x \in E'$; so by definition there is a sequence $x_n \in E \setminus \{x\}$ such that $x_n \rightarrow x$. Well, since $x_n \rightarrow x$ and $x_n \in E$, by assumption $x \in E$. Thus $E' \subseteq E$, and E is closed. \square

There is a complementary notion to *closed*, called *open*.

DEFINITION 5.12. *A subset E of a metric space is called **open** if, for any point $x \in E$, there is a ball $B_r(x)$ (for some $r > 0$) with $B_r(x) \subseteq E$.*

EXAMPLE 5.13. (1) The set $E = (-1, 0] \cup \mathbb{N}$ from Example 5.8 is *not* open. Consider the point $0 \in E$. For any $0 < r < 1$, the ball $B_r(0) = (-r, r)$ contains some points (for example $\frac{r}{2}$) that are in $(0, 1)$, and hence not in E . Similarly, any of the points in \mathbb{N} are in E but none is contained in a ball contained in E . So E is not open.

(2) On the other hand, the set $U = (0, 1)$ is open. Indeed, let $x \in U$. Let's consider two cases: either $0 < x < \frac{1}{2}$ or $\frac{1}{2} \leq x < 1$. In the former case, the ball $B_x(x) = (0, 2x)$ is contained in $U = (0, 1)$; in the latter case, the ball $B_{1-x}(x) = (2x - 1, 1)$ is contained in U . So every point of x is contained in a ball inside U , showing that U is open.

(3) Let (X, d) be a discrete metric space. If $x \in X$, then by Example 5.4(3), $B_1(x) = \{x\}$. Thus, every singleton point in a discrete metric space is an open set. On the other hand, by Example 5.6, there are no non-eventually-constant sequences converging to any point x , which means every point is isolated. That is: X has no limit points, which means that (vacuously) X contains all its limit points. So X is also closed.

(4) Consider the empty set \emptyset in any metric space. It is both open and closed. Indeed, the definitions of "open" and "closed" each start with "for every point in the set..." and since there are no points in \emptyset to check the condition, it follows that the condition holds vacuously.

Example 5.13(2) has a nice, important generalization to any metric space. Not that $(0, 1)$ is itself a ball in \mathbb{R} : it is the ball $B_{1/2}(1/2)$. The fact is, any ball is open.

PROPOSITION 5.14. *Let (X, d) be a metric space, let $x \in X$, and let $r > 0$. Then the ball $B_r(x)$ is open in X .*

PROOF. Let $y \in B_r(x)$. This means $d(x, y) < r$. Hence, there is some $\epsilon > 0$ so that $d(x, y) = r - \epsilon$. I claim that $B_\epsilon(y) \subset B_r(x)$. Indeed, suppose that $z \in B_\epsilon(y)$, meaning that $d(z, y) < \epsilon$. Then

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

so $z \in B_r(x)$. We have thus shown that, for any $y \in B_r(x)$, there is a ball $B_\epsilon(y) \subset B_r(x)$. That is: $B_r(x)$ is open. \square

2. Lecture 15: February 23, 2016

We referred to *open* and *closed* as complementary properties. That doesn't mean that any set is either open or closed: for example, the set $(-1, 0]$ considered above is neither open nor closed. But they concepts are complementary, in the following precise sense.

PROPOSITION 5.15. *Let (X, d) be a metric space. A subset $E \subseteq X$ is open if and only if $E^c = X \setminus E$ is closed.*

Since $(E^c)^c = E$, it follows similarly that E is closed iff E^c is open. In the proof we will use the characterization of *closed* given in Proposition 5.11.

PROOF. Suppose E is open. Let (x_n) be a sequence in E^c that converges to some point $x \in X$. We want to show that $x \in E^c$; to produce a contradiction, we therefore assume that $x \notin E^c$, meaning $x \in E$. Since E is open, by definition there is some $r > 0$ so that $B_r(x) \subseteq E$. On the other hand, since $x_n \rightarrow x$, there is certainly some N so that $d(x_N, x) < r$. Thus $x_N \in B_r(x) \subseteq E$, which means that $x_N \in E$. But we assumed that $x_N \in E^c$, so this is a contradiction. Therefore $x \in E^c$. This shows that any convergent sequence in E^c has limit in E^c , which shows that E^c is closed.

Conversely, suppose E^c is closed. Let $x \in E$. We want to show that there is some $r > 0$ so that $B_r(x) \subseteq E$; to produce a contradiction, we therefore assume that there is no such r . That means that, for say $r = \frac{1}{n}$, $B_{1/n}(x) \not\subseteq E$, which means precisely that there is some point $x_n \notin E$ such that $x_n \in B_{1/n}(x)$. So, we have produced a sequence $x_n \in E^c$ such that $d(x_n, x) < \frac{1}{n}$, meaning that $x_n \rightarrow x$. Thus x is the limit of a sequence in E^c , and so by assumption $x \in E^c$. This contradicts the assumption that $x \in E$. Therefore there must be some $r > 0$ so that $B_r(x) \subseteq E$, and so E^c is closed. \square

Let us make a few more definitions that pertain to the local properties of open and closed sets.

DEFINITION 5.16. *Let (X, d) be a metric space, and let $E \subseteq X$.*

- (1) The **closure** of E is the set $\overline{E} = E \cup E'$.
- (2) A point $x \in E$ is called an **interior point** if there is some $r > 0$ with $B_r(x) \subseteq E$. The set of all interior points of E is called the **interior** of E , and is denoted $\overset{\circ}{E}$.
- (3) The **boundary** of E is the set $\partial E = \overline{E} \setminus \overset{\circ}{E}$.

REMARK 5.17. Following the proof of Proposition 5.11, \overline{E} can alternatively be described as the set of all limits of convergent sequences in E .

EXAMPLE 5.18. Consider again the set $E = (-1, 0] \cup \mathbb{N}$ in \mathbb{R} , considered in Examples 5.8 and 5.10. We've shown that the points in $(-\infty, -1)$ and $(0, \infty)$ are not limit points (the points $1, 2, 3, \dots$ are in E but are isolated); on the other hand, we've shown that the points $[-1, 0]$ are all limit points. Thus $E' = [-1, 0]$, and so $\overline{E} = E \cup E' = [-1, 0] \cup \mathbb{N}$. We've also shown in Example 5.13(1) that there is no ball centered at 0 contained in E ; similarly, there are no balls centered at the points $1, 2, 3, \dots$ contained in E , so these are not interior points. On the other hand, an argument very similar to Example 5.13(2) shows that the points in $(-1, 0)$ are interior points. So $\overset{\circ}{E} = (-1, 0)$. Finally, this shows that $\partial E = \overline{E} \setminus \overset{\circ}{E} = \{-1, 0, 1, 2, \dots\}$.

EXAMPLE 5.19. Let \mathbb{Q} denote the rational numbers as a subset of the metric space \mathbb{R} . By Theorem 1.17(2) (the density of \mathbb{Q} in \mathbb{R}), given any real numbers $a < b$ there is a rational number $q \in \mathbb{Q}$ with $a < q < b$. In particular, fix $x \in \mathbb{R}$; then for $n \in \mathbb{N}$ there is a rational number

q_n with $x + \frac{1}{2n} < q_n < x + \frac{1}{n}$. In particular, we have $\frac{1}{2n} < |q_n - x| < \frac{1}{n}$. This shows that $q_n \rightarrow x$ but $q_n \neq x$ for any n ; thus $x \in \mathbb{Q}'$. So every real number is a limit point of \mathbb{Q} , and so $\overline{\mathbb{Q}} = \mathbb{Q} \cup \mathbb{Q}' = \mathbb{Q} \cup \mathbb{R} = \mathbb{R}$. (This is another way of saying \mathbb{Q} is dense in \mathbb{R} ; in general, we say a subset $E \subseteq X$ is **dense** in a metric space X if $\overline{E} = X$.)

On the other hand, let $r > 0$, and let $q \in \mathbb{Q}$. The number $x = q + \frac{r}{\sqrt{2}}$ is $< x + r$, which shows that $x \in (q - r, q + r) = B_r(q)$. But $x \notin \mathbb{Q}$: indeed, we can solve $\sqrt{2} = \frac{r}{x - q}$, and so if x were rational $\sqrt{2}$ would also be rational, which we know it is not. Thus $B_r(q) \not\subseteq \mathbb{Q}$ for any $r > 0$. This shows q is not interior to \mathbb{Q} . This holds for any $q \in \mathbb{Q}$, and so, in fact, $\overset{\circ}{\mathbb{Q}} = \emptyset$.

THEOREM 5.20. *Let (X, d) be a metric space, and let $E \subseteq X$.*

(1) \overline{E} is closed; E is closed iff $E = \overline{E}$.

(2) $\overset{\circ}{E}$ is open; E is open iff $E = \overset{\circ}{E}$.

PROOF. We begin with item 1. Let (x_n) be a sequence in \overline{E} with limit x . We wish to show $x \in \overline{E}$. If $x \in E \subseteq \overline{E}$ we are done, so assume $x \notin E$. For each x_n , either $x_n \in E$ or $x_n \in E'$. In the latter case, by definition of E' there is some other sequence $y_k \in E$ such that $y_k \rightarrow x_n$; in particular, we can choose some k large enough that $d(y_k, x_n) < \frac{1}{n}$. So, we can define a new sequence (x'_n) as follows: if $x_n \in E$ then $x'_n = x_n$; if $x_n \notin E'$, then $x'_n = y_k$ as above, so $x'_n \in E$ and $d(x_n, x'_n) < \frac{1}{n}$. Then we have $d(x'_n, x) \leq d(x'_n, x_n) + d(x_n, x) < \frac{1}{n} + d(x_n, x) \rightarrow 0$, and so $x'_n \rightarrow x$. As $x'_n \in E$ and $x \notin E$, it follows that x is a limit point of E , and so $x \in E' \subseteq E \cup E' = \overline{E}$. Thus \overline{E} is closed under limits; by Proposition 5.11, it follows that \overline{E} is closed, as claimed. Now, by definition E is closed iff $E' \subseteq E$, and this happens iff $\overline{E} = E \cup E' = E$, proving the second point.

For item 2, let $x \in \overset{\circ}{E}$; thus, there is some ball $B_r(x)$ contained in E . But by Proposition 5.14, the ball $B_r(x)$ is open, which means all its points are interior points; thus $B_r(x) \subseteq \overset{\circ}{E}$. So, any point in $\overset{\circ}{E}$ is interior to $\overset{\circ}{E}$, which shows that $\overset{\circ}{E}$ is open. By definition $\overset{\circ}{E} \subseteq E$ for any set E ; thus $\overset{\circ}{E} = E$ iff $E \subseteq \overset{\circ}{E}$, which is the statement that every point of E is an interior point, which is precisely the definition of E being open. \square

EXAMPLE 5.21. Let $E \subset \mathbb{R}$ be nonempty and bounded above. Then $\alpha = \sup E$ exists. By definition, $\alpha - \frac{1}{n}$ is not an upper bound for E for any $n \in \mathbb{N}$, which shows that there is an element $x_n \in E$ with $\alpha - \frac{1}{n} < x_n \leq \alpha$. This shows that $x_n \rightarrow \alpha$. By Remark 5.17, it follows that $\alpha \in \overline{E}$: the supremum is always in the closure. On the other hand, if there were some $r > 0$ with $B_r(\alpha) \subseteq E$, then, for example, $\alpha + \frac{r}{2} \in E$. Since $\alpha + \frac{r}{2} > \alpha$, this contradicts α being an upper bound for E . Thus, α is not in $\overset{\circ}{E}$. That is: $\sup E \in \overline{E} \setminus \overset{\circ}{E} = \partial E$.

3. Lecture 16: February 25, 2016

Now we come to an important concept you may not have encountered before: *compactness*.

DEFINITION 5.22. Let (X, d) be a metric space. A subset $K \subseteq X$ is called **compact** if every sequence (x_n) in K has a convergent subsequence whose limit is in K .

- EXAMPLE 5.23. (1) Let $a < b$ be real numbers, and consider the set $K = [a, b]$. The Bolzano-Weierstrass Theorem for \mathbb{R} (Theorem 2.25) is precisely the statement that $[a, b]$ is compact.
- (2) On the other hand, $E = [a, b)$ is not compact: the sequence $x_n = b - \frac{b-a}{n}$ is in E , but converges to $b \notin E$, therefore all of its subsequences converge to b , and hence none of them converge in E . Similarly, an unbounded interval like $[0, \infty)$ is not compact: for example the sequence $x_n = n$ has no convergent subsequences at all.
- (3) Let (X, d) be a discrete metric space. If K is a *finite* subset of X , say $K = \{y_1, y_2, \dots, y_m\}$, then K is compact. Indeed, if (x_n) is any sequence in K , then there must be some (perhaps many) y_j so that $x_n = y_j$ for infinitely many n (by the pigeonhole principle). That means exactly that there is an increasing sequence n_k with $x_{n_k} = y_j$ for all k , which means $x_{n_k} \rightarrow y_j \in K$. Thus K is compact. On the other hand if $E \subseteq X$ is infinite, it is not compact: for then we can find an infinite sequence $x_1, x_2, x_3, \dots \in E$ all distinct. Thus, any subsequence also has all distinct terms, which means it is not eventually constant. By Example 5.6, this means no subsequence converges.

Now, there is an alternate definition of compactness which is the only one used in Rudin; we refer to it as **topological compactness**, given in Definition 5.24 below. First, let us highlight the fact that Definition 5.22 was the *original* definition of compact, and predated the so-called “modern” definition by almost a century. Bolzano was already using our definition of compactness in 1817, although it would not be until 1906 that Definition 5.22 was written down formally (by Fréchet). It was around this time that Lebesgue proved (as a useful lemma) that Definition 5.24 also characterizes compactness; indeed, as we will see, it is a very useful tool. Much later, in 1929, the Russian school (led by Alexandrov and Urysohn) *redefined* compactness as what we are calling topological compactness. Our definition of the word *compact* is now often called **sequentially compact**.

DEFINITION 5.24. Let (X, d) be a metric space. Let $K \subseteq X$ be a subset. An **open cover** of K is a collection (finite or infinite) of open set \mathcal{C} in X such that every point in K is in at least one $U \in \mathcal{C}$: that is $K \subseteq \bigcup \mathcal{C}$. We call K **topologically compact** if, given any open cover \mathcal{C} of K , there is a finite sub cover: that is, there are finitely many $U_1, \dots, U_m \in \mathcal{C}$ such that $K \subseteq U_1 \cup \dots \cup U_m$.

EXAMPLE 5.25. Consider the interval $(0, 1]$. We have already seen this is not compact. It is also not topologically compact. Indeed, consider the sets $U_n = (\frac{1}{n}, 2)$ for $n \in \mathbb{N}$. If $x \in (0, 1)$ then $x > 0$ and so there is some $n \in \mathbb{N}$ with $\frac{1}{n} < x$. Therefore $x \in (\frac{1}{n}, 1) \subset (\frac{1}{n}, 2) = U_n$. This shows that the collection $\mathcal{C} = \{U_n : n \in \mathbb{N}\}$ is an open cover of $(0, 1]$. Now, consider *any* finite collection of sets from \mathcal{C} : $U_{n_1}, U_{n_2}, \dots, U_{n_k}$ for some $k \in \mathbb{N}$. Note that $\frac{1}{m} < \frac{1}{\ell}$ when $m > \ell$, and so $U_\ell \subset U_m$ in this case. What that means is that, if we let $m = \max\{n_1, \dots, n_k\}$ then $U_{n_1} \cup \dots \cup U_{n_k} = U_m = (\frac{1}{m}, 2)$. But then this does *not* cover $(0, 1]$: there are points $x \in (0, 1]$ with $x < \frac{1}{m}$. Thus, no finite subcover of \mathcal{C} will cover all of $(0, 1]$. The existence of such an open cover without any finite subcover shows that $(0, 1]$ is not topologically compact.

THEOREM 5.26. *Let K be a set in a subset of a metric space. Then K is sequentially compact iff K is topologically compact.*

We will not prove Theorem 5.26 here; this is the sort of thing that will be covered in an undergraduate topology course (such as Math 190). Rudin chooses to use the more abstract topological definition of compactness (for historical reasons that I find unsatisfactory), and this has the effect of both making everything more abstract, and also making all the proofs *harder than necessary*. We will stick exclusively with *sequential compactness*. This means all our proofs will be different from Rudin's – and generally shorter and easier to understand!

In Example 5.23(2), the absence of the point b from $[a, b)$ makes the set non-compact. Note that b is in the closure of $[a, b)$. This highlights the following proposition.

PROPOSITION 5.27. *Compact sets are closed. Also, if K is compact and $F \subseteq K$ is closed, then F is compact.*

PROOF. Suppose K is compact. Let (x_n) be a sequence in K which converges. By compactness, there is some subsequence (x_{n_k}) that converges to a point in K . But we know that every subsequence of (x_n) converges to $\lim x_n$, and hence $\lim x_n \in K$. Thus, K is closed under limits, and so K is closed.

Now, let F be a closed subset of a compact set K . Let (y_n) be any sequence in F . Then (y_n) is a sequence in K , and hence by compactness there is a subsequence (y_{n_k}) that converges in K . Note that $y_{n_k} \in F$ for each k , and hence since F is closed it follows that $\lim y_{n_k} \in F$. Thus, any sequence in F has a convergent subsequence with limit in F ; i.e. F is compact. \square

DEFINITION 5.28. *Let E be a subset of a metric space. The **diameter** of E , denoted $\text{diam}(E)$ is defined to be*

$$\text{diam}(E) = \sup\{d(x, y) : x, y \in E\}.$$

*Note: this might well be $+\infty$. If $\text{diam}(E) < +\infty$, we call E **bounded**; otherwise E is **unbounded**.*

EXAMPLE 5.29. (1) $\text{diam}(B_r(x)) \leq 2r$ for any ball in a metric space. But it could be less: for example in a discrete metric space with at least two elements, $\text{diam}(B_r(x)) = 0$ if $r \leq 1$ and $= 1$ if $r > 1$.

(2) In \mathbb{R} , $\text{diam}(0, 1) = \text{diam}(0, 1] = \text{diam}[0, 1) = \text{diam}[0, 1] = 1$.

(3) In \mathbb{R} , $\text{diam}(\mathbb{N}) = \infty$. Indeed, $d(0, n) = n$ so $\sup\{d(x, y) : x, y \in \mathbb{N}\} \geq n$ for every n .

Note: if E is a bounded set, with diameter $\delta > 0$, then for any point $x \in E$, $E \subseteq B_{2\delta}(x)$ (or $B_{1.0001\delta}(x)$, or $B_{\delta+0.0001}(x)$, etc.) Conversely, suppose there is some x in the metric space and some $r > 0$ with $E \subseteq B_r(x)$. Since $\text{diam}(B_r(x)) \leq 2r$, it follows that $\text{diam}(E) \leq 2r$. So, to say E is bounded is the same as saying it is contained in some ball.

PROPOSITION 5.30. *Compact sets are bounded.*

PROOF. We prove the contrapositive: *unbounded sets are not compact*. Let E be unbounded, and fix a point $x_0 \in E$. Consider the set of balls $B_n(x_0)$ for $n \in \mathbb{N}$. By assumption, $E \not\subseteq B_n(x_0)$ for any n , so we can choose a point $x_n \in E$ with $d(x_0, x_n) \geq n$.

In fact, the sequence (x_n) has no convergent subsequences. For let x be any point in the metric space. Let $n \in \mathbb{N}$ be large enough that $N > d(x_0, x)$. Then for $n \geq N + 1$, we have by the triangle inequality

$$d(x_n, x) \geq d(x_n, x_0) - d(x_0, x) \geq n - d(x_0, x) \geq 1 + N - d(x_0, x) > 1.$$

That is: for any point x in the metric space, eventually x_n never comes within distance 1 of x . It follows that no subsequence of (x_n) can converge to x . Since this holds for any x , it follows that (x_n) has no convergent subsequences. Since (x_n) is a sequence in E , this means E is not compact. \square

Thus, we have seen that compact sets are closed and bounded. One of the biggest theorems of this course, the Heine–Borel Theorem, states that the converse is true *in Euclidean space*.

THEOREM 5.31 (Heine–Borel). *Let $m \in \mathbb{N}$. A subset of \mathbb{R}^m is compact iff it is closed and bounded.*

PROOF. Let $K \subset \mathbb{R}^m$. If K is compact, then by Propositions 5.27 and 5.30 K is closed and bounded. We must prove the converse. Suppose K is a closed and bounded subset of \mathbb{R}^m . Let (x_n) be a sequence in K . We may write it in terms of its components

$$x_n = (x_n^1, x_n^2, \dots, x_n^m).$$

Consider first the sequence $(x_n^1)_{n=1}^\infty$ in \mathbb{R} . Note that

$$|x_n^1| \leq |x_n| = d(x_n, x_1) \leq \text{diam}(K).$$

So the sequence (x_n^1) is a bounded sequence in \mathbb{R} . By Theorem 2.25 (the Bolzano–Weierstrass Theorem for \mathbb{R}), there is a subsequence $x_{n_k}^1$ that converges. Now we proceed as in the proof of Theorem 3.15 (the Bolzano–Weierstrass Theorem for \mathbb{C}). Consider the subsequence $x_{n_k}^2$. Again we have $|x_{n_k}^2| \leq \text{diam}(K)$ is bounded, so by the Bolzano–Weierstrass Theorem for \mathbb{R} , it possesses a further subsequence $x_{n_{k_\ell}}^2$ that is convergent. Note that $x_{n_{k_\ell}}^1$ is a subsequence of the convergent subsequence $x_{n_k}^1$, so it is also convergent. Now we proceed to select a further convergent subsubsubsequence that makes $x_{n_{k_\ell s}}^3$ converge, and so forth. The notation becomes ridiculous, but in the end (after m steps) we produce a single set of indices $1 \leq \ell_1 < \ell_2 < \dots$ such that all of the components $(x_{\ell_n}^1, x_{\ell_n}^2, \dots, x_{\ell_n}^m)$ converge as $n \rightarrow \infty$. We now follow the proof of Proposition 3.13 exactly to see that convergence in \mathbb{R}^m is equivalent to convergence of each component separately, and so we conclude that the subsequence (x_{ℓ_n}) converges to some element $x \in \mathbb{R}^m$. Finally, note that $x_{\ell_n} \in K$ by assumption, and K is closed; thus the point x is also in K . This shows that every sequence in K has a convergent subsequence with limit in K , concluding the proof that K is compact. \square

4. Lecture 17: February 29, 2016

So, closed intervals $[a, b]$ are compact (as we already knew), as are sets like $[0, 1] \cup [2, 3] \cup \{4, 5, 6, 7, 8\}$. In fact, there are much more complicated closed and bounded sets in \mathbb{R} (e.g. the Cantor set of Example 5.34 below). Let's emphasize that the Heine–Borel Theorem is exclusively about the metric spaces \mathbb{R}^d ; it does not apply in general.

EXAMPLE 5.32. (1) Let (X, d) be a discrete metric space. Then for any two points $x, y \in X$, either $d(x, y) = 0$ or $d(x, y) = 1$. Thus, for any subset $E \subseteq X$, $\text{diam}(E) \leq 1$, so E is bounded. We have also shown that any subset E is closed. However, if E is an infinite set, it is not compact, cf. Example 5.23(3). So any infinite discrete metric spaces contains closed and bounded sets that are not compact.

(2) For a less contrived example, consider again $B[0, 1]$, the set of all bounded, real-valued functions on $[0, 1]$, which is a metric space with respect to the metric

$$d(f, g) = \sup_{x \in [0, 1]} |f(x) - g(x)|.$$

Consider the functions

$$f_n(x) = \begin{cases} 1, & x \geq \frac{1}{n} \\ 0, & x < \frac{1}{n} \end{cases}.$$

All of these functions are in $B[0, 1]$. We can also compute the function $f_n - f_m$; assuming $m > n$ we have

$$f_m(x) - f_n(x) = \begin{cases} 0, & x \geq \frac{1}{n} \\ 1, & \frac{1}{m} \leq x < \frac{1}{n} \\ 0, & x < \frac{1}{m} \end{cases}.$$

This shows that $d(f_n, f_m) = \sup_x |f_n(x) - f_m(x)| = 1$ for any $m \neq n$! Thus, the sequence (f_n) cannot have a convergent subsequence: no two terms in the sequence are ever closer to (or farther from) each other than 1.

Here is an important property of compact sets. This is the generalization of the nested intervals property that we used in the construction of \mathbb{R} .

PROPOSITION 5.33. *Let K_1, K_2, K_3, \dots be nonempty compact sets in a metric space, and suppose they are nested: $K_{n+1} \subseteq K_n$ for all n . Then $\bigcap_n K_n$ is a nonempty compact set. If, in addition, $\text{diam}(K_n) \rightarrow 0$ as $n \rightarrow \infty$, then $\bigcap_n K_n$ consists of exactly one point.*

PROOF. Since $K_n \neq \emptyset$ for any n , we can choose a point $x_n \in K_n$ for each n . By the nested property, $x_n \in K_1$ for each n . Thus, (x_n) is a sequence in the compact set K_1 , and therefore it has a convergent subsequence x_{n_k} with a limit $x \in K_1$. Now, for any $m \in \mathbb{N}$, the tail subsequence $(x_{n_k})_{k=m}^\infty$ also converges to x ; but this is a sequence of terms in K_{n_m} , which is closed, and so $x \in K_{n_m}$. This holds for every m . Finally, for any n , there is $n_m > n$, and therefore $K_{n_m} \subseteq K_n$; thus, $x \in K_n$ for every n , which shows that $x \in \bigcap_n K_n$. This intersection is therefore nonempty. It is an intersection of compact sets, therefore it is compact (Exercise 1 on HW9).

For the second claim, let $x, y \in \bigcap_n K_n$. Fix $\epsilon > 0$; since $\text{diam}(K_n) \rightarrow 0$, there is some n with $\text{diam}(K_n) < \epsilon$. Thus, since $x, y \in K_n$, $d(x, y) \leq \text{diam}(K_n) < \epsilon$. So $0 \leq d(x, y) < \epsilon$ for all $\epsilon > 0$; it follows that $d(x, y) = 0$ and so $x = y$. That is: there is at most one point in the intersection. As we've shown the intersection is nonempty, this proves that it consists of exactly one point, as claimed. \square

EXAMPLE 5.34 (Cantor set). The unit interval $K_0 = [0, 1]$ is compact. Now remove the “middle third” and let $K_1 = [0, \frac{1}{3}] \cup [\frac{2}{3}, 1]$; this set is also compact, and $K_1 \subset K_0$. Now repeat this: remove the middle third from each of the two intervals in K_1 , producing $K_2 = [0, \frac{1}{9}] \cup [\frac{2}{9}, \frac{1}{3}] \cup [\frac{2}{3}, \frac{7}{9}] \cup [\frac{8}{9}, 1]$. Again, this set is compact, and $K_2 \subset K_1$. We can repeat this “delete middle thirds” process indefinitely. Note a certain self-similarity: K_2 has two pieces, each of which looks like K_1 shrunk down by a factor of $\frac{1}{3}$. In fact, we can inductively define

$$K_n = \frac{1}{3}K_{n-1} \cup \left(\frac{2}{3} + \frac{1}{2}K_{n-1} \right).$$

All of these sets are finite collections of closed, bounded intervals, so they are all compact, and they are nested $K_{n+1} \subset K_n$. Hence, by Proposition 5.33, the set $C = \bigcap_n K_n$ is a nonempty compact set. This set is called the **Cantor set**.

What can we say about this set? Well, what is the length of the longest interval in it? Note that K_n consists of 2^n intervals, each of length $\frac{1}{3^n}$. Since the length of the intersection of two intervals is \leq the length of either interval, and since $C \subset K_n$ for every n , this means C contains no intervals of length $> \frac{1}{3^n}$ for any $n \in \mathbb{N}$; but $\frac{1}{3^n} \rightarrow 0$ as $n \rightarrow \infty$, and therefore C contains no intervals of length > 0 . This proves that $\overset{\circ}{C} = \emptyset$. Indeed, if x were an interior point of C , that would mean $B_r(x) \subseteq C$ for some $r > 0$; but $B_r(x) = (x - r, x + r)$ is an interval of length $2r > 0$, which we know is not contained in C . Thus no point is interior to C . At the same time, C is compact, so it is closed. Thus $\overline{C} = C$, and so $\partial C = \overline{C} \setminus \overset{\circ}{C} = C$ – the Cantor set is its own boundary.

That also happens for discrete sets: if K consists entirely of isolated points, then K is closed and $\overset{\circ}{K} = \emptyset$, so $\partial K = K$. But the Cantor set is the *opposite* of a discrete set: it contains *no* isolated points, so C consists entirely of limit points, $C' = C$. To see this, fix $x \in C$; so $x \in K_n$ for every n . Now K_n is a collection of disjoint closed intervals, so there is some interval $I_n \subset K_n$ with $x \in I_n$. Either x is in the interior of this interval or it is one of the endpoints; either way, there is *one* endpoint x_n of I_n with $x_n \neq x$. Now, from the construction of C , the endpoints of all the intervals are in C , so $x_n \in C$. Also, as $\text{diam}(I_n) = \frac{1}{3^n} \rightarrow 0$ and $x, x_n \in I_n$, we have $d(x, x_n) \rightarrow 0$. This $x_n \rightarrow x$, but $x_n \neq x$ for any n , and $x_n \in C$; this proves that $x \in C'$. Since x was an arbitrary element of C , this means $C \subseteq C'$, and since C is closed, we have $C' \subseteq C$, so $C = C'$.

CHAPTER 6

Limits and Continuity

1. Lecture 18: March 3, 2016

We now begin to study functions. We have, of course, been studying functions (for example sequences, which are functions with domain \mathbb{N}); now we will concentrate on metric properties of functions. So we will set things up in terms of functions between metric spaces.

DEFINITION 6.1. *Let X and Y be metric spaces. Let $E \subseteq X$, and let $x_0 \in E'$ be a limit point of E . Let $L \in Y$. Now, for any function $f: E \rightarrow Y$, we say $f(x)$ **tends to L as $x \rightarrow x_0$** , or **the limit as $x \rightarrow x_0$ of $f(x)$ is L** , in symbols*

$$\lim_{x \rightarrow x_0} f(x) = L$$

if: given any sequence (x_n) in $E \setminus \{x_0\}$ that converges $x_n \rightarrow x_0$, it follows that the sequence $(f(x_n))$ in Y converges $f(x_n) \rightarrow L$.

This is a more general kind of limit than the limit of a sequence: we are letting the argument “tend to” a limit point through a set that may be quite different from \mathbb{N} . Our definition makes use of our knowledge of limits of sequences. This is useful, for example, in establishing some of the basic properties of limits. For example:

LEMMA 6.2. *Limits are unique: if $\lim_{x \rightarrow x_0} f(x) = L_1$ and $\lim_{x \rightarrow x_0} f(x) = L_2$, then $L_1 = L_2$.*

PROOF. Since x_0 is a limit point, there is a sequence (x_n) with $x_n \neq x_0$ for any n and $x_n \rightarrow x_0$. By definition of $\lim_{x \rightarrow x_0} f(x) = L_1$, this means that the sequence $f(x_n)$ converges to L_1 ; by definition of $\lim_{x \rightarrow x_0} f(x) = L_2$, this means that $f(x_n)$ converges to L_2 . Thus, by uniqueness of limits of sequences, $L_1 = L_2$. \square

REMARK 6.3. (1) If we had not included in the definition the fact that x_0 is a limit point, this argument would fail. Indeed, if x_0 is an isolated point, vacuously it holds that $\lim_{x \rightarrow x_0} f(x) = L$ for all L .

(2) On the other hand, we might try to modify the definition of limit so that this wouldn't happen: we could, for example, insist that $f(x_n) \rightarrow L$ for any sequence x_n that converges to x_0 , even if it does hit x_0 at some times. But this would rule out some of our intuition about limits, as the following example shows.

EXAMPLE 6.4. Consider the function $f: \mathbb{R} \rightarrow \mathbb{R}$ given by

$$f(x) = \begin{cases} 0, & x \neq 0 \\ 1, & x = 0 \end{cases}.$$

We know from our calculus intuition that $\lim_{x \rightarrow 0} f(x) = 0$. Indeed, we can verify this from Definition 6.1: if x_n is any sequence in $\mathbb{R} \setminus \{0\}$, then $f(x_n) = 0$ for all n , and the constant sequence 0 does indeed converge to 0.

On the other hand, suppose we had left out the $x_n \neq x_0$ clause in Definition 6.1, and insisted that $f(x_n) \rightarrow L$ for every sequence $x_n \rightarrow x_0$. In this scenario, the function above would have no limit at 0. Indeed, we could take the sequence $x_n = \frac{1}{n}$ if n is even and $x_n = 0$ if n is odd. Then the sequence $f(x_n) = (0, 1, 0, 1, 0, 1, \dots)$ has no limit.

This illustrates the fundamental idea of limits: a limit is where a function is *going* as you *approach* the limit point; it is unrelated to the actual *value* of the function at that point (if it is even defined).

We can use our theory of limits of sequences to calculate many limits. For example, if the range space for the function is the familiar \mathbb{C} , we have the following echo of the limit theorems for \mathbb{C} -sequences:

THEOREM 6.5 (Limit Theorems). *Let $f, g: X \rightarrow \mathbb{C}$, and let x_0 be a limit point in X . If $\lim_{x \rightarrow x_0} f(x) = L$ and $\lim_{x \rightarrow x_0} g(x) = M$, then*

$$\lim_{x \rightarrow x_0} [f(x) + g(x)] = L + M, \quad \text{and} \quad \lim_{x \rightarrow x_0} [f(x) \cdot g(x)] = L \cdot M.$$

PROOF. Let (x_n) be any sequence in $X \setminus \{x_0\}$ that converges to x_0 . By assumption, $f(x_n) \rightarrow L$ and $g(x_n) \rightarrow M$. Thus, by the limit theorems for sequences in \mathbb{C} , cf. Theorem 2.27, $f(x_n) + g(x_n) \rightarrow L + M$ and $f(x_n) \cdot g(x_n) \rightarrow L \cdot M$. This is precisely what it means to say that $\lim_{x \rightarrow x_0} [f(x) + g(x)] = L + M$ and $\lim_{x \rightarrow x_0} [f(x) \cdot g(x)] = L \cdot M$. \square

There is an equivalent definition of limit which does not explicitly rely on sequences. This definition is one of the crowning achievements of 19th Century mathematics. The calculus was built on an intuitive understanding of limits in the minds of Newton and Leibnitz (and others), but it wasn't until Weierstrass came up with this modern definition that analysis of functions was finally put on rigorous footing.

THEOREM 6.6. *Let (X, d_X) and (Y, d_Y) be metric spaces, let $E \subseteq X$, let $f: E \rightarrow Y$ be a function, let $x_0 \in E'$, and let $L \in Y$. Then $\lim_{x \rightarrow x_0} f(x) = L$ if and only if the following holds true:*

$$\forall \epsilon > 0 \exists \delta > 0 \text{ s.t. } \forall x \in E, x \in B_\delta(x_0) \setminus \{x_0\} \implies f(x) \in B_\epsilon(L).$$

I.e.

$$\forall \epsilon > 0 \exists \delta > 0 \text{ s.t. } \forall x \in E, 0 < d_X(x, x_0) < \delta \implies d_Y(f(x), L) < \epsilon. \quad (6.1)$$

In words: to say $f(x)$ tends to L as x tends to x_0 means that, for any tolerance $\epsilon > 0$, no matter how small, there is some (potentially even smaller) tolerance $\delta > 0$ so that, if x is δ -close to x_0 (but not *equal* to x_0), then $f(x)$ is ϵ -close to L .

PROOF. First, suppose that (6.1) holds. Let $x_n \in E \setminus \{x_0\}$ be a sequence converging to x_0 . Fix $\epsilon > 0$, and let $\delta > 0$ be the corresponding δ . Now, as $x_n \rightarrow x_0$, there is some $N \in \mathbb{N}$ so that, for $n \geq N$, $d_X(x_n, x_0) < \delta$. It follows from (6.1) that $d_Y(f(x_n), L) < \epsilon$ for all $n \geq N$. This proves that $f(x_n) \rightarrow L$. Thus, we have shown that $\lim_{x \rightarrow x_0} f(x) = L$ by definition.

Conversely, suppose (6.1) fails to hold. This means that there exists some $\epsilon > 0$ so that, for all $\delta > 0$, there is some point $x_\delta \in B_\delta(x_0) \setminus \{x_0\}$ such that $f(x_\delta)$ is *not* in $B_\epsilon(L)$. In particular, do this with $\delta = \frac{1}{n}$: for each $n \in \mathbb{N}$, choose some $x_n \in B_{\frac{1}{n}}(x_0)$ such that $d_Y(f(x_n), L) \geq \epsilon$. On the one hand, since $0 < d_X(x_n, x_0) < \frac{1}{n} \rightarrow 0$, we have $x_n \rightarrow x_0$ but $x_n \neq x_0$. On the other hand, since $d_Y(f(x_n), L) \geq \epsilon$ for all n , this means that the sequence $f(x_n)$ does not converge to L . By definition, this means that the statement $\lim_{x \rightarrow x_0} f(x) = L$ is false. \square

EXAMPLE 6.7. Let us work directly from the ϵ - δ definition of (6.1) to show that $\lim_{x \rightarrow 2} x^2 = 4$. Here the domain and range metric spaces are both \mathbb{R} . Fix $\epsilon > 0$. We want to guarantee that $|x^2 - 4| < \epsilon$. Write this as $|x - 2||x + 2| < \epsilon$. We want to choose $\delta > 0$ and force $0 < |x - 2| < \delta$, meaning $2 - \delta < x < 2 + \delta$. So, as long as we assure that $\delta \leq 2$, this means that $0 \leq x \leq 4$, in which case $|x - 2||x + 2| \leq 6|x - 2|$. Thus, it suffices to make sure that $6|x - 2| < \epsilon$, which is to say $|x - 2| < \epsilon/6$. This tells us how to choose δ .

So, starting fresh: Let $\epsilon > 0$. Choose $\delta = \epsilon/6$ if this is < 2 , or $\delta = 2$ otherwise. Then, so long as $0 < |x - 2| < \delta$, we have $0 \leq 2 - \delta < x < 2 + \delta \leq 4$, and so

$$|x^2 - 4| = |x + 2||x - 2| \leq 6|x - 2| < 6 \cdot \frac{\epsilon}{6} = \epsilon.$$

Thus, by (6.1), we have proven that $\lim_{x \rightarrow 2} x^2 = 4$.

On the other hand, if we refer to Theorem 6.5, we see that this follows from the fact that $\lim_{x \rightarrow 2} x = 2$ (which is easy to verify by either definition of limit) and therefore $\lim_{x \rightarrow 2} x \cdot x = 2 \cdot 2 = 4$. Similar considerations show that $\lim_{x \rightarrow x_0} f(x) = f(x_0)$ holds for any point $x_0 \in \mathbb{R}$ (or \mathbb{C}) if f is a polynomial, for example.

2. Lecture 19: March 8, 2016

In Example 6.7, what we showed is that the function $f(x) = x^2$ satisfies $\lim_{x \rightarrow 2} f(x) = f(2)$. We should recognize this as saying that f is *continuous at 2*.

DEFINITION 6.8. Let X, Y be metric spaces, $E \subseteq X$, and $f: E \rightarrow Y$. Let $x_0 \in E'$ be a limit point. Say that f is **continuous** at x_0 if

$$\lim_{x \rightarrow x_0} f(x) = f(x_0).$$

Note that this only defines continuity at limit points: we have left undefined what it would mean for f to be continuous at an isolated point of its domain of definition. Indeed, what should we mean by saying that a function is continuous on the set \mathbb{N} ? This is, to some degree, up to debate. The standard answer is to say this is a vacuous condition: *every* function is continuous on a discrete set.

Now, consider again Example 6.7. To use the definition of limit, we assumed that $d(x, x_0) = |x - 2| > 0$ (as limits are about where you're going, not where you get to). However, observe that this requirement was never used in the proof. That is generically true in limits of continuous functions, as the next result demonstrates.

PROPOSITION 6.9. Let X, Y be metric spaces and $f: X \rightarrow Y$. Let $x_0 \in X'$. Then f is continuous at x_0 if and only if for every sequence (x_n) in X with $\lim_{n \rightarrow \infty} x_n = x_0$, it follows that $\lim_{n \rightarrow \infty} f(x_n) = f(x_0)$. Similarly, f is continuous at x_0 if and only if

$$\forall \epsilon > 0 \exists \delta > 0 \text{ s.t. } \forall x \in E, d_X(x, x_0) < \delta \implies d_Y(f(x), f(x_0)) < \epsilon.$$

That is: we need not assume that the sequence (x_n) never hits x_0 ; and we need not remove x_0 from the δ -ball in the ϵ - δ definition of the limit. In fact, with these assumption no longer required, there is no reason to assume $x_0 \in X'$; this definition makes perfect sense for isolated points as well, so we take it more generally as the definition of continuity. From this more general definition, it follows that any function is continuous at an isolated point of its domain (as you should work out).

PROOF. Suppose that $x_n \rightarrow x_0$ implies $f(x_n) \rightarrow f(x_0)$ in general; then in particular this holds if we also assume that $x_n \neq x_0$ for any n , which means that $\lim_{x \rightarrow x_0} f(x) = f(x_0)$ by definition. Thus f is continuous at x_0 . Conversely (in contrapositive form), suppose there is some sequence $x_n \rightarrow x_0$ such that $f(x_n) \not\rightarrow f(x_0)$. This means there is some $\epsilon > 0$ so that $d(f(x_n), f(x_0)) \geq \epsilon$ for infinitely many n . So let n_1, n_2, n_3, \dots be these infinitely many indices where, for each k , $d(f(x_{n_k}), f(x_0)) \geq \epsilon$; then $(x_{n_k})_{k=1}^{\infty}$ is a sequence in X that converges to x_0 (as a subsequence of a sequence x_n which converge to x_0), but $f(x_{n_k}) \not\rightarrow f(x_0)$ and, moreover, since $f(x_{n_k}) \neq f(x_0)$ for any k , it follows that $x_{n_k} \neq x_0$ for any k . This shows, from the definition, that it is false that $\lim_{x \rightarrow x_0} f(x) = f(x_0)$, completing the proof of the first statement.

The proof of the equivalence of the ϵ - δ statement is similar and left to the reader. \square

The point is: when the putative limit is the *value* of the function at the limit point, there is no reason to exclude the limit point from consideration: where you are going and where you get to are the same in this case!

EXAMPLE 6.10. Let (X, d) be a metric space, and let $y \in X$. Then the function $f(x) = d(x, y)$ is continuous at every point in X . Indeed, fix $x \in X$, and let (x_n) be a sequence in X with $x_n \rightarrow x$. Then

$$d(x_n, y) \leq d(x_n, x) + d(x, y)$$

and so $d(x_n, y) - d(x, y) \leq d(x_n, x)$. But also

$$d(x, y) \leq d(x, x_n) + d(x_n, y)$$

and so $d(x, y) - d(x_n, y) \leq d(x, x_n) = d(x_n, x)$. Together, these give

$$0 \leq |d(x_n, y) - d(x, y)| \leq d(x_n, x).$$

Since $x_n \rightarrow x$, $d(x_n, x) \rightarrow 0$ by definition, and so by the squeeze theorem $|d(x_n, y) - d(x, y)| \rightarrow 0$, meaning that $f(x_n) = d(x_n, y) \rightarrow d(x, y) = f(x)$. This shows that f is continuous at x_0 .

We would be remiss if we did not include some examples of *discontinuous* functions.

EXAMPLE 6.11. Let $f: \mathbb{R} \rightarrow \mathbb{R}$ be the function


$$f(x) = \begin{cases} 0, & x \notin \mathbb{Q} \\ 1, & x \in \mathbb{Q} \end{cases}.$$

Then f (sometimes called *Dirichlet's function*) is not continuous at *any* point. Indeed, fix $x \in \mathbb{R}$. For any $\delta > 0$, the ball $B_\delta(x) = (x - \delta, x + \delta)$ contains both rational and irrational numbers. So, if $x \in \mathbb{Q}$, choose some $y \notin \mathbb{Q}$ in the ball, and we have $|f(x) - f(y)| = |1 - 0| = 1$; if $x \notin \mathbb{Q}$, choose some $y \in \mathbb{Q}$ in the ball, and we have $|f(x) - f(y)| = |0 - 1| = 1$. In any case, we see that for any $\delta > 0$ there are points $y \in B_\delta(x)$ so that $|f(y) - f(x)| = 1$, so we can never force $f(y)$ to be in, for example, $B_{\frac{1}{2}}(f(x))$. This shows f is discontinuous at x , for any x .

EXAMPLE 6.12. Consider the following function $f: [0, 1] \rightarrow [0, 1]$, sometimes called the *popcorn function*:

$$f(x) = \begin{cases} 0, & x \notin \mathbb{Q} \\ \frac{1}{q}, & x = \frac{p}{q} \text{ in lowest terms} \end{cases}.$$

The graph of this function looks like this:



In fact, f is discontinuous at all rational points, but it is actually continuous at all irrational points. Indeed, let $x = \frac{p}{q}$ be rational, and let $x_n = x + \frac{\sqrt{2}}{n}$ for all n large enough that this is in $[0, 1]$; then $x_n \rightarrow x$. Then $x_n \notin \mathbb{Q}$ meaning that $f(x_n) = 0$; but $f(x) = \frac{1}{q} \neq 0$, so $f(x_n) \not\rightarrow f(x)$. On the other hand, let $x \notin \mathbb{Q}$; we want to show that f is continuous at x , meaning

$\lim_{y \rightarrow x} f(y) = f(x) = 0$. Fix $\epsilon > 0$, and choose some integer $n \in \mathbb{N}$ with $\frac{1}{n} < \epsilon$. As $f(x) \geq 0$ for all x , it suffices to show that $f(y) < \frac{1}{n}$ for all y sufficiently close to x . Well, if y is a point in $[0, 1]$ where $f(y) \geq \frac{1}{n}$, then $y \in \mathbb{Q}$ and, when written in lowest terms, $y = \frac{p}{q}$ with $q \leq n$. There are only finitely many such rational numbers, and x (which is irrational) is not one of them. Thus, we can define $\delta = \min\{|x - y| : y = \frac{p}{q} \text{ in lowest terms, with } q \leq n\}$; then for $|y - x| < \delta$, it follows that $f(y) < \frac{1}{n} < \epsilon$, proving that $\lim_{y \rightarrow x} f(y) = 0$, and so f is continuous at x .

In Examples 6.11 and 6.12, we looked at the set of points where a function is continuous. That is: if $f: X \rightarrow Y$ is a function and $E \subseteq X$, say that f **is continuous on** E if, for each $x \in E$, f is continuous at x .

EXAMPLE 6.13. Let $X = (0, 1)$, and let $f: (0, 1) \rightarrow \mathbb{R}$ be the function $f(x) = \frac{1}{x}$. Then f is continuous on its whole domain: for every $x \in (0, 1)$, f is continuous at x . We could see this by applying the limit theorems; but let's use this as an opportunity to practice our ϵ - δ proofs. Fix $\epsilon > 0$. We want to guarantee that, when y is close to x , we have $|\frac{1}{x} - \frac{1}{y}| < \epsilon$. That is

$$\frac{1}{x} - \epsilon < \frac{1}{y} < \frac{1}{x} + \epsilon.$$

We only need this to hold for all sufficiently small $\epsilon > 0$, so it's fine to assume ϵ is small enough that $\frac{1}{x} - \epsilon > 0$. Thus we can reciprocate to get

$$\frac{1}{1/x - \epsilon} > y > \frac{1}{1/x + \epsilon}.$$

Now, subtract x from both sides and we have

$$\frac{-\epsilon x}{1/x + \epsilon} = \frac{1}{1/x + \epsilon} - x < y - x < \frac{1}{1/x - \epsilon} - x = \frac{\epsilon x}{1/x - \epsilon}.$$

This shows us how to choose δ : we define

$$\delta = \min \left\{ \frac{\epsilon x}{1/x + \epsilon}, \frac{\epsilon x}{1/x - \epsilon} \right\} = \frac{\epsilon x}{1/x + \epsilon}. \quad (6.2)$$

Then, reversing the above steps, we have that for any $y \in B_\delta(x)$, we have $|y - x| < \frac{\epsilon x}{1/x + \epsilon} < \frac{\epsilon x}{1/x - \epsilon}$, and this gives in particular the above two inequalities that can be reversed to say $|\frac{1}{x} - \frac{1}{y}| < \epsilon$. So we have proved that there is a $\delta > 0$ for any given $\epsilon > 0$ (as long as $\epsilon < \frac{1}{x}$; otherwise, if $\epsilon \geq \frac{1}{x} \geq 1$, we could take δ to be something silly and big), proving continuity at x .

