Chapter 3

Differential Calculus in $\mathbb{R}^n$

In the first chapter, we learned about metric spaces, their topology, and specifically, about metric spaces of real-valued functions. Note that we have almost not dealt with differential and integral calculus. The reason is that derivatives and integrals are associated with limits of differences and sums, which are not pertinent to general metric spaces. The notions of derivatives and integrals can be defined for functions that take values in normed spaces, which, as we learned, are a subclass of metric spaces. In this chapter we will develop the differential calculus of functions between *finite-dimensional* normed spaces, which are all isomorphic to a Euclidean space. Differential calculus can also be constructed for infinite-dimensional normed spaces, but this is beyond the scope of the present course (physics students who studied analytical mechanics encountered **functional derivatives**, which are derivatives of mappings between functions and real numbers).

3.1 Differentiability and derivatives

In this chapter we consider functions $f : \mathbb{R}^k \to \mathbb{R}^m$. More generally, we will consider functions $f : A \to \mathbb{R}^m$, where $A \subset \mathbb{R}^k$; we will always assume that A has a non-empty interior, so that there exist points $a \in A$ that have an open neighborhood in the domain of f.

We will usually denote points in A by $a, b, \dots$; we will commonly use the symbols $x, y, \dots$ to denote vectors in $\mathbb{R}^k$ connecting two points in A. If a is an interior point of A, there there exists a sufficiently small $r > 0$, such that $a + tx \in A$ for for every $x \in \mathbb{R}^k$ and $t \in (-r, r)$.

Each function $f : A \to \mathbb{R}^m$ is a collection of m functions $f_j : A \to \mathbb{R}$, each depending on k variables:

$$f(a) = (f_1(a_1, \dots, a_k), \dots, f_m(a_1, \dots, a_k)).$$

We will denote vectors in $\mathbb{R}^k$ as column vectors, whereas row vectors will represent linear functionals on vectors, acting via multiplication; sometimes, we will not do so for notational convenience.

Example: Consider the function $f : \mathbb{R}^2 \to \mathbb{R}^3$,

$$f\begin{pmatrix} a_1 \\ a_2 \end{pmatrix} = \begin{pmatrix} a_1^2 \sin a_2 \\ \sqrt{a_1/a_2} \\ a_1 \end{pmatrix}.$$

The vector-valued function f has three component, given by

$$f_1(a) = a_1^2 \sin a_2 \qquad f_2(a) = \sqrt{a_1/a_2} \qquad \text{and} \qquad f_3(a) = a_1.$$

▲▲▲

By default, we endow $\mathbb{R}^n$ with the Euclidean norm, which we will denote by $\|\cdot\|_n$, to have an explicit mention to the dimension of the space. Since $\mathbb{R}^k$ and $\mathbb{R}^m$ are metric spaces, we have a well-defined notion of continuity. A function $f : \mathbb{R}^k \to \mathbb{R}^m$ is continuous at $a \in \mathbb{R}^k$. if for every sequence a_n converging to a, i.e.,

$$\lim_{n\to\infty} a_n = a \qquad \text{in } \mathbb{R}^k,$$

we have

$$\lim_{n\to\infty} f(a_n) = f(a) \qquad \text{in } \mathbb{R}^m.$$

Moreover, since convergence in $\mathbb{R}^m$ amounts to the convergence of each component, f is continuous at a if $a_n \to a$ implies that

$$\lim_{n\to\infty} f_j(a_n) = f_j(a) \qquad j = 1, \dots, m.$$

The most elementary functions from $\mathbb{R}^k$ to $\mathbb{R}^m$ are **linear maps**, which are represented by m-by-k matrices acting on vectors in $\mathbb{R}^k$ via matrix-vector multiplication. That is, a linear map $T : \mathbb{R}^k \to \mathbb{R}^m$ has a representation

$$T\begin{pmatrix} a_1 \\ \vdots \\ a_k \end{pmatrix} = \begin{pmatrix} T_{11} & \cdots & T_{1k} \\ \vdots & \ddots & \vdots \\ T_{m1} & \cdots & T_{mk} \end{pmatrix} \begin{pmatrix} a_1 \\ \vdots \\ a_k \end{pmatrix}.$$

If e_i denotes the i-th unit vector in $\mathbb{R}^k$, then

$$T(e_i) = \begin{pmatrix} T_{1i} \\ \vdots \\ T_{mi} \end{pmatrix},$$

so that $(e_j, T(e_i)) = T_{ji}$.

