Notes on Vector Calculus

Dinakar Ramakrishnan

March, 2010

Chapter 1

Subsets of Euclidean space, vector fields, and continuity

Introduction

The aims of this course are the following:

(i) Extend the main results of one-variable Calculus to higher dimensions

(ii) Explore new phenomena which are non-existent in the one-dimensional case

Regarding the first aim, a basic step will be to define notions of continuity and differentiability in higher dimensions. These are not as intuitive as in the one-dimensional case. For example, given a function $f : \mathbb{R} \to \mathbb{R}$, we can say that f is continuous if one can draw its graph Γ_f without lifting the pen (resp. chalk) off the paper (resp. blackboard). For any $n \geq 1$, we can still define the graph of a function (here called a *scalar field*) $f : \mathbb{R}^n \to \mathbb{R}$ to be

$$\Gamma(f) := \{(x, y) \in \mathbb{R}^n \times \mathbb{R} \,|\, y = f(x)\},$$

where x denotes the vector $(x_1, \ldots, x_n)$ in $\mathbb{R}^n$. Since $\mathbb{R}^n \times \mathbb{R}$ is just $\mathbb{R}^{n+1}$, we can think of $\Gamma(f)$ as a subset of the $(n + 1)$-dimensional space. But the graph will be n-dimensional, which is hard (for non-constant f)to form a picture of, except possibly for $n = 2$; even then it cannot be drawn on a plane like a blackboard or a sheet of paper. So one needs to define basic notions such as continuity by a more formal method. It will be beneficial to think of a lot of examples in dimension 2, where one has some chance of forming a mental picture.

Integration is also subtle. One is able to integrate *nice* functions f on closed rectangular boxes R, which naturally generalize the closed interval $[a, b]$ in $\mathbb{R}$, and when there is spherical symmetry, also over closed balls in $\mathbb{R}^n$. Here f being nice means that f is bounded on R and continuous outside a *negligible set*. But it is problematic to define integrals of even continuous functions over arbitrary subsets Y of $\mathbb{R}^n$, even when they are *bounded*, i.e., can be enclosed in a rectangular box. However, when Y is compact, i.e., closed and bounded, one can integrate continuous functions f over it, at least when f vanishes on the *boundary* of Y.

The second aim is more subtle than the first. Already in the plane, one is interested in *line integrals*, i.e., integrals over (nice) curves C, of *vector fields*, i.e., vectors of scalar fields, and we will be interested in knowing when the integrals depend only on the *beginning and end points* of the curve. This leads to the notion of *conservative fields*, which is very important also for various other subjects like Physics and Electrical Engineering. Somehow the point here is to not blindly compute such integrals, but to exploit the (beautiful) *geometry* of the situation.

This chapter is concerned with defining the basic structures which will be brought to bear in the succeeding chapters. We start by reviewing the *real numbers*.

1.1 Construction and properties of real numbers

This section intends to give some brief background for some of the basic facts about real numbers which we will use in this chapter and in some later ones.

Denote by $\mathbb{Z}$ the set of integers $\{0, \pm 1, \pm 2, \ldots\}$ and by $\mathbb{Q}$ the set of rational numbers $\{\frac{a}{b} | a, b \in \mathbb{Z}, b \neq 0\}$. As usual we identify $\frac{a}{b}$ with $\frac{ma}{mb}$ for any integer $m \neq 0$. To be precise, the rational numbers are equivalence classes of ordered pairs (a, b) of integers, $b \neq 0$, with (a, b) being considered equivalent to (c, d) iff $ad = bc$. The rational numbers can be added, subtracted and multiplied, just like the integers, but in addition we can also divide any rational number x by another (non-zero) y. The properties make $\mathbb{Q}$ into a *field* (as one says in *Algebra*), not to be confused (at all) with a *vector field* which we will introduce below. One defines the absolute value of any rational number x as $\mathrm{sgn}(x)x$, where $\mathrm{sgn}(x)$ denotes the sign of x. Then it is easy to see that $|xy| = |x||y|$ and $|x+y| \leq |x|+|y|$ (triangle inequality). There are other absolute values one can define on $\mathbb{Q}$, and they satisfy a stronger inequality than the triangle inequality. For this reason, some call the one above the *archimedean* absolute value.

The real numbers are not so easy to comprehend, not withstanding the fact that we have been using them at will and with ease, and their construction is quite subtle. It was a clever ploy on the part of mathematicians to call them real numbers as it makes people feel that they are *real* and so easy to understand. Some irrational numbers do come up in geometry, like the ubiquitous π, which is the area of a unit circle, and the quadratic irrational $\sqrt{2}$, which is the length of the hypotenuse of a right triangle with two unit sides. However, many irrational real numbers have no meaning except as limits of nice sequences of rational numbers or as points on a continuum.

What constitutes a nice sequence? The basic criterion has to be that the sequence looks like it has a chance of converging, though we should be able to see that without knowing anything about the limit. The precise condition below was introduced by the nineteenth century French mathematician Cauchy, whence the term *Cauchy sequence*.