3. Lecture 20: March 10, 2016

In Example 6.13, we showed explicitly that the function $f(x) = \frac{1}{x}$ is continuous at every point $x \in (0, 1)$. But note: the δ we had to choose for each $\epsilon > 0$ in (6.2) *depends on x* as well as ϵ . This will generically be true. Look at the ϵ - δ definition of continuity: a function f is continuous on a set E if

$$\forall x \in E \forall \epsilon > 0 \exists \delta > 0 \text{ s.t. } \forall y \in E, d_X(x, y) < \delta \implies d_Y(f(x), f(y)) < \epsilon. \quad (6.3)$$

Having chosen a x and $\epsilon > 0$, we must then find a suitable $\delta = \delta(x, \epsilon)$. In Example 6.13, not only does δ depend on x , but it does so in a bad way: as $x \rightarrow 0$, for given $\epsilon > 0$, the $\delta \rightarrow 0$ as well (quite fast, in fact: the numerator is shrinking and the denominator is growing). The closer x is to 0, the smaller δ must be to get the same control over the function. So, while the function is continuous, there is a lack of uniformity in how continuous it is. (Note: we have shown this δ works; one might ask whether a larger, possibly more uniform δ could work just as well. The answer is no: it is not hard to show in this example that the δ in (6.2) is the largest possible δ for the given x and ϵ ; it is called the *modulus of continuity* of the function.)

DEFINITION 6.14. *Let X, Y be metric spaces, $E \subseteq X$, and $f: E \rightarrow Y$. Call f **uniformly continuous** on E if*

$$\forall \epsilon > 0 \exists \delta > 0 \text{ s.t. } \forall x, y \in E, d_X(x, y) < \delta \implies d_Y(f(x), f(y)) < \epsilon. \quad (6.4)$$

Compare (6.4) with (6.3). The difference appears subtle: just the placement of the quantifier $\forall x$. This makes a world of difference: (6.4) says that, not only is f continuous at each point x , but one can choose a $\delta = \delta(\epsilon)$ that is **uniform**: it need not depend on x . This outlaws behavior like the function $f(x) = \frac{1}{x}$ near 0.

EXAMPLE 6.15. Let $f(x) = 2x$ on \mathbb{R} . Then $|f(x) - f(y)| = 2|x - y|$, so for any $\epsilon > 0$, we may let $\delta = \epsilon/2$; then if $|x - y| < \delta = \epsilon/2$, it follows that $|f(x) - f(y)| = 2|x - y| < 2\delta = \epsilon$. This shows that f is continuous at all points; moreover, we may choose $\delta = \delta(\epsilon) = \epsilon/2$ uniformly over all $x, y \in \mathbb{R}$. Thus, f is uniformly continuous on \mathbb{R} .

EXAMPLE 6.16. Let $f(x) = x^2$ on $[0, \infty)$. We want to make $|x^2 - y^2|$ small. We have $|x^2 - y^2| = (x + y)|x - y|$. Thus, in order for $|x^2 - y^2| < \epsilon$, we must have $|x - y| < \frac{\epsilon}{x+y}$ (these are equivalent). But this shows that f is not uniformly continuous. Indeed, in order for $|x - y| < \delta$ to imply that $|x - y| < \frac{\epsilon}{x+y}$, we must have $\delta \leq \frac{\epsilon}{x+y}$; and there is no positive number $\delta = \delta(\epsilon)$ that is $\leq \frac{\epsilon}{x+y}$ for all $x, y > 0$.

In Examples 6.13 and 6.16, we saw continuous functions on the intervals $(0, 1)$ and $[0, \infty)$ that are not uniformly continuous. In both cases, the non-uniformity was manifest by uncontrolled growth near the “edge”. As it turns out, if the domain of the continuous function is compact, this cannot happen. That will be our final big theorem of this class.

THEOREM 6.17. *Let X, Y be metric spaces, $K \subseteq X$ compact, and $f: K \rightarrow Y$ a continuous function. Then f is uniformly continuous.*

PROOF. Suppose, for a contradiction, that f is not uniformly continuous on K . Negating Definition 6.14, this means

$$\exists \epsilon > 0 \forall \delta > 0 \exists x, y \in K \text{ s.t. } d_X(x, y) < \delta, \text{ but } d_Y(f(x), f(y)) \geq \epsilon.$$

That is: there is a positive number $\epsilon > 0$ so that, for every positive number $\delta > 0$, we can find two points x and y that are within distance δ of each other, but such that $f(x)$ and $f(y)$ are at

least ϵ apart. So, let's do this with $\delta = \frac{1}{n}$ for any given positive integer: we can find x_n, y_n with $d_X(x_n, y_n) < \frac{1}{n}$, and yet $d_Y(f(x_n), f(y_n)) \geq \epsilon$.

Now, we use the compactness of the domain K : the sequence (x_n) has a convergent subsequence (x_{n_k}) with a limit $x \in K$. Consider, now, the corresponding subsequence y_{n_k} ; this has a convergent subsequence $y_{n_{k_\ell}}$ with a limit $y \in K$. Now, $x_{n_{k_\ell}}$ is a subsequence of x_{n_k} which converges to x , hence $x_{n_{k_\ell}} \rightarrow x$ as well. But we also have

$$d_X(x_{n_{k_\ell}}, y_{n_{k_\ell}}) < \frac{1}{n_{k_\ell}} < \frac{1}{\ell} \rightarrow 0.$$

Hence, it follows from the triangle inequality that $x = y$. On the other hand, by their very construction, the points $x_{n_{k_\ell}}$ and $y_{n_{k_\ell}}$ all satisfy

$$d_Y(f(x_{n_{k_\ell}}), f(y_{n_{k_\ell}})) \geq \epsilon. \quad (6.5)$$

But $x_{n_{k_\ell}} \rightarrow x$ and so, since f is continuous, $f(x_{n_{k_\ell}}) \rightarrow f(x)$; similarly, $y_{n_{k_\ell}} \rightarrow y = x$, and so by continuity $f(y_{n_{k_\ell}}) \rightarrow f(y) = f(x)$. Thus, by Problem 4 on Exam 2,

$$d_Y(f(x_{n_{k_\ell}}), f(y_{n_{k_\ell}})) \rightarrow d_Y(f(x), f(x)) = 0.$$

This contradicts (6.5). Thus, we have proven that f is, in fact, uniformly continuous. \square

Theorem 6.17 is typically the best way to prove uniform continuity of a function. For example: any polynomial is continuous on \mathbb{R} , but, as we saw in Example 6.16, they need not be uniformly continuous. By Theorem 6.17, polynomial functions on compact intervals $[a, b]$ are automatically uniformly continuous. What's more: once you know a function is uniformly continuous on a set K , it is then automatically uniformly continuous on any subset $E \subseteq K$ (the same $\delta = \delta(\epsilon)$ that works on all of K also works on all of $E \subseteq K$). So, for example, polynomials are uniformly continuous on all bounded intervals (a, b) , $[a, b]$, etc. Similarly, the function $f(x) = \frac{1}{x}$ of Example 6.13 is uniformly continuous on $[\alpha, 1]$ for any $\alpha > 0$. We could see this directly from (6.2), since the modulus of continuity

$$\delta = \delta(x, \epsilon) = \frac{\epsilon x}{1/x + \epsilon}$$

decreases as x decreases; it follows that the uniform $\delta = \delta(\alpha, \epsilon)$ will work for all $x \geq \alpha$. However, this gets smaller as α shrinks, and if we include all of $(0, 1]$ in the domain, there is *no* uniform δ . For an alternate proof of the non-uniformity in this example, see HW10.4.

Here is another very useful property of continuous functions on compact sets.

PROPOSITION 6.18. *Let X, Y be metric spaces, $K \subseteq X$ compact, and $f: K \rightarrow Y$ continuous (hence uniformly continuous). Then the image $f(K) \subseteq Y$ is compact.*

To be clear: $f(K)$ denotes the image of f on K :

$$f(K) = \{f(x) \in Y : x \in K\} = \{y \in Y : \exists x \in K \text{ s.t. } y = f(x)\}.$$

PROOF. Let (y_n) be any sequence in $f(K)$. By definition of $f(K)$, for each y_n , there exists some (or potentially many) $x_n \in K$ such that $y_n = f(x_n)$. Since K is compact, it then follows that the sequence (x_n) has a convergent subsequence (x_{n_k}) with limit $x \in K$. Since f is continuous, it then follows that $f(x_{n_k}) \rightarrow f(x)$ as $k \rightarrow \infty$. Since $x \in K$, $f(x) \in f(K)$. Thus, the subsequence $y_{n_k} = f(x_{n_k})$ of $y_n = f(x_n)$ converges in $f(K)$. We have thus shown that every sequence in $f(K)$ has a convergent subsequence with limit in $f(K)$; that is, $f(K)$ is compact. \square

COROLLARY 6.19 (Extreme Values Theorem). *Let K be a nonempty compact metric space, and $f: K \rightarrow \mathbb{R}$. Then f attains its maximum and minimum values on K .*

Corollary 6.19 is a standard result stated in calculus classes, usually in the special case that $K = [a, b]$ is a compact interval in \mathbb{R} .

PROOF. By Proposition 6.18, $f(K)$ is compact. In particular, it is closed and bounded. It is also nonempty since K is nonempty (so $f(K)$ contains $f(x)$ for any $x \in K$). Thus, by the least upper bound property of \mathbb{R} , the set $f(K) \subset \mathbb{R}$ has a supremum M and an infimum m . Now, for any $n \in \mathbb{N}$, $M - \frac{1}{n} < M$, which means that $M - \frac{1}{n}$ is not an upper bound for $f(K)$; thus, there is some $y_n \in f(K)$ with $M - \frac{1}{n} < y_n \leq M$. Hence, by the Squeeze Theorem, $y_n \rightarrow M$. By definition of $f(K)$, there exists some $x_n \in K$ with $y_n = f(x_n)$. Since K is compact, there is a convergent subsequence (x_{n_k}) of (x_n) , with limit $x \in K$. Since f is continuous, $y_{n_k} = f(x_{n_k}) \rightarrow f(x)$. But y_{n_k} is a subsequence of y_n which converges to M ; thus $f(x) = M$. We have therefore found a $x \in K$ for which $f(x) = M = \sup f(K)$. That is: $\sup f(K) = \max f(K)$, and the maximum is achieved at the point x . A very similar argument shows there is a point x' with $f(x') = m = \inf f(K)$, completing the proof. \square

Part 2

Math 140B

CHAPTER 7

More on Continuity

1. Lecture 1: March 29, 2016

Presently, we return to the definition(s) of continuity, and consider more purely topological characterizations. As such, we will not be dealing with continuity at a point, but instead continuity on a set (or most of the time on the whole domain metric space), which simply means pointwise continuity at each point in the domain. To begin, we note that continuity is well-behaved under composition.

PROPOSITION 7.1. *Let X, Y, Z be metric spaces, and suppose $f: X \rightarrow Y$ and $g: Y \rightarrow Z$ are continuous functions. Then the composite function $g \circ f: X \rightarrow Z$ is continuous.*

PROOF 1. Here we see very clearly the power of the sequential definition of continuity (Definition 6.8). Let $x \in X$, and let (x_n) be a sequence that converges to x . Since f is continuous on X , hence at x , it follows that $f(x_n) \rightarrow f(x)$. Since $(f(x_n))$ is a sequence in Y that converges to $f(x)$, and since g is continuous on Y , hence at $f(x)$, it follows that $g(f(x_n)) \rightarrow g(f(x))$. But $g(f(x_n)) = (g \circ f)(x_n)$, and $g(f(x)) = (g \circ f)(x)$. It follows that $g \circ f$ is continuous at (an arbitrarily chosen) $x \in X$. \square

PROOF 2. It is also possible to prove the proposition using the ϵ - δ definition of continuity (Proposition 6.9). Fix $x \in X$ and $\epsilon > 0$. Since g is continuous on Y , hence at $f(x)$, there is some $\delta > 0$ so that $g(B_\delta(f(x))) \subseteq B_\epsilon(g(f(x)))$. Now, since f is continuous on X , hence at x , there is some $\delta' > 0$ so that $f(B_{\delta'}(x)) \subseteq B_\delta(f(x))$. Now, given two subsets $A, B \subseteq Y$ with $A \subseteq B$, it follows that $g(A) \subseteq g(B)$ by definition. Thus

$$(g \circ f)(B_{\delta'}(x)) = g(f(B_{\delta'}(x))) \subseteq g(B_\delta(f(x))) \subseteq B_\epsilon(g(f(x))) = B_\epsilon((g \circ f)(x)).$$

Thus, for each $\epsilon > 0$ we can choose $\delta' > 0$ with $(g \circ f)(B_{\delta'}(x)) \subseteq B_\epsilon((g \circ f)(x))$, which shows $g \circ f$ is continuous at (an arbitrarily chosen) x . \square

In Proof 2 above, we used a property of set mappings (that $A \subseteq B \implies f(A) \subseteq f(B)$). In order to state and prove our next theorem, we need a more thorough understanding of the behavior of set mapping; the following discussion is purely set theoretic.

DEFINITION 7.2. *Let X be a set. Denote by 2^X the **power set** of X : the set of all subsets of X . Let Y be another set, and suppose $f: X \rightarrow Y$ is a function. Then there is an induced function (also denoted f) from 2^X to 2^Y : for any subset $A \subseteq X$, $f(A) = \{f(x) : x \in A\}$. It also induces a reverse map $f^{-1}: 2^Y \rightarrow 2^X$: for $B \subseteq Y$, $f^{-1}(B) = \{x \in X : f(x) \in B\}$.*

Note: it is not necessary for f to be one-to-one in order for f^{-1} to exist as a set mapping: it always exists. In fact, f^{-1} generally has better properties than f as a set mapping.

LEMMA 7.3. *Let X and Y be sets, and $f: X \rightarrow Y$.*

(1) *For any $B \subseteq Y$, $f^{-1}(B^c) = f^{-1}(B)^c$.*

- (2) For any $B_1, B_2 \subseteq Y$, $f^{-1}(B_1 \cup B_2) = f^{-1}(B_1) \cup f^{-1}(B_2)$, and $f^{-1}(B_1 \cap B_2) = f^{-1}(B_1) \cap f^{-1}(B_2)$.
- (3) For any $A_1, A_2 \subseteq X$, $f(A_1 \cup A_2) = f(A_1) \cup f(A_2)$, while in general $f(A_1 \cap A_2) \subseteq f(A_1) \cap f(A_2)$.

PROOF. (1) An element $x \in X$ is in $f^{-1}(B)^c$ if and only if $x \notin f^{-1}(B)$. By definition, this is the same as saying that $f(x) \notin B$, which is to say that $f(x) \in B^c$, or equivalently $x \in f^{-1}(B^c)$, as desired.

(2) An element $x \in X$ is in $f^{-1}(B_1 \cup B_2)$ iff $f(x) \in B_1 \cup B_2$. If $f(x) \in B_1$ then $x \in f^{-1}(B_1)$; if $f(x) \in B_2$ then $x \in f^{-1}(B_2)$; altogether, this means that $f(x) \in B_1 \cup B_2$ iff $x \in f^{-1}(B_1) \cup f^{-1}(B_2)$, as desired. The argument for intersections is similar.

(3) For the first statement, let $y \in f(A_1) \cup f(A_2)$. So either there is an $x_1 \in A_1$ with $f(x_1) = y$, or there is an $x_2 \in A_2$ with $f(x_2) = y$. In either case, $y \in f(A_1 \cup A_2)$, as desired. Conversely, if $y \in f(A_1 \cup A_2)$, then there is some $x \in A_1 \cup A_2$ with $y = f(x)$. Then either $x \in A_1$, in which case $y = f(x) \in f(A_1)$, or $x \in A_2$, in which case $y = f(x) \in f(A_2)$; thus $y \in f(A_1) \cup f(A_2)$, as desired.

For the second statement, let $y \in f(A_1 \cap A_2)$; so there is some $x \in A_1 \cap A_2$ with $f(x) = y$. Since $x \in A_1$, it follows that $y = f(x) \in f(A_1)$; since $x \in A_2$, it follows that $y = f(x) \in f(A_2)$; thus $y \in f(A_1) \cap f(A_2)$, as desired. □

REMARK 7.4. Nothing like item (1) holds for the forward set mapping in general. On the one hand, if f is a bijection with inverse function g , then $f(A) = g^{-1}(A)$ (you should untwist the definitions to check this), and so in this case the forward set mapping has all the same nice properties as the inverse. On the other hand, suppose X has more than one element, and f is a constant map $f(x) = y_0$ for all $x \in X$. Then for any $A \subseteq X$, $f(A^c) = \{y_0\}$, while $f(A)^c = \{y_0\}^c$, so the two sets are not only not equal, they are complementary (meaning in general there are not even any consistent inclusions of one into the other).

Similarly, the inclusion $f(A_1 \cap A_2) \subseteq f(A_1) \cap f(A_2)$ is, in general, not an equality if f is not a bijection. For example, take again a constant function $f(x) = y_0$, defined on a set X containing at least two distinct points x_1 and x_2 . Then $f(\{x_1\}) \cap f(\{x_2\}) = \{y_0\}$, but $\{x_1\} \cap \{x_2\} = \emptyset$, so $f(\{x_1\} \cap \{x_2\}) = \emptyset \subsetneq \{y_0\}$. So, we see, in general the forward set mapping is not as well-behaved as the inverse set mapping with respect to the Boolean operations; i.e. the inverse set mapping is always a Boolean homomorphism.

The nice behavior of the inverse setting mapping helps to explain why it appears in the following topological characterization of continuity, instead of the forward set mapping.

THEOREM 7.5. *Let X and Y be metric spaces. A function $f: X \rightarrow Y$ is continuous if and only if the preimage of any open set is open; i.e. for all open sets $V \subseteq Y$, $f^{-1}(V)$ is open in X .*

PROOF. First, suppose f is continuous. Let $V \subseteq Y$ be an open set, and let $x \in f^{-1}(V)$; this means that $f(x) \in V$. Since V is open, there is some $\epsilon > 0$ with $B_\epsilon(f(x)) \subseteq V$. Since f is continuous, there is some $\delta > 0$ with $f(B_\delta(x)) \subseteq B_\epsilon(f(x)) \subseteq V$. But that means that $B_\delta(x) \subseteq f^{-1}(V)$. Hence, every $x \in f^{-1}(V)$ is an interior point of $f^{-1}(V)$, and so $f^{-1}(V)$ is open.

Conversely, suppose we know that $f^{-1}(V)$ is open for every open $V \subseteq Y$. Let $x \in X$ and fix $\epsilon > 0$. By assumption, $f^{-1}(B_\epsilon(f(x)))$ is open. Since x is a point in this open preimage, that

means there is some $\delta > 0$ with $B_\delta(x) \subseteq f^{-1}(B_\epsilon(f(x)))$. Hence $f(B_\delta(x)) \subseteq B_\epsilon(f(x))$, which shows that f is continuous at (an arbitrarily chosen) $x \in X$. \square

REMARK 7.6. In the last step of the above proof, we essentially used the fact that $f(f^{-1}(B)) \subseteq B$, which is left as an exercise for you to verify. The reverse containment $f^{-1}(f(A)) \supseteq A$ also holds true in general; in both cases, the containments are generally not equalities.

The above proof is quite natural, and really shows that the ϵ - δ definition of continuity fundamentally says that preimages of open sets are open. However, combining Theorem 7.5 with Lemma 7.3(1) gives an alternate characterization which is less clear (and therefore more interesting).

COROLLARY 7.7. *Let X and Y be metric spaces. A function $f: X \rightarrow Y$ is continuous if and only if the preimage of any closed set is closed; i.e. for all closed sets $C \subseteq Y$, $f^{-1}(C)$ is closed in X .*

PROOF. By Theorem 7.5, it suffices to show the preimages of closed sets are closed if and only if preimages of open sets are open. But this follows from 7.3(1) together with Proposition 5.15: the complement of an open set is closed. Precisely: suppose the preimage of any open set is open. Let C be a closed set in Y . Then C^c is open, and so $f^{-1}(C^c)$ is open. But $f^{-1}(C^c) = f^{-1}(C)^c$, and so $f^{-1}(C)$ is a set whose complement is open, meaning that it is closed. The converse is very similar. \square

EXAMPLE 7.8. Consider the metric space $\mathbb{S} = \{e^{i\theta} : \theta \in \mathbb{R}\} \subset \mathbb{C}$, the unit circle in the complex plane. As a subset of \mathbb{C} , it inherits the Euclidean \mathbb{C} -metric. but this might not be the most natural metric to use. For example, the distance between the points $(1, 0)$ and $(-1, 0)$ is 2 in the \mathbb{C} -metric; but if we are thinking of distance as the shortest path in the circle, then the distance should be π . To try to define this intrinsic metric, we'd like to have a correspondence between $e^{i\theta}$ and θ . Of course, we cannot do this for all $\theta \in \mathbb{R}$, but we can restrict the allowed θ to be in $[0, 2\pi)$, and define

$$f(e^{i\theta}) = \theta, \quad 0 \leq \theta < 2\pi.$$

This is well-defined (by the polar decomposition: every nonzero complex number z has a *unique* decomposition $z = |z|e^{i\theta}$ for some $\theta \in [0, 2\pi)$). But is this a continuous function $\mathbb{S} \rightarrow \mathbb{R}$? Naively, we might write $f(u) = -i \ln u$ and expect this means it is continuous. However, consider the closed set $[\pi, 10] \subset \mathbb{R}$; the preimage of this set is

$$f^{-1}([\pi, 10]) = f^{-1}([\pi, 2\pi)) = f^{-1}([\pi, 2\pi) \cup [2\pi, 10]) = f^{-1}([\pi, 2\pi)) \cup f^{-1}([2\pi, 10]).$$

Since the image of f does not intersect $[2\pi, 10]$, this second preimage is empty, and so $f^{-1}([\pi, 10]) = \{e^{i\theta} : \pi \leq \theta < 2\pi\}$, which is the bottom half of \mathbb{S} , including the point $(-1, 0)$ but not including the point $(1, 0)$. This set is not closed in \mathbb{S} : the sequence $e^{-\pi i/n}$ lives in this bottom half, and converges to $(1, 0)$ as $n \rightarrow \infty$. Thus, f is not continuous.

REMARK 7.9. It is, indeed, possible to define a metric on \mathbb{S} that represents this “length of shortest path” intuition; but this example shows that one cannot draw a global correspondence between \mathbb{S} and an interval in \mathbb{R} to make it work. Rather, one must work only locally. This is a topic for a course in differentiable manifolds.

Example 7.8 illustrates a very interesting point. The function f in that example is a bijection; let's consider its inverse $g: [0, 2\pi) \rightarrow \mathbb{S}$, which is simply $g(\theta) = e^{i\theta}$. This function *is* continuous (although we will not prove this until we study sequences and series of functions, in the next chapter). So it is possible for a continuous bijection to have an inverse that is not continuous.

Indeed, in this example, g “glues together” the points 0 and 2π at the boundary of its domain, and so its inverse must “rip them apart”, which is discontinuous.

This kind of pathology does not happen, however, if the domain of the bijection is compact.

THEOREM 7.10. *Let X and Y be metric spaces, and suppose X is compact. If $f: X \rightarrow Y$ is a continuous bijection, then $f^{-1}: Y \rightarrow X$ is continuous.*

PROOF. To avoid confusing notation, denote the inverse function $f^{-1} = g$. Let $C \subseteq X$ be closed. The preimage $g^{-1}(C)$ is equal to the forward image $f(C)$. Since X is compact and $C \subseteq X$ is closed, C is compact (cf. Proposition 5.27). Now, by Proposition 6.18, it follows that $f(C)$ is compact, and therefore closed (again by Proposition 5.27). Thus, $g^{-1}(C)$ is closed for every closed C , and so by Corollary 7.7, g is continuous. \square

A continuous bijection whose inverse is continuous is called a **homeomorphism**. The function $g(\theta) = e^{i\theta}$ from $[0, 2\pi)$ onto \mathbb{S} is a continuous bijection, but it is not a homeomorphism. In topology, homeomorphisms are the basic “isomorphisms”; they tell you when two structures are topologically indistinguishable. The circle \mathbb{S} and the half-open interval $[0, 2\pi)$ are topologically distinguishable (this is intuitively clear, but proving that there exists no homeomorphism between them requires developing tools beyond the scope of this course).

2. Lecture 2: March 31, 2016

Before proceeding with more discussion continuity, we return to a purely topological notion.

DEFINITION 7.11. Let X metric space X . Two subsets A and B in X are called **separated** if $\overline{A} \cap B = A \cap \overline{B} = \emptyset$; that is, no point in A is in the closure of B , and no point in B is in the closure of A . The metric space X is called **connected** if it is not the union of two nonempty separated sets.

EXAMPLE 7.12. The two intervals $(0, 1)$ and $(2, 3)$ are separated in \mathbb{R} ; in fact, their closures $[0, 1]$ and $[2, 3]$ are disjoint (which is nominally stronger than separation). Similarly, the two intervals $(0, 1)$ and $(1, 2)$ are separated: $[0, 1]$ does not intersect $(1, 2)$, and $(0, 1)$ does not intersect $[1, 2]$. However, $(0, 1)$ and $[1, 2)$ are *not* separated: the closure $[0, 1]$ does intersect $[1, 2)$.

Two separated sets are, of course, disjoint, but as the example points out, disjointness is not sufficient to imply separation.

In \mathbb{R} , connected sets can be easily characterized; they are precisely the sets we've been most concerned with: intervals.

PROPOSITION 7.13. A subset $E \subseteq \mathbb{R}$ is connected if and only if it is an interval: i.e. if and only if it has the property that, for all $x, y, z \in E$, if $x < z < y$ and $x, y \in E$, then $z \in E$.

More precisely, the property in Proposition 7.13 should be called the *intermediate value property* (of subsets in an ordered set). It is equivalent to insisting that E is an interval (with one or both endpoints included or not, or possibly infinite). Indeed, it is straightforward to see that intervals have the intermediate value property. The converse is a case analysis. For example, suppose $E \neq \emptyset$ has finite sup and inf. If $\sup E = \inf E$ then E consists of a single point which is an interval. Otherwise, $\inf E < \sup E$; let z be in between. Since $z < \sup E$, there exists some point $y \in E$ with $z < y \leq \sup E$. Similarly, since $z > \inf E$, there exists some point $x \in E$ with $\inf E \leq x < z$. Thus $x < z < y$, and by the intermediate value property, $z \in E$. This shows that $(\inf E, \sup E) \subseteq E$. Of course E contains no points bigger than $\sup E$ or smaller than $\inf E$; so this shows E is one of the four intervals whose closure is $(\inf E, \sup E)$ (with neither, either one, or both of the endpoints included). The argument is similar when one or both of the sup and inf are infinite.

PROOF OF PROPOSITION 7.13. First, suppose that E does not have the intermediate value property: there are points $x < z < y$ with $x, y \in E$ but $z \notin E$. Define $A = (-\infty, z) \cap E$ and $B = (z, \infty) \cap E$; since $z \notin E$, it follows that $E = A \cup B$. Since $x \in A$ and $y \in B$, both sets are nonempty. Now, z is an upper bound for A , so $\sup A \leq z$. This means that $\overline{A} \subseteq (-\infty, z]$; this follows from the squeeze theorem: if $a_n \in A$ and $a_n \rightarrow a$ then, since $a_n < z$ for all n , $a \leq z$. Hence, since $B \subseteq (z, \infty)$, it follows that $\overline{A} \cap B = \emptyset$. A similar argument shows that $A \cap \overline{B} = \emptyset$. Thus $E = A \cup B$ with A and B nonempty and separated; thus E is not connected.

For the converse, suppose that E is not connected, and let $E = A \cup B$ with A and B nonempty and separated. Choose $x \in A$ and $y \in B$; we'll assume $x < y$ (if they're reverse ordered, just rename them). Set $z = \sup(A \cap [x, y])$. By Example 5.21, $z \in \overline{A}$, and therefore $z \notin B$ (since A and B are separated). Since $y \in B$, $z \neq y$. We now consider two cases.

- Suppose $z \notin A$. Since $x \in A$, $z \neq x$. But $z \in [x, y]$, therefore $x < z < y$. But $z \notin A$ and $z \notin B$, so $z \notin E = A \cup B$. Therefore, E does not have the intermediate value property.
- Suppose $z \in A$. Since A and B are separated, $z \notin \overline{B}$, so $z \in \overline{B}^c$, which is an open set. Thus there is some $\epsilon > 0$ so that $B_\epsilon(z) \subseteq \overline{B}^c$, and so $z' = z + \frac{\epsilon}{2}$ is not in \overline{B} , therefore not

in B ; therefore $z' \neq y$. Also $z' > z$ which is an upper bound for A , so $z' \notin A$. Thus we have $x \leq z < z' < y$, and $z' \notin E = A \cup B$. Therefore E does not have the intermediate value property. □

Connectedness is a topological property: it is invariant under homeomorphisms. In fact, it is invariant under any continuous map.

THEOREM 7.14. *Let X and Y be metric spaces, and let $f: X \rightarrow Y$ be continuous. If $E \subseteq X$ is connected, then $f(E)$ is connected.*

PROOF. We argue the contrapositive: suppose that $f(E)$ is not connected. That is, $f(E)$ a union of two nonempty separated sets: $f(E) = A \cup B$ where $A, B \neq \emptyset$ and $\overline{A} \cap B = A \cap \overline{B} = \emptyset$. Then let $G = E \cap f^{-1}(A)$ and $H = E \cap f^{-1}(B)$. Since $E \subseteq f^{-1}(f(E)) = f^{-1}(A \cup B) = f^{-1}(A) \cup f^{-1}(B)$, it follows that $E = G \cup H$. We will show that G and H are nonempty and separated, and therefore E is not connected.

First, note that G and H are nonempty. Indeed, if G were empty, then $f^{-1}(A)$ would not intersect E , so no element of E is mapped into A . Since $f(E) = A \cup B$, this would mean that $f(E) = B$, and since $A \cap B = \emptyset$, that would imply $A = \emptyset$, contradicting the hypothesis. A similar argument shows that H is nonempty.

Now, since $A \subseteq \overline{A}$, $G \subseteq f^{-1}(A) \subseteq f^{-1}(\overline{A})$. Since \overline{A} is closed, $f^{-1}(\overline{A})$ is closed (this is the only point in the proof where continuity is used!), and it follows that $\overline{G} \subseteq f^{-1}(\overline{A})$, which is to say that $f(\overline{G}) \subseteq \overline{A}$. On the other hand, using Lemma 7.3(3), we have $f(H) = f(E \cap f^{-1}(B)) \subseteq f(E) \cap f(f^{-1}(B)) \subseteq (A \cup B) \cap B = B$. Thus, applying Lemma 7.3(3) once more,

$$f(\overline{G} \cap H) \subseteq f(\overline{G}) \cap f(H) \subseteq \overline{A} \cap B = \emptyset.$$

Since $\overline{G} \cap H$ is a set in the domain of f , it follows that $\overline{G} \cap H = \emptyset$. An entirely analogous argument demonstrates that $G \cap \overline{H} = \emptyset$.

Thus E is a union of two nonempty separated sets, and so E is not connected. □

Combining Theorem 7.14 with Proposition 7.13 (characterizing connected sets in \mathbb{R}) yields another big calculus theorem: the Intermediate Value Theorem.

COROLLARY 7.15 (Intermediate Value Theorem). *Let $a < b$ be in \mathbb{R} $f: [a, b] \rightarrow \mathbb{R}$ be continuous. For any y between $f(a)$ and $f(b)$, there exists and $x \in [a, b]$ where $f(x) = y$.*

PROOF. By Proposition 7.13, the interval $[a, b]$ is connected. Thus, by Theorem 7.14, $f([a, b])$ is also connected. Note that $f(a)$ and $f(b)$ are two points in $f([a, b])$; hence, by Proposition 7.13, if y is between $f(a)$ and $f(b)$, then $y \in f([a, b])$, which is precisely to say that there exists an $x \in [a, b]$ where $f(x) = y$. □

3. Lecture 3: April 5, 2016

It is tempting to think that the Intermediate Value Theorem actually characterizes continuous functions, but this is not so. To be precise: say that a function $f: [a, b] \rightarrow \mathbb{R}$ has the intermediate value property if, given any c, d with $a < c < d < b$, and any value y between $f(c)$ and $f(d)$, there is a point $x \in [c, d]$ with $f(x) = y$. Corollary 7.15 shows that any continuous function has the intermediate value property; but there are discontinuous functions that do as well (e.g. Example 7.18 below). To see how this can happen, we now turn to further discuss what kinds of discontinuities functions can have, continuing the discussion from Examples 6.11 and 6.12.

EXAMPLE 7.16. A piecewise continuous function defined on an interval in \mathbb{R} can have a “jump discontinuity” (see Definition 7.20 below), where the two functions on the two pieces do not “match up” at the transition point. For example, take $f: [0, 2] \rightarrow [0, 1]$ to be the functions defined piecewise by $f(x) = x$ for $0 \leq x < 1$ and $f(x) = x - 1$ for $1 \leq x \leq 2$. Then f is not continuous at $x = 1$: consider the two sequences $x_n = 1 - \frac{1}{n}$ and $y_n = 1 + \frac{1}{n}$, both of which converge to 1. Then $f(x_n) = x_n \rightarrow 1$ as $n \rightarrow \infty$, while $f(y_n) = y_n - 1 \rightarrow 0$ as $n \rightarrow \infty$. It therefore cannot be true that all sequences converging to 1 have images that converge to $f(1)$, since these two sequences, both converging to 1, have different image limits. (In this case $f(y_n) \rightarrow 0 = f(1)$; we will see below that f is *right-continuous*.)

Any function with a jump discontinuity will fail to have the intermediate value property as well. In the above example, consider the interval $[0.9, 1.1]$. The image of f here is $f([0.9, 1.1]) = f([0.9, 1)) \cup f([1, 1.1]) = [0.9, 1) \cup [0, 0.1]$, which does not include any points in the interval $(0.1, 0.9)$. But $f(0.9) = 0.9$ and $f(1.1) = 0.1$. So in this extreme example, *no point* between $f(0.9)$ and $f(1.1)$ is in the image of f on $[0.9, 1.1]$.


EXAMPLE 7.17. If we don't restrict ourselves to a one dimensional domain, there is plenty of room for weirder discontinuities. For example, consider the function $f: \mathbb{R}^2 \rightarrow \mathbb{R}$ defined by

$$f(x, y) = \begin{cases} \frac{xy}{x^2 + y^2} & \text{if } (x, y) \neq (0, 0) \\ 0 & \text{if } (x, y) = (0, 0) \end{cases}$$

This is a rational function, whose denominator only vanishes at $(0, 0)$; therefore, by the limit theorems, f is continuous on $\mathbb{R}^2 \setminus \{(0, 0)\}$. But the function is not continuous at $(0, 0)$. Consider the two sequences $a_n = (1/n, 1/n)$ and $b_n = (-1/n, 1/n)$. Both converge to $(0, 0)$, but

$$f(a_n) = \frac{\frac{1}{n} \cdot \frac{1}{n}}{\left(\frac{1}{n}\right)^2 + \left(\frac{1}{n}\right)^2} = \frac{1}{2}, \quad f(b_n) = \frac{-\frac{1}{n} \cdot \frac{1}{n}}{\left(-\frac{1}{n}\right)^2 + \left(\frac{1}{n}\right)^2} = -\frac{1}{2}.$$

Hence, there is no limit, since these two sequences, both converging to $(0, 0)$, have different image limits. In fact, notice that for any $m \in \mathbb{R}$, f is constant along the line $y = mx$: $f(x, mx) = \frac{x \cdot mx}{x^2 + (mx)^2} = \frac{m}{1+m^2}$. The range of the function $m \mapsto \frac{m}{1+m^2}$ is $[-\frac{1}{2}, \frac{1}{2}]$ (the endpoints are achieved at $m = \pm 1$ as computed above; the function is continuous and so all values between $\pm \frac{1}{2}$ are achieved by the Intermediate Value Theorem; and no other values are achieved since $0 \leq (1 - m)^2 = 1 + m^2 - 2m$, so $2m \leq 1 + m^2$). The graph of the function can be thought of as a “spiral / helical slide” wrapping around a vertical pole at the origin. While it is difficult to even state what an “intermediate value property” should mean in the context of a two-dimensional domain, it is clear that this function does not have a “jump” discontinuity: there are no gaps in the range of its values near 0. Below is a rendering of the graph of f on $[-1, 1]^2$; you should use software like Maple or Mathematica to explore this function further.




EXAMPLE 7.18. We can find examples of non-jump discontinuities for functions of a single real variable, too. The classic example is $f(x) = \sin(\frac{1}{x})$ for $x \neq 0$ and $f(0) = 0$ (or any value in $[-1, 1]$). A rendering of the graph of f on $[-\frac{2}{\pi}, \frac{2}{\pi}]$ can be found on the following page.

We have not formally introduced the sin function yet; if you want, you can replace it with $f(x) = \zeta(\frac{1}{x})$ for any periodic continuous function $\zeta: \mathbb{R} \rightarrow [-1, 1]$ with $\zeta(0) = 0$ to see a very similar picture (e.g. $\zeta(x) = x$ for $-1 \leq x \leq 1$ and $\zeta(x) = 2 - x$ for $1 \leq x \leq 3$, and then on any interval of the form $[4n - 1, 4n + 3]$ for $n \in \mathbb{Z}$ define $f(x) = \zeta(x - 4n)$). Since $x \mapsto \frac{1}{x}$ is continuous on $\mathbb{R} \setminus \{0\}$, and sin is continuous on \mathbb{R} , f is continuous on $\mathbb{R} \setminus \{0\}$ by Proposition 7.1. However, f is not continuous at 0. For example, let $x_n = \frac{1}{2n\pi + \pi/2}$ and $y_n = \frac{1}{2n\pi + 3\pi/2}$. Then $x_n \rightarrow 0$ and $y_n \rightarrow 0$, but $f(x_n) = \sin(2n\pi + \pi/2) = 1$ and $f(y_n) = \sin(2n\pi + 3\pi/2) = -1$, so f does not have a limit at 0. (Both of these sequences approach 0 from within $(0, \infty)$, so f does not have a “right limit” either.)

However, f *does* have the intermediate value property on any interval in \mathbb{R} . Since f is continuous on $\mathbb{R} \setminus \{0\}$, this is only interesting for intervals that include 0 in their interior, so let $a < 0 < b$. Consider the point $c = \frac{1}{1/b + 2\pi}$. Then $c < b$, and $f(c) = \sin(1/b + 2\pi) = \sin(1/b) = f(b)$. What’s more, on the interval $[1/b, 1/b + 2\pi]$, the function sin achieves its full range of values, $[-1, 1]$, and so $f([c, b]) = [-1, 1]$. This shows that $f([a, b]) \supset f([c, b]) = [-1, 1]$ as well. Since the range of f is $[-1, 1]$, any value y in between $f(a)$ and $f(b)$ is in $[-1, 1]$ and therefore is also in $f([a, b])$. Thus, f is a discontinuous function with the intermediate value property.

Let’s formalize the insights from the above examples into a characterization of different kinds of discontinuities. First we formally define left and right limits.

FIGURE 1. The graph of $\sin(\frac{1}{x})$.

DEFINITION 7.19. Let $a < b$ be in \mathbb{R} , and let $f: [a, b] \rightarrow \mathbb{R}$ be a function. Let $x_0 \in (a, b)$. Say that

$$\lim_{x \rightarrow x_0^-} f(x) = L$$

if, for any sequence (x_n) in $[a, x_0)$ with $x_n \rightarrow x_0$, $f(x_n) \rightarrow L$. For shorthand, we write $f(x_0^-) = L$. Similarly, we say

$$\lim_{x \rightarrow x_0^+} f(x) = R$$

if, for any sequence (y_n) in $(x_0, b]$ with $y_n \rightarrow x_0$, $f(y_n) \rightarrow R$. For shorthand, we write $f(x_0^+) = R$.

Equivalently: $f(x_0^-) = L$ means that the function $f|_{[a, x_0]}$ has limit L as $x \rightarrow x_0$, and $f(x_0^+) = R$ means that the function $f|_{[x_0, b]}$ has limit R as $x \rightarrow x_0$. You should verify that $\lim_{x \rightarrow x_0} f(x) = L$ is equivalent to $f(x_0^-) = f(x_0^+) = L$.

DEFINITION 7.20. Let $a < b$ in \mathbb{R} , and let $x_0 \in (a, b)$. A function $f: [a, b] \rightarrow \mathbb{R}$ is said to have a jump discontinuity at x_0 if $f(x_0^-)$ and $f(x_0^+)$ exists, but $f(x_0^-) \neq f(x_0^+)$. In this case, we say that f has a jump of size $f(x_0^+) - f(x_0^-)$.

If f is discontinuous at x_0 , but does not have a jump discontinuity at x_0 , we say f has a non-jump discontinuity at x_0 .

REMARK 7.21. Rudin uses the term “simple discontinuity” for jumps, and “discontinuity of the second kind” for non-jump discontinuities. This terminology is silly, for two reasons. First,

no working analyst has used that terminology in at least 60 years. Second, and more importantly, the distinction between the two kinds of discontinuities is very much like dividing the world into bananas and non-bananas; jump discontinuities are rare. Thus “jumps” vs. “non-jumps” is a more accurate division.

REMARK 7.22. It is almost correct to characterize non-jump discontinuities as those where either the left or the right limit does not exist. The one kind of exception to this is a kind of degenerate discontinuity, where $\lim_{t \rightarrow x_0} f(t)$ exists, but is not equal to $f(x_0)$. (For example: $f(t) = 0$ for all $t \neq x_0$, but $f(x_0) = 1$.) Such a discontinuity does not count as a jump discontinuity by the above definition, and indeed it should not: the function is not jumping anywhere. (This is an important distinction, for example in the theory of paths of stochastic processes.) But it is also quite a bit milder than the discontinuity of Example 7.18. Such examples are called *removable discontinuities*. So non-jump discontinuities are partitioned into removable ones, and the truly bad ones, which we might call *oscillatory discontinuities*.

In example 7.16, the function has a jump discontinuity of size -1 at 0. In example 7.18, the function has a non-jump discontinuity at 0. In Example 6.11 (Dirichlet’s function, the indicator function of the rationals), *every* point is a non-jump discontinuity.

There is one important class of functions (which we will use extensively later in this course, when we study integration theory) in which only jump discontinuities can occur.

DEFINITION 7.23. Let $a < b$ in \mathbb{R} . a function $f: [a, b] \rightarrow \mathbb{R}$ is said to be monotone increasing or nondecreasing if, for all $x < y$ in $[a, b]$, $f(x) \leq f(y)$; it is said to be monotone decreasing or nonincreasing if, for all $x < y$ in $[a, b]$, $f(x) \geq f(y)$. If f is either monotone increasing or monotone decreasing, f is called monotone.

The functions in Examples 6.11, 6.12, and 7.18 are *not* monotone; all of them oscillate wildly. In all cases, the discontinuities were not jump discontinuities. Example 7.16 is also not monotone, but it is milder (just one oscillation), and its discontinuity is a jump discontinuity. In general, monotone functions can be discontinuous (e.g. like Example 7.16, but with a positive jump: $f(x) = x$ for $0 \leq x < 1$ and $f(x) = x + 1$ for $1 \leq x \leq 2$), but that’s as bad as they get.

PROPOSITION 7.24. Let $a < b$ in \mathbb{R} , and let $f: (a, b) \rightarrow \mathbb{R}$ be a monotone function. Then $f(x-)$ and $f(x+)$ exists for each $x \in (a, b)$; thus, all discontinuities of f are jump discontinuities. If f is monotone increasing, then for all $x \in (a, b)$,

$$\sup_{a < t < x} f(t) = f(x-) \leq f(x) \leq f(x+) = \inf_{x < t < b} f(t). \quad (7.1)$$

Also $f(x+) \leq f(y-)$ whenever $a < x < y < b$. For monotone decreasing f , the inequalities (and the sup and inf) are reversed.

PROOF. We assume f is monotone increasing; the proof for the monotone decreasing case is analogous. But assumption, for any $t < x$, $f(t) \leq f(x)$. So the set $\{f(t) : a < t < x\}$ is bounded above by $f(x)$, which shows that $\alpha = \sup_{a < t < x} f(t) \leq f(x)$. Fix $\epsilon > 0$. Since α is the *least* upper bound, there must exist some $x' \in (a, x)$ with $\alpha - \epsilon < f(x') \leq \alpha$. Since f is monotone increasing, if $x' < t < x$, $f(x') \leq f(t) \leq \alpha$. Combining these last two inequalities yields

$$|f(t) - \alpha| < \epsilon, \quad \text{for } x' < t < x.$$

Hence, if $t_n \in (a, x)$ and $t_n \rightarrow x$, there is some N so that $x' < t_n < x$ for $n \geq N$, and so $|f(t_n) - \alpha| < \epsilon$ for $n \geq N$; this shows that $f(x-) = \lim_{t \rightarrow x-} f(t) = \alpha$, as claimed. The proof that $f(x) \leq f(x+) = \inf_{x < t < b} f(t)$ is very similar; thus we have proven (7.1) true.

Now, if $a < x < y < b$, then since $f(s) \geq f(y)$ for all $s > y$, we have

$$f(x+) = \inf_{x < t < b} f(t) = \inf_{x < t < y} f(t).$$

Applying the left-hand-side of (7.1) to the interval (x, y) instead of (a, b) shows that

$$f(x+) = \inf_{x < t < y} f(t) \leq \sup_{x < t < y} f(t) = f(y-)$$

which concludes the proof of the second inequality. The proofs for monotone decreasing functions are analogous. \square

Differentiation of Functions of a Real Variable

1. Lecture 4: April 7, 2016

Almost everything that was discussed in Chapters 5, 6, and 7 applied quite generally to functions / sequences in metric spaces. In this chapter, we will focus on concepts that are exclusive to functions defined on (intervals in) \mathbb{R} . Some of these concepts can be extended to more general spaces, but not the level of generality we've seen so far: taking derivatives involves more underlying structure than most metric spaces afford.

Let $a < b$ in \mathbb{R} , and suppose $f: [a, b] \rightarrow \mathbb{R}$ is a function. For any $x, y \in [a, b]$ with $x \neq y$, we may define the *difference quotient* $(DQf)(x, y)$ as follows:

$$(DQf)(x, y) \equiv \frac{f(x) - f(y)}{x - y}.$$

EXAMPLE 8.1. If $f(x) = x$, then $(DQf)(x, y) = \frac{x-y}{x-y} = 1$ for all $x \neq y$. If $f(x) = x^2$, we can factor to find that

$$(DQf)(x, y) = \frac{x^2 - y^2}{x - y} = \frac{(x - y)(x + y)}{x - y} = x + y, \quad x \neq y.$$

In general, for power functions $f(x) = x^k$, since $x^k - y^k = (x - y)(x^{k-1} + x^{k-2}y + x^{k-3}y^2 + \dots + xy^{k-2} + y^{k-1})$, we see that $(DQf)(x, y)$ is a polynomial (of homogeneous degree $k - 1$) in x and y . It is, a priori, only defined when $x \neq y$.

We will use the difference quotient DQf to define the derivative of f , by taking limits, as you no doubt recall. Before we do so, it is instructive to note a few nice properties that the difference quotient has. (Note that *every* function has a difference quotient.)

LEMMA 8.2. Let $a < b$ in \mathbb{R} , and let $f, g: [a, b] \rightarrow \mathbb{R}$ be functions. For any $x \neq y$ in $[a, b]$, we have the following.

