

Analysis in $\mathbb{R}^n$ II

Math 20400, Spring 2025

Shuli Chen

based on Boller & Sally and the lecture notes by Ao Sun

Contents

1	Differentiation	5
1.1	Differentiation in one variable	5
1.1.1	Basic definition and properties	5
1.1.2	Mean value theorem	8
1.2	Differential calculus in $\mathbb{R}^n$	12
1.2.1	Review of Euclidean space	12
1.2.2	Directional derivative	13
1.2.3	Differentiable functions from $\mathbb{R}^n$ to $\mathbb{R}^m$	15
1.3	The derivative as a matrix of partial derivatives	21
1.3.1	More about matrices and linear maps	24
1.4	The mean value theorem in $\mathbb{R}^n$	30
1.5	Higher order partial derivatives and Taylor's theorem	33
1.6	Hypersurfaces and tangent hyperplanes in $\mathbb{R}^n$	37
1.6.1	General definitions	37
1.6.2	Special case of a graph	40
1.7	Hessian test for max-min problems	41
1.8	Lagrange multipliers	45
1.8.1	Special case: when we have a parametrization	45
1.8.2	Lagrange multipliers	46
1.9	Implicit and inverse function theorems	50
1.9.1	Inverse function theorem	50
1.9.2	Implicit function theorem	56
2	Integration in one variable	61
2.1	The Riemann Integral in One Variable: Definitions	61

2.2	The Riemann Integral in One Variable: Properties	64
2.3	The Fundamental Theorem of Calculus and Its Consequences	66
2.4	Principal Value Integrals	69

1. Differentiation

1.1 Differentiation in one variable

1.1.1 Basic definition and properties

Let us recall some basic concepts. $\mathbb{R}$ is the real number field.

Definition 1.1.1 — differentiability, derivative. Suppose $f : [a, b] \rightarrow \mathbb{R}$ is a single variable function, we say f is *differentiable* at a point $x \in (a, b)$ if there exists $L \in \mathbb{R}$ such that

$$\lim_{h \rightarrow 0} \frac{f(x+h) - f(x)}{h} = L.$$

The number L is called the *derivative* of f at x , and is denoted by $f'(x)$ or $Df(x)$.

R If we use the ε - δ language to write down this definition, a function f is differentiable at x if and only if there exists $L \in \mathbb{R}$, such that for any $\varepsilon > 0$, there exists $\delta > 0$, such that for any $0 < |h| < \delta$ and $x+h \in [a, b]$,

$$\left| \frac{f(x+h) - f(x)}{h} - L \right| \leq \varepsilon.$$

Theorem 1.1.1 — differentiability implies continuity. Suppose $f : [a, b] \rightarrow \mathbb{R}$ is differentiable at a point $x \in (a, b)$. Then f is continuous at x .

Proof. Take $\varepsilon = 1$. Then there exists a $\delta > 0$ such that $|f(x+h) - f(x) - f'(x)h| < \varepsilon|h| = |h|$ whenever $|h| < \delta$. It follows from the triangle inequality that

$$|f(x+h) - f(x)| < |h| + |f'(x)||h| = (1 + |f'(x)|)|h|$$

when $|h| < \delta$. This implies that

$$\lim_{h \rightarrow 0} f(x+h) = f(x),$$

which implies that f is continuous at x . ■




Figure 1.1: Weierstrass' example

Differentiability is a local (i.e. pointwise) property of functions, that is, it is possible for a function to be differentiable at one point and nowhere else.

■ **Example 1.1** Let $f : \mathbb{R} \rightarrow \mathbb{R}$ be defined by

$$f(x) = \begin{cases} x^2 & \text{if } x \text{ is rational} \\ 0 & \text{if } x \text{ is irrational.} \end{cases}$$

This function is discontinuous at every nonzero x , hence f is not differentiable at any nonzero x . On the other hand, f is continuous at $x = 0$, and we can calculate $f'(0) = 0$, so f is differentiable at 0. Hence f is differentiable at $x = 0$ and nowhere else. ■

Exercise 1.1 Consider the function f_r defined as follows. Let $r \geq 1$ and set

$$f_r(x) = \begin{cases} \frac{1}{q^r} & \text{if } x = \frac{p}{q} \neq 0 \text{ and} \\ 0 & \text{if } x = 0 \text{ or } x \text{ is irrational.} \end{cases}$$

1. Show that for any $r \geq 1$, $f_r(x)$ is continuous at 0 and the irrational numbers and is discontinuous at the nonzero rationals.
2. If $1 \leq r \leq 2$, show that f_r is not differentiable at any irrational point.
3. For which r is f_r differentiable at $x = 0$?

On the other hand, a continuous function may be far from a differentiable function. We have the following very pathological example.

■ **Example 1.2 — Weierstrass' example.** Define

$$f_1(x) = \begin{cases} x & \text{if } 0 \leq x \leq 1/2 \\ 1 - x & \text{if } 1/2 \leq x \leq 1 \end{cases}$$

and extend periodically to $\mathbb{R}$ by $f_1(x + 1) = f_1(x)$. Then f_1 is a continuous function.

Then define for all $n \geq 2$, $f_n(x) = \frac{1}{2} f_{n-1}(2x)$. Let $S_m(x) = \sum_{n=1}^m f_n(x)$. Then S_m is a continuous function on $\mathbb{R}$. ■

- Exercise 1.2** 1. Show that the sequence $(S_m)_{m \in \mathbb{N}^+}$ converges uniformly to a continuous function S .
2. Show that S is not differentiable at dyadic rationals, i.e., rational numbers of the form $p/2^i$ for some integer p and some $i \geq 0$.
3. Let $x \in \mathbb{R}$ be a point that is not a dyadic rational. Define $x_m = p_m/2^m$ and $y_m = (p_m + 1)/2^m$ for $p_m \in \mathbb{Z}$ so that $x_m < x < y_m$ and $\lim x_m = x = \lim y_m$. Notice that by definition, each f_m is differentiable at x , and $|f'_m(x)| = 1$. Thus each S_m is differentiable at x , and

$$|S'_m(x) - S'_{m+1}(x)| = |f'_{m+1}(x)| = 1. \quad (*)$$

Prove that

$$\frac{S(y_m) - S(x)}{y_m - x} < S'_m(x) < \frac{S(x_m) - S(x)}{x_m - x},$$

and use this inequality together with $(*)$ to show S is not differentiable at point x . ■

We said above that differentiability is a pointwise property. Generally speaking, we will assume that a function is not only differentiable at a point, but at all points in a neighborhood of a given point. It is rarely the case that we deal with functions which are differentiable only at a single point.

Next, we review some basic properties of derivatives.

Theorem 1.1.2 Suppose $f : [a, b] \rightarrow \mathbb{R}$ and $g : [a, b] \rightarrow \mathbb{R}$ are both differentiable at $x \in (a, b)$. Then for any $\alpha \in \mathbb{R}$, $\alpha f + g$ is differentiable at x . Also, the product fg and the quotient $\frac{f}{g}$ are differentiable at x (for $\frac{f}{g}$ we must have $g(x) \neq 0$). We have:

1. (Linear property) $(\alpha f + g)'(x) = \alpha f'(x) + g'(x)$;
2. (Product property) $(fg)'(x) = f(x)g'(x) + f'(x)g(x)$;
3. (Quotient property) $\left(\frac{f}{g}\right)'(x) = \frac{f'(x)g(x) - f(x)g'(x)}{g^2(x)}$.

Proof. If you forget about the proof, you can check any rigorous calculus textbook to see a proof. See e.g, Rudin Theorem 5.3. ■

We want to pay special attention to the derivative of a composition of functions, sometimes known as the chain rule. The proof here takes a little care and will require even more care in several variables.

Theorem 1.1.3 — Chain Rule. Let f be differentiable at a point a and let g be differentiable at $f(a)$. Then $g \circ f$ is differentiable at a and $D(g \circ f)(a) = (Dg)(f(a))Df(a)$.

Proof. Let $b = f(a)$, $L = Df(a)$ and $M = Dg(b)$. Set

$$F(x) = f(x) - f(a) - L(x - a),$$

$$G(y) = g(y) - g(b) - M(y - b),$$

$$H(x) = (g \circ f)(x) - (g \circ f)(a) - ML(x - a).$$

By hypothesis, we know that

$$\lim_{x \rightarrow a} \frac{|F(x)|}{|x - a|} = \lim_{y \rightarrow b} \frac{|G(y)|}{|y - b|} = 0.$$

To prove the theorem, we must show that

$$\lim_{x \rightarrow a} \frac{|H(x)|}{|x - a|} = 0.$$

Notice that $H(x) = G(f(x)) + MF(x)$. Now,

$$\frac{|MF(x)|}{|x - a|} = |M| \frac{|F(x)|}{|x - a|} \rightarrow 0$$

as $x \rightarrow a$. It remains to show $\frac{|G(f(x))|}{|x - a|} \rightarrow 0$. For this remaining term, it follows from above that given $\varepsilon > 0$, there exists a $\delta > 0$ such that $|y - b| < \delta$ implies $|G(y)| < \varepsilon|y - b|$. The continuity of f at a implies that given this δ , there exists a $\delta' > 0$ such that $|f(x) - b| < \delta$ when $|x - a| < \delta'$. Hence, if $|x - a| < \delta'$, we have $|G(f(x))| < \varepsilon|f(x) - b|$. But $|f(x) - b| \leq |F(x)| + |L||x - a|$, so

$$\frac{|G(f(x))|}{|x - a|} \leq \frac{\varepsilon|f(x) - b|}{|x - a|} \leq \varepsilon \frac{|F(x)| + |L||x - a|}{|x - a|} = \varepsilon \frac{|F(x)|}{|x - a|} + \varepsilon|L| \rightarrow 0$$

as $x \rightarrow a$. ■

1.1.2 Mean value theorem

Let us review one of the most important theorems in single variable calculus.

Suppose that $[a, b]$ is a closed interval in $\mathbb{R}$ and $f : [a, b] \rightarrow \mathbb{R}$ is continuous. Assume f is differentiable on (a, b) .

R Since f is continuous, we know that f assumes a (global) maximum and minimum value on $[a, b]$. Note that the global maximum and minimum could occur at the boundary points a, b . So f could have no local maximum or minimum on (a, b) , as the function $f(x) = x$ on $[a, b]$ shows.

Theorem 1.1.4 — Local extrema are stationary. Suppose that f satisfies the hypotheses above and f assumes a local maximum or minimum at a point $c \in (a, b)$. Then $f'(c) = 0$.

Proof. Assume that f has a local maximum at c . There exists $\varepsilon > 0$ such that if $|x - c| < \varepsilon$ then $f(x) \leq f(c)$. It follows that if $x \in (c - \varepsilon, c)$ then $\frac{f(x) - f(c)}{x - c} \geq 0$ and if $x \in (c, c + \varepsilon)$ then $\frac{f(x) - f(c)}{x - c} \leq 0$. Thus, in evaluating $f'(c) = \lim_{x \rightarrow c} \frac{f(x) - f(c)}{x - c}$, the former inequality tells us that this limit, if it exists, must be greater than or equal to 0, and the latter shows that if it exists, it must be less than or equal to 0. Since we know by hypothesis that f is differentiable at c , the limit does exist and can only equal 0. The proof for a local minimum is similar. ■

Theorem 1.1.5 — Mean Value Theorem. Let $f : [a, b] \rightarrow \mathbb{R}$ be continuous and suppose that f is differentiable on (a, b) .

1. (Rolle's Theorem) If $f(a) = f(b)$, then there exists $c \in (a, b)$ such that $f'(c) = 0$.
2. (Mean Value Theorem) In any case, there exists $c \in (a, b)$ such that $f'(c)(b - a) = f(b) - f(a)$.
3. (Generalized Mean Value Theorem) If g satisfies the same hypotheses as f , then there exists $c \in (a, b)$ such that $(f(b) - f(a))g'(c) = (g(b) - g(a))f'(c)$.

Proof. Notice that (2) implies (1), and (3) implies (2) by taking $g(x) = x$. Thus we only need to prove only (3). To prove 3, set $h(x) = (f(b) - f(a))g(x) - (g(b) - g(a))f(x)$. Then h is continuous on $[a, b]$ and differentiable on (a, b) . Note that $h(a) = f(b)g(a) - g(b)f(a) = h(b)$.

If h is constant, then $h' \equiv 0$, so we are done. If not, then the global max and min are different, so $h(a) = h(b)$ cannot be both the max and the min. Thus h assumes a max or min at some point $c \in (a, b)$. Theorem 1.1.4 then says that $h'(c) = 0$. The conclusion follows. ■

Corollary 1.1.6 Suppose f is continuous on $[a, b]$ and differentiable on (a, b) . If $f'(x) = 0$ for all $x \in (a, b)$, then f is constant.

Proof. Given any two points x and y in $[a, b]$ with $x < y$, there exists a point $c \in (x, y)$ such that $f(y) - f(x) = (y - x)f'(c) = 0$. Hence, $f(x) = f(y)$. ■

Corollary 1.1.7 Suppose f is continuous on $[a, b]$ and differentiable on (a, b) . If $f'(c) > 0$ for all $c \in (a, b)$, then f is monotonic increasing.

Proof. For $x, y \in [a, b]$ with $x < y$, there exists $c \in (x, y)$ such that $f(y) - f(x) = (y - x)f'(c) > 0$. ■

Corollary 1.1.8 Suppose f is continuous on $[a, b]$ and differentiable on (a, b) . If $f'(c) < 0$ for all $c \in (a, b)$, then f is monotonic decreasing.

Proof. Exercise. ■

Corollary 1.1.9 — L'Hôpital's Rule. Let $I = (a, b)$ be any open interval in $\mathbb{R}$ and suppose f and g are differentiable on I . Let $c \in I$, and suppose $\lim_{x \rightarrow c} f(x) = \lim_{x \rightarrow c} g(x) = 0$ and $\lim_{x \rightarrow c} \frac{f'(x)}{g'(x)}$ exists. Then $\lim_{x \rightarrow c} \frac{f(x)}{g(x)} = \lim_{x \rightarrow c} \frac{f'(x)}{g'(x)}$.

Proof. Exercise. ■

Exercise 1.3 Prove a similar result for limits as x approaches ∞ . ■

■ **Example 1.3** The derivative could be non-continuous. Let

$$f(x) = \begin{cases} x^2 \sin(1/x) & \text{when } x \neq 0 \\ 0 & \text{when } x = 0. \end{cases}$$

Then

$$f'(x) = \begin{cases} 2x \sin(1/x) - \cos(1/x) & \text{when } x \neq 0 \\ 0 & \text{when } x = 0. \end{cases}$$

Then $f'(0)$ exists, but f' is not continuous at 0. ■

On the other hand, although the derivative may not be continuous, it does have the intermediate value property. Recall that continuous functions satisfies an intermediate value theorem.)

Theorem 1.1.10 — Intermediate value theorem for derivatives. Let $f : [a, b] \rightarrow \mathbb{R}$ be continuous and let f be differentiable on (a, b) . Suppose that $(c, d) \subset (a, b)$, $f'(c) < 0$, and $f'(d) > 0$. Then there exists a point $x \in (c, d)$ so that $f'(x) = 0$.

Proof. Since f is continuous on $[c, d]$, it assumes both a maximum and minimum value. Since $f'(c) < 0$, there exists a point $x \in (c, d)$ such that $f(x) < f(c)$ and, since $f'(d) > 0$, there exists a point $y \in (c, d)$ such that $f(y) < f(d)$. Hence the minimum does not occur at either c or d . Theorem 1.1.4 then says there exists a point $z \in (c, d)$ so that $f'(z) = 0$. The conclusion follows. ■

There is no reason to stop at one derivative. Once we get started, we can continue taking derivatives as long as the function allows us.

Definition 1.1.2 — C^k functions. We will use the following notation: suppose $f : (a, b) \rightarrow \mathbb{R}$ is a function. If f is continuous, we write $f \in C^0((a, b))$. If f is differentiable, and its derivative is continuous on (a, b) , then we write $f \in C^1((a, b))$. In general, if f is k -times differentiable and all its m -th order derivatives for $1 \leq m \leq k$ are continuous on (a, b) , then we write $f \in C^k((a, b))$. We will use $f^{(k)}$ to denote the k -th order derivative of f . In particular, $f^{(0)} = f$. We say that $f \in C^\infty(a, b)$ if f has derivatives of all orders in (a, b) . If U is any open set in $\mathbb{R}$, the expression $C^k(U)$ is defined similarly.

R As Example 1.3 shows, f being differentiable on (a, b) alone doesn't guarantee that $f \in C^1((a, b))$. We also need f' to be continuous on (a, b) .

R Notice that by Theorem 1.1.1, if $f^{(k)}$ exists on the whole (a, b) (which means $f^{(k-1)}$ is differentiable on the whole (a, b)), then $f^{(k-1)}$ is continuous. So for $f \in C^\infty(a, b)$, automatically all its derivatives are continuous.

Exercise 1.4 Recall $(e^x)' = e^x$. Let

$$f(x) = \begin{cases} e^{-\frac{1}{x^2}} & \text{for } x \neq 0 \\ 0 & \text{when } x = 0 \end{cases}.$$

Show $f \in C^\infty(\mathbb{R})$ in the following steps:

1. Use induction to show for $x \neq 0$, $f^{(k)}(x) = P_k(1/x)e^{-1/x^2}$ for some polynomial P_k for all $k \geq 0$.
2. Using L'Hôpital's rule, or anything you wish, show that $f^{(k)}(0) = 0$ for all $k \geq 0$. ■

The mean value theorem can be viewed as an approximation theorem - it describes how a general differentiable function f can be approximated by a linear function. In fact, we can interpret the mean value theorem as

$$f(x) \approx f(a) + f'(a)(x - a).$$

In general, we can approximate a function with continuous k -th order derivatives by a degree k polynomial, which is known as the *Taylor polynomial*.

Corollary 1.1.11 — Taylor's Theorem. Suppose $f \in C^{k+1}(a, b)$ and $x_0 \in (a, b)$. Then, for any $x \in (a, b)$, we can write

$$\begin{aligned} f(x) &= f(x_0) + f'(x_0)(x - x_0) + \frac{f''(x_0)}{2!}(x - x_0)^2 + \cdots \\ &\quad + \frac{f^{(k)}(x_0)}{k!}(x - x_0)^k + \frac{1}{(k+1)!}f^{(k+1)}(c)(x - x_0)^{k+1} \end{aligned}$$

where c is a point between x and x_0 .

Definition 1.1.3 — Taylor expansion. What Taylor's theorem allows us to do is to approximate a C^{k+1} function in the neighborhood of a point by a polynomial of degree k . The polynomial in x

$$P_k(x) := f(x_0) + f'(x_0)(x - x_0) + \frac{f''(x_0)}{2!}(x - x_0)^2 + \cdots + \frac{f^{(k)}(x_0)}{k!}(x - x_0)^k$$

is the k -th order Taylor polynomial for f at x_0 , and

$$R_k(x, x_0) := \frac{1}{(k+1)!} f^{(k+1)}(c)(x - x_0)^{k+1}$$

is the k -th order remainder term.

Usually the remainder term $R_k(x, x_0)$ is reasonably small because of the $(k+1)!$ in the denominator. The expansion in Taylor's theorem is called the *Taylor expansion of f about x_0* .

Proof of Corollary 1.1.11. Without loss of generality, we assume that $x > x_0$. For $t \in [x_0, x]$, define a polynomial $T_k(x, t)$ in the variable x by

$$T_k(x, t) = f(t) + f'(t)(x - t) + \frac{f''(t)}{2!}(x - t)^2 + \cdots + \frac{f^{(k)}(t)}{k!}(x - t)^k.$$

Define the remainder term $R_k(x, t) = f(x) - T_k(x, t)$. Notice that $R_k(x, x) = f(x) - f(x) = 0$. Our goal is to find $c \in (x_0, x)$ such that $R_k(x, x_0) = \frac{1}{(k+1)!} f^{(k+1)}(c)(x - x_0)^{k+1}$.

Note that $R_k(x, x_0) \in C^{k+1}(a, b)$, and the derivatives with respect to x satisfy $R_k^{(i)}(x_0, x_0) = 0$ for $0 \leq i \leq k$, and $R_k^{(k+1)}(x, x_0) = f^{(k+1)}(x)$ for all $x \in (a, b)$. For $t \in [x_0, x]$, $x \in (a, b)$, $x > x_0$, set

$$Q(t) = (x - t)^{k+1} \frac{R_k(x, x_0)}{(x - x_0)^{k+1}} - R_k(x, t).$$

It is clear that Q is continuous on $[x_0, x]$, differentiable on (x_0, x) . We also find compute $Q(x_0) = R_k(x, x_0) - R_k(x, x_0) = 0$, and $Q(x) = -R_k(x, x) = 0$. Using product rule, we can compute that the derivative of $R_k(x, t) = f(x) - T_k(x, t)$ with respect to t is $-\frac{(x-t)^k}{k!} f^{(k+1)}(t)$. So

$$Q'(t) = -(k+1)(x-t)^k \frac{R_k(x, x_0)}{(x-x_0)^{k+1}} + \frac{(x-t)^k}{k!} f^{(k+1)}(t).$$

Since $Q(x_0) = Q(x) = 0$, by Rolle's Theorem applied to the function Q on (x_0, x) , there exists $c \in (x_0, x)$ such that $Q'(c) = 0$. Thus

$$(k+1)(x-c)^k \frac{R_k(x, x_0)}{(x-x_0)^{k+1}} = \frac{(x-c)^k}{k!} f^{(k+1)}(c),$$

so $R_k(x, x_0) = \frac{1}{(k+1)!} f^{(k+1)}(c)(x - x_0)^{k+1}$ as desired. ■

Further Reading

- *Rudin*: Chapter 5
- *Tao*: Analysis I, Chapter 10

1.2 Differential calculus in $\mathbb{R}^n$

1.2.1 Review of Euclidean space

Before we start the discussion of the differentiation of multivariable functions, we need to understand the space that we will be working on. The n -dimensional Euclidean space, denoted by $\mathbb{R}^n$, is the Cartesian product of n copies of real numbers as a set. There are three important structures on the Euclidean space: the vector space structure, the normed space structure, and the inner product space structure.

Vector space structure

We begin with the vector space structure of $\mathbb{R}^n$. Apart from thinking n -tuples of real numbers as points, we can also consider n -tuples of real numbers as n -dimensional vectors giving the data of a direction and a magnitude.

Definition 1.2.1 Let $\mathbb{R}^n = \{(x_1, x_2, \dots, x_n) \mid x_i \in \mathbb{R}, i = 1, 2, \dots, n\}$ be the *points* in $\mathbb{R}^n$. Let $\mathbb{E}^n = \{v_1 \mathbf{e}_1 + v_2 \mathbf{e}_2 + \dots + v_n \mathbf{e}_n \mid v_i \in \mathbb{R}, i = 1, 2, \dots, n\}$, where $\mathbf{e}_1 = (1, 0, 0, \dots, 0)$, $\mathbf{e}_2 = (0, 1, 0, \dots, 0)$, $\dots$, $\mathbf{e}_n = (0, 0, 0, \dots, 0, 1)$, be the *vectors* in $\mathbb{R}^n$.

$\mathbb{R}^n$ is a ($\mathbb{R}$ -) vector space. This means that we can define the following two operations, addition and scalar multiplication:

Definition 1.2.2 Given $\mathbf{x} = (x_1, x_2, \dots, x_n)$, $\mathbf{y} = (y_1, y_2, \dots, y_n)$ and $c \in \mathbb{R}$, we define

$$\begin{aligned}\mathbf{x} + \mathbf{y} &:= (x_1 + y_1, x_2 + y_2, \dots, x_n + y_n) \\ c\mathbf{x} &:= (c \cdot x_1, c \cdot x_2, \dots, c \cdot x_n).\end{aligned}$$

These operations should satisfy the following axioms. Suppose $\mathbf{u}, \mathbf{v}, \mathbf{w}$ are elements of $\mathbb{R}^n$, a, b are elements of $\mathbb{R}$.

1. Associativity of vector addition: $\mathbf{u} + (\mathbf{v} + \mathbf{w}) = (\mathbf{u} + \mathbf{v}) + \mathbf{w}$
2. Commutativity of vector addition: $\mathbf{u} + \mathbf{v} = \mathbf{v} + \mathbf{u}$
3. Identity element of vector addition: There exists an element $\mathbf{0} \in V$, called the zero vector, such that $\mathbf{v} + \mathbf{0} = \mathbf{v}$ for all $\mathbf{v} \in V$.
4. Inverse elements of vector addition: For every $\mathbf{v} \in V$, there exists an element $-\mathbf{v} \in V$, called the additive inverse of $\mathbf{v}$, such that $\mathbf{v} + (-\mathbf{v}) = \mathbf{0}$.
5. Compatibility of scalar multiplication with field multiplication $a(b\mathbf{v}) = (ab)\mathbf{v}$
6. Identity element of scalar multiplication $1\mathbf{v} = \mathbf{v}$
7. Distributivity of scalar multiplication with respect to vector addition: $a(\mathbf{u} + \mathbf{v}) = a\mathbf{u} + a\mathbf{v}$
8. Distributivity of scalar multiplication with respect to field addition: $(a + b)\mathbf{v} = a\mathbf{v} + b\mathbf{v}$

In general, a vector space can be defined over any field, not necessarily $\mathbb{R}$. Throughout this course, we will be only focused on vector spaces over $\mathbb{R}$. We also have a special name for the zero element.