Definition. *A sequence $\{x_1, x_2, \ldots, x_n, \ldots\}$ of rational numbers is Cauchy iff we can find, for every $\epsilon > 0$ in $\mathbb{Q}$, a positive integer N such that $|x_n - x_m| < \epsilon$, for all $n, m \geq N$.*

Simply put, a sequence is Cauchy iff the terms eventually bunch up around each other. This behavior is clearly necessary to have a limit.

Possibly the simplest (non-constant) Cauchy sequence of rational numbers is $\{\frac{1}{n}|n \geq 1\}$. This sequence has a limit in $\mathbb{Q}$, namely 0. (Check it!) An example of a sequence which is not Cauchy is given by $\{x_n\}$, with $x_n = \sum_{k=1}^{n} \frac{1}{k}$. It is not hard to see that this sequence diverges; some would say that the limit is $+\infty$, which lies outside $\mathbb{Q}$.

There are many sequences of rational numbers which are Cauchy but do *not* have a limit in $\mathbb{Q}$. Two such examples $\{x_n|n \geq 1\}$ are given by the following: (i) $x_n = \sum_{k=1}^{n} \frac{1}{k^2}$, and (ii) $x_n = \sum_{k=1}^{n} \frac{1}{k!}$. The first sequence converges to $\pi^2/6$ and the second to e, both of which are not rational (first proved by a mathematician named Lambert around 1766). In fact, these numbers are even *transcendental*, which means they are not roots of polynomials $f(X) = a_0 + a_1X + \cdots + a_nX^n$ with $a_0, a_1, \ldots, a_n \in \mathbb{Q}$. Numbers like $\sqrt{2}$ and $i = \sqrt{-1}$ are *algebraic numbers*, though irrational, and are not transcendental. Of course i is not a "real number".

Take your favorite decimal expansion such as the one of $\pi = 3.1415926.....$ This can be viewed as a Cauchy sequence $\{3, 31/10, 314/100, 3141/1000, ...\}$; you may check that such a sequence is Cauchy no matter what the digits are! Again, for "most" expansions you can come up with, the limit is not in $\mathbb{Q}$. Recall that a decimal expansion represents (converges to) a rational number if and only if it is periodic after some point.

Essentially, the construction of $\mathbb{R}$ consists of formally adjoining to $\mathbb{Q}$ the missing limits of Cauchy sequences of rational numbers. How can we do this in a logically satisfactory manner? Since every Cauchy sequence should give a real number and every real number should arise from a Cauchy sequence, we start with the set X of all Cauchy sequences of rational numbers. On this set we can even define addition and multiplication by $\{x_1, x_2, ..\} + \{y_1, y_2, ...\} := \{x_1 + y_1, x_2 + y_2, ..\}$ etc. and also division if the denominator sequence stays away from zero (there is something to be checked here, namely that for example $\{x_n + y_n\}$ is again Cauchy but this is not so hard). It seems that X with these operations is already something like the real numbers but there is a problem: Two sequences might have the same limit, think of the zero sequence $\{0, 0, 0..\}$ and the reciprocal sequence $\{1/n | n \geq 1\}$, and such sequences should give the same real number, which is 0 in this case. So the last step is to introduce on X an *equivalence relation*: Declare $x = \{x_n\}, y = \{y_n\}$ in X to be equivalent if for any $\epsilon > 0$ in $\mathbb{Q}$, there exists $N > 0$ such that $|x_n - y_n| < \epsilon$ for all $n > N$ (Check that this is indeed an equivalence relation). The equivalence classes of sequences in X are declared to constitute the set $\mathbb{R}$ of all "real numbers". The rational numbers naturally form a subset of $\mathbb{R}$ by viewing $q \in \mathbb{Q}$ as the *class* of the *constant sequence* $\{q, q, \ldots, q, \ldots\}$; note that in this class we have many other sequences such as $\{q + 1/n\}$, $\{q + 2^{-n}\}$, $\{q + 1/n!\}$). Besides $\mathbb{Q}$ we obtain a lot of new numbers to play with. The real number represented by the sequence (ii) above, for example, is called e. When we say we have an $x \in \mathbb{R}$ we think of *some* Cauchy sequence $\{x_1, x_2, ..\}$ representing x, for example its decimal expansion. But note that our definition of $\mathbb{R}$ is in no way tied to the base ten, only to the rational numbers. Now one has to check that the addition, multiplication and division defined above pass to the equivalence classes, i.e. to $\mathbb{R}$. This is a doable exercise. One can also introduce an order on X which passes to equivalences classes: $\{x_n\} < \{y_n\}$ if there is some N so that $x_n < y_n$ for all $n > N$. So the notion of positivity for rational numbers carries over to the reals. One has the triangle inequality

$$|x + y| \leq |x| + |y|, \quad \text{forall} \quad x, y \in \mathbb{R}.$$

where $|x| = x$, resp. $-x$ if $x \geq 0$, resp. $x < 0$.

Here are *the* key facts about the real numbers on which much of Calculus is based.

Theorem 1 *(a) (Completeness of $\mathbb{R}$) Every Cauchy sequence of* real *numbers has a limit in $\mathbb{R}$.*

(b) (Density of the rationals) Every real number is the limit of a Cauchy sequence of rational numbers.