- (1) $DQ(f + g)(x, y) = DQf(x, y) + DQg(x, y)$.
- (2) $DQ(fg)(x, y) = f(x)DQg(x, y) + DQf(x, y)g(y)$.
- (3) If $g(x) \neq 0$ for $x \in [a, b]$, then $DQ\left(\frac{1}{g}\right)(x, y) = -\frac{1}{g(x)g(y)}DQg(x, y)$.

PROOF. (1) is an immediate consequence of the distributive properties of addition and multiplication, and is left to the reader. For (2), we calculate as follows:

$$\begin{aligned} f(x)g(x) - f(y)g(y) &= f(x)g(x) - f(x)g(y) + f(x)g(y) - f(y)g(y) \\ &= f(x)(g(x) - g(y)) + (f(x) - f(y))g(y). \end{aligned}$$

Dividing through by $x - y$ yields the result. For (3), we compute

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

□

REMARK 8.3. Since $fg = gf$, (2) implies that we also have $DQ(fg)(x, y) = f(y)DQg(x, y) + g(x)DQf(x, y)$. One can see this as well in the proof of (2), by introducing the opposite cross terms $-f(y)g(x) + f(y)g(x)$.

You should recognize the statements of Lemma 8.2 as versions the usual rules of differentiation; e.g. (2) is the product rule. It is illuminating to note that these are simply algebraic rules that hold for difference quotients, having nothing to do with the limits involved in taking derivatives. In that vein, we can also prove a difference quotient version of the chain rule; this is particularly instructive, since it highlights a technical difficulty we will have to overcome in proving the chain rule for derivatives.

LEMMA 8.4. *Let $a < b$ in \mathbb{R} , and let $g: [a, b] \rightarrow \mathbb{R}$ be a function. Suppose f is a real valued function defined on the range of g ; then $f \circ g: [a, b] \rightarrow \mathbb{R}$ is a function. If $x \neq y$ in $[a, b]$, and if $g(x) \neq g(y)$, then*

$$DQ(f \circ g)(x, y) = DQf(g(x), g(y)) \cdot DQg(x, y).$$

PROOF. Since $g(x) \neq g(y)$, we can multiply and divide through by it, to see that

$$DQ(f \circ g)(x, y) = \frac{f \circ g(x) - f \circ g(y)}{x - y} = \frac{f(g(x)) - f(g(y))}{g(x) - g(y)} \frac{g(x) - g(y)}{x - y}$$

and, by the definition of DQ , this gives the desired result. \square

It is possible that $g(x) = g(y)$, even if $x \neq y$; in this case, $DQf(g(x), g(y))$ is not defined, since $DQf(u, v)$ is only defined when $u \neq v$. Of course, we would like to define it even on the diagonal $u = v$ by taking the limit as $v \rightarrow u$; we will presently do this to define the derivative. But the reader should be wary of the above proof: it is possible that $g(x) = g(y)$ for x and y arbitrarily close to each other, in which case it will take some care to make sense of the limit; we will do this below in Proposition 8.12.

We now come to the definition of the derivative.

DEFINITION 8.5. *Let $a < b$ in \mathbb{R} , and let $f: [a, b] \rightarrow \mathbb{R}$. For any point $x \in [a, b]$, say that f is differentiable at x if the limit $\lim_{t \rightarrow x} DQf(x, t)$ exists; in this case, we call this limit the derivative of f at x , and denote it*

$$f'(x) = \frac{df}{dx} \equiv \lim_{t \rightarrow x} \frac{f(x) - f(t)}{x - t}.$$

REMARK 8.6. Some authors insist that $x \in (a, b)$ for this definition. As stated above, we highlight the fact that differentiability (like continuity) depends on the domain of the function. At $x = a$, the above limit is actually a right limit $\lim_{t \rightarrow x^+} DQf(x, t)$, and we might call $f'(a)$ (if it exists) the *right derivative* of f at a . If f is actually defined on a larger interval including a in its interior, then it is possible that f is not differentiable at a , even if it is right differentiable there. Similar comments apply to *left* differentiability at b .

Differentiability is a regularity property of a function at a point. It is stronger than continuity.

PROPOSITION 8.7. *Let $a < b$ in \mathbb{R} , and let $f: [a, b] \rightarrow \mathbb{R}$. If $x \in [a, b]$ and f is differentiable at x , then f is continuous at x .*

PROOF. Let (x_n) be a sequence in $[a, b] \setminus \{x\}$ with $x_n \rightarrow x$. Then we can multiply and divide by $x_n - x$, and so

$$f(x) - f(x_n) = \frac{f(x) - f(x_n)}{x - x_n} (x - x_n) = DQf(x, x_n)(x - x_n).$$

By assumption, $f'(x)$ exists, which means that $\lim_{n \rightarrow \infty} DQf(x, x_n) = f'(x)$. We also have $\lim_{n \rightarrow \infty} (x - x_n) = 0$. Thus, by the limit theorems,

$$\lim_{n \rightarrow \infty} [f(x_n) - f(x)] = \lim_{n \rightarrow \infty} DQf(x, x_n) \cdot \lim_{n \rightarrow \infty} (x - x_n) = f'(x) \cdot 0 = 0.$$

This shows $\lim_{n \rightarrow \infty} f(x_n) = f(x)$, for any such sequence. By definition, this means $\lim_{t \rightarrow x} f(t) = f(x)$, which is the definition of continuity of f at x . \square

REMARK 8.8. In the case $x = a$, differentiability (which really means *right* differentiability) only implies right continuity; similarly differentiability at $x = b$ (which really means *left* differentiability) only implies left continuity.

The converse of Proposition 8.7 is very false.

EXAMPLE 8.9. Consider the function $f(x) = |x|$ defined on \mathbb{R} . This function is continuous on \mathbb{R} , and in particular at 0: if $x_n \rightarrow 0$ then $|x_n| \rightarrow 0$. Now, for $x \neq 0$, $DQf(0, t) = \frac{0-|t|}{0-t} = \frac{|t|}{t}$. This is either $+1$ if $t > 0$ or -1 if $t < 0$. So take, for example, the sequence $t_n = \frac{(-1)^n}{n}$, which tends to 0. The sequence $DQf(0, t_n) = (-1)^n$ does not converge as $n \rightarrow \infty$. This means that the limit $f'(0)$ does not exist, so f is not differentiable at 0.

REMARK 8.10. In Example 8.9, the function f fails to be differentiable only at a single point. It was long thought (by Newton and others prior to the 19th Century) that non-differentiable points of continuous functions were all of this nature: that continuous functions would have to be differentiable except at a *discrete* set of points. This is extremely far from true. Later on, we will see examples of functions that are continuous everywhere on \mathbb{R} , but differentiable *nowhere*. In a sense, such functions are the most important, to the modern theory of analysis and probability.

Let us now state the well-known differentiation rules that follow from Lemma 8.2.

PROPOSITION 8.11. *Let $a < b$ in \mathbb{R} , and let $f, g: [a, b] \rightarrow \mathbb{R}$. Suppose f and g are differentiable at a point $x \in [a, b]$. The following hold true.*

- (1) $f + g$ is differentiable at x , and $(f + g)'(x) = f'(x) + g'(x)$.
- (2) fg is differentiable at x , and $(fg)'(x) = f'(x)g(x) + f(x)g'(x)$.
- (3) If $g(x) \neq 0$, then $\frac{1}{g}$ is differentiable at x , and $(\frac{1}{g})'(x) = -\frac{g'(x)}{g(x)^2}$.

PROOF. (1) follow immediately from Lemma 8.2(1), using the limit theorem (a limit of a sum is the sum of the limits). For (2), using Lemma 8.2(2), we have

$$\begin{aligned} (fg)'(x) &= \lim_{t \rightarrow x} DQ(fg)(x, t) = \lim_{t \rightarrow x} [f(x)DQg(x, t) + DQf(x, t)g(t)] \\ &= f(x) \lim_{t \rightarrow x} DQg(x, t) + \lim_{t \rightarrow x} DQf(x, t) \cdot \lim_{t \rightarrow x} g(t) \end{aligned}$$

The limits of the difference quotients are, by definition, $f'(x)$ and $g'(x)$. By Proposition 8.7, g is continuous at x , and therefore $\lim_{t \rightarrow x} g(t) = g(x)$, yielding the desired formula. For (3), we note (again by Proposition 8.7) that g is continuous at x . Since $g(x) \neq 0$, it follows that, for some $\delta > 0$, $g(t) \neq 0$ for all $t \in (x - \delta, x + \delta) \cap [a, b]$; indeed, from the ϵ - δ definition of continuity at x , this follows by taking any $\epsilon < |g(x)|$. Restricting to this neighborhood of x , the function $\frac{1}{g}$ is well defined, and we may apply Lemma 8.2(3) and the limit theorems to find that

$$\left(\frac{1}{g}\right)'(x) = \lim_{t \rightarrow 0} DQg(x, t) = \lim_{t \rightarrow 0} \left(-\frac{1}{g(x)g(t)} DQg(x, t)\right) = -\frac{1}{g(x)} \cdot \lim_{t \rightarrow x} \frac{1}{g(t)} \cdot \lim_{t \rightarrow x} DQg(x, t).$$

Since g is continuous and nonvanishing on a neighborhood of x , $\frac{1}{g}$ is continuous at x , and the first limit is $\frac{1}{g(x)}$; the second limit is $g'(x)$ by definition. This concludes the proof. \square

Part (3) above is a special case of the so-called *quotient rule*, which states that

$$\left(\frac{f}{g}\right)'(x) = \frac{f'(x)g(x) - f(x)g'(x)}{g(x)^2}$$

under the conditions stated above. This follows immediately by combining (2) and (3), noting that $\frac{f}{g} = f \cdot \frac{1}{g}$; the details are left to the reader.

We now come to the chain rule.

PROPOSITION 8.12. *Let $a < b$ in \mathbb{R} , and let $g: [a, b] \rightarrow \mathbb{R}$. Suppose $x \in [a, b]$ and g is differentiable at x . Let f be defined on the range of g , and in particular on a neighborhood of $g(x)$, and suppose that f is differentiable at $g(x)$. Then $f \circ g$ is differentiable at x , and $(f \circ g)'(x) = f'(g(x)) \cdot g'(x)$.*

As mentioned above, the proof is technically more challenging than might at first seem necessary. The straightforward approach would be to take the statement of Lemma 8.4 and apply the limit theorems to derive the stated formula:

$$\lim_{t \rightarrow x} DQ(f \circ g)(x, t) = \lim_{t \rightarrow x} DQf(g(x), g(t)) \cdot \lim_{t \rightarrow x} Dg(x, t) = \lim_{t \rightarrow x} DQf(g(x), g(t)) \cdot g'(x)$$

For the first limit, we would like to say that since g is continuous at x , $g(t) \rightarrow g(x)$ as $t \rightarrow x$, and so this limit is the same as $\lim_{s \rightarrow g(x)} Df(g(x), s) = f'(g(x))$. The problem with this is that the quantity $Df(g(x), g(t))$ is only well-defined if $g(x) \neq g(t)$, and it is perfectly possible for g to have the property that $g(t) = g(x)$ for many t s, arbitrarily close to x (e.g. g could be constant near x ; or, technically worse, g could oscillate fast near x). Therefore, $\lim_{t \rightarrow x} DQf(g(x), g(t))$ may not be defined since, for $\lim_{t \rightarrow x} u(t)$ to make sense, u must be defined on $(x - \delta, x + \delta) \setminus \{x\}$ (intersected with $[a, b]$). Therefore, we must be much more careful in proving the chain rule.

PROOF. By definition, $\lim_{t \rightarrow x} DQg(x, t) = g'(x)$; so if we set $u_x(t) = DQg(x, t) - g'(x)$, we have $u_x(t) \rightarrow 0$ as $t \rightarrow x$. Similarly, setting $y = g(x)$, we have $\lim_{s \rightarrow y} DQf(y, s) = f'(y)$; so if we set $v_y(s) = DQf(y, s) - f'(s)$, we have $v_y(s) \rightarrow 0$ as $s \rightarrow y$. Now, unraveling the definition of the difference quotient, we have

$$\begin{aligned} g(x) - g(t) &= (x - t)DQg(x, t) = (x - t)[g'(x) + u_x(t)] \\ f(y) - f(s) &= (y - s)DQf(y, s) = (y - s)[f'(y) + v_y(s)]. \end{aligned}$$

Here t is in a neighborhood of x (where g is defined) and s is in a neighborhood of y (where f is defined). Composing, we have $f \circ g(x) - f \circ g(t) = f(y) - f(g(t))$. Since g is continuous, when t is in a small enough neighborhood of x , $s = g(t)$ is in the given neighborhood of $y = g(x)$, and so

$$\begin{aligned} f \circ g(x) - f \circ g(t) &= f(y) - f(s) = (y - s)[f'(y) + v_y(s)] \\ &= (g(x) - g(t))[f'(y) + v_y(s)] \\ &= (x - t)[g'(x) + u_x(t)][f'(y) + v_y(s)] \end{aligned}$$

For all $t \neq x$, we can then divide through by $x - t$ and we see that

$$DQ(f \circ g)(x, t) = [g'(x) + u_x(t)][f'(y) + v_y(s)] = [g'(x) + u_x(t)][f'(y) + v_y(g(t))].$$

As $t \rightarrow x$, $u_x(t) \rightarrow 0$. Also, g is continuous at x , so as $t \rightarrow x$, $g(t) \rightarrow g(x) = y$, and so $v_y(g(t)) \rightarrow 0$. It now follows from the limit theorems that $\lim_{t \rightarrow x} DQ(f \circ g)(x, t) = [g'(x) + 0][f'(y) + 0] = g'(x)f'(g(x))$, as desired. \square

2. Lecture 5: April 12, 2016

As you'll recall from calculus, one of the main applications of derivatives is in the study of extrema.

DEFINITION 8.13. *Let X, Y be metric spaces, and let $f: X \rightarrow Y$. A point $x \in X$ is called a local maximizer of f if there is a positive radius $\delta > 0$ so that, for all $t \in B_\delta(x)$, $f(t) \leq f(x)$; in this case the value $f(x)$ is called a local maximum. Similarly, x is a local minimizer if there is a $\delta > 0$ so that, for all $t \in B_\delta(x)$, $f(t) \geq f(x)$; in this case the value $f(x)$ is called local minimum. The local maxima and minima of f are called its local extrema, and the points at which they occur are local extremizers.*

For a $f: \mathbb{R} \rightarrow \mathbb{R}$, local extrema of f can be determined by locating the points x where $f'(x) = 0$.

THEOREM 8.14. *Let $a < b$ in \mathbb{R} , and suppose $f: (a, b) \rightarrow \mathbb{R}$ be a function. If f has a local extremum at $x \in (a, b)$, and if $f'(x)$ exists, then $f'(x) = 0$.*

PROOF. We will assume x is a local maximizer; the argument for a local minimizer is analogous. By assumption, there is a $\delta > 0$ such that $a < x - \delta < x + \delta < b$ and $f(x) \geq f(t)$ for all $t \in (x - \delta, x + \delta)$. In particular, this means that if $x - \delta < t < x$, then $x - t > 0$ while $f(x) - f(t) \geq 0$, and so $DQf(x, t) \geq 0$ in this interval; likewise, if $x < t < x + \delta$, then $x - t < 0$ while $f(x) - f(t) \geq 0$, so here $DQf(x, t) \leq 0$. Thus, using the squeeze theorem, we have

$$\lim_{t \rightarrow x^-} DQf(x, t) \leq 0 \quad \text{and} \quad \lim_{t \rightarrow x^+} DQf(x, t) \geq 0.$$

By assumption $f'(x) = \lim_{t \rightarrow x} DQf(x, t) = \lim_{t \rightarrow x^-} DQf(x, t) = \lim_{t \rightarrow x^+} DQf(x, t)$; thus, this common value is both ≥ 0 and ≤ 0 . It follows that $f'(x) = 0$, as claimed. \square

As you'll recall, points x where $f'(x) = 0$ are called critical points. The content of Theorem 8.14 is that, if a function is known to be differentiable at all points in its domain, then its local extrema all occur at critical points. This is one of the core tools in calculus, and we will use it in many theoretical applications as well. In order for it to be useful, of course, we must know that our function is differentiable not just at certain points but everywhere: for example, the function $f(x) = |x|$ attains its global minimum (which is its only local extremum) at the point $x = 0$, where f is not differentiable. We will thus be most interested, for now, in functions that are differentiable everywhere.

Like continuity, differentiability is a local property: we speak of differentiability of a function at a point x in its domain (and this only depends on the behavior of f in an arbitrarily small neighborhood of x). Also like continuity, we can then boost this to a global property, by talking about f being differentiable on a set (i.e. differentiable at all points of a given set). In the case of derivatives, this produces a *new function*.

DEFINITION 8.15. *Let $a < b$ in \mathbb{R} , and let $f: [a, b] \rightarrow \mathbb{R}$ be differentiable (at all points of its domain). We let f' denote the new function $f': [a, b] \rightarrow \mathbb{R}$ whose value at x is $f'(x) = \lim_{t \rightarrow x} DQf(x, t)$.*

Here are two examples that show how f' can be poorly behaved (at points) even when f is quite well behaved. Both examples take for granted the functions \sin and \cos , which are differentiable on \mathbb{R} and satisfy $\sin' = \cos$ and $\cos' = -\sin$. We will formally develop these later this quarter.

EXAMPLE 8.16. Define a function $f: \mathbb{R} \rightarrow \mathbb{R}$ as follows:

$$f(x) = \begin{cases} x \sin \frac{1}{x} & x \neq 0 \\ 0 & x = 0 \end{cases}$$

First of all, note that f is continuous at all $x \neq 0$ since it is a composition of continuous functions there; and at $x = 0$, note that $|\sin \frac{1}{x}| \leq 1$ for all x , so $|f(x)| \leq |x| \rightarrow 0 = f(0)$ as $x \rightarrow 0$. So f is continuous on \mathbb{R} . What's more, using the rules of differentiation, we can compute that f is differentiable at all points $x \neq 0$, and for such points

$$f'(x) = \sin \frac{1}{x} + x \cos \frac{1}{x} \cdot \left(-\frac{1}{x^2}\right) = \sin \frac{1}{x} - \frac{1}{x} \cos \frac{1}{x}.$$

However, at 0, we have

$$DQf(0, t) = \frac{0 - f(t)}{0 - t} = \sin \frac{1}{t}$$

and this function has no limit as $t \rightarrow 0$ (cf. Example 7.18). Thus, $f'(0)$ does not exist. Also, the formula for $f'(x)$ for $x \neq 0$ involves terms like $\sin \frac{1}{x}$ and $\frac{1}{x} \cos \frac{1}{x}$ which are at least as badly behaved as $\sin \frac{1}{x}$. So there is no way to extend the function f' to be defined at 0 in such a way that it will be continuous: f' has a non-jump discontinuity at 0.




FIGURE 1. The graph of $x \sin(\frac{1}{x})$; the envelope is given by the lines $y = \pm x$.

EXAMPLE 8.17. Now consider the following function:

$$g(x) = \begin{cases} x^2 \sin \frac{1}{x} & x \neq 0 \\ 0 & x = 0 \end{cases}$$

An argument very similar to the one in Example 8.16 shows that g is continuous on \mathbb{R} , and differentiable everywhere except possibly at 0; here

$$g'(x) = 2x \sin \frac{1}{x} - \cos \frac{1}{x}.$$

Now, looking at $x = 0$ specifically, we have

$$DQg(0, t) = \frac{0 - f(t)}{0 - t} = \frac{t^2 \sin \frac{1}{t}}{t} = t \sin \frac{1}{t} = f(t) \rightarrow 0 \text{ as } t \rightarrow 0$$

where f is the function from Example 8.16, which is continuous at 0 as shown above. Thus, g is actually differentiable on all of \mathbb{R} , with $f'(0) = 0$. But the formula for g' at points other than 0 shows that g is not continuous on \mathbb{R} : it has a non-jump discontinuity at 0. Ergo: it is possible for a function to be continuous and differentiable everywhere, but for its derivative to have a (bad) discontinuity somewhere.




FIGURE 2. The graph of $x^2 \sin(\frac{1}{x})$; the envelope is given by the curves $y = \pm x^2$.

If we want to use tools like Theorem 8.14, functions like the one in Example 8.16 are out; functions like the one in Example 8.17, on the other hand, are eligible, despite the bad behavior of the derivative as a function.

We now come to the most important theorem on all of calculus: the Mean Value Theorem.

THEOREM 8.18 (Mean Value Theorem). *Let $a < b$ in \mathbb{R} , and let $f: [a, b] \rightarrow \mathbb{R}$ be a continuous function. Suppose also that f is differentiable on (a, b) . Then there is a point $x \in (a, b)$ such that $f'(x) = DQf(a, b)$; i.e.*

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

Notice that we do not even need to assume that f is (one-sided) differentiable at the endpoints; it need only be continuous there.

PROOF. Define a new function $h(t) = (b - a)f(t) - t[f(b) - f(a)]$. First, notice that

$$h(a) = (b - a)f(a) - a[f(b) - f(a)] = bf(a) - af(b),$$

$$h(b) = (b - a)f(b) - b[f(b) - f(a)] = bf(a) - af(b).$$

So $h(a) = h(b)$.

Now, by the differentiation rules, h is differentiable on (a, b) and $h'(t) = (b-a)f'(t) - [f(b) - f(a)]$, and so our goal is to show that $h'(x) = 0$ for some $x \in (a, b)$. First, consider the case that h is constant: then $h'(x) = 0$ for all $x \in (a, b)$, and we're done. Otherwise, there is some point $t \in (a, b)$ such that $h(t) \neq h(a) = h(b)$. Thus, either $h(t) > h(a)$ or $h(t) < h(a)$. For the moment, we assume the former: there is some point $t \in (a, b)$ where $h(t) > h(a)$.

Since f is continuous on $[a, b]$, so is h by the limit theorems. Since $[a, b]$ is compact, by the Extreme Values theorem, there is a point $x \in [a, b]$ where $h(x) = \max\{f(t) : t \in [a, b]\}$. As $h(a) = h(b)$ and this value is *not* the maximum (by assumption that $h(t) > h(a)$ for some $t \in (a, b)$), it follows that $x \in (a, b)$. As h is differentiable in (a, b) , by Theorem 8.14 it follows that $h'(x) = 0$, and this shows the desired conclusion in this case.

For the case that we only know that $h(t) < h(a)$ for some $t \in (a, b)$, the argument is analogous, using the minimum, rather than the maximum, of h . \square

To demonstrate why we called the Mean Value Theorem the most important theorem in calculus, note the following corollary which encapsulates most of the material (on curve sketching, etc.) which follows immediately from it.

COROLLARY 8.19. *Let $a < b$ in \mathbb{R} , and let $f : [a, b] \rightarrow \mathbb{R}$ be continuous and differentiable on (a, b) .*

- (1) *If $f' \geq 0$ then f is monotone increasing.*
- (2) *If $f' = 0$ then f is constant.*
- (3) *If $f' \leq 0$ then f is monotone decreasing.*

PROOF. All three parts of the proof follow from the fact that, for any $x_1, x_2 \in [a, b]$, by the Mean Value Theorem there is a point $x \in (x_1, x_2)$ such that

$$f(x_2) - f(x_1) = (x_2 - x_1)f'(x).$$

For example, in case (2) where $f'(x) = 0$ for all x , this shows $f(x_2) - f(x_1) = 0$ for all $x_1, x_2 \in [a, b]$, which is precisely to say that f is constant. The other two cases are similar. \square

In addition to practical applications like curve sketching, the mean value theorem allows us to see that, while derivative functions f' can be quite irregular (e.g. Example 8.17, there are constraints. While derivative functions need not be continuous, they always have the intermediate value property.

PROPOSITION 8.20 (Darboux). *Let $a < b$ in \mathbb{R} , and suppose $f : (a, b) \rightarrow \mathbb{R}$ is differentiable. Then f has the intermediate value property on (a, b) : for any $x_1 < x_2$ in (a, b) , if y be a real number between $f'(x_1)$ and $f'(x_2)$, then there is a point $x \in (x_1, x_2)$ such that $f'(x) = y$.*

PROOF. Without loss of generality, we suppose $f'(x_1) < y < f'(x_2)$; the reverse ordering case is very similar. Set $g(t) = f(t) - yt$. Since f is differentiable on (a, b) , it is continuous on $[x_1, x_2]$, and so therefore is g . As $[x_1, x_2]$ is compact, by the Extreme Values theorem, g attains its minimum value at some point $x \in [x_1, x_2]$. Now, g is also differentiable on (a, b) , and $g'(t) = f'(t) - y$. By assumption, $g'(x_1) = f'(x_1) - y < 0$ and $g'(x_2) = f'(x_2) - y > 0$.

We claim that $x_1 < x < x_2$. To prove this, we just need to show that neither x_1 nor x_2 is a minimizer for g . Indeed, consider x_1 : we have $\lim_{t \rightarrow x_1} DQg(x_1, t) = g'(x_1) < 0$; i.e.

$$\lim_{t \rightarrow x_1} \frac{g(t) - g(x_1)}{t - x_1} < 0.$$

It follows that, for all t sufficiently close to x_1 , $DQg(x_1, t) < 0$. In particular, for $t > x_1$ and sufficiently small, multiplying through by the positive $t - x_1$ yields $g(t) - g(x_1) < 0$. This shows that $g(x_1)$ is not the minimum of g on $[x_1, x_2]$. An analogous argument at x_2 , using the fact that $g'(x_2) > 0$, shows that x_2 is not a minimizer for g . Thus, $x \in (x_1, x_2)$ as claimed.

Thus, by the Mean Value Theorem (or in fact by Theorem 8.14, which is equivalent to it), $g'(x) = 0$. This shows that $0 = g'(x) = f'(x) - y$, and so $f'(x) = y$ as claimed. \square

REMARK 8.21. This theorem was proved by Darboux in the late 19th Century. Until that point, it was widely believed that only continuous functions could have the intermediate value property. But Darboux proved Proposition 8.20, and then gave examples (like Example 8.17 to show that this is not the case.

COROLLARY 8.22. *If f is differentiable, then f' cannot have any jump discontinuities.*

PROOF. As mentioned just after Example 7.16, any function with a jump discontinuity must fail to have the intermediate value property (in a neighborhood of the jump). \square

3. Lecture 6: April 14, 2016

As we saw in the previous lecture, a differentiable function f can have a derivative f' which is not very regular. The function f' cannot be arbitrary (as Darboux's theorem shows, f' must at least have the intermediate value property), but f' certainly need not be continuous. In many circumstances, it is natural to assume that f' is continuous. For example, in many calculus applications, we use the *second* derivative, meaning we actually assume that the function f' is differentiable (ergo continuous).

DEFINITION 8.23. *Let $a < b$ in \mathbb{R} , and let $f: [a, b] \rightarrow \mathbb{R}$ be differentiable. If f' is a continuous function on $[a, b]$, we say f is continuously differentiable, and write $f \in C^1[a, b]$. More generally, for a positive integer k , we say $f \in C^k[a, b]$ if $f^{(j)}$ is continuous for $0 \leq j \leq k$, where $f^{(j)}$ is defined recursively by $f^{(j)} = (f^{(j-1)})'$, and $f^{(0)} = f$.*

Example 8.17 shows that a function f can be differentiable (on all of \mathbb{R} , even) without being C^1 . The difference is subtle but important. It is analogous to the difference between continuous and uniformly continuous functions: continuity is a local property boosted to a global property pointwise, while uniform continuity is a truly global property. Similarly, being differentiable is a local property boosted to a global property pointwise, while being C^1 is a truly global property. (Warning: this is not a perfect analogy. It is, for example, *not* true that a differentiable function is continuously differentiable on a compact interval.)

Recall our motivation for introducing differentiability. Continuity of f at a point x is the statement that $f(x+t) - f(x)$ tends to 0 as $h \rightarrow 0$, but at what rate? If f is differentiable, the answer is that there is a correction factor, the linear function $h \mapsto f'(x)t$, so that

$$f(x+t) - f(x) - f'(x)t = o(t) \quad (8.1)$$

where the notation $\alpha(t) = \beta(t) + o(t)$ means that $\lim_{t \rightarrow 0} \frac{\alpha(t) - \beta(t)}{t} = 0$. Thus, differentiability implies that, up to a linear correction, $f(x+t)$ is closer to $f(x)$ than any linear function: the difference goes to 0 faster than t as $t \rightarrow 0$. This leads us to ask what happens if there are higher derivatives? Is the difference even smaller, modulo higher order corrections? The answer is yes, which is the statement of Taylor's theorem.

THEOREM 8.24. *[Taylor's Theorem] Let $a < b$ in \mathbb{R} , let $k \in \mathbb{N}$, and let $f: (a, b) \rightarrow \mathbb{R}$ be C^{k-1} , such that $f^{(k)}$ exists (but need not be differentiable) on (a, b) . Define the Taylor polynomial $T_x^{k-1}f$ by*

$$(T_x^{k-1}f)(t) = f(x) + f'(x)t + \frac{1}{2}f''(x)t^2 + \cdots + \frac{1}{(k-1)!}f^{(k-1)}(x)t^{k-1} = \sum_{j=0}^{k-1} \frac{f^{(j)}(x)}{j!}t^j.$$

For each t such that $x+t \in (a, b)$, there exists a point ξ between x and $x+t$ such that

$$f(x+t) = (T_x^{k-1}f)(t) + \frac{1}{k!}f^{(k)}(\xi)t^k.$$

The theorem is often stated using the variable $y = x+t$ as $f(y) = (T_x^{k-1}f)(y-x) + \frac{1}{k!}f^{(k)}(\xi)(x-y)^k$ for some point ξ between x and y . This is natural in the sense that the polynomial $y \mapsto (T_x^{k-1}f)(y-x)$ can be described as the unique degree $k-1$ polynomial whose derivatives of orders $\leq k-1$ at x match those of f . Note also that, in the special case $k=1$, the statement is just

$$f(y) = f(x) + f'(\xi)(y-x) \quad \text{for some } \xi \text{ between } x \text{ and } y$$

which is exactly the statement of the Mean Value Theorem. We will see below that the proof of Taylor's theorem is just repeated application of Taylor's theorem.

The last term involving $f^{(k)}(\xi)$ is called the k -remainder term, often written as $(R^k f)(x, y)$, so we have

$$f(y) = \sum_{j=0}^{k-1} \frac{f^{(j)}(x)}{j!} (y-x)^j + (R^k f)(x, y).$$

This says that f is well-approximated by the Taylor polynomial $T_x^{k-1} f$, so long as the remainder term $R^k f(x, y)$ can be shown to be small. From the form of the remainder term in the theorem, this amounts to having good control on the k th derivative $f^{(k)}$ (at arbitrary points, since all we know is that ξ is between x and y). Indeed, we have

$$\frac{f(x+t) - (T_x^{k-1} f)(t)}{t^k} = \frac{1}{k!} f^{(k)}(\xi)$$

and so, as long as we have good uniform control on $f^{(k)}$, this gives us the alluded-to generalization of (8.1): if, for example, $|f^{(k)}(\xi)| \leq M$ for all ξ , then we'll have

$$f(x+t) = T_x^{k-1} f(t) + o(t^k).$$

PROOF OF THEOREM 8.24. For any t_0 such that $x+t_0 \in (a, b)$, define α_0 by the equation

$$f(x+t_0) = (T_x^{k-1} f)(t_0) + \frac{t_0^k}{k!} \alpha_0.$$

Our goal is to show that $\alpha_0 = f^{(k)}(\xi)$ for some ξ between x and $x+t_0$. Well, on the interval between 0 and t_0 , consider the function

$$g(t) = (T_x^{k-1} f)(t) + \frac{t^k}{k!} \alpha_0 - f(x+t), \quad \text{for } t \text{ between } 0 \text{ and } t_0.$$

As f and $T_x^{k-1} f$ are C^{k-1} and have k th derivatives, the same holds true of g . By definition of α_0 , $g(t_0) = 0$. Also, $(T_x^{k-1} f)(0) = f(x)$ by definition, so $g(0) = 0$. Therefore, by the mean value theorem there is a point t_1 between 0 and t_0 such that $g'(t_1) = \frac{g(t_0) - g(0)}{t_0} = 0$. But by construction of Taylor polynomials, and the fact that $\frac{d}{dt} t^k k t^{k-1} = 0$ at $t = 0$, we have $g'(0) = (T_x^{k-1} f)'(0) - f'(x) = 0$, and so we can apply the Mean Value Theorem to g' to find a point t_2 between 0 and t_1 where $g''(t_2) = 0$. We may continue this way finding $t_1, t_2, t_3, \dots, t_k$, with t_j between 0 and t_{j-1} , such that $g^{(j)}(t_j) = 0$. At the last step, since $T_x^{k-1} f$ is a polynomial of degree $k-1$, its k th derivative is 0, and also $\frac{d^k}{dt^k} \frac{1}{k!} t^k = 1$, so

$$0 = g^{(k)}(t_k) = (T_x^{k-1} f)^{(k-1)}(t_k) + \alpha_0 - f^{(k)}(x+t_k) = \alpha_0 - f^{(k)}(x+t_k).$$

Setting $\xi = x+t_k$ concludes the proof. □

Taylor's theorem is one of the most ubiquitously useful results in analysis, and we will use it frequently. As a first application, we now use it to understand another very powerful computational tool: *l'Hôpital's rule*. This deals with limits of the form

$$\lim_{y \rightarrow x} \frac{f(y)}{g(y)}$$

where $\lim_{y \rightarrow x} f(y) = \lim_{y \rightarrow x} g(y) = 0$. (There are some other cases as well, which we will discuss below.) To understand what happens to such a limit, let us suppose for the moment that f and g are C^2 in a neighborhood of x . These constraints are much stronger than required, as we'll see in the

proof below; but with these assumptions, we can use Taylor's theorem to understand what happens. Since f and g are continuous at x , we have $f(x) = \lim_{y \rightarrow x} f(y) = 0$ and $g(x) = \lim_{y \rightarrow x} g(y) = 0$. We then have, for sufficiently small t ,

$$\begin{aligned} f(x+t) &= f(x) + f'(x)t + \frac{1}{2}f''(\xi)t^2 = f'(x)t + \frac{1}{2}f''(\xi)t^2 \\ g(x+t) &= g(x) + g'(x)t + \frac{1}{2}g''(\eta)t^2 = g'(x)t + \frac{1}{2}g''(\eta)t^2 \end{aligned}$$

for some ξ and η between x and $x+t$. Thus

$$\frac{f(x+t)}{g(x+t)} = \frac{f'(x) + \frac{1}{2}f''(\xi)t}{g'(x) + \frac{1}{2}g''(\eta)t}.$$

Since we assumed $f, g \in C^2$, the functions f'' and g'' are continuous. As $t \rightarrow 0$, since ξ and η are between x and $x+t$, we see that $\xi \rightarrow x$ and $\eta \rightarrow x$, and so $f''(\xi) \rightarrow f''(x)$ and $g''(\xi) \rightarrow g''(x)$. By the limit theorems, then we see that if $g'(x) \neq 0$, then

$$\lim_{y \rightarrow x} \frac{f(y)}{g(y)} = \lim_{t \rightarrow 0} \frac{f(x+t)}{g(x+t)} = \frac{f'(x)}{g'(x)}.$$

This is the result of L'Hôpital's rule: the limit of a ratio of functions that is indeterminate is equal to the limit of ratio of derivatives (provided it is not indeterminate). If $f'(x) = g'(x) = 0$, we could then apply the same reasoning with higher derivatives (assuming the functions are C^3 , for example) to get $\frac{f''(x)}{g''(x)}$, and so forth.

4. Lecture 7: April 19, 2016

Having used Taylor's theorem to explore *why* L'Hôpital's rule makes sense (and prove it for nice enough functions), we now go about proving that it holds much more generally: we do not need the functions to be C^2 , or even C^1 , merely differentiable; and they need not be differentiable at the point x , only in a (one-sided) neighborhood of x .

THEOREM 8.25 (L'Hôpital's Rule). *Let $a < b$ in \mathbb{R} , and let $f, g: (a, b) \rightarrow \mathbb{R}$ be differentiable functions, with $g(x) \neq 0$ and $g'(x) \neq 0$ for all $x \in (a, b)$. Suppose that*

$$\lim_{x \rightarrow a^+} \frac{f'(x)}{g'(x)} = L$$

exists in \mathbb{R} . If $\lim_{x \rightarrow a^+} f(x) = \lim_{x \rightarrow a^+} g(x) = 0$, or if $\lim_{x \rightarrow a^+} |g(x)| = \infty$, then

$$\lim_{x \rightarrow a^+} \frac{f(x)}{g(x)} = L.$$

The analogous statement holds with $\lim_{x \rightarrow a^+}$ replaced with $\lim_{x \rightarrow b^-}$.

To prove L'Hôpital's rule in this general form, we first need an extended version of the Mean Value Theorem that deals with ratios of functions.

LEMMA 8.26 (Extended Mean Value Theorem). *Let $a < b$ in \mathbb{R} , and let $f, g: [a, b] \rightarrow \mathbb{R}$ be continuous functions, such that f and g are differentiable on (a, b) . Assume that $g(a) \neq g(b)$. Then there is a point $x \in (a, b)$ where $g'(x) \neq 0$, and*

$$\frac{f(b) - f(a)}{g(b) - g(a)} = \frac{f'(x)}{g'(x)}.$$

The requirement that $g(b) \neq g(a)$ is simply so that the ratio on the left makes sense; the lemma can be stated more generally without this assumption, and without the conclusion that $g'(x) \neq 0$, in the form $[f(b) - f(a)]g'(x) = [g(b) - g(a)]f'(x)$. It is tempting to try to prove the lemma by applying the Mean Value Theorem to f and g separately, but this will only show that the ratio on the left is equal to $f'(x)/g'(y)$ for two (not necessarily equal) points x, y . Instead, we just follow the precise outline of the proof of the Mean Value Theorem 8.18.

PROOF. Define $h(t) = [f(b) - f(a)]g(t) - [g(b) - g(a)]f(t)$. If h is constant, then its derivative is 0, and the result is that the conclusion holds at every x . So assume h is not constant, and wlog assume there is some point $t \in (a, b)$ where $h(t) > h(a)$. The function h is continuous on the compact interval $[a, b]$, so achieves its maximum at some point x in this closed interval; because there is some t with $h(t) > h(a)$, we know $x \neq a$. Also, a quick calculation shows that $h(b) = h(a) = f(b)g(a) - f(a)g(b)$, and so $x \neq b$ either. Since h is differentiable in (a, b) , and h achieves its maximum at $x \in (a, b)$, it follows that $h'(x) = 0$, which yields the desired conclusion. \square

We can now proceed to prove L'Hôpital's rule.

PROOF OF THEOREM 8.25. We first treat the case $\lim_{x \rightarrow a^+} f(x) = \lim_{x \rightarrow a^+} g(x) = 0$. Here, we may extend the functions to be defined on $[a, b]$ by defining $f(a) = g(a) = 0$, and this (by definition) makes them continuous on this interval. Let (x_n) be a sequence in (a, b) such that $x_n \rightarrow a$ as $n \rightarrow \infty$. Then f and g are continuous on $[a, x_n]$, and differentiable on (a, x_n) ,

with $g'(x) \neq 0$ for all x in this interval. By the extended Mean Value Theorem, there is a point $t_n \in (a, x_n)$ such that

$$\frac{f'(t_n)}{g'(t_n)} = \frac{f(x_n) - f(a)}{g(x_n) - g(a)} = \frac{f(x_n)}{g(x_n)}.$$

By the Squeeze theorem, $t_n \rightarrow a$, and so by assumption $\lim_{n \rightarrow \infty} f'(t_n)/g'(t_n) \rightarrow L$. Hence, the function $r(x) = f(x)/g(x)$ satisfies $r(x_n) \rightarrow L$ for every sequence (x_n) in (a, b) for which $x_n \rightarrow a$; by definition of right limit, the conclusion follows.

Now consider the case that $\lim_{x \rightarrow a^+} |g(x)| = \infty$. The idea here is related, but a little more complicated, and is more amenable to an ϵ - δ proof. By assumption $\lim_{x \rightarrow a^+} \frac{f(x)}{g(x)} = L$; thus, for any $\epsilon > 0$, there is some $\delta_0 > 0$ so that

$$\left| \frac{f'(t)}{g'(t)} - L \right| < \frac{\epsilon}{4} \quad \text{for all } t \in (a, a + \delta_0). \quad (8.2)$$

(We freely assume δ_0 is small enough that $a + \delta_0 < b$, of course.) Fix some $x_0 \in (a, a + \delta_0)$. Since $|g(x)| \rightarrow \infty$ as $x \rightarrow a$, there is some $\delta_1 > 0$ so that, for $x \in (a, a + \delta_1)$, $|g(x)| > |g(x_0)|$; in particular, $g(x) \neq g(x_0)$ for $x \in (a, a + \delta_1)$. Now, both f and g are continuous on $[a, x_0]$, and differentiable on (a, x_0) . It follows from the Extended Mean Value Theorem that for each $x \in (a, a + \delta_1)$, there is a point $t_0 = t(x, x_0)$ in (x, x_0) (and therefore in $(a, a + \delta_0)$) such that

$$\frac{f(x) - f(x_0)}{g(x) - g(x_0)} = \frac{f'(t_0)}{g'(t_0)}.$$

We are interested in $\frac{f(x)}{g(x)}$, not the more complicated quotient above; but since $|g(x)|$ grows without bound, the two are very close. Dividing through top and bottom by $g(x)$ gives

$$\frac{f'(t_0)}{g'(t_0)} = \frac{\frac{f(x)}{g(x)} - \frac{f(x_0)}{g(x_0)}}{1 - \frac{g(x_0)}{g(x)}}; \quad \text{therefore} \quad \frac{f(x)}{g(x)} = \left(1 - \frac{g(x_0)}{g(x)}\right) \frac{f'(t_0)}{g'(t_0)} + \frac{f(x_0)}{g(x)}. \quad (8.3)$$

We now have all the pieces in place. Before proceeding formally, let's outline how this works. From (8.2), $\frac{f'(t_0)}{g'(t_0)}$ is close to L . Since $|g(x)| \rightarrow \infty$, for x close to a , $1 - \frac{g(x_0)}{g(x)}$ is close to 1, and $\frac{f(x_0)}{g(x)}$ is close to 0. This shows that for such x , $\frac{f(x)}{g(x)}$ is close to L , as desired.

Now let's make this precise. Working from (8.3), if $a < x < a + \min\{\delta_0, \delta_1\}$, then $t_0 \in (a, a + \delta_0)$, and we have

$$\begin{aligned} \left| \frac{f(x)}{g(x)} - L \right| &\leq \left| \left(1 - \frac{g(x_0)}{g(x)}\right) \frac{f'(t_0)}{g'(t_0)} - L \right| + \left| \frac{f(x_0)}{g(x)} \right| \\ &= \left| \left(1 - \frac{g(x_0)}{g(x)}\right) \left(\frac{f'(t_0)}{g'(t_0)} - L\right) - L \frac{g(x_0)}{g(x)} \right| + \left| \frac{f(x_0)}{g(x)} \right| \\ &\leq \left|1 - \frac{g(x_0)}{g(x)}\right| \left| \frac{f'(t_0)}{g'(t_0)} - L \right| + |L| \left| \frac{g(x_0)}{g(x)} \right| + \left| \frac{f(x_0)}{g(x)} \right| \\ &< \left(1 + \left| \frac{g(x_0)}{g(x)} \right| \right) \frac{\epsilon}{4} + |L| \left| \frac{g(x_0)}{g(x)} \right| + \left| \frac{f(x_0)}{g(x)} \right|, \end{aligned} \quad (8.4)$$

where we used (8.2) in (8.4). As $|g(x)| \rightarrow \infty$ as $x \rightarrow a+$, and since $|f(x_0)|$ and $|g(x_0)|$ are fixed finite numbers, we can find some $\delta_2 > 0$ so that, for all $x \in (a, a + \delta_2)$,

$$\frac{|g(x_0)|}{|g(x)|} < 1 \quad \text{and} \quad |L| \frac{|g(x_0)|}{|g(x)|} + \frac{|f(x_0)|}{|g(x)|} < \frac{\epsilon}{2}.$$

The first bound shows that the first term in (8.4) is $< \frac{\epsilon}{2}$, and the second bound shows that the last terms in (8.4) are $< \frac{\epsilon}{2}$. Thus

$$\left| \frac{f(x)}{g(x)} - L \right| < \frac{\epsilon}{2} + \frac{\epsilon}{2} = \epsilon \quad \text{for } x \in (a, a + \min\{\delta_0, \delta_1, \delta_2\}).$$

As $\epsilon > 0$ was chosen arbitrarily, this shows that $\lim_{x \rightarrow a+} \frac{f(x)}{g(x)} = L$, as desired. \square

REMARK 8.27. The same holds true with $a = -\infty$ or $b = +\infty$, and the proof is extremely similar; the details are left to the bored reader.

Let us now conclude our discussion of differentiation by considering how much of what we've developed applies to vector-valued functions of a real variable. First, the definitions are essentially the same.

DEFINITION 8.28. Let $a < b$ in \mathbb{R} , let $d \in \mathbb{N}$, and let $\mathbf{f}: (a, b) \rightarrow \mathbb{R}^d$ be a functions. Call \mathbf{f} differentiable at $t_0 \in [a, b]$ if

$$\mathbf{f}'(t_0) = \lim_{t \rightarrow t_0} \frac{\mathbf{f}(t) - \mathbf{f}(t_0)}{t - t_0}$$

exists. Say that \mathbf{f} is differentiable if $\mathbf{f}'(t_0)$ exists for each t_0 in the domain of \mathbf{f} ; in this case, \mathbf{f}' is a function $[a, b] \rightarrow \mathbb{R}^d$ as well. If this function is continuous, we say $\mathbf{f} \in C^1$. More generally, for $k \in \mathbb{N}$, we say $\mathbf{f} \in C^k$ if \mathbf{f} has k continuous derivatives. If $\mathbf{f} \in C^k$ for all k , we say $\mathbf{f} \in C^\infty$, and call \mathbf{f} smooth.

It is often customary to use the variable name t for vector-valued functions, as we often think of \mathbf{f} as tracing out a curve in space, with t tracking the flow of time. It is then very important to note that smoothness properties of \mathbf{f} are different from smoothness properties of the range $\mathbf{f}[a, b]$: it is perfectly possible to construct a smooth function $\mathbf{f}: [0, 1] \rightarrow \mathbb{R}^2$ such that the image $\mathbf{f}[0, 1]$ is a curve with a right angle in it. We may discuss this later, time permitting.

PROPOSITION 8.29. A function $\mathbf{f} = (f_1, f_2, \dots, f_d)$ is differentiable at t_0 if and only if each of its component functions f_1, f_2, \dots, f_d is differentiable at t_0 , in which case

$$\mathbf{f}'(t_0) = (f_1'(t_0), f_2'(t_0), \dots, f_d'(t_0)).$$

More generally, a function \mathbf{f} is C^k iff its component functions are C^k .

PROOF. We just write out that

$$\mathbf{f}(t) - \mathbf{f}(t_0) = (f_1(t) - f_1(t_0), f_2(t) - f_2(t_0), \dots, f_d(t) - f_d(t_0))$$

and so, dividing by $t - t_0$, we see that the difference quotient is a vector whose j th component is $\frac{f_j(t) - f_j(t_0)}{t - t_0}$. Since a limit of a vector sequence exists iff the limits of all its components exist, in which case the limit is the vector of component limits (as proved, e.g., in the case $d = 2$ in Proposition 3.13), the first result follows. The same applies to C^k functions by noting that continuity obeys the same structure (a vector-valued function is continuous iff its component functions are all continuous), and then iterating. \square

EXAMPLE 8.30. As a special case, consider functions $\mathbf{f}: \mathbb{R} \rightarrow \mathbb{R}^2$, where we identify $\mathbb{R}^2 \cong \mathbb{C}$; then we may write

$$\mathbf{f}(t) = f_1(t) + if_2(t).$$

For example, taking it on faith that \cos and \sin are differentiable functions satisfying $\sin' = \cos$ and $\cos' = -\sin$, consider the function

$$\mathbf{e}(t) = \cos t + i \sin t.$$

Both component functions are smooth, and therefore \mathbf{e} is a smooth function. Indeed, this function traces out the unit circle at unit angular speed in the counter-clockwise direction. Note that

$$\mathbf{e}'(t) = -\sin t + i \cos t = i(i \sin t + \cos t) = i\mathbf{e}(t).$$

Thinking back to your knowledge of differential equations, if u is a smooth function satisfying $u'(t) = au(t)$ for some constant a , then $u(t) = e^{at}u(0)$. In our case $\mathbf{e}(0) = 1 + 0i = 1$, so we should expect that $\mathbf{e}(t) = e^{it}$. This is, in fact, true, and we will prove it more satisfactorily later this quarter.