Definition 1.2.3 $\mathbf{0} \in \mathbb{R}^n$ is called the origin.

Given a vector space V , a *basis* of V is a collection of vectors $\mathbf{v}_1, \dots, \mathbf{v}_n$ such that any vector in V can be expressed by an unique *linear combination* of these vectors, namely, for any vector $\mathbf{u} \in V$, there exist an unique array of numbers $a_1, a_2, \dots, a_n$ such that

$$\mathbf{u} = \sum_{j=1}^n a_j \mathbf{v}_j$$

In particular, this implies that $\mathbf{v}_1, \dots, \mathbf{v}_n$ are linearly independent, namely if

$$0 = \sum_{j=1}^n b_j \mathbf{v}_j$$

then all the $b_j = 0$.

In the Euclidean space, there is a natural basis given by the Cartesian coordinate. Again, we will use e_j to denote the vector $(0, \dots, 0, 1, 0, \dots, 0)$, which has entry 1 in the j -th position and 0 elsewhere. One can check that $\{e_1, \dots, e_n\}$ is a basis of $\mathbb{R}^n$.

Metric space structure

In the Euclidean space, we can also measure the length of a vector, or the distance between two points. To do this, the vector space structure of $\mathbb{R}^n$ is not enough; we also need to introduce the notion of a norm:

Definition 1.2.4 — norm. If $\mathbf{v} = (v_1, \dots, v_n) \in \mathbb{R}^n$, then the *norm* or *length* of $\mathbf{v}$ is defined by

$$\|\mathbf{v}\| := (v_1^2 + v_2^2 + \dots + v_n^2)^{1/2}.$$

The definition of the norm leads to the usual definition of Euclidean distance between two points in $\mathbb{R}^n$.

Definition 1.2.5 — Euclidean distance. Given two points $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$, the *Euclidean distance* between $\mathbf{x}$ and $\mathbf{y}$ is given by

$$d(\mathbf{x}, \mathbf{y}) := \|\mathbf{x} - \mathbf{y}\|.$$

This is a metric because it satisfies the three axioms of a metric space. The topological structure induced by this metric is crucial. We can then define the notion of open sets, closed sets, etc.

inner product space structure

$\mathbb{R}^n$ is an inner product space, and we can define an inner product on $\mathbb{R}^n$ as follows:

Definition 1.2.6 — dot product (also called scalar product, inner product). Let $\mathbf{v} = (v_1, \dots, v_n)$ and $\mathbf{w} = (w_1, \dots, w_n)$ be two vectors in $\mathbb{R}^n$. The *dot product* of $\mathbf{v}$ and $\mathbf{w}$ is

$$\mathbf{v} \cdot \mathbf{w} := \sum_{i=1}^n v_i w_i.$$

The dot product is also denoted by $\langle \mathbf{v}, \mathbf{w} \rangle$.

This is an inner product because it satisfies the following axioms:

1. (symmetry) $\langle \mathbf{v}, \mathbf{w} \rangle = \langle \mathbf{w}, \mathbf{v} \rangle$;
2. (linearity) $\langle a\mathbf{u} + b\mathbf{v}, \mathbf{w} \rangle = a\langle \mathbf{u}, \mathbf{w} \rangle + b\langle \mathbf{v}, \mathbf{w} \rangle$;
3. (positive-definiteness) $\langle \mathbf{v}, \mathbf{v} \rangle \geq 0$, and $\langle \mathbf{v}, \mathbf{v} \rangle = 0$ iff $\mathbf{v} = \mathbf{0}$.

Using the dot product we can measure the angle between two vectors.

■ **Example 1.4 — angle.** Let θ be the angle between $\mathbf{v}$ and $\mathbf{w}$, then $\langle \mathbf{v}, \mathbf{w} \rangle = \|\mathbf{v}\| \|\mathbf{w}\| \cos \theta$. In particular, for two non-zero vectors $\mathbf{v}$ and $\mathbf{w}$, $\mathbf{v} \perp \mathbf{w}$ (It reads $\mathbf{v}$ is *perpendicular* to $\mathbf{w}$), i.e., $\theta = \frac{\pi}{2}$, if and only if $\langle \mathbf{v}, \mathbf{w} \rangle = 0$. ■

1.2.2 Directional derivative

Now suppose our function is not a single variable function, but a multivariable function. Namely, the function we want to study is $f : \mathbb{R}^n \rightarrow \mathbb{R}$, where $n \geq 2$. We will use $f(x) = f(x_1, x_2, \dots, x_n)$ to denote the value of the function at $x = (x_1, x_2, \dots, x_n)$.

Recall that the intuition behind the derivative of a single variable function: it is the limit of the quotient $\frac{f(x+h)-f(x)}{h}$, where h is the “small increment”. So the derivative of a multivariable function should be the same thing: the limit of the quotient $\frac{f(x+h)-f(x)}{h}$, where h is the small increment.

Recall the vector space structure of the Euclidean space: if we want to add something to a vector x , then it must also be a vector. Next, recall the metric structure gives us a way to measure the length of a vector. These motivate the following definition.

Definition 1.2.7 — directional derivative. Let $U \subset \mathbb{R}^n$ be an open set, and let $f : U \rightarrow \mathbb{R}$. Let $x \in U$ and $\mathbf{v} \in \mathbb{R}^n$ be a unit vector. If the following limit exists

$$\lim_{t \rightarrow 0} \frac{f(x + t\mathbf{v}) - f(x)}{t},$$

then we call it the *directional derivative* of f at x in the direction $\mathbf{v}$. It is denoted by $D_{\mathbf{v}}f(x)$.

In the particular case when $\mathbf{v} = e_j$, a standard basis vector, we obtain the *partial derivative in the j th direction*

$$\frac{\partial f}{\partial x_j}(x) = D_j f(x) = D_{e_j} f(x) = \lim_{t \rightarrow 0} \frac{f(x + te_j) - f(x)}{t}.$$

R A particular notation: if we use (x, y) to denote the coordinate of $\mathbb{R}^2$, then we will use $D_x f$ or $\frac{\partial f}{\partial x}$ to denote $D_1 f$ and $D_y f$ or $\frac{\partial f}{\partial y}$ to denote $D_2 f$. If we use (x, y, z) to denote the coordinate of $\mathbb{R}^3$, we have similar notations. If we use $(x_1, \dots, x_n)$ to denote the coordinate of $\mathbb{R}^n$, then we may use $\frac{\partial f}{\partial x_i}$ to denote $D_i f$.

The reason that the function is defined on an open set can be seen from the definition: to calculate the directional derivative, we need to move the point in the direction we would like to calculate the derivative a little bit. If the domain is an open set then this allows the expression $f(x + h\mathbf{v})$ is still meaningful when h is very small.

■ **Example 1.5** Suppose $f : \mathbb{R}^2 \rightarrow \mathbb{R}$ is defined by $f(x, y) = \mathcal{D}(x)$, where $\mathcal{D}$ is the Dirichlet function

$$\mathcal{D}(x) = \begin{cases} 1 & \text{if } x \text{ is rational} \\ 0 & \text{if } x \text{ is irrational.} \end{cases}$$

Then one can check that $D_x f(x, y)$ does not exist for all $x, y \in \mathbb{R}^2$, while $D_y f(x, y)$ exists and equals 0 for all $x, y \in \mathbb{R}^2$. This example implies that a function can have directional derivatives in some directions meanwhile have no directional derivatives in some other directions. ■

The following propositions for directional derivatives can be checked just like the derivatives of single variable functions. We leave them as exercises.

Proposition 1.2.1 Let $U \subset \mathbb{R}^n$ be an open set, and let $f : U \rightarrow \mathbb{R}$, $g : U \rightarrow \mathbb{R}$. Let $x \in U$ and $\mathbf{v} \in \mathbb{R}^n$ be a unit vector. If f and g both have directional derivatives with respect to $\mathbf{v}$ at x_0 , then

1. (Linear property) $D_{\mathbf{v}}(\alpha f + \beta g)(x_0) = \alpha D_{\mathbf{v}}f(x_0) + \beta D_{\mathbf{v}}g(x_0)$;
2. (Product property) $D_{\mathbf{v}}(fg)(x_0) = f(x_0)D_{\mathbf{v}}g(x_0) + g(x_0)D_{\mathbf{v}}f(x_0)$;
3. (Quotient property) $D_{\mathbf{v}}\left(\frac{f}{g}\right)(x_0) = \frac{g(x_0)D_{\mathbf{v}}f(x_0) - f(x_0)D_{\mathbf{v}}g(x_0)}{g^2(x_0)}$.

It is natural to expect that D is also linear in $\mathbf{v}$. However, the following example suggests that in general, this is not true.

■ **Example 1.6 — a function where all the partial derivatives exist, but some directional derivatives don't.** Suppose $f : \mathbb{R}^2 \rightarrow \mathbb{R}$ is defined by

$$f(x, y) = \begin{cases} \frac{xy}{x^2+y^2} & \text{if } (x, y) \neq (0, 0) \\ 0 & \text{if } (x, y) = (0, 0). \end{cases}$$

Then one can compute that $D_x f(0, 0) = 0$, and $D_y f(0, 0) = 0$. And notice that D_x, D_y also exist at all nonzero points, so D_x, D_y are defined everywhere on $\mathbb{R}^2$. However, for any $\mathbf{v} = (v_1, v_2)$ such that $v_1 v_2 \neq 0$,

$$D_{\mathbf{v}} f(0, 0) = \lim_{t \rightarrow 0} \frac{v_1 v_2}{t(v_1^2 + v_2^2)}$$

does not exist. This example suggests that even though all the partial derivatives exist on the whole $\mathbb{R}^2$, some directional derivatives may not exist even at a point. ■

This example shows that the definition of “differentiable functions” is actually subtle.

1.2.3 Differentiable functions from $\mathbb{R}^n$ to $\mathbb{R}^m$

Basic definitions

From the Taylor expansion, we see that $f(x_0) + f'(x_0)(x - x_0)$ is the best linear function to approximate $f(x)$ near $f(x_0)$. The derivative of single variable functions can be interpreted as the best linear approximation of the function. Similarly, we can define the derivative of a multivariable function as the best linear approximation.

We first need to recall the definition of a linear map.

Definition 1.2.8 — linear transformation, linear map, or linear operator. A function $T : \mathbb{R}^n \rightarrow \mathbb{R}^m$ is called a *linear transformation, linear map, or linear operator* if

1. $T(\mathbf{v}_1 + \mathbf{v}_2) = T(\mathbf{v}_1) + T(\mathbf{v}_2)$ for all $\mathbf{v}_1, \mathbf{v}_2 \in V$,
2. $T(\alpha \mathbf{v}) = \alpha T(\mathbf{v})$ for all $\alpha \in \mathbb{R}$ and $\mathbf{v} \in V$.

R Throughout this course, we will generally use the word “map” or “mapping” to denote a function $\mathbb{R}^n \rightarrow \mathbb{R}^m$.

R For a linear map T , we often omit the parentheses and use the notation Tx to denote $T(x)$.

Lemma 1.2.2 Let $T : \mathbb{R}^n \rightarrow \mathbb{R}^m$ be a linear map. Then T is continuous. ■

Proof. Exercise. ■

Definition 1.2.9 — operator norm. Let $T : \mathbb{R}^n \rightarrow \mathbb{R}^m$ be a linear map. Let $S^{n-1} = \{x \in \mathbb{R}^n \mid \|x\| = 1\}$. This set is closed and bounded, so it is a compact set by Heine-Borel. Since a continuous function attains a maximum on a compact set, T achieves a maximum on S^{n-1} . Let $\|T\|$ denote this maximum value. This is called the *operator norm* of T .

R We use the same symbol $\|\cdot\|$ to indicate the operator norm for a linear transformation as for the Euclidean norm in $\mathbb{R}^n$, but the interpretation will always be clear from the context.

Lemma 1.2.3 Let $T : \mathbb{R}^n \rightarrow \mathbb{R}^m$ be a linear map. Then $\|Th\| \leq \|T\|\|h\|$ for any $h \in \mathbb{R}^n$. Furthermore, for any $M > 0$, we have $\|T\| = \max_{\|h\|=M} \frac{\|Th\|}{M}$.

Proof. Exercise. ■

Exercise 1.5 Let $T : \mathbb{R}^n \rightarrow \mathbb{R}^m$ be a linear map.

1. Show that T is continuous.
2. Show that $\|T\| = 0$ if and only if $T = 0$.
3. Show $\|Th\| \leq \|T\|\|h\|$ for any $h \in \mathbb{R}^n$. Furthermore, for any $M > 0$, show we have $\|T\| = \max_{\|h\|=M} \frac{\|Th\|}{M}$. (Note the norm notation is used in three different ways in this statement!)

Definition 1.2.10 — differentiability, derivative. Suppose $U \subset \mathbb{R}^n$ is an open set. A function $f : U \rightarrow \mathbb{R}^m$ is *differentiable at* $x \in U$ if there is a linear map $T : \mathbb{R}^n \rightarrow \mathbb{R}^m$ such that

$$\lim_{\substack{h \rightarrow 0 \\ h \in \mathbb{R}^n}} \frac{\|f(x+h) - f(x) - Th\|}{\|h\|} = 0.$$

Notice that the norm sign in the numerator denotes the standard Euclidean norm in $\mathbb{R}^m$ while the norm sign in the denominator denotes the standard Euclidean norm in $\mathbb{R}^n$. We write $T = Df(x)$ and we call this the *derivative* of f at x . We say that f is *differentiable on* U if f is differentiable at each point in U .

Similar to the one dimensional case,

$$f(x_0) + Df(x_0)(x - x_0)$$

provides the “best linear approximation” to the function f at the point $x_0 \in \mathbb{R}^n$. If a function is already linear, then the best linear approximation is the function itself:

Proposition 1.2.4 — linear maps are differentiable. If $f : \mathbb{R}^n \rightarrow \mathbb{R}^m$ is a linear map and $x \in \mathbb{R}^n$, then f is differentiable at x and $Df(x) = f$.

Proof. This basically follows from the definition. Since f is linear, we have $f(x+h) = f(x) + f(h)$. Thus

$$\lim_{\substack{h \rightarrow 0 \\ h \in \mathbb{R}^n}} \frac{\|f(x+h) - f(x) - f(h)\|}{\|h\|} = 0.$$

Connection with directional derivatives

When $m = 1$ and f is a real-valued function, we can see how differentiability connects to the directional derivatives we defined earlier.

Proposition 1.2.5 Suppose U is an open set in $\mathbb{R}^n$ and $f : U \rightarrow \mathbb{R}$ is differentiable at a point $x_0 \in U$. Let $\mathbf{v} = (v_1, \dots, v_n) \in \mathbb{R}^n$ be a unit vector. Then $D_{\mathbf{v}}f(x_0)$, the directional derivative of f at x_0 , exists, and

$$D_{\mathbf{v}}f(x_0) = (Df(x_0))(\mathbf{v}) = \sum_{j=1}^n v_j D_j f(x_0).$$

Recall in Example 1.6, if the function f is not differentiable, then D might not be linear in the first argument $\mathbf{v}$. However, when f is differentiable, using the above expression of directional derivatives, we can easily check the linearity of D in the first argument $\mathbf{v}$.

Proposition 1.2.6 Suppose $U \subset \mathbb{R}^n$ is an open set. A function $f : U \rightarrow \mathbb{R}$, $x_0 \in U$. Suppose f is differentiable at x_0 . Then for any $\mathbf{v}, \mathbf{u} \in \mathbb{R}^n$ and $a, b \in \mathbb{R}$,

$$D_{a\mathbf{u}+b\mathbf{v}}f(x_0) = aD_{\mathbf{u}}f(x_0) + bD_{\mathbf{v}}f(x_0).$$

Exercise 1.6

1. For a fixed k with $1 \leq k \leq n$, let $p_k : \mathbb{R}^n \rightarrow \mathbb{R}$ be defined by $p_k(x_1, x_2, \dots, x_n) = x_k$. Show that p_k is differentiable at any point and that $Dp_k(x) = p_k$.
2. Find $D_{\mathbf{v}}p_k(x)$ for any unit vector $\mathbf{v}$ in $\mathbb{R}^n$.

Definition 1.2.11 — projection. The map p_k in the above exercise is called the *projection onto the k^{th} coordinate*. More generally, if $m \leq n$, we can pick indices $1 \leq i_1 < i_2 < \dots < i_m \leq n$ and define a projection $p : \mathbb{R}^n \rightarrow \mathbb{R}^m$ by $p(x_1, x_2, \dots, x_n) = (x_{i_1}, x_{i_2}, \dots, x_{i_m})$.

Be careful: Although the differentiability of f implies the directional derivatives exist for all directions, the converse is not true. For differentiability, one must approach the given point from every direction along all sorts of paths, while the directional derivative is taken along straight-line paths through the point.

■ **Example 1.7 — a non-differentiable function where all the directional derivatives exist.** Suppose $f : \mathbb{R}^2 \rightarrow \mathbb{R}$ is defined as follows

$$f(x, y) = \begin{cases} \frac{y^2}{x} & x \neq 0 \\ 0 & x = 0. \end{cases}$$

For any $\mathbf{v} = (v_1, v_2) \in \mathbb{R}^2$, $D_{\mathbf{v}}f(0, 0)$ exists. In fact, if $v_1 = 0$,

$$\lim_{h \rightarrow 0} \frac{f(0, hv_2) - f(0, 0)}{h} = 0,$$

and if $v_1 \neq 0$,

$$\lim_{h \rightarrow 0} \frac{f(hv_1, hv_2) - f(0, 0)}{h} = \lim_{h \rightarrow 0} \frac{h^2 v_2^2}{h^2 v_1} = \frac{v_2^2}{v_1}.$$

However, f is not differentiable at $(0, 0)$. In fact, if we take a sequence of vectors $\{\mathbf{v}_n := (\frac{1}{n^2}, \frac{1}{n})\}_{n \in \mathbb{N}^+}$, for any linear function $A = (a_1, a_2) : \mathbb{R}^2 \rightarrow \mathbb{R}$, we have

$$\frac{f(\mathbf{v}_n) - f(0, 0) - A(\mathbf{v}_n)}{\|\mathbf{v}_n\|} = \frac{1 - \frac{a_1}{n^2} - \frac{a_2}{n}}{\sqrt{\frac{1}{n^4} + \frac{1}{n^2}}}$$

which is tending to $\frac{1}{0}$, hence

$$\lim_{n \rightarrow \infty} \frac{f(\mathbf{v}_n) - f(0, 0) - A(\mathbf{v}_n)}{\|\mathbf{v}_n\|}$$

turns to infinity and does not exist. But $\|\mathbf{v}_n\| \rightarrow 0$ as $n \rightarrow \infty$. This implies that Df does not exist at 0 , although all the directional derivatives exist. ■

The above example is very pathological. We also recall from Example 1.6 that even if all the partial derivatives exist in a neighborhood of some point x_0 , the function could still be non-differentiable at that point. Nevertheless, if we require the partial derivatives of f to be *continuous* in a neighborhood of x_0 , then f is differentiable at x_0 .

Theorem 1.2.7 Suppose $U \subset \mathbb{R}^n$ is open, $f : U \rightarrow \mathbb{R}$, $x_0 \in U$. Suppose there exists a open neighborhood $V \subset U$ of x_0 such that for all $j = 1, 2, \dots, n$, $D_j f$ exist and are continuous in V , then f is differentiable at x_0 .

R Note that V is a open neighborhood of all point inside it, so this just implies f is differentiable on the whole open set V .

Proof. For simplicity, let us assume V is a ball centered at $x_0 = (x_1, x_2, \dots, x_n)$. Suppose $\mathbf{v} = (v_1, v_2, \dots, v_n) \in \mathbb{R}^n$ satisfies that $x_0 + \mathbf{v} \in V$ and $\mathbf{v} \neq 0$. (We want to have a ball around x_0 because we want to make sure when we are later writing the expression $\sup_{\|\mathbf{w}\| \leq \|\mathbf{v}\|} |D_k g(\mathbf{u})|$, the vector $x_0 + \mathbf{w}$ is also in V (hence in U) so that the expression $D_k g(u)$ makes sense.) Let $A = (D_1 f(x_0), \dots, D_n f(x_0))$. Then $D_j A = D_j f(x_0)$. Notice that at this moment we don't know $A = Df(x_0)$ because we haven't proved that f is differentiable at x_0 . Let $g(\mathbf{v}) = f(x_0 + \mathbf{v}) - A\mathbf{v}$. Then it suffices to show that

$$\lim_{\mathbf{v} \rightarrow 0} \frac{|g(\mathbf{v}) - g(\mathbf{0})|}{\|\mathbf{v}\|} = 0.$$

We can write

$$\begin{aligned} & g(\mathbf{v}) - g(\mathbf{0}) \\ &= g(v_1, \dots, v_n) - g(0, \dots, 0) \\ &= [g(v_1, \dots, v_n) - g(v_1, \dots, v_{n-1}, 0)] + [g(v_1, \dots, v_{n-1}, 0) - g(v_1, \dots, v_{n-2}, 0, 0)] \\ &\quad + \dots + [g(v_1, 0, \dots, 0) - g(0, \dots, 0)]. \end{aligned}$$

For the term

$$g(v_1, \dots, v_k, v_{k+1}, 0, \dots, 0) - g(v_1, \dots, v_k, 0, 0, \dots, 0),$$

fixing $v_1, \dots, v_k$, we can use the mean value theorem to show that there exists c_{k+1} between 0 and v_{k+1} , such that

$$g(v_1, \dots, v_k, v_{k+1}, 0, \dots, 0) - g(v_1, \dots, v_k, 0, 0, \dots, 0) = D_{k+1} g(v_1, \dots, v_k, c_{k+1}, 0, \dots, 0) v_{k+1}.$$

Note that even when $v_{k+1} = 0$, this equality holds. Then we have

$$\|g(\mathbf{v}) - g(\mathbf{0})\| \leq \sum_{j=1}^n |v_j| \sup_{\|\mathbf{u}\| \leq \|\mathbf{v}\|} |D_j g(\mathbf{u})|.$$

It is straightforward to check that for any $1 \leq j \leq n$, $|v_j| \leq \|\mathbf{v}\|$. So we have

$$\frac{\|g(\mathbf{v}) - g(\mathbf{0})\|}{\|\mathbf{v}\|} \leq \sum_{j=1}^n \sup_{\|\mathbf{u}\| \leq \|\mathbf{v}\|} |D_j g(\mathbf{u})|.$$

By the continuity of $D_j f$, we have

$$\begin{aligned} \lim_{\mathbf{u} \rightarrow 0} |D_j g(\mathbf{u})| &= \lim_{\mathbf{u} \rightarrow 0} |D_j f(x_0 + \mathbf{u}) - D_j A(\mathbf{u})| \\ &= \lim_{\mathbf{u} \rightarrow 0} |D_j f(x_0 + \mathbf{u}) - D_j f(x_0)| \\ &= 0. \end{aligned}$$

Then we use the comparison of limits to conclude that

$$\lim_{\mathbf{v} \rightarrow 0} \frac{\|g(\mathbf{v}) - g(\mathbf{0})\|}{\|\mathbf{v}\|} = 0.$$

■

R Notice that we have this chain of implication where each arrow cannot go the other way:

$$\begin{aligned} & \text{all partial derivatives } D_j f \text{ exist and are continuous in a neighborhood of } x_0 \\ \implies & f \text{ is differentiable at } x_0 \\ \implies & \text{all directional derivatives } D_v f(x_0) \text{ exist at } x_0 \\ \implies & \text{all partial derivatives } D_j f(x_0) \text{ exist at } x_0 \end{aligned}$$

Basic properties of the derivative

There are many theorems about derivatives of functions of several variables which are analogous to those in one variable.

Theorem 1.2.8 — differentiability implies continuity. Suppose U is an open set in $\mathbb{R}^n$ and $f : U \rightarrow \mathbb{R}^m$ is differentiable at a point $x_0 \in U$. Then f is continuous at x_0 .

Proof. Take $\varepsilon = 1$. Then there exists a $\delta > 0$ such that

$$\|f(x_0 + h) - f(x_0) - Df(x_0)h\| < \varepsilon \|h\| = \|h\|$$

whenever $\|h\| < \delta$. It follows from the triangle inequality that

$$\|f(x_0 + h) - f(x_0)\| \leq \|Df(x_0)h\| + \|h\| \leq \|Df(x_0)\| \cdot \|h\| + \|h\| = (\|Df(x_0)\| + 1)\|h\|$$

when $\|h\| < \delta$. Taking $h \rightarrow 0$ shows $\lim_{h \rightarrow 0} f(x_0 + h) = f(x_0)$, so f is continuous. ■

Proposition 1.2.9 — uniqueness of derivative. If U is an open set in $\mathbb{R}^n$ and $f : U \rightarrow \mathbb{R}^m$ is differentiable, then the derivative of f is unique.