More precisely, part (a) says that given $\{x_n\} \subset \mathbb{R}$ so that $\forall\ \epsilon > 0,\ \exists\ N$ such that $\forall\ n, m > N, \quad |x_n - x_m| < \epsilon$, then there exists $x \in \mathbb{R}$ so that $\forall\ \epsilon > 0,\ \exists\ N$ such that $\forall\ n > N, \quad |x - x_n| < \epsilon$. In other words, if we repeat the completion process by which we obtained $\mathbb{R}$ from $\mathbb{Q}$ (i.e. start with the set of Cauchy sequences of *real* numbers and pass to equivalence classes) we end up with the real numbers again.

Part (b) says that given any real number x and an $\epsilon > 0$, we can find infinitely many rational numbers in the interval $(x - \epsilon, x + \epsilon)$. In particular, we have the following **very useful fact**:

Given any pair x, y of real numbers, say with $x < y$, we can find a rational number z such that $x < z < y$, regardless of how small $x - y$ is.

Proof of Theorem 1. Part (b) holds by construction, and it suffices to prove part (a). If $\{x_n\} \subset \mathbb{R}$ is a Cauchy sequence we represent each x_n as a Cauchy sequence of *rational* numbers $x_{n,i}$, say. We can find $f(n) \in \mathbb{N}$ so that $|x_{n,i} - x_{n,j}| < 2^{-n}$ for $i, j \geq f(n)$ because $x_{n,i}$ is Cauchy for fixed n. Then the sequence $x_{n,f(n)} \in \mathbb{Q}$ is Cauchy. Indeed, given ϵ make n, m large enough so that $2^{-n} < \epsilon/3$, $2^{-m} < \epsilon/3$ and $|x_n - x_m| < \epsilon/3$. Now unravel the meaning of this last statement. It means that for k large enough we have $|x_{n,k} - x_{m,k}| < \epsilon/3$. (Note that k depends on m, n and should perhaps be denoted $k(m, n)$.) Enlarge k further, if necessary, to make it also bigger than $f(n)$ and $f(m)$. Then we have

$$|x_{n,f(n)} - x_{m,f(m)}| < |x_{n,f(n)} - x_{n,k}| + |x_{n,k} - x_{m,k}| + |x_{m,k} - x_{m,f(m)}| < 2^{-n} + \epsilon/3 + 2^{-m} < \epsilon$$

for large enough n, m. The real number x represented by the Cauchy sequence $x_{n,f(n)}$ is the limit of the sequence x_n. To see this, given ϵ take n large so that $2^{-n} < \epsilon/2$ and $|x_{k,f(k)} - x_{n,f(n)}| < \epsilon/2$ for $k \geq n$. If we also have $k > f(n)$ then

$$|x_{k,f(k)} - x_{n,k}| \leq |x_{k,f(k)} - x_{n,f(n)}| + |x_{n,f(n)} - x_{n,k}| < \epsilon.$$

By definition this means that $|x - x_n| < \epsilon$.

Done.

The completeness of $\mathbb{R}$ has a number of different but equivalent formulations. Here is one of them. Call a subset A of real numbers *bounded from above* if there is some $y \in \mathbb{R}$ such that $x \leq y$ for every x in A. Such a y is called an *upper bound* for A. We are interested in knowing if there is a least upper bound. For example, when A is the interval $(a, 5)$ and $a < 5$, A has a least upper bound, namely the number 5. On the other hand, if $a \geq 5$, the set A is empty, and it has no least upper bound. The **least upper bound** is also called the **supremum**, when it exists, and denoted *lub* or *sup*. It is easy to see that the least upper bound has to be unique if it exists.

Theorem 2 *Let A be a* **non-empty** *subset of $\mathbb{R}$ which is bounded from above. Then A has a least upper bound.*

Proof. As we have already mentioned, this property is actually equivalent to the completeness of $\mathbb{R}$. Let's first deduce it from completeness. Since A is bounded from above, there exists, by definition, some upper bound $b_1 \in \mathbb{R}$. Pick any $a_1 \in A$, which exists because A is non-empty. Consider the midpoint $z_1 = (a_1 + b_1)/2$ between a_1 and b_1. If z_1 is an upper bound of A, put $b_2 = z_1$ and $a_2 = a_1$. Otherwise, there is an $a_2 \in A$ so that $a_2 > z_1$ and we put $b_2 = b_1$. In both cases we have achieved the inequality $|b_2 - a_2| \leq \frac{1}{2}|b_1 - a_1|$. Next consider the mid-point z_2 between a_2 and b_2 and define a_3, b_3 by the same procedure. Continuing thus "ad infinitum" (to coin a favorite phrase of Fermat), we arrive at *two* sequences of real numbers, namely $\{a_n | n \geq 1\}$ and $\{b_n | n \geq 1\}$. If we put $c = |x_1 - y_1|$, it is easy to see that by construction,

$$|a_n - b_n| \ \leq \ \frac{c}{2^{n-1}},$$

and that

$$|a_n - a_{n+1}| \ \leq \ |a_n - b_n| \ \leq \ \frac{c}{2^{n-1}}$$

and that the same holds for $|b_n - b_{n+1}|$. Consequently, both of these sequences are Cauchy and have the same limit $z \in \mathbb{R}$, say. Now we claim that z is the *lub* of A. Indeed, since z is the limit of the upper bounds b_n, it must also be an upper bound. On the other hand, since it is also the limit of the numbers a_n lying in A, any smaller number than z cannot be an upper bound of A. Done.