While the definitions and notation of derivatives apply equally well to vector-valued functions, the major theorems do *not* apply. For example, consider the Mean Value Theorem. One might think that one could show that $\frac{\mathbf{f}(b) - \mathbf{f}(a)}{b - a} = \mathbf{f}'(\xi)$ for some $\xi \in (a, b)$, the usual way one might approach this (componentwise) fails. The Mean Value Theorem applied to each component f_j asserts that there is some point $\xi_j \in (a, b)$ where $\frac{f_j(b) - f_j(a)}{b - a} = f_j'(\xi_j)$; but the point ξ_j can certainly depend on j , and there is no obvious way to guarantee that a single point ξ will work for all j . In fact, this is just not true.

EXAMPLE 8.31. Continuing with Example 8.30, notice that $\mathbf{e}(2\pi) - \mathbf{e}(0) = 1 - 1 = 0$. However, $\mathbf{e}'(t) = i\mathbf{e}(t)$, and so $\|\mathbf{e}'(t)\|^2 = \sin^2 t + \cos^2 t = 1$ for any t . Thus, there is no point ξ where $\mathbf{e}'(\xi) = 0$, and so

$$\frac{\mathbf{e}(2\pi) - \mathbf{e}(0)}{2\pi - 0} \neq \mathbf{e}'(\xi) \quad \text{for any } \xi \in (0, 2\pi).$$

The Mean Value Theorem also featured prominently in the proof of L'Hôpital's Rule. In general for vector-valued functions \mathbf{f} and \mathbf{g} , it does not even make *sense* to take the ratio $\frac{\mathbf{f}}{\mathbf{g}}$ (one cannot divide by a vector). In the case of \mathbb{C} -valued functions, we can divide (in the same sense that we can for real-valued functions: so long as the denominator itself does not vanish at points where we are dividing). Even so, L'Hôpital's Rule fails here.

EXAMPLE 8.32. Consider the two complex-valued functions $f(t) = t$ and $g(t) = te^{i/t}$. That is, f is nominally complex-valued, $f(t) = t + i0$, while $g(t) = t(\cos \frac{1}{t} + i \sin \frac{1}{t})$. Of course $\lim_{t \rightarrow 0} f(t) = 0$. In Example 8.16, we showed that $t \sin \frac{1}{t} \rightarrow 0$ as $t \rightarrow 0$; a very similar proof shows that $t \cos \frac{1}{t} \rightarrow 0$ as $t \rightarrow 0$. Thus $\lim_{t \rightarrow 0} g(t) = 0$ as well. So $\frac{f(t)}{g(t)}$ is a $\frac{0}{0}$ type indeterminate form as $t \rightarrow 0$, where both top and bottom are differentiable away from 0. If L'Hôpital's Rule held precisely as in Theorem 8.25, we would expect to see $\lim_{t \rightarrow 0} \frac{f(t)}{g(t)} = \lim_{t \rightarrow 0} \frac{f'(t)}{g'(t)}$. We can compute these derivatives (for $t \neq 0$): $f'(t) = 1$, while

$$g'(t) = \frac{d}{dt} \left(t \cos \frac{1}{t} \right) + i \frac{d}{dt} \left(t \sin \frac{1}{t} \right) = \left(\cos \frac{1}{t} - \frac{1}{t} \sin \frac{1}{t} \right) + i \left(\sin \frac{1}{t} + \frac{1}{t} \cos \frac{1}{t} \right).$$

Notice that

$$|g'(t)|^2 = \left(\cos \frac{1}{t} + \frac{1}{t} \sin \frac{1}{t}\right)^2 + \left(\sin \frac{1}{t} - \frac{1}{t} \cos \frac{1}{t}\right)^2 = \left(\cos^2 \frac{1}{t} + \sin^2 \frac{1}{t}\right) \left(1 + \frac{1}{t^2}\right) = 1 + \frac{1}{t^2} > 0$$

and hence $g'(t)$ is never equal to 0. The reciprocal can be computed: since $\frac{1}{z} = \frac{\bar{z}}{|z|^2}$,

$$\frac{f'(t)}{g'(t)} = \frac{1}{g'(t)} = \frac{1}{1 + 1/t^2} \left[\left(\cos \frac{1}{t} + \frac{1}{t} \sin \frac{1}{t}\right) - i \left(\sin \frac{1}{t} - \frac{1}{t} \cos \frac{1}{t}\right) \right].$$

Let's look at the real part for now. Simplifying, we have

$$\frac{t^2}{t^2 + 1} \left(\cos \frac{1}{t} + \frac{1}{t} \sin \frac{1}{t}\right) = \frac{1}{t^2 + 1} \left(t^2 \cos \frac{1}{t} + t \sin \frac{1}{t}\right).$$

At $t \rightarrow 0$, $\frac{1}{t^2+1} \rightarrow 1$, and both terms inside the brackets tend to 0 (cf. Example 8.17). Similar arguments show that the second term in $\frac{f'(t)}{g'(t)}$ tends to 0; so $\lim_{t \rightarrow 0} \frac{f'(t)}{g'(t)} = 0$.

However, we can compute directly (using computations just like the ones above, or by following our nose with the exponential notation) that

$$\frac{f(t)}{g(t)} = \frac{t}{te^{i/t}} = e^{-i/t} = \cos \frac{1}{t} - i \sin \frac{1}{t}$$

and as we know (cf. Example 7.18), this limit does not exist. So the statement of L'Hôpital's rule is simply false for complex-valued functions.

We have claimed that the Mean Value Theorem is the key tool of calculus; and L'Hôpital's rule is a very powerful computational tool. Are we then lost when it comes to calculus of vector-valued functions of a real variable? No. The saving grace is that, while the Mean Value Theorem does not hold, it does hold as an *inequality* in general, and that is enough for most theoretical applications.

THEOREM 8.33 (Mean Value Inequality for Vector-Valued Functions). *Let $a < b$ in \mathbb{R} , $d \in \mathbb{N}$, and let $\mathbf{f}: [a, b] \rightarrow \mathbb{R}^d$ be continuous, and differentiable on (a, b) . Then there exists a point $\xi \in (a, b)$ where*

$$\|\mathbf{f}(b) - \mathbf{f}(a)\| \leq (b - a) \|\mathbf{f}'(\xi)\|.$$

Note that this result does not contradict Example 8.31: in that example, the left-hand-side of the inequality is 0, which is certainly \leq the non-negative right-hand-side.

To prove Theorem 8.33, we need one key inequality which is one of the most important inequalities in all of analysis.

LEMMA 8.34 (The Cauchy-Schwarz Inequality). *Let $d \in \mathbb{N}$, and let $\mathbf{v}, \mathbf{w} \in \mathbb{R}^d$. Then $|\mathbf{v} \cdot \mathbf{w}| \leq \|\mathbf{v}\| \|\mathbf{w}\|$.*

PROOF. In this finite-dimensional setting, it is possible to prove the inequality by squaring both sides and comparing terms directly. But it is instructive to give a cleaner proof even here. Consider the function $p(t) = \|\mathbf{v} - t\mathbf{w}\|^2$. This function is always ≥ 0 . Expanding it out, we have

$$\begin{aligned} p(t) &= (\mathbf{v} - t\mathbf{w}) \cdot (\mathbf{v} - t\mathbf{w}) = \mathbf{v} \cdot \mathbf{v} - 2t\mathbf{v} \cdot \mathbf{w} + t^2\mathbf{w} \cdot \mathbf{w} \\ &= \|\mathbf{v}\|^2 - 2t\mathbf{v} \cdot \mathbf{w} + t^2\|\mathbf{w}\|^2 \\ &= at^2 + bt + c \end{aligned}$$

where $a = \|\mathbf{w}\|^2$, $b = -2\mathbf{v} \cdot \mathbf{w}$, and $c = \|\mathbf{v}\|^2$. Thus p is a quadratic polynomial, with leading coefficient $a > 0$. Since we know $p(t) \geq 0$ for all t , it follows that p cannot have two distinct real roots (if so, the function would be strictly negative between those roots). Thus, either p has no real roots, or a double-root. From the quadratic formula, it thus follows that the discriminant $b^2 - 4ac \leq 0$. That is:

$$0 \geq b^2 - 4ac = (-2\mathbf{v} \cdot \mathbf{w})^2 - 4\|\mathbf{w}\|^2\|\mathbf{v}\|^2.$$

This shows that $(\mathbf{v} \cdot \mathbf{w})^2 \leq \|\mathbf{v}\|^2\|\mathbf{w}\|^2$; taking square roots yields the result. \square

REMARK 8.35. This proof relied on nothing more than the fact that the norm $N(\mathbf{v}) = \|\mathbf{v}\|$ “polarizes” in terms of the bilinear form $B(\mathbf{v}, \mathbf{w}) = \mathbf{v} \cdot \mathbf{w}$: $N(\mathbf{v})^2 = B(\mathbf{v}, \mathbf{v})$. Indeed, the theorem holds in this level of generality: if B is a symmetric bilinear form on any vector space for which $B(\mathbf{v}, \mathbf{v}) \geq 0$ for all \mathbf{v} , then

$$|B(\mathbf{v}, \mathbf{w})| \leq B(\mathbf{v}, \mathbf{v})^{1/2} B(\mathbf{w}, \mathbf{w})^{1/2}.$$

This generalization is especially handy in many infinite-dimensional settings.

PROOF OF THEOREM 8.33. Set $\mathbf{v} = \mathbf{f}(b) - \mathbf{f}(a) \in \mathbb{R}^d$. Consider the real-valued function

$$\varphi(t) = \mathbf{v} \cdot \mathbf{f}(t) = \sum_{j=1}^d v_j f_j(t), \quad t \in [a, b].$$

As each component $v_j f_j$ is a continuous function on $[a, b]$, differentiable on (a, b) , the same applies to the sum, and so φ is the kind of function to which the Mean Value Theorem applies. Thus, there is a point $\xi \in (a, b)$ where $\varphi'(\xi) = \frac{\varphi(b) - \varphi(a)}{b - a}$. That is to say

$$(b - a)\varphi'(\xi) = \varphi(b) - \varphi(a) = \mathbf{v} \cdot \mathbf{f}(b) - \mathbf{v} \cdot \mathbf{f}(a) = \mathbf{v} \cdot [\mathbf{f}(b) - \mathbf{f}(a)] = \mathbf{v} \cdot \mathbf{v} = \|\mathbf{v}\|^2.$$

Also, note that

$$\varphi'(\xi) = \left. \frac{d}{dt} \sum_{j=1}^d v_j f_j(t) \right|_{t=\xi} = \sum_{j=1}^d v_j f_j'(\xi) = \mathbf{v} \cdot \mathbf{f}'(\xi).$$

Thus we have

$$(b - a)\mathbf{v} \cdot \mathbf{f}'(\xi) = \|\mathbf{v}\|^2.$$

Taking absolute values and applying the Cauchy-Schwarz inequality,

$$\|\mathbf{v}\|^2 = \|\|\mathbf{v}\|^2\| = |(b - a)\mathbf{v} \cdot \mathbf{f}'(\xi)| \leq (b - a)\|\mathbf{v}\|\|\mathbf{f}'(\xi)\|.$$

If $\mathbf{v} = \mathbf{f}(b) - \mathbf{f}(a) = \mathbf{0}$, there is nothing to prove; otherwise we can divide out one factor of $\|\mathbf{v}\|$ to yield

$$\|\mathbf{f}(b) - \mathbf{f}(a)\| = \|\mathbf{v}\| \leq (b - a)\|\mathbf{f}'(\xi)\|$$

which is the desired conclusion. \square

COROLLARY 8.36. *If $\mathbf{f} \in C^1(a, b)$, then*

$$\|\mathbf{f}(s) - \mathbf{f}(t)\| \leq \sup_{s \leq \xi \leq t} \|\mathbf{f}'(\xi)\| |s - t|, \quad \text{for all } s, t \in (a, b).$$

This allows for a vector-valued version of Lipschitz functions (cf. Homework 3): we define the Lipschitz norm of such a function to be

$$\|\mathbf{f}\|_{\text{Lip}} = \sup_{s \neq t} \frac{\|\mathbf{f}(s) - \mathbf{f}(t)\|}{|s - t|}$$

and we call a vector-valued function \mathbf{f} *Lipschitz* if $\|\mathbf{f}\|_{\text{Lip}} < \infty$. Corollary 8.36 shows that if $\mathbf{f} \in C^1$ with a bounded derivative then \mathbf{f} is Lipschitz, and $\|\mathbf{f}\|_{\text{Lip}} \leq \sup_{\xi} \|\mathbf{f}'(\xi)\|$; in fact, these are equal for the same reason they are in the scalar-valued case covered on the homework exercise (since $\|\mathbf{f}'(t)\|$ is very close to $\frac{\|\mathbf{f}(s) - \mathbf{f}(t)\|}{|s - t|}$ when s is close to t). One can then follow with a comparable theory of Hölder continuous vector-valued functions as well; this is important, but for now it is left to the reader's imagination.

One final remark: since Taylor's theorem was proved simply by iterating the Mean Value Theorem, in the vector-valued setting, one cannot expect a Taylor expansion with a remainder term analogous to the one in Theorem 8.24. One can, instead, formulate a *Taylor inequality*, where

$$\|\mathbf{f}(x + t) - T_x^{k-1}\mathbf{f}(x)\| \leq \frac{1}{k!} \|\mathbf{f}^{(k)}(\xi)\| |t|^k, \quad \text{for some point } \xi \text{ between } x \text{ and } x + t.$$

This is just as useful for approximations; but in truth, most applications of Taylor's theorem in the vector-valued setting are just as amenable to applying the Theorem to each component of \mathbf{f} separately and working from there.

CHAPTER 9

Integration

1. Lecture 8: April 21, 2016

You likely saw some version of the actual definition of the integral in your calculus class: it was presented as a “limit of Riemann sums”. This is not quite accurate: it is not a limit in the sense that we’ve defined in this course. (It is a much more complicated kind of limit.) There are several approaches to making sense of this rigorously. We are going to present the largely historically-accurate version here developed by Riemann somewhat non-rigorously, and really properly developed by Darboux (indeed, in some sources it is called the Darboux integral).

DEFINITION 9.1. Let $a < b$ in \mathbb{R} . A **partition** Π of $[a, b]$ is a finite set of points $\{t_0, t_1, \dots, t_n\}$ where

$$a = t_0 < t_1 < t_2 < \dots < t_{n-1} < t_n = b.$$

For $1 \leq j \leq n$, denote $\Delta t_j = x_j - x_{j-1}$.

A word on notation: we will usually have a fixed partition around, and the letter n will be used consistently to mean the index of the largest partition point (unless otherwise stated).

DEFINITION 9.2. Let $a < b$ in \mathbb{R} . Given a bounded function $f: [a, b] \rightarrow \mathbb{R}$, and a partition $\Pi = \{t_0, t_1, \dots, t_n\}$ of $[a, b]$, we can define **upper and lower sums** of f on Π :

$$U(f, \Pi) \equiv \sum_{j=1}^n \sup_{t_{j-1} \leq t \leq t_j} f(t) \cdot \Delta t_j$$
$$L(f, \Pi) \equiv \sum_{j=1}^n \inf_{t_{j-1} \leq t \leq t_j} f(t) \cdot \Delta t_j.$$

Note that this only makes sense for bounded f ; if f is not bounded, then on at least one partition interval $[t_{j-1}, t_j]$ f is unbounded, and so at least one of the terms in either $U(f, \Pi)$ or $L(f, \Pi)$ will be $\pm\infty$; but there could be more than one term that is $\pm\infty$, perhaps with opposite signs, and so the sums may not even be defined. Thus, we restrict our attention to bounded functions for now; we will later discuss how to extend integration to *some* unbounded functions.

DEFINITION 9.3. Let $a < b$ in \mathbb{R} , and let $f: [a, b] \rightarrow \mathbb{R}$. Define the **upper and lower Darboux integrals** of f as follows:

$$U(f) = \inf_{\Pi} U(f, \Pi), \quad L(f) = \sup_{\Pi} L(f, \Pi).$$

The inf and sup are taken over all partitions of $[a, b]$. If $U(f) = L(f)$, we say that f is **Riemann integrable** (or **Darboux integrable**), and denote this common value by

$$\int_a^b f = \int_a^b f(t) dt \equiv U(f) = L(f).$$

REMARK 9.4. The textbook uses the notation

$$U(f) = \overline{\int_a^b} f(t) dt \quad \text{and} \quad L(f) = \underline{\int_a^b} f(t) dt.$$

We will probably never use this notation.

Before continuing, let's discuss where these definitions come from. $U(f, \Pi)$ and $L(f, \Pi)$ are approximations of the "area under the curve" of the graph of f : you partition the domain $[a, b]$ into a finite collection of adjacent intervals, and on each interval you approximate the function f by a constant one. Which constant should you choose? The upper sum has you choose the largest (or more precisely supremal) value on that interval, while the lower sum chooses the smallest (or more precisely infimal) value. There are other possibilities: the original scheme by Riemann is to choose another set of points $t_1^*, t_2^*, \dots, t_n^*$ with $t_{j-1} \leq t_j^* \leq t_j$, and then approximate f by the constant $f(t_j^*)$ on the interval $[t_{j-1}, t_j]$. If the points $\{t_j^*\}$ are denoted together as Π^* , we might denote this sum as

$$R(f, \Pi, \Pi^*) = \sum_{j=1}^n f(t_j^*) \Delta t_j$$

which is the kind of "Riemann sum" you saw in calculus. In all cases, we then define an approximate integral to be the (signed) area of the rectangles with base length Δt_j and height given by the appropriate constant approximation of the function. Notice that, by definition of sup and inf on each interval,

$$L(f, \Pi) \leq R(f, \Pi, \Pi^*) \leq U(f, \Pi)$$

and so if $U(f) = L(f)$, the Riemann sums must also "converge" to this common value.

How do we compute the area under the actual curve, assuming that even makes sense? In the Riemann scheme, we then start changing the partition Π (and the associated points Π^*) to make the maximum width of any Δt_j smaller; this is supposed to give a better approximation, and then we take the "limit" as this width goes to 0. In the Darboux approach (which we are following), notice that $U(f, P)$ is definitely an over-estimate for the "actual" area: if you change the partition Π (for example by refining it to add a new point splitting one of the intervals into two), you can only decrease the value of the sum, since the sup of f on the two sub-intervals cannot be larger than the sup on the whole interval: for $t_{j-1} \leq s \leq t_j$,

$$\sup_{t_{j-1} \leq t \leq t_j} f(t) \cdot (t_j - t_{j-1}) \geq \sup_{t_{j-1} \leq t \leq s} f(t) \cdot (s - t_{j-1}) + \sup_{s \leq t \leq t_j} f(t) \cdot (t_j - s). \quad (9.1)$$

(You should draw a quick picture to see why this is true.) Hence, replacing Π with a refined partition Π' will yield $U(f, \Pi') \leq U(f, \Pi)$. A similar argument shows that $L(f, \Pi') \geq L(f, \Pi)$. This is why we define the upper Darboux integral to be the *infimum* of the upper sums over all partitions, and likewise the lower Darboux integral is the *supremum* of the lower sums over all partitions.

Of course, we need to make sure these infima and suprema make sense. Let's summarize this with the following lemma.

LEMMA 9.5. *Let $a < b$ in \mathbb{R} , and let $f: [a, b] \rightarrow \mathbb{R}$ be a bounded function, with $|f(t)| \leq M$ for all $t \in [a, b]$. Then for any partitions Π_1 and Π_2 of $[a, b]$,*

$$-M(b-a) \leq L(f, \Pi_1) \leq U(f, \Pi_2) \leq M(b-a).$$

Ergo

$$-M(b-a) \leq L(f) \leq U(f) \leq M(b-a).$$

PROOF. Let $\Pi_1 = \{s_0 < s_1 < \cdots < s_m\}$ and $\Pi_2 = \{t_0 < t_1 < \cdots < t_n\}$. Since $f(t) \leq M$ for all t , M is an upper bound for $\{f(t) : t_{j-1} \leq t \leq t_j\}$ for each j , and hence $\sup_{t_{j-1} \leq t \leq t_j} f(t) \leq M$. Thus

$$U(f, \Pi_2) = \sum_{j=1}^n \sup_{t_{j-1} \leq t \leq t_j} f(t) \cdot \Delta t_j \leq \sum_{j=1}^n M \Delta t_j = M \sum_{j=1}^n (t_j - t_{j-1}) = M(b - a).$$

(The last sum is telescoping.) A similar proof, using the fact that $f(t) \geq -M$ for all $t \in [a, b]$, shows that $L(f, \Pi_1) \geq -M(b - a)$. For the middle inequality, we use an important trick: we introduce a new partition Π which is the *common refinement* of Π_1 and Π_2 : $\Pi = \Pi_1 \cup \Pi_2 = \{s_0, s_1, \dots, s_m, t_0, t_1, \dots, t_n\}$ (not written in order here). For this common partition (whose points we'll refer to as u_j), we note that, for each j ,

$$\inf_{u_{j-1} \leq t \leq u_j} f(t) \leq \sup_{u_{j-1} \leq t \leq u_j} f(t).$$

Multiplying both sides by the positive number Δu_j and summing up yields $L(f, \Pi) \leq U(f, \Pi)$. Then, using (9.1), we see (by induction) that

$$L(f, \Pi_1) \leq L(f, \Pi) \leq U(f, \Pi) \leq U(f, \Pi_2).$$

This concludes the proof of the first chain of inequalities. For the second: since $U(f, \Pi_2) \leq M(b - a)$ for all Π_2 , it follows that $U(f) \leq M(b - a)$; a similar argument shows that $L(f, \Pi_1) \geq -M(b - a)$. For the middle inequality, we hold Π_2 fixed: since $L(f, \Pi_1) \leq U(f, \Pi_2)$ for all Π_1 , it follows that $L(f) = \sup_{\Pi_1} L(f, \Pi_1) \leq U(f, \Pi_2)$. Thus $L(f)$ is a lower bound for $\{U(f, \Pi_2) : \Pi_2\}$, and so $U(f) = \inf_{\Pi_2} U(f, \Pi_2) \geq L(f)$, as desired. \square

Thus, $U(f)$ and $L(f)$ are well-defined for any bounded function f , and ordered $L(f) \leq U(f)$. The question is whether they're equal. The answer is: certainly not always.

EXAMPLE 9.6. Consider Dirichlet's function from Example 6.11: the indicator function of the rationals.

$$f(x) = \begin{cases} 0, & x \notin \mathbb{Q} \\ 1, & x \in \mathbb{Q} \end{cases}.$$

Taking this function for $x \in [0, 1]$ let us compute the upper and lower integrals. First, fix any partition $\Pi = \{0 = t_0 < t_1 < \cdots < t_n = 1\}$ of $[0, 1]$. On any interval $[t_{j-1}, t_j]$, since $t_{j-1} < t_j$, we know there are both rational and irrational points in the interval. It follows that

$$\sup_{t_{j-1} \leq t \leq t_j} f(t) = 1, \quad \inf_{t_{j-1} \leq t \leq t_j} f(t) = 0.$$

Thus

$$U(f, \Pi) = \sum_{j=1}^n 1 \cdot \Delta t_j = (1 - 0) = 1, \quad L(f, \Pi) = \sum_{j=1}^n 0 \cdot \Delta t_j = 0.$$

So $U(f, \Pi)$ and $L(f, \Pi)$ do not depend on Π , and therefore taking appropriate sup and inf, we see that $U(f) = 1$ while $L(f) = 0$. Ergo, f is not Riemann integrable: $U(f) \neq L(f)$.

The question of when it is true that $U(f) = L(f)$ is a delicate one having to do with the continuity properties of the function f . We will explore and answer this question fully. Before we do, it pays to be a little more general right away (without adding any abstraction), and talk about the *Riemann-Stieltjes* integral. That is our next topic.

2. Lecture 9: April 26, 2016

The generalization of the Riemann / Darboux integral we will now develop allows for a “weighting” of the domain. In the construction of the integral, when approximating the area under the graph of a function by rectangles, we may compute the “area” of each partition rectangle over $[t_{j-1}, t_j]$ by declaring the length of this interval is not necessarily Δt_j but instead some (well-behaved) function of this width. In fact, a wide variety of weight functions are possible to produce meaningful theories of integration, but to mimic precisely the upper and lower integral construction outlined in the previous lecture, we restrict ourselves to monotone increasing weight functions.

DEFINITION 9.7. *Let $a < b$ in \mathbb{R} , and let $\alpha: [a, b] \rightarrow \mathbb{R}$ be a monotone increasing function. (In particular, since $\alpha([a, b]) = [\alpha(a), \alpha(b)]$, α is also bounded.) Given a partition $\Pi = \{a = t_0 < t_1 < \dots < t_n = b\}$, define $\Delta\alpha_j = \alpha(t_j) - \alpha(t_{j-1})$ (which is ≥ 0 since α is monotone increasing).*

*Let $f: [a, b] \rightarrow \mathbb{R}$ be a bounded function. Define the **upper sum** and **lower sum** of f relative to Π and α as*

$$U(f, \Pi, \alpha) \equiv \sum_{j=1}^n \sup_{t_{j-1} \leq t \leq t_j} f(t) \cdot \Delta\alpha_j$$

$$L(f, \Pi, \alpha) \equiv \sum_{j=1}^n \inf_{t_{j-1} \leq t \leq t_j} f(t) \cdot \Delta\alpha_j$$

An argument exactly like the one in Lemma 9.5 shows that, if $|f(t)| \leq M$ for all t then for any partitions Π_1 and Π_2 , with $\Pi^* = \Pi_1 \cup \Pi_2$, we have

$$-M(\alpha(b) - \alpha(a)) \leq L(f, \Pi_1, \alpha) \leq L(f, \Pi^*, \alpha) \leq U(f, \Pi^*, \alpha) \leq U(f, \Pi_2, \alpha) \leq M(\alpha(b) - \alpha(a)). \quad (9.2)$$

In particular, the definition / proposition makes sense.

PROPOSITION 9.8. *Let $a < b$ in \mathbb{R} , let $\alpha: [a, b] \rightarrow \mathbb{R}$ be monotone increasing, and let $f: [a, b] \rightarrow \mathbb{R}$ be bounded. Define*

$$U(f, \alpha) \equiv \inf_{\Pi} U(f, \Pi, \alpha), \quad \text{and} \quad L(f, \alpha) = \sup_{\Pi} L(f, \Pi, \alpha).$$

Then

$$-M(\alpha(b) - \alpha(a)) \leq L(f, \alpha) \leq U(f, \alpha) \leq M(\alpha(b) - \alpha(a)).$$

*If $L(f, \alpha) = U(f, \alpha)$, we say that f is **Riemann-Stieltjes integrable** with respect to α , and write $f \in \mathcal{R}(\alpha)$. In this case, we denote the common value by*

$$\int_a^b f d\alpha = \int_a^b f(t) d\alpha(t) \equiv U(f, \alpha) = L(f, \alpha).$$

REMARK 9.9. We will tend not to use the second notation $\int_a^b f(t) d\alpha(t)$, since the t is just a dummy variable and carries no independent meaning here.

Taking $\alpha(x) = x$, this reduces to the Riemann integral theory discussed in the previous lecture. In this case, we write simply $f \in \mathcal{R}$. We might write the integral in this case as $\int_a^b f(t) dt$, but more likely just as $\int_a^b f$.

We now give a quantitative characterization of Riemann-Stieltjes integrability.

LEMMA 9.10. *Let $a < b$ in \mathbb{R} , let $f: [a, b] \rightarrow \mathbb{R}$ be bounded, and let $\alpha: [a, b] \rightarrow \mathbb{R}$ be monotone increasing. Then $f \in \mathcal{R}(\alpha)$ if and only if for each $\epsilon > 0$ there is a partition Π of $[a, b]$ such that*

$$U(f, \Pi, \alpha) - L(f, \Pi, \alpha) < \epsilon. \quad (9.3)$$

PROOF. First, suppose that (9.3) holds; so fix $\epsilon > 0$ and let Π be a partition verifying that inequality. By definition $L(f, \Pi, \alpha) \leq L(f, \alpha) \leq U(f, \alpha) \leq U(f, \Pi, \alpha)$, and so (9.3) implies that $0 \leq U(f, \alpha) - L(f, \alpha) < \epsilon$ for each $\epsilon > 0$; thus $U(f) = L(f)$, as desired.

Conversely, suppose $f \in \mathcal{R}(\alpha)$, and let $\epsilon > 0$. Since $\int f d\alpha = \sup_{\Pi} L(f, \Pi, \alpha)$, there is some partition Π_1 so that $L(f, \Pi_1, \alpha) > \int f d\alpha - \frac{\epsilon}{2}$, and there is some partition Π_2 so that $U(f, \Pi_2, \alpha) < \int f d\alpha + \frac{\epsilon}{2}$. Let $\Pi^* = \Pi_1 \cup \Pi_2$ be the common refinement. Then by (9.2), we have

$$U(\Pi^*, f, \alpha) \leq U(\Pi_2, f, \alpha) < \int f d\alpha + \frac{\epsilon}{2} < L(f, \Pi_1, \alpha) + \frac{\epsilon}{2} + \frac{\epsilon}{2} \leq L(f, \Pi^*, \alpha) + \epsilon.$$

Thus (9.3) is verified by the partition Π^* , concluding the proof. \square

With the characterization of Lemma 9.10, we can now easily see that continuous functions are always Riemann-Stieltjes integrable.

THEOREM 9.11. *Let $a < b$ in \mathbb{R} , let $f: [a, b] \rightarrow \mathbb{R}$ be continuous, and let $\alpha: [a, b] \rightarrow \mathbb{R}$ be monotone increasing. Then $f \in \mathcal{R}(\alpha)$.*

PROOF. First, if α is constant on $[a, b]$, then $\Delta\alpha_j = 0$ on any interval $[t_{j-1}, t_j]$, and so $U(f, \Pi, \alpha) = L(f, \Pi, \alpha) = 0$ for all Π ; thus $U(f, \alpha) = L(f, \alpha) = \int_a^b f d\alpha = 0$. Otherwise, we must have $\alpha(a) < \alpha(b)$. Now, since f is continuous on the compact interval $[a, b]$, it is uniformly continuous there, so we may choose $\delta > 0$ such that, for any $x, y \in [a, b]$ with $|x - y| < \delta$, it follows that $|f(x) - f(y)| < \frac{\epsilon}{\alpha(b) - \alpha(a)}$.

Now, fix a partition $\Pi = \{a = t_0 < t_1 < \dots < t_n = b\}$ of $[a, b]$ for which $\Delta t_j = t_j - t_{j-1} < \delta$ for all j ; for example, fix n with $\frac{b-a}{n} < \delta$ and use the even partition $\Pi = \{a, a + \frac{b-a}{n}, a + 2\frac{b-a}{n}, \dots, a + (n-1)\frac{b-a}{n}, b\}$. Since f is continuous on each interval $[t_{j-1}, t_j]$, there are points x_j and y_j such that

$$\sup_{t_{j-1} \leq t \leq t_j} f(t) = f(x_j) \quad \text{and} \quad \inf_{t_{j-1} \leq t \leq t_j} f(t) = f(y_j).$$

Since $t_{j-1} \leq x_j, y_j \leq t_j$ and $t_j - t_{j-1} < \delta$ it follows that $|x_j - y_j| < \delta$, and so $f(x_j) - f(y_j) = |f(x_j) - f(y_j)| < \frac{\epsilon}{\alpha(b) - \alpha(a)}$. Thus

$$U(f, \Pi, \alpha) - L(f, \Pi, \alpha) = \sum_{j=1}^n [f(x_j) - f(y_j)] \Delta\alpha_j < \sum_{j=1}^n \frac{\epsilon}{\alpha(b) - \alpha(a)} \Delta\alpha_j.$$

Factoring out the constant $\frac{\epsilon}{\alpha(b) - \alpha(a)}$, we have just the telescoping sum $\sum_{j=1}^n \Delta\alpha_j = \alpha(b) - \alpha(a)$, and so we see that with this partition, $U(f, \Pi, \alpha) - L(f, \Pi, \alpha) < \epsilon$. As we can do this for any $\epsilon > 0$, by Lemma 9.10, it follows that $f \in \mathcal{R}(\alpha)$, as desired. \square

So, we now know how to integrate continuous functions. In particular, taking $\alpha(x) = x$, this gives us the usual Riemann integral of continuous functions, which is the main object of study in integral calculus. However, the class of functions that can be integrated is much larger than continuous functions. For example:

THEOREM 9.12. *Let $a < b$ in \mathbb{R} , let $f: [a, b] \rightarrow \mathbb{R}$ be bounded, and let $\alpha: [a, b] \rightarrow \mathbb{R}$ be monotone increasing. Suppose that the set of points in $[a, b]$ where f is discontinuous is finite, and at each such point α is continuous. Then $f \in \mathcal{R}(\alpha)$.*

PROOF. As in the proof of Theorem 9.11, if α is constant there is nothing to prove, so we freely assume $\alpha(b) < \alpha(a)$. Fix $\epsilon > 0$, and let $M = \sup |f|$. Let $D = \{x_1, x_2, \dots, x_d\}$ be the set of points in $[a, b]$ where f is not continuous. By assumption α is continuous at each x_j , so there is some $\eta_j > 0$ such that $|\alpha(s) - \alpha(t)| < \frac{\epsilon}{4Md}$ for all $s, t \in (x_j - \eta_j, x_j + \eta_j) \cap [a, b]$. Let $\eta \leq \min\{\eta_1, \dots, \eta_d\}$ be small enough that $(x_j - \eta, x_j + \eta) \subset (a, b)$ for all j (except if possibly $x_j = a$ or $x_j = b$). Now define points $u_j < v_j$ as follows: if $x_j = a$ then $u_j = a$ and $v_j = a + \frac{\epsilon}{2}$; if $x_j = b$ then $u_j = b - \frac{\eta}{2}$; and if $x_j \in (a, b)$ then $u_j = x_j - \frac{\eta}{2}$ and $v_j = x_j + \frac{\eta}{2}$. By definition of η , it follows that $\alpha(v_j) - \alpha(u_j) < \frac{\epsilon}{4Md}$ for $1 \leq j \leq d$, and so

$$\sum_{j=1}^d [\alpha(v_j) - \alpha(u_j)] < \frac{\epsilon}{4M}.$$

Now, let $K = [a, b] \setminus \bigcup_{j=1}^d (u_j, v_j)$. This is a closed subset of $[a, b]$ and so is compact. Because no point in D is in K , f is continuous on K , and hence uniformly continuous there. So we may choose $\delta > 0$ so that, whenever $s, t \in K$ with $|s - t| < \delta$, it follows that

$$|f(s) - f(t)| < \frac{\epsilon}{2(\alpha(b) - \alpha(a))}.$$

Now form a partition Π of $[a, b]$ as follows: Π contains all the points u_j and v_j for $1 \leq j \leq d$; it contains no points in any of the intervals (u_j, v_j) ; and for any point t_i in Π not of the form v_j , we have $t_i - t_{i-1} < \delta$.

We now expand $U(f, \Pi, \alpha) - L(f, \Pi, \alpha)$:

$$U(f, \Pi, \alpha) - L(f, \Pi, \alpha) = \sum_i \left[\sup_{t_{i-1} \leq t \leq t_i} f(t) - \inf_{t_{i-1} \leq t \leq t_i} f(t) \right] \Delta\alpha_i.$$

We break this sum into two kinds of terms, dividing the indices into $i \in I_1$ and $i \in I_2$: I_1 consists of those i for which $t_i = v_j$ for some j , and I_2 consists of all the others. For $i \in I_1$, by construction $t_{i-1} = u_j$. For these terms we make the estimate that $\sup_{t_{i-1} \leq t \leq t_i} f(t) - \int_{t_{i-1} \leq t \leq t_i} f(t) \leq 2M$ and so we get

$$\begin{aligned} \sum_{i \in I_1} \left[\sup_{t_{i-1} \leq t \leq t_i} f(t) - \inf_{t_{i-1} \leq t \leq t_i} f(t) \right] \Delta\alpha_i &\leq 2M \sum_{i \in I_1} \Delta\alpha_i \\ &= 2M \sum_{i \in I_1} [\alpha(t_i) - \alpha(t_{i-1})] \\ &= 2M \sum_{j=1}^d [\alpha(v_j) - \alpha(u_j)] \\ &< 2M \cdot \frac{\epsilon}{4M} = \frac{\epsilon}{2}. \end{aligned}$$

Now, for those $i \in I_2$, we have constructed Π so that $t_i - t_{i-1} < \delta$, which means that for any $s, t \in [t_{i-1}, t_i]$ $|f(s) - f(t)| \leq \frac{\epsilon}{2[\alpha(b) - \alpha(a)]}$, and so this also holds true for the difference between the supremal and infimal values of f on the interval. Summing up these terms yields

$$\begin{aligned} \sum_{i \in I_2} \left[\sup_{t_{i-1} \leq t \leq t_i} f(t) - \inf_{t_{i-1} \leq t \leq t_i} f(t) \right] \Delta\alpha_i &\leq \sum_{i \in I_2} \frac{\epsilon}{2[\alpha(b) - \alpha(a)]} \cdot [\alpha(t_i) - \alpha(t_{i-1})] \\ &\leq \frac{\epsilon}{2[\alpha(b) - \alpha(a)]} \sum_i [\alpha(t_i) - \alpha(t_{i-1})] = \frac{\epsilon}{2}. \end{aligned}$$

Thus $U(f, \Pi, \alpha) - L(f, \Pi, \alpha) < \frac{\epsilon}{2} + \frac{\epsilon}{2} = \epsilon$, and so by Lemma 9.10, it follows that $f \in \mathcal{R}(\alpha)$. \square

REMARK 9.13. If both f and α are discontinuous at a point, it is possible that $f \notin \mathcal{R}(\alpha)$, regardless of how continuous f is elsewhere. You will work with such an example on this week's homework.

REMARK 9.14. It is important to note that Theorem 9.12 applies to discontinuities in general, not just jump discontinuities. In particular, we now know that the function $f(t) = \sin \frac{1}{t}$ of Example 7.18 is Riemann-Stieltjes integrable (with respect to any monotone increasing integrator α) on any compact interval, including 0 or not. And the proof shows us why. Even though the function oscillates wildly near 0, it is bounded by 1 in absolute value, and so for any small $\delta > 0$, the total contribution of the upper or lower sums over any partition from points in $(-\delta, \delta)$ is no bigger than $\alpha(\delta) - \alpha(-\delta)$, which is very small so long as α is continuous at 0.

REMARK 9.15. It is natural to wonder how much further this can be taken, in terms of allowing irregular f to be integrated. For example, the above proof does not immediately generalize to the case where f has countably infinitely many discontinuities. It is a delicate matter in general to settle which f are in $\mathcal{R}(\alpha)$ for a particular α . The most important case where $\alpha(x) = x$ (the Riemann integral), however, is completely understood. In that case, even if f has (at most) countably infinitely many discontinuities, it is still Riemann integrable. In that case, the exact criterion for Riemann integrability is *continuous almost everywhere*: a bounded function f is in \mathcal{R} iff, for every $\epsilon > 0$, there is a countable collection of open intervals I_1, I_2, I_3, \dots with $\sum_j \text{length}(I_j) < \epsilon$ such that the set of discontinuities of f is contained in $\bigcup_j I_j$. (I.e. the set of discontinuities has “measure 0”)

3. Lecture 10: April 28, 2016

We now turn to the basic properties of the Riemann-Stieltjes integral, all of which reflect its nature as a kind of limit sum. We state them as a sequence of lemmas. In all cases, $a < b$ are real numbers, $f, f_1, f_2: [a, b] \rightarrow \mathbb{R}$ are bounded functions, and $\alpha, \alpha_1, \alpha_2: [a, b] \rightarrow \mathbb{R}$ are monotone increasing functions.

LEMMA 9.16. *If $f_1, f_2 \in \mathcal{R}(\alpha)$ and $c \in \mathbb{R}$, then $f_1 + f_2 \in \mathcal{R}(\alpha)$ and $cf_1 \in \mathcal{R}(\alpha)$, with*

$$\int_a^b (f_1 + f_2) d\alpha = \int_a^b f_1 d\alpha + \int_a^b f_2 d\alpha \quad \text{and} \quad \int_a^b cf_1 d\alpha = c \int_a^b f_1 d\alpha.$$

PROOF. Fix a partition $\Pi = \{a = t_0 < t_1 < \cdots < t_n = b\}$. To save notation, let $I_j = [t_{j-1}, t_j]$. Since

$$\inf_{I_j} f_1 + \inf_{I_j} f_2 \leq \inf_{I_j} (f_1 + f_2) \leq \sup_{I_j} (f_1 + f_2) \leq \sup_{I_j} f_1 + \sup_{I_j} f_2$$

multiplying by $\Delta\alpha_j$ and summing yields

$$\begin{aligned} L(f_1, \Pi, \alpha) + L(f_2, \Pi, \alpha) &\leq L(f_1 + f_2, \Pi, \alpha) \\ &\leq U(f_1 + f_2, \Pi, \alpha) \leq U(f_1, \Pi, \alpha) + U(f_2, \Pi, \alpha). \end{aligned} \quad (9.4)$$

Since $f_j \in \mathcal{R}(\alpha)$ for $j = 1, 2$, there are partitions Π_j such that $U(f_j, \Pi_j, \alpha) - L(f_j, \Pi_j, \alpha) < \frac{\epsilon}{2}$. Letting $\Pi^* = \Pi_1 \cup \Pi_2$ as usual, we then have $U(f_j, \Pi^*, \alpha) \leq U(f_j, \Pi_j, \alpha)$ and $L(f_j, \Pi^*, \alpha) \geq L(f_j, \Pi_j, \alpha)$, so $U(f_j, \Pi^*, \alpha) - L(f_j, \Pi^*, \alpha) < \frac{\epsilon}{2}$. Adding up and applying (9.4) (with Π^* in place of Π) yields

$$\begin{aligned} &U(f_1 + f_2, \Pi^*, \alpha) - L(f_1 + f_2, \Pi^*, \alpha) \\ &\leq U(f_1, \Pi^*, \alpha) + U(f_2, \Pi^*, \alpha) - L(f_1, \Pi^*, \alpha) - L(f_2, \Pi^*, \alpha) < \frac{\epsilon}{2} + \frac{\epsilon}{2} = \epsilon. \end{aligned}$$

Hence $f_1 + f_2 \in \mathcal{R}(\alpha)$. What's more, since $L(f_j, \Pi^*, \alpha) \leq \int f_j d\alpha \leq U(f_j, \Pi^*, \alpha)$, it follows that $U(f_j, \Pi^*, \alpha) < \int f_j d\alpha + \frac{\epsilon}{2}$ and $L(f_j, \Pi^*, \alpha) > \int f_j d\alpha - \frac{\epsilon}{2}$. Thus, applying (9.4) again, we have

$$\int f_1 d\alpha + \int f_2 d\alpha - \epsilon < L(f_1 + f_2, \Pi^*, \alpha) \leq U(f_1 + f_2, \Pi^*, \alpha) < \int f_1 d\alpha + \int f_2 d\alpha + \epsilon.$$

Taking sup and inf as appropriate shows that $\int (f_1 + f_2) d\alpha$ is distance $< \epsilon$ away from $\int f_1 d\alpha + \int f_2 d\alpha$ for each $\epsilon > 0$, and this establishes the first equality.

For the second, we note that for $c > 0$, we simply have $U(cf_1, \Pi, \alpha) = cU(f_1, \Pi, \alpha)$ and $L(cf_1, \Pi, \alpha) = cL(f_1, \Pi, \alpha)$ for any partition Π , and hence if we choose a Π so that $U(f_1, \Pi, \alpha) - L(f_1, \Pi, \alpha) < \frac{\epsilon}{c}$, then $U(cf_1, \Pi, \alpha) - L(cf_1, \Pi, \alpha) = c[U(f_1, \Pi, \alpha) - L(f_1, \Pi, \alpha)] < \epsilon$, so $cf_1 \in \mathcal{R}(\alpha)$, and moreover $\int cf_1 d\alpha = U(cf_1, \alpha) = \inf_{\Pi} U(cf_1, \Pi, \alpha) = \inf_{\Pi} cU(f_1, \Pi, \alpha) = c \cdot \inf_{\Pi} U(f_1, \Pi, \alpha) = cU(f, \alpha) = c \int f_1 d\alpha$. If, on the other hand, $c < 0$, then multiplication by c interchanges sup and inf and so $U(cf_1, \Pi, \alpha) = cL(f_1, \Pi, \alpha)$ and $L(cf_1, \Pi, \alpha) = cU(f_1, \Pi, \alpha)$. A very similar argument now yields the result. If $c = 0$, there is nothing to prove as both sides are 0. \square

LEMMA 9.17. *If $f_1 \leq f_2$ on $[a, b]$ and $f_1, f_2 \in \mathcal{R}(\alpha)$, then $\int f_1 d\alpha \leq \int f_2 d\alpha$.*

PROOF. The inequality $f_1(t) \leq f_2(t)$ implies that $\sup_{I_j} f_1 \leq \sup_{I_j} f_2$ for all intervals I_j , and hence $U(f_1, \Pi, \alpha) \leq U(f_2, \Pi, \alpha)$ for any partition Π . Taking \inf_{Π} now yields $\int f_1 d\alpha = U(f_1, \alpha) \leq U(f_2, \alpha) = \int f_2 d\alpha$. \square

LEMMA 9.18. *If $a < c < b$ and $f \in \mathcal{R}(\alpha)$ on $[a, b]$, then it is in $\mathcal{R}(\alpha)$ on $[a, c]$ and on $[c, b]$, and*

$$\int_a^b f d\alpha = \int_a^c f d\alpha + \int_c^b f d\alpha.$$

PROOF. Set $f_1 = f\mathbb{1}_{[a,c]}$ and $f_2 = f\mathbb{1}_{[c,b]}$. Then both f_1 and f_2 are in $\mathcal{R}(\alpha)$. Indeed, if Π is a partition for which $U(f, \Pi, \alpha) - L(f, \Pi, \alpha) < \epsilon$, then we may (possibly) refine it by taking $\Pi^* = \Pi \cup \{c\}$ and noting that we still have $U(f, \Pi^*, \alpha) - L(f, \Pi^*, \alpha) < \epsilon$. Let m be the index in $\Pi^* = \{a = t_0 < t_1 < \dots < t_n = b\}$ with $t_m = c$, and let $I_j = [t_{j-1}, t_j]$; then

$$\begin{aligned} U(f_1, \Pi^*, \alpha) - L(f_1, \Pi^*, \alpha) &= \sum_{j=1}^n [\sup_{I_j} f - \inf_{I_j} f] \Delta\alpha_j \\ &= \sum_{j=1}^m [\sup_{I_j} f - \inf_{I_j} f] \Delta\alpha_j \leq U(f, \Pi^*, \alpha) - L(f, \Pi^*, \alpha) < \epsilon \end{aligned}$$

since the omitted terms are all ≥ 0 . A similar calculation shows that $U(f_2, \Pi^*, \alpha) - L(f_2, \Pi^*, \alpha)$ is the sum of terms of index $\geq m$ which is \leq the full sum, hence $< \epsilon$. So both f_1, f_2 are in $\mathcal{R}(\alpha)$, and by Lemma 9.16,

$$\int_a^b f d\alpha = \int_a^b f_1 d\alpha + \int_a^b f_2 d\alpha.$$

It is left to the reader to establish (using very similar arguments) that $\int_a^b f_1 d\alpha = \int_a^c f d\alpha$ and $\int_a^b f_2 d\alpha = \int_c^b f d\alpha$, concluding the proof. \square

LEMMA 9.19. *If $f \in \mathcal{R}(\alpha_1)$ and $f \in \mathcal{R}(\alpha_2)$ then $f \in \mathcal{R}(\alpha_1 + \alpha_2)$ and $f \in \mathcal{R}(c\alpha_1)$ for $c > 0$, with*

$$\int_a^b f d(\alpha_1 + \alpha_2) = \int_a^b f d\alpha_1 + \int_a^b f d\alpha_2 \quad \text{and} \quad \int_a^b f d(c\alpha_1) = c \int_a^b f d\alpha_1.$$

PROOF. Here we simply note that

$$\begin{aligned} \Delta(\alpha_1 + \alpha_2)_j &= (\alpha_1 + \alpha_2)(t_j) - (\alpha_1 + \alpha_2)(t_{j-1}) \\ &= [\alpha_1(t_j) - \alpha_1(t_{j-1})] + [\alpha_2(t_j) - \alpha_2(t_{j-1})] = \Delta(\alpha_1)_j + \Delta(\alpha_2)_j, \end{aligned}$$

and similarly $\Delta(c\alpha_1)_j = c\Delta(\alpha_1)_j$, for each j . Since all these increments are ≥ 0 , it then follows that

$$U(f, \Pi, \alpha_1 + \alpha_2) = U(f, \Pi, \alpha_1) + U(f, \Pi, \alpha_2) \quad \text{and} \quad U(f, \Pi, c\alpha_1) = cU(f, \Pi, \alpha_1).$$

Taking \inf_{Π} yields we get $U(f, \alpha_1 + \alpha_2) = U(f, \alpha_1) + U(f, \alpha_2) = \int f d\alpha_1 + \int f d\alpha_2$ and $U(f, c\alpha_1) = cU(f, \alpha_1) = c \int f d\alpha_1$. Similar considerations with lower sums show that these two are equal to the given linear combinations, concluding the proof. \square

LEMMA 9.20. *Let $f \in \mathcal{R}(\alpha)$. Let $M = \sup f$ and $m = \inf f$, and suppose $\phi: [m, M] \rightarrow \mathbb{R}$ is continuous. Then $h = \phi \circ f$ is in $\mathcal{R}(\alpha)$.*

PROOF. Fix $\epsilon > 0$. Denote by $C = \sup |\phi|$; we assume it is > 0 , otherwise $\phi = 0$ and the statement of the lemma is silly; similarly, we assume $\alpha(b) > \alpha(a)$. As ϕ is uniformly continuous on $[m, M]$, there is some $\delta > 0$ so that $|\phi(x) - \phi(y)| < \frac{\epsilon}{2[\alpha(b) - \alpha(a)]}$ whenever $|x - y| < \delta$. For later convenience, we will make sure to select $\delta \leq \frac{\epsilon}{4C}$.