Proof. Suppose T and T' are derivatives of f at $x \in U$. Then, for $h \in \mathbb{R}^n$,

$$\begin{aligned} \frac{\|Th - T'h\|}{\|h\|} &= \frac{\|(f(x+h) - f(x) - Th) - (f(x+h) - f(x) - T'h)\|}{\|h\|} \\ &\leq \frac{\|f(x+h) - f(x) - Th\|}{\|h\|} + \frac{\|f(x+h) - f(x) - T'h\|}{\|h\|}. \end{aligned}$$

Thus,

$$\frac{\|Th - T'h\|}{\|h\|} \rightarrow 0 \quad \text{as } h \rightarrow 0.$$

But $T - T'$ is a linear map, so by the properties of the operator norm (Lemma 1.2.3),

$$\max_{\|h\|=M} \frac{\|Th - T'h\|}{\|h\|} = \|T - T'\|.$$

Because this is a constant independent of M for $M \neq 0$, we must have $\|T - T'\| = 0$, that is, $T = T'$. ■

Proposition 1.2.10 — linearity of derivative. Let U be an open set in $\mathbb{R}^n$ and $f, g : U \rightarrow \mathbb{R}^m$ be differentiable at $x \in U$. Then $f + g$ is differentiable at $x \in U$, and $D(f + g)(x) = Df(x) + Dg(x)$.

Proof. For $x \in U$,

$$\begin{aligned} & \frac{\|(f + g)(x + h) - (f + g)(x) - (Df(x) + Dg(x))h\|}{\|h\|} \\ & \leq \frac{\|f(x + h) - f(x) - Df(x)h\|}{\|h\|} + \frac{\|g(x + h) - g(x) - Dg(x)h\|}{\|h\|}. \end{aligned}$$

Both expressions on the right go to zero as h goes to zero. ■

The reader might ask about the product rule. The question depends on the type of multiplication being used. Let us try the easiest case in which f and g map U to $\mathbb{R}$.

Proposition 1.2.11 — product rule for functions that map to $\mathbb{R}$. If U is an open set in $\mathbb{R}^n$, and $f : U \rightarrow \mathbb{R}$ and $g : U \rightarrow \mathbb{R}$ are real-valued functions which are differentiable at $x \in U$, then fg is differentiable at $x \in U$, and $D(fg)(x) = f(x)Dg(x) + g(x)Df(x)$.

Proof. Before starting the proof, we observe that the expression $f(x)Dg(x)$ really makes sense, since $f(x)$ is a real scalar and $Dg(x)$ is a linear map, and we can always multiply linear maps by scalars. And notice that $f(x), g(x)$ are just real numbers. Thus

$$\begin{aligned} & \frac{1}{\|h\|} |f(x+h)g(x+h) - f(x)g(x) - (f(x)Dg(x) + g(x)Df(x))h| \\ &= \frac{1}{\|h\|} |(f(x+h)[g(x+h) - g(x)] - f(x)Dg(x)h) + (g(x)[f(x+h) - f(x)] - g(x)Df(x)h)| \\ &\leq \frac{1}{\|h\|} \left(|(f(x+h) - f(x))(g(x+h) - g(x))| + |f(x)(g(x+h) - g(x) - Dg(x)h)| \right. \\ &\quad \left. + |g(x)(f(x+h) - f(x) - Df(x)h)| \right). \end{aligned}$$

By the definition of the derivative, the second and third terms vanish as $h \rightarrow 0$. For the first term, we have

$$\frac{1}{\|h\|} |(f(x+h) - f(x))(g(x+h) - g(x))| = \frac{|f(x+h) - f(x)|}{\|h\|} \cdot |g(x+h) - g(x)|.$$

This is going to zero because the first term is bounded and the second term is going to zero by the continuity of g . Indeed, we can see the first term is bounded by

$$\begin{aligned} \frac{|f(x+h) - f(x)|}{\|h\|} &= \frac{|f(x+h) - f(x) - Df(x)h + Df(x)h|}{\|h\|} \\ &\leq \frac{|f(x+h) - f(x) - Df(x)h| + |Df(x)h|}{\|h\|} \\ &= \frac{|f(x+h) - f(x) - Df(x)h|}{\|h\|} + \frac{|Df(x)h|}{\|h\|} \\ &\leq \frac{|f(x+h) - f(x) - Df(x)h|}{\|h\|} + \|Df(x)\| \\ &\rightarrow 0 + \|Df(x)\| = \|Df(x)\|. \end{aligned}$$

■

We now turn to the Chain Rule which is a very important theorem in the calculus of several variables. The reader will note that the proof is pretty much the same as the proof in one variable.

Theorem 1.2.12 — Chain Rule. Suppose that U is an open set in $\mathbb{R}^n$ and $f : U \rightarrow \mathbb{R}^m$ is differentiable on U . Take a point $a \in U$ and let V be an open set in $\mathbb{R}^m$ such that $f(a) \in V$. Suppose $g : V \rightarrow \mathbb{R}^p$ is differentiable on V . Then $g \circ f : U \rightarrow \mathbb{R}^p$ is differentiable at a , and $D(g \circ f)(a) = Dg(f(a)) \circ Df(a)$.

Proof. Let $b = f(a)$, $L = Df(a)$, and $M = Dg(b)$. Set

$$\phi(x) = f(x) - f(a) - L(x - a),$$

$$\psi(y) = g(y) - g(b) - M(y - b),$$

and

$$\rho(x) = (g \circ f)(x) - (g \circ f)(a) - ML(x - a).$$

Then, our hypotheses tell us that

$$\lim_{x \rightarrow a} \frac{\|\phi(x)\|}{\|x - a\|} = \lim_{y \rightarrow b} \frac{\|\psi(y)\|}{\|y - b\|} = 0,$$

and we need to show that

$$\lim_{x \rightarrow a} \frac{\|\rho(x)\|}{\|x - a\|} = 0.$$

But from the definitions we have

$$\rho(x) = \psi(f(x)) + M\phi(x)$$

Then

$$\frac{\|M\phi(x)\|}{\|x - a\|} \leq \|M\| \frac{\|\phi(x)\|}{\|x - a\|} \rightarrow 0$$

as $x \rightarrow a$. By the above limit and the continuity of f , for any $\varepsilon > 0$, there is $\delta > 0$ such that $\|x - a\| < \delta$ ensures $\|\psi(f(x))\| < \varepsilon\|f(x) - b\|$. But $\|f(x) - b\| \leq \|\phi(x)\| + \|L\|\|x - a\|$, so

$$\frac{\|\psi(f(x))\|}{\|x - a\|} \rightarrow 0$$

as $x \rightarrow a$ as well. ■

1.3 The derivative as a matrix of partial derivatives

Definition 1.3.1 If $f : \mathbb{R}^n \rightarrow \mathbb{R}^m$, we can write $f(x) = (f_1(x), f_2(x), \dots, f_m(x))$. Then, for each k , $1 \leq k \leq m$, the map $f_k : \mathbb{R}^n \rightarrow \mathbb{R}$ is called the k -th *component function* of f .

A consequence of the chain rule is the following:

Corollary 1.3.1 If $f : \mathbb{R}^n \rightarrow \mathbb{R}^m$ is differentiable, then the component function f_k is differentiable for each k .

Proof. Observe that $f_k = p_k \circ f$ and use the chain rule. Here p_k is the projection onto the k th coordinate we defined earlier. ■

Conversely, we have

Lemma 1.3.2 If each f_k is differentiable on an open set $U \subset \mathbb{R}^n$, then $f = (f_1, f_2, \dots, f_m)$ is differentiable on U and $Df = (Df_1, Df_2, \dots, Df_m)$.

Proof. Exercise. ■

Recall that for a function $f : U \rightarrow \mathbb{R}$ that map to $\mathbb{R}$, the existence of the partial derivatives at a point, or even in the neighborhood of a point, does not assure the differentiability of the function at the point. However, in 1.2.7, we showed that if the partial derivatives of f all exist and are continuous in a neighborhood of a point, then the function is differentiable at that point. Using the previous lemma, we can extend this result to general functions from $\mathbb{R}^n$ to $\mathbb{R}^m$.

Theorem 1.3.3 Let U be an open set in $\mathbb{R}^n$ and let $f : U \rightarrow \mathbb{R}^m$ be a function with the property that $D_j f_i$ is continuous on U for all i, j . Then f is differentiable on U and $Df = (Df_1, Df_2, \dots, Df_m)$.

Proof. By Theorem 1.2.7, each component function f_k is differentiable. Then by the previous lemma, f is differentiable and $Df = (Df_1, Df_2, \dots, Df_m)$. ■

R This theorem is very useful for showing that some function is differentiable. Usually, it is indirect to find some linear map and check the definition of differentiability, but easier to compute the partial derivatives. And in many cases these partial derivatives are just given by some nice familiar functions so we know they are continuous. Then the previous theorem allows us to conclude the function is differentiable.

Note that here each $Df_i(x)$ is a linear transformation (i.e. linear map) from $\mathbb{R}^n$ to $\mathbb{R}$ and $Df(x)$ is a linear transformation from $\mathbb{R}^n$ to $\mathbb{R}^m$. Recall from linear algebra that linear transformations from $\mathbb{R}^n$ to $\mathbb{R}^m$ can be represented by m by n matrices with respect to any bases we choose in the respective spaces. In particular we can choose the standard basis $\{e_j\}$ for the Euclidean spaces. Then, given a linear map $T : \mathbb{R}^n \rightarrow \mathbb{R}^m$, $T(e_j)$ is a vector in $\mathbb{R}^m$, and if we put all the n many vectors $T(e_j)$ together, we get an m by n matrix:

$$[T(e_1) \quad \dots \quad T(e_n)].$$

Using this idea, we can also put the derivative in a matrix form, and we have the following definition:

Definition 1.3.2 — Jacobi matrix. Let $U \subset \mathbb{R}^n$ and $f = (f_1, f_2, \dots, f_m) : U \rightarrow \mathbb{R}^m$ be differentiable at $x_0 \in U$. The *Jacobi matrix* of f at x_0 , denoted by $\mathcal{J}_f(x_0)$, is the m by n matrix

$$\mathcal{J}_f(x_0) := \begin{bmatrix} \frac{\partial f_1}{\partial x_1} & \frac{\partial f_1}{\partial x_2} & \dots & \frac{\partial f_1}{\partial x_n} \\ \frac{\partial f_2}{\partial x_1} & \frac{\partial f_2}{\partial x_2} & \dots & \frac{\partial f_2}{\partial x_n} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{\partial f_m}{\partial x_1} & \frac{\partial f_m}{\partial x_2} & \dots & \frac{\partial f_m}{\partial x_n} \end{bmatrix}.$$

That is, if we use a_{ij} to denote the entry in the i -th row and j -th column of the matrix $\mathcal{J}_f(x_0)$, then $a_{ij} = D_j f_i(x_0) = \frac{\partial f_i}{\partial x_j}$.

This allows us to make computations more explicitly. We can abuse the notation and just write $Df(x_0) = \mathcal{J}_f(x_0)$ as a matrix. For $h = (h_1, \dots, h_n) \in \mathbb{R}^n$, we have

Proposition 1.3.4 If $U \subset \mathbb{R}^n$ and $f : U \rightarrow \mathbb{R}^m$ is differentiable at x_0 , then for $h = (h_1, \dots, h_n) \in \mathbb{R}^n$, we have

$$Df(x_0)(h) = \mathcal{J}_f(x_0)h = \begin{bmatrix} \frac{\partial f_1}{\partial x_1} & \frac{\partial f_1}{\partial x_2} & \dots & \frac{\partial f_1}{\partial x_n} \\ \frac{\partial f_2}{\partial x_1} & \frac{\partial f_2}{\partial x_2} & \dots & \frac{\partial f_2}{\partial x_n} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{\partial f_m}{\partial x_1} & \frac{\partial f_m}{\partial x_2} & \dots & \frac{\partial f_m}{\partial x_n} \end{bmatrix} \begin{bmatrix} h_1 \\ \vdots \\ h_n \end{bmatrix} = \begin{bmatrix} \sum_{j=1}^n \frac{\partial f_1}{\partial x_j} h_j \\ \vdots \\ \sum_{j=1}^n \frac{\partial f_m}{\partial x_j} h_j \end{bmatrix}.$$

In the case that $m = 1$, $\mathcal{J}_f(x_0)$ is a $1 \times n$ matrix, so it is a row vector. But since usually we view vectors as a column vector, we make the following definition:

Definition 1.3.3 The vector-valued function $(D_1f(x), D_2f(x), \dots, D_nf(x))$ is called the *gradient vector* and is denoted $\nabla f(x)$.

In this case $Df(x)$ is a $1 \times n$ row vector, and $\nabla f(x)$ is a $n \times 1$ column vector. Then from our previous discussions and this definition, we have

Proposition 1.3.5 If $U \subset \mathbb{R}^n$ and $f : U \rightarrow \mathbb{R}$ is differentiable at x_0 , then for $h = (h_1, \dots, h_n) \in \mathbb{R}^n$,

$$Df(x_0)h = D_1f(x_0)h_1 + D_2f(x_0)h_2 + \dots + D_nf(x_0)h_n = \nabla f(x_0) \cdot h.$$

Proposition 1.3.6 — chain rule in terms of matrices. Suppose that U is an open set in $\mathbb{R}^n$ and $f : U \rightarrow \mathbb{R}^m$ is differentiable on U . Take a point $a \in U$ and let V be an open set in $\mathbb{R}^m$ such that $f(a) \in V$. Suppose $g : V \rightarrow \mathbb{R}^p$ is differentiable on V . Using the matrix notation, we can express the chain rule as

$$D(g \circ f)(x_0) = Dg(f(x_0))Df(x_0),$$

where $D(g \circ f)(x_0)$ is an p by n matrix, $Dg(f(x_0))$ is a p by m matrix and $Df(x_0)$ is a m by n matrix, and $Dg(f(x_0))Df(x_0)$ is the matrix multiplication.

With this matrix representation, it becomes much easier to compute the derivative of a function given by explicit formulas.

■ **Example 1.8** Let $f : \mathbb{R}^3 \rightarrow \mathbb{R}^2$ be defined by

$$f(x, y, z) = (x^2y + e^{xz}, \sin(xyz)).$$

Then, f is differentiable on all of $\mathbb{R}^3$, and its Jacobian matrix is

$$Df(x, y, z) = \begin{bmatrix} 2xy + ze^{xz} & x^2 & xe^{xz} \\ yz \cos(xyz) & xz \cos(xyz) & xy \cos(xyz) \end{bmatrix}.$$

Thus, evaluating at $(1, 0, -1)$, we obtain

$$Df(1, 0, -1) = \begin{bmatrix} -e^{-1} & 1 & e^{-1} \\ 0 & -1 & 0 \end{bmatrix}.$$

■

■ **Example 1.9 — a non-differentiable function where all the partial derivatives are bounded.**

Let $f : \mathbb{R}^2 \rightarrow \mathbb{R}$ be defined by

$$f(x, y) = \begin{cases} \frac{xy}{\sqrt{x^2 + y^2}}, & \text{if } (x, y) \neq (0, 0), \\ 0, & \text{if } (x, y) = (0, 0). \end{cases}$$

On $\mathbb{R}^2 \setminus \{0\}$, the function is differentiable, and the gradient is

$$\nabla f = (D_x f, D_y f) = \left(\frac{y^3}{(x^2 + y^2)^{3/2}}, \frac{x^3}{(x^2 + y^2)^{3/2}} \right).$$

Notice that this vector is bounded: each entry is bounded above by 1.

At $(x, y) = (0, 0)$, the partial derivatives exist, and we find that $D_x f(0, 0) = D_y f(0, 0) = 0$.

However, this function is not differentiable at $(0, 0)$. If it were differentiable there, then

$$\lim_{(x,y) \rightarrow (0,0)} \frac{|f(x, y) - f(0, 0) - Df(0, 0) \cdot (x, y)|}{\sqrt{x^2 + y^2}} = 0.$$

However, taking $x = y$ shows

$$\lim_{(x,x) \rightarrow (0,0)} \frac{|f(x,x) - f(0,0)|}{\sqrt{x^2 + x^2}} = \frac{1}{2},$$

so the function is not differentiable at $(0,0)$.

Also, we note that the partial derivatives $D_x f, D_y f$ are not continuous at the origin (Check that!) ■

■ **Example 1.10 — polar coordinates.** The polar coordinate system provides an important parametrization of $\mathbb{R}^2$ (excluding a small set). Suppose $U \subset \mathbb{R}^2$ is the set $(0, \infty) \times (0, 2\pi)$. We use r to denote the variable in $(0, \infty)$ and θ to denote the variable in $(0, 2\pi)$. Then we define a map $f : U \rightarrow \mathbb{R}^2$ by

$$f(r, \theta) = (r \cos \theta, r \sin \theta).$$

The Jacobian matrix of f is

$$\mathcal{J}_f(r, \theta) = \begin{bmatrix} \cos \theta & -r \sin \theta \\ \sin \theta & r \cos \theta \end{bmatrix}.$$

■ **Example 1.11** Suppose $g : \mathbb{R}^2 \rightarrow \mathbb{R}$ is a differentiable function, and $f : (0, \infty) \times (0, 2\pi) \rightarrow \mathbb{R}^2$ is the polar parametrization. Then, using the chain rule, we have

$$\begin{bmatrix} \frac{\partial(g \circ f)}{\partial r}(r, \theta) \\ \frac{\partial(g \circ f)}{\partial \theta}(r, \theta) \end{bmatrix} = \begin{bmatrix} \frac{\partial g}{\partial x}(f(r, \theta)) & \frac{\partial g}{\partial y}(f(r, \theta)) \end{bmatrix} \begin{bmatrix} \cos \theta & -r \sin \theta \\ \sin \theta & r \cos \theta \end{bmatrix}.$$

Expanding the matrix product, we get:

$$\begin{bmatrix} \cos \theta \frac{\partial g}{\partial x}(f(r, \theta)) + \sin \theta \frac{\partial g}{\partial y}(f(r, \theta)) \\ -r \sin \theta \frac{\partial g}{\partial x}(f(r, \theta)) + r \cos \theta \frac{\partial g}{\partial y}(f(r, \theta)) \end{bmatrix}.$$

Sometimes, people abuse the notation and write:

$$\frac{\partial}{\partial r} = \cos \theta \frac{\partial}{\partial x} + \sin \theta \frac{\partial}{\partial y},$$

and

$$\frac{\partial}{\partial \theta} = -r \sin \theta \frac{\partial}{\partial x} + r \cos \theta \frac{\partial}{\partial y}.$$

■

1.3.1 More about matrices and linear maps

Recall the Cauchy-Schwarz inequality we learned in Math 203 and also from some linear algebra course:

Theorem 1.3.7 — Cauchy-Schwarz inequality. For all $\mathbf{v}, \mathbf{w} \in \mathbb{R}^n$, we have

$$|\langle \mathbf{v}, \mathbf{w} \rangle| \leq \|\mathbf{v}\| \|\mathbf{w}\|.$$

Moreover, if $|\langle \mathbf{u}, \mathbf{v} \rangle| = \|\mathbf{u}\| \|\mathbf{v}\|$ then $\mathbf{u}, \mathbf{v}$ are in the same or opposite direction. That is, if $\mathbf{v} \neq \mathbf{0}$, then $\mathbf{u} = \frac{\langle \mathbf{u}, \mathbf{v} \rangle}{\|\mathbf{v}\|^2} \mathbf{v}$.

Suppose $A : \mathbb{R}^n \rightarrow \mathbb{R}^m$ is a linear map. Recall the operator norm $\|A\| = \max_{\|\mathbf{v}\|=1} \|A\mathbf{v}\|$.

Using the matrix representation of a linear map and the Cauchy-Schwarz inequality, we can also give an upper bound of $\|A\|$.

Proposition 1.3.8 Suppose

$$A = \begin{bmatrix} a_{11} & a_{12} & \cdots & a_{1n} \\ a_{21} & a_{22} & \cdots & a_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ a_{m1} & a_{m2} & \cdots & a_{mn} \end{bmatrix}.$$

Then

$$\|A\| \leq \sqrt{\sum_{i=1}^m \sum_{j=1}^n a_{ij}^2}.$$

Proof. Suppose $\mathbf{v} \in \mathbb{R}^n$ such that $\|\mathbf{v}\| = 1$. Then

$$A\mathbf{v} = \begin{bmatrix} \sum_{j=1}^n a_{1j}v_j \\ \vdots \\ \sum_{j=1}^n a_{mj}v_j \end{bmatrix},$$

so

$$\|A\mathbf{v}\|^2 = \sum_{i=1}^m \left(\sum_{j=1}^n a_{ij}v_j \right)^2.$$

For each i , the Cauchy-Schwarz inequality applied to the two vectors $(a_{i1}, \dots, a_{in})$ and $(v_1, \dots, v_n)$ shows that

$$\left(\sum_{j=1}^n a_{ij}v_j \right)^2 \leq \left(\sum_{j=1}^n a_{ij}^2 \right) \left(\sum_{j=1}^n v_j^2 \right).$$

Since $\sum_{j=1}^n v_j^2 = \|\mathbf{v}\|^2 = 1$, we get

$$\left(\sum_{j=1}^n a_{ij}v_j \right)^2 \leq \sum_{j=1}^n a_{ij}^2.$$

Thus

$$\|A\mathbf{v}\|^2 = \sum_{i=1}^m \left(\sum_{j=1}^n a_{ij}v_j \right)^2 \leq \sum_{i=1}^m \sum_{j=1}^n a_{ij}^2.$$

Since this holds for all $\|\mathbf{v}\| = 1$, we have $\|A\|^2 \leq \sum_{i=1}^m \sum_{j=1}^n a_{ij}^2$ as desired. $\blacksquare$

Using matrices, we can also give a more precise characterization of the operator norm. For this we need to recall a few more notions from linear algebra. If these notions sound unfamiliar to you, it's your opportunity to refresh your linear algebra knowledge.

For a matrix A , we let A_{ij} be the (i, j) entry; that is, the entry in the i -th row and j -th column.

Definition 1.3.4 — transpose. The *transpose* of an m by n matrix $A = (a_{ij})$ is the n by m matrix A^T obtained by switching the columns and rows. That is, the entries are given by $(A^T)_{ij} = A_{ji}$.

If $A = A^T$ then we say A is *symmetric*. (Then $A_{ij} = A_{ji}$ for all i, j , and necessarily $m = n$.)

■ **Example 1.12**

$$A = \begin{bmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \end{bmatrix}$$

which is a 2×3 matrix. Its transpose is

$$A^T = \begin{bmatrix} 1 & 4 \\ 2 & 5 \\ 3 & 6 \end{bmatrix}$$

which is a 3×2 matrix. ■

■ **Example 1.13** Let

$$A = \begin{bmatrix} 2 & -1 & 0 \\ -1 & 3 & 1 \\ 0 & 1 & 4 \end{bmatrix},$$

then A is a symmetric matrix. ■

For a vector $\mathbf{v} \in \mathbb{R}^n$, we can view it as an $n \times 1$ matrix. Then $\mathbf{v}^T$ is a $1 \times n$ matrix. Then using matrix multiplication, we have the following relationship:

Lemma 1.3.9 For $\mathbf{v}, \mathbf{w} \in \mathbb{R}^n$,

$$\langle \mathbf{v}, \mathbf{w} \rangle = \sum_{i=1}^n v_i w_i = \mathbf{v}^T \mathbf{w} = \mathbf{w}^T \mathbf{v}.$$

Corollary 1.3.10 Let A be a symmetric $n \times n$ matrix. Then

$$\langle A\mathbf{v}, \mathbf{w} \rangle = (A\mathbf{v})^T \mathbf{w} = \mathbf{v}^T A \mathbf{w} = \langle \mathbf{v}, A\mathbf{w} \rangle \quad \forall \mathbf{v}, \mathbf{w} \in \mathbb{R}^n.$$

Proposition 1.3.11 — properties of matrix transpose. Let A and B be matrices and c be a scalar.

1. $(A^T)^T = A$.
2. $(A + B)^T = A^T + B^T$.
3. $(cA)^T = cA^T$.
4. $(AB)^T = B^T A^T$.

By induction, this result extends to the general case of multiple matrices, so $(A_1 A_2 \dots A_k)^T = A_k^T \dots A_2^T A_1^T$.

■ **Example 1.14** Let $U \subset \mathbb{R}^n$ be open, and $g : U \rightarrow \mathbb{R}$ be a function differentiable on U . Notice that for g , in the matrix notation, the derivative $Dg(x)$ is a $1 \times n$ row vector, while we define the gradient $\nabla g(x)$ as an $n \times 1$ column vector. In other words, $\nabla g(x) = (Dg(x))^T$.

Thus if we have functions $f : U \rightarrow \mathbb{R}^m$, $g : \mathbb{R}^m \rightarrow \mathbb{R}$, using the chain rule for the derivative

$$D(g \circ f)(x_0) = Dg(f(x_0))Df(x_0),$$

we get chain rule for the gradient by taking transpose:

$$\nabla(g \circ f)(x_0) = (D(g \circ f)(x_0))^T = \left(Dg(f(x_0))Df(x_0) \right)^T = (Df(x_0))^T \nabla g(f(x_0)).$$

■

Definition 1.3.5 — eigenvalue, eigenvector. Suppose A is an n by n matrix. λ is an *eigenvalue* of A if there exists a vector $0 \neq \mathbf{v}$ such that $A\mathbf{v} = \lambda\mathbf{v}$. The vector $\mathbf{v}$ is called an *eigenvector*.