We only sketch of the proof of the converse. Given a Cauchy sequence $\{x_n\}$ in $\mathbb{R}$ consider the set

$$A = \{y \in \mathbb{R} | \text{ the set } \{n | x_n \leq y\} \text{ is finite, possibly empty}\}.$$

Then A is nonempty because there is N so that $|x_n - x_m| < 1$ for $n, m \geq N$. If $y = x_N - 2$ then $\{n | x_n \leq y\} \subseteq \{1, 2, ..., N\}$ is finite, so $y \in A$. One can then show that a least upper bound of A, which will in fact be unique, is also a limit of the sequence $\{x_n\}$.

QED

Finally a word about the **geometric representation of real numbers**. Real numbers x can be represented in a one-to-one fashion by the points $P = P(x)$ on a line, called the *real line* such that the following hold: (i) if $x < y$, $P(y)$ is situated strictly to the right of $P(x)$; and (ii) $|x - y|$ is the distance between $P(x)$ and $P(y)$. In particular, for any pair of real numbers x, y, the mid-point between $P(x)$ and $P(y)$, which one can find by a ruler and compass, corresponds to a unique real number z such that $|x - y| = 2|x - z| = 2|y - z|$. It is customary to identify the numbers with the corresponding points, and simply write x to denote both. Note that the notions of line and distance here are classical; in modern, set-theory-based mathematics one simply defines a line as some set of points that is in one-to-one correspondence with the real numbers.

1.2 The norm in $\mathbb{R}^n$

Consider the n-dimensional Euclidean space

$$\mathbb{R}^n = \{x = (x_1, x_2, \cdots, x_n) \mid x_i \in \mathbb{R}, \forall i\},$$

equipped with the *inner product* (also called the *scalar product*)

$$\langle x, y \rangle \ = \ \sum_{j=1}^{n} x_j y_j,$$

for all $x = (x_1, \ldots, x_n)$ and $y = (y_1, \ldots, y_n)$ in $\mathbb{R}^n$, and the *norm* (or *length*) given by

$$||x|| \ = \ \langle x, x \rangle^{1/2} \ \geq \ 0.$$

Note that $\langle x, y \rangle$ is linear in each variable, and that $||x|| = 0$ iff $x = 0$.

Basic Properties:

(i) $||cx|| = |c| \, ||x||$, for all $c \in \mathbb{R}$ and $x \in \mathbb{R}^n$.

(ii) (triangle inequality) $||x + y|| \leq ||x|| + ||y||$, for all $x, y \in \mathbb{R}^n$.

(iii) (Cauchy-Schwarz inequality) $|\langle x, y \rangle| \leq ||x|| ||y||$, for all x, y in $\mathbb{R}^n$.

Proof. Part (i) follows from the definition. We claim that part (ii) follows from part (iii). Indeed, by the bilinearity and symmetry of $\langle ., . \rangle$,

$$\langle x + y, x + y \rangle \ = \ \langle x, x \rangle + \langle y, y \rangle + 2\langle x, y \rangle.$$

By the Cauchy-Schwartz inequality, $|\langle x, y \rangle| \leq ||x|| ||y||$, whence the claim.

It remains to prove part (iii). Since the assertion is trivial if x or y is zero, we may assume that $x \neq 0$ and $y \neq 0$. If $w = \alpha x + \beta y$, with $\alpha, \beta \in \mathbb{R}$, we have

$$(*) \qquad 0 \ \leq \ \langle w, w \rangle \ = \ \alpha^2 \langle x, x \rangle + 2\alpha\beta \langle x, y \rangle + \beta^2 \langle y, y \rangle.$$

Since this holds for all α, β, we are free to choose them. Put

$$\alpha = \langle y, y \rangle, \ \beta = -\langle x, y \rangle.$$

Dividing $(*)$ by α, we obtain the inequality

$$0 \le \langle x, x\rangle\langle y, y\rangle - \langle x, y\rangle^2.$$

Done.

The *scalar product* $\langle ., .\rangle$ and norm $||.||$ can both be defined on $\mathbb{C}^n$ as well. For this we set

$$\langle x, y\rangle = \sum_{j=1}^{n} x_j \overline{y}_j,$$

where $\overline{z}$ denotes, for any $z \in \mathbb{C}$, the *complex conjugate* of z. Here $\langle x, y\rangle$ is linear in the first variable, but *conjugate linear* in the second, i.e., $\langle \alpha x, y\rangle = \alpha\langle x, y\rangle$, while $\langle x, \alpha y\rangle = \overline{\alpha}\langle x, y\rangle$, for any complex scalar α. In French such a product will be said to be *sesquilinear* (meaning "one and a half" linear). In any case, note that $\langle x, x\rangle$ is a non-negative real number, which is positive iff $x \neq 0$. So again it makes good sense to say that the *norm* (or *length*) of any $x \in \mathbb{C}^n$ to be $||x|| = \sqrt{\langle x, x\rangle}$. It is a routine exercise to verify that the *basic properties* (i), (ii), (iii) above continue to hold in this case.