Since $f \in \mathcal{R}(\alpha)$, there is a partition Π of $[a, b]$ such that $U(f, \Pi, \alpha) - L(f, \Pi, \alpha) < \delta^2$. Now the difference $U(h, \Pi, \alpha) - L(h, \Pi, \alpha)$ is equal to

$$\sum_j \left[\sup_{I_j} h - \inf_{I_j} h \right] \Delta\alpha_j.$$

We break up the sum into two parts: $j \in J_1 \sqcup J_2$, where $j \in J_1$ iff $\sup_{I_j} f - \inf_{I_j} f < \delta$, and $j \in J_2$ iff this difference is $\geq \delta$. The choice of δ shows that, for $j \in J_1$, $\sup_{I_j} h - \inf_{I_j} h < \frac{\epsilon}{2[\alpha(b) - \alpha(a)]}$. For $j \in J_2$, on the other hand, the best we can say in general is that $\sup_{I_j} h - \inf_{I_j} h \leq 2 \sup |h| = 2C$. But we've arranged things so that the total α -length of the intervals indexed by J_2 is small: because $\sup_{I_j} f - \inf_{I_j} f \geq \delta$ for $j \in J_2$,

$$\delta \sum_{j \in J_2} \Delta\alpha_j \leq \sum_{j \in J_2} \left[\sup_{I_j} f - \inf_{I_j} f \right] \Delta\alpha_j = U(f, \Pi, \alpha) - L(f, \Pi, \alpha) < \delta^2$$

and so $\sum_{j \in J_2} \Delta\alpha_j < \delta$. Thus, adding up,

$$\begin{aligned} U(h, \Pi, \alpha) - L(h, \Pi, \alpha) &= \sum_{j \in J_1} \left[\sup_{I_j} h - \inf_{I_j} h \right] \Delta\alpha_j + \sum_{j \in J_2} \left[\sup_{I_j} h - \inf_{I_j} h \right] \Delta\alpha_j \\ &< \frac{\epsilon}{2[\alpha(b) - \alpha(a)]} \sum_{j \in J_1} \Delta\alpha_j + 2C \sum_{j \in J_2} \Delta\alpha_j \\ &\leq \frac{\epsilon}{2[\alpha(b) - \alpha(a)]} \sum_j \Delta\alpha_j + 2C\delta \\ &< \frac{\epsilon}{2} + 2C \cdot \frac{\epsilon}{4C} = \epsilon. \end{aligned}$$

This shows that $h \in \mathcal{R}(\alpha)$, as desired. \square

REMARK 9.21. The above proof shows no simple connection between the value of $\int \phi \circ f \, d\alpha$ and $\int f \, d\alpha$, and indeed there is no simple connection.

LEMMA 9.22. *If $f, g \in \mathcal{R}(\alpha)$, then so is fg .*

PROOF. By Lemma 9.16, $f \pm g$ are both in $\mathcal{R}(\alpha)$. The function $\phi(x) = x^2$ is continuous, and so by Lemma 9.20, $(f + g)^2$ and $(f - g)^2$ are both in $\mathcal{R}(\alpha)$. Thus, by Lemma 9.16 again, so is

$$\frac{1}{4} [(f + g)^2 - (f - g)^2] = fg.$$

\square

LEMMA 9.23. *If $f \in \mathcal{R}(\alpha)$, then $|f| \in \mathcal{R}(\alpha)$, and $\left| \int f \, d\alpha \right| \leq \int |f| \, d\alpha$.*

PROOF. The function $\phi(x) = |x|$ is continuous, and so by Lemma 9.20, $|f| \in \mathcal{R}(\alpha)$. Now, let $\sigma = \text{sgn} \left(\int f \, d\alpha \right)$ (so $\sigma = 1$ if the integral is ≥ 0 and $\sigma = -1$ if the integral is < 0 .) Then by Lemma 9.16,

$$\left| \int f \, d\alpha \right| = \sigma \int f \, d\alpha = \int (\sigma f) \, d\alpha.$$

For each t , we have $f(t) \leq |f(t)|$ and $-f(t) \leq |f(t)|$; thus $\sigma f \leq f$. Thus, by Lemma 9.17,

$$\left| \int f \, d\alpha \right| = \int (\sigma f) \, d\alpha \leq \int |f| \, d\alpha.$$

□

4. Lecture 11: May 3, 2016

We now know how to integrate (or more precisely that we *can* integrate) a reasonably large class of functions, against arbitrary monotone increasing integrators. One may then reasonably ask: what benefit have we gained by including more general integrator weights α ? To partly answer this, consider the following example.

EXAMPLE 9.24. Let $a < b$ in \mathbb{R} , and fix some $s \in (a, b)$. Let α_s be the monotone function $\alpha_s(t) = \mathbb{1}_{(s, b]}(t)$, i.e. $\alpha_s(t) = 0$ if $t \leq s$ and $\alpha_s(t) = 1$ if $t > s$. If f is continuous on $[a, b]$ then

$$\int_a^b f d\alpha_s = f(s).$$

To see this, fix any partition $\Pi = \{a = t_0 < t_1 < \dots < t_n = b\}$. Let $m \geq 1$ be the unique index such that $t_{m-1} < s \leq t_m$. For any $j \neq m$, α_s is constant on $[t_{j-1}, t_j]$ and so $\Delta\alpha_j = 0$, while $\Delta\alpha_m = 1 - 0 = 1$. Thus

$$U(f, \Pi, \alpha_s) = \sup_{t_{m-1} \leq t \leq t_m} f(t), \quad L(f, \Pi, \alpha_s) = \inf_{t_{m-1} \leq t \leq t_m} f(t).$$

Since $s \in [t_{m-1}, t_m]$, $\inf_{t_{m-1} \leq t \leq t_m} f(t) \leq f(s) \leq \sup_{t_{m-1} \leq t \leq t_m} f(t)$. It follows that $L(f, \alpha_s) \leq f(s) \leq U(f, \alpha_s)$. Since f is continuous, Theorem 9.11 implies that $L(f, \alpha_s) = U(f, \alpha_s)$, and so this common value must be $\int f d\alpha_s = f(s)$ as claimed. (A more careful argument shows that this holds true even if we only assume that f is continuous at s .)

Example 9.24 gives a rigorous treatment of a “delta function”. Physicists love to use delta functions: a “function” $\delta(t)$ with the property that, for $s \in (a, b)$,

$$\int_a^b f(t)\delta(t-s) dt = f(s).$$

In fact, there is no such function $\delta(t)$ which Riemann integrates a function by evaluating it at a point. Instead, $\delta(t-s) dt$ must be interpreted as the Riemann–Stieltjes integrator $d\alpha_s(t)$.

We can use the additivity of the integral to generalize Example 9.24, and put discrete infinite sums on the same footing as integrals, and treat them all as one kind of object. If $(s_n)_{n=1}^{\infty}$ is an increasing (possibly finite) sequence in (a, b) , and if (c_n) is a nonnegative sequence such that $\sum c_n < \infty$, we can define a *pure step function*

$$\alpha(t) = \sum_{n=1}^{\infty} c_n \alpha_{s_n}(t).$$

This is a monotone increasing function on $[a, b]$. It is constant on $[a, a + s_1]$, taking value c_1 ; its value on $(a + s_1, a + s_2]$ is $c_1 + c_2$; and so forth.

LEMMA 9.25. Let α be the pure step function above. If f is continuous on $[a, b]$, then

$$\int_a^b f d\alpha = \sum_{n=1}^{\infty} c_n f(s_n).$$

PROOF. Since $|\alpha_s(t)| \leq 1$ for any s , by the comparison test, the series $\alpha(t) = \sum_{n=1}^{\infty} c_n \alpha_{s_n}(t)$ converges for all t . The function α defined is monotone increasing, since $c_n \geq 0$ for all n , and $\alpha_s(t) \leq \alpha_s(t')$ whenever $t < t'$. Hence, it makes sense to compute $\int_a^b f d\alpha$. (Note also that $\alpha(a) = 0$ while $\alpha(b) = \sum_n c_n$.) Similarly, since f is bounded on $[a, b]$, the sum $\sum_{n=1}^{\infty} c_n f(s_n)$ also converges by the comparison test, so all quantities presented make sense.

Now, let $\epsilon > 0$, and let $M = \sup |f|$ (and assume $M > 0$ to avoid silliness). By the Cauchy criterion for convergence of $\sum_n c_n$, there is some $N \in \mathbb{N}$ so that $\sum_{n>N} c_n < \frac{\epsilon}{M}$. We then break up α accordingly: $\alpha(t) = \alpha_1(t) + \alpha_2(t)$, where

$$\alpha_1(t) = \sum_{n=1}^N c_n \alpha_{s_n}(t), \quad \alpha_2(t) = \sum_{n=N+1}^{\infty} c_n \alpha_{s_n}(t).$$

By Lemma 9.19, we have

$$\int f d\alpha = \int f d\alpha_1 + \int f d\alpha_2.$$

The first integral, by induction on Lemma 9.19, has the value

$$\int f d\alpha_1 = \sum_{n=1}^N c_n \int f d\alpha_{s_n} = \sum_{n=1}^N c_n f(s_n).$$

Thus, by Lemma 9.23 and Lemma 9.17,

$$\left| \int f d\alpha - \int f d\alpha_1 \right| = \left| \int f d\alpha_2 \right| \leq \int |f| d\alpha_2 \leq \int M d\alpha_2 = M[\alpha_2(b) - \alpha_2(a)].$$

Since (s_n) is an increasing sequence, $s_{N+1} > a$, and so $\alpha_{s_n}(a) = 0$ for $n > N$; thus $\alpha_2(a) = 0$. On the other hand $\alpha_{s_n}(b) = 1$ for each n , and so $\alpha_2(b) = \sum_{n=N+1}^{\infty} c_n < \frac{\epsilon}{M}$ by construction. Thus the above inequality shows that

$$\left| \int f d\alpha - \sum_{n=1}^N c_n f(s_n) \right| = \left| \int f d\alpha - \int f d\alpha_1 \right| < \frac{\epsilon}{M} \cdot M = \epsilon.$$

Since we can do this for each $\epsilon > 0$, and since we know $\sum_{n=1}^N c_n f(s_n)$ converges as $N \rightarrow \infty$, it follows that it must converge to $\int f d\alpha$ as claimed. \square

Thus, if we use a pure step function as our integrator, we unify the theory of infinite series and integration. On the flip side, what happens if α is the opposite of a purely discrete function: what if α is differentiable? In that case, we will see that integration with respect to α actually involved the derivative α' (and this will motivate the connection afterward to the Fundamental Theorem of Calculus). As you should be used to by now, anything involving derivatives will involve the Mean Value Theorem, which will mean evaluating the derivative of the integrand on $[t_{j-1}, t_j]$ at a (random) point $\eta_j \in (t_{j-1}, t_j)$. As such, it will be convenient in the next several results to connect our formulation of the integral, in terms of upper and lower sums, to Riemann(–Stieltjes) sums.

LEMMA 9.26. *Let f be in $\mathcal{R}(\alpha)$ on $[a, b]$, let $\epsilon > 0$, and let Π be a partition for which $U(f, \Pi, \alpha) - L(f, \Pi, \alpha) < \epsilon$ (cf. Lemma 9.10). With $\Pi = \{a = t_0 < t_1 < t_2 < \dots < t_n = b\}$ and $I_j = [t_{j-1}, t_j]$, let $\eta_j \in I_j$. Then*

$$\left| \sum_{j=1}^n f(\eta_j) \Delta\alpha_j - \int_a^b f d\alpha \right| < \epsilon.$$

PROOF. By definition $\inf_{I_j} f \leq f(\eta_j) \leq \sup_{I_j} f$ for each j . Multiplying by the non-negative numbers $\Delta\alpha_j$ and summing up shows that

$$L(f, \Pi, \alpha) \leq \sum_{j=1}^n f(\eta_j) \Delta\alpha_j \leq U(f, \Pi, \alpha). \quad (9.5)$$

Of course, we also have

$$L(f, \Pi, \alpha) \leq \int_a^b f d\alpha \leq U(f, \Pi, \alpha). \quad (9.6)$$

So the Riemann–Stieltjes sum in question and the integral $\int f d\alpha$ are both in the interval from $L(f, \Pi, \alpha)$ up to $U(f, \Pi, \alpha)$. By assumption this interval has width $< \epsilon$. Thus, the distance between the Riemann–Stieltjes sum and the integral is $< \epsilon$, as claimed. \square

THEOREM 9.27. *Let $\alpha: [a, b] \rightarrow \mathbb{R}$ be monotone increasing, and suppose that α is differentiable on $[a, b]$ with $\alpha' \in \mathcal{R}$. Then for any bounded $f: [a, b] \rightarrow \mathbb{R}$, $f \in \mathcal{R}(\alpha)$ if and only if $f\alpha' \in \mathcal{R}$, and*

$$\int f d\alpha = \int f\alpha'.$$

This theorem is often summarized by writing $d\alpha(t) = \alpha'(t) dt$.

PROOF. Fix $\epsilon > 0$, and let $M = \sup |f|$ (assume $M > 0$ to avoid silliness). Since $\alpha' \in \mathcal{R}$, there is a partition $\Pi = \{a = t_0 < t_1 < \cdots < t_n = b\}$ of $[a, b]$ such that $U(\alpha', \Pi) - L(\alpha', \Pi) < \frac{\epsilon}{3M}$. By the Mean Value Theorem (applied to α), for each j there is a point $\xi_j \in (t_{j-1}, t_j)$ such that $\Delta\alpha_j = \alpha(t_j) - \alpha(t_{j-1}) = \alpha'(\xi_j) \Delta t_j$. Now, fix any points $\eta_j \in I_j$; then the Riemann–Stieltjes sum at the points η_j can nearly be expressed as a Riemann sum:

$$\sum_{j=1}^n f(\eta_j) \Delta\alpha_j = \sum_{j=1}^n f(\eta_j) \alpha'(\xi_j) \Delta t_j.$$

We would like to convert this to a proper Riemann sum at the points η_j ; to do so, we need to compare the points η_j to ξ_j , which results in a correction factor

$$\begin{aligned} \left| \sum_{j=1}^n f(\eta_j) \Delta\alpha_j - \sum_{j=1}^n f(\eta_j) \alpha'(\eta_j) \Delta t_j \right| &= \left| \sum_{j=1}^n f(\eta_j) [\alpha'(\xi_j) - \alpha'(\eta_j)] \Delta t_j \right| \\ &\leq \sum_{j=1}^n |f(\eta_j)| |\alpha'(\xi_j) - \alpha'(\eta_j)| \Delta t_j \\ &\leq M \sum_{j=1}^n |\alpha'(\xi_j) - \alpha'(\eta_j)| \Delta t_j. \end{aligned}$$

Now, since $\xi_j, \eta_j \in I_j$, it follows that $|\alpha'(\xi_j) - \alpha'(\eta_j)| \leq \sup_{I_j} \alpha' - \inf_{I_j} \alpha'$, and thus

$$\sum_{j=1}^n |\alpha'(\xi_j) - \alpha'(\eta_j)| \Delta t_j \leq \sum_{j=1}^n \left[\sup_{I_j} \alpha' - \inf_{I_j} \alpha' \right] \Delta t_j = U(\alpha', \Pi) - L(\alpha', \Pi) < \frac{\epsilon}{3M}.$$

Thus, we see that, for any points $\eta_j \in I_j$,

$$\left| \sum_{j=1}^n f(\eta_j) \Delta\alpha_j - \sum_{j=1}^n f(\eta_j) \alpha'(\eta_j) \Delta t_j \right| < \frac{\epsilon}{3}. \quad (9.7)$$

In particular, this shows that

$$\sum_{j=1}^n f(\eta_j) \Delta\alpha_j < \sum_{j=1}^n (f\alpha')(\eta_j) \Delta t_j + \frac{\epsilon}{3} \leq \sum_{j=1}^n \sup_{I_j} (f\alpha') \Delta t_j + \frac{\epsilon}{3} = U(f\alpha', \Pi) + \frac{\epsilon}{3}$$

holds for every choice of points $\eta_j \in I_j$, and so taking sup again we find that $U(f, \Pi, \alpha) < U(f\alpha', \Pi) + \frac{\epsilon}{3}$. Now, doing the same argument again in the other order shows that

$$\sum_{j=1}^n (f\alpha')(\eta_j) \Delta t_j < \sum_{j=1}^n f(\eta_j) \Delta \alpha_j + \frac{\epsilon}{3} \leq U(f, \Pi, \alpha) + \frac{\epsilon}{3}$$

and so $U(f\alpha', \Pi) < U(f, \Pi, \alpha) + \frac{\epsilon}{3}$ as well. Combining these shows that $|U(f, \Pi, \alpha) - U(f\alpha', \Pi)| < \frac{\epsilon}{3}$. An entirely analogous argument shows that $|L(f, \Pi, \alpha) - L(f\alpha', \Pi)| < \frac{\epsilon}{3}$ as well.

Thus, we have shown that, for any partition Π for which $U(\alpha', \Pi) - L(\alpha', \Pi) < \frac{\epsilon}{3M}$, it follows that $|U(f, \Pi, \alpha) - U(f\alpha', \Pi)| < \frac{\epsilon}{3}$ and $|L(f, \Pi, \alpha) - L(f\alpha', \Pi)| < \frac{\epsilon}{3}$. If we find such a partition Π , the antecedent will also hold true for any refinement of Π .

Now, to conclude the proof: suppose that $f \in \mathcal{R}(\alpha)$; then let Π' be a partition for which $U(f, \Pi', \alpha) - L(f, \Pi', \alpha) < \frac{\epsilon}{3}$. If Π is a partition for which $U(\alpha', \Pi) - L(\alpha', \Pi) < \frac{\epsilon}{3M}$, take $\Pi^* = \Pi \cup \Pi'$; then we have $U(\alpha', \Pi^*) - L(\alpha', \Pi^*) < \frac{\epsilon}{3M}$. Thus $|U(f, \Pi^*, \alpha) - U(f\alpha', \Pi^*)| < \frac{\epsilon}{3}$ and $|L(f, \Pi^*, \alpha) - L(f\alpha', \Pi^*)| < \frac{\epsilon}{3}$. Since Π^* refines Π' , it is also true that $U(f, \Pi^*, \alpha) - L(f, \Pi^*, \alpha) < \frac{\epsilon}{3}$. Thus

$$\begin{aligned} & U(f\alpha', \Pi^*) - L(f\alpha', \Pi^*) \\ & \leq |U(f\alpha', \Pi^*) - U(f, \Pi^*, \alpha)| + |U(f, \Pi^*, \alpha) - L(f, \Pi^*, \alpha)| + |L(f, \Pi^*, \alpha) - L(f\alpha', \Pi^*)| \\ & < \frac{\epsilon}{3} + \frac{\epsilon}{3} + \frac{\epsilon}{3} = \epsilon. \end{aligned}$$

Thus $f\alpha' \in \mathcal{R}$ as well; and the fact that we can find a partition Π^* such that $|U(f, \Pi^*, \alpha) - U(f\alpha', \Pi^*)| < \frac{\epsilon}{3}$ for each $\epsilon > 0$ shows that $|\int f d\alpha - \int f\alpha'| < \frac{\epsilon}{3}$ for all $\epsilon > 0$, thus the two integrals are equal as claimed.

An entirely analogous argument beginning from the assumption that $f\alpha' \in \mathcal{R}$ shows that $f \in \mathcal{R}(\alpha)$ and that the two integrals are the same, concluding the proof. \square

This brings us to one of the most important tool for actually *computing* integrals.

THEOREM 9.28 (The Change of Variables Formula). *Let $a < b$ and $c < d$ in \mathbb{R} , and let $\varphi: [c, d] \rightarrow [a, b]$ be a strictly increasing surjective function. Let $\alpha: [a, b] \rightarrow \mathbb{R}$ be monotone increasing, and let $f \in \mathcal{R}(\alpha)$. Define $\beta, g: [c, d] \rightarrow \mathbb{R}$ by $\beta = \alpha \circ \varphi$ and $g = f \circ \varphi$. Then β is monotone increasing, $g \in \mathcal{R}(\beta)$, and*

$$\int_c^d g d\beta = \int_a^b f d\alpha.$$

PROOF. Since φ is strictly increasing, it is one-to-one; since it is surjective onto $[a, b]$, there is an inverse map $\varphi^{-1}: [a, b] \rightarrow [c, d]$. This gives a bijection between partitions Π of $[a, b]$ and partitions Θ of $[c, d]$: the correspondence is $t \in \Theta$ iff $\varphi(t) \in \Pi$, and so we write $\varphi(\Theta) = \Pi$. Writing out the definitions, we see that

$$U(g, \Theta, \beta) = U(f, \Pi, \alpha) \quad \text{and} \quad L(g, \Theta, \beta) = L(f, \Pi, \alpha).$$

Thus, finding a partition Π so that $U(f, \Pi, \alpha) - L(f, \Pi, \alpha) < \epsilon$ is equivalent to finding a partition $\Theta = \varphi^{-1}(\Pi)$ for which $U(g, \Theta, \beta) - L(g, \Theta, \beta) < \epsilon$, and thus $g \in \mathcal{R}(\beta)$. Because the map $\Pi \rightarrow \Theta$ is a bijection of partitions, $\int f d\alpha = \inf_{\Pi} U(f, \Pi, \alpha) = \inf_{\Theta} U(g, \Theta, \beta) = \int g d\beta$. \square

Theorem 9.28 may not look like the change of variables theorem you remember from calculus, but it is actually the generalization of it to the Riemann–Stieltjes integral. To restore the theorem you recall exactly, take the special case $\alpha(x) = x$. Then $\beta = \alpha \circ \varphi = \varphi$. If we further assume that φ is differentiable and $\varphi' \in \mathcal{R}$, then Theorem 9.27 shows that $d\beta(x) = \varphi'(x) dx$, and so

$$\int_a^b f(u) du = \int_a^b f d\alpha = \int_c^d f \circ \varphi d\beta = \int_c^d f(\varphi(x)) \varphi'(x) dx \quad (9.8)$$

which is the statement you learned in calculus.

REMARK 9.29. In calculus, you may have stated this theorem without the requirement that φ is strictly increasing. It is possible to generalize this, by reinterpreting the symbol \int_a^b to include *orientation*. Suppose, for example, that φ is strictly *decreasing*. Then $\beta = \alpha \circ \varphi$ is monotone decreasing, and therefore not the kind of integrator we know how to use. We could redo everything in this chapter so far for monotone decreasing integrators, and it would all work similarly, with appropriate minus signs thrown in. This can be accounted for by introducing the new (familiar from calculus) notation that if $a < b$ then $\int_b^a \equiv -\int_a^b$; this is what we mean by adding an orientation to the integral. With this in hand, everything works the same for monotone decreasing integrators, including Theorem 9.28. What's more, by employing Lemma 9.18, suitably reinterpreted in terms of the new orientation concept, we can even handle the case that φ is *piecewise* strictly monotone: there are finitely many points $a = x_0 < x_1 < x_2 < \cdots < x_p = b$ such that φ is strictly monotone on each interval (x_{j-1}, x_j) . We can even allow φ to be flat on some of the intervals (since the integral will just be 0 there). These are actually the kinds of functions for which the calculus change of variables formula works as stated above. Alternatively, one can restrict a little further to functions $\varphi \in C^1$ such that φ' has only finitely many zeroes (this is true for all the usual functions studied in calculus).

5. Lecture 12: May 5, 2016

We now come to the central result of calculus, appropriately called the Fundamental Theorem of Calculus: the Riemann integral is (more or less) the inverse of the derivative.

THEOREM 9.30 (Fundamental Theorem of Calculus). *Let $a < b$ in \mathbb{R} , and let $f \in \mathcal{R}$ on $[a, b]$.*

(a) *For $a \leq x \leq b$, define $F(x)$ by*

$$F(x) = \int_a^x f(t) dt.$$

Then F is Lipschitz continuous on $[a, b]$. Moreover, if f is continuous at a point $x_0 \in [a, b]$, then F is differentiable at x_0 , with $F'(x_0) = f(x_0)$.

(b) *If there exists a differentiable function $G: [a, b] \rightarrow \mathbb{R}$ such that $G' = f$, then $\int_a^b f(t) dt = G(b) - G(a)$.*

PROOF. For part (a), since $f \in \mathcal{R}$ it is bounded, say $|f| \leq M$. Then for any $x, y \in [a, b]$, say $x < y$; then by Lemma 9.18 and 9.23,

$$F(y) = \int_a^y f = \int_a^x f + \int_x^y f = F(x) + \int_x^y f$$

and so by Lemma 9.17

$$|F(y) - F(x)| = \left| \int_x^y f \right| \leq \int_x^y |f| \leq \int_x^y M = M(y - x).$$

This shows that F is Lipschitz continuous, with Lipschitz constant $\leq M$. (In fact, the Lipschitz constant is precisely $\sup |f|$.)

Now, suppose f is continuous at x_0 . Fix $\epsilon > 0$, and choose $\delta > 0$ so that $|f(t) - f(x_0)| < \epsilon$ whenever $|t - x_0| < \delta$. Then for $x_0 \leq t < x_0 + \delta$, we have

$$\frac{F(t) - F(x_0)}{t - x_0} - f'(x_0) = \frac{1}{t - x_0} \int_{x_0}^t f - f'(x_0) = \frac{1}{t - x_0} \int_{x_0}^t [f - f(x_0)]$$

where we have used the fact that $\int_{x_0}^t f(x_0) = f(x_0)(t - x_0)$. Because $|t - x_0| < \delta$, $|f - f(x_0)| < \epsilon$ on $[x_0, t]$, and so

$$\left| \frac{F(t) - F(x_0)}{t - x_0} - f'(x_0) \right| \leq \frac{1}{t - x_0} \int_{x_0}^t |f - f(x_0)| \leq \frac{1}{t - x_0} \int_{x_0}^t \epsilon = \epsilon.$$

An analogous argument shows that the difference quotient $DQF(t, x_0)$ is distance less than ϵ from $f'(x_0)$ in the case $x_0 - \delta < t < x_0$ as well. Thus, we have shown that $F'(x_0) = \lim_{t \rightarrow x_0} DQF(t) = f'(x_0)$, as claimed.

For part (b), fix $\epsilon > 0$, and let $\Pi = \{a = t_0 < t_1 < \dots < t_n = b\}$ be a partition for which $U(f, \Pi) - L(f, \Pi) < \epsilon$. By the Mean Value Theorem, for each j there is a point $\xi_j \in (t_{j-1}, t_j)$ such that $G(t_j) - G(t_{j-1}) = F'(\xi_j) \Delta t_j = f(\xi_j) \Delta t_j$. Thus, reversing the telescoping sum,

$$G(b) - G(a) = \sum_{j=1}^n [G(t_j) - G(t_{j-1})] = \sum_{j=1}^n f(\xi_j) \Delta t_j.$$

By Lemma 9.26, by the choice of Π , we have

$$\left| \int_a^b f - [G(b) - G(a)] \right| = \left| \int_a^b f - \sum_{j=1}^n f(\xi_j) \Delta t_j \right| < \epsilon.$$

As this holds for every $\epsilon > 0$, the result follows. \square

In summary: if f is continuous, then $F(x) = \int_a^x f(t) dt$ is an anti-derivative (a differentiable function with $F' = f$), and moreover if G is any anti-derivative, then $\int_a^b f = G(b) - G(a)$. Any two antiderivatives differ by a constant: if $F' = G'$ on $[a, b]$, then $(F - G)' = F' - G' = 0$. So $F - G$ is a differentiable function whose derivative is 0, which means it is constant by Corollary 8.19(2). So this is consistent: if $F = G + c$, then $F(b) - F(a) = G(b) - G(a)$.

In light of Theorem 9.30, the special case of the Change of Variables Formula in (9.8) is a straightforward consequence of the chain rule (under continuity assumptions). Indeed, let f be continuous and suppose φ is differentiable with $\varphi' \in \mathcal{R}$ on $[c, d]$. Let F be an anti-derivative of f (as in Theorem 9.30(a)), so that $F' = f$. Then $F \circ \varphi$ is differentiable, and by the chain rule $(F \circ \varphi)' = (F' \circ \varphi)\varphi' = (f \circ \varphi)\varphi'$. Employing the Fundamental Theorem of Calculus, we get

$$\int_c^d f(\varphi(x))\varphi'(x) dx = \int_c^d (F \circ \varphi)'(x) dx = F \circ \varphi(d) - F \circ \varphi(c).$$

Note that this holds regardless of whether φ is increasing (or even piecewise monotone). So as long as $\varphi(c) = a$ and $\varphi(d) = b$, we then have

$$\int_c^d f(\varphi(x))\varphi'(x) dx = F(b) - F(a) = \int_b^a f(u) du$$

again by the Fundamental Theorem of Calculus. This condition certainly holds if φ is strictly increasing from $[c, d]$ onto $[a, b]$, but this is not required; it could oscillate infinitely often as it fills out the interval. But this approach required the assumption that f is continuous, rather than just Riemann integrable; if one wants more general Riemann integrable functions f , the previous approach (which required φ be at least piecewise monotone) is required.

In the same light, let us now use the Fundamental Theorem of Calculus to turn the product rule into a powerful computational (and theoretical) tool for Riemann integration.

THEOREM 9.31 (Integration by Parts). *Let $a < b$ in \mathbb{R} , and let $f, g: [a, b] \rightarrow \mathbb{R}$ be differentiable functions with $f', g' \in \mathcal{R}$. Then fg' and $f'g$ are both in \mathcal{R} , and*

$$\int_a^b f(t)g'(t) dt = f(b)g(b) - f(a)g(a) - \int_a^b f'(t)g(t) dt.$$

PROOF. The chain rule gives $(fg)' = fg' + f'g$. By assumption f' and g' are both in \mathcal{R} , and so are f and g since they are differentiable (hence continuous). Thus, by Lemma 9.22, fg' and $f'g$ are both in \mathcal{R} , and by the Fundamental Theorem of Calculus and 9.16,

$$f(b)g(b) - f(a)g(a) = \int_a^b (fg)'(t) dt = \int_a^b f(t)g'(t) dt + \int_a^b f'(t)g(t) dt.$$

Subtracting yields the result. \square

As with the Change of Variables formula, we can wonder if there is a version for more general integrators. Indeed, suppose that g is differentiable, and g' is strictly increasing. Then Theorem 9.27 shows that $g'(t) dt = dg(t)$, and so we can rewrite Integration by Parts in the form

$$\int_a^b f dg = f(b)g(b) - f(a)g(a) - \int_a^b f'(t)g(t) dt.$$

In fact, this formula holds true even if g is not differentiable, as you will prove on your homework.

To conclude this chapter on integration, we consider an application to *curves*.

DEFINITION 9.32. Let $a < b$ in \mathbb{R} and let $d \in \mathbb{N}$. A (parametrized) curve in \mathbb{R}^d , with parameter domain $[a, b]$, is a continuous function $\gamma: [a, b] \rightarrow \mathbb{R}^d$. If $\gamma(a) = \gamma(b)$, it is called a closed curve. If γ is one-to-one on (a, b) , it is called a simple curve.

Note: a curve is more than just a path traced out in space (which is the image $\gamma([a, b])$); it also includes the information of the parametrization $t \mapsto \gamma(t)$, which is usually thought of as specifying the position of a particle at time t as it moves through space.

Our present goal is to measure the *length* of a curve. To that end, we begin by approximating the curve by pieces whose lengths we know how to measure: line segments. Fix a partition $\Pi = \{a = t_0 < t_1 < \dots < t_n = b\}$ of $[a, b]$. We replace γ by the path which passes through the points $\gamma(t_0), \gamma(t_1), \gamma(t_2), \dots, \gamma(t_n)$, and is a straight line segment between successive points. The length of each such line segment is just the Euclidean length of the difference vector: $|\gamma(t_j) - \gamma(t_{j-1})|$. Hence, we approximate the length of γ by

$$\Lambda(\gamma, \Pi) \equiv \sum_{j=1}^n |\gamma(t_j) - \gamma(t_{j-1})|.$$

Note that if we refine the partition used, the triangle inequality makes this length shrink: for any point $s \in (t_{j-1}, t_j)$, then

$$|\gamma(t_j) - \gamma(t_{j-1})| \leq |\gamma(t_j) - \gamma(s)| + |\gamma(s) - \gamma(t_{j-1})|.$$

It follows (by induction) that if Π^* is a refinement of Π then $\Lambda(\gamma, \Pi) \leq \Lambda(\gamma, \Pi^*)$. Indeed, this matches up with the fact (known as the isoperimetric inequality) that the shortest path between two points is the straight-line path. Thus, if we want to take a “limit” making partitions finer and finer, we ought to define the length of the curve γ by

$$\Lambda(\gamma) \equiv \sup_{\Pi} \Lambda(\gamma, \Pi).$$

This sup may well be infinite. A first guess at such an example would be something like the graph of the function $f(x) = \frac{1}{x}$, which has a vertical asymptote. But there is no way to parametrize this curve continuously on a *closed* interval, which is needed in our definition of curve. Nevertheless, there are continuous curves on closed intervals that have infinite length.

EXAMPLE 9.33. Consider the curve $\gamma: [0, 1] \rightarrow \mathbb{R}^2$ which traces out the graph of the function $f(x) = x \sin \frac{1}{x}$ (with $f(0) = 0$) of Example 8.16:

$$\gamma(t) = (t, f(t)).$$

Since f is continuous on $[0, 1]$, as is the identity function, it follows that γ is continuous on $[0, 1]$, and hence is a curve according to the above definition. Now, for each n , consider the partition (with points written in the reverse of the usual order) $\Pi_n = \{1 = t_0 > t_1 > \dots > t_{n-1} > t_n = 0\}$ where

for $0 < j < n$, $t_j = (\frac{\pi}{2} + (n-1)\pi)^{-1}$. So $\frac{1}{t_1} = \frac{\pi}{2}$, $\frac{1}{t_2} = \frac{3\pi}{2}$, and so forth through $\frac{1}{t_{n-1}} = \frac{(2n-3)\pi}{2}$. At all of these points, $\sin(\frac{1}{t_{j-1}}) = \pm 1$, with the sign changing from one term to the next. Hence, the lengths of adjacent increments (when neither j nor $j-1$ is 0 or n) are given by

$$|\gamma(t_j) - \gamma(t_{j-1})| = |(t_j, \pm t_j) - (t_{j-1}, \mp t_{j-1})| = \sqrt{(t_j - t_{j-1})^2 + (t_j + t_{j-1})^2} > t_j + t_{j-1} > t_j.$$

Now, $t_j = \frac{2}{(2j-1)\pi} > \frac{1}{4(j-1)}$ for $0 < j < n$, and so in the range $2 \leq j \leq n-1$ (where neither j nor $j-1$ is 0 or n), we have

$$\Lambda(\gamma, \Pi_n) > \sum_{j=2}^{n-1} |\gamma(t_j) - \gamma(t_{j-1})| > \frac{1}{4} \sum_{j=2}^{n-1} \frac{1}{j-1} = \frac{1}{4} \sum_{k=1}^{n-2} \frac{1}{k}.$$

Since $\sum_{k=1}^{\infty} \frac{1}{k} = +\infty$, it follows that $\Lambda(\gamma) = \sup_{\Pi} \Lambda(\gamma, \Pi) = +\infty$.

So the length of the curve traced out by the graph of our favorite pathological continuous function is infinite.

We are interested in curves whose length is finite; these are called *rectifiable*. (The act of *rectifying* a curve is to “unravel” it into a straight line without stretching.) There is a large class of curves that are rectifiable, and with (nominally) computable lengths: C^1 curves. To prove this, we first need to briefly extend the integral to curves (i.e. vector-valued functions of a real variable).

DEFINITION 9.34. Let $\mathbf{f}: [a, b] \rightarrow \mathbb{R}^d$ be a function $\mathbf{f} = (f_1, \dots, f_n)$ where each component f_j is Riemann integrable on $[a, b]$; we still denote this by $\mathbf{f} \in \mathcal{R}$. The integral $\int_a^b \mathbf{f}$ is the vector defined by componentwise integration:

$$\int_a^b \mathbf{f} \equiv \left(\int_a^b f_1, \dots, \int_a^b f_d \right).$$

As with derivatives, any theorem about integrals of scalar-valued functions extends immediately to vector-valued functions, so long as it applies separately to the components. For example: the Fundamental Theorem of Calculus still holds: if $\mathbf{f} \in \mathcal{R}$ and $\mathbf{F}' = \mathbf{f}$ (meaning that $\mathbf{F} = (F_1, \dots, F_d)$ is differentiable and $F'_j = f_j$ for each j), then $\int_a^b \mathbf{f} = \mathbf{F}(b) - \mathbf{F}(a)$. One result about integrals that does not obviously carry over to the vector-valued case is Lemma 9.23: $|\int f| \leq \int |f|$. If we try to apply this componentwise, we get

$$\left| \int \mathbf{f} \right| = \sqrt{\sum_{j=1}^n \left| \int f_j \right|^2} \leq \sqrt{\sum_{j=1}^n \left(\int |f| \right)^2}$$

but this is not related in any clear way to

$$\int |\mathbf{f}| = \int \sqrt{\sum_{j=1}^n |f_j|^2}.$$

Nonetheless, these two are comparable in precisely the same manner.

LEMMA 9.35. Let $\mathbf{f} \in \mathcal{R}$ on $[a, b]$. Then $|\mathbf{f}| \in \mathcal{R}$ as well, and

$$\left| \int_a^b \mathbf{f} \right| \leq \int_a^b |\mathbf{f}|.$$

PROOF. By Lemma 9.22, $f_j^2 \in \mathcal{R}$ for each j , and then by induction on Lemma 9.22, $|\mathbf{f}|^2 = f_1^2 + \cdots + f_d^2$ is in \mathcal{R} . Since $x \mapsto \sqrt{x}$ is continuous on $[0, \infty)$, it then follows from Lemma 9.20 that $|\mathbf{f}| = \sqrt{f_1^2 + \cdots + f_d^2}$ is in \mathcal{R} as claimed. To prove the inequality, we take a hint from the proof of the vector Mean Value inequality (Theorem 8.33: let $\mathbf{v} = \int \mathbf{f}$, and note that

$$|\mathbf{v}|^2 = \mathbf{v} \cdot \mathbf{v} = \mathbf{v} \cdot \int \mathbf{f} = \sum_{j=1}^d v_j \int f_j = \int \left(\sum_{j=1}^d v_j f_j \right) = \int \mathbf{v} \cdot \mathbf{f}$$

where we used Lemma 9.16 in the penultimate equality. Applying the Cauchy-Schwarz inequality, we have $\mathbf{v} \cdot \mathbf{f} \leq |\mathbf{v}||\mathbf{f}|$. Since we know that $|\mathbf{f}| \in \mathcal{R}$, it then follows from Lemma 9.17 that $\int \mathbf{v} \cdot \mathbf{f} \leq \int |\mathbf{v}||\mathbf{f}| = |\mathbf{v}| \int |\mathbf{f}|$. Thus $|\mathbf{v}|^2 \leq |\mathbf{v}| \int |\mathbf{f}|$; so either $\mathbf{v} = \mathbf{0}$ (in which case $|\mathbf{v}| = 0 \leq \int |\mathbf{f}|$) or we can cancel one $|\mathbf{v}| > 0$ to find $|\mathbf{v}| \leq \int |\mathbf{f}|$, which is the desired inequality. \square

REMARK 9.36. We can define $\int \mathbf{f} d\alpha$ for general increasing α in the same manner, and the above proof shows that the inequality of Lemma 9.35 holds in this general setting; but we will not have occasion to use it beyond the standard Riemann integral case.

With these facts about integration of vector-valued functions at hand, we can now prove that C^1 curves are rectifiable.

THEOREM 9.37. *Let γ be a C^1 curve on $[a, b]$, meaning γ' is continuous on $[a, b]$. Then γ is rectifiable, and*

$$\Lambda(\gamma) = \int_a^b |\gamma'(t)| dt.$$

PROOF. Let $\Pi = \{a = t_0 < t_1 < \cdots < t_n = b\}$ be a partition. Applying the Fundamental Theorem of Calculus, we have

$$\gamma(t_j) - \gamma(t_{j-1}) = \int_{t_{j-1}}^{t_j} \gamma'(t) dt$$

and thus, applying Lemma 9.35 we have

$$\Lambda(\gamma, \Pi) = \sum_{j=1}^n |\gamma(t_j) - \gamma(t_{j-1})| = \sum_{j=1}^n \left| \int_{t_{j-1}}^{t_j} \gamma'(t) dt \right| \leq \sum_{j=1}^n \int_{t_{j-1}}^{t_j} |\gamma'(t)| dt = \int_a^b |\gamma'(t)| dt.$$

(The last equality follows from Lemma 9.18, collapsing the telescoping sum.) Hence, $\int_a^b |\gamma'(t)| dt$ is an upper bound for $\Lambda(\gamma, \Pi)$ over all Π , and thus $\Lambda(\gamma) = \sup_{\Pi} \Lambda(\gamma, \Pi) \leq \int_a^b |\gamma'(t)| dt$. We are left only to prove the reverse inequality.

Fix $\epsilon > 0$. Since γ' is continuous on the compact interval $[a, b]$, it is uniformly continuous, and so there is some $\delta > 0$ so that $|\gamma'(s) - \gamma'(t)| < \frac{\epsilon}{2(b-a)}$ whenever $|s - t| < \delta$. Let Π be any partition with $\Delta t_j < \delta$ for each j . Then for any $t \in [t_{j-1}, t_j]$, since $|t - t_j| < \delta$, it follows that $||\gamma'(t)| - |\gamma'(t_j)|| < \frac{\epsilon}{2(b-a)}$; in particular, $|\gamma'(t)| \leq |\gamma'(t_j)| + \frac{\epsilon}{2(b-a)}$. Thus

$$\int_{t_{j-1}}^{t_j} |\gamma'(t)| dt \leq \int_{t_{j-1}}^{t_j} \left(|\gamma'(t_j)| + \frac{\epsilon}{2(b-a)} \right) dt = (|\gamma'(t_j)| + \frac{\epsilon}{2(b-a)}) \Delta t_j. \quad (9.9)$$

For the first term here, we make the following clever estimate:

$$\begin{aligned} |\gamma'(t_j)|\Delta t_j &= \left| \int_{t_{j-1}}^{t_j} \gamma'(t_j) dt \right| = \left| \int_{t_{j-1}}^{t_j} [\gamma'(t) + \gamma'(t_j) - \gamma'(t)] dt \right| \\ &\leq \left| \int_{t_{j-1}}^{t_j} \gamma'(t) dt \right| + \left| \int_{t_{j-1}}^{t_j} [\gamma'(t_j) - \gamma'(t)] dt \right| \\ &\leq |\gamma(t_j) - \gamma(t_{j-1})| + \frac{\epsilon}{2(b-a)}\Delta t_j \end{aligned}$$

where we've applied the Fundamental Theorem of Calculus again to the first term, and in the second term used Lemma 9.35 and then the fact that $|\gamma'(t_j) - \gamma'(t)| < \frac{\epsilon}{2(b-a)}$ again. Combining this with (9.9) and summing (using Lemma 9.18) gives

$$\begin{aligned} \int_a^b |\gamma'(t)| dt &= \sum_{j=1}^n \int_{t_{j-1}}^{t_j} |\gamma'(t)| dt \leq \sum_{j=1}^n \left(|\gamma(t_j) - \gamma(t_{j-1})| + 2\frac{\epsilon}{2(b-a)}\Delta t_j \right) = \Lambda(\gamma, \Pi) + \epsilon \\ &\leq \Lambda(\gamma) + \epsilon. \end{aligned}$$

Since this holds true for all $\epsilon > 0$, it follows that $\int_a^b |\gamma'(t)| dt \leq \Lambda(\gamma)$, as desired. \square

REMARK 9.38. In the above proof, the continuity of γ' was not needed to show that γ is rectifiable; indeed, this holds whenever $\gamma' \in \mathcal{R}$, and the inequality $\Lambda(\gamma) \leq \int_a^b |\gamma'(t)| dt$ holds true. But this inequality may be strict if γ' is Riemann integrable but not continuous.

Sequences and Series of Functions

1. Lecture 13: May 10, 2016

Although we have been tacitly using many of the familiar functions from calculus (like \cos , \sin , \exp) and assuming regularity properties of them, we have yet to formally define and develop these functions rigorously. To do so requires an understand of limits of sequences and series of functions, to which this chapter is devoted.

For each $n \in \mathbb{N}$, suppose f_n be a real- or complex-valued function, defined on some set X which does not vary with n ; for us X will usually be an interval in \mathbb{R} . We can then talk about the *pointwise* limit of these functions (should it exist): for each $x \in X$, we consider the sequence $(f_n(x))_{n=1}^{\infty}$, which is a sequence in \mathbb{R} or \mathbb{C} . (Indeed, we could consider more general cases where the f_n take values in some common metric space Y .)

DEFINITION 10.1. *Let (f_n) be a sequence of functions as described above. If the sequences $(f_n(x))$ converge for each $x \in X$, we define a new function $f(x) \equiv \lim_{n \rightarrow \infty} f_n(x)$. The function f is called the pointwise limit of f_n .*

The questions that will be important to us concern whether common properties of the functions f_n carry over to the limit function f , should it exist; e.g. continuity, differentiability, integrability, etc. In all cases, the answer will be a resounding *no* without further assumptions. Let's consider some examples.