The set of all eigenvectors of A corresponding to the same eigenvalue λ , together with the zero vector, is called an *eigenspace*. The dimension of the eigenspace corresponding to λ is called the (geometric) *multiplicity* of λ .

If a set of eigenvectors of A forms a basis of $\mathbb{R}^n$, then this basis is called an *eigenbasis*.

R Geometrically, a matrix's eigenvectors are those vectors that are only stretched or shrunk, with neither rotation nor shear. The corresponding eigenvalue is the factor by which an eigenvector is stretched or shrunk. If the eigenvalue is negative, the eigenvector's direction is reversed.

■ **Example 1.15 — identity matrix.** We let I_n be the $n \times n$ identity matrix, i.e., the matrix where the diagonal entries are all 1 and all the other entries are 0. Then it has n eigenvalues all equal to 1, and every nonzero vector is an eigenvector. ■

■ **Example 1.16 — diagonal matrix.** Let A be a diagonal matrix. This means only the diagonal entries a_{jj} are nonzero. Then we can check that $Ae_j = a_{jj}e_j$, so each e_j is an eigenvector with eigenvalue a_{jj} . ■

■ **Example 1.17 — 2×2 case.** For a 2×2 matrix $A = \begin{bmatrix} a & b \\ c & d \end{bmatrix}$, the eigenvalues of A are the roots of the *characteristic polynomial*

$$\lambda^2 - (a + d)\lambda + (ad - bc) = 0.$$

Notice that this quadratic equation could have two distinct real roots, a repeated real root, or no real roots (in this case two complex roots). ■

■ **Example 1.18 — rotation matrix.** A rotation matrix in $\mathbb{R}^2$ for an angle θ is given by:

$$R(\theta) = \begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix}$$

This matrix rotates a vector counterclockwise by an angle θ when multiplied by the vector.

The characteristic polynomial is

$$\lambda^2 - 2 \cos \theta \cdot \lambda + 1.$$

Since the discriminant is $\Delta = 4 \cos^2 \theta - 4 \leq 0$, when $\theta \neq k\pi$, we see that there are no real roots. As a result, there is no real eigenvalue of $R(\theta)$. Geometrically, this makes sense because when $\theta \neq k\pi$, no vector's direction is preserved or reversed, so there aren't any eigenvectors. ■

Definition 1.3.6 An n by n square matrix A is called *invertible* (also called nonsingular or nondegenerate) if there exists an n by n square matrix B such that

$$AB = BA = I_n.$$

If this is the case, then the matrix B is uniquely determined by A , and is called the (multiplicative) inverse of A , denoted by A^{-1} .

Proposition 1.3.12 Let A be an $n \times n$ matrix. The following statements are equivalent:

1. A is invertible.
2. The linear transformation mapping x to Ax is bijective; that is, the equation $Ax = b$ has exactly one solution for each b in $\mathbb{R}^n$.

3. The transpose A^T is an invertible matrix.
4. The columns of A form a basis of $\mathbb{R}^n$.
5. The rows of A form a basis of $\mathbb{R}^n$.
6. The number 0 is not an eigenvalue of A .

■ **Example 1.19** The rotation matrix $R(\theta)$ is invertible. Indeed, since geometrically this matrix just rotates a vector counterclockwise by an angle θ , its inverse is rotation by an angle $-\theta$. Thus $(R(\theta))^{-1} = R(-\theta)$. ■

■ **Example 1.20** Let

$$A = \begin{bmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \\ 7 & 8 & 9 \end{bmatrix},$$

then A is not invertible.

We can see this by noticing the columns are not linearly independent: If we let $\mathbf{c}_1, \mathbf{c}_2, \mathbf{c}_3$ to denote the columns, then we see that $\mathbf{c}_2 - \mathbf{c}_1 = (1, 1, 1) = \mathbf{c}_3 - \mathbf{c}_2$. ■

■ **Example 1.21** For a diagonal matrix A , it is invertible if and only if each $a_{jj} \neq 0$. And A^{-1} is the diagonal matrix with diagonal entries a_{jj}^{-1} . ■

Proposition 1.3.13 — properties of matrix inverse. Let A and B be invertible matrices and c be a nonzero scalar.

1. $(A^{-1})^{-1} = A$.
2. $(cA)^{-1} = c^{-1}A^{-1}$.
3. $(A^T)^{-1} = (A^{-1})^T$.
4. $(AB)^{-1} = B^{-1}A^{-1}$.

By induction, this result extends to the general case of multiple matrices, so $(A_1 A_2 \dots A_k)^{-1} = A_k^{-1} \dots A_2^{-1} A_1^{-1}$.

In general, for real $n \times n$ matrix, it will have at most n real eigenvalues, counted with multiplicity, and we have seen as above that even for the 2×2 case, there might be no real eigenvalue at all. However, if we restrict our attention to symmetric matrices, the situation becomes nicer, and we have:

Theorem 1.3.14 — spectral theorem. Let A be a symmetric $n \times n$ matrix. Then A has n real eigenvalues $\lambda_1, \dots, \lambda_n$ counted with multiplicity, and has an corresponding orthonormal eigenbasis $\{\mathbf{v}_1, \dots, \mathbf{v}_n\}$ with $A\mathbf{v}_i = \lambda_i \mathbf{v}_i$.

R Here, orthonormal just means each vector has unit norm $\|\mathbf{v}_i\| = 1$, and different ones are orthogonal $\langle \mathbf{v}_i, \mathbf{v}_j \rangle = 0$ for $i \neq j$.

Definition 1.3.7 — positive definite, negative definite. A symmetric $n \times n$ matrix A is

- *positive definite* if for all $\mathbf{v} \in \mathbb{R}^n \setminus \{0\}$, $\mathbf{v}^T A \mathbf{v} > 0$.
- *positive semi-definite* if for all $\mathbf{v} \in \mathbb{R}^n \setminus \{0\}$, $\mathbf{v}^T A \mathbf{v} \geq 0$.
- *negative definite* if for all $\mathbf{v} \in \mathbb{R}^n \setminus \{0\}$, $\mathbf{v}^T A \mathbf{v} < 0$.
- *negative semi-definite* if for all $\mathbf{v} \in \mathbb{R}^n \setminus \{0\}$, $\mathbf{v}^T A \mathbf{v} \leq 0$.

■ **Example 1.22** For any $m \times n$ matrix A , the $n \times n$ matrix $A^T A$ is symmetric. Indeed, we check that $(A^T A)^T = A^T (A^T)^T = A^T A$.

Further, it is positive semi-definite: for any $v \in \mathbb{R}^n$,

$$v^T(A^T A)v = (v^T A^T)(Av) = (Av)^T(Av) = \langle Av, Av \rangle = \|Av\|^2 \geq 0.$$

■ **Example 1.23** With the spectral theorem, we can compute $v^T Av$ using the eigenvalues $\lambda_1, \dots, \lambda_n$ and the orthonormal eigenbasis $\mathbf{v}_1, \dots, \mathbf{v}_n$.

We can write $v = \sum_{i=1}^n a_i \mathbf{v}_i$ for some $a_i \in \mathbb{R}$, so

$$\|v\|^2 = \langle v, v \rangle = \left\langle \sum_{j=1}^n a_j \mathbf{v}_j, \sum_{i=1}^n a_i \mathbf{v}_i \right\rangle = \sum_{j=1}^n \sum_{i=1}^n a_i a_j \langle \mathbf{v}_i, \mathbf{v}_j \rangle = \sum_{i=1}^n a_i^2.$$

and

$$\begin{aligned} v^T Av &= \langle v, Av \rangle \\ &= \left\langle \sum_{j=1}^n a_j \mathbf{v}_j, A \left(\sum_{i=1}^n a_i \mathbf{v}_i \right) \right\rangle \\ &= \left\langle \sum_{j=1}^n a_j \mathbf{v}_j, \sum_{i=1}^n a_i A \mathbf{v}_i \right\rangle \\ &= \left\langle a_j \sum_{j=1}^n \mathbf{v}_j, \sum_{i=1}^n a_i \lambda_i \mathbf{v}_i \right\rangle \\ &= \sum_{j=1}^n \sum_{i=1}^n \lambda_i a_i a_j \langle \mathbf{v}_i, \mathbf{v}_j \rangle \\ &= \sum_{i=1}^n \lambda_i a_i^2. \end{aligned}$$

Using the above computation, we see that we can characterize the definiteness of a symmetric matrix using eigenvalues:

Proposition 1.3.15 A symmetric matrix A is

- *positive definite* if and only if all its eigenvalues are > 0 .
- *positive semi-definite* if and only if all its eigenvalues are ≥ 0 .
- *negative definite* if and only if all its eigenvalues are < 0 .
- *negative semi-definite* if and only if all its eigenvalues are ≤ 0 .

Finally, we are able to give a characterization of the operator norm using eigenvalues.

Theorem 1.3.16 — operator norm in terms of eigenvalue. Let $A : \mathbb{R}^n \rightarrow \mathbb{R}^m$ be a linear map. We view it as a matrix and also let A to denote the matrix. Then $\|A\|$ equals the square root of the largest eigenvalue of $A^T A$.

Proof. Recall $\|A\| = \max_{\|\mathbf{v}\|=1} \|A\mathbf{v}\|$. For any unit vector $\mathbf{v}$, we compute

$$\|A\mathbf{v}\|^2 = \langle A\mathbf{v}, A\mathbf{v} \rangle = (A\mathbf{v})^T A\mathbf{v} = \mathbf{v}^T (A^T A) \mathbf{v}.$$

Then if we denote the eigenvalues of $A^T A$ by $\lambda_1 \geq \dots \geq \lambda_n$ and the orthonormal eigenbasis by $\mathbf{v}_1, \dots, \mathbf{v}_n$, by the computations above we know each $\lambda_i \geq 0$. Further, if we write $\mathbf{v} = \sum_{i=1}^n a_i \mathbf{v}_i$,

then $1 = \sum_{i=1}^n a_i^2$, and $\mathbf{v}^T (A^T A) \mathbf{v} = \sum_{i=1}^n \lambda_i a_i^2$. Then we have

$$\|A\mathbf{v}\|^2 = \sum_{i=1}^n \lambda_i a_i^2 \leq \sum_{i=1}^n \lambda_1 a_i^2 = \lambda_1.$$

On the other other hand, if we choose $\mathbf{v} = \mathbf{v}_1$, then $\|A\mathbf{v}_1\|^2 = \lambda_1$. Thus we have

$$\|A\| = \max_{\|\mathbf{v}\|=1} \|A\mathbf{v}\| = \|A\mathbf{v}_1\| = \sqrt{\lambda_1}$$

as desired. ■

1.4 The mean value theorem in $\mathbb{R}^n$

Next we consider the Mean Value Theorem for functions of several variables.

Recall that the mean value theorem for single variable functions: if $f : [a, b] \rightarrow \mathbb{R}$ is continuous on $[a, b]$ and differentiable on (a, b) , then there exists $c \in (a, b)$ such that $f(b) - f(a) = f'(c) \cdot (b - a)$.

As in the case of functions of one variable, the Mean Value Theorem relates the average rate of change of a function in a specified direction to the instantaneous rate of change at a particular point as measured by the derivative. In the case of a function of one variable, there is little choice about the so-called specified direction. But, when we consider functions from $\mathbb{R}^n$ to $\mathbb{R}^m$, we find that it is necessary to specify a direction both in the domain and in the range of the function in order to give a proper interpretation to the Mean Value Theorem.

We first look at the domain. For two distinct points $x, y \in \mathbb{R}^n$, there are many ways to go from one to another. Thus we introduce the special maps called differentiable curves.

Definition 1.4.1 — differentiable curve, tangent vector. Suppose $(a, b) \subset \mathbb{R}$. A map $\gamma : (a, b) \rightarrow \mathbb{R}^n$ is called a *differentiable curve* if γ is differentiable as a function and $D\gamma \neq 0$ on (a, b) . The n -vector $D\gamma(s)$ is called the *tangent vector* of γ at $\gamma(s)$.

It is natural to define a *curve with end points*, namely we consider $\gamma : [a, b] \rightarrow \mathbb{R}^n$, where γ is continuous on $[a, b]$ and differentiable on (a, b) and $D\gamma \neq 0$ on (a, b) .

R Sometimes, people say the image of f is a curve. During this course, we will view the map itself as the “curve”.

Then a curve $\gamma : [a, b] \rightarrow \mathbb{R}^n$ connects the two points $x = \gamma(a)$ and $y = \gamma(b)$ in $\mathbb{R}^n$.

■ **Example 1.24** The simplest curve is of course the straight-line connecting two points. Let $x, y \in \mathbb{R}^n$ be distinct. Then we can define $\gamma : [0, 1] \rightarrow \mathbb{R}^n$ by

$$\gamma(t) = ty + (1 - t)x.$$

We can compute $D\gamma(t) = y - x$. ■

■ **Example 1.25** Suppose $\gamma : [0, 2\pi] \rightarrow \mathbb{R}^2$ is defined by

$$\gamma(\theta) = (\cos \theta, \sin \theta).$$

Then the curve is the unit circle. We can compute $D\gamma(\theta) = (-\sin \theta, \cos \theta)$. ■

■ **Example 1.26** Suppose $\gamma : \mathbb{R} \rightarrow \mathbb{R}^3$ is defined by

$$\gamma(s) = (\cos s, \sin s, s).$$

Then the curve is the helix. It is known to be the model of DNA. We can compute $D\gamma(s) = (-\sin s, \cos s, 1)$. ■

With the notion of a differentiable curve, we can formulate the Mean Value Theorem for functions from $\mathbb{R}^n$ to $\mathbb{R}$:

Theorem 1.4.1 — Mean value theorem for multivariable functions to $\mathbb{R}$. Suppose $U \subset \mathbb{R}^n$ is open, and $f : U \rightarrow \mathbb{R}$ is a differentiable function. Suppose $\gamma : [a, b] \rightarrow U$ is a curve with endpoints $x = \gamma(a)$ and $y = \gamma(b)$, where $x, y \in U$. Then there exists $s \in (a, b)$ such that

$$f(y) - f(x) = Df(\gamma(s))D\gamma(s)(b - a).$$

Proof. This is an application of the mean value theorem in one variable and the chain rule.

Define $F : [a, b] \rightarrow \mathbb{R}$ by $F(t) = f(\gamma(t))$.

Then by the chain rule, F as the composition of f and γ is continuous on $[a, b]$ and differentiable on (a, b) , so by the mean value theorem, there exists $s \in (a, b)$ such that $F(b) - F(a) = (F'(s))(b - a)$.

The left-hand side is $f(y) - f(x)$, and by the chain rule,

$$F'(s) = Df(\gamma(s))D\gamma(s).$$

■

■ **Example 1.27** Let $U = \mathbb{R}^2$, and let $f : \mathbb{R}^2 \rightarrow \mathbb{R}$ be defined by

$$f(x, y) = x^2 + y^2.$$

Let the curve $\gamma : [0, 1] \rightarrow \mathbb{R}^2$ be defined as

$$\gamma(t) = (t, t^2), \quad \text{for } t \in [0, 1].$$

Then the endpoints of the curve are $x = \gamma(0) = (0, 0)$, $y = \gamma(1) = (1, 1)$.

By the above theorem, there exists some $s \in (0, 1)$ such that

$$f(\gamma(1)) - f(\gamma(0)) = Df(\gamma(s))D\gamma(s).$$

Here, $f(\gamma(0)) = 0$, and $f(\gamma(1)) = 2$. The derivative $Df(x, y)$ is a 1×2 row vector

$$Df(x, y) = \left[\frac{\partial f}{\partial x} \quad \frac{\partial f}{\partial y} \right] = [2x \quad 2y],$$

so $Df(\gamma(s)) = [2s \quad 2s^2]$.

The derivative $D\gamma(t)$ is a 2×1 column vector

$$D\gamma(t) = \begin{bmatrix} \frac{d}{dt}t \\ \frac{d}{dt}t^2 \end{bmatrix} = \begin{bmatrix} 1 \\ 2t \end{bmatrix}.$$

Thus, at $s \in (0, 1)$, we have

$$2 = 2s + 4s^3.$$

The mean value theorem tells us the above equation has a solution in $(0, 1)$. Indeed, solving the above equation, we find $s \approx 0.59$. ■

The most natural choice of a curve is of course the straight-line, and in that case we have the following theorem:

Theorem 1.4.2 Let U be an open set in $\mathbb{R}^n$, and let $f : U \rightarrow \mathbb{R}$ be differentiable. Let x, y be two distinct points in U such that the line segment joining x to y lies entirely in U . Then, there exists $\xi \in (0, 1)$ such that

$$f(y) - f(x) = Df(z)(y - x),$$

where $z = (1 - \xi)x + \xi y$.

Proof. In the previous theorem, take $\gamma : [0, 1] \rightarrow \mathbb{R}^n$ be defined by $\gamma(t) = ty + (1 - t)x$. Then $D\gamma(t) = y - x$. ■

Recall that we learned the notion of a convex set in Math 203:

Definition 1.4.2 — convex set. Let $A \subset \mathbb{R}^n$. We say that A is *convex* if, for any $x, y \in A$, we have

$$\{ty + (1 - t)x \mid t \in [0, 1]\} \subset A.$$

Corollary 1.4.3 — Mean value theorem in a convex set. Let U be a convex open set in $\mathbb{R}^n$, and let $f : U \rightarrow \mathbb{R}$ be differentiable. Let x, y be two distinct points in U . Then, there exists $\xi \in (0, 1)$ such that

$$f(y) - f(x) = Df(z)(y - x),$$

where $z = (1 - \xi)x + \xi y$.

Similar to the one-dimensional case, we also have:

Corollary 1.4.4 Let U be a convex open set in $\mathbb{R}^n$ and let $f : U \rightarrow \mathbb{R}$ be differentiable. If $Df(x) = 0$ for all $x \in U$, then f is constant on U .

Proof. Exercise. ■

Next we want to consider the general case when the target is $\mathbb{R}^m$. In this case, in addition to having a direction implicitly chosen for us in the domain (either along a curve or in the straight-line direction), we must also explicitly choose a direction in $\mathbb{R}^m$ in order to make sense of the mean value property.

Theorem 1.4.5 — Mean value theorem for multivariable functions to $\mathbb{R}^m$. Let U be an open subset of $\mathbb{R}^n$ and let $f : U \rightarrow \mathbb{R}^m$ be differentiable on U . For any two distinct points $x, y \in U$ such that the line segment joining x to y lies entirely in U , and any vector $\mathbf{v} \in \mathbb{R}^m$, there exists $\xi \in (0, 1)$ such that

$$\mathbf{v} \cdot (f(y) - f(x)) = \mathbf{v} \cdot (Df(z)(y - x)),$$

where $z = (1 - \xi)x + \xi y$.

R We note in connection to the discussion above that if $\mathbf{v}$ is a unit vector, the expression $\mathbf{v} \cdot (f(y) - f(x))$ is the component of $f(y) - f(x)$ in the direction of the vector $\mathbf{v}$. A similar statement is true for the expression $\mathbf{v} \cdot (Df(z)(y - x))$.

If we write $f = (f_1, \dots, f_m)$ and take $\mathbf{v} = e_j$, then this theorem just reduces to the mean value theorem for f_j .

Corollary 1.4.6 Let U be a convex open set in $\mathbb{R}^n$ and let $f : U \rightarrow \mathbb{R}^m$ be differentiable. If $Df(x) = 0$ for all $x \in U$, then f is constant on U .

The significance of the multivariable version of the Mean Value Theorem is that the direction vector $\mathbf{v}$ is arbitrary. However, in general, there is no single ξ that will satisfy the conclusion of the Mean Value Theorem for all $\mathbf{v}$ simultaneously. This is shown in the following exercise.

Exercise 1.7 Show that the function $f : \mathbb{R} \rightarrow \mathbb{R}^2$ given by $f(x) = (x^2, \sin x)$ does not satisfy the property that, given any $x, y \in \mathbb{R}$ with $x < y$, there exists $z \in (x, y)$ such that $f(y) - f(x) = Df(z)(y - x)$. ■

1.5 Higher order partial derivatives and Taylor's theorem

The next natural question is, “What about second derivatives?”

Recall that we can identify a linear map $T : \mathbb{R}^n \rightarrow \mathbb{R}^m$ with a $m \times n$ matrix, so the space of all such linear maps can be identified with the space of all such matrices. But an $m \times n$ matrix has mn entries and we are free to choose each of these entries, so the space of all $m \times n$ matrices can be naturally identified with $\mathbb{R}^{mn}$.

Definition 1.5.1 We denote the space of all linear maps from $\mathbb{R}^n$ to $\mathbb{R}^m$ is by $\mathcal{L}(\mathbb{R}^n, \mathbb{R}^m)$. We have $\mathcal{L}(\mathbb{R}^n, \mathbb{R}^m) \cong \mathbb{R}^{mn}$.

Now if $f : \mathbb{R}^n \rightarrow \mathbb{R}^m$ is differentiable at a point x , then $Df(x) \in \mathcal{L}(\mathbb{R}^n, \mathbb{R}^m) \cong \mathbb{R}^{mn}$. Since at each point x where f is differentiable we get a linear map, we think of the symbol Df as a map from a subset U of $\mathbb{R}^n$ to $\mathcal{L}(\mathbb{R}^n, \mathbb{R}^m) \cong \mathbb{R}^{mn}$.

Then if we keep doing this and take an l -th order derivative of a function f , we get a map from a subset U of $\mathbb{R}^n$ to $\mathbb{R}^{mnl}$. This is a space with huge dimensions and is hard to work with. In general, we can only write down a higher order derivative in coordinates.

Thus, we look at partial derivatives. For a function $f : \mathbb{R}^n \rightarrow \mathbb{R}$, the partial derivative $D_j f$ can be thought of as a function from $\mathbb{R}^n$ to $\mathbb{R}$. Then we can take the partial derivatives of partial derivatives.

Definition 1.5.2 — mixed partial derivative. Consider a function $f = (f_1, f_2, \dots, f_m)$ from an open set in $\mathbb{R}^n$ to $\mathbb{R}^m$. Then, we define the *mixed partial derivative*

$$D_{ij} f_k = D_i (D_j f_k)$$

assuming that the $D_j f_k$ has partial derivatives for all j, k .

For higher order derivatives, we use the notation

$$D_{i_1, i_2, \dots, i_l} f_k(x_0) = D_{i_1} (D_{i_2} (\dots D_{i_l} f_k)) (x_0).$$

In general, it would be hard to understand higher-order derivatives (actually, it is even hard to write down such notations!). We will use second-order derivatives as examples to understand higher-order derivatives.

Recall the approximation form of partial derivatives:

$$D_1 f(x) h \approx f(x + h e_1) - f(x).$$

Then we should have

$$\begin{aligned} D_{21} f(x) h^2 &\approx D_1 f(x + h e_2) h - D_1 f(x) h \\ &\approx (f(x + h e_2 + h e_1) - f(x + h e_2)) - (f(x + h e_1) - f(x)) \\ &= f(x + h e_2 + h e_1) - f(x + h e_2) - f(x + h e_1) + f(x). \end{aligned}$$

Similarly, we have

$$D_{12}f(x) h^2 \approx f(x + he_2 + he_1) - f(x + he_2) - f(x + he_1) + f(x).$$

So we expect that $D_{12}f = D_{21}f$. Of course, we should be careful with limits in higher dimensions, and such an identity is not always true.

Exercise 1.8 Let $f(x, y) = \begin{cases} 0 & \text{when } (x, y) = (0, 0), \\ \frac{x^3y - xy^3}{x^2 + y^2} & \text{otherwise.} \end{cases}$

Show that f is differentiable. Show that $D_{12}f(0, 0)$, $D_{21}f(0, 0)$ exist, but $D_{12}f(0, 0) \neq D_{21}f(0, 0)$. ■

Nevertheless, we have the following theorem:

Theorem 1.5.1 Let $U \subset \mathbb{R}^n$ be open and let $f : U \rightarrow \mathbb{R}$ be differentiable. Suppose that $D_{ij}f$ and $D_{ji}f$ exist and are continuous on U . Then $D_{ij}f(x) = D_{ji}f(x)$ for all $x \in U$.

Proof. The proof adapts the approximation idea that we discussed above. To do this, we will use the mean value theorem. The idea is similar to the proof of Theorem 1.2.7.

Let us prove that $D_{12}f(x) = D_{21}f(x)$, and the proof for other indices is the same.