We may define a sequence of vectors $v_1, v_2, \ldots, v_m, \ldots$ in $\mathbb{R}^n$ a Cauchy sequence iff for every $\varepsilon > 0$, we can find an $N > 0$ such that for all $m, r > N$, $||v_m - v_r|| < \varepsilon$. In other words, the vectors in the sequence eventually become bunched up together, as tightly as one requires.

Theorem 3 $\mathbb{R}^n$ *is complete with respect to the norm* $||\ ||$.

Idea of Proof. The inequality

$$|x_i| = \sqrt{x_i^2} \le \sqrt{x_1^2 + \cdots + x_n^2} = ||x||$$

shows that the components of any $||\ ||$-Cauchy sequence in $\mathbb{R}^n$ are ordinary Cauchy sequences in $\mathbb{R}$. Hence we are reduced to part (a) of Theorem 1.

1.3 Basic Open Sets in $\mathbb{R}^n$

There are (at least) two types of basic "open" sets in $\mathbb{R}^n$, one *round*, using open balls, and the other *flat*, using rectangular boxes.

For each $a \in \mathbb{R}^n$ and $r > 0$ set

$$B_a(r) = \{x \in \mathbb{R}^n \mid ||x - a|| < r\},$$

and call it the *open ball* of radius r and center a; and

$$\bar{B}_a(r) = \{x \in \mathbb{R}^n \mid ||x - a|| \leq r\},$$

the *closed ball* of radius r and center a. One also defines the *sphere*

$$S_a(r) = \{x \in \mathbb{R}^n | \, ||x - a|| = r\}.$$

A *closed rectangular box* in $\mathbb{R}^n$ is of the form

$$[a, b] \; = \; [a_1, b_1] \times [a_2, b_2] \times \cdots \times [a_n, b_n] \; = \; \{x \in \mathbb{R}^n \mid \; x_i \in [a_i, b_i], \, \forall i\},$$

for some $a = (a_1, \cdots, a_n) \in \mathbb{R}^n$ and $b = (b_1, \cdots, b_n) \in \mathbb{R}^n$.

An *open rectangular box* is of the form

$$(a, b) \; = \; (a_1, b_1) \times (a_2, b_2) \times \cdots \times (a_n, b_n) \; = \{x \in \mathbb{R}^n \mid \forall i \, x_i \in (a_i, b_i)\}.$$

Definition. A basic open set in $\mathbb{R}^{2n}$ is (either) an open ball or an open rectangular box.

You can use either or both. One gets the same answers for our problems below. In this vein, observe that every open ball contains an open rectangular box, and conversely.

Remark. It is important to note that given any pair of basic open sets V_1, V_2, we can find a nonempty basic open set W contained in their intersection $V_1 \cap V_2$ if this intersection is non-empty.

1.4 Open and closed sets

Given any subset X of $\mathbb{R}^n$, let us denote by X^c the complement $\mathbb{R}^n - X$ in $\mathbb{R}^n$. Clearly, the complement of the empty set $\emptyset$ is all of $\mathbb{R}^n$.

Let A be a subset of $\mathbb{R}^n$ and let y be a point in $\mathbb{R}^n$. Then there are three possibilities for y relative to A.

(IP) There exists a basic open set U containing y such that $U \subseteq A$.

(EP) There exists a basic open set U centered at y which lies completely in the complement of A. i.e., in $\mathbb{R}^n - A$.

(BP) Every basic open set centered at y meets both A and A^c.

In case (IP), y is called an **interior point of** A. In case (EP), y is called an **exterior point of** A. In case (BP), y is called a **boundary point of** A. Note that in case (IP) $y \in A$, in case (EP) $y \notin A$, and in case (BP) y may or may not belong to A.)

Definition. *A set A in $\mathbb{R}^n$ is open if and only if every point of A is an interior point.*

Explicitly, this says: "Given any $z \in A$, we can find a basic open set U containing z such that $U \subseteq A$."

Definition. $A \subseteq \mathbb{R}^n$ is closed if its complement is open.

Lemma 1 *A subset A of $\mathbb{R}^n$ is closed iff it contains all of its boundary points.*

Proof. Let y be a boundary point of A. Suppose y is not in A. Then it belongs to A^c, which is open. So, by the definition of an open set, we can find a basic open set U containing y with $U \subseteq A^c$. Such a U does not meet A, contradicting the condition (BP). So A must contain y.

Conversely, suppose A contains all of its boundary points, and consider any z in A^c. Then z has to be an interior point or a boundary point of A^c. But the latter possibility does not arise as then z would also be a boundary point of A and hence belong to A (by hypothesis). So z is an interior point of A^c. Consequently, A^c is open, as was to be shown.