EXAMPLE 10.2. For $x \in \mathbb{R}$ and $n \in \mathbb{N}$, define

$$f_n(x) = \frac{1}{(1+x^2)^n}.$$

This is a rational function, and the denominator vanishes nowhere, so f_n is C^∞ on \mathbb{R} . Note that $f_n(0) = 1$ for all n , so $\lim_{n \rightarrow \infty} f_n(0) = 1$. But for $x \neq 0$, $1+x^2 > 1$, and so $\frac{1}{(1+x^2)^n} \rightarrow 0$. Thus

$$f(x) = \lim_{n \rightarrow \infty} f_n(x) = \begin{cases} 1 & x = 0 \\ 0 & x \neq 0 \end{cases}$$

Thus, (f_n) has a pointwise limit, but even though the functions f_n are all smooth, the limit function is not even continuous.

EXAMPLE 10.3. Taking for granted that \sin and \cos are differentiable functions with $\sin' = \cos$, define functions f_n on \mathbb{R} by

$$f_n(x) = \frac{\sin(nx)}{\sqrt{n}}.$$

These functions are smooth, and since \sin is bounded, $\lim_{n \rightarrow \infty} f_n(x) = 0$ for all x , so the limit function is also smooth. But note that

$$f'_n(x) = \sqrt{n} \cos(nx).$$




FIGURE 1. The function $f_n(x) = \frac{1}{(1+x^2)^n}$, for $n = 1, 5, 10, 50$.

The behavior of these functions is not nearly as nice. For example, $f'_n(0) = \sqrt{n} \rightarrow +\infty$ as $n \rightarrow \infty$. Even worse, $f'_n(\pi) = \sqrt{n}(-1)^n$ oscillates without bound. Evaluating at other points can produce even worse behavior. So, while f_n is differentiable and $f_n \rightarrow 0$ pointwise, the derivatives f'_n do not converge pointwise to anything.




FIGURE 2. The function $f_n(x) = \frac{\sin(nx)}{\sqrt{n}}$, for $n = 1, 5, 10, 50$.

EXAMPLE 10.4. Let $f_n(x) = nx(1-x^2)^n$ for $0 \leq x \leq 1$. Note that $f_n(0) = 0$ for all n . For $0 < x \leq 1$, the sequence $x(1-x^2)^n$ converges to 0 exponentially fast, and therefore even

multiplying by n , we still have $f_n(x) \rightarrow 0$ as $n \rightarrow \infty$ (cf. Proposition 2.30(4)). Thus, the pointwise limit of (f_n) is the constant function 0 on $[0, 1]$. Nevertheless, we may easily compute that

$$\int_0^1 f_n(x) dx = n \int_0^1 x(1-x^2)^n dx = n \int_0^1 \frac{1}{2} u^n du = \frac{n}{2(n+1)}$$

where we made the change of variables $u = 1 - x^2$. Thus

$$0 = \int_0^1 \lim_{n \rightarrow \infty} f_n(x) dx \neq \lim_{n \rightarrow \infty} \int_0^1 f_n(x) dx = \frac{1}{2}.$$

So although the functions f_n are all integrable, as is their pointwise limit function, the limit of the integrals is not the integral of the limit. Using similar arguments, note that $g_n(x) = n^2 f_n(x)$ still satisfies $\lim_{n \rightarrow \infty} g_n(x) = 0$ for all x , and yet following the above we have $\int_0^1 g_n(x) dx = \frac{n^2}{2(n+1)} \rightarrow +\infty$ as $n \rightarrow \infty$; so we can even have the total area blowing up, even as the function converges to 0.




FIGURE 3. The function $f_n(x) = nx(1-x^2)^n$, for $n = 1, 5, 10, 100, 500$.

What's going on in all these examples? To understand at a lower level, let's consider the question of continuity of the limit function: when is it true for continuous function f_n with pointwise limit function f that $\lim_{t \rightarrow x} f(t) = f(x)$ for each x ? Writing this out further, since $\lim_{t \rightarrow x} f_n(t) = f_n(x)$, we are asking that

$$\lim_{t \rightarrow x} \lim_{n \rightarrow \infty} f_n(t) = \lim_{t \rightarrow x} f(t) = f(x) = \lim_{n \rightarrow \infty} f_n(x) = \lim_{n \rightarrow \infty} \lim_{t \rightarrow x} f_n(t).$$

I.e. we are asking whether we can *interchange the limit*. One might naïvely expect to be able to change the order of limits at will, but this is not so. Indeed, a limit is characterized by a quantified sentence “for all $\epsilon > 0$, there exists a $\delta > 0$ or $N \in \mathbb{N} \dots$ ”. Thus, a double limit involves two such sentences in order, and we know that changing the order of quantifiers can really change the meaning of the sentence. Indeed, we can easily exhibit examples of double limits (even for sequences of real numbers) where the order matters big time.

EXAMPLE 10.5. Consider the double sequence $a_{m,n} = \frac{m}{m+n}$, for $m, n \geq 1$.

For fixed m , $\lim_{n \rightarrow \infty} a_{m,n} = \lim_{n \rightarrow \infty} \frac{m}{m+n} = 0$. For fixed n , $\lim_{m \rightarrow \infty} a_{m,n} = \lim_{m \rightarrow \infty} \frac{m}{m+n} = 1$.

Thus

$$0 = \lim_{m \rightarrow \infty} \lim_{n \rightarrow \infty} a_{m,n} \neq \lim_{n \rightarrow \infty} \lim_{m \rightarrow \infty} a_{m,n} = 1.$$

It is actually not surprising that regularity properties of functions do not generally survive under pointwise limit. If (f_n) is a sequence of functions, and $f(x) = \lim_{n \rightarrow \infty} f_n(x)$ for each x , the limit has no way of “seeing” any properties (local or global) of the functions; it is only paying attention to each x separately, decoupled from the others. If we want the limit function f to reflect any properties of the functions f_n , we will need a stronger form of convergence. In fact, we already introduced it in Example 5.2(4).

DEFINITION 10.6. Let (f_n) be a sequence of real-valued functions defined on a common set E . We say that (f_n) converges to f uniformly, $f_n \rightarrow_u f$, iff $\sup_{x \in E} |f_n(x) - f(x)|$ converges to 0 as $n \rightarrow \infty$. I.e. $f_n \rightarrow_u f$ iff for each $\epsilon > 0$, there is an $N \in \mathbb{N}$ so that, for all $n \geq N$ and all $x \in E$, $|f_n(x) - f(x)| < \epsilon$.

Like comparing continuity to uniform continuity, uniform convergence differs from pointwise convergence by a reordering of quantifiers: here we must be able to choose N independently of x , so we have $\forall \epsilon > 0 \exists N \forall n \geq N \forall x$, as opposed to the pointwise definition which begins with $\forall x \forall \epsilon > 0 \exists N \forall n \geq N$.

REMARK 10.7. If the functions (f_n) in the sequence are all bounded, then we are talking about metric convergence here: on the space $B(E)$ of bounded real-valued functions defined on E , $d_u(f, g) = \sup_{x \in E} |f(x) - g(x)|$ is a metric, as discussed in Example 5.2(4); then $f_n \rightarrow_u f$ simply means $d_u(f_n, f) \rightarrow 0$, meaning that $f_n \rightarrow f$ in the metric space. However, functions f_n need not be bounded in order to converge uniformly. For example, $f_n(x) = x + \frac{1}{n}$ converges uniformly to $f(x) = x$, since $f_n(x) - f(x) = \frac{1}{n}$ which converges to 0 uniformly in x . But $d_u(f_n, 0) = \infty$ for all n , so d_u is not a well-defined metric on any space containing the functions f_n (and 0).

Motivated by the above remark, the following result essentially says that d_u is a Cauchy-complete metric.

PROPOSITION 10.8. A sequence (f_n) of real-valued functions on E converges uniformly to some function if and only if the sequence is uniformly-Cauchy: for each $\epsilon > 0$, there is an $N \in \mathbb{N}$ so that, for all $n, m \geq N$ and for all $x \in E$, $|f_n(x) - f_m(x)| < \epsilon$.

PROOF. First suppose we know there exists a function f on E with $f_n \rightarrow_u f$. Choose N so that, for all $n \geq N$ and all $x \in E$, $|f_n(x) - f(x)| < \frac{\epsilon}{2}$. Then for $n, m \geq N$, and all $x \in E$,

$$\begin{aligned} |f_n(x) - f_m(x)| &= |f_n(x) - f(x) + f(x) - f_m(x)| \\ &\leq |f_n(x) - f(x)| + |f_m(x) - f(x)| < \frac{\epsilon}{2} + \frac{\epsilon}{2} = \epsilon. \end{aligned}$$

Hence, (f_n) is uniformly Cauchy, as claimed. For the converse, suppose (f_n) is uniformly Cauchy. Fix any $t \in E$. Then $|f_n(t) - f_m(t)| \leq \sup_{x \in E} |f_n(x) - f_m(x)|$, and since the latter is Cauchy, it follows that $(f_n(t))$ is a Cauchy sequence in \mathbb{R} . Since \mathbb{R} is Cauchy complete, it follows that there is a real number $f(t)$ such that $f_n(t) \rightarrow f(t)$. Thus we have found a pointwise limit for (f_n) . We must prove that $f_n \rightarrow_u f$. To that end, fix $\epsilon > 0$, and let N be chosen so that $|f_n(x) - f_m(x)| < \frac{\epsilon}{2}$

for all $n, m \geq N$. Since $f_m(x) \rightarrow f(x)$, there is some N_x so that $|f_m(x) - f(x)| < \frac{\epsilon}{2}$ for all $m \geq N_x$. Thus, taking $n \geq N$ and $m \geq \max\{N, N_x\}$, we have

$$|f_n(x) - f(x)| \leq |f_n(x) - f_m(x)| + |f_m(x) - f(x)| < \frac{\epsilon}{2} + \frac{\epsilon}{2} = \epsilon.$$

Although the set of allowed m depends on x , this makes no difference to the final statement: we see that when $n \geq N$ (which is independent of x) $|f_n(x) - f(x)| < \epsilon$. Thus $f_n \rightarrow_u f$. \square

2. Lecture 14: May 12, 2016

In Example 10.2, we saw that a sequence of continuous (even smooth) functions can have a pointwise limit that is discontinuous. That is: pointwise limits of functions do not preserve limits in general. The following theorem shows that uniform convergence *does* preserve limits, and hence continuity.

THEOREM 10.9. *Let E be a subset of a metric space, and let $f_n \rightarrow_u f$ on E . Let $x \in E'$, and suppose that the following limits exist for all n :*

$$\lim_{t \rightarrow x} f_n(t) = a_n.$$

then the sequence (a_n) converges to $\lim_{t \rightarrow x} f(t)$. In particular, if f_n are all continuous at a point x , then their uniform limit f is also continuous at x .

PROOF. Fix $\epsilon > 0$, and choose N so that $|f_n(t) - f_m(t)| < \frac{\epsilon}{2}$ for all $t \in E$ and all $n, m \geq N$. By the Squeeze Theorem, it follows that $|a_n - a_m| \leq \frac{\epsilon}{2} < \epsilon$ for all $n, m \geq N$. Thus, we have proved that (a_n) is a Cauchy sequence in \mathbb{R} , and therefore converges to some limit a . Then we have, for each $t \in E$,

$$|f(t) - a| \leq |f(t) - f_n(t)| + |f_n(t) - a_n| + |a_n - a|.$$

Since $f_n \rightarrow_u f$, we can choose N_1 such that $|f(t) - f_n(t)| < \frac{\epsilon}{3}$ for all $n \geq N_1$ and $t \in E$. As $a_n \rightarrow a$, we can choose N_2 so that $|a_n - a| < \frac{\epsilon}{3}$ for all $n \geq N_2$. Now, for any given $n \geq \max\{N_1, N_2\}$, since $\lim_{t \rightarrow x} f_n(t) = a_n$, we can choose some $\delta > 0$ so that for all $t \in B_\delta(x) \setminus \{x\}$, $|f_n(t) - a_n| < \frac{\epsilon}{3}$. Thus, for this choice of δ , we see that $|f(t) - a| < \frac{\epsilon}{3} + \frac{\epsilon}{3} + \frac{\epsilon}{3} = \epsilon$ when $0 < |t - x| < \delta$, and we have shown that $\lim_{t \rightarrow x} f(t) = a$ as claimed.

For the final statement: if f_n is continuous at x , then $a_n = \lim_{t \rightarrow x} f_n(t) = f_n(x)$. The conclusion is then that $f_n(x)$ converges to $\lim_{t \rightarrow x} f(t)$; but by definition $f_n(x) \rightarrow f(x)$, so $f(x) = \lim_{t \rightarrow x} f(t)$, and ergo f is continuous at x . \square

We can now put together Proposition 10.8 and Theorem 10.9, with Example 5.2(4), to yield the following.

COROLLARY 10.10. *Let X be a metric space, and let $C_b(X)$ denote the set of real-valued bounded continuous functions on X . Then $d_u(f, g) = \sup_{x \in X} |f(x) - g(x)|$ is a complete metric on $C_b(X)$.*

PROOF. We know that d_u is a metric on the set of bounded functions, and therefore on the subset $C_b(X)$. Let (f_n) be a Cauchy sequence in this metric space. By Proposition 10.8, (f_n) converges uniformly to some function f . Since the f_n are all continuous, so is f by Theorem 10.9. Moreover, since $f_n \rightarrow_u f$, there is some N so that $\sup_x |f_N(x) - f(x)| < 1$, and so $|f(x)| < |f_N(x)| + 1$ for all $x \in X$. But f_N is bounded by, say, M , and so $|f(x)| \leq M + 1$, so f is also bounded. Thus there is a limit f of (f_n) in $C_b(X)$, proving that this is a complete metric space. \square

REMARK 10.11. The set $C_b(X)$ plays a central role in analysis and probability theory. It is also worth noting that it is a vector space: if $f, g \in C_b(X)$ then so is $f + g$, and $\lambda f \in C_b(X)$ for all $\lambda \in \mathbb{R}$. The metric $d_u(f, g) = \sup_x |f(x) - g(x)|$ comes from a *norm* on this vector space: the uniform norm $\|f\|_u = \sup_x |f(x)|$; then $d_u(f, g) = \|f - g\|_u$. (Another common notation for the uniform norm is $\|f\|_\infty$.) In addition to the properties of a metric, the norm also has homogeneity: $\|\lambda f\|_u = |\lambda| \|f\|_u$. We will discuss this and other norms on functions further in later lectures.

Let us examine again Example 10.2, where a sequence $f_n(x) = \frac{1}{(1+x^2)^n}$ of continuous functions converges pointwise to a discontinuous function $f = \mathbb{1}_{\{0\}}$. By Theorem 10.9, it must be true that the convergence is non-uniform. Indeed, let $\epsilon > 0$. For $x \neq 0$, we want to find the smallest N_x so that $|f_n(x) - f(x)| = f_n(x) < \epsilon$. So we want

$$\frac{1}{(1+x^2)^n} < \epsilon \iff (1+x^2)^n > \frac{1}{\epsilon} \iff n > \frac{\ln(1/\epsilon)}{\ln(1+x^2)}.$$

Hence, the smallest N_x is

$$N_x = \left\lceil \frac{\ln(1/\epsilon)}{\ln(1+x^2)} \right\rceil.$$

As x approaches 0, $1+x^2 \rightarrow 1$, and $\ln(1+x^2) \rightarrow 0$; hence $N_x \rightarrow \infty$ as $x \rightarrow 0$. Therefore, there can be no uniform N that works for all x . We see directly that (f_n) does not converge uniformly.

On the other hand, it is entirely possible for a sequence of functions (f_n) to converge pointwise but *non-uniformly* to a continuous function. Example 10.4 demonstrates this: the functions $f_n(x) = nx(1-x^2)^n$ tend to 0 for each x , so the limit is continuous; however, using calculus, we can easily check that the only critical points of the polynomial f_n are $x = \pm \frac{1}{\sqrt{2n+1}}$, and that furthermore the maximum value of f_n on $[0, 1]$ is achieved at the critical point in that interval, where the value is

$$f_n\left(\frac{1}{\sqrt{2n+1}}\right) = \frac{n}{\sqrt{2n+1}} \left(\frac{2n}{2n+1}\right)^n.$$

Now

$$\left(\frac{2n}{2n+1}\right)^n = \left(1 - \frac{1}{2n+1}\right)^n \rightarrow e^{-1/2}$$

and $\frac{n}{\sqrt{2n+1}} \rightarrow \infty$, so we see that $\sup_{0 \leq x \leq 1} f_n(x) \rightarrow \infty$ as $n \rightarrow \infty$ (which is suggested by Figure 3). Thus, $\sup_x |f_n(x) - 0| \rightarrow \infty$, which means that $f_n \not\rightarrow_u 0$, even though the limit 0 turns out to be continuous. We can see this phenomenon with a simpler example.

EXAMPLE 10.12. Let $f_n: (0, 1) \rightarrow \mathbb{R}$ be given by $f_n(x) = \frac{1}{nx+1}$. For each n , f_n is a continuous function (it is continuous on \mathbb{R} except at $-1/n$ which is not in $(0, 1)$). For each fixed x , $f_n(x) \rightarrow 0$ as $n \rightarrow \infty$, so the pointwise limit of these continuous functions is, again, continuous. Nevertheless, note that $f_n(1/n) = \frac{1}{2}$ for all n , so $\sup_{x \in (0,1)} |f_n(x) - 0| \geq \frac{1}{2}$; in fact, f_n extends continuously to $[0, 1]$, and on this interval the maximum is $f_n(0) = 1$, so $\sup_{x \in (0,1)} f_n(x) = 1$. In any case, this means that $f_n \not\rightarrow_u 0$ uniformly on $(0, 1)$.

There is one situation where the converse of Theorem 10.9 does hold: that pointwise convergence of a sequence of continuous functions to a continuous function implies uniform convergence. This next result is ‘‘Dini’s monotone convergence theorem’’.

THEOREM 10.13 (Dini’s Monotone Convergence Theorem). *Let K be a compact set, and let (f_n) be a sequence of real-valued continuous functions on K . Suppose that the pointwise limit $f = \lim_{n \rightarrow \infty} f_n$ is continuous, and furthermore that $f_n(x) \geq f_{n+1}(x)$ for all $x \in K$. Then $f_n \rightarrow_u f$ on K .*

Note: since $f_n \rightarrow_u f$ iff $-f_n \rightarrow_u -f_n$, we could instead assume that $f_n \leq f_{n+1}$ and get the same conclusion.

PROOF. Set $g_n = f_n - f$. Then $g_n \rightarrow 0$ pointwise, g_n is continuous, and $g_n \geq g_{n+1}$. Our goal is to prove that $g_n \rightarrow_u 0$. Fix $\epsilon > 0$, and let $K_n \subset K$ be the set of points $K_n \ni x$ where $g_n(x) \geq \epsilon$; i.e. $K_n = g_n^{-1}([\epsilon, \infty))$. Since g_n is continuous and $[\epsilon, \infty)$ is closed, it follows that K_n




FIGURE 4. The function $f_n(x) = \frac{1}{nx+1}$ on $x \in (0, 1)$, for $n = 1, 2, 5, 10, 100$.

is a closed subset of the compact set K , so K_n is compact. Now, since $g_n \geq g_{n+1}$, if $g_{n+1}(x) \geq \epsilon$, then $g_n(x) \geq \epsilon$, meaning that $K_{n+1} \subseteq K_n$. So (K_n) is a nested sequence of compact sets. If they were all nonempty, Proposition 5.33 would imply that $\bigcap_n K_n$ is nonempty. However, given any fixed $x \in K$, we know that $g_n(x) \rightarrow 0$, which means that $g_n(x) < \epsilon$ for all large n , which means that $x \notin K_n$ for all large n . Thus, $\bigcap_n K_n = \emptyset$, and so it follows that the $K_n = \emptyset$ for some n ; since they are nested, this means there is some N so that $K_n = \emptyset$ for all $n \geq N$. This means that, for all $x \in K$, $g_n(x) < \epsilon$ for all $n \geq N$; note that this N is uniform in x . As $g_n(x)$ is a decreasing sequence with limit 0, it follows that $g_n(x) \geq 0$, and so we have $|g_n(x) - 0| < \epsilon$ for all $n \geq N$ and $x \in K$. This says precisely that $g_n \rightarrow_u 0$. \square

REMARK 10.14. Note that the sequence $f_n(x) = \frac{1}{nx+1}$ of Example 10.12 is actually a decreasing sequence of continuous functions, with a continuous limit, and yet the convergence is not uniform; this highlights the necessity of the compactness of the domain. Indeed, in that example, the functions f_n all extend continuously to the compact set $K = [0, 1]$, but the pointwise limit function on that domain is $\lim_{n \rightarrow \infty} f_n = \mathbb{1}_{\{0\}}$ (since $f_n(0) = 1$ for all n), which is not continuous.

Now we have a good idea of when continuity passes to the limit, and why Example 10.2 failed to transfer continuity to its limit: in general, one should have uniform convergence to pass continuity to the limit. Let us now reconsider Example 10.4, $f_n(x) = nx(1-x^2)^n$, which converges pointwise to 0 on $[0, 1]$, a perfectly integrable function, but $\int_0^1 f_n \not\rightarrow \int_0^1 0$. As we saw above, the convergence isn't uniform. In fact, if we don't have uniform convergence, it can be worse still: it's possible for the pointwise limit of a sequence of integrable functions to be non-integrable.

EXAMPLE 10.15. The rational numbers \mathbb{Q} are countable; in particular, it is possible to enumerate those rational numbers in $[0, 1]$: $\mathbb{Q} \cap [0, 1] = \{q_1, q_2, q_3, \dots\}$. Consider the function $f_n = \sum_{j=1}^n \mathbb{1}_{\{q_j\}}$. That is:

$$f_n(x) = \begin{cases} 1, & x \in \{q_1, q_2, \dots, q_n\} \\ 0, & \text{otherwise} \end{cases}$$

The function f_n is continuous at all points other than $\{q_1, q_2, \dots, q_n\}$, which is a finite set; therefore, by Theorem 9.12, $f_n \in \mathcal{R}$. Also, for any $\epsilon > 0$, build the partition Π_n containing the points 0, 1, and all points $q_j \pm \frac{\epsilon}{2n}$ (possibly shrinking smaller if any q_j is closer to 0 or 1 than $\frac{\epsilon}{2n}$). Then on the n intervals $[q_j - \frac{\epsilon}{2n}, q_j + \frac{\epsilon}{2n}]$, the supremum of f_n is 1 and the infimum is 0; on all the other intervals, f_n is constantly 0, and so we have

$$U(f_n, \Pi_n) = n \cdot \left(\frac{\epsilon}{2n} - \frac{-\epsilon}{2n} \right) = \epsilon, \quad L(f_n, \Pi_n) = 0.$$

Since we can do this for any $\epsilon > 0$, it follows that $U(f_n) = L(f_n) = 0$, so $\int_0^1 f_n = 0$ for each n .

However, notice that $\lim_{n \rightarrow \infty} f_n = \mathbb{1}_{\mathbb{Q} \cap [0, 1]}$ is Dirichlet's function of Example 6.11. Indeed, for any rational number $q \in \mathbb{Q} \cap [0, 1]$, there is some $m \in \mathbb{N}$ with $q = q_m$, and so $f_n(q) = 1$ for all $n \geq m$; on the other hand, if x is irrational, $f_n(x) = 0$ for all n . As we showed in Example 9.6, the limit function $\lim_{n \rightarrow \infty} f_n$ is not Riemann integrable, even though all the functions f_n are, with integral 0. So this is not a matter of “mass escaping to infinity” as in Example 10.4; pointwise limits simply do not preserve integrability in general.

3. Lecture 15: May 17, 2016

As with continuity, to avoid the kinds of pathologies in Example 10.15, in general we require uniform convergence for integrals to play nicely with limits.

THEOREM 10.16. *Let $a < b$ in \mathbb{R} , let $\alpha: [a, b] \rightarrow \mathbb{R}$ be monotone increasing, and suppose $f_n \in \mathcal{R}(\alpha)$ for each n . If $f_n \rightarrow_u f$ on $[a, b]$, then $f \in \mathcal{R}(\alpha)$, and*

$$\lim_{n \rightarrow \infty} \int_a^b |f_n - f| d\alpha = 0.$$

In particular, it follows that

$$\lim_{n \rightarrow \infty} \int_a^b f_n d\alpha = \int_a^b f d\alpha.$$

PROOF. For each n , let $\epsilon_n = \sup_{x \in [a, b]} |f_n(x) - f(x)|$; by assumption, $\epsilon_n \rightarrow 0$ as $n \rightarrow \infty$. Note that

$$f_n(x) - \epsilon_n \leq f(x) \leq f_n(x) + \epsilon_n, \quad \forall x \in [a, b].$$

Since f_1 is bounded, this means f is bounded, so its upper and lower sums are well-defined. Given any partition Π , taking sup and inf of all terms on each subinterval and multiplying by $\Delta\alpha_j$, we get

$$L(f_n - \epsilon_n, \Pi, \alpha) \leq L(f, \Pi, \alpha) \leq U(f, \Pi, \alpha) \leq U(f_n + \epsilon_n, \Pi, \alpha).$$

Also, since $L(g, \Pi, \alpha)$ is clearly linear in g , using the usual telescoping sum we find that Taking \sup_{Π} in the first inequality, \inf_{Π} in the third inequality, and using the fact that $L(f, \alpha) \leq U(f, \alpha)$, we have

$$L(f_n - \epsilon_n, \alpha) \leq L(f, \alpha) \leq U(f, \alpha) \leq U(f_n + \epsilon_n, \alpha). \quad (10.1)$$

Since $f_n \pm \epsilon_n \in \mathcal{R}(\alpha)$, the upper and lower integrals are equal, and by linearity of the integral we have

$$\int_a^b (f_n \pm \epsilon_n) d\alpha = \int_a^b f_n d\alpha \pm \epsilon_n[\alpha(b) - \alpha(a)]. \quad (10.2)$$

Putting (10.1) and (10.2) together, we see that $L(f, \alpha)$ and $U(f, \alpha)$ are both in the interval

$$\left[\int_a^b f_n d\alpha - \epsilon_n[\alpha(b) - \alpha(a)], \int_a^b f_n d\alpha + \epsilon_n[\alpha(b) - \alpha(a)] \right]. \quad (10.3)$$

In particular, it follows that the difference between these two is \leq the length of the interval, so

$$0 \leq U(f, \alpha) - L(f, \alpha) \leq 2\epsilon_n[\alpha(b) - \alpha(a)].$$

This is true for all n , and $\epsilon_n \rightarrow 0$ as $n \rightarrow \infty$, so it follows that $U(f, \alpha) - L(f, \alpha) = 0$, i.e. $f \in \mathcal{R}(\alpha)$. Now, since $\int_a^b f d\alpha = U(f, \alpha)$ is in the interval in (10.3), we have

$$\left| \int_a^b f d\alpha - \int_a^b f_n d\alpha \right| \leq \epsilon_n[\alpha(b) - \alpha(a)].$$

Since $\epsilon_n \rightarrow 0$, it follows that $\int_a^b f_n d\alpha \rightarrow \int_a^b f d\alpha$ as claimed. But we also easily see the stronger claim: since $|f_n(x) - f(x)| \leq \epsilon_n$ for all x ,

$$\int_a^b |f_n - f| d\alpha \leq \int_a^b \epsilon_n d\alpha = \epsilon_n[\alpha(b) - \alpha(a)] \rightarrow 0.$$

This concludes the proof. □

REMARK 10.17. Referring to Homework 7, this means that the metric distance $d_1(f_n, f) \rightarrow 0$, provided that $f_n \rightarrow_u f$; i.e. uniform convergence is stronger than \mathcal{R}^1 convergence.

As we noted before proving Dini's monotone convergence theorem, continuity can be preserved under limits even without uniform convergence. Similarly, it is possible for the integral to be preserved under pointwise limits even if uniform convergence fails. In general, uniform convergence is the best tool to prove that the integral interchanges with limits, but there are some tools that work in a more general framework. We state a result here that is due to Arzelà, originally proved in 1885.

THEOREM 10.18 (Arzelà's Montone and Dominated Convergence Theorems). *Let (f_n) be a sequence of functions in \mathcal{R} on $[a, b]$, and suppose that there is a pointwise limit $f_n \rightarrow f$ such that $f \in \mathcal{R}$ on $[a, b]$. Suppose additionally that either*

- (a) (f_n) is monotone: $(f_n(x))_{n=1}^\infty$ is either monotone increasing for all x or monotone decreasing for all x ; or
- (b) (f_n) is uniformly bounded: there is some constant M so that $|f_n(x)| \leq M$ for all $x \in [a, b]$.

Then $\int_a^b |f_n - f| d\alpha \rightarrow 0$ as $n \rightarrow \infty$, and in particular $\int_a^b f_n d\alpha \rightarrow \int_a^b f d\alpha$.

Theorem 10.18 can be a powerful tool for analyzing limits of integrals, but it is important to note that you *must know that the pointwise limit function is integrable a priori*. As Example 10.15 shows, this can fail, even when the sequence f_n is monotone or uniformly bounded (both of which are true for that example). Theorem 10.16, with stronger conditions, gives the integrability of the limit function for free, and so we will rely on this in most cases (when uniform convergence actually holds).

We will not prove Theorem 10.18 presently. Although it is possible to give a proof based only on what we have developed so far (essentially a complicated iterative argument using Dini's monotone convergence theorem), it is quite intricate and tricky. More's the point, this theorem really belongs in the domain of measure theory: in the theory of the stronger Lebesgue integral, both the statement and the proof of the theorem are much easier (the Lebesgue integrability of the limit function comes for free, and the proof is only a few lines long once one has the power of measure theory at hand).

Now that we have seen that continuity and integration behave well under uniform limits, we turn our attention to derivatives. However, let us consider again Example 10.3, $f_n(x) = \frac{\sin(nx)}{\sqrt{n}}$. We saw that this sequence tends to 0, and in fact it tends to 0 *uniformly*: $|f_n(x)| \leq \frac{1}{\sqrt{n}}$, so we can make $|f_n(x)| < \epsilon$ for all x by choosing n (uniformly in x) bigger than $\frac{1}{\epsilon^2}$. But, as we saw in that example, the derivative f'_n is very badly behaved, not having a pointwise limit as $n \rightarrow \infty$. So, it will certainly not be true that $f_n \rightarrow_u f$ tells us much about the limit of f'_n , or even whether the limit exists.

In fact, more or less the best we can do in general is to cheat and do it *backwards*: if we already know that f'_n converges uniformly, then it nearly follows that f_n converges uniformly to some differentiable f , and then f' is the uniform limit of the f'_n . Even this is not true as stated, though: for example you could have $f_n = n$ is constant for each n ; then $f'_n = 0$ certainly converges uniformly to 0, but f_n does not converge at all. If we rule out this kind of uniform blow up by insisting that $f_n(x_0)$ converges at least at one point x_0 , we do get a true theorem.

THEOREM 10.19. *Let $f_n \in C^1[a, b]$ be a sequence of continuously differentiable functions. Suppose there is at least one point $x_0 \in [a, b]$ where $f_n(x_0)$ converges. If $f'_n \rightarrow_u g$ on $[a, b]$, then there is a C^1 function f on $[a, b]$ with $f_n \rightarrow_u f$, and $f' = g$.*

In fact, a stronger theorem holds true: we need not assume that the functions f_n are C^1 , only differentiable; then the conclusion is that the function f is differentiable and $f' = g$. The proof of this stronger theorem is much more complicated (though still reasonable); we choose the slightly weaker statement to give a very nice proof based on the Fundamental Theorem of Calculus.

PROOF. Let us agree to the convention that $\int_t^s f = -\int_s^t f$ if $s < t$. By the Fundamental Theorem of Calculus, for each $x \in [a, b]$ we have

$$\int_{x_0}^x f'_n(t) dt = f_n(x) - f_n(x_0).$$

Since $f'_n \rightarrow_u g$ on $[a, b]$, we also have $f'_n \rightarrow_u g$ on $[x_0, x]$ (or $[x, x_0]$ if $x < x_0$), and so by Theorem 10.16, it follows that $\int_{x_0}^x f'_n \rightarrow \int_{x_0}^x g$. By assumption $\lim_{n \rightarrow \infty} f_n(x_0) = y_0$ exists. Thus, we have

$$f_n(x) = f_n(x_0) + \int_{x_0}^x f'_n(t) dt \rightarrow y_0 + \int_{x_0}^x g(t) dt. \quad (10.4)$$

By Theorem 10.9, since f'_n is continuous, so is the uniform limit g . Hence, by the Fundamental Theorem of Calculus, the function $F(x) = \int_{x_0}^x g(t) dt$ is C^1 with $F' = g$. Set $f = y_0 + F$, which is also C^1 with $f' = F' = g$, and (10.4) says that $f_n \rightarrow f$ pointwise.

Thus f_n converges pointwise to a C^1 function f with $f' = g$. We are left to show that the convergence is uniform. Again we apply the Fundamental Theorem of Calculus: for each x ,

$$f_n(x) - f(x) = \left[f_n(x_0) + \int_{x_0}^x f'_n(t) dt \right] - \left[y_0 + \int_{x_0}^x g(t) dt \right] = [f_n(x_0) - y_0] + \int_{x_0}^x [f'_n - g].$$

Thus

$$\sup_{x \in [a, b]} |f_n(x) - f(x)| \leq |f_n(x_0) - y_0| + \sup_{x \in [a, b]} \left| \int_{x_0}^x [f'_n - g] \right|.$$

The first term tends to 0 by assumption; for the second term, let $\epsilon_n = \sup_{x \in [a, b]} |f'_n(x) - g(x)|$. Then

$$\left| \int_{x_0}^x [f'_n - g] \right| \leq \int_{x_0}^x |f'_n - g| \leq \int_{x_0}^x \epsilon_n = \epsilon_n |x - x_0| \leq (b - a)\epsilon_n.$$

This is true uniformly in x . So, since $\epsilon_n \rightarrow 0$ by assumption, we have

$$\sup_{x \in [a, b]} |f_n(x) - f(x)| \leq |f_n(x_0) - y_0| + (b - a)\epsilon_n \rightarrow 0.$$

Thus $f_n \rightarrow_u f$, concluding the proof. \square

Theorem 10.19 is really a theorem about *antiderivatives*, not derivatives. In general, without a priori knowledge of the uniform convergence of the derivative f'_n , we cannot use convergence properties of f_n to conclude anything about f'_n . (We will see a little later how extreme this statement is: we can construct a sequence of C^∞ functions f_n that converges uniformly to a function that is *nowhere differentiable*!) There is one important large class of functions for which we *can* use Theorem 10.19 to conclude smoothness. To get there, we must first talk a little about *series* of functions.

DEFINITION 10.20. Let $f_n: E \rightarrow \mathbb{R}$ be a sequence of real-valued functions defined on a common set E . For each $x \in E$, let $s_n(x) = \sum_{j=1}^n f_j(x)$. If the sequence $s_n: E \rightarrow \mathbb{R}$ converges pointwise, we denote the limit $\lim_{n \rightarrow \infty} s_n = \sum_n f_n$. If s_n converge uniformly, we say that the series $\sum_n f_n$ converges uniformly.

EXAMPLE 10.21. Consider the functions $f_n(x) = \frac{x^2}{(1+x^2)^n}$; these are the same as Example 10.2, modified with a factor of x^2 . We showed in that Example that $\frac{1}{(1+x^2)^n}$ converges to 0 pointwise except at 0, where it converges to 1; hence, here $f_n(x) \rightarrow 0$ pointwise for all x , so there is a chance that the series converges. Consider the series $\sum_{n=0}^{\infty} f_n$. For $x = 0$, this is just the sum of 0, so we get 0; for $x \neq 0$, we have x^2 times a geometric series with ratio $r = \frac{1}{1+x^2} < 1$, so the series converges:

$$\sum_{n=0}^{\infty} f_n(x) = x^2 \sum_{n=0}^{\infty} \left(\frac{1}{1+x^2} \right)^n = x^2 \cdot \frac{1}{1 - \frac{1}{1+x^2}} = 1 + x^2, \quad \text{for } x \neq 0.$$

Thus, the series converges pointwise to the function $(1+x^2)\mathbb{1}_{x \neq 0}$. From Theorem 10.9, it follows that the series does *not* converge uniformly: the functions f_n are all continuous, hence so are the partial sums s_n , but the limit function is discontinuous at 0.

EXAMPLE 10.22. On the other hand, consider the straightforward geometric series: let $f_n(x) = x^n$. For $|x| < 1$, we know that $\sum_{n=0}^{\infty} f_n(x) = \frac{1}{1-x} \equiv s(x)$. The limit function is not continuous at 1, and so we can conclude that the convergence of $\sum_n x^n$ is not uniform on the whole interval $(-1, 1)$. Indeed, we have

$$s(x) - \sum_{j=0}^n x^j = \sum_{j=n+1}^{\infty} x^j = x^{n+1} \sum_{k=0}^{\infty} x^k = x^{n+1} s(x)$$

and so

$$\left| s(x) - \sum_{j=0}^n x^j \right| = \frac{|x|^{n+1}}{1-x}.$$

The supremum of this function on $(-1, 1)$ is ∞ for all n , so the series does not converge uniformly. However, if we restrict the domain to $(-1, 1 - \delta]$ for any $\delta \in (0, 1)$, we have $1 - x \geq \delta$ and $|x|^n \leq (1 - \delta)^n$, so

$$\sup_{-1 < x \leq 1 - \delta} \left| s(x) - \sum_{j=0}^n x^j \right| \leq \frac{(1 - \delta)^n}{\delta}$$

and this uniform bound tends to 0. Thus, the series $\sum_n x^n$ converges uniformly on $(-1, 1 - \delta]$ for each δ .

Example 10.22 demonstrates an important, simple technique, which goes under the name *the Weierstrass M-test*.

LEMMA 10.23 (Weierstrass M-test). Let $f_n: E \rightarrow \mathbb{R}$ be real-valued functions. Suppose that, for each n , there is a constant $0 \leq M_n < \infty$ so that $f_n(x) \leq M_n$ for all $x \in E$. If $\sum_n M_n < \infty$, then $\sum_n f_n$ converges uniformly.

PROOF. Let $s_n(x) = \sum_{j=1}^n f_j(x)$. For any $m > n$, we have

$$|s_m(x) - s_n(x)| = \left| \sum_{j=n+1}^m f_j(x) \right| \leq \sum_{j=n+1}^m |f_j(x)| \leq \sum_{j=n+1}^m M_j.$$

Since the sequence $S_n = \sum_{j=1}^n M_j$ converges, it follows that it is Cauchy, so for any $\epsilon > 0$ there is some N such that $|S_m - S_n| < \epsilon$ for $m > n \geq N$. But $S_m - S_n = \sum_{j=n+1}^m M_j$, and so we see that the sequence of functions $(s_n)_{n=1}^\infty$ is uniformly Cauchy (since N was chosen independently of x). The result now follows from Proposition 10.8. \square

In Example 10.22, the functions $f_n(x) = x^n$ are bounded only by 1 on $(-1, 1)$, and the series $\sum_n 1$ is not convergent, so the Weierstrass M -test does not apply. But if we restrict the domain to $[-1 + \delta, 1 - \delta]$, we have $|f_n(x)| = |x|^n \leq (1 - \delta)^n$, and the series $\sum_n (1 - \delta)^n = \frac{1}{1 - (1 - \delta)} = \frac{1}{\delta}$ is convergent for $\delta > 0$; so we *do* get uniform convergence on this interval. (The explicit computation in Example 10.22 showed that we can even include the interval close to -1 ; this does not follow from the Weierstrass M -test.)

4. Lecture 16: May 19, 2016

Example 10.22 gives a case where a uniform limit of smooth functions is smooth. We will shortly explore what features of this example make that possible, but it is important to stress that *smoothness of a uniform limit is not guaranteed by smoothness of the sequence of functions*. To demonstrate just how far from true that would be, we have the following example which shocked the mathematical world when it was discovered by Weierstrass in 1872.

THEOREM 10.24. *Let $0 < a < 1$, and choose a positive odd integer b large enough that $\frac{\pi}{ab-1} < \frac{2}{3}$ (i.e. $ab > 1 + \frac{3\pi}{2}$). For example, take $a = \frac{1}{2}$ and $b = 11$. Define the function $W : \mathbb{R} \rightarrow \mathbb{R}$ by*

$$W(x) = \sum_{n=0}^{\infty} a^n \cos(b^n \pi x).$$

Then W is uniformly continuous on \mathbb{R} , but is differentiable at no point.

A complete proof of Theorem 10.24 (which is 3 pages long) is posted to the course website as a separate document. It gets a little involved in the precise details, but the key idea is as follows: first, W is continuous because the partial sums ($n = 0$ up to $n = N$) are continuous (in fact they are C^∞), and the series converges uniformly because $|a^n \cos(b^n \pi x)| \leq |a|^n$ which is summable (cf. the Weierstrass M -test); hence the uniform limit function W is continuous by Theorem 10.9. However, the terms have very high frequency oscillations (with amplitudes damping out), which makes the derivative larger and larger, much like in Example 10.3.

Rudin presents a simplified example of a continuous, nowhere-differentiable function which is also constructed as a uniformly convergent series. Its disadvantage is that the terms in the series are not differentiable: they have increasingly dense sets of sharp corners which might lead one to believe it is this proliferation of non-smooth points that results in the non-smoothness of the limit functions. This is not the case; the problem is the oscillations, which can be accomplished with smooth functions, cf. Theorem 10.24 above. We will nevertheless present Rudin's example here for completeness.

THEOREM 10.25. *Let $\varphi : \mathbb{R} \rightarrow \mathbb{R}$ be the continuous function defined by $\varphi(x) = |x|$ for $-1 \leq x \leq 1$, and extended periodically to satisfy $\varphi(x+2) = \varphi(x)$ for all $x \in \mathbb{R}$. Then define*

$$R(x) = \sum_{n=0}^{\infty} \left(\frac{3}{4}\right)^n \varphi(4^n x).$$

The function R is continuous on \mathbb{R} , but nowhere differentiable.

PROOF. The continuity of φ yields the continuity of all the partial sums, and then the uniform convergence of the series via the Weierstrass M -test yields continuity of R , as described above for Weierstrass's function W . We will now show, for any given $x \in \mathbb{R}$, that R is not differentiable at x . To do so, we will construct a sequence $\delta_m \rightarrow 0$ so that

$$\lim_{m \rightarrow \infty} \left| \frac{R(x + \delta_m) - R(x)}{\delta_m} \right| = \infty. \quad (10.5)$$

Well, for any $\delta \in \mathbb{R}$,

$$\left| \frac{R(x + \delta) - R(x)}{\delta} \right| = \left| \sum_{n=0}^{\infty} \left(\frac{3}{4}\right)^n \frac{\varphi(4^n(x + \delta)) - \varphi(4^n x)}{\delta} \right|.$$

Denote this inside difference quotient as $\gamma_n(\delta)$:

$$\gamma_n(\delta) = \frac{\varphi(4^n x + 4^n \delta) - \varphi(4^n x)}{\delta}.$$

If $4^n \delta$ is an even integer, then since φ is 2-periodic, we'll have $\gamma_n(\delta) = 0$. In fact, for a given m , we define

$$\delta_m = \pm \frac{1}{2} \cdot \frac{1}{4^m}$$

(where the sign \pm will be chosen depending on x in a moment). Then $|4^n \delta_m| = 2^{2(n-m)-1}$, so $4^n \delta_m$ is an even integer for all $n > m$, and so $\gamma_n(\delta_m) = 0$ for $n > m$. This means that

$$\left| \frac{R(x + \delta_m) - R(x)}{\delta_m} \right| = \left| \sum_{n=0}^{\infty} \left(\frac{3}{4}\right)^n \gamma_n(\delta_m) \right| = \left| \sum_{n=0}^m \left(\frac{3}{4}\right)^n \gamma_n(\delta_m) \right|.$$

We now apply the reverse triangle inequality to get

$$\left| \frac{R(x + \delta_m) - R(x)}{\delta_m} \right| \geq \left(\frac{3}{4}\right)^m |\gamma_m(\delta_m)| - \sum_{n=0}^{m-1} \left(\frac{3}{4}\right)^n |\gamma_n(\delta_m)|. \quad (10.6)$$

For the positive term $|\gamma_m(\delta_m)|$, note that $4^m \delta_m = \pm \frac{1}{2}$, so

$$|\gamma_m(\delta_m)| = \frac{|\varphi(4^m x \pm \frac{1}{2}) - \varphi(4^m x)|}{|\delta_m|}.$$

Since the interval $[4^m x - \frac{1}{2}, 4^m x + \frac{1}{2}]$ has length 1, it contains at most one integer in its interior. We now choose the sign of δ_m so that there is no integer between $4^m x \pm \frac{1}{2}$ and $4^m x$. Note that the function φ is linear with slope ± 1 on any interval between two integers; thus, with this choice, $\varphi(4^m x \pm \frac{1}{2}) - \varphi(4^m x) = \pm \frac{1}{2}$. Thus

$$|\gamma_m(\delta_m)| = \frac{|\pm \frac{1}{2}|}{|\delta_m|} = 4^m.$$

For the other (negative) terms in (10.6), we use the following fact: for all $s, t \in \mathbb{R}$, $|\varphi(s) - \varphi(t)| \leq |s - t|$. When there is no integer between s and t this is true with equality (as noted above); if there is an integer between them, this follows from the reverse triangle inequality $||s| - |t|| \leq |s - t|$. Thus, in general, we have

$$|\gamma_n(\delta)| = \frac{|\varphi(4^n x + 4^n \delta) - \varphi(4^n x)|}{|\delta|} \leq \frac{|4^n \delta|}{|\delta|} = 4^n.$$

So the negative terms in (10.6) all satisfy

$$-|\gamma_n(\delta_m)| \geq -4^n.$$

Thus, (10.6) yields

$$\left| \frac{R(x + \delta_m) - R(x)}{\delta_m} \right| \geq \left(\frac{3}{4}\right)^m \cdot 4^m - \sum_{n=0}^{m-1} \left(\frac{3}{4}\right)^n \cdot 4^n = 3^m - \sum_{n=0}^{m-1} 3^n = \frac{1}{2}(3^m + 1).$$

In summary, there is a sequence $\delta_m \rightarrow 0$ so that $\left| \frac{R(x + \delta_m) - R(x)}{\delta_m} \right| \rightarrow \infty$. It follows that $R'(x)$ does not exist. \square

Let us turn back, now, to a setting where uniformly convergent series of smooth functions do, in the end, turn out to be smooth. We can generalize the idea of Example 10.22 to study *power series*, the biggest and most important class of functions in Calculus.

DEFINITION 10.26. Let (a_n) be a sequence of real (or complex) numbers. The associated **power series** is the function defined by

$$f(x) = \sum_{n=0}^{\infty} a_n x^n$$

defined on the set of those x for which the series converges.

Of course, the definition might be silly: it might be that there are no x (other than 0, for which the series becomes finite $f(0) = a_0$) converges. For example, $f(x) = \sum_{n=0}^{\infty} n^n x^n$ has this property: no matter how small $|x|$ is, $n^n |x|^n = (n|x|)^n \rightarrow \infty$ as $n \rightarrow \infty$ (since, by the Archimidean property, $n|x| > 1$ for all large n). However, making use of the Root Test, we can find a huge collection of nontrivial power series.

THEOREM 10.27. Let (a_n) be a sequence in \mathbb{R} or \mathbb{C} , and let $R = \frac{1}{\limsup |a_n|^{1/n}} > 0$ (defined to be ∞ if the \limsup is 0). Then the power series $\sum_n c_n x^n$ converges uniformly on any compact subset of $(-R, R)$.

The number R is called the *radius of convergence* of the power series.

PROOF. Let $K \subset (-R, R)$ be a compact subset. Since K is closed and bounded, there is some $\sigma < R$ so that $K \subseteq [-\sigma, \sigma]$ (this follows from the fact that $|K|$ is closed and bounded, and since $(-R, R)$ is open, $\sigma = \sup |K| < R$). For all $x \in K$, we therefore have $|a_n x^n| \leq |a_n| \sigma^n$. Note that $\limsup |a_n \sigma^n|^{1/n} = \sigma \limsup |a_n|^{1/n}$. If $R = \infty$, this is 0; otherwise it is $\frac{\sigma}{R} < 1$. Hence, by the Root Test, the series $\sum_n |a_n| \sigma^n$ converges. The result now follows from the Weierstrass M -test. \square

We can now prove the power series are differentiable, and in fact C^∞ .