For $(y_1, y_2) \neq (0, 0)$, define

$$g(x, y_1, y_2) = \frac{f(x + y_2e_2 + y_1e_1) - f(x + y_2e_2) - f(x + y_1e_1) + f(x)}{y_1y_2}.$$

We fix y_1 and then define a single-variable function $\alpha(z_2)$ to measure the error in the e_2 direction:

$$\alpha(z) = f(x + ze_2 + y_1e_1) - f(x + ze_2) - f(x + y_1e_1) + f(x).$$

The function $\alpha(z)$ is differentiable because f is. Then the single-variable Mean Value Theorem applied to $\alpha(z)$ shows that there exists z_2 between 0 and y_2 such that

$$\alpha(y_2) - \alpha(0) = \alpha'(z_2)y_2.$$

Thus

$$g(x, y_1, y_2)y_1y_2 - 0 = (D_2f(x + z_2e_2 + y_1e_1) - D_2f(x + z_2e_2))y_2,$$

so

$$g(x, y_1, y_2)y_1 = D_2f(x + z_2e_2 + y_1e_1) - D_2f(x + z_2e_2).$$

Then we fix this z_2 and consider the function

$$\beta(z) = D_2f(x + z_2e_2 + ze_1) - D_2f(x + z_2e_2).$$

It is differentiable because D_{12} exists and is continuous. Thus we can apply the single-variable Mean Value Theorem to it and find z_1 between 0 and y_1 such that

$$\beta(y_1) - \beta(0) = \beta'(z_1)y_1.$$

Thus

$$g(x, y_1, y_2)y_1 - 0 = D_{12}f(x + z_2e_2 + z_1e_1)y_1,$$

so

$$g(x, y_1, y_2) = D_{12}f(x + z_2e_2 + z_1e_1).$$

Likewise, if we consider D_{21} instead, then we find that there exist w_1 between 0 and y_1 , and w_2 between 0 and y_2 , such that

$$g(x, y_1, y_2) = D_{21}f(x + w_2e_2 + w_1e_1).$$

In particular, if we let $(y_1, y_2) \rightarrow (0, 0)$ with $y_1 \neq 0$, $y_2 \neq 0$, then by the continuity of $D_{12}f$ and $D_{21}f$, we conclude that

$$D_{12}f(x) = D_{21}f(x).$$

■

Similar to the one-variable case, we can define the class C^k .

Definition 1.5.3 — C^k functions. Let $U \subset \mathbb{R}^n$ be an open set and $f : U \rightarrow \mathbb{R}^m$. The function f is said to be in $C^k(U, \mathbb{R}^m)$ (or C^k on U) if f has all continuous partial derivatives up to and including order k on U . The function f is said to be in $C^\infty(U, \mathbb{R}^m)$ or *smooth* on U if f has all partial derivatives of all orders on U .

Now, suppose $U \subset \mathbb{R}^n$ and $f : U \rightarrow \mathbb{R}$ is C^k on U . Let $\ell \leq k$ and let $i_1, i_2, \dots, i_\ell$ be a collection of integers between 1 and n . Then it is not hard to show that for any $\sigma \in S_\ell$, the symmetric group on ℓ letters,

$$D_{i_1, i_2, \dots, i_\ell} f = D_{i_{\sigma(1)}, i_{\sigma(2)}, \dots, i_{\sigma(\ell)}} f.$$

That is, mixed partial derivatives of a C^k function are symmetric under permutation of the indices.

As a result, for a C^k function f , we don't need to worry about the order we take the partial derivatives, and every partial derivative of order k can be written in the form

$$D_1^{\alpha_1} D_2^{\alpha_2} \cdots D_n^{\alpha_n} f,$$

where the α_i are nonnegative integers and

$$\sum_{i=1}^n \alpha_i = k.$$

There is an interesting combinatorial problem that arises here: how many ways can I write down $D_1^{\alpha_1} D_2^{\alpha_2} \cdots D_n^{\alpha_n} f$?

Exercise 1.9 This is an exercise in high school algebra.

Consider the polynomial in n variables:

$$p(x_1, x_2, \dots, x_n) = (x_1 + x_2 + \cdots + x_n)^k.$$

Show that, upon expanding this into monomials, the coefficient of the term

$$x_1^{k_1} x_2^{k_2} \cdots x_n^{k_n}, \quad \text{where } \sum_{i=1}^n k_i = k,$$

is given by

$$\frac{k!}{k_1! k_2! \cdots k_n!}.$$

This expression is called a *multinomial coefficient* and is denoted by

$$\binom{k}{k_1, k_2, \dots, k_n}.$$

Now, with these preparations, we want to generalize Taylor's theorem to multivariable functions. Again, we first use the idea of approximation to understand the correct formulation of the Taylor theorem.

Suppose f is a sufficiently smooth (i.e., C^{k+1}) function. Then, we can expand f using the single-variable Taylor theorem, applied one variable at a time. For example, we would get:

$$\begin{aligned}
 f(x_1 + y_1, \dots, x_n + y_n) &\approx \sum_{0 \leq \alpha_1 \leq k} \frac{y_1^{\alpha_1}}{\alpha_1!} D_1^{\alpha_1} f(x_1, x_2 + y_2, \dots, x_n + y_n) \\
 &\approx \sum_{0 \leq \alpha_2 \leq k_2} \sum_{0 \leq \alpha_1 \leq k_1} \frac{y_1^{\alpha_1} y_2^{\alpha_2}}{\alpha_1! \alpha_2!} D_2^{\alpha_2} D_1^{\alpha_1} f(x_1, x_2, \dots, x_n + y_n) \\
 &\vdots \\
 &\approx \sum_{0 \leq \alpha_n \leq k_n} \dots \sum_{0 \leq \alpha_1 \leq k_1} \frac{y_1^{\alpha_1} y_2^{\alpha_2} \dots y_n^{\alpha_n}}{\alpha_1! \alpha_2! \dots \alpha_n!} D_n^{\alpha_n} \dots D_1^{\alpha_1} f(x_1, x_2, \dots, x_n)
 \end{aligned}$$

Here $k_1 + \dots + k_n = k$.

To express this more compactly, we introduce the multi-index notation:

Let $\alpha = (\alpha_1, \alpha_2, \dots, \alpha_n)$, where each $\alpha_i \in \mathbb{Z}_{\geq 0}$. Then:

- $|\alpha| := \alpha_1 + \alpha_2 + \dots + \alpha_n$
- $D^\alpha := D_1^{\alpha_1} D_2^{\alpha_2} \dots D_n^{\alpha_n}$
- $x^\alpha := x_1^{\alpha_1} x_2^{\alpha_2} \dots x_n^{\alpha_n}$
- $\alpha! := \alpha_1! \alpha_2! \dots \alpha_n!$

With this notation, I can write

$$f(x_1 + y_1, \dots, x_n + y_n) \approx \sum_{|\alpha| \leq k} \frac{y^\alpha}{\alpha!} D^\alpha f(x).$$

Now we can state the Taylor theorem for functions $f : \mathbb{R}^n \rightarrow \mathbb{R}$:

Theorem 1.5.2 — Taylor's Theorem for Multivariable Functions. Suppose $U \subset \mathbb{R}^n$ is an open convex set, and $f : U \rightarrow \mathbb{R}$ is a C^{k+1} function. Then for any $x, y \in U$, $x \neq y$, write $y = x + th$, where h is a unit vector in $\mathbb{R}^n$, and $t > 0$. Then there exists $s \in \mathbb{R}$, $0 < s < t$, such that

$$f(y) = \sum_{|\alpha| \leq k} \frac{D^\alpha f(x)}{\alpha!} h^\alpha t^{|\alpha|} + \sum_{|\alpha| = k+1} \frac{D^\alpha f(x + sh)}{\alpha!} h^\alpha t^{k+1}$$

Proof. Just as the multivariate mean value theorem is the consequence of the chain rule and the single-variable case, this is the consequence of the chain rule and the single-variable Taylor's Theorem. Consider the function $F(a) = f(x + ah)$. Then

$$F'(a) = Df(x + ah)h = \sum_{i=1}^n D_i f(x + ah) h_i.$$

So

$$\begin{aligned}
 F^{(k)}(a) &= \sum_{i_k=1}^n \dots \sum_{i_1=1}^n D_{i_k, \dots, i_1} f(x + ah) \cdot h_{i_1} \dots h_{i_k} \\
 &= \sum_{|\alpha|=k} \frac{k!}{\alpha!} D^\alpha f(x + ah) \cdot h^\alpha.
 \end{aligned}$$

Here we used Exercise 1.9 to say there are $\frac{k!}{\alpha!}$ different ways to write down $D^\alpha f(x + ah)$. Apply Corollary 1.1.11 to $F(t)$, and we get

$$F(t) = F(0) + F'(0)t + \frac{F''(0)}{2!}t^2 + \cdots + \frac{F^{(k)}(0)}{k!}t^k + \frac{1}{(k+1)!}F^{(k+1)}(s)t^{k+1}$$

for some $0 < s < t$, so plugging in what $F^{(k)}(0)$ is yields the result. ■

Definition 1.5.4 In the above theorem, $\sum_{|\alpha| \leq k} \frac{D^\alpha f(x)}{\alpha!} h^\alpha t^{|\alpha|}$ is called the *Taylor polynomial of degree k* for the function f at x in the direction h .

$$R_k(x, y) := \sum_{|\alpha|=k+1} \frac{D^\alpha f(x + sh)}{\alpha!} h^\alpha t^{k+1}$$

is called the *remainder term*.

1.6 Hypersurfaces and tangent hyperplanes in $\mathbb{R}^n$

1.6.1 General definitions

The derivative of a function of one variable $f : \mathbb{R} \rightarrow \mathbb{R}$ that is differentiable at a point $x = c$ is the slope of the tangent line to the curve $y = f(x)$ at $(c, f(c))$. The equation of the line tangent to the curve at that point can be written $y = f(c) + f'(c)(x - c)$.

Now, instead of a function, we can more generally think about curves in $\mathbb{R}^2$. Before we've seen curves defined explicitly by parametrization; here we can also see curves defined implicitly as a level curve of a two-variable function to $\mathbb{R}$.

■ **Example 1.28** For example, consider the unit circle in $\mathbb{R}^2$. Before we can view it as the image of $\gamma : [0, 2\pi] \rightarrow \mathbb{R}^2$, $\gamma(\theta) = (\cos \theta, \sin \theta)$.

Here we view it implicitly as the level curve $\{(x, y) \in \mathbb{R}^2 \mid F(x, y) = 1\}$ of the function $F(x, y) = x^2 + y^2$. ■

The function $y = f(x)$ may be considered as the level curve $\{(x, y) \in \mathbb{R}^2 \mid F(x, y) = 0\}$ of the function $F : \mathbb{R}^2 \rightarrow \mathbb{R}$ given by $F(x, y) = f(x) - y$.

We now assume that $F : \mathbb{R}^2 \rightarrow \mathbb{R}$ is any C^1 function. Consider the level curve $F(x, y) = k$ in $\mathbb{R}^2$. Let (x_0, y_0) be a point on this curve, and assume that the gradient $\nabla F(x_0, y_0) \neq (0, 0)$. The tangent line to the curve at the point (x_0, y_0) is the line through the point perpendicular to $\nabla F(x_0, y_0)$. The equation of this line is then

$$\nabla F(x_0, y_0) \cdot (x - x_0, y - y_0) = 0,$$

or equivalently,

$$D_1 F(x_0, y_0)(x - x_0) + D_2 F(x_0, y_0)(y - y_0) = 0.$$

R In the case $F(x, y) = f(x) - y$, the equation of the tangent line is the same as that given in the first paragraph above.

Now we see that the level curve of one equation in $\mathbb{R}^2$ gives a curve. Now we want to generalize this idea to higher dimensions. We consider a C^1 function $F : \mathbb{R}^N \rightarrow \mathbb{R}$.

Definition 1.6.1 — smooth hypersurface. Let $F : \mathbb{R}^N \rightarrow \mathbb{R}$ be a C^1 function. A smooth hypersurface in $\mathbb{R}^N$ is a set of points defined by the equation

$$F(x_1, x_2, \dots, x_N) = k$$

with the property that $\nabla F \neq 0$ at each point in this set.

■ **Example 1.29** Here are some examples:

1. If $F : \mathbb{R}^N \rightarrow \mathbb{R}$ is a linear map, then $F(x) = k$ gives a *hyperplane*. This is the higher dimensional analog of a line in $\mathbb{R}^2$ and a plane in $\mathbb{R}^3$.

If we write F as $F(x) = [a_1 \ a_2 \ \dots \ a_N]x$, then $\nabla F = (a_1, a_2, \dots, a_N)$, and the equation $F(x) = k$ is just

$$a_1x_1 + \dots + a_Nx_N = k.$$

2. Let $F : \mathbb{R}^N \rightarrow \mathbb{R}$ be defined by

$$F(x_1, x_2, \dots, x_N) = x_1^2 + x_2^2 + \dots + x_N^2.$$

Then the smooth hypersurface defined by $F(x_1, x_2, \dots, x_N) = 1$ is the *unit sphere* in $\mathbb{R}^N$. This is the higher dimensional analog of the unit circle in $\mathbb{R}^2$.

3. Let $F : \mathbb{R}^3 \rightarrow \mathbb{R}$ be defined by

$$F(x, y, z) = x^2 + y^2 - z.$$

Then the equation $F(x, y, z) = 0$ gives a smooth hypersurface in $\mathbb{R}^3$ that is called a *paraboloid*. See Figure 1.2. This is the analog of the parabola $\{(x, y) \in \mathbb{R}^2 \mid x^2 - y = 0\}$ in $\mathbb{R}^2$.

4. Let $F : \mathbb{R}^3 \rightarrow \mathbb{R}$ be defined by

$$F(x, y, z) = x^2 + y^2 - z^2.$$

Then the equation $F(x, y, z) = 1$ defines a *one-sheeted hyperboloid* in $\mathbb{R}^3$. See Figure 1.3. This is the analog of the hyperbola $\{(x, y) \in \mathbb{R}^2 \mid x^2 - y^2 = 1\}$ in $\mathbb{R}^2$.

5. Let $F : \mathbb{R}^3 \rightarrow \mathbb{R}$ be defined by

$$F(x, y, z) = x^2 + y^2 - z^2.$$

Then the equation $F(x, y, z) = -1$ defines a *two-sheeted hyperboloid* in $\mathbb{R}^3$. See Figure 1.4. This is the analog of the hyperbola $\{(x, y) \in \mathbb{R}^2 \mid x^2 - y^2 = -1\}$ in $\mathbb{R}^2$. ■

Definition 1.6.2 — tangent hyperplane. Let S be a smooth hypersurface in $\mathbb{R}^N$ defined by a function $F : \mathbb{R}^N \rightarrow \mathbb{R}$, and let $c = (c_1, c_2, \dots, c_N) \in S$. We define the *tangent hyperplane* to S at c to be the hyperplane through the point c normal to the vector $\nabla F(c)$; that is, the hyperplane given by the equation

$$\nabla F(c) \cdot (x - c) = 0.$$

■ **Example 1.30** Here are some examples:

1. If S is a hyperplane $a_1x_1 + \dots + a_Nx_N = k$, then for the point c to be on S , we have $a_1c_1 + \dots + a_Nc_N = k$. Then the equation for the hyperplane is

$$(a_1, \dots, a_N) \cdot (x - c) = 0,$$

so simplifying we get

$$a_1x_1 + \dots + a_Nx_N = (a_1, \dots, a_N) \cdot c = k,$$




Figure 1.2: Paraboloid




Figure 1.3: One-sheeted hyperboloid




Figure 1.4: Two-sheeted hyperboloid

which is exactly the equation for S itself.

Thus for a hyperplane S , the tangent hyperplane to S at any point $c \in S$ is itself.

2. For S the unit sphere, $\nabla F(x) = (2x_1, \dots, 2x_n)$. Thus at $c = (1, 0, \dots, 0)$, the tangent hyperplane is given by

$$(1, 0, \dots, 0) \cdot (x - c) = 0,$$

and simplifying we get the equation is

$$x_1 = 1.$$

3. For the paraboloid defined by $x^2 + y^2 - z = 0$, we have $\nabla F(x, y, z) = (2x, 2y, -1)$. Thus at $c = (0, 0, 0)$, the tangent hyperplane is given by

$$(0, 0, -1) \cdot ((x, y, z) - (0, 0, 0)) = 0,$$

and simplifying we get the equation is

$$z = 0.$$

4. For the one-sheeted hyperboloid defined by $x^2 + y^2 - z^2 = 1$, we have $\nabla F(x, y, z) = (2x, 2y, -2z)$. Thus at $c = (1, 1, 1)$, the tangent hyperplane is given by

$$(2, 2, -2) \cdot ((x, y, z) - (1, 1, 1)) = 0,$$

and simplifying we get the equation is

$$x + y - z = 1.$$

5. For the two-sheeted hyperboloid defined by $x^2 + y^2 - z^2 = -1$, we have $\nabla F(x, y, z) = (2x, 2y, -2z)$. Thus at $c = (0, 0, 1)$, the tangent hyperplane is given by

$$(0, 0, -2) \cdot ((x, y, z) - (0, 0, 1)) = 0,$$

and simplifying we get the equation is

$$z = 1.$$

■

1.6.2 Special case of a graph

In some cases, the discussion above can be made simpler. Suppose that the hypersurface defined by

$$F(x_1, x_2, \dots, x_N) = 0$$

may be realized as the graph of a C^1 function $f : \mathbb{R}^n \rightarrow \mathbb{R}$, where $N = n + 1$. Specifically, we suppose the surface can be defined by

$$x_{n+1} = f(x_1, x_2, \dots, x_n),$$

that is, as the level set

$$\{(x_1, x_2, \dots, x_n, x_{n+1}) \in \mathbb{R}^{n+1} \mid F(x_1, x_2, \dots, x_n, x_{n+1}) = 0\},$$

where

$$F(x_1, x_2, \dots, x_n, x_{n+1}) = f(x_1, x_2, \dots, x_n) - x_{n+1}.$$

A point C on this hypersurface can be written as $C = (c, f(c))$, where $c = (c_1, c_2, \dots, c_n)$ is a point in $\mathbb{R}^n$.

Using the ideas developed above, we note that the tangent hyperplane to this hypersurface at the point C is defined by

$$\nabla F(C) \cdot (X - C) = 0,$$

where $X = (x_1, \dots, x_n, x_{n+1})$. Note that $\nabla F(C) = (\nabla f(c), -1)$. Writing the equation out in coordinates, we get

$$D_1 f(c)(x_1 - c_1) + D_2 f(c)(x_2 - c_2) + \dots + D_n f(c)(x_n - c_n) - (x_{n+1} - f(c)) = 0,$$

which can be rearranged to give

$$x_{n+1} = f(c) + D_1 f(c)(x_1 - c_1) + D_2 f(c)(x_2 - c_2) + \dots + D_n f(c)(x_n - c_n).$$

■ **Example 1.31** Looking at our previous examples, we see that the paraboloid defined by

$$F(x, y, z) = x^2 + y^2 - z = 0$$

is the graph of $f(x, y) = x^2 + y^2$.

Then the tangent hyperplane at a point $C = (c_1, c_2, f(c))$ is given by

$$z = c_1^2 + c_2^2 + 2c_1(x - c_1) + 2c_2(y - c_2).$$

■

1.7 Hessian test for max-min problems

Recall that for a single variable function, the extreme points are critical points, and we can use the second-order derivative test to determine whether the critical points are maximum or minimum. We can generalize this tool to multi-variable worlds. In this section, we will look at multi-variable functions that map to $\mathbb{R}$. Before then we need to introduce some notations.

Definition 1.7.1 — global maximum/minimum. Let $B \subset \mathbb{R}^n$ be any set, and let $f : B \rightarrow \mathbb{R}$ be a function. If there exists $p \in B$ such that

$$f(x) \leq f(p) \quad \text{for all } x \in B,$$

then we say that $f(p)$ is the *global maximum* of f on B .

Similarly, if there exists $q \in B$ such that

$$f(x) \geq f(q) \quad \text{for all } x \in B,$$

then we say that $f(q)$ is the *global minimum* of f on B .

Definition 1.7.2 — local maximum/minimum. Let $B \subset \mathbb{R}^n$ be any set, and let $f : B \rightarrow \mathbb{R}$ be a function. We say that f assumes a *local maximum* at a point $p \in B$ if there exists $r > 0$ such that

$$f(x) \leq f(p) \quad \text{for all } x \in B_r(p) \cap B.$$

Similarly, we say that f assumes a *local minimum* at a point $q \in B$ if there exists $r > 0$ such that

$$f(x) \geq f(q) \quad \text{for all } x \in B_r(q) \cap B.$$

Definition 1.7.3 — extremum. If p is either a local (respectively, global) minimum or maximum we call it's a local (respectively, global) *extremum*.

Note in particular that if $p \in B$ is a point at which f attains a global maximum, then f automatically has a local maximum at p , and similarly for minima. This implies that our search for global maxima and minima will begin with the search for local maxima and minima, a search to which we may naturally apply the tools of analysis.

Note further that global maxima and global minima need not exist. If they do exist, they need not be unique. Indeed, a global maximum or minimum may even occur at an infinite number of points as in the case of a constant function. Finally, note that we do know of one very important special case in which global maxima and minima are guaranteed to exist: namely, when B is compact and $f : B \rightarrow \mathbb{R}$ is continuous.

Let us begin our search for the local maxima and minima of a function $f : B \rightarrow \mathbb{R}$. Similar to the one-dimensional case, a local maximum or minimum can be found in one of two places: in the interior of B , which is an open subset of $\mathbb{R}^n$, or on the boundary of B , which is typically a hypersurface or the intersection of several hypersurfaces.

Note the natural correspondence with the one-variable case, where maxima and minima can occur at three types of points: critical points ($f' = 0$), singular points (f' does not exist), and endpoints.

We begin the analysis by considering $f : U \rightarrow \mathbb{R}$, where U is an open subset of $\mathbb{R}^n$.

Theorem 1.7.1 Let U be an open set in $\mathbb{R}^n$, and $f : U \rightarrow \mathbb{R}$. If f has a local extremum at p , and f is differentiable at p , then $Df(p) = 0$.

Proof. We prove the case that p is a local maximum point, and the local minimum point case is similar. We only need to show that $D_{\mathbf{v}}f(p) = 0$ for all $\mathbf{v} \in \mathbb{R}^n$. Suppose p is the maximum point in an open ball $B_r(p)$.

$$D_{\mathbf{v}}f(p) = \lim_{h \rightarrow 0} \frac{f(p + h\mathbf{v}) - f(p)}{h}.$$

Because for h sufficiently small such that $p + h\mathbf{v} \in B_r(p)$, we have $f(p + h\mathbf{v}) - f(p) \leq 0$, so

$$\lim_{h \rightarrow 0^+} \frac{f(p + h\mathbf{v}) - f(p)}{h} \geq 0, \quad \lim_{h \rightarrow 0^-} \frac{f(p + h\mathbf{v}) - f(p)}{h} \leq 0.$$

Therefore, $D_{\mathbf{v}}f(p) = 0$. ■

Definition 1.7.4 — critical point. As in the one-variable case, a point where Df vanishes is called a *critical point* of f , and f need not have either a maximum or minimum at a critical point.

■ **Example 1.32** Even in the single variable case, 0 is a critical point of $f(x) = x^3$, but it is neither local minimum or maximum. ■

■ **Example 1.33** The function $f(x, y) = xy$ has vanishing derivative at the origin, but has neither maximum nor minimum there, since f is positive in the first and third quadrants and negative in the second and fourth quadrants. ■

■ **Example 1.34** Let

$$f(x, y) = \begin{cases} 2xy^2/(x^2 + y^4) & (x, y) \neq (0, 0) \\ 0 & (x, y) = (0, 0) \end{cases}$$

Show that f has all partial derivatives equal to zero at the origin, yet f is not even continuous there. ■

Once we have identified the critical points of a function f , we might then ask if there is a convenient way to determine whether f actually assumes a local maximum or minimum value at these points. Of course, we can check the behavior of the function in a neighborhood of the critical point directly using inequalities. On the other hand, similar to the one-variable case, we can consider the second-derivative test. But now, since the function has more variables, we need to put the second derivatives in a matrix.

Definition 1.7.5 — Hessian matrix. Suppose $U \subset \mathbb{R}^n$ is an open set, and $f \in C^2(U; \mathbb{R})$. For $x \in U$, the *Hessian matrix* of f at x is the $n \times n$ matrix

$$\text{Hess } f(x) := D^2 f(x) = \begin{bmatrix} D_{11}f(x) & D_{12}f(x) & \cdots & D_{1n}f(x) \\ D_{21}f(x) & D_{22}f(x) & \cdots & D_{2n}f(x) \\ \vdots & \vdots & \ddots & \vdots \\ D_{n1}f(x) & D_{n2}f(x) & \cdots & D_{nn}f(x) \end{bmatrix}$$

The Hessian is a symmetric matrix, namely the (i, j) -entry equals the (j, i) -entry.

■ **Example 1.35** Consider the function $f(x, y) = xy$. The Hessian matrix at any point (x, y) is $\begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$. ■

Recall that in the one-variable case, at a critical point x , if the second derivative is positive, then it's a local minimum; if the second derivative is negative, then it's a local maximum. The analogous result holds in the multi-variable case as well.