Examples. (1) Basic open sets are open: Indeed, let y belong to the open ball $B_a(r) = \{x \mid ||x - a|| < r\}$. Then, since $||\, y - a\, || < r$, the number $r' = \frac{1}{2}(r - ||\, y - a\, ||)$ is positive, and the open ball $B_y(r')$ around y is completely contained in $B_{x_0}(r)$. The case of open rectangular boxes is left as an easy exercise.

(2) The empty set ϕ and $\mathbb{R}^n$ are both open and closed.

Since they are complements of each other, it suffices to check that they are both open, which is clear from the definition.

(3) Let $\{W_\alpha\}$ be a (possibly infinite and perhaps uncountable) collection of open sets in $\mathbb{R}^n$. Then their union $W = \cup_\alpha W_\alpha$ is also open.

Indeed, let $y \in W$. Then $y \in W_\alpha$ for some index α, and since W_α is open, there is an open set $V \subseteq W_\alpha$ containing y. Then we are done as $y \in V \subseteq W_\alpha \subseteq W$.

(4) Let $\{W_1, W_2, \cdots, W_n\}$ be a **finite** collection of open sets. Then their intersection $W = \bigcap_{i=1}^{n} W_i$ is open.

Proof. Let $y \in W$. Then $y \in W_i$, $\forall_i$. Since each W_i is open, we can find a basic open set V_i such that $y \in V_i \subseteq W_i$. Then, by the remark at the end of the previous section, we can find a basic open set U contained in the intersection of the V_i such that $y \in U$. Done.

Warning. The intersection of an infinite collection of open sets need not be open, already for $n = 1$, as shown by the following (counter)example. Put, for each $k \geq 1$, $W_k = \left(-\frac{1}{k}, \frac{1}{k}\right)$. Then $\cap_k W_k = \{0\}$, which is not open.

(5) Any **finite** set of points $A = \{P_1, \ldots, P_r\}$ is closed.

Proof. For each j, let U_j denote the complement of P_j (in $\mathbb{R}^n$). Given any z in U_j, we can easily find a basic open set V_j containing z which avoids P_j. So U_j is open, for each j. The complement of A is simply $\cap_{j=1}^{r} U_j$, which is then open by (4).

More generally, one can show, by essentially the same argument, that a *finite* union of closed sets is again closed.

It is important to remember that there are many sets A in $\mathbb{R}^n$ which are neither open nor closed. For example, look at the half-closed, half-open interval $[0, 1)$ in $\mathbb{R}$.

1.5 Compact subsets of $\mathbb{R}^n$.

It is easy to check that the closed balls $\overline{B}_a(r)$ and the closed rectangular boxes $[a, b]$ are indeed closed. But they are more than that. They are also bounded in the obvious sense. This leads to the notion of "compactness".

Definition. An **open covering** of a set A in $\mathbb{R}^n$ is a collection $\mathcal{U} = \{V_\alpha\}$ of open sets in $\mathbb{R}^n$ such that

$$A \subseteq \cup_\alpha V_\alpha.$$

In other words, each V_α in the collection is open and any point of A belongs to some V_α. Note that $\mathcal{U}$ may be infinite, possibly uncountable.

Clearly, any subset A of $\mathbb{R}^n$ admits an open covering. Indeed, we can just take $\mathcal{U}$ to be the singleton $\{\mathbb{R}^n\}$.

A *subcovering* of an opencovering $\mathcal{U} = \{V_\alpha\}$ of a set A is a subcollection $\mathcal{U}'$ of $\mathcal{U}$ such that any point of A belongs to some set in $\mathcal{U}'$.

Definition. *A set A in $\mathbb{R}^n$ is compact if and only if* **any** *open covering $\mathcal{U} = \{V_\alpha\}$ of A contains a* **finite** *subcovering.*

Example of a set which is not compact: Let $A = (0,1)$ be the open interval in $\mathbb{R}$. Look at $\mathcal{U} = \{W_1, W_2, \cdots\}$ where $W_m = (\frac{1}{m}, 1 - \frac{1}{m})$, for each m. We claim that $\mathcal{U}$ is an open covering of A. Indeed, each W_m is clearly open and moreover, given any number $x \in A$, then $x \in W_m$ for some m. But *no* finite subcollection can cover A, which can be seen as follows. Suppose A is covered by $W_{m_1}, \ldots, W_{m_r}$ for some r, with $m_1 < m_2 < \ldots < m_r$. Then $W_{m_j} \subset W_{m_r}$ for each j, while the point $1/(m_r + 1)$ belongs to A, but not to W_{m_r}, leading to a contradiction. Hence A is not compact.

Example of a set which is compact: Let A be a finite set $\subseteq \mathbb{R}^n$. Then A is compact. Prove it!