THEOREM 10.28. Let $f(x) = \sum_{n=0}^{\infty} a_n x^n$ be a power series with radius of convergence R . Then $f \in C^\infty(-R, R)$, and for all k $f^{(k)}$ is given by the following power series, which also has radius of convergence R :

$$f^{(k)}(x) = \sum_{n=k}^{\infty} n(n-1) \cdots (n-k+1) a_n x^{n-k}. \quad (10.7)$$

PROOF. We begin with the case $k = 1$. Define

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

for those x for which the series converges. (We can reindex to make this an explicit power series, but it's more convenient not to do so at the moment.) Applying the Root Test, the series converges provided $\limsup |n a_n x^{n-1}|^{1/n}$ converges. If $x = 0$ this is clear; if $|x| > 0$, then by Proposition 2.30,

$$|x^{n-1}|^{1/n} = |x|^{\frac{n-1}{n}} = |x|^{1+\frac{1}{n-1}} \rightarrow |x|.$$

By the same proposition, $n^{1/n} \rightarrow 1$. Thus, applying Proposition 2.29,

$$\limsup_{n \rightarrow \infty} |na_n x^{n-1}|^{1/n} \leq \limsup_{n \rightarrow \infty} n^{1/n} \cdot \limsup_{n \rightarrow \infty} |x|^{1+\frac{1}{n-1}} \cdot \limsup_{n \rightarrow \infty} |a_n|^{1/n} = \frac{|x|}{R}.$$

Hence, if $|x| < R$, the the Root Test guarantees that the series defining $g(x)$ converges. What's more, following the proof of Theorem 10.27, if $\sigma < R$, for any $x \in [-\sigma, \sigma]$ we have $|na_n x^{n-1}| \leq n|a_n| \sigma^{n-1}$, and since $\sigma/R < 1$, the above shows that the series $\sum_n na_n \sigma^{n-1}$ converges. Thus, by the Weierstrass M -test, the series g converges uniformly on $[-\sigma, \sigma]$. Thus g is the uniform limit of the partial sums $g_n(x) = \sum_{j=1}^n ja_j x^{j-1}$. These are polynomials; note that that $g_n = f'_n$ where $f_n(x) = \sum_{j=0}^n a_j x^j$. By Theorem 10.27, f_n converges uniformly to f on $[-\sigma, \sigma]$. So we have $C^1[-\sigma, \sigma]$ functions f_n that converge at many points to f , and $f_n \in C^1[-\sigma, \sigma]$ with $f'_n \rightarrow_u g$. By Theorem 10.19, it follows that $f \in C^1[-\sigma, \sigma]$, with $f' = g$.

We can now proceed by induction. Suppose we have shown that $f \in C^k$, with $f^{(k)}$ given by (10.7). Let us reindex the sum to give

$$f^{(k)}(x) = \sum_{m=0}^{\infty} n(n-1) \cdots (m+1) a_{m+k} x^m$$

by assumption having the same radius of convergence R . Set $b_m = n(n-1) \cdots (m+1) a_{m+k}$; then $h = f^{(k)}$ is a new power series $h(x) = \sum_{m=0}^{\infty} b_m x^m$ with radius of convergence R , and so applying the base case above to h , we find that h is C^1 on $(-R, R)$ with derivative $h'(x) = \sum_{m=1}^{\infty} m b_m x^{m-1}$. But this means that f is C^{k+1} on $(-R, R)$ with derivative

$$f^{(k+1)}(x) = h'(x) = \sum_{m=1}^{\infty} m \cdot [n(n-1) \cdots (m+1) a_{m+k}] x^{m-1}$$

and reindexing again this gives

$$f^{(k+1)}(x) = \sum_{n=k+1}^{\infty} n(n-1) \cdots (n-k) a_n x^{n-k-1}$$

which is precisely (10.7) at $k \rightarrow k+1$. This concludes the induction, and the proof. \square

REMARK 10.29. Theorem 10.28 says that power series can be differentiated “term by term”: we can pretend the series terminates finitely and differentiate it precisely as if it were a polynomial.

Functions that are given by a convergent power series are called **analytic**. Actually, we are being too restrictive here: everything above is for power series centered at 0; we can just as well discuss power series like $\sum_n a_n (x - x_0)^n$ centered at any point $x_0 \in \mathbb{R}$. A function is called analytic on a domain if, at each point x_0 in the domain, it has a power series expansion centered at x_0 with positive radius of convergence. The preceding two theorems are about functions analytic at 0. It is a theorem (which we could cover at this point, but won't) that if f has a power series centered at x_0 with radius of convergence R , then f is analytic on $(x_0 - R, x_0 + R)$ (and perhaps beyond): i.e. we can “re-center” the power series. For our present purposes, we will be content only with power series centered at 0.

5. Lecture 17: May 24, 2016

Theorem 10.28 says that analytic functions are C^∞ . The converse is generally false: there are (many) C^∞ functions that are not analytic. To see examples, we first need to connect analyticity with Taylor's theorem.

COROLLARY 10.30. *If f is analytic in a neighborhood of 0, then it is given by its Taylor series*

$$f(x) = \sum_{n=0}^{\infty} \frac{f^{(n)}(0)}{n!} x^n.$$

In particular, analytic functions have unique power series expansions.

PROOF. Let f be given by a convergent power series $f(x) = \sum_{n=0}^{\infty} a_n x^n$. By Theorem 10.28, we may compute $f^{(k)}(0) = k!a_k$. This yields the result. \square

REMARK 10.31. Regarding uniqueness: the above shows that if two power series agree on a neighborhood of 0, then they have the same coefficients and so are equal everywhere on their common domain of convergence. In fact, much more than this is true: it turns out that if two power series agree on any set that has a limit point, then they are equal everywhere. This is a connectedness result that we could cover here, but is best left to a course on complex variables.

Combining this with Taylor's theorem, we can get a sense of what should be true in order for a C^∞ function to be analytic.

PROPOSITION 10.32. *Let $f \in C^\infty(-R, R)$. Suppose that, for each $\sigma \in (0, R)$,*

$$\frac{1}{k!} \sup_{|x|, |\xi| \leq \sigma} |f^{(k)}(\xi)x^k| \rightarrow 0 \quad \text{as } k \rightarrow \infty.$$

Then f is analytic on $(-R, R)$.

PROOF. By Taylor's theorem, for each $x \in [-\sigma, \sigma]$ and each $k \in \mathbb{N}$, there is some point $\xi(x, k)$ between 0 and x such that

$$f(x) = \sum_{n=0}^{k-1} \frac{f^{(n)}(0)}{n!} x^n + \frac{1}{k!} f^{(k)}(\xi(x, k))x^k.$$

Subtracting, taking absolute values, and taking sup over $|x| \leq \sigma$ yields

$$\sup_{|x| \leq \sigma} \left| f(x) - \sum_{n=0}^{k-1} \frac{f^{(n)}(0)}{n!} x^n \right| \leq \frac{1}{k!} \sup_{|x| \leq \sigma} |f^{(k)}(\xi(x, k))x^k| \leq \frac{1}{k!} \sup_{|x|, |\xi| \leq \sigma} |f^{(k)}(\xi)x^k| \rightarrow 0.$$

Thus, the Taylor polynomials $T_0^{k-1}f$ converge uniformly to f on $[-\sigma, \sigma]$ for any $\sigma < R$. Since $T_0^{k-1}f$ also converges to the Taylor series, it follows that f is given by its convergent Taylor series on $(-R, R)$ as claimed. \square

The condition of Proposition 10.32 is stronger than absolutely necessary for a function to be analytic, but in practice it is the easiest way to verify that a given function is analytic.

EXAMPLE 10.33. Suppose that $E: (-R, R) \rightarrow \mathbb{R}$ is a function which satisfies

$$E(x+y) = E(x)E(y) \quad \text{for all } x, y \in (-R, R) \text{ such that } x+y \in (-R, R).$$

Note then that for any $x \in (-R, R)$, $E(x) = E(x+0) = E(x)E(0)$, so $E(x)(E(0) - 1) = 0$ for all x . Thus, either $E(x) = 0$ for all x , or $E(0) = 1$. In the latter (non-boring) case, let us make the

additional assumption that E is differentiable at 0. (This assumption must be added: it does not follow from the functional equation $E(x+y) = E(x)E(y)$.) Then $E'(0) = \lim_{t \rightarrow 0} \frac{E(t)-1}{t}$ exists. We can then compute for any $x \in (-R, R)$, taking t small enough that $x+t \in (-R, R)$,

$$E'(x) = \lim_{t \rightarrow 0} \frac{E(x+t) - E(x)}{t} = \lim_{t \rightarrow 0} \frac{E(x)E(t) - E(x)}{t} = E(x) \lim_{t \rightarrow 0} \frac{E(t) - 1}{t} = E'(0)E(x).$$

So E is differentiable at all points in its domain, and setting $\lambda = E'(0)$, E satisfies the differential equation $E'(x) = \lambda E(x)$ for all x . Iterating this, we see that $E''(x) = \lambda^2 E(x)$, $E'''(x) = \lambda^3 E(x)$, and in general $E^{(k)}(x) = \lambda^k E(x)$. Since E is defined on $(-R, R)$, this shows that $E \in C^k(-R, R)$ for all k , so in fact $E \in C^\infty(-R, R)$. What's more, for any $\sigma < R$, we have

$$\frac{1}{k!} \sup_{|x|, |\xi| \leq \sigma} |E^{(k)}(\xi)x^k| = \frac{|\lambda|^k}{k!} \sup_{|x|, |\xi| \leq \sigma} |E(\xi)x^k| = \frac{(|\lambda|\sigma)^k}{k!} \sup_{|\xi| \leq \sigma} |E(\xi)|.$$

Since E is $C^\infty(-R, R)$, it is certainly continuous on the compact interval $[-\sigma, \sigma]$, hence bounded there. What's more, the sequence $\frac{(|\lambda|\sigma)^k}{k!}$ tends to 0 as $k \rightarrow \infty$ (this follows from arguments similar to the proof of Lemma 4.20). Hence, by Proposition 10.32, E is analytic in $(-R, R)$.

We can now compute the Taylor series expansion of E : we have $E^{(n)}(0) = \lambda^n E(0) = \lambda^n$, and so

$$E(x) = \sum_{n=0}^{\infty} \frac{\lambda^n}{n!} x^n.$$

Note that the coefficients $a_n = \frac{\lambda^n}{n!}$ satisfy $\frac{a_{n+1}}{a_n} = \frac{\lambda^{n+1}}{(n+1)!} \cdot \frac{n!}{\lambda^n} = \frac{\lambda}{n+1} \rightarrow 0$, and hence by Lemma 4.14, $\limsup |a_n|^{1/n} = 0$ as well. By Theorem 10.27, it follows that the power series expansion for E actually converges uniformly on all of \mathbb{R} (i.e. the radius of converge R can be taken to be $R = \infty$).

In the special case $\lambda = 1$, we thus define **the exponential function**

$$\exp(x) = \sum_{n=0}^{\infty} \frac{x^n}{n!}.$$

This function is analytic on \mathbb{R} ; the more general function E above is then simply $E(x) = \exp(\lambda x)$. Comparing to Example 4.19, note that the number e is given by $e = \exp(1)$. We began this discussion with the property $E(x+y) = E(x)E(y)$, and we concluded that any function with this property that happens to be differentiable at 0 must be analytic on \mathbb{R} and have the form $E(x) = \exp(\lambda x)$ for some $\lambda \in \mathbb{R}$ (or, recalling that we discarded the boring case $E(x) = 0$, we could have the constant 0 function as well). That doesn't prove that the function \exp has this property (a priori, it could be true that there is no such differentiable function); but actually, \exp does have the given property. It is possible to prove this by manipulating power series, but it is much more elegant to prove this using a differential equation approach; this is an exercise on Homework 9.

Hence, we do know that $\exp(x+y) = \exp(x)\exp(y)$ for all $x, y \in \mathbb{R}$. From this, it then follows that $\exp(n) = e^n$ for positive integers n ; then $\exp(-n)e^n = \exp(-n+n) = 1$, so $\exp(-n) = e^{-n}$; and moreover $(\exp(1/m))^m = \exp(m \cdot 1/m) = \exp(1) = e$, so $\exp(1/m) = e^{1/m}$. Combining these shows that if $q \in \mathbb{Q}$ then $\exp(q) = e^q$. Since both \exp and $x \mapsto e^x$ are continuous functions, that agree on the dense set \mathbb{Q} in \mathbb{R} , we see that $\exp(x) = e^x$.

EXAMPLE 10.34. The functions \cos and \sin are defined geometrically as follows: for $\theta \in [0, 2\pi)$, $(\cos \theta, \sin \theta)$ is the point on the unit circle such that the arclength of the circle curve from

$(0, 0)$ to $(\cos \theta, \sin \theta)$ is θ . (This is how angles are actually defined: angle θ means arclength θ on a unit circular arc.) From this definition, it is fun and not difficult to prove the following facts:

- $|\sin \theta| \leq 1$ and $|\cos \theta| \leq 1$, and $\cos^2 \theta + \sin^2 \theta = 1$ (by definition: $(\cos \theta, \sin \theta)$ is a point on the unit circle).
- $\lim_{\theta \rightarrow 0} (\cos \theta, \sin \theta) = (1, 0)$. (Follows from the fact that the circle is a continuous curve.)
- $\lim_{\theta \rightarrow 0} \frac{\sin \theta}{\theta} = 1$, and $\lim_{\theta \rightarrow 0} \frac{1 - \cos \theta}{\theta} = 0$. (The first follows by the Squeeze Theorem applied to the inequality $\cos \theta < \frac{\sin \theta}{\theta} < 1$ for θ small, which can be deduced from a triangle diagram.)
- If $\theta + \phi \in [0, 2\pi)$, then $\sin(\theta + \phi) = \sin \theta \cos \phi + \cos \theta \sin \phi$ and $\cos(\theta + \phi) = \cos \theta \cos \phi - \sin \theta \sin \phi$. (These follow from more intricate triangle diagrams.)
- Extending \cos and \sin to all of \mathbb{R} by declaring them to be 2π -periodic, all of the above properties continue to hold for the extended functions.

From these facts (whose fun proofs have been known since Ancient Greece, and you should lookup on the interwebz), it is straightforward to deduce that \sin and \cos are differentiable functions on \mathbb{R} , with $\sin' = \cos$ and $\cos' = -\sin$. For example, we have

$$\begin{aligned} \sin'(\theta) &= \lim_{\phi \rightarrow 0} \frac{\sin(\theta + \phi) - \sin \theta}{\phi} = \lim_{\phi \rightarrow 0} \frac{\sin \theta \cos \phi + \cos \theta \sin \phi - \sin \theta}{\phi} \\ &= \sin \theta \lim_{\phi \rightarrow 0} \frac{\cos \phi - 1}{\phi} + \cos \theta \lim_{\phi \rightarrow 0} \frac{\sin \phi}{\phi} = \cos \theta. \end{aligned}$$

We therefore have a repeating pattern of derivatives: $\sin^{(k)}$ and $\cos^{(k)}$ are all among the four functions $\{\pm \sin, \pm \cos\}$. In particular, this means that all derivatives exist, so \sin and \cos are C^∞ , and moreover all the derivatives are bounded by 1. Thus, applying Proposition 10.32, we have for any $\sigma \in \mathbb{R}$

$$\frac{1}{k!} \sup_{|x|, |\xi| \leq \sigma} |\sin^{(k)}(\xi)x^k| \leq \frac{\sigma^k}{k!} \rightarrow 0, \quad \text{and} \quad \frac{1}{k!} \sup_{|x|, |\xi| \leq \sigma} |\cos^{(k)}(\xi)x^k| \leq \frac{\sigma^k}{k!} \rightarrow 0.$$

Thus, \sin and \cos are analytic on \mathbb{R} . We can then compute their Taylor series: since $\sin(0) = 0$ and $\cos(0) = 1$, this gives the familiar

$$\begin{aligned} \sin x &= x - \frac{x^3}{3!} + \frac{x^5}{5!} - \dots = \sum_{n=0}^{\infty} (-1)^n \frac{x^{2n+1}}{(2n+1)!} \\ \cos x &= 1 - \frac{x^2}{2!} + \frac{x^4}{4!} - \dots = \sum_{n=0}^{\infty} (-1)^n \frac{x^{2n}}{(2n)!}. \end{aligned}$$

In both cases, the coefficient a_k of x^k is either 0 or $\pm \frac{1}{k!}$, and it follows that $\limsup |a_k|^{1/k} = 0$; so, as with the exponential function, the Taylor series of the trigonometric functions \cos and \sin have radius of convergence ∞ .

REMARK 10.35. Power series make sense for complex variables as well, since the series $\sum_n a_n z^n$ can be made sense of (in terms of the complex modulus) as a limit in a metric space. In that wider context, the power series of $\exp z$, $\cos z$, and $\sin z$ all make perfect sense as convergent series in \mathbb{C} . Using the fact that the sequence i^n follows the pattern $(1, i, -1, -i)$ repeated,

which mirrors the derivative pattern of \cos and \sin , it is easy to check from the power series that

$$e^{i\theta} = \cos \theta + i \sin \theta.$$

The special case $\theta = \pi$ gives Euler's famous formula $e^{i\pi} = -1$, or $e^{i\pi} + 1 = 0$. More generally, one can check from the power series that, for all $z \in \mathbb{C}$, $\exp(iz) = \cos(z) + i \sin(z)$. Manipulating this, one can express \cos and \sin in terms of \exp :

$$\cos(z) = \frac{1}{2}[\exp(iz) + \exp(-iz)], \quad \sin(z) = \frac{1}{2i}[\exp(iz) - \exp(-iz)].$$

Thus, the trigonometric functions and the exponential function are really just different (complex) linear combinations of the same function. A full understanding of this topic requires a deeper development of complex variables, which we leave to a separate course on the subject.

EXAMPLE 10.36. Consider the function

$$f(x) = \begin{cases} e^{-1/x} & x > 0 \\ 0 & x \leq 0. \end{cases}$$

This function is C^∞ on $(0, \infty)$ and $(-\infty, 0)$. For $x < 0$, $f^{(k)}(x) = 0$ for all k , of course. For $x > 0$, it is a little messy to compute all the derivatives; the first few are

$$f'(x) = \frac{1}{x^2}e^{-1/x}, \quad f''(x) = \left(-\frac{2}{x^3} + \frac{1}{x^4}\right)e^{-1/x}, \quad f^{(3)}(x) = \left(\frac{6}{x^4} - \frac{6}{x^5} + \frac{1}{x^6}\right)e^{-1/x}.$$

In general, there is a polynomial p_k of degree $2k$ so that $f^{(k)}(x) = p_k(1/x)e^{-1/x}$. This is easily proved by induction. The base case $k = 0$ is immediate, and the cases $k = 1, 2, 3$ are explicitly done above. Assuming we've proved it up to level k , we compute that $f^{(k+1)}(x)$ is equal to

$$\frac{d}{dx}f^{(k)}(x) = \frac{d}{dx}[p_k(1/x)e^{-1/x}] = -\frac{p'_k(1/x)}{x^2}e^{-1/x} + p_k(1/x)\frac{e^{-1/x}}{x^2} = \frac{p_k(1/x) - p'_k(1/x)}{x^2}e^{-1/x}.$$

Set $p_{k+1}(t) = t^2[p_k(t) - p'_k(t)]$; then we've shown that $f^{(k+1)}(x) = p_{k+1}(1/x)e^{-1/x}$. Since p_k is a polynomial of degree $2k$, p'_k is a polynomial of degree $2k - 1$, and so p_{k+1} is a polynomial of degree $2k + 2 = 2(k + 1)$, demonstrating the inductive claim.

Thus, we have seen that $f^{(k)}(x)$ is a (finite) linear combination of terms of the form $x^{-m}e^{-1/x}$ for $x \neq 0$. Now, from the power series expansion for \exp , we have for $t > 0$

$$e^t = \sum_{n=0}^{\infty} \frac{t^n}{n!} > \frac{t^{n+1}}{(n+1)!}.$$

Thus $e^{-t} = \frac{1}{e^t} < \frac{(m+1)!}{t^{m+1}}$ for any m , when $t > 0$. Taking $t = 1/x$ shows that

$$x^{-m}e^{-1/x} = \frac{e^{-1/x}}{x^m} < \frac{(m+1)!}{(1/x)^{m+1}} \cdot \frac{1}{x^m} = (m+1)!x$$

and this tends to 0 as $x \rightarrow 0+$. This $\lim_{x \rightarrow 0+} f^{(k)}(x) = 0$ for all k ; also, since $f^{(k)}(x) = 0$ for $x < 0$, $\lim_{x \rightarrow 0-} f^{(k)}(x) = 0$. In addition

$$\lim_{x \rightarrow 0+} \frac{f^{(k)}(x) - 0}{x} = \lim_{x \rightarrow 0+} \frac{1}{x} p_k(1/x) e^{-1/x} = 0$$

using the same argument as above; and of course the left limit of the difference quotient is also 0. Thus, $f^{(k)}$ is actually continuous at 0, with value 0, for all k .

So we see that $f \in C^\infty(\mathbb{R})$, and that $f^{(k)}(0) = 0$ for all k . Its Taylor series centered at 0 is thus $(T_0 f)(x) = 0$. But this does not converge to the function's value for any $x > 0$: $f(x) > 0$ for all $x > 0$. Thus, although the radius of convergence of the Taylor series is ∞ , it does not actually converge to the value of the function on *any* interval around 0 – the remainder term does not tend to 0.

There are several important points to glean from this example. First, we see f is an example of a C^∞ function that is *not analytic*. Second, this demonstrates that it is very important, when trying to prove that a function is analytic, to show that the remainder term in Taylor's theorem tends to 0; it is *not enough* to check that the Taylor series has positive radius of convergence, as it may not in fact converge to the function! Third: there exist C^∞ functions that are 0 on an interval; this was not clear before this example. In fact, we can finesse this and produce an example of a C_c^∞ function: a function which is C^∞ smooth, and also “compactly-supported”, meaning that it is 0 outside some compact interval. Indeed, with f as above, just take

$$\psi(x) = f(x)f(1-x).$$

When $x \leq 0$, $f(x) = 0$ so $\psi(x) = 0$; when $x \geq 1$, $f(1-x) = 0$ so $\psi(x) = 0$; on the other hand, $\psi(x) > 0$ for $x \in (0, 1)$, and of course $\psi \in C^\infty$. This is a typical example of a “bump function”, and the existence of such smooth, compactly supported functions is of utmost important to analysis and geometry.




FIGURE 5. The function f on $[0, \frac{1}{2}]$ (left), and the bump function ψ on $[0, 1]$ (right). Notice that $f(x)$ is extremely flat as $x \rightarrow 0$.

REMARK 10.37. Example 10.36 gives a function that is C^∞ but fails to be analytic at a point. It is, however, analytic at all other points, as the function

$$\frac{1}{x} = \frac{1}{\lambda} \cdot \frac{1}{1 - (1-x/\lambda)} = \frac{1}{\lambda} \sum_{n=0}^{\infty} (1-x/\lambda)^n = \sum_{n=0}^{\infty} (-1)^n \lambda^{n-1} (x-\lambda)^n$$

has a convergent power series expansion with radius of convergence $\frac{1}{\lambda}$ centered at any point $\lambda > 0$; and since \exp is analytic, so is the composition $\exp(-1/x)$ (as can easily be proved by manipulating power series algebraically).

In the spirit of pathologies we've studied in this class, we might wonder if one can produce a function that is C^∞ but not analytic at *any point*. Indeed, such functions exist, and they are abundant (there is a precise sense in which “most” C^∞ functions are nowhere analytic). For example, it can be shown that the series

$$f(x) = \sum_{n=0}^{\infty} e^{-2^{n/2}} \cos(2^n x)$$

converges to C^∞ function on \mathbb{R} such that the radius of convergence of the Taylor series of f centered at any point is 0. For a non-contrived example from probability theory, let X_n be an infinite sequence of independent random variables each having the uniform distribution on $[0, 1]$. Set

$$X = \sum_{n=1}^{\infty} 2^{-n} X_n.$$

Then X is a random variable taking values in $[0, 1]$. In 1966, J. Fabius showed that the cumulative distribution function $F(x) = \mathbb{P}(X \leq x)$ is a C^∞ function which is analytic nowhere. Indeed, he showed that F is differentiable on $(0, 1)$ with $F'(x) = 2F(2x)$ when $0 \leq x \leq \frac{1}{2}$ (and $F(1-x) = 1 - F(x)$ for $\frac{1}{2} \leq x \leq 1$). Iterating this shows that F is C^∞ , similarly to how we showed \exp is C^∞ ; but the scalings by 2 inside and outside the function have the effect that the Taylor series coefficients of F centered at any point (other than 0 and 1) blow up so fast that the radius of convergence is always 0. You can find a nice picture of Fabius's function on Wikipedia.

6. Lecture 18: May 26, 2016

We've now seen a large class of functions (analytic functions) that are uniform limits of smooth functions, and are themselves smooth. This is to be contrasted with Weierstrass's function, cf. Theorem 10.24, which is a uniform limit of smooth functions, yet is itself nowhere differentiable. Now, analytic functions are uniform limits of polynomials, so you might be tempted to believe that polynomials are somehow special with regard to uniform limits. In fact, nothing could be farther from the truth: we will soon see that *every* continuous function (including Weierstrass's nowhere differentiable function) on a compact interval is a uniform limit of polynomials. To prove this, we will first develop a ubiquitous and useful tool called *convolution* which is widely used to approximate functions by smoother functions. Convolution is built using integration; as a first step, we present the following results on continuity and differentiability of functions given by integrating a two-variable kernel.

LEMMA 10.38. *Let $K \subseteq \mathbb{R}^d$ be a compact set, $a < b$ in \mathbb{R} , and $\alpha: [a, b] \rightarrow \mathbb{R}$ be monotone increasing. Let $f: K \times [a, b] \rightarrow \mathbb{R}$ be a continuous function. Define $F: K \rightarrow \mathbb{R}$ by*

$$F(x) = \int_a^b f(x, t) d\alpha(t).$$

Then F is a continuous function.

The proof of this Lemma is an exercise on Homework 9.

Now, consider the special case that K is also a compact interval, and suppose that $f(x, t)$ is differentiable in the x variable. Under mild additional conditions, it turns out that F is then differentiable in the x variable, and we can compute its derivative by differentiating under the integral sign.

THEOREM 10.39. *Let $u < v$ and $a < b$ in \mathbb{R} , and let $\alpha: [a, b] \rightarrow \mathbb{R}$ be monotone increasing. Let $f: [u, v] \times [a, b] \rightarrow \mathbb{R}$ be a continuous function. Suppose also that for each fixed $t \in [a, b]$, the function $s \mapsto f(s, t)$ is differentiable on (u, v) , and that the function $f_s(s, t) = \frac{\partial}{\partial s} f(s, t)$ is continuous. Then $F(s) = \int_a^b f(s, t) d\alpha(t)$ is differentiable on (u, v) , and $F'(s) = \int_a^b f_s(s, t) d\alpha(t)$. I.e.*

$$\frac{d}{ds} \int_a^b f(s, t) d\alpha(t) = \int_a^b \frac{\partial}{\partial s} f(s, t) d\alpha(t).$$

PROOF. For $s \in (u, v)$, choose $\delta > 0$ small enough that $[s - \delta, s + \delta] \subset (u, v)$. Then for $|h| < \delta$, $f(s + h, t)$ is well-defined, and we have for $0 < |h| \leq \delta$

$$\frac{F(s + h) - F(s)}{h} = \int_a^b \frac{f(s + h, t) - f(s, t)}{h} d\alpha(t).$$

Now, define a function $g: [s - \delta, s + \delta] \times [a, b] \rightarrow \mathbb{R}$ by

$$g(h, t) = \int_0^1 f_s(s + rh, t) dr.$$

Notice that $g(0, t) = f_s(s, t)$, while for $0 < |h| \leq \delta$, since f_s is continuous, we may apply the Fundamental Theorem of Calculus and the chain rule to compute that

$$g(h, t) = \frac{f(s + h, t) - f(s, t)}{h}.$$

The function $(h, t, r) \mapsto f_s(s + rh, t)$ is continuous, and so Applying Lemma 10.38 with r as the integration variable, we conclude that g is continuous. We can then apply Lemma 10.38 once more, this time to g with t as the variable of integration, with the conclusion that $h \mapsto \int_a^b g(h, t) d\alpha(t)$ is continuous, and in particular at $h = 0$. Thus

$$\lim_{h \rightarrow 0} \frac{F(s+h) - F(s)}{h} = \lim_{h \rightarrow 0} \int_a^b g(h, t) d\alpha(t) = \int_a^b g(0, t) d\alpha(t) = \int_a^b f_s(s, t) d\alpha(t).$$

This concludes the proof. \square

Now, let us define the convolution of two functions.

DEFINITION 10.40. *Let $f, g: \mathbb{R} \rightarrow \mathbb{R}$ be functions that are both Riemann integrable on all compact intervals. Then for each fixed y , the function $x \mapsto f(x - y)$ and $x \mapsto g(x - y)$ are also Riemann integrable on all compact intervals, as can be readily checked. Thus the function $x \mapsto f(x - y)g(y)$ is Riemann integrable on all compact intervals. Since one of f and g is 0 outside some compact set, the same is true of the function $x \mapsto f(x - y)g(y)$. If this integrand is 0 outside $[a, b]$, the **convolution** $f * g: \mathbb{R} \rightarrow \mathbb{R}$ is defined by*

$$(f * g)(x) = \int_a^b f(x - y)g(y) dy.$$

*The value of $(f * g)(x)$ does not depend on which $a < b$ are chosen, provided that $y \mapsto f(x - y)g(y)$ is 0 outside $[a, b]$.*

The (closure of) the set where a function is nonzero is called its **support**. The last comment is that we are really just integrating over the support of $y \mapsto f(x - y)g(y)$; if we integrate over a larger interval, it makes no difference to the value since the integrand is 0 outside its support.

One property of convolution that may not be obvious from the definition is that it doesn't matter what order you convolve the functions.

LEMMA 10.41. *For f, g as in Definition 10.40, and any $x \in \mathbb{R}$, $(f * g)(x) = (g * f)(x)$.*

PROOF. Fix x , and choose $a < b$ so that $f(x - y)g(y) = 0$ for $y \notin [a, b]$. In the integral defining $(f * g)(x)$, we make the change of variables $t = x - y$. Then $dt = -dy$, and we have

$$(f * g)(x) = \int_a^b f(x - y)g(y) dy = \int_{x-a}^{x-b} f(t)g(t - x) (-dt) = \int_{x-b}^{x-a} g(t - x)f(t) dt.$$

Since $f(x - y)g(y) = 0$ when $y \notin [a, b]$, it follows that $g(t - x)f(t) = 0$ when $t \notin [x - b, x - a]$. Since the definition of convolution does not depend on which interval outside the support of the integrand is used, it follows that the last integral is equal to $(g * f)(x)$. \square

So convolution is a symmetric operation on functions, which produces a new function. The question is: what kind of function is $f * g$? Let's consider a few examples.

EXAMPLE 10.42. For $\delta > 0$, let $f_\delta = \frac{1}{2\delta} \mathbb{1}_{[-\delta, \delta]}$. This function is 0 outside the compact interval $[-\delta, \delta]$, and is Riemann integrable on all compact intervals. If g is any Riemann integrable function, we compute that

$$(f_\delta * g)(x) = \int_{-\delta}^{\delta} g(x - y)f_\delta(y) dy = \int_{-\delta}^{\delta} g(x - y) \frac{1}{2\delta} dy = \frac{1}{2\delta} \int_{x-\delta}^{x+\delta} g(t) dt$$

where we have made the change of variables $x - y = t$ in the last step. So: $f_\delta * g$ is the "running average" of g : for each x , it replaces $g(x)$ with the average value g takes on $[-\delta, \delta]$. Notice that if

g is continuous, the Fundamental Theorem of Calculus shows that $f_\delta * g$ is differentiable at all x : letting $G(x) = \int_0^x g(t) dt$, we have

$$(f_\delta * g)(x) = \frac{G(x + \delta) - G(x - \delta)}{2\delta}.$$

It follows then that $\lim_{\delta \rightarrow 0} (f_\delta * g)(x) = g(x)$ for all x . In fact,

$$\begin{aligned} \sup_{x \in \mathbb{R}} |(f_\delta * g)(x) - g(x)| &= \sup_{x \in \mathbb{R}} \left| \frac{1}{2\delta} \int_{x-\delta}^{x+\delta} g(t) dt - g(x) \right| = \sup_{x \in \mathbb{R}} \left| \frac{1}{2\delta} \int_{x-\delta}^{x+\delta} [g(t) - g(x)] dt \right| \\ &\leq \frac{1}{2\delta} \sup_{x \in \mathbb{R}} \int_{x-\delta}^{x+\delta} |g(t) - g(x)| dt. \end{aligned}$$

Now, suppose that g is uniformly continuous. Fix $\epsilon > 0$, and let $\delta_0 > 0$ be (uniformly) small enough that $|g(t) - g(s)| < \epsilon$ whenever $|s - t| < \delta_0$. Then for all $\delta \leq \delta_0$, we have

$$\sup_{x \in \mathbb{R}} |(f_\delta * g)(x) - g(x)| \leq \frac{1}{2\delta} \sup_{x \in \mathbb{R}} \int_{x-\delta}^{x+\delta} \epsilon dt = \epsilon.$$

Thus, we see that $(f_\delta * g) \rightarrow_u g$ as $\delta \rightarrow 0$ in this case. (We have thus far only defined uniform convergence for a sequence of functions, but the definition makes sense with the natural modifications for a family of functions depending on a continuous parameter like δ . Alternatively, you could interpret the above as saying that the sequence, say, $f_{1/n} * g$ converges uniformly to g .)

Note, in this example, that if we start with a continuous g , the convolution $f_\delta * g$ is actually differentiable, even though f_δ is discontinuous, and g need not be differentiable. While we don't always get a "step up" like this, it is true that the convolution of two functions is always at least as smooth as the smoothest of the two.

LEMMA 10.43. *Let $f: \mathbb{R} \rightarrow \mathbb{R}$ be a C^k function, and let $g: [a, b] \rightarrow \mathbb{R}$ be continuous; extend g to \mathbb{R} by setting $g = 0$ outside $[a, b]$. Then $f * g$ is C^k on \mathbb{R} , and $(f * g)^{(k)} = f^{(k)} * g$.*

The proof of Lemma 10.43 is an exercise on Homework 10.

Example 10.42 shows that convolution can be used to produce an averaging operation on a function, and this is more or less the way we will always use it. In fact, we can generalize the same properties to a wide class of function sequences called *approximate identities*.

LEMMA 10.44. *Let $\psi_n: \mathbb{R} \rightarrow \mathbb{R}$ for $n \in \mathbb{N}$ be a sequence of functions that are non-negative, all supported in a given compact interval, such that $\int \psi_n(y) dy = 1$ for all n , and such that*

$$\text{for each } \delta > 0, \lim_{n \rightarrow \infty} \int_{|y| > \delta} \psi_n(y) dy = 0. \quad (10.8)$$

*(Such sequences are called approximate identities.) If f is a continuous function on \mathbb{R} , then $\psi_n * f$ converges uniformly to f on compact subsets of \mathbb{R} .*

One easy way to achieve the conditions of an approximate identity sequence is to let the support interval of ψ_n shrink to 0, and then normalize the functions to each have integral 1, such as in Example 10.42. But, as we will see below, it is sometimes convenient not to assume this; we only need that the total mass of ψ_n is asymptotically concentrated in an arbitrarily small neighborhood of 0.

PROOF. Let $r > 0$ be large enough that all the supports of all the ψ_n are contained in $[-r, r]$. Since $\int_{-r}^r \psi_n(y) dy = 1$, we can write $f(x) = \int_{-r}^r f(x)\psi_n(y) dy$ for any n . Then

$$|(\psi_n * f)(x) - f(x)| = \left| \int_{-r}^r \psi_n(y)[f(x-y) - f(x)] dy \right| \leq \int_{-r}^r \psi_n(y)|f(x-y) - f(x)| dy.$$

Now, let $K \subset \mathbb{R}$ be compact; then $K \subseteq [-s, s]$ for some $s > 0$. For all $x \in K$, therefore, $x - y$ is contained in $[-r - s, r + s]$ for all $|y| \leq r$. Since f is continuous on \mathbb{R} , it is uniformly continuous on $[-r - s, r + s]$. Fix $\epsilon > 0$, and select a uniform $\delta > 0$ (with $\delta < r + s$) so that $|f(s) - f(t)| < \frac{\epsilon}{2}$ for $|s - t| < \delta$ with $s, t \in [-r - s, r + s]$. Then we can breakup the integral

$$\int_{-r}^r \psi_n(y)|f(x-y) - f(x)| dy = \int_{-\delta}^{\delta} \psi_n(y)|f(x-y) - f(x)| dy + \int_{\delta < |y| \leq r} \psi_n(y)|f(x-y) - f(x)| dy.$$

For the first term, $x - y$ and x are contained in $[-r - s, r + s]$ and $|(x - y) - x| = |y| < \delta$. Thus, $|f(x - y) - f(x)| < \frac{\epsilon}{2}$ in that term. For the second term, let $M = \max\{f(t) : |t| \leq r + s\}$; then $|f(x - y) - f(x)| \leq 2M$ for $|y| \leq r$. In total, then, we have

$$\begin{aligned} \sup_{x \in K} |(\psi_n * f)(x) - f(x)| &\leq \int_{-\delta}^{\delta} \psi_n(y) \cdot \frac{\epsilon}{2} dy + \int_{\delta < |y| \leq r} \psi_n(y) \cdot 2M dy \\ &\leq \frac{\epsilon}{2} + 2M \int_{|y| > \delta} \psi_n(y) dy. \end{aligned}$$

By (10.8), we may choose N (uniform in x) large enough that the remaining integral is $< \frac{\epsilon}{4M}$ for $n \geq N$. This shows that $\psi_n * f \rightarrow_u f$ on K , as desired. \square

By Lemma 10.43, if the approximate identity sequence ψ_n consists of all smooth functions (say C^∞), then the functions $\psi_n * f$ will all be C^∞ as well, no matter how rough f is. This is a very important technique: it is always possible to approximate any continuous function f uniformly by smooth functions.

7. Lecture 19: May 31, 2016

We will momentarily use Lemma 10.44 to approximate any continuous function uniformly by polynomials. To see how, we note that convolving any function f with a polynomial p yields a polynomial. In fact, we only need p to be equal to a polynomial on a certain compact set determined by the support of f ; p can be anything we like outside this. To make things simple, we take the support of f to be $[0, 1]$.

LEMMA 10.45. *If f is Riemann integrable and supported in $[0, 1]$, and if $p: [-1, 1] \rightarrow \mathbb{R}$ is equal to a polynomial on this domain, then $p * f$ is a polynomial on $[0, 1]$.*

PROOF. Because p is a polynomial on $[-1, 1]$, by applying the binomial theorem, we see that for $x, y \in [-1, 1]$, $(x, y) \mapsto p(x - y)$ is a polynomial in two variables. In particular, we can expand

$$p(x - y) = \sum_{k=0}^d a_k(y)x^k$$

where d is the degree of p , and a_k are single-variable polynomials. Now, by definition $(p * f)(x) = \int_0^1 p(x - y)f(y) dy$; for $x, y \in [0, 1]$, the difference $x - y$ is in $[-1, 1]$, where p is a polynomial. Then using the linearity of the integral, we have

$$(p * f)(x) = \int_0^1 p(x - y)f(y) dy = \int_0^1 \sum_{k=0}^d a_k(y)x^k f(y) dy = \sum_{k=0}^d x^k \int_0^1 a_k(y)f(y) dy.$$

The polynomial $a_k(\cdot)$ is Riemann integrable on $[0, 1]$, as is f , so it follows that these integrated coefficients are finite real numbers, and this shows that for $x \in [0, 1]$ $x \mapsto p * f(x)$ is a polynomial function. \square

REMARK 10.46. In general, if f is supported in a given compact set K , for this proof to work we see that p has to equal a polynomial on $K - K$: the set of all points of the form $x - y$ where $x, y \in K$. If p is simply a polynomial everywhere, this is certainly true, but it will be advantageous to allow p to equal a polynomial in some interval and be identically 0 outside that interval.

Combining Lemmas 10.44 and 10.45 allows us to prove the Weierstrass approximation theorem.

THEOREM 10.47 (Weierstrass Approximation Theorem). *Let $a < b$ in \mathbb{R} , and let $f \in C[a, b]$. There exists a sequence of polynomials p_n with $p_n \rightarrow_u f$.*

PROOF. We will use an approximate identity sequence, and invoke Lemma 10.44. To do so requires extending f to be continuous on all of \mathbb{R} . We take care of this, at the same time shifting all the action into the unit interval, by translating and dilating. Indeed, if we define

$$\tilde{f}(x) = f(a + x(b - a)) - (f(a) + x[f(b) - f(a)]),$$

then \tilde{f} is continuous on $[0, 1]$ and $\tilde{f}(0) = \tilde{f}(1) = 0$. Hence we can extend \tilde{f} continuously to \mathbb{R} by setting it equal to 0 outside $[0, 1]$. Once we find polynomials \tilde{p}_n uniformly approximating \tilde{f} , it follows that the polynomials $\tilde{q}_n(x) = \tilde{p}_n(x) + f(a) + x[f(b) - f(a)]$ uniformly approximate $f(a + x(b - a))$, and by translating and dilating back, $p_n(x) = \tilde{q}_n(\frac{x-a}{b-a})$ uniformly approximate f , as the reader can easily verify. We will henceforth replace the notation \tilde{f} with f , and assume that f is continuous on \mathbb{R} and supported in $[0, 1]$.

The proof now simply consists in showing there exists an approximate identity sequence consisting of polynomials on $[-1, 1]$. There are many; the most well known is the sequence of *Bernstein polynomials*, defined by

$$\psi_n(x) = \frac{1}{c_n}(1-x^2)^n \mathbb{1}_{[-1,1]}(x), \quad \text{where } c_n \equiv \int_{-1}^1 (1-x^2)^n dx.$$

By Lemma 10.45, $\psi_n * f$ is a polynomial function on $[0, 1]$ for each n . It is immediately verifiable




FIGURE 6. The Bernstein polynomials ψ_n , for $n = 1, 2, 5, 10, 100$.

than ψ_n is non-negative, supported in $[-1, 1]$ for all n , and normalized so that $\int_{-1}^1 \psi_n(y) dy = 1$. To show that $\{\psi_n\}$ is an approximate identity, we need to show all the mass of ψ_n concentrates near 0. To that end, fix $\delta \in (0, 1)$. We note that

$$\int_0^1 (1-x^2)^n dx \geq \int_0^\delta (1-x^2)^n dx \geq \int_0^\delta \frac{x}{\delta} (1-x^2)^n dx$$

since $\frac{x}{\delta} \leq 1$ for $x \in [0, \delta]$. By the same reasoning, we have

$$\int_\delta^1 (1-x^2)^n dx \leq \int_\delta^1 \frac{x}{\delta} (1-x^2)^n dx.$$

Since $\psi_n(x) = \psi_n(-x)$, we therefore have

$$\int_{|x|>\delta} \psi_n(x) dx = 2 \int_\delta^1 \frac{1}{c_n} (1-x^2)^n dx \leq \frac{2}{\delta c_n} \int_\delta^1 x(1-x^2)^n dx.$$

Similarly

$$c_n = \int_{-1}^1 (1-x^2)^n dx = 2 \int_0^1 (1-x^2)^n dx \geq 2 \int_0^\delta \frac{x}{\delta} (1-x^2)^n dx.$$

Putting these together gives

$$\int_{|x|>\delta} \psi_n(x) \leq \frac{\int_\delta^1 x(1-x^2)^n dx}{\int_0^\delta x(1-x^2)^n dx}.$$

Making the change of variables $u = 1 - x^2$ gives $du = -2x dx$, and so this ratio becomes

$$\frac{\int_{\delta}^1 x(1-x^2)^n dx}{\int_0^{\delta} x(1-x^2)^n dx} = \frac{\int_0^{1-\delta^2} u^n du}{\int_{1-\delta^2}^1 u^n du} = \frac{(1-\delta^2)^{n+1}}{1-(1-\delta^2)^{n+1}}.$$

The reader can now quickly check that this tends to 0 as $n \rightarrow \infty$.

Thus, we have shown that ψ_n is an approximate identity sequence. Invoking Lemma 10.44, it follows that $p_n = \psi_n * f$ converges to f uniformly on $[0, 1]$. But, as shown above following Lemma 10.45, p_n is a polynomial function on $[0, 1]$. This completes the proof. \square

The Weierstrass approximation theorem is a very useful tool for theoretical computations. But in many instances it would be better to work with a different class of functions more adapted to the problem at hand. The question then becomes: what's so special about polynomials that allows them to uniformly approximate continuous functions? Can we find other families of nice functions that uniformly approximate continuous functions?

This was answered with a resounding *yes!*, by Marshall Stone in 1937. Stone was a towering intellect at the University of Chicago who attracted other world-class faculty there and turned that department into one of the best in the country; the period when he was most active there was called the "Stone Age".

Here are the properties of the family of polynomial functions on a compact set that turn out to be responsible for their uniform density in the continuous functions.

DEFINITION 10.48. *Let \mathcal{F} be a collection of real (or complex) valued functions all defined on some set X .*

- \mathcal{F} is an **algebra** if it is closed under pointwise addition, scalar multiplication, and pointwise multiplication. That is: given $f, g \in \mathcal{F}$ and $\lambda \in \mathbb{R}$ (or \mathbb{C}), the functions $f + g$, λf , and fg are also in \mathcal{F} .
- If the functions in \mathcal{F} are \mathbb{C} -valued, say that \mathcal{F} is **self-adjoint** if it is closed under complex conjugation: for each $f \in \mathcal{F}$, \bar{f} is also in \mathcal{F} .
- We say \mathcal{F} **separates points** if, for any pair $x \neq y$ in X , there is some function $f \in \mathcal{F}$ such that $f(x) \neq f(y)$.
- \mathcal{F} is said to **vanish nowhere** if there is no single point $x \in X$ such that $f(x) = 0$ for all $f \in \mathcal{F}$.

These properties all hold for polynomials: polynomials are in fact the smallest algebra of functions that contain both constant functions and the function $f(x) = x$; they separate points since the function $f(x) = x$ does this for any pair of distinct points; and they vanish nowhere since they contain constant non-zero functions. Stone realized that these properties are all that is needed to guarantee that a family \mathcal{F} of functions can be used to uniformly approximate all continuous functions on a compact metric space X . His original proof was quite complicated; here we will present a simplified for of his proof that he published in *Mathematics Magazine* in 1948. First, let us use the language of metric spaces from the beginning.

Let X be a compact metric space. By the extreme values theorem, any function $f: X \rightarrow \mathbb{R}$ (or \mathbb{C}) is bounded (for the complex case, just consider the real-valued function $|f|$, which is continuous as well, so achieves its maximum on X). Thus, the uniform "metric"

$$d_u(f, g) = \sup_{x \in X} |f(x) - g(x)|$$

is a genuine metric on $C(X)$. Then saying that a sequence f_n of functions in $C(X)$ uniformly approximates a function $f \in C(X)$ means that $d_u(f_n, f) \rightarrow 0$ as $n \rightarrow \infty$; in other words, $f_n \rightarrow f$ in the metric space $(C(X), d_u)$. Hence, in this language, the Weierstrass approximation theorem says that every function $f \in C[a, b]$ is the uniform limit of some sequence p_n of polynomials on $[a, b]$: i.e. in the metric space $(C[a, b], d_u)$, the set of polynomials is a *dense set*. We usually summarize this by calling them *uniformly dense*. In other words, if $\mathcal{P}[a, b]$ denotes the set of polynomial functions, then the *uniform closure* of $\mathcal{P}[a, b]$ (the closure of this set in the metric space $(C[a, b], d_u)$) is call of $C[a, b]$.