Lemma 1.7.2 Let $U \subset \mathbb{R}^n$ be an open set, and let $f \in C^2(U; \mathbb{R})$. If $\text{Hess } f(p)$ is positive definite at a point $p \in U$, then $\text{Hess } f(x)$ is positive definite for x in a neighborhood $B_r(p)$ of p , and similarly when $\text{Hess } f(p)$ is negative definite.

Proof. It is enough to prove the statement about positive definiteness. Let $\lambda > 0$ be the smallest eigenvalue of $\text{Hess } f(x)$, then for any vector $\mathbf{v} = (v_1, \dots, v_n)$, we have


$$\sum_{i=1}^n \sum_{j=1}^n (D_{ij}f(p))v_i v_j = \mathbf{v}^T (\text{Hess } f(p)) \mathbf{v} \geq \lambda \|\mathbf{v}\|^2.$$

Since f is C^2 , the functions $D_{ij}f$ are continuous in x . Thus for x sufficiently close to p , $D_{ij}f(x)$ is very close to $D_{ij}f(p)$, and we have

$$\mathbf{v}^T (\text{Hess } f(x)) \mathbf{v} = \sum_{i=1}^n \sum_{j=1}^n (D_{ij}f(x))v_i v_j \geq \frac{1}{2} \lambda \|\mathbf{v}\|^2.$$

This shows $\text{Hess } f(x)$ is also positive definite. ■

Theorem 1.7.3 — Second-derivative test for extrema. Let $U \subset \mathbb{R}^n$ be an open set, and let $f \in C^2(U; \mathbb{R})$. Suppose p is a critical point. Then if $\text{Hess } f(p)$ is positive definite, f has a



minimum at p ; if $\text{Hess } f(p)$ is negative definite, f has a maximum at p .

Proof. It is enough to prove the statement about minima, since we can replace f by $-f$ for maxima.

Choose $r > 0$ such that $\text{Hess } f(x)$ is positive definite for $x \in B_r(p)$.

Given any unit vector h , any $0 < t < r$, by the previous lemma, the Taylor formula for f at p , to second order, is

$$f(p + th) = f(p) + \sum_{i=1}^n (D_i f(p)) h_i t + \frac{1}{2} \sum_{i=1}^n \sum_{j=1}^n D_{ij} f(p + sh) h_i h_j t^2$$

for some $s \in [0, t]$. Since p is a critical point, $Df(p) = 0$. Since $\text{Hess } f(p + sh)$ is positive definite, we have $\sum_{i=1}^n \sum_{j=1}^n D_{ij} f(p + sh) h_i h_j > 0$.

Thus $f(p + th) > f(p)$, which proves that f assumes a local minimum at p . ■

Exercise 1.10 Let $f : U \subset \mathbb{R}^n \rightarrow \mathbb{R}$ be a C^2 function. Let $p \in U$ be a critical point of f , and suppose that $\text{Hess } f(p)$ has at least one positive eigenvalue $\lambda > 0$ and at least one negative eigenvalue $\lambda' < 0$. Show that f does not have a local extremum at p . ■

For $f : U \subset \mathbb{R}^n \rightarrow \mathbb{R}$, the critical points of the type in the above exercise are referred to as *saddle points*.

Lastly, we just want to discuss about the notation $\text{Hess } f(x) = D^2 f(x)$ for $f : \mathbb{R}^n \rightarrow \mathbb{R}$ and give you a chain rule for Hessian.

It is tempting to think $D^2 f(x)$ is just the derivative of $Df(x)$. However, to do this, Df needs to be a function from $\mathbb{R}^n$ to $\mathbb{R}^n$! Is this the case? Well, Df is a function that maps x to an $1 \times n$ matrix, which is a row vector. In other words, Df maps from $\mathbb{R}^n$ to the space of $1 \times n$ matrices. So it isn't quite right without further identifications.

Instead, if we take the gradient vector, $\nabla f(x) = (Df(x))^T$, then because it is a column vector, we see that ∇f maps $x \in \mathbb{R}^n$ to some vector in $\mathbb{R}^n$. Thus ∇f is a function from $\mathbb{R}^n$ to $\mathbb{R}^n$. Then we can safely take the derivative of it; we see that because $\nabla f : \mathbb{R}^n \rightarrow \mathbb{R}^n$, we have that $D(\nabla f)(x)$ will be an $n \times n$ matrix. What is it? If we compute, we see that the (i, j) -th entry is given by $D_j((\nabla f)_i)(x) = D_j D_i f(x)$. Thus we see that exactly we have

$$D(\nabla f)(x) = (\text{Hess } f(x))^T.$$

But if f is C^2 , then Hessian is symmetric, so we have

Proposition 1.7.4 — Hessian in terms of the gradient. Let $U \subset \mathbb{R}^n$ be an open set, and let $f \in C^2(U; \mathbb{R})$. Let $x \in U$. Then

$$\text{Hess } f(x) = D(\nabla f)(x).$$

Why do we care about this? It's because this will make chain rule to work in the correct way. Recall that the chain rule for derivative and the gradient works differently due to the transpose.

■ **Example 1.36 — chain rule for Hessian.** Let $A : \mathbb{R}^n \rightarrow \mathbb{R}^m$ be a linear map. Let $f : \mathbb{R}^m \rightarrow \mathbb{R}$, and let $g : \mathbb{R}^n \rightarrow \mathbb{R}$ be the composition $g = f \circ A$. Then we want to write the Hessian of g in terms of the Hessian of f . For this, we first use chain rule for the gradient to get

$$\nabla g(x) = A^T(\nabla f(Ax)).$$

Then we take the derivative of ∇g and use chain rule for the derivative to get

$$\text{Hess } g(x) = D(\nabla g)(x) = A^T D(\nabla f(Ax)) = A^T (D(\nabla f)(Ax))A = A^T \text{Hess } f(Ax)A.$$

If you are not careful, you might get expressions like $\text{Hess } f(Ax)AA$ which doesn't make sense (the dimensions of the matrices don't match). ■

1.8 Lagrange multipliers

So far we have learned how to find local extremum in an open set. We now turn to the case where we wish to find the extreme values of a function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ restricted to a hypersurface in $\mathbb{R}^n$ or the intersection of several hypersurfaces. We will first deal with the case of a single hypersurface S . The equation for the hypersurface S is sometimes referred to as a *constraint*, and the overall technique we develop here is called optimization with constraints.

1.8.1 Special case: when we have a parametrization

In the case that S can be realized as the graph of a function $g : \mathbb{R}^{n-1} \rightarrow \mathbb{R}$, we are reduced to the problem of finding extreme values of the function $h : \mathbb{R}^{n-1} \rightarrow \mathbb{R}$ defined by

$$h(x_1, \dots, x_{n-1}) = f(x_1, \dots, x_{n-1}, g(x_1, \dots, x_{n-1}))$$

on the open set $\mathbb{R}^{n-1}$, which we have done in the previous section using Hessian.

Even if the equation defining S cannot be solved explicitly for one of the variables, it may still be possible to describe S by an unconstrained set of $n - 1$ variables by parametrization. While we will not pursue a full discussion of the notion of parametrization here, we will give a typical example of this approach.

■ **Example 1.37** Let $S = \{(x, y, z) \mid x^2 + y^2 + z^2 = 1\}$ be the unit sphere in $\mathbb{R}^3$. Suppose we wish to find the maxima and minima of the function $f(x, y, z) = x + y + z$ on S . We may parametrize S using spherical coordinates as follows. Recall that, in general, spherical coordinates are given by

$$x = \rho \cos \theta \sin \phi, \quad y = \rho \sin \theta \sin \phi, \quad z = \rho \cos \phi,$$

where $\rho \in (0, \infty)$ is the distance to the origin, $\theta \in [0, 2\pi)$ is the azimuthal angle and $\phi \in [0, \pi]$ is the polar angle.

In a homework exercise, we computed the Jacobi matrix of the change of variable map $(\rho, \theta, \phi) \mapsto (x, y, z)$ is given by

$$J = \begin{bmatrix} \sin \phi \cos \theta & -\rho \sin \phi \sin \theta & \rho \cos \phi \cos \theta \\ \sin \phi \sin \theta & \rho \sin \phi \cos \theta & \rho \cos \phi \sin \theta \\ \cos \phi & 0 & -\rho \sin \phi \end{bmatrix}.$$

In the specific case at hand, however, we have $\rho = 1$, so S is parametrized by the two unconstrained variables θ and ϕ .

Written in spherical coordinates, the function f becomes

$$f_0(\theta, \phi) = \cos \theta \sin \phi + \sin \theta \sin \phi + \cos \phi.$$

This parameterization causes problems at the two poles $(x, y, z) = (0, 0, \pm 1)$ where $\sin \phi = 0$. At these two points, the Jacobi matrix of the change of variables map above is not invertible, and we cannot determine whether $(0, 0, \pm 1)$ are critical with this parameterization.

We have

$$\frac{\partial f_0}{\partial \theta} = -\sin \theta \sin \phi + \cos \theta \sin \phi = \sin \phi (\cos \theta - \sin \theta),$$

and

$$\frac{\partial f_0}{\partial \phi} = \cos \theta \cos \phi + \sin \theta \cos \phi - \sin \phi = \cos \phi (\cos \theta + \sin \theta) - \sin \phi.$$

Note that $\frac{\partial f_0}{\partial \theta} = 0$ when $\cos \theta = \sin \theta$ or $\sin \theta = 0$. The former occurs when $\cos \theta = \sin \theta = \pm \frac{\sqrt{2}}{2}$. In this case, $\cos \theta + \sin \theta = \pm \sqrt{2}$, so $\frac{\partial f_0}{\partial \phi} = 0$ when $\pm \sqrt{2} \cos \theta - \sin \theta = 0$, that is, when $\tan \theta = \pm \sqrt{2}$.

We have thus shown that points with $\cos \theta = \sin \theta = \pm \frac{\sqrt{2}}{2}$ and $\tan \theta = \pm \sqrt{2}$ (where the signs are the same) are among the critical points of f .

■

1.8.2 Lagrange multipliers

In the above example, we can see that even if we have a parametrization, it might not be so nice to do the computation. And in many cases we don't have a nice parametrization of the hypersurface. Thus we want to approach this extreme value problem via another approach, which is the method of *Lagrange multipliers*.

The idea is as follows. Let the hypersurface S be the level set of a C^1 function $g : \mathbb{R}^n \rightarrow \mathbb{R}$. We have already observed that $\nabla g(x)$ is normal to the tangent hyperplane at any point $x \in S$. Now suppose that a function $f : \mathbb{R}^n \rightarrow \mathbb{R}$, when restricted to the hypersurface S , has an extreme value at the point $a \in S$. Then, for any C^1 curve $\gamma : \mathbb{R} \rightarrow S$ with $\gamma(0) = a$, the function $f \circ \gamma : \mathbb{R} \rightarrow \mathbb{R}$ has an extreme value at 0, and hence

$$0 = (f \circ \gamma)'(0) = \nabla f(\gamma(0)) \cdot \gamma'(0) = \nabla f(a) \cdot \gamma'(0).$$

In other words, $\nabla f(a)$ is normal to the tangent vector of the curve γ at the point a . Since this is true for any such curve γ , $\nabla f(a)$ is normal to the tangent hyperplane of S at a . Since $\nabla g(a)$ is also normal to the tangent hyperplane of S at a , this implies that $\nabla f(a)$ is proportional to $\nabla g(a)$, so we can write

$$\nabla f(a) = \lambda \nabla g(a)$$

for some $\lambda \in \mathbb{R}$. The scalar λ is known as a *Lagrange multiplier*.

Now let's have a formal statement of the theorem.

Theorem 1.8.1 — Lagrange multiplier. Let U be an open set in $\mathbb{R}^n$ and let $f : U \rightarrow \mathbb{R}$ be C^1 . Let $g : U \rightarrow \mathbb{R}$ also be C^1 , and let S be the hypersurface defined by $g(x) = 0$.

If $f|_S$ has a local maximum or minimum at a point $a \in S$, and $\nabla g(a) \neq 0$, then there exists $\lambda \in \mathbb{R}$ such that

$$\nabla f(a) = \lambda \nabla g(a).$$

For the proof, we need to make the intuitive argument we give rigorous, and this requires results in the next section.

Now let us study the example above using this new idea.

■ **Example 1.38** Let $S = \{(x, y, z) \mid x^2 + y^2 + z^2 = 1\}$ be the unit sphere in $\mathbb{R}^3$. Suppose we wish to find the maxima and minima of the function $f(x, y, z) = x + y + z$ on S . We observe first that $\nabla f(x, y, z) = (1, 1, 1)$.

The surface S can be realized as the level set $g(x, y, z) = 0$ for the function $g(x, y, z) = x^2 + y^2 + z^2 - 1$, which has gradient $\nabla g(x, y, z) = (2x, 2y, 2z)$. To identify the potential points $a = (x_0, y_0, z_0)$ on the surface S where f attains an extreme value, we set up the Lagrange multiplier equation $\nabla f(a) = \lambda \nabla g(a)$, which becomes

$$(1, 1, 1) = \lambda(2x_0, 2y_0, 2z_0).$$

This gives us three equations in our four unknowns, namely,

$$2\lambda x_0 = 1, \quad 2\lambda y_0 = 1, \quad 2\lambda z_0 = 1.$$

But we must not forget that we have the original constraint, namely, that $g(a) = 0$. Hence, there is a fourth equation,

$$x_0^2 + y_0^2 + z_0^2 = 1.$$

Solving this system of four equations in four unknowns gives $x_0 = y_0 = z_0 = \pm \frac{1}{\sqrt{3}}$.

The point $a_1 = \left(\frac{1}{\sqrt{3}}, \frac{1}{\sqrt{3}}, \frac{1}{\sqrt{3}}\right)$ is the point where f attains its maximum value of $\sqrt{3}$ when restricted to S , and the point $a_2 = \left(-\frac{1}{\sqrt{3}}, -\frac{1}{\sqrt{3}}, -\frac{1}{\sqrt{3}}\right)$ is the point where f attains its minimum value of $-\sqrt{3}$ when restricted to S .

Finally, it is worth noting that in this case, the precise value of λ is irrelevant and was merely used in an auxiliary fashion to help find the potential extrema. ■

R Suppose now the question becomes that we wish to find the global maxima and minima of the function $f(x, y, z) = x + y + z$ on the closed ball $\{(x, y, z) \mid x^2 + y^2 + z^2 \leq 1\}$.

Then we need to proceed in three steps. First, we compute the derivative of f to find the critical points of f in the open ball $\{(x, y, z) \mid x^2 + y^2 + z^2 < 1\}$. Second, we use the Lagrange multiplier method as above to find the extreme values of f on $S = \{(x, y, z) \mid x^2 + y^2 + z^2 = 1\}$. Lastly, we compare these values to find the global maxima and minima.

■ **Example 1.39** Let us use the method of Lagrange multipliers to determine the points on the ellipse

$$x^2 + 4y^2 = 4$$

that are closest to and farthest from the point $(1, 0)$.

The function that we want to optimize is the distance function. However, this function has the square root which makes computation annoying. Instead, we look at the square of the distance from the point (x, y) to $(1, 0)$ given by

$$f(x, y) = (x - 1)^2 + y^2$$

and we wish to optimize this function subject to the constraint

$$g(x, y) = x^2 + 4y^2 - 4 = 0.$$

We have

$$\nabla f(x, y) = (2(x - 1), 2y), \quad \nabla g(x, y) = (2x, 8y),$$

so we consider the equations

$$2(x - 1) = \lambda(2x), \quad 2y = \lambda(8y), \quad x^2 + 4y^2 = 4.$$

The second equation tells either $y = 0$ or $\lambda = \frac{1}{4}$. In the former case, we can find that $(x, y) = (\pm 2, 0)$. In the latter case, we can find that $(x, y) = (\frac{4}{3}, \pm \frac{\sqrt{5}}{3})$.

We evaluate $f(x, y) = (x - 1)^2 + y^2$ at each point:

$$\begin{aligned} f(2, 0) &= (2 - 1)^2 + 0 = 1, \\ f(-2, 0) &= (-2 - 1)^2 + 0 = 9, \\ f\left(\frac{4}{3}, \pm \frac{\sqrt{5}}{3}\right) &= \left(\frac{1}{3}\right)^2 + \left(\frac{\sqrt{5}}{3}\right)^2 = \frac{1 + 5}{9} = \frac{2}{3}. \end{aligned}$$

Then we can see that $(\frac{4}{3}, \pm \frac{\sqrt{5}}{3})$ are the nearest point and $(-2, 0)$ is the farthest point. ■

So far we learned how to find the extreme values for a function restricted to a hypersurface in $\mathbb{R}^n$, which is an $(n - 1)$ -dimensional object. What if we want to solve this problem on some lower dimensional objects in $\mathbb{R}^n$? Recall that the intersection of two curves in $\mathbb{R}^2$ are points, the intersection of two surfaces in $\mathbb{R}^3$ will generally be curves, and the intersection of three surfaces in $\mathbb{R}^3$ will generally be points. In general, we expect the intersection of m hypersurfaces in $\mathbb{R}^n$ to be some $(n - m)$ -dimensional object.

In this case, we also have a max-min theorem. Here we only give the statement without the proof.

Theorem 1.8.2 — Lagrange multipliers with several constraints. Let U be an open set in $\mathbb{R}^n$ and let $f : U \rightarrow \mathbb{R}$ be C^1 . Let $g_1, \dots, g_m : \mathbb{R}^n \rightarrow \mathbb{R}$ be m C^1 functions, with $m < n$.

For each $1 \leq i \leq m$, let S_i be the hypersurface $S_i = \{x \in \mathbb{R}^n \mid g_i(x) = 0\}$. Let $S = \bigcap_{i=1}^m S_i$ be the intersection.

If $f|_S$ has a local maximum or minimum at $a \in S$, and the vectors $\nabla g_1(a), \dots, \nabla g_m(a)$ form a linearly independent set, then there exist m numbers $\lambda_1, \dots, \lambda_m$ satisfying the equation

$$\nabla f(a) = \lambda_1 \nabla g_1(a) + \dots + \lambda_m \nabla g_m(a).$$

Here, the scalars $\lambda_1, \dots, \lambda_m$ are called *Lagrange multipliers*.

R The condition that the vectors $\nabla g_1(a), \dots, \nabla g_m(a)$ form a linearly independent set is really necessary in the above theorem, as seen by the example below. In the following example, the intersection S is still some $n - m$ dimensional object, but $\nabla g_1(a), \dots, \nabla g_m(a)$ fail to be linearly independent.

■ **Example 1.40** Consider the following example in $\mathbb{R}^2$: take $f(x, y) = y$, and take $g_1(x, y) = (x - 1)^2 + y^2 - 1$, $g_2(x, y) = (x + 1)^2 + y^2 - 1$.

Then $S_1 = \{g_1(x, y) = 0\}$, $S_2 = \{g_2(x, y) = 0\}$ are two circles, and $S = S_1 \cap S_2 = \{(0, 0)\}$ is a single point.

Then if we want to minimize f on S , of course we get that f attains the minimum on the single point $(0, 0)$. However, we see that

$$\nabla f(0, 0) = (0, 1)$$

$$\nabla g_1(0, 0) = (-2, 0)$$

$$\nabla g_2(0, 0) = (2, 0),$$

and $\nabla f(0, 0)$ is not a linear combination of $\nabla g_1(0, 0)$ and $\nabla g_2(0, 0)$. ■

■ **Example 1.41** Now we work out an example where we use Lagrange multipliers with several constraints.

Given the line defined by

$$L(x, y) = y - (x + 4) = 0,$$

and the circle defined by

$$C(x, y) = x^2 + y^2 - 1 = 0,$$

we wish to find the minimum distance between a point (x_1, y_1) on the line and a point (x_2, y_2) on the circle. Drawing a picture shows this line and this circle doesn't overlap, so the distance will be positive.

Similar as before, because the distance function has a square root, we choose to minimize the squared distance

$$d(x_1, y_1, x_2, y_2) = (x_1 - x_2)^2 + (y_1 - y_2)^2$$

subject to the constraints

$$g_1(x_1, y_1, x_2, y_2) = y_1 - x_1 - 4 = 0, \quad g_2(x_1, y_1, x_2, y_2) = x_2^2 + y_2^2 - 1 = 0.$$

We compute the gradients:

$$\nabla d(x_1, y_1, x_2, y_2) = (2(x_1 - x_2), 2(y_1 - y_2), -2(x_1 - x_2), -2(y_1 - y_2)),$$

$$\nabla g_1 = (-1, 1, 0, 0), \quad \nabla g_2 = (0, 0, 2x_2, 2y_2).$$

Using Lagrange multipliers, at an extreme point, we have

$$\nabla d = \lambda_1 \nabla g_1 + \lambda_2 \nabla g_2,$$

so the Lagrange multipliers equation together with the two constraints gives us six equations with six unknowns:

$$2(x_1 - x_2) = \lambda_1(-1) \tag{1.1}$$

$$2(y_1 - y_2) = \lambda_1 \tag{1.2}$$

$$-2(x_1 - x_2) = \lambda_2 \cdot 2x_2 \tag{1.3}$$

$$-2(y_1 - y_2) = \lambda_2 \cdot 2y_2 \tag{1.4}$$

$$y_1 - x_1 - 4 = 0 \tag{1.5}$$

$$x_2^2 + y_2^2 - 1 = 0. \tag{1.6}$$

To solve, from (1) and (2) we see that $x_1 - x_2 = -(y_1 - y_2)$, or equivalently, $x_1 + y_1 = x_2 + y_2$. Using this, we add (3) and (4), and find $\lambda_2(x_2 + y_2) = 0$. Thus either $\lambda_2 = 0$, or $x_1 + y_1 = x_2 + y_2 = 0$.

If $\lambda_2 = 0$, then plugging in (3) and (4) we get that $x_1 = x_2$, $y_1 = y_2$. This implies the point $(x_1, y_1) = (x_2, y_2)$ lies on both the line and the circle, but this is impossible by looking the picture.

Indeed, we can also check this by plugging $x_1 = x_2$, $y_1 = y_2$ and (5) into (6). Then we get $x_2 = x_1 = y_1 - 4 = y_2 - 4$, so (6) becomes $(y_2 - 4)^2 + y_2^2 = 1$. Expanding we get $2y_2^2 - 8y_2 + 15 = 0$, which have no solution because the discriminant $\Delta = b^2 - 4ac = 8^2 - 4 \cdot 2 \cdot 15 = 64 - 120 < 0$.

So we know $\lambda_2 \neq 0$. Then we must have $x_1 + y_1 = x_2 + y_2 = 0$. Plugging $x_1 = -y_1$ into (5) gives $2y_1 - 4 = 0$, so $y_1 = 2$, and $x_1 = -y_1 = -2$. Thus $(x_1, y_1) = (-2, 2)$.

Plugging $y_2 = -x_2$ into (6) gives $2x_2^2 = 1$, so $x_2 = \pm \frac{\sqrt{2}}{2}$. Thus $(x_2, y_2) = (\frac{\sqrt{2}}{2}, -\frac{\sqrt{2}}{2})$ or $(-\frac{\sqrt{2}}{2}, \frac{\sqrt{2}}{2})$.

We can then compute the distances from $(-2, 2)$ to $(\frac{\sqrt{2}}{2}, -\frac{\sqrt{2}}{2})$ and $(-\frac{\sqrt{2}}{2}, \frac{\sqrt{2}}{2})$ or simply look at the picture. We find that $(x_1, y_1) = (-2, 2)$ and $(x_2, y_2) = (-\frac{\sqrt{2}}{2}, \frac{\sqrt{2}}{2})$ gives the closest distance between the line and the circle.

(There is just one small issue: we didn't justify that a global minimum exists. Notice that because the line is not compact, we don't get the existence of a global minimizer automatically. One thing we can do is to say if we go very far along the line, the distance will get longer, so we just need to consider a compact middle portion of my line, where the compactness guarantees the existence of the minimizer.)

■

1.9 Implicit and inverse function theorems

1.9.1 Inverse function theorem

Definition 1.9.1 — inverse function. Let X, Y be sets. Let $f : X \rightarrow Y$ be a function. Then f is *invertible* if there exists a function g from Y to X such that $g(f(x)) = x$ for all $x \in X$ and $f(g(y)) = y$ for all $y \in Y$.

If f is invertible, then there is exactly one function g satisfying this property. The function g is called the *inverse* of f , and is usually denoted as f^{-1} .

R The function f is invertible if and only if it is bijective. This is because the condition $g(f(x)) = x$ for all $x \in X$ implies that f is injective, and the condition $f(g(y)) = y$ for all $y \in Y$ implies that f is surjective.

■ **Example 1.42** The linear function $f : \mathbb{R} \rightarrow \mathbb{R}$ defined by $f(x) = x$ is invertible. The inverse function $f^{-1} = f$ is itself. (Don't be confused by the notation and think f^{-1} is $1/x$.) ■

■ **Example 1.43** The cubic function $f : \mathbb{R} \rightarrow \mathbb{R}$ defined by $f(x) = x^3$ is bijective hence invertible. The inverse function is $f^{-1}(y) = \sqrt[3]{y}$. ■

■ **Example 1.44** The exponential function $\exp : \mathbb{R} \rightarrow (0, \infty)$ and the logarithmic function $\log : (0, \infty) \rightarrow \mathbb{R}$ are inverses of each other. We have $e^{\log x} = x$ and $\log e^x = x$. ■

■ **Example 1.45** Let $f : [0, \infty) \rightarrow [0, \infty)$ be defined by $f(x) = x^2$. Then f is bijective, and hence has an inverse $g(y) = \sqrt{y}$ for all $y \in [0, \infty)$.