Theorem 4 (Heine–Borel) *The closed interval $[a,b]$ in $\mathbb{R}$ is compact for any pair of real numbers a, b with $a < b$.*

Proof. Let $\mathcal{U} = \{V_\alpha\}$ be an open covering of $[a,b]$. Let us call a subset of $[a,b]$ *good* if it can be covered by a finite number of V_α. Put

$$J = \{x \in [a,b] | \, [a,x] \quad \text{is} \quad \text{good}\}.$$

Clearly, a belongs to J, so J is non-empty, and by Theorem 1 of section 2.1, J has a least upper bound; denote it by z. Since b is an upper bound of J, $z \leq b$. We claim that z lies in J. Indeed, pick any open set V_β in $\mathcal{U}$ which contains z. (This is possible because $\mathcal{U}$ is an open covering of $[a,b]$ and z lies in this interval.) Then V_β will contain points to the left of z; call one of them y. Then, by the definition of *lub*, y must lie in J and consequently, $[a,y]$ is covered by a finite subcollection $\{V_{\alpha_1}, \ldots, V_{\alpha_r}\}$ of $\mathcal{U}$. Then $[a,z]$ is covered by $\mathcal{U}' = \{V_\beta, V_{\alpha_1}, \ldots, V_{\alpha_r}\}$, which is still finite; so $z \in J$. In fact, z has to be b. Otherwise, the open set V_β will also contain points to the right of z, and if we call one of them t, say, $[a,t]$ will be covered by $\mathcal{U}'$, implying that t lies in J, contradicting the fact that z is an upper bound of J. Thus b lies in J, and the theorem is proved.

Call a subset A in $\mathbb{R}^n$ **bounded** if we can enclose it in a closed rectangular box.

Theorem 5 *Let A be a subset of $\mathbb{R}^n$ which is closed and bounded. Then A is compact.*

Corollary 1 *Closed balls and spheres in $\mathbb{R}^n$ are compact.*

Remark. It can be shown that the converse of this theorem is also true, i.e., any compact set in $\mathbb{R}^n$ is closed and bounded.

Proof of Theorem 5. The first step is to show that **any closed rectangular box** $R = [a, b]$ **in** $\mathbb{R}^n$ **is compact**: When $n = 1$, this is just the Heine-Borel theorem. So let $n > 1$ and assume by induction that the assertion holds in dimension $< n$. Now we can write

$$R = [a_1, b_1] \times R', \quad \text{with} \quad R' = [a_2, b_2] \times \ldots \times [a_n, b_n].$$

Let $\mathcal{U} = \{W_\alpha\}$ be an open covering of R. Then, for each $y = t \times y'$ in R with t in $[a_1, b_1]$ and $y' = (y_2, \ldots, y_n)$ in R', there is a open set $W_\alpha(y)$ in $\mathcal{U}$ containing this point. By the openness of $W_\alpha(y)$, we can then find an interval $(c_1, d_1) \subset \mathbb{R}$ and an open rectangular box $(c', d') \subset \mathbb{R}^{n-1}$ such that $t \times y \in (c_1, d_1) \times (c', d') \subset W_\alpha(y)$. Then the collection of the sets (c', d') for $y' \in R'$ and t fixed covers R'. Since R' is compact by the induction hypothesis, we can find a finite set, call it $\mathcal{V}'$, of the (c', d') whose union covers R'. Let $I(t)$ denote the intersection of the corresponding finite collection of open intervals (c_1, d_1), which is open (cf. the previous section) and contains t. Then the collection $\{I(t)\}$ is an open covering of $[a_1, b_1]$. By Heine-Borel, we can then extract a finite subcovering, say $\mathcal{V}$, of $[a_1, b_1]$. It is now easy to see that $\mathcal{V} \times \mathcal{V}'$ is a finite subcovering of $[a, b]$. For any $(c_1, d_1) \in \mathcal{V}$ and $(c', d') \in \mathcal{V}'$ we have $(c_1, d_1) \times (c', d') \subseteq W_\alpha$ for some α so R is contained in a union of finitely many W_α's.

The next step is to consider any closed, bounded set A in $\mathbb{R}^n$. Pick any open covering $\mathcal{U}$ of A. Since A is bounded, we can enclose it completely in a closed rectangular box $[a, b]$. Since A is closed, its complement A^c is open. Thus $\mathcal{U} \cup \{A^c\}$ is an open covering of $[a, b]$. Then, by the first step, a finite subcollection, say $\{V_{\alpha_1}, \ldots, V_{\alpha_r}, A^c\}$ covers $[a, b]$. Then, since $A^c \cap A = \emptyset$, the (finite) collection $\{V_{\alpha_1}, \ldots, V_{\alpha_r}\}$ covers A. Done.

1.6 Vector Fields and Continuity

We are interested in functions:

$$f : \mathcal{D} \to \mathbb{R}^m$$

with $\mathcal{D} \subseteq \mathbb{R}^n$, called the **domain** of f. The **image** (or **range**) of f is $f(\mathcal{D}) \subset \mathbb{R}^m$. Such an f is called a **vector field**. When $m = 1$, one says **scalar field** instead, for obvious reasons.

For each $j \leq m$, we write $f_j(x)$ for the jth coordinate of $f(x)$. Then f is completely determined by the collection of scalar fields $\{f_j | j = 1, \ldots, m\}$, called the **component fields**.