THEOREM 10.49 (Stone–Weierstrass Approximation Theorem). *Let \mathcal{A} be an algebra of real (or complex) valued functions on a compact metric space X . Suppose that \mathcal{A} is self-adjoint, separates points, and vanishes nowhere. Then \mathcal{A} is uniformly dense in $C(X)$.*

Note: if \mathcal{A} consists of real valued functions, it is automatically self-adjoint, since $\bar{f} = f$ for any real valued function. In fact, the complex case follows fairly easily from the real case, as follows.

PROOF OF \mathbb{C} -VALUED STONE-WEIERSTRASS, ASSUMING \mathbb{R} -VALUED CASE. Let \mathcal{A} satisfy all the conditions of the theorem. Consider the set of real valued functions in \mathcal{A} :

$$\mathcal{A}_{\mathbb{R}} = \{u \in \mathcal{A} : u \text{ is real valued}\}.$$

Since \mathcal{A} is closed under sum and product, so is $\mathcal{A}_{\mathbb{R}}$ (because the sum and product of real valued functions are real valued), and $\mathcal{A}_{\mathbb{R}}$ is also closed under scalar multiplication by \mathbb{R} : if $\lambda \in \mathbb{R}$ and $u \in \mathcal{A}_{\mathbb{R}}$, then $u \in \mathcal{A}$ so $\lambda u \in \mathcal{A}$, but this function is also real valued, so it is in $\mathcal{A}_{\mathbb{R}}$; thus $\mathcal{A}_{\mathbb{R}}$ is an algebra.

We will now show that $\mathcal{A}_{\mathbb{R}}$ separates points and vanishes nowhere; to do so, we need to use the self-adjointness of \mathcal{A} . Since $\bar{f} \in \mathcal{A}$ for each $f \in \mathcal{A}$, and since \mathcal{A} is a vector space, \mathcal{A} also contains $\operatorname{Re} f$ and $\operatorname{Im} f$ for each $f \in \mathcal{A}$:

$$\operatorname{Re} f = \frac{f + \bar{f}}{2}, \quad \operatorname{Im} f = \frac{f - \bar{f}}{2i}.$$

Thus $\mathcal{A}_{\mathbb{R}}$ contains $\operatorname{Re} f$ and $\operatorname{Im} f$ for all $f \in \mathcal{A}$. (Without this assumption, $\mathcal{A}_{\mathbb{R}}$ could be very small; indeed, it could consist only of the 0 function.) Now, let $x \neq y$ in X ; then there is some $f \in \mathcal{A}$ such that $f(x) \neq f(y)$. This means that either $\operatorname{Re} f(x) \neq \operatorname{Re} f(y)$ or $\operatorname{Im} f(x) \neq \operatorname{Im} f(y)$; in either case, there is some element of $\mathcal{A}_{\mathbb{R}}$ that assigns different values to x and y , so $\mathcal{A}_{\mathbb{R}}$ separates points. Similarly, given any $x \in X$, there is some $f \in \mathcal{A}$ with $f(x) \neq 0$; thus one of $\operatorname{Re} f(x)$ and $\operatorname{Im} f(x)$ is non-zero, and so some element of $\mathcal{A}_{\mathbb{R}}$ does not vanish at x . Thus $\mathcal{A}_{\mathbb{R}}$ vanishes nowhere.

Hence $\mathcal{A}_{\mathbb{R}}$ is a real algebra that separates points and vanishes nowhere; once we have proven the Stone–Weierstrass approximation theorem for such real algebras, it will follow that $\mathcal{A}_{\mathbb{R}}$ is uniformly dense in the real-valued continuous functions on X . But then for any continuous function $f: X \rightarrow \mathbb{C}$, taking $u = \operatorname{Re} f$ and $v = \operatorname{Im} f$, we can find $u_n \in \mathcal{A}_{\mathbb{R}}$ with $u_n \rightarrow_u u$ and $v_n \in \mathcal{A}_{\mathbb{R}}$ with $v_n \rightarrow_u v$; thus $f_n \equiv u_n + iv_n \rightarrow_u u + iv = f$, and since \mathcal{A} is a \mathbb{C} algebra containing u_n and v_n , $f_n = u_n + iv_n \in \mathcal{A}$. Hence \mathcal{A} is uniformly dense in the continuous complex valued functions on X . \square

This leaves us to prove the real case: that any algebra of functions that separates points and vanishes nowhere is uniformly dense in $C(X)$. We will prove this in a series of lemmas. The first is that the uniform closure of an algebra is, itself, an algebra.

LEMMA 10.50. *Let \mathcal{A} be an algebra of bounded real-valued functions on a compact set X . Then the uniform closure $\overline{\mathcal{A}}$ of \mathcal{A} is also an algebra.*

PROOF. Let $f, g \in \overline{\mathcal{A}}$, and let $\lambda \in \mathbb{R}$. By definition of uniform closure, there are sequences $f_n, g_n \in \mathcal{A}$ with $f_n \rightarrow_u f$ and $g_n \rightarrow_u g$. Since \mathcal{A} is an algebra, $f_n + \lambda g_n \in \mathcal{A}$ for each n . Now

$$\begin{aligned} d_u(f_n + \lambda g_n, f + \lambda g) &= \sup_{x \in X} |(f_n(x) + \lambda g_n(x)) - (f(x) + \lambda g(x))| \\ &\leq \sup_{x \in X} |f_n(x) - f(x)| + \lambda \sup_{x \in X} |g_n(x) - g(x)| \\ &= d_u(f_n, f) + \lambda d_u(g_n, g). \end{aligned}$$

Since $d_u(f_n, f) \rightarrow 0$ and $d_u(g_n, g) \rightarrow 0$, it then follows that $d_u(f_n + \lambda g_n, f + \lambda g) \rightarrow 0$; hence $f + \lambda g$ is in the uniform closure of \mathcal{A} . This shows that $\overline{\mathcal{A}}$ is closed under addition and scalar multiplication.

For pointwise multiplication, we first make the following observation: there is a uniform (in n) constant M so that $\sup_n |f_n| < M$. Indeed, since $f_n \rightarrow_u f$, there is some N so that $\sup |f_n - f| < 1$ for all $n \geq N$; thus $|f_n(x) - f(x)| < 1$ and so using the reverse triangle inequality $|f_n(x)| \leq 1 + |f(x)|$, so in fact $\sup |f_n| < 1 + \sup |f|$. Thus, since all the f_n and f are by assumption bounded, we can take $M = \max\{\sup |f_1|, \sup |f_2|, \dots, \sup |f_N|, 1 + \sup |f|\}$. With that in hand, we estimate

$$|f_n g_n - f g| = |f_n g_n - f_n g + f_n g - f g| \leq |f_n g_n - f_n g| + |f_n g - f g| \leq |f_n| |g_n - g| + |f_n - f| |g|.$$

Fix a constant M as above so that $\sup |f_n| < M$ for all n ; then

$$\sup |f_n g_n - f g| \leq M \sup |g_n - g| + \sup |f_n - f| \sup |g|.$$

Since $g_n \rightarrow_u g$ and $f_n \rightarrow_u f$, and since M and $\sup |g| < \infty$, it follows that $\sup |f_n g_n - f g| \rightarrow 0$, and so $f_n g_n \rightarrow_u f g$. As \mathcal{A} is an algebra, $f_n g_n \in \mathcal{A}$ for each n , and this proves that the product $f g$ is in the uniform closure of \mathcal{A} . Thus $\overline{\mathcal{A}}$ is closed under pointwise multiplication, and we have completed the proof that $\overline{\mathcal{A}}$ is an algebra. \square

REMARK 10.51. The above proof applies equally well to complex valued functions, but we will only need the statement for real valued functions.

The real part of Theorem 10.49 is the statement that a real subalgebra of $C(X)$ that separates points and vanishes nowhere is dense in $C(X)$. In light of Lemma 10.50, we can rephrase this as follows: *If $\mathcal{A} \subseteq C(X)$ is a uniformly closed algebra that separates points and vanishes nowhere, then $\mathcal{A} = C(X)$.* This is the statement we will now aim to prove. We need two more lemmas first.

LEMMA 10.52. *Let X be a compact metric space, and let $\mathcal{A} \subseteq C(X)$ be a uniformly closed algebra. Then for each $f \in \mathcal{A}$, $|f| \in \mathcal{A}$ as well. Consequently, if $f, g \in \mathcal{A}$, then $\max\{f, g\}$ and $\min\{f, g\}$ are in \mathcal{A} as well.*

A brief note on notation: $\max\{f, g\}$ is the function $\max\{f, g\}(x) = \max\{f(x), g(x)\}$; similar for \min . These functions are continuous whenever f and g are; this is not hard to see directly, but it will also follow easily from their representations in the following proof.

PROOF. Let $M = \sup_{x \in X} |f(x)|$. Applying the Weierstrass approximation Theorem 10.47 to the function $a(x) = |x|$ on the compact interval $[-M, M]$, we can find a sequence of polynomials p_n such that $\sup_{|t| \leq M} ||t| - p_n(t)| \rightarrow 0$. Note that, since $|0| = 0$, it follows that $p_n(0) \rightarrow 0$; hence, replacing p_n with $p_n - p(0)$ if necessary, we may assume that $p_n(0) = 0$ for all n . It follows that $p_n \circ f$ is in \mathcal{A} : indeed, if $p_n(t) = a_1 t + a_2 t^2 + \dots + a_n t^n$, then $p_n \circ f = a_1 f + a_2 f^2 + \dots + a_n f^n$,

and since \mathcal{A} is an algebra containing f , this is also an element of \mathcal{A} . Moreover, for any $x \in X$, $f(x) \in [-M, M]$, and so $|p_n(f(x)) - |f(x)|| \leq \sup_{|t| \leq M} ||t| - p_n(t)|$; hence

$$\sup_{x \in X} |p_n(f(x)) - |f(x)|| \leq \sup_{|t| \leq M} ||t| - p_n(t)| \rightarrow 0.$$

So $|f|$ is the uniform limit of the function $p_n \circ f$, which are all in \mathcal{A} ; since \mathcal{A} is uniformly closed, it follows that $|f| \in \mathcal{A}$ as claimed.

The second statement follows from the fact that

$$\max\{f, g\} = \frac{f + g + |f - g|}{2} \quad \text{and} \quad \min\{f, g\} = \frac{f + g - |f - g|}{2}$$

which can be easily checked with a case analysis. \square

REMARK 10.53. An algebra of functions which is also closed under the max and min operations is sometimes called a *lattice*. Note that $|f| = \max\{f, 0\} - \min\{f, 0\}$, so being a lattice is equivalent to being closed under absolute value.

We need one final lemma about spaces of functions that vanish nowhere and separate points.

LEMMA 10.54. *Let \mathcal{A} be an algebra of functions that separates points and vanishes nowhere. Then for any $x \neq y$ in \mathcal{A} and any $a, b \in \mathbb{R}$, there is a function $f_{y \rightarrow b}^{x \rightarrow a} \in \mathcal{A}$ with $f_{y \rightarrow b}^{x \rightarrow a}(x) = a$ and $f_{y \rightarrow b}^{x \rightarrow a}(y) = b$.*

PROOF. Let $V = \{(f(x), f(y)) : f \in \mathcal{A}\} \subseteq \mathbb{R}^2$; our goal is to show that $V = \mathbb{R}^2$. First note that, since \mathcal{A} is a vector space, so is V : for any points \mathbf{v} and \mathbf{w} in V and any $\lambda \in \mathbb{R}$, there are functions $f, g \in \mathcal{A}$ with $\mathbf{v} = (f(x), f(y))$ and $\mathbf{w} = (g(x), g(y))$, and so

$$\mathbf{v} + \lambda \mathbf{w} = (f(x), f(y)) + \lambda(g(x), g(y)) = ((f + \lambda g)(x), (f + \lambda g)(y)).$$

As \mathcal{A} is a vector space, $f + \lambda g \in \mathcal{A}$, and so $\mathbf{v} + \lambda \mathbf{w} \in V$.

Thus $V \subseteq \mathbb{R}^2$ is a vector space. We assume for a contradiction that it is a strict subspace. Since \mathcal{A} vanishes nowhere, there is some function $f \in \mathcal{A}$ with $f(x) \neq 0$ and also some function g with $g(y) \neq 0$; hence V is not contained in either coordinate axis of \mathbb{R}^2 , and V is not the 0 space. If $V \neq \mathbb{R}^2$, it is therefore 1 dimensional, and is spanned by some vector $V = \text{span}\{(s, t)\}$ where $s, t \neq 0$. In particular, there is some function $f \in \mathcal{A}$ with $(f(x), f(y)) = (s, t)$. Since \mathcal{A} is an algebra, $f^2 \in \mathcal{A}$ as well, and notice that $V \ni (f^2(x), f^2(y)) = (s^2, t^2)$. Thus $(s^2, t^2) \in \text{span}\{(s, t)\}$. It follows that

$$0 = \det \begin{bmatrix} s & t \\ s^2 & t^2 \end{bmatrix} = st^2 - s^2t = st(s - t).$$

Since $st \neq 0$, it follows that $s = t$. But that means that $(f(x), f(y)) \in \text{span}\{(s, s)\}$ for all $f \in \mathcal{A}$; in particular, $f(x) = f(y)$ for all $f \in \mathcal{A}$. This contradicts the assumption that \mathcal{A} separates points. Hence, it must be that in fact $V = \mathbb{R}^2$, as claimed. \square

We are finally in a position to complete the proof of Theorem 10.49.

PROOF OF THE REAL PART OF THEOREM 10.49. Fix some $F \in C(X)$. Fix a point $x \in X$. For any point $y \in X$, by Lemma 10.54, there is a function $f_y^x = f_{y \rightarrow F(y)}^{x \rightarrow F(x)} \in \mathcal{A}$ that takes the same values on x, y as F . Fix $\epsilon > 0$. Since F and f_y^x are continuous at y , there is a neighborhood V_y of y so that $f_y^x(u) > F(y) - \frac{\epsilon}{2}$ and $F(u) < F(y) + \frac{\epsilon}{2}$ for all $u \in V_y$. In particular

$$f_y^x(u) > F(u) - \epsilon \quad \text{for all } u \in V_y.$$

Now, X is compact, and $\{V_y : y \in X\}$ covers X , so there is a finite subcover $V_{y_1}, V_{y_2}, \dots, V_{y_m}$. Now, define

$$f^x = \max\{f_{y_1}^x, f_{y_2}^x, \dots, f_{y_m}^x\}.$$

By (induction on) Lemma 10.52, $f^x \in \mathcal{A}$. By construction, $f_{y_j}^x > F - \epsilon$ on V_{y_j} , and hence $f^x \geq f_{y_j}^x > F - \epsilon$ on V_{y_j} for each j ; since the V_{y_j} cover X , it follows that

$$f^x(y) > F(y) - \epsilon \quad \text{for all } y \in X.$$

All the function f_y^x satisfy $f_y^x(x) = F(x)$, and so also $f^x(x) = F(x)$. Since f^x and F are continuous, there is a neighborhood V_x of x such that $f^x(u) < F(x) + \epsilon$ for all $u \in V_x$. Now, again, since $\{U_x : x \in X\}$ form a cover of X , and X is compact, there is a finite subcover $U_{x_1}, U_{x_2}, \dots, U_{x_k}$. Define

$$f = \min\{f^{x_1}, f^{x_2}, \dots, f^{x_k}\}.$$

Again by (induction on) Lemma 10.52, $f \in \mathcal{A}$. Since $f^{x_j} < F + \epsilon$ on U_{x_j} for each j , it follows that $f \leq f^{x_j} < F + \epsilon$ on U_j for all j . Since the U_j cover X , it follows that

$$f(y) < F(y) + \epsilon \quad \text{for all } y \in X. \quad (10.9)$$

But also $f(y) = \min\{f^{x_1}(y), \dots, f^{x_k}(y)\}$ and $f^{x_j}(y) > F(y) - \epsilon$ for each j , so we also have

$$f(y) > F(y) - \epsilon \quad \text{for all } y \in X. \quad (10.10)$$

Combining (10.9) and (10.10), we see that for each $\epsilon > 0$, there is a function $f \in \mathcal{A}$ with $|f - F| < \epsilon$. Taking this for each $\epsilon = \frac{1}{n}$, we can therefore construct a sequence of functions $f_n \in \mathcal{A}$ with $d_u(f_n, F) < \frac{1}{n}$; hence $f_n \rightarrow_u F$. But $f_n \in \mathcal{A}$, and \mathcal{A} is uniformly closed; thus the uniform limit F is in \mathcal{A} . As F was an arbitrary function in $C(X)$, we have thus proven that $C(X) \subseteq \mathcal{A}$; the reverse containment is clear (since \mathcal{A} is assumed to be an algebra of continuous functions), and the proof is complete. \square

EXAMPLE 10.55. Let $d \in \mathbb{N}$, and let $X \subset \mathbb{R}^d$ be a compact set. Then the set $\mathcal{P}(X)$ of polynomial functions (in d variables) on X is a real algebra (it is the smallest algebra containing the functions 1 and x_j for $1 \leq j \leq d$) that vanishes nowhere (since it contains all constants) and separates points (since it contains the functions x_1, x_2, \dots, x_d). Thus, by the Stone–Weierstrass approximation theorem, $\mathcal{P}(X)$ is uniformly dense in $C(X)$. This can be proved directly following a multivariate generalization of our proof of the Weierstrass Approximation Theorem 10.47 with an approximate identity sequence of d -variable polynomials. But it is not necessary to do this now that we have the Stone–Weierstrass theorem.

EXAMPLE 10.56. Let $[a, b]$ be a compact interval, and let $\kappa: [a, b] \rightarrow \mathbb{R}$ be any continuous one-to-one function. Then κ alone separates points: for any $x \neq y$ in $[a, b]$, $\kappa(x) \neq \kappa(y)$. Let $\mathcal{A} = \{p \circ \kappa : p \text{ is a polynomial}\}$. Then \mathcal{A} is an algebra (it is the smallest algebra containing 1 and κ) which separates points (since κ does) and vanishes nowhere (since it contains the constant function 1). Hence, by the Stone–Weierstrass theorem, \mathcal{A} is uniformly dense in $C[a, b]$. For example: the set of all functions of the form $a_0 + a_1 e^{\lambda x} + a_2 e^{2\lambda x} + \dots + a_m e^{m\lambda x}$ for some $\lambda \neq 0$ and $a_0, \dots, a_m \in \mathbb{R}$ can be used to uniformly approximate any continuous function on compact interval, because $x \mapsto e^{\lambda x}$ is one-to-one for any $\lambda \neq 0$.

EXAMPLE 10.57. Consider the set \mathcal{T} of *trigonometric polynomials* defined on $[-\pi, \pi)$:

$$\mathcal{T} = \left\{ x \mapsto \sum_{k=-m}^m a_k e^{ikx} : a_k \in \mathbb{C} \right\}$$

This set can be expressed as the set of complex polynomials in e^{ix} and e^{-ix} , and is therefore an algebra. It vanishes nowhere since it contains the constant function 1, and it separates points of $[-\pi, \pi)$ since the function $x \mapsto e^{ix} = (\cos x, \sin x)$ is one-to-one on this interval. This algebra is also self-adjoint: for $f(x) = \sum_{k=-m}^m a_k e^{ikx}$, $\overline{f(x)} = \sum_{k=-m}^m \overline{a_k} e^{-ikx} = \sum_{k=-m}^m \overline{a_{-k}} e^{ikx} \in \mathcal{T}$. Hence, by the Stone–Weierstrass approximation theorem, \mathcal{T} is uniformly dense in $C[-\pi, t]$ for any $t < \pi$. It is not, however, uniformly dense in $C[-\pi, \pi]$, since all elements $f \in \mathcal{T}$ satisfy $f(-\pi) = f(\pi)$. In fact, \mathcal{T} is uniformly dense in the continuous 2π -periodic functions on \mathbb{R} : those continuous functions f satisfying $f(x) = f(x + 2\pi)$ for all $x \in \mathbb{R}$.

Noting that the map $x \mapsto e^{ix}$ is continuous from $[-\pi, \pi]$ onto the unit circle \mathbb{S} , and is one-to-one except at the endpoints, we should really think of \mathcal{T} as an algebra of functions on \mathbb{S} , which is compact; the fact that \mathbb{S} “wraps around” takes care of the periodicity requirement automatically. It might be tempting to think that \mathcal{T} is then just the restriction of the algebra of 2-variable polynomials on the plane to the compact subset $\mathbb{S} \subset \mathbb{R}^2$, but this is not so: \mathcal{T} must contain *both positive and negative powers* of e^{ix} in order to be closed under conjugation. If this is omitted, and we consider the subalgebra \mathcal{T}_+ of those functions of the form $\sum_{k=0}^m a_k e^{ikx}$, the resulting algebra is *not* uniformly dense in $C(\mathbb{S})$, as Homework 10 asks you to show.

REMARK 10.58. The trigonometric polynomials of the last example are uniformly dense in $C(\mathbb{S})$. The uniform metric is stronger than the “ L^2 -metric” discussed in some homework exercises: if f_n is a sequence of trigonometric polynomials uniformly approximating $f \in C(\mathbb{S})$, we also have

$$d_2(f_n, f)^2 = \int_{-\pi}^{\pi} |f_n(x) - f(x)|^2 dx \rightarrow 0.$$

The nice thing about this metric is that it respects the structure of the trigonometric polynomials very nicely. It turns out that, in the d_2 metric, approximating functions using \mathcal{T} closely resembles power series: one can find an approximating sequence $f_n \in \mathcal{T}$ where f_n is of “degree” n (meaning containing e^{ikx} for $|k| \leq n$) and such that $f_n - f_{n-1}$ contains only scalar $e^{\pm nx}$. This leads the way to series representations of continuous (and more general) functions on \mathbb{S} using trigonometric polynomials: these are called *Fourier series*, and form a powerful underpinning of huge swaths of analysis, and applied science (e.g. signal processing). Unfortunately, we will not have time to do this subject justice in this course, and it will have to be left to a future course.

8. Lecture 20: June 2, 2016

We began these notes with the construction of \mathbb{R} : a complete metric space. One of our earliest results was that the field \mathbb{Q} of rational numbers is dense in \mathbb{R} . Later, we defined compactness (in a general metric space), and proved that compact sets are closed and bounded; and then we proved the Heine–Borel theorem: that the compact subsets of the metric space \mathbb{R} are precisely those that are closed and bounded.

In the last few lectures at the end of these notes, we’ve been considering the metric space $C(X)$ of continuous functions on a compact metric space X , equipped with the uniform metric d_u . In Corollary 10.10, we showed that this metric space is complete. (In fact, we showed a more general result: even if X is not compact, the metric space $C_b(X)$ of bounded continuous functions on X is complete in the d_u metric; when X is compact, $C_b(X) = C(X)$ by the extreme values theorem.) In the previous lecture, we proved the Stone–Weierstrass approximation theorem, which gives a characterization of a class of dense subsets of $C(X)$: like \mathbb{Q} in \mathbb{R} , they have an algebraic character.

The topic of this, our final lecture, is the analog of the Heine–Borel theorem for the metric space $C(X)$. Since $C(X)$ is a metric space, we know that compact subsets must be closed and bounded. To be clear: this means closed and bounded in terms of d_u ; for emphasis, we may refer to these properties as *uniformly closed* and *uniformly bounded*.

DEFINITION 10.59. *Let X be compact, and let $\mathcal{F} \subseteq C(X)$ be a subset. Say \mathcal{F} is **uniformly bounded** if there is a constant M so that, for all $f \in \mathcal{F}$, $\sup_{x \in X} |f(x)| \leq M$. Say that \mathcal{F} is **uniformly closed** if, given any sequence (f_n) in \mathcal{F} such that there is a function $f \in C(X)$ with $f_n \rightarrow_u f$, it follows that $f \in \mathcal{F}$.*

Again, these are just the usual metric space notions of “bounded” and “closed” with respect to the metric d_u . (Since $C(X)$ is complete, we can also characterize *uniformly closed* as equivalent to *uniformly Cauchy*: any Cauchy sequence in \mathcal{F} has a uniform limit in \mathcal{F} .)

Propositions 5.27 and 5.30 show that any compact set in $C(X)$ must be uniformly closed and uniformly bounded. However, unlike the the case of compact sets in \mathbb{R} (cf. the Heine–Borel theorem), these two conditions alone are *not* enough to guarantee compactness in $C(X)$. We can show this using an example similar in spirit to Example 5.32(2).

EXAMPLE 10.60. Let $\overline{B}_u(X) = \{f \in C(X) : \sup_{x \in X} |f(x)| \leq 1\}$ denote the closed unit ball in $C(X)$. This set is uniformly bounded by 1, and it is closed: if $f_n \in \overline{B}_u(X)$ and $f_n \rightarrow_u f$, then $f_n(x) \rightarrow f(x)$ for each x , and since $|f_n(x)| \leq 1$, it follows from the Squeeze Theorem that $|f(x)| \leq 1$ for each $x \in X$, so $\sup_{x \in X} |f(x)| \leq 1$ and $f \in \overline{B}_u(X)$.

However, for most X , $\overline{B}_u(X)$ is not compact. Take the example $X = [0, 1]$, and let $f_n(x) = \frac{1}{nx+1}$ as in Example 10.12. These functions are all in $\overline{B}_u(X)$ since $\sup_{x \in [0,1]} |f_n(x)| = f_n(0) = 1$. But (f_n) has *no uniformly convergent subsequence*. Indeed, if it did, say $f_{n_k} \rightarrow_u f$, then by Theorem 10.9, the uniform limit f would be continuous on $[0, 1]$. However, we have shown (cf. Example 10.12) that f_n converges pointwise to the function $\mathbb{1}_{\{0\}}$ on $[0, 1]$, and hence the subsequence f_{n_k} also converges pointwise to this discontinuous function. Since uniform convergence implies pointwise convergence, it follows that $f = \mathbb{1}_{\{0\}}$, which contradicts Theorem 10.9.

So in $C(X)$, even closed balls tend not to be compact; being closed and bounded is simply not enough to imply compactness. It turns out that there is one additional, natural condition which is needed: *equicontinuity*.

DEFINITION 10.61. Let $\mathcal{F} \subseteq C(X)$ be a collection of functions. We say \mathcal{F} is **equicontinuous** if, for each $\epsilon > 0$, there is a $\delta > 0$ so that for all $x, y \in X$ and all $f \in \mathcal{F}$, $|f(x) - f(y)| < \epsilon$ whenever $d(x, y) < \delta$.

Equicontinuity is the ultimate in uniformity of continuity. We upgraded the definition of continuity of a function to uniform continuity by insisting that the δ window for each ϵ tolerance can be chosen independently of the point x under consideration; this new notion of equicontinuity says that, not only that, we must be able to pick *the same* δ window for all functions $f \in \mathcal{F}$ simultaneously.

REMARK 10.62. Some textbooks refer to the property of Definition 10.61 as *uniform equicontinuity*. One might also define equicontinuity of a family $\mathcal{F} \subseteq C(X)$ at a point $x \in X$ to mean that, for each $\epsilon > 0$, there is a $\delta = \delta(\epsilon, x)$ so that, for all $y \in B_\delta(x)$ and all $f \in \mathcal{F}$, $|f(y) - f(x)| < \epsilon$; then \mathcal{F} could be called equicontinuous on X if it is equicontinuous at each point. If X is not a compact metric space, this is strictly weaker than our definition, as it allows δ to vary with x (with no necessary positive lower bound), while it is uniform in \mathcal{F} . However, since we are primarily concerned with compact X , it will turn out to make no difference: just as continuity implies uniform continuity on a compact set, so too pointwise equicontinuity implies uniform equicontinuity on a compact set.

EXAMPLE 10.63. (1) If $\mathcal{F} \subset C(X)$ is a finite collection of *uniformly* continuous functions, $\mathcal{F} = \{f_1, f_2, \dots, f_m\}$, then \mathcal{F} is equicontinuous: for each $\epsilon > 0$, and each $1 \leq j \leq m$, we can choose a $\delta_j > 0$ so that $|f_j(x) - f_j(y)| < \epsilon$ whenever $d(x, y) < \delta_j$; then $\delta \equiv \min\{\delta_1, \delta_2, \dots, \delta_m\} > 0$ works for all the functions in \mathcal{F} simultaneously.

(2) Fix $\alpha \in (0, 1]$, and consider the set $C^\alpha(X)$ of α -Hölder continuous functions on X : $f \in C^\alpha(X)$ iff $\|f\|_{C^\alpha(X)} < \infty$, where

$$\|f\|_{C^\alpha(X)} = \sup_{x \neq y} \frac{|f(x) - f(y)|}{d(x, y)^\alpha}.$$

(When $\alpha = 1$, this gives the Lipschitz functions on X ; this should not be confused with continuously differentiable functions, also denoted C^1 , which don't even make sense on a general metric space X .)

The set $C^\alpha(X) \subset C(X)$ is generally *not* equicontinuous. For example, taking $X = [0, 1]$, for each $\lambda > 0$, the function $f_\lambda(x) = \lambda x$ is Lipschitz with Lipschitz constant λ , hence also C^α for all $\alpha \in (0, 1]$. (The reader should compute that $\|f_\lambda\|_{C^\alpha[0,1]} = \lambda$ for all $\alpha \in (0, 1]$.) However, we can compute the largest δ that works for f_λ : in order for $\epsilon > |f_\lambda(x) - f_\lambda(y)| = \lambda|x - y|$, we must take $|x - y| < \epsilon/\lambda$, and so the largest δ for f_λ is $\delta = \epsilon/\lambda$. Since $\inf_{\lambda > 0} \epsilon/\lambda = 0$, there is no uniform δ that works for all $f \in C^\alpha[0, 1]$.

However, if we uniformly bound the size of the α -Hölder norm, the set becomes equicontinuous. For fixed $M \in (0, \infty)$, let $C_M^\alpha(X)$ denote the set of functions $f \in C^\alpha(X)$ with $\|f\|_{C^\alpha(X)} \leq M$. Then for $f \in C_M^\alpha(X)$, $|f(x) - f(y)| \leq M d(x, y)^\alpha$ for all x, y . Hence, for any $\epsilon > 0$, take $\delta = (\frac{\epsilon}{M})^{1/\alpha}$, which is uniform over X and $C_M^\alpha(X)$. For any $x, y \in X$ and $f \in C_M^\alpha(X)$, if $d(x, y) < \delta$ then

$$|f(x) - f(y)| \leq M d(x, y)^\alpha < M \delta^\alpha = M \left[\left(\frac{\epsilon}{M} \right)^{1/\alpha} \right]^\alpha = \epsilon.$$

This shows $C_M^\alpha(X)$ is equicontinuous.

- (3) The collection of functions $\{f_n\}$ on $[0, 1]$ with $f_n(x) = \frac{1}{nx+1}$ is not equicontinuous. It is possible to show this directly, by computing the modulus of continuity (the largest δ that works uniformly for a given $\epsilon > 0$) of f_n and showing that it tends to 0 as $n \rightarrow \infty$; but this is really quite complicated. In fact, we will soon prove our main theorem relating equicontinuity to compactness in $C(X)$, and from this and Example 10.60 it will follow that $\{f_n\}$ is not equicontinuous.

Our final main theorem is the following characterization of compact sets in $C(X)$.

THEOREM 10.64 (Arzelà–Ascoli). *Let X be a compact metric space, and consider the metric space $C(X)$ equipped with the uniform metric. A subset $\mathcal{K} \subset C(X)$ is compact if and only if \mathcal{K} is uniformly closed, uniformly bounded, and equicontinuous.*

This theorem was first proved in the context of the metric space $X = [0, 1]$ (or any compact interval in \mathbb{R}). The sufficiency of equicontinuity was proved by Ascoli in 1884; in 1895, Arzelà proved that equicontinuity is also a necessary condition, and gave the first complete proof that would be considered rigorous by today’s standards. Theorem 10.64 as stated, on a compact metric space, was first proved in this level of generality by Fréchet in 1906.

We begin by proving the “only if” direction: if $\mathcal{K} \subset C(X)$ is compact, then it is uniformly closed, uniformly bounded, and equicontinuous. We know the first two conditions must hold since compact sets are closed and bounded in *any* metric space; hence, we need only show equicontinuity.

PROOF OF (\implies) IN THEOREM 10.64. We argue by contradiction: suppose \mathcal{K} is compact, but not equicontinuous. This means there is some $\epsilon > 0$ so that, for all $\delta > 0$, we can find points $x, y \in X$ and some function $f \in \mathcal{K}$ such that $d(x, y) < \delta$ and yet $|f(x) - f(y)| \geq \epsilon$. We let $\delta = \frac{1}{n}$ for $n \in \mathbb{N}$, and select such x_n, y_n, f_n ; so (x_n) and (y_n) are sequences in X , and (f_n) is a sequence of functions in \mathcal{K} . Now, X is compact, so the sequences (x_n) and (y_n) have convergent subsequences; by choosing them successively, we can find a common subsequential index set $\{n_k\}$ so that $x_{n_k} \rightarrow x$ and $y_{n_k} \rightarrow y$ for some $x, y \in X$. Since $d(x_{n_k}, y_{n_k}) < \frac{1}{n_k} \rightarrow 0$, it follows that in fact $x = y$.

Now, \mathcal{K} is compact in $(C(X), d_u)$, and so the sequence (f_{n_k}) has a uniformly convergent subsequence $f_{n_{k_m}} \equiv g_m$; that is, there is some $g \in \mathcal{K}$ with $g_m \rightarrow_u g$. Let $u_m = x_{n_{k_m}}$ and $v_m = y_{n_{k_m}}$; then $u_m \rightarrow x$ and $v_m \rightarrow x$, and we have

$$|g_m(u_m) - g_m(v_m)| = |f_{n_{k_m}}(x_{n_{k_m}}) - f_{n_{k_m}}(y_{n_{k_m}})| \geq \epsilon \quad \text{for all } m.$$

It then follows that

$$\epsilon \leq |g_m(u_m) - g_m(v_m)| \leq |g_m(u_m) - g(x)| + |g(x) - g_m(v_m)| \quad \text{for all } m. \quad (10.11)$$

Since $g_m \rightarrow_u g$ and $u_m \rightarrow x$, it follows (from a homework problem) that $g_m(u_m) \rightarrow g(x)$. Similarly $g_m(v_m) \rightarrow g(x)$. Thus $|g_m(u_m) - g(x)| + |g(x) - g_m(v_m)| \rightarrow 0$; this contradicts (10.11), concluding the proof. \square

We now proceed to prove the “if” direction of Theorem 10.64. To do so, we need two general lemmas. The first, which could have been done back in Chapter 5, shows that compact metric spaces have countable dense subsets, generalizing the density of the rationals in any compact interval.

LEMMA 10.65. *Let X be a compact metric space. There exists a countable set $E \subseteq X$ such that $\bar{E} = X$.*

PROOF. For each $n \in \mathbb{N}$, consider the collection $\mathcal{U}_n = \{B_{1/n}(x) : x \in X\}$. Each \mathcal{U}_n is an open cover of X (as each x is contained in, for example, $B_{1/n}(x)$). Since X is compact, there is a finite subcover $B_{1/n}(x_n^1), B_{1/n}(x_n^2), \dots, B_{1/n}(x_n^{k_n})$.

Let $E = \{x_n^j : n \in \mathbb{N}, 1 \leq j \leq k_n\}$. Then E is countable: it is a countable union (over n) of the finite sets $\{x_n^1, x_n^2, \dots, x_n^{k_n}\}$. Now fix $\epsilon > 0$, and let $x \in X$. If $n \in \mathbb{N}$ with $\frac{1}{n} < \epsilon$, then since $\{B_{1/n}(x_j^n) : 1 \leq j \leq k_n\}$ cover X , there is some j with $x \in B_{1/n}(x_j^n)$. That is: $d(x, x_j^n) < \frac{1}{n} < \epsilon$. This shows that, for each $\epsilon > 0$ and $x \in X$, there is some point $t \in E$ ($t = x_j^n$) with $d(x, t) < \epsilon$. Thus E is dense in X . \square

The second required lemma is one of the several general results that go under the name "Cantor diagonalization".

LEMMA 10.66. *Let (f_n) be a sequence of real (or complex) valued functions on a countable set E . Suppose that (f_n) is pointwise bounded: for each $t \in E$, the sequence $(f_n(t))_{n \in \mathbb{N}}$ is bounded. Then there is a uniform (over E) increasing index set $\{n_k\}_{k \in \mathbb{N}}$ of positive integers such that the subsequence $(f_{n_k}(t))$ that converges for all $t \in E$.*

PROOF. The only tricky thing about this proof is making the notation readable. To do so, we slightly reformulate how we think about a subsequence. A sequence (a_n) in \mathbb{R} is a function $a: \mathbb{N} \rightarrow \mathbb{R}$, where we write $a(n) = a_n$. A subsequence of (a_n) is usually thought of as a new sequence (b_k) (meaning a function $b: \mathbb{N} \rightarrow \mathbb{R}$) which is defined by $b_k = a_{n_k}$ for some increasing infinite sequence of integers $n_1 < n_2 < n_3 < \dots$. We could just as well think of this subsequence as the restriction of the function a to the infinite subset $\Gamma = \{n_1, n_2, n_3, \dots\} \subseteq \mathbb{N}$. This is the perspective we take presently.

To begin, enumerate the countable set $E = \{t_1, t_2, t_3, \dots\}$. By assumption, $(f_n(t_1))$ is a bounded sequence in \mathbb{R} (or \mathbb{C}), so it follows from the Heine–Borel theorem that it has a convergent subsequence: this means there is an infinite subset $\Gamma_1 \subseteq \mathbb{N}$ so that the function $\Gamma_1 \ni n \mapsto f_n(t_1)$ is a convergent sequence. Now, consider the subsequence $\Gamma_1 \ni n \mapsto f_n(t_2)$. This is a subsequence of the full sequence $(f_n(t_2))$, which is bounded; hence it is a bounded sequence, and again by the Heine–Borel theorem, it has a convergent subsequence. This means there is an infinite subset $\Gamma_2 \subseteq \Gamma_1$ so that the function $\Gamma_2 \ni n \mapsto f_n(t_2)$ is a convergent sequence. Proceeding this way, by induction, we produce a nested sequence of infinite sets $\mathbb{N} \supseteq \Gamma_1 \supseteq \Gamma_2 \supseteq \Gamma_3 \supseteq \dots$ so that the function $\Gamma_k \ni n \mapsto f_n(t_k)$ is a convergent sequence.

Now, we define an increasing sequence of positive integers $\{n_k\}$ as follows: $n_1 = \min \Gamma_1$, $n_2 = \min(\Gamma_2 \cap (n_1, \infty))$, $n_3 = \min(\Gamma_3 \cap (n_2, \infty))$, and in general $n_k = \min(\Gamma_k \cap (n_{k-1}, \infty))$: i.e. n_k is the smallest element of Γ_k that is strictly bigger than n_{k-1} ; such an element exists since Γ_k is infinite (otherwise, n_{k-1} would be an upper bound for Γ_k , which means Γ_k would be finite). (Note: since $\Gamma_k \subseteq \Gamma_{k-1}$, $\min \Gamma_k \geq \min \Gamma_{k-1}$, but it is possible that the two have the same minimum, and we must have $n_k > n_{k-1}$ for each k .)

Now, by construction, $n_k \in \Gamma_k$, and since the sets are nested, this means $n_k \in \Gamma_m$ for $m \leq k$. For any $t \in E$, there is some m with $t = t_m$; thus, for all $k \geq m$, $n_k \in \Gamma_m$. This means that $(f_{n_k}(t_m))_{k=m}^\infty$ is a subsequence of $\Gamma_m \ni n \mapsto f_n(t_m)$, which is convergent; thus $(f_{n_k}(t_m))_{k=m}^\infty$ is convergent for each m , as desired. \square

REMARK 10.67. The only way pointwise boundedness of the (f_n) on E was used was to guarantee that each of the successive subsequences had a convergent subsequence. In fact, we

could have simply stated this weaker condition as the premise to the lemma: it holds whenever each sequence $(f_n(t))$ has the property that every subsequence of it has a *further* subsequence that is convergent. Since we will only need this lemma in the context of (uniformly bounded) functions, we are content with the stronger hypothesis.

We are now in a position to prove the forward direction of the Arzelà–Ascoli theorem.

PROOF OF (\Leftarrow) IN THEOREM 10.64. Let (f_n) be a sequence in \mathcal{K} . By Lemma 10.65, there is a countable dense subset $E \subseteq X$. Since \mathcal{K} is uniformly bounded, the functions f_n are uniformly bounded and hence pointwise bounded on E . Thus, by Lemma 10.66, there is a fixed subsequence (f_{n_k}) that converges on all of E . Let $g_k = f_{n_k}$; we will prove that (g_k) is uniformly Cauchy on X . It then follows, since $C(X)$ is complete, that $g_k \rightarrow_u g$ for some $g \in C(X)$, and since \mathcal{K} is assumed to be uniformly closed, $g \in \mathcal{K}$; this will thus conclude the proof that \mathcal{K} is (sequentially) compact.

Fix $\epsilon > 0$. By equicontinuity of \mathcal{K} , choose a uniform $\delta > 0$ so that for all $x, y \in X$ and $f \in \mathcal{K}$, $|f(x) - f(y)| < \frac{\epsilon}{3}$ whenever $d(x, y) < \delta$; in particular, for all k $|g_k(x) - g_k(y)| < \frac{\epsilon}{3}$ in this case. Now, consider the collection $\{B_\delta(t) : t \in E\}$. This is actually an open cover of all of X : since E is dense in X , for any $x \in X$ there is some $t \in E$ with $d(x, t) < \delta$, and thus $t \in B_\delta(t)$. Since X is compact, there are finitely many points $t_1, \dots, t_m \in E$ with $\{B_\delta(t_j) : 1 \leq j \leq m\}$ covering X . Since g_k converges on E , $(g_k(x))$ is Cauchy for each $x \in E$, so there are N_1, N_2, \dots, N_m so that $|g_k(x_j) - g_\ell(x_j)| < \frac{\epsilon}{3}$ for $k, \ell \geq N_j$; take $N = \max\{N_1, \dots, N_m\}$.

Now, let $x \in X$, and let $k, \ell \geq N$. There is some $j \in \{1, \dots, m\}$ with $d(x_j, x) < \delta$, and it follows that $|g_k(x) - g_k(x_j)| < \frac{\epsilon}{3}$ for all k . Hence

$$\begin{aligned} |g_k(x) - g_\ell(x)| &\leq |g_k(x) - g_k(x_j) + g_k(x_j) - g_\ell(x_j) + g_\ell(x_j) - g_\ell(x)| \\ &\leq |g_k(x) - g_k(x_j)| + |g_k(x_j) - g_\ell(x_j)| + |g_\ell(x_j) - g_\ell(x)| \\ &< \frac{\epsilon}{3} + |g_k(x_j) - g_\ell(x_j)| + \frac{\epsilon}{3} \end{aligned}$$

for all k, ℓ . The middle term is also $< \frac{\epsilon}{3}$ whenever $k, \ell \geq N$. This is true for all x , and so in fact $\sup_x |g_k(x) - g_\ell(x)| < \epsilon$ whenever $k, \ell \geq N$. This shows that (g_k) is uniformly Cauchy, completing the proof. \square

As a side note: the Arzelà–Ascoli theorem is often stated with weaker conditions: \mathcal{K} is compact in $C(X)$ iff it is equicontinuous, uniformly closed, and *pointwise* bounded: the last condition means that, for each x , the set of numbers $\{f(x) : f \in \mathcal{K}\}$ is bounded. (In other words: there is a function $\phi : X \rightarrow [0, \infty)$, not necessarily bounded, so that $|f(x)| \leq \phi(x)$ for each $x \in X$ and each $f \in \mathcal{K}$.) In fact, in the context of an equicontinuous family of functions on a compact metric space, this already implies uniform boundedness.

PROPOSITION 10.68. *Let $\mathcal{F} \subseteq C(X)$ be an equicontinuous family of functions that is pointwise bounded. Then in fact \mathcal{F} is uniformly bounded.*

PROOF. Using equicontinuity with $\epsilon = 1$, let $\delta > 0$ be chosen uniformly so that for all $f \in \mathcal{F}$ and all $x, y \in X$, $d(x, y) < \delta$ implies that $|f(x) - f(y)| < 1$. Now, the collection $\{B_\delta(x) : x \in X\}$ is an open cover of X , and hence contains a finite subcover: there is a finite set of points x_1, x_2, \dots, x_m with $\bigcup_{j=1}^m B_\delta(x_j) = X$. Now, since \mathcal{F} is pointwise bounded, there are constants $C_1, C_2, \dots, C_m < \infty$ so that $|f(x_j)| \leq C_j$ for all $f \in \mathcal{F}$. Let $C = \max\{C_1, \dots, C_m\}$. Then for any $x \in X$, there is some j with $d(x_j, x) < \delta$, and so by equicontinuity $|f(x) - f(x_j)| < 1$ for all $f \in \mathcal{F}$. But this means that $|f(x)| \leq |f(x_j)| + 1 \leq C_j + 1 \leq C + 1$ for all $x \in X$ and all $f \in \mathcal{F}$; thus \mathcal{F} is uniformly bounded by $C + 1$. \square

EXAMPLE 10.69. Consider again the functions $f_n(x) = \frac{1}{nx+1}$ on $[0, 1]$. They are continuous and uniformly bounded. But, as shown in Example 10.60, there is no uniformly convergent subsequence of (f_n) . It then follows that the set $\{f_n\}$ is uniformly closed: for let (g_k) be a sequence in the set $\{f_n\}$ that converges uniformly to some $g \in C(X)$. Since $g_k \in \{f_n\}$, there is some n_k with $g_k = f_{n_k}$. Suppose that $\limsup_{k \rightarrow \infty} n_k = \infty$; then there is a subsequence of the n_k that tends to ∞ , which means there is a subsequence of f_n that converges uniformly to g – a contradiction. Thus $\limsup_{k \rightarrow \infty} n_k < \infty$; but since the n_k are integers, this means that the set $\{n_k : k \in \mathbb{N}\}$ is finite. As the uniform distance between any two elements of f_n is strictly positive, it follows that if $g_k \rightarrow g$, we must have n_k eventually constant, in which case $g \in \{f_n\}$. Thus $\{f_n\}$ contains all its uniform limits, and so it is uniformly closed.

Hence, the set $\{f_n\}$ is uniformly bounded and uniformly closed, and yet there is a sequence in this set (namely the full sequence (f_n)) that possesses no uniformly convergent subsequences. By the Arzelà–Ascoli theorem, it follows that $\{f_n\}$ is not equicontinuous.

REMARK 10.70. The argument in the preceding example actually shows in great generality that, to prove a set of uniformly bounded functions is not equicontinuous, it suffices to show that it contains a sequence with no convergent subsequences. That sequence will then be uniformly closed and uniformly bounded, and ergo not equicontinuous by the Arzelà–Ascoli theorem.

EXAMPLE 10.71. Let X be a compact metric space, let $0 < \alpha \leq 1$ and $M > 0$, and let $C_M^\alpha(X)$ be the space of α -Hölder continuous functions with Hölder constant $\leq M$, cf. Example 10.63(2). Since X is compact, it has finite diameter; thus, fixing some $x_0 \in X$, we have for all $x \in X$

$$|f(x)| \leq |f(x_0)| + |f(x) - f(x_0)| \leq |f(x_0)| + Md(x, x_0)^\alpha \leq |f(x_0)| + M \operatorname{diam}(X)^\alpha.$$

This shows that $C_M^\alpha(X)$ is uniformly bounded. It is also uniformly closed: if $f_n \in C_M^\alpha(X)$ and $f_n \rightarrow_u f$ (or even if $f_n \rightarrow f$ just pointwise), we have $|f(x) - f(y)| = \lim_{n \rightarrow \infty} |f_n(x) - f_n(y)| \leq Md(x, y)^\alpha$ by the Squeeze Theorem. Thus $f \in C_M^\alpha(X)$. We also showed, in Example 10.63(2), that $C_M^\alpha(X)$ is equicontinuous. Hence, by the Arzelà–Ascoli theorem, this set is compact.

In particular, this shows that if (f_n) is any sequence of Hölder continuous functions with uniformly bounded Hölder constants, then (f_n) possesses a uniformly convergent subsequence.