But if we look at $f(x) = x^2$ defined on the whole real line, then it is not bijective and hence has no inverse. ■

R For the 1-d case, it is very easy to see the inverse function via graphs. The graph of the inverse function f^{-1} is the graph of f reflected across the line $y = x$.

■ **Example 1.46** Let A be an invertible (also called nonsingular) $n \times n$ matrix. Then viewing A as a linear map from $\mathbb{R}^n$ to $\mathbb{R}^n$ given by $x \mapsto Ax$, it is an invertible function with inverse being the matrix inverse A^{-1} .




Figure 1.5: graph of the function and its inverse

Recall that we learned about many characterizations of invertible matrices. One characterization is that an $n \times n$ matrix A is invertible if and only if the determinant of A is nonzero. For a 2×2 matrix, this is probably the easiest condition to check. For $A = \begin{bmatrix} a & b \\ c & d \end{bmatrix}$, the determinant is $\det A = ad - bc$.

■

Let's first look at the single-variable case. Let $U \subset \mathbb{R}$ be an open set and $f : U \rightarrow \mathbb{R}$ is C^1 . In general f is not bijective onto its image. Nevertheless, if the derivative is nonzero, then f is locally invertible.

Theorem 1.9.1 — inverse function theorem for the one-variable case. Let $U \subset \mathbb{R}$ be an open set and $f : U \rightarrow \mathbb{R}$ is C^1 . Take a point $x_0 \in U$. Suppose $f'(x_0) \neq 0$. Then there exists an open interval $I \subset \mathbb{R}$ around x_0 and an open interval $J \subset \mathbb{R}$ around $f(x_0)$, such that f maps I bijectively onto J , and the inverse function $f^{-1} : J \rightarrow I$ is also C^1 , with

$$(f^{-1})'(y) = \frac{1}{f'(x)}$$

for $y = f(x) \in J$.

Proof. Since f is C^1 , the derivative f' is continuous. Since $f'(x_0) \neq 0$, by continuity there exists an open interval $I \ni x_0$ such that $f' \neq 0$ and $\geq \frac{1}{2}|f'(x_0)|$ on I .

We saw earlier in the chapter (Corollary 1.1.7 and Corollary 1.1.8) that $f' \neq 0$ on I implies f is monotonic on I . This, of course, implies that f is one-to-one on I . Moreover, $f(I)$ is an open interval J contained in $\mathbb{R}$. Thus $f : I \rightarrow J$ is bijective, and we thus have the inverse $f^{-1} : J \rightarrow I$.

For any $y = f(x) \in J$, and $y' = f(x') \in J$, by the mean value theorem, $y' - y = f(x') - f(x) = f'(s)(x' - x)$ for some $s \in J$, so $|y' - y| = |f'(s)||x' - x| \geq \frac{1}{2}|f'(x_0)||x - x'|$.

Thus as y' approaches y , x' approaches x , so

$$(f^{-1})'(y) = \lim_{y' \rightarrow y} \frac{f^{-1}(y') - f^{-1}(y)}{y' - y} = \lim_{x' \rightarrow x} \frac{x' - x}{f(x') - f(x)} = \frac{1}{f'(x)}.$$

This shows f^{-1} is differentiable, and $(f^{-1})'$ is continuous because f' is. ■

■ **Example 1.47** Even if $f' = 0$ somewhere, it is still possible for f to have an inverse function. For example, we can take $f(x) = x^3$. It is just in this case the inverse will not be C^1 everywhere; we see that $(f^{-1})'(0)$ doesn't exist. ■

- R** If we look at the proof, we see that if $f'(x) \neq 0$ on the whole U , then f will be monotone on the whole U so an inverse exists on the whole U and we get a global inverse. This will not happen in the multi-variable case.

Exercise 1.11 Show the above theorem fails if f is not C^1 . That is, find an example of a function from $\mathbb{R}$ to $\mathbb{R}$ where the derivative is nonzero at a point but the function is not invertible in any neighborhood of that point. ■

Now we want to have a similar result for multi-variable functions f . First we see that the dimension of the domain and the target should match; thus we should be looking at functions from $\mathbb{R}^n$ to $\mathbb{R}^n$. Next, we want to generalize the 1d condition $f'(x_0) \neq 0$ to higher dimensions. Intuitively, if a function is C^1 , then it locally “behaves like” the derivative (which is a linear map). Since we want the function to be locally invertible, the natural condition would be to require $Df(x_0)$ to be invertible.

Theorem 1.9.2 — inverse function theorem. Let U be an open set in $\mathbb{R}^n$ and let $f : U \rightarrow \mathbb{R}^n$ be C^1 .

Let $x_0 \in U$ be such that $Df(x_0)$ is invertible. Then there exists a neighborhood V of x_0 such that:

1. $f : V \rightarrow f(V)$ is a bijection;
2. $f(V)$ is an open set in $\mathbb{R}^n$;
3. $f^{-1} : f(V) \rightarrow V$ is C^1 and

$$Df^{-1}(f(x)) = (Df(x))^{-1} \quad \text{for } x \in V.$$

- R** For item 1, by definition of the image, $f : V \rightarrow f(V)$ is automatically onto. Thus to show f is a bijection, we only need to show f is one-to-one.
- R** For item 2, we want $f(V)$ to be open because we want the domain of f^{-1} to be an open set. We see that this conclusion is sort of obvious in the 1d case. But in higher dimensions, because open sets can be more complicated, this requires a proof.
- R** The above theorem only guarantees the existence of a local inverse, but not the global inverse! This is different from the one-variable case. Even if $Df(x)$ is invertible for all $x \in U$, it doesn't mean we can take V to be the entire U . A counterexample is given by the function $f : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ defined by $f(x, y) = (e^x \cos y, e^x \sin y)$.

■ **Example 1.48** By item 3 of the theorem, $Df^{-1}(f(x)) = (Df(x))^{-1}$ is the matrix inverse of $Df(x)$. In general it takes some time to compute the matrix inverse. But when $n = 2$, there is a nice formula for the matrix inverse. Suppose $A = \begin{bmatrix} a & b \\ c & d \end{bmatrix}$. Then $A^{-1} = \frac{1}{ad-bc} \begin{bmatrix} d & -b \\ -c & a \end{bmatrix}$. ■

To prove the theorem, we try to generalize the proof of the single-variable case. We thus need the following lemma.

Lemma 1.9.3 Let $U \subset \mathbb{R}^n$ be open and let $f : U \rightarrow \mathbb{R}^n$ be of class C^1 . Take $x_0 \in U$ and suppose that $Df(x_0)$ is invertible. Then there exists a neighborhood W of x_0 and a constant $c > 0$ such that

$$\|f(y) - f(x)\| \geq c\|y - x\| \quad \text{for all } x, y \in W.$$

Proof. We first see that this is true for linear maps. For any linear map $T : \mathbb{R}^n \rightarrow \mathbb{R}^n$, using an inequality of the operator norm, we know that

$$\|T(y) - T(x)\| = \|T(y - x)\| \leq \|T\| \cdot \|y - x\|.$$

If T is invertible, then we can apply this result to T^{-1} , and get

$$\|T(y) - T(x)\| \geq \|T^{-1}\|^{-1} \cdot \|y - x\|.$$

(Exercise: Why is this?)

Since $Df(x_0)$ is an invertible matrix, we can take $c = \|Df(x_0)^{-1}\|^{-1}/2$. Suppose that $f_1, f_2, \dots, f_n$ are the components of f . Of course, these are C^1 , so there exists a convex neighborhood W of x_0 such that

$$\|Df_i(y) - Df_i(x_0)\| \leq \frac{c}{\sqrt{n}} \quad \text{for all } y \in W \text{ and for all } i.$$

Now, the Mean Value Theorem implies that, given $x, y \in W$, there exists a point ξ_i on the line segment joining x and y such that

$$f_i(y) - f_i(x) = Df_i(\xi_i)(y - x).$$

(Think about why we need to do this for each f_i instead of for the function f .)

Consequently, for $x, y \in W$ and each i , we have

$$|f_i(y) - f_i(x) - Df_i(x_0)(y - x)| = |(Df_i(\xi_i) - Df_i(x_0))(y - x)| \leq \|Df_i(\xi_i) - Df_i(x_0)\| \cdot \|y - x\| \leq \frac{c}{\sqrt{n}} \|y - x\|.$$

It follows immediately that

$$\|f(y) - f(x) - Df(x_0)(y - x)\| = \sqrt{\sum_{i=1}^n |f_i(y) - f_i(x) - Df_i(x_0)(y - x)|^2} \leq c \|y - x\|.$$

Now, using the triangle inequality, we get

$$\|f(y) - f(x)\| \geq \|Df(x_0)(y - x)\| - \|f(y) - f(x) - Df(x_0)(y - x)\| \geq (\|Df(x_0)^{-1}\|^{-1} - c) \|y - x\|.$$

Since $c = \frac{1}{2} \|Df(x_0)^{-1}\|^{-1}$, we conclude that

$$\|f(y) - f(x)\| \geq \frac{1}{2} \|Df(x_0)^{-1}\|^{-1} \|y - x\|$$

for all $x, y \in W$. ■

With this lemma, we can prove item 1 in the inverse function theorem. Recall that we only need to show f is one-to-one on some neighborhood.

Corollary 1.9.4 Let $U \subset \mathbb{R}^n$ be open and let $f : U \rightarrow \mathbb{R}^n$ be of class C^1 . Take $x_0 \in U$ and suppose that $Df(x_0)$ is invertible. Then there exists a neighborhood W of x_0 such that $f|_W$ is one-to-one.

Proof. Exercise. We can take the same W as in the previous lemma. ■

We can also prove item 2 in the inverse function theorem.

Corollary 1.9.5 Let $U \subset \mathbb{R}^n$ be open and let $f : U \rightarrow \mathbb{R}^n$ be of class C^1 . Take $x_0 \in U$ and suppose that $Df(x_0)$ is invertible. Then there exists a neighborhood V of x_0 such that $f(V)$ is open.

Proof. Using the previous lemma, we can pick a neighborhood W of x_0 such that $W \subset U$ and for some constant $c > 0$,

$$\|f(y) - f(x)\| \geq c\|y - x\| \quad \text{for all } x, y \in W,$$

and furthermore, $Df(x)$ is invertible for all $x \in W$.

Let V be any open ball such that the closed ball $\bar{V}$ is contained in W , and let $S = \partial V$. Given a point $x \in V$ with $y = f(x)$, we notice that $y \notin f(S)$ because f is one-to-one on W by the previous corollary. Since $f(S)$ is compact, the distance from y to $f(S)$, which we denote by d , is greater than zero.

To show that $f(V)$ is open, we establish that $f(V)$ contains the ball $B_{d/2}(y)$. To see this, take $z \in B_{d/2}(y)$. Our goal is to show $z \in f(V)$, i.e., there exists some x_1 such that $z = f(x_1)$.

By definition we have $\|z - y\| < d/2$. Moreover, the distance from z to $f(\bar{V})$, which equals

$$\inf_{x' \in \bar{V}} \|z - f(x')\|,$$

is less than $d/2$ because $\inf_{x' \in \bar{V}} \|z - f(x')\| \leq \|z - f(x)\| = \|z - y\| < d/2$. Since the distance from y to $f(S)$ is d , it follows from the triangle inequality that the distance from z to $f(S)$ is greater than $d/2$.

For $x' \in \bar{V}$, define the function

$$g(x') = \|z - f(x')\|^2 = \sum_{i=1}^n (z_i - f_i(x'))^2.$$

We want to minimize this function. Because $\bar{V}$ is compact, we know that the minimum exists, so there exists $x_1 \in \bar{V}$ such that

$$\|z - f(x_1)\|^2 = d(z, f(V))^2 < (d/2)^2.$$

From the previous inequalities, it follows that $x_1 \notin S$, so that $x_1 \in V$. Thus, the minimum of g occurs at x_1 .

This implies that x_1 is a critical point of g , so for each $1 \leq j \leq n$,

$$0 = D_j g(x_1) = -2 \sum_{i=1}^n (z_i - f_i(x_1)) D_j f_i(x_1).$$

Writing this in a matrix form, we get

$$0 = Df(x_1)(z - f(x_1)),$$

and since $Df(x_1)$ is invertible, we conclude that $z = f(x_1)$.

Since $z \in B_{d/2}(y)$ was arbitrary, it follows that $B_{d/2}(y) \subset f(V)$, and hence $f(V)$ is open. ■

Lastly, we show item 3 in the inverse function theorem.

Proof of Theorem 1.9.2, item 3. So far, we have proved that $f : V \rightarrow f(V)$ is one-to-one and that $f(V)$ is open. We consider the map

$$f^{-1} : f(V) \rightarrow V.$$

We want to prove that it is C^1 .

We first notice that f^{-1} is continuous: Take $y, z \in f(V)$, then $f^{-1}(y), f^{-1}(z) \in V$. By Lemma 1.9.3, we see that $\|y - z\| = \|f(f^{-1}(y)) - f(f^{-1}(z))\| \geq c\|f^{-1}(y) - f^{-1}(z)\|$, so as $z \rightarrow y$, we also have $f^{-1}(z) \rightarrow f^{-1}(y)$ showing f^{-1} is continuous.

Then, to show f^{-1} is C^1 , we just need to show $Df^{-1}(f(x))$ exists and is continuous.

A formal computation using the chain rule applied to the function $f^{-1} \circ f$ shows that $Df^{-1}(f(x))$, if it exists, is equal to $(Df(x))^{-1}$.

To complete the proof, we take $y = f(x) \in f(V)$ and consider the difference quotient

$$Q := \frac{\|f^{-1}(z) - f^{-1}(y) - Df(x)^{-1}(z - y)\|}{\|z - y\|}.$$

To show f^{-1} is differentiable, we want to show this goes to zero as $z \rightarrow y$. We can write this as

$$\frac{\|Df(x)^{-1} [Df(x)(f^{-1}(z) - f^{-1}(y)) - (z - y)]\|}{\|z - y\|},$$

which is less than or equal to

$$\|Df(x)^{-1}\| \cdot \frac{\|Df(x)(f^{-1}(z) - f^{-1}(y)) - (z - y)\|}{\|z - y\|} = \|Df(x)^{-1}\| \cdot \frac{\|Df(x)(f^{-1}(z) - x) - (z - y)\|}{\|z - y\|}.$$

Now write $z = f(w)$ for $w \in V$. Using the inequality of Lemma 1.9.3, we have $\|z - y\| = \|f(w) - f(x)\| \geq c\|w - x\|$,

This shows the above quantity is less than or equal to

$$\frac{\|Df(x)^{-1}\|}{c} \cdot \frac{\|Df(x)(w - x) - (z - y)\|}{\|w - x\|} = \frac{\|Df(x)^{-1}\|}{c} \cdot \frac{\|Df(x)(w - x) - (f(w) - f(x))\|}{\|w - x\|}$$

Finally, since f^{-1} is continuous, as $z \rightarrow y$, we have $w = f^{-1}(z)$ goes to $x = f^{-1}(y)$. Then the quantity $\frac{\|Df(x)(w - x) - (f(w) - f(x))\|}{\|w - x\|}$ goes to zero because f is differentiable. Together this shows the expression above goes to zero, so the quotient Q we start with also goes to zero, showing f is differentiable.

Lastly, the function Df^{-1} is continuous because it is the composition of two continuous functions, the inverse and Df : Df is continuous, and, by the exercise below, the map $A \mapsto A^{-1}$ is continuous under the operator norm in the space of invertible matrices. ■

Exercise 1.12 Prove that the map $f : GL_n(\mathbb{R}) \rightarrow GL_n(\mathbb{R})$ given by $f(A) = A^{-1}$ is continuous under the operator norm.

Here, $GL_n(\mathbb{R})$ is called the general linear group of degree n , which is the space of $n \times n$ invertible matrices, equipped with the operator norm.

Hint: Note that when A, B are two invertible matrices,

$$A^{-1} - B^{-1} = B^{-1}BA^{-1} - B^{-1}AA^{-1} = B^{-1}(B - A)A^{-1}.$$

1.9.2 Implicit function theorem

Closely related to the inverse function theorem is the implicit function theorem.

Suppose we have a C^1 function $F : \mathbb{R}^n \rightarrow \mathbb{R}$ and consider the hypersurface $S = \{F(x) = 0\}$. We learned that if S is a graph (that is, $F(x) = x_n - f(x_1, \dots, x_{n-1})$ for some $f : \mathbb{R}^{n-1} \rightarrow \mathbb{R}$), then this makes the computations more explicit and simple. Notice that in this case, $\nabla F = (-\nabla f, 1)$ is always nonzero.

We know in general this will not be the case, but perhaps locally we can write the solution set of $F(x) = 0$ as a graph in some small neighborhood if we have $\nabla F(x) \neq 0$.

In particular, if at a point $a \in S$, $\frac{\partial F}{\partial x_j}(a) \neq 0$, then we expect to be able to find a function $f : \mathbb{R}^{n-1} \rightarrow \mathbb{R}$ such that for a point $x \in S$ near a , we can write $x_j = f(x_1, \dots, x_{j-1}, x_{j+1}, \dots, x_n)$.

■ **Example 1.49** Let $F : \mathbb{R}^2 \rightarrow \mathbb{R}$ be defined by

$$F(x, y) = y^2 + x^2 - 1.$$

Here, the equation $F(x, y) = 0$ gives us the unit circle in $\mathbb{R}^2$. To write the circle locally as a graph, we notice that for $x \neq \pm 1$, we can solve for y in terms of x in a neighborhood, and we get

$$y = \sqrt{1 - x^2}, \quad \text{for } -1 < x < 1, y > 0$$

or

$$y = -\sqrt{1 - x^2}, \quad \text{for } -1 < x < 1, y < 0.$$

Note that, in either case, we have $F(x, y) = 0$. Also,

$$\frac{\partial F}{\partial y} = 2y \neq 0 \quad \text{since } y \neq 0.$$

At the points $(\pm 1, 0)$, the circle is not locally a graph of the x -coordinate, but it is a graph of the y -coordinate: we see that the gradient is $\nabla F = (\pm 2, 0)$ so $\frac{\partial F}{\partial x} \neq 0$, and we can solve for x in terms of y similarly. ■

We make this idea precise in the following theorem. For ease of notation, let's write the domain as $W \subset \mathbb{R} \times \mathbb{R}^n$ and a point in this domain as (y, x) , where $y \in \mathbb{R}$ is the variable that we want to write as a graph of $x \in \mathbb{R}^n$.

Theorem 1.9.6 — Implicit function theorem for functions from $\mathbb{R} \times \mathbb{R}^n \rightarrow \mathbb{R}$. Let W be an open set in $\mathbb{R} \times \mathbb{R}^n$. Let $F : W \rightarrow \mathbb{R}$ be a C^1 function. Let $(y_0, x_0) \in W$ be a point in this open set such that $F(y_0, x_0) = 0$, and

$$\frac{\partial F}{\partial y}(y_0, x_0) \neq 0.$$

Then, there exist open sets $V \subset \mathbb{R}$ and $U \subset \mathbb{R}^n$, such that $(y_0, x_0) \in V \times U$, and for every $x \in U$, there exists a unique $\varphi(x)$ in V such that $F(\varphi(x), x) = 0$, and $\varphi : U \subset \mathbb{R}^n \rightarrow \mathbb{R}$ is C^1 on U . Furthermore,

$$\frac{\partial \varphi}{\partial x_j}(x) = -\frac{\frac{\partial F}{\partial x_j}(\varphi(x), x)}{\frac{\partial F}{\partial y}(\varphi(x), x)}.$$

Proof. We can either prove it directly, or use the inverse function theorem to prove it. Here we will choose the second option.

From our map $F : W \rightarrow \mathbb{R}$ we define the map

$$G : W \subset \mathbb{R} \times \mathbb{R}^n \rightarrow \mathbb{R} \times \mathbb{R}^n \quad \text{by} \quad G(y, x) = (F(y, x), x).$$

The Jacobi matrix of G at (y_0, x_0) is:

$$DG(y_0, x_0) = \begin{bmatrix} \frac{\partial F}{\partial y}(y_0, x_0) & \frac{\partial F}{\partial x_1}(y_0, x_0) & \cdots & \frac{\partial F}{\partial x_n}(y_0, x_0) \\ 0 & 1 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & 1 \end{bmatrix}.$$

Since $\frac{\partial F}{\partial y}(y_0, x_0) \neq 0$, this matrix is invertible. (This is because this is an upper triangular matrix and the eigenvalues are simply the diagonal entries and are all nonzero. The determinant is the product of the diagonal entries and is nonzero.) Thus, by the Inverse Function Theorem, there exists a neighborhood $W' \subset \mathbb{R}^{n+1}$ of (y_0, x_0) such that we can invert G on W' to get an inverse $G^{-1} : G(W') \rightarrow W'$ which is a C^1 function. Then we can choose neighborhood $U' \subset \mathbb{R}^n$ of x_0 and $V \subset \mathbb{R}$ of y_0 such that $V \times U' \subset W'$.

Since $F(y_0, x_0) = 0$, we get $G(y_0, x_0) = (0, x_0)$. For points $(0, x) \in G(V \times U')$, we get

$$G^{-1}(0, x) = (y, x), \quad \text{where } F(y, x) = 0.$$

Let $U = \{x \in U' \mid (0, x) \in G(V \times U')\} \subset \mathbb{R}^n$. This is a nonempty because $x_0 \in U$. It is open because $G(V \times U')$ is open (it is the inverse image of $V \times U'$ under the continuous function G^{-1} , hence open).

Define $\varphi : U \subset \mathbb{R}^n \rightarrow \mathbb{R}$ by letting $\varphi(x)$ be the first component y of $G^{-1}(0, x)$. Then, φ is C^1 because G^{-1} is, and $F(\varphi(x), x) = 0$ for all x near x_0 .

To compute the derivative of φ , differentiate the identity $F(\varphi(x), x) = 0$ with respect to x_j using chain rule, we get

$$\frac{\partial F}{\partial y}(\varphi(x), x) \cdot \frac{\partial \varphi}{\partial x_j}(x) + \frac{\partial F}{\partial x_j}(\varphi(x), x) = 0.$$

(You can also write down the chain rule using matrices and figure out what the entry for x_j is.)

Solving for $\frac{\partial \varphi}{\partial x_j}(x)$, we obtain:

$$\frac{\partial \varphi}{\partial x_j}(x) = -\frac{\partial F}{\partial x_j}(\varphi(x), x) \Big/ \frac{\partial F}{\partial y}(\varphi(x), x).$$

■

The above theorem tells when we can write functions that map to $\mathbb{R}$ as a graph. To put it another way, it also tells us when we can solve a variable in terms of the other variables.

Now we want to generalize this to higher dimensional target. Suppose we have a C^1 function $F : \mathbb{R}^N \rightarrow \mathbb{R}^m$. We consider the set $S = \{z \in \mathbb{R}^N \mid F(z) = \mathbf{0}\}$. If we let $S_i = \{z \in \mathbb{R}^N \mid F_i(z) = 0\}$, then this is a hypersurface if $\nabla F_i \neq 0$, and $S = \bigcap_{i=1}^m S_i$ is the intersection. In general we expect this to be an $(N - m)$ dimensional object, so we want to know if we can locally parametrize it using $N - m$ parameters. In particular, we can ask if we can parametrize it using $N - m$ coordinates.

For ease of notation, let $N = m + n$ and $\mathbb{R}^N = \mathbb{R}^m \times \mathbb{R}^n$. For a function $F : \mathbb{R}^m \times \mathbb{R}^n \rightarrow \mathbb{R}^m$, for simplicity let's just ask if we can parametrize the zero set in terms of the last n variables. At a point $(y, x) \in \mathbb{R}^m \times \mathbb{R}^n$, we can write

$$DF(y, x) = \left[\begin{array}{ccc|ccc} \frac{\partial F_1}{\partial y_1}(y, x) & \cdots & \frac{\partial F_1}{\partial y_m}(y, x) & \frac{\partial F_1}{\partial x_1}(y, x) & \cdots & \frac{\partial F_1}{\partial x_n}(y, x) \\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ \frac{\partial F_m}{\partial y_1}(y, x) & \cdots & \frac{\partial F_m}{\partial y_m}(y, x) & \frac{\partial F_m}{\partial x_1}(y, x) & \cdots & \frac{\partial F_m}{\partial x_n}(y, x) \end{array} \right] = [L(y, x) \mid R(y, x)].$$

A natural condition is to require the matrix $L(y_0, x_0)$ to be invertible. Indeed, if this is the case, then the rows of $L(y_0, x_0)$ will be linearly independent, so the rows of $DF(y_0, x_0)$ are also linearly independent. But the rows of $DF(y_0, x_0)$ are exactly the directives of F_i 's, so this is saying the gradient vectors $\nabla F_i(y_0, x_0)$ are linearly independent. If you recall, this is exactly the condition we need in the Lagrange multipliers for several constraints. In that case, it guarantees S is nice. In our case, we also expect that this condition shows S is nice.