Definition. Let $a \in \mathcal{D}$, $b \in \mathbb{R}^m$. Then b is the **limit** of $f(x)$ as x tends to a, denoted

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

if the following holds: *For any $\epsilon > 0$ there is a $\delta > 0$ so that for all $x \in \mathcal{D}$ satisfying $0 < ||x - a|| < \delta$, we have $||f(x) - b|| < \epsilon$.*

This is just the like the definition of limits for functions on the real line but with the absolute value replaced by the norm. In dimension 1, the existence of a limit is equivalent to having a right limit and a left limit, and then having the two limits being equal. In higher dimensions, one can approach a point a from infinitely many directions, and one way to think of it will be to start with an open neighborhood of a and then shrinking it in many different ways to the point a. So the existence of a limit is more stringent a condition here.

The definition of continuity is now literally the same as in the one variable case:

Definition. Let a be in the domain $\mathcal{D}$. Then $f(x)$ is **continuous** at $x = a$ if and only if $\lim_{x \to a} f(x) = f(a)$.

Remarks:

a) A vector field f is continuous at a iff each component field f_j is continuous, for $j = 1, \ldots, m$.

b) In these definitions we need not assume that a is an interior point of $\mathcal{D}$. For example, a could be a boundary point of a domain $\mathcal{D}$ which is a closed box. In an extreme case a could also be the only point of $\mathcal{D}$ in some open ball. In this case continuity becomes an empty condition; every function f is continuous at such a "discrete" point a.

Examples:

(1) Let $f(x, y, z) = ((x^2 + y^2 x)z^2, xy + yz)$, $\mathcal{D} = \mathbb{R}^3$. Then $f : \mathbb{R}^3 \to \mathbb{R}^2$ is continuous at any a in $\mathbb{R}^n$.

More generally, **any** polynomial function is n variables is continuous everywhere. Rational functions $\dfrac{P(x_1, \cdots, x_n)}{Q(x_1, \cdots, x_n)}$ are continuous at all $x = (x_1, \cdots, x_n)$ where $Q(x) \neq 0$.

(2) f is a polynomial function of sines, cosines and exponentials.

It is reasonable to ask at this point what all this has to do with open sets and compact sets. We answer this in the following two lemmas. We call a subset of $\mathcal{D} \subseteq \mathbb{R}^n$ open, resp. closed, if it is the intersection of an open, resp. closed set of $\mathbb{R}^n$ with $\mathcal{D}$.

Lemma 2 *Let $f : \mathcal{D} \to \mathbb{R}^m$ be a vector field. Then f is continuous at every point $a \in \mathcal{D}$ if and only if the following holds: For every open set W of $\mathbb{R}^m$, its inverse image $f^{-1}(W) := \{x \in \mathcal{D} | f(x) \in W\} \subseteq \mathcal{D}$ is open.*

Warning: f continuous does *not* mean that the image of an open set is open. Take, for instance, the constant function $f : \mathbb{R} \to \mathbb{R}$, given by $f(x) = 0$, for all $x \in \mathbb{R}$. This is a continuous map, but any open interval (a, b) in $\mathbb{R}$ gets squished to a point, and $f((a, b))$ is not open.

Proof. Let W be an open set in $\mathbb{R}^m$ and $a \in f^{-1}(W)$. Then choose ϵ so that $B_{f(a)}(\epsilon) \subseteq W$ which is possible since W is open. By continuity there is a δ so that $f(B_a(\delta) \cap \mathcal{D}) \subseteq B_{f(a)}(\epsilon) \subseteq W$ which just means that $B_a(\delta) \cap \mathcal{D} \subseteq f^{-1}(W)$. Since a is arbitrary we find that $f^{-1}(W)$ is open. Conversely, if f satisfies this condition then $f^{-1}(B_f(a)(\epsilon))$ is open since $B_f(a)(\epsilon)$ is open. Hence we can find a small ball $B_a(\delta) \cap \mathcal{D} \subseteq f^{-1}(B_f(a)(\epsilon))$ around $a \in f^{-1}(B_f(a)(\epsilon))$ which implies that f is continuous at a.

Remark: This Lemma shows that the notion of continuity does not depend on the particular norm function used in its definition, only on the collection of open sets defined via this norm function (recall the equivalent ways of using boxes or balls to define open sets).

Lemma 3 *Let $f : \mathbb{R}^n \to \mathbb{R}^m$ be continuous. Then, given any compact set C of $\mathbb{R}^n$, $f(C)$ is compact.*

Proof. Let C be a compact subset of $\mathbb{R}^n$. Pick any open covering $\mathcal{U} = \{W_\alpha\}$ of $f(C)$. Then by the previous lemma, if we put $V_\alpha = f^{-1}(W_\alpha)$, each V_α is an open subset of $\mathbb{R}^n$. Then the collection $\{V_\alpha\}$ will be an open covering of C. By the compactness of C, we ca then find a finite subcovering $\{V_{\alpha_1}, \ldots, V_{\alpha_r}\}$ of C. Since each W_α is simply $f(V_\alpha)$, $f(C)$ will be covered by the finite collection $\{W_{\alpha_1}, \ldots, W_{\alpha_r}\}$. Done.

As a consequence, any continuous image of a closed rectangular box or a closed ball or a sphere will be compact.