Theorem 1.9.7 — Implicit function theorem. Let W be an open set in $\mathbb{R}^m \times \mathbb{R}^n$. Let $F : W \rightarrow \mathbb{R}^m$ be a C^1 function. Let $(y_0, x_0) \in W$ be a point in this open set such that $F(y_0, x_0) = \mathbf{0}$, and suppose that the $m \times m$ matrix

$$L(y_0, x_0) = \left[\frac{\partial F_i}{\partial y_j}(y_0, x_0) \right]_{i,j=1,\dots,m}$$

is invertible.

Then, there exist open sets $V \subset \mathbb{R}^m$ and $U \subset \mathbb{R}^n$, such that $(y_0, x_0) \in V \times U$, and for every $x \in U$, there exists a unique $\varphi(x)$ in V such that $F(\varphi(x), x) = \mathbf{0}$, and $\varphi : U \subset \mathbb{R}^n \rightarrow \mathbb{R}^m$ is C^1 on U . Furthermore,

$$D\varphi(x) = -(L(\varphi(x), x))^{-1}(R(\varphi(x), x)),$$

where $L(\varphi(x), x)$ is the $m \times m$ matrix and $R(\varphi(x), x)$ is the $m \times n$ matrix as above.

Proof. Again we use the inverse function theorem to prove it.

From our map $F : W \rightarrow \mathbb{R}^m$ we define the map

$$G : W \subset \mathbb{R}^m \times \mathbb{R}^n \rightarrow \mathbb{R}^m \times \mathbb{R}^n \quad \text{by} \quad G(y, x) = (F(y, x), x).$$

The Jacobi matrix of G at (y_0, x_0) is:

$$DG(y_0, x_0) = \left[\begin{array}{ccc|ccc} \frac{\partial F_1}{\partial y_1}(y_0, x_0) & \cdots & \frac{\partial F_1}{\partial y_m}(y_0, x_0) & \frac{\partial F_1}{\partial x_1}(y_0, x_0) & \cdots & \frac{\partial F_1}{\partial x_n}(y_0, x_0) \\ \vdots & & \vdots & \vdots & & \vdots \\ \frac{\partial F_m}{\partial y_1}(y_0, x_0) & \cdots & \frac{\partial F_m}{\partial y_m}(y_0, x_0) & \frac{\partial F_m}{\partial x_1}(y_0, x_0) & \cdots & \frac{\partial F_m}{\partial x_n}(y_0, x_0) \\ 0 & \cdots & 0 & 1 & \cdots & 0 \\ \vdots & & \vdots & \vdots & \ddots & \vdots \\ 0 & \cdots & 0 & 0 & \cdots & 1 \end{array} \right].$$

The upper left matrix is precisely $L(y_0, x_0)$. If you recall your knowledge of the determinant, we have $\det DG(y_0, x_0) = \det L(y_0, x_0) \neq 0$, so $DG(y_0, x_0)$ is invertible.

Then we can apply the Inverse Function Theorem, and get that there exists a neighborhood $U' \subset \mathbb{R}^n$ of x_0 and $V \subset \mathbb{R}^m$ of y_0 such that we can invert $G|_{V \times U'}$ to get a C^1 inverse $G^{-1} : G(V \times U') \rightarrow V \times U'$. Let $U = \{x \in U' \mid (0, x) \in G(V \times U')\} \subset \mathbb{R}^n$. This is again nonempty open.

The rest of the proof is very similar to that of Theorem 1.9.6 and is left as an exercise. ■

■ **Example 1.50** We can use the implicit function theorem to tell when we can solve for some variables in terms of the other ones. Consider the system of equations

$$3x + \sin(y) - z + u^2 = 0$$

$$-y + 2z^2 + u = 0$$

$$-3z + 2u = 0$$

We want to show that we can solve for x, y, u in terms of z . For this, we just need to form the matrix

$\begin{bmatrix} 3 & \cos(y) & 2u \\ 0 & -1 & 1 \\ 0 & 0 & 2 \end{bmatrix}$ and notice it's invertible. (Again, this is a triangular matrix so the determinant is just the product of the diagonal entries.) ■

2. Integration in one variable

2.1 The Riemann Integral in One Variable: Definitions

We begin with a discussion of the Riemann integral in $\mathbb{R}$. The basic idea for integration is to calculate signed areas under curves.

We start with a closed, bounded interval $[a, b] \subset \mathbb{R}$ and a bounded function $f : [a, b] \rightarrow \mathbb{R}$. As we proceed through the definition, we will derive conditions on the function f that ensure the existence of the Riemann integral.

The first definition is about how we chop up the interval.

Definition 2.1.1 — partition, mesh. A *partition* P of the interval $[a, b]$ is a finite set of points

$$P = \{a_0, a_1, \dots, a_k\}$$

such that

$$a = a_0 < a_1 < \dots < a_{k-1} < a_k = b.$$

The *mesh* of the partition P is defined by

$$|P| = \max_{1 \leq i \leq k} (a_i - a_{i-1}).$$

Let f be a bounded function on $[a, b]$ with $m \leq f(x) \leq M$ for all $x \in [a, b]$.

Definition 2.1.2 — upper sum, lower sum. Let $[a, b] \subset \mathbb{R}$, and let $f : [a, b] \rightarrow \mathbb{R}$ be a bounded function. If P is a partition of $[a, b]$, we set

$$m_i = \inf_{x \in [a_{i-1}, a_i]} f(x), \quad \text{and} \quad M_i = \sup_{x \in [a_{i-1}, a_i]} f(x).$$




Figure 2.1: lower sum, upper sum

We define the *lower sum* of f on $[a, b]$ relative to the partition P to be

$$L(f, P) = \sum_{i=1}^k m_i(a_i - a_{i-1}),$$

and the *upper sum* of f on $[a, b]$ relative to P to be

$$U(f, P) = \sum_{i=1}^k M_i(a_i - a_{i-1}).$$

R It is not hard to see that given a partition P of $[a, b]$, we have

$$m(b - a) \leq L(f, P) \leq U(f, P) \leq M(b - a).$$

Given a partition, if we divide it into finer intervals, we get a refinement.

Definition 2.1.3 — refinement. Let $P = \{a_j \mid j = 0, 1, 2, \dots, k\}$ be a partition of $[a, b]$. A partition $P' = \{a'_j \mid j = 0, 1, 2, \dots, k'\}$, where $k' \geq k$, is called a *refinement* of P if $P \subseteq P'$.




Figure 2.2: refinement

R Let P be a partition of $[a, b]$, and let P' be a refinement of P . Then

$$L(f, P) \leq L(f, P') \leq U(f, P') \leq U(f, P).$$

Lemma 2.1.1 Let $f : [a, b] \rightarrow \mathbb{R}$ be a bounded function, and let P_1 and P_2 be any two partitions of $[a, b]$. Then $L(f, P_1) \leq U(f, P_2)$.

Proof. Let $P' = P_1 \cup P_2$. Then P' is a refinement of both P_1 and P_2 . Then

$$L(f, P_1) \leq L(f, P') \leq U(f, P') \leq U(f, P_2).$$

■

We are now ready to define the lower and upper integrals of f on $[a, b]$.

Definition 2.1.4 — lower integral, upper integral. Let $f : [a, b] \rightarrow \mathbb{R}$ be a bounded function.

1. We define the *lower integral* of f on $[a, b]$ to be

$$\int_a^b f = \sup\{L(f, P) \mid P \text{ is a partition of } [a, b]\}.$$

2. We define the *upper integral* of f on $[a, b]$ to be

$$\overline{\int_a^b} f = \inf\{U(f, P) \mid P \text{ is a partition of } [a, b]\}.$$

R Note that the set of lower sums is bounded above by $M(b - a)$, and the set of upper sums is bounded below by $m(b - a)$. Therefore, the lower and upper integrals exist for any bounded function.

Lemma 2.1.2 Let $f : [a, b] \rightarrow \mathbb{R}$ be a bounded function. Then $\int_a^b f \leq \overline{\int_a^b} f$.

Proof. Exercise. ■

Definition 2.1.5 Let f be a bounded function on an interval $[a, b]$. We say that f is *Riemann integrable* on $[a, b]$ if

$$\int_a^b f = \overline{\int_a^b f},$$

and we denote this common value by

$$\int_a^b f,$$

which we call the *Riemann integral* of f on $[a, b]$.

We often refer to a function that is Riemann integrable as simply *integrable*.

■ **Example 2.1** It easily follows from the definition that any constant function $f = c$ on $[a, b]$ is integrable, and $\int_a^b f = c(b - a)$. ■

■ **Example 2.2** The Dirichlet function $f : [0, 1] \rightarrow \mathbb{R}$ is defined by

$$f(x) = \begin{cases} 1, & \text{if } x \in \mathbb{Q} \cap [0, 1], \\ 0, & \text{if } x \in [0, 1] \setminus \mathbb{Q}. \end{cases}$$

That is, f is one at every rational number and zero at every irrational number. It is not integrable. ■

2.2 The Riemann Integral in One Variable: Properties

One useful criterion that follows directly from the definition is the following.

Theorem 2.2.1 Let f be a bounded function on an interval $[a, b]$. Then f is integrable on $[a, b]$ if and only if, given $\varepsilon > 0$, there exists a partition P of $[a, b]$ such that

$$U(f, P) - L(f, P) < \varepsilon.$$

Proof. This useful criterion follows directly from the definitions of integrability and the upper and lower integrals. ■

What's more, we also have the following proposition.

Proposition 2.2.2 Let $f : [a, b] \rightarrow \mathbb{R}$ be a bounded function on $[a, b]$.

1. f is integrable on $[a, b]$ if and only if, given $\varepsilon > 0$, there exists $\delta > 0$ such that if P is a partition with $|P| < \delta$, then

$$U(f, P) - L(f, P) < \varepsilon.$$

2. Suppose f is integrable on $[a, b]$, and we have a sequence of partitions $(P_k)_{k \in \mathbb{N}}$ such that $\lim_{k \rightarrow \infty} |P_k| = 0$. Then

$$\int_a^b f = \lim_{k \rightarrow \infty} L(f, P_k) = \lim_{k \rightarrow \infty} U(f, P_k).$$

Using this criterion, we have

Theorem 2.2.3 Suppose $f : [a, b] \rightarrow \mathbb{R}$ is continuous. Then f is Riemann integrable on $[a, b]$.

Proof. Fix $\varepsilon > 0$. Since f is continuous on $[a, b]$, a compact set in $\mathbb{R}$, it follows that f is uniformly continuous on $[a, b]$. So there is a $\delta > 0$ such that $|x_1 - x_2| < \delta$ implies $|f(x_1) - f(x_2)| < \frac{\varepsilon}{b-a}$.

Choose a partition P such that $|P| < \delta$. Then

$$U(f, P) - L(f, P) = \sum_{i=1}^k (M_i - m_i)(x_i - x_{i-1}) < \frac{\varepsilon}{b-a} \sum_{i=1}^k (x_i - x_{i-1}) = \varepsilon.$$

■

Proposition 2.2.4 Suppose that $f, g : [a, b] \rightarrow \mathbb{R}$ are integrable on $[a, b]$, and $\alpha \in \mathbb{R}$. Then:

1. (Linearity) $\int_a^b (f + g) = \int_a^b f + \int_a^b g$, $\int_a^b \alpha f = \alpha \int_a^b f$.
2. (Additivity) If $c \in [a, b]$, then f is integrable on $[a, c]$ and $[c, b]$, and

$$\int_a^b f = \int_a^c f + \int_c^b f.$$

3. (Monotonicity) If $f(x) \leq g(x)$ for all $x \in [a, b]$, then

$$\int_a^b f \leq \int_a^b g.$$

Exercise 2.1 1. Show if f is integrable, then $|f|$ is integrable, and

$$\left| \int_a^b f \right| \leq \int_a^b |f|.$$

2. Show the converse to part (1) is not true by considering the following function on $[0, 1]$:

$$f(x) = \begin{cases} 1, & \text{if } x \in \mathbb{Q} \cap [0, 1], \\ -1, & \text{if } x \in [0, 1] \setminus \mathbb{Q}. \end{cases}$$

Show $|f|$ is integrable, but f is not.

■

■ **Example 2.3** Let $f(x) = x^2$ on $[-1, 1]$. Then f is integrable on $[-1, 1]$ because f is continuous, and we compute $\int_{-1}^1 f(x) dx$ as follows. By Proposition 2.2.2, it suffices to consider partitions whose mesh goes to zero, and we make our computation easier by choosing regular partitions. That is, we let each sub-interval to have the same length; for each $k \in \mathbb{N}$, we let

$$P_k = \left\{ -1, -1 + \frac{1}{k}, -1 + \frac{2}{k}, \dots, -1 + \frac{2(k-1)}{k}, 1 \right\}.$$

Note that $|P_k| = \frac{1}{k}$. Furthermore, it suffices to compute the upper integral. We have

$$U(f, P_k) = \sum_{i=1}^{2k} M_i \cdot \frac{1}{k} = 2 \sum_{i=1}^k M_{k+i} \cdot \frac{1}{k}$$

using the symmetry, where $M_{k+i} = \left(\frac{i}{k}\right)^2$, so

$$U(f, P_k) = 2 \sum_{i=1}^k \left(\frac{i}{k}\right)^2 \cdot \frac{1}{k} = \frac{2}{k^3} \sum_{i=1}^k i^2 = \frac{2}{k^3} \cdot \frac{k(k+1)(2k+1)}{6}.$$

Here we used the identity $\sum_{i=1}^k i^2 = \frac{k(k+1)(2k+1)}{6}$, which can be checked by induction. Then,

$$\int_{-1}^1 f(x) dx = \lim_{k \rightarrow \infty} \frac{2}{k^3} \cdot \frac{k(k+1)(2k+1)}{6} = \frac{2}{3}.$$

■

2.3 The Fundamental Theorem of Calculus and Its Consequences

The next theorem is often regarded as the most important theorem in integral calculus in one variable. The fundamental theorem of calculus links the concept of differentiating a function with the concept of integrating a function. Roughly speaking, the two operations can be thought of as inverses of each other.

Theorem 2.3.1 — Fundamental Theorem of Calculus. Let f be Riemann integrable on $[a, b]$. For $x \in [a, b]$, define

$$F(x) = \int_a^x f(t) dt.$$

If f is continuous at a point $x \in (a, b)$, then F is differentiable at x , and

$$F'(x) = f(x).$$

Proof. Fix $\varepsilon > 0$. To show F is differentiable, we consider the difference quotient

$$\frac{F(x+h) - F(x)}{h} = \frac{1}{h} \left(\int_a^{x+h} f(t) dt - \int_a^x f(t) dt \right).$$

Assume first that $h > 0$. Then the difference quotient is equal to

$$\frac{1}{h} \int_x^{x+h} f(t) dt.$$

Since f is continuous at x , we can choose h small enough so that $|f(x+t) - f(x)| < \varepsilon$ for $t \in [0, h]$. Then

$$(f(x) - \varepsilon)h < \int_x^{x+h} f(t) dt < (f(x) + \varepsilon)h.$$

So

$$f(x) - \varepsilon < \frac{1}{h} \int_x^{x+h} f(t) dt < f(x) + \varepsilon.$$

Since ε was arbitrary, it follows that the difference quotient converges to $f(x)$. For $h < 0$, the argument is similar. ■

Corollary 2.3.2 Suppose f is continuous on $[a, b]$, and G is continuous on $[a, b]$ and differentiable on (a, b) with $G'(x) = f(x)$ for $x \in (a, b)$. Then $\int_a^b f = G(b) - G(a)$. Moreover, if $F = \int_a^x f$ is defined as above, then $F(x) = G(x) - G(a)$.

Proof. Since $G' = F'$ on (a, b) , we have $(G - F)' = 0$, so it follows from Corollary 1.1.6 that $G - F$ is constant: there exists a constant C such that $F(x) - G(x) = C$. It is clear that $C = F(a) - G(a) = -G(a)$. ■



The function G is often referred to as a primitive or an antiderivative of f . A substantial portion of many single-variable calculus courses consists of a search for primitives.

■ **Example 2.4** Let $f(x) = x^2$ on $[-1, 1]$.

Since we know $(x^n)' = nx^{n-1}$, an antiderivative of $f = x^2$ is $\frac{1}{3}x^3$. Thus $\int_{-1}^1 f(x) dx = [\frac{1}{3}x^3]_{-1}^1 = \frac{1}{3} - (-\frac{1}{3}) = \frac{2}{3}$. ■

For the next corollary, we need to introduce the traditional notation for the integral. That is, if f is integrable on $[a, b]$, we write $\int_a^b f = \int_a^b f(x)dx = \int_a^b f(t)dt$, where x , t , and any other symbol are called “dummy” variables.

Corollary 2.3.3 — Change of variables in $\mathbb{R}$. Let f be continuous on $[a, b]$, and let $\varphi : [c, d] \rightarrow [a, b]$ be continuous on $[c, d]$ and C^1 on (c, d) such that $\varphi(c) = a$, $\varphi(d) = b$, and $\varphi'(x) > 0$ for all $x \in (c, d)$. Then

$$\int_a^b f(u) du = \int_c^d f(\varphi(x))\varphi'(x) dx.$$

Proof. The proof uses the fundamental theorem of calculus and chain rule.

Define

$$F(u) := \int_a^u f(t) dt.$$

Since f is continuous on $[a, b]$, the function F is differentiable on (a, b) with $F'(u) = f(u)$.

Now define

$$G(x) := F(\varphi(x)) \quad \text{for } x \in [c, d].$$

Since $\varphi \in C^1((c, d))$ and $F \in C^1([a, b])$, the composition $G(x) = F(\varphi(x))$ is differentiable on (c, d) , and by the chain rule:

$$G'(x) = F'(\varphi(x)) \cdot \varphi'(x) = f(\varphi(x)) \cdot \varphi'(x).$$

Then:

$$\int_c^d f(\varphi(x))\varphi'(x) dx = \int_c^d G'(x) dx = G(d) - G(c) = F(\varphi(d)) - F(\varphi(c)).$$

But $\varphi(c) = a$, $\varphi(d) = b$, so:

$$F(\varphi(d)) - F(\varphi(c)) = F(b) - F(a) = \int_a^b f(t) dt.$$

Hence,

$$\int_c^d f(\varphi(x))\varphi'(x) dx = \int_a^b f(t) dt.$$

■

The change of variables theorem is what allows the classic integration technique known as “integration by substitution.”

■ **Example 2.5** For example, suppose we want to evaluate $\int_0^1 \sqrt{x+1} dx$. The function $g(x) = \sqrt{x+1}$ is not among the small number of functions that have an obvious antiderivative. However, we can write $g(x) = f(\varphi(x))\varphi'(x)$, where $f(u) = \sqrt{u}$, $\varphi(x) = x+1$, and $\varphi'(x) = 1$.

$$\int_0^1 \sqrt{x+1} dx = \int_1^2 \sqrt{u} du = \int_1^2 u^{1/2} du = \left[\frac{2}{3} u^{3/2} \right]_1^2 = \frac{2}{3} (2^{3/2} - 1) = \frac{2}{3} (2\sqrt{2} - 1).$$

■

Corollary 2.3.4 — Integration by parts. Suppose f and g are C^1 on $[a, b]$. Then

$$\int_a^b f(x)g'(x)dx = f(b)g(b) - f(a)g(a) - \int_a^b f'(x)g(x)dx.$$

Proof. Let $H(x) = f(x)g(x)$. Then $H'(x) = f'(x)g(x) + f(x)g'(x)$ by the product rule.

Then by Corollary 2.3.2, we have $\int_a^b H'(x) = H(b) - H(a)$. Writing this out and using linearity gives

$$\int_a^b f'(x)g(x) + \int_a^b f(x)g'(x) = f(b)g(b) - f(a)g(a).$$

■

In application, to choose f and g , choose an f that gets simpler when you differentiate it and a g that doesn't get any more complicated when you integrate it.

■ **Example 2.6** Suppose we want to evaluate $\int_0^1 xe^x dx$. We let $f(x) = x$ and $g'(x) = e^x$. Then $f'(x) = 1$ and $g(x) = e^x$, so

$$\int_0^1 xe^x dx = f(1)g(1) - f(0)g(0) - \int_0^1 f'(x)g(x) = e^1 - \int_0^1 e^x = e^1 - [e^x]_0^1 = e^1 - e^1 + e^0 = 1.$$

■

Corollary 2.3.5 — Mean value theorem for integrals. Let f be continuous on $[a, b]$. Then there exists a point $c \in (a, b)$ such that $f(c) = \frac{1}{b-a} \int_a^b f$.

 The quantity $\frac{1}{b-a} \int_a^b f$ is called the *average value* of f on $[a, b]$.

Proof. Let $F(x) = \int_a^x f$. Then $F' = f$. Since f is continuous, we have that F is C^1 . Apply the mean value theorem for derivatives to F , we get

$$F(b) - F(a) = F'(c)(b - a)$$

for some $c \in (a, b)$. Thus

$$\int_a^b f = f(c)(b - a).$$

■

2.4 Principal Value Integrals

So far, in this chapter, we have worked with bounded, real-valued functions defined on a closed, bounded interval in $\mathbb{R}$. This is sort of restrictive, and we want to extend the definition of integrals to both unbounded functions and unbounded intervals.

First, we consider the case of a bounded, real-valued function f defined on an unbounded interval of the form $[a, \infty)$.

Definition 2.4.1 — principal value integral (or improper integral). Suppose that $f : [a, \infty) \rightarrow \mathbb{R}$ is a bounded function, and suppose further that for all $b > a$, the function f is integrable on $[a, b]$. The *principal value integral* (also called *improper integral*)

$$\int_a^\infty f$$

is defined by

$$\lim_{b \rightarrow \infty} \int_a^b f.$$

Similarly, we can define the principal value integral $\int_{-\infty}^b f$.

■ **Example 2.7** Let $f(x) = \frac{1}{x^\alpha}$ on $[1, \infty)$, where $\alpha > 0$. Then, for any $b > 1$, f is integrable on $[1, b]$, and

$$\int_1^b \frac{1}{x^\alpha} dx = \begin{cases} \frac{b^{1-\alpha}-1}{1-\alpha}, & \text{if } \alpha \neq 1, \\ \log b, & \text{if } \alpha = 1. \end{cases}$$

If $\alpha > 1$, then the principal value integral exists and is equal to $\frac{1}{\alpha-1}$. If $\alpha \leq 1$, the limit diverges, so the principal value integral does not exist. ■

The case of the principal value integral of a bounded real-valued function on $\mathbb{R}$ is more subtle.

Suppose f is a bounded function on $\mathbb{R}$, and suppose further that f is integrable on any closed interval $[a, b]$. There are two potential impediments to f being integrable on $\mathbb{R}$: the first is the behavior of f toward $+\infty$, and the second is the behavior of f toward $-\infty$. We ensure that neither of these impediments is too great by defining the principal value integral of f on $\mathbb{R}$ as follows.

Definition 2.4.2 Suppose that $f : \mathbb{R} \rightarrow \mathbb{R}$ is a bounded function, and suppose further that f is integrable on any closed bounded interval in $\mathbb{R}$. The principal value integral of f on $\mathbb{R}$ is defined as

$$\int_{-\infty}^\infty f = \int_{-\infty}^0 f + \int_0^\infty f,$$

provided that the latter two principal value integrals exist.

■ **Example 2.8** If the principal value integral of f on $\mathbb{R}$ exists, then $\int_{-\infty}^\infty f = \lim_{a \rightarrow \infty} \int_{-a}^a f$.

On the other hand, we know $\lim_{a \rightarrow \infty} \int_{-a}^a \sin x = 0$ because sine is odd, but the principal value integral $\int_{-\infty}^\infty \sin x$ does not exist. ■

Next, we consider unbounded functions on bounded intervals.

Definition 2.4.3 Let $f : (a, b] \rightarrow \mathbb{R}$ be a function. Suppose further that for all $c \in (a, b)$, f is bounded and integrable on $[c, b]$. The principal value integral of f on $(a, b]$ is

$$\int_a^b f = \lim_{c \rightarrow a^+} \int_c^b f,$$

provided this limit exists.

R Similarly, we can define $\int_a^b f$ for $f : [a, b) \rightarrow \mathbb{R}$.

■ **Example 2.9** Let $f(x) = \frac{1}{x^\alpha} = x^{-\alpha}$ on the interval $(0, 1]$, where $\alpha > 0$. Then, for any $c \in (0, 1)$, f is integrable on $[c, 1]$, and

$$\int_c^1 \frac{1}{x^\alpha} dx = \begin{cases} \frac{1-c^{1-\alpha}}{1-\alpha}, & \text{if } \alpha \neq 1, \\ -\log c, & \text{if } \alpha = 1. \end{cases}$$

If $\alpha < 1$, then the principal value integral from 0 to 1 exists and is equal to

$$\int_0^1 \frac{1}{x^\alpha} dx = \frac{1}{1-\alpha}.$$

If $\alpha \geq 1$, the limit diverges, so the principal value integral does not exist. ■