We denote the set of linear transformations from $\mathbb{R}^k$ to $\mathbb{R}^m$ by $\mathrm{Hom}(\mathbb{R}^k, \mathbb{R}^m)$ (a linear transformation is a homomorphism with respect to the group structure of the vector spaces). Recall that $\mathrm{Hom}(\mathbb{R}^k, \mathbb{R}^m)$ is a normed space with respect to the **operator norm** (הנורמה האופרטורית),

$$\|T\|_{k,m} = \max_{\|x\|_k = 1} \|Tx\|_m.$$

The operator norm on matrices is induced by the norm for vectors; we will always assume the vector norm to be the Euclidean norm (even though, as we recall, all the norms on $\mathbb{R}^n$ are equivalent). Note our choice of notation $\|\cdot\|_{k,m}$, which reminds us that we have an operator norm between a k-dimensional space and an m-dimensional space.

TA material 3.1 Obtain an explicit expression for the operator norm $\|\cdot\|_{k,m}$.

Example: Real-valued functions, $f : \mathbb{R} \to \mathbb{R}$, are best visualized by their **graphs**, as one-dimensional curves embedded in the plane. One coordinate of the plane—the **abscissa**—represents the value of the independent variable, whereas the second coordinate—the **ordinate**—represents the value of its image. Similarly, real-valued functions $f : \mathbb{R}^2 \to \mathbb{R}$ are visualized by their graph which is a (two-dimensional) surface embedded in $\mathbb{R}^3$ (see Figure 3.1). In the same way, the graph of a real-valued function $f : \mathbb{R}^n \to \mathbb{R}$ is an n-dimensional **manifold** (יריעה) embedded in $\mathbb{R}^{n+1}$. ▲▲▲

Back to general functions $f : \mathbb{R}^k \to \mathbb{R}^m$, our goal is to define a notion of differentiability. To this end, let's first recall how we define derivatives of functions $f : \mathbb{R} \to \mathbb{R}$. The most common definition for the derivative of f at a point $a \in \mathbb{R}$ is

$$f'(a) = \lim_{x \to 0} \frac{f(a+x) - f(a)}{x}.$$

Let now $f : \mathbb{R}^k \to \mathbb{R}^m$; if we replace the denominator by the norm of x, we may obtain a meaningful expression, in which case, the limit would be an element

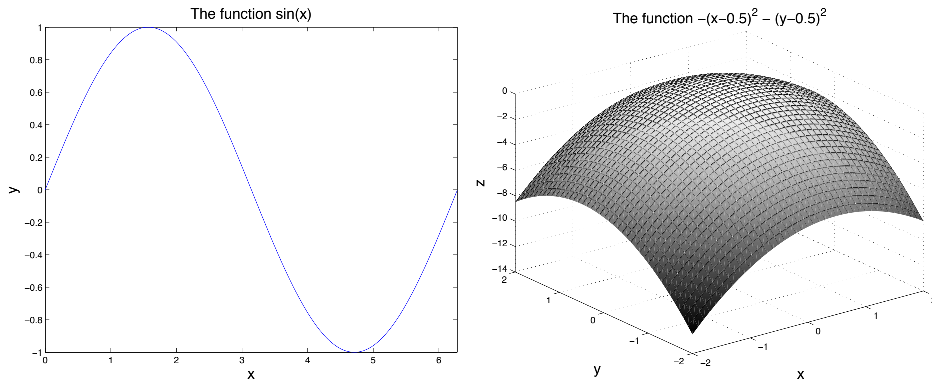

Figure 3.1: The graphs of real-valued functions $\mathbb{R} \to \mathbb{R}$ (left) and $\mathbb{R}^2 \to \mathbb{R}$ (right).

of $\mathbb{R}^m$. It turns out, however, that this is not the natural generalization of the derivative.

Recall that the derivative is (the limit of) a ratio between a variation in the image and a variation in the domain,

$$\Delta f \simeq f'(a)\, \Delta x.$$

Viewed under this angle, $f'(a)$ is a linear transformation from $\mathbb{R}$ to $\mathbb{R}$. In the multivariate case, $\Delta x \in \mathbb{R}^k$ and $\Delta f \in \mathbb{R}^m$, hence the derivative should be a linear transformation from $\mathbb{R}^k$ to $\mathbb{R}^m$.

In fact, an equivalent definition of the derivative of a univariate function is the following: we say that f is differentiable at a if there exists a real number T such that

$$\lim_{x\to 0} \frac{f(a+x) - f(a) - Tx}{|x|} = 0,$$

and this number T, which can be shown to be unique, is the derivative of f at a. Note that although a and x are both real-numbers, they play different roles: a is a point in $\mathbb{R}$, whereas x is a **displacement** in $\mathbb{R}$ (in a more general context, x is said to be an element to the space **tangent** (מרחב משיק) to $\mathbb{R}$ at the point a).

Having formulated differentiability this way, we think of the number T as a linear transformation, converting the displacement x from the point a, into a displacement Tx from the image $f(a)$. This interpretation of the derivative can be generalized in a natural way into mappings from $\mathbb{R}^k$ to $\mathbb{R}^m$:

*Definition 3.1 Let $A \subset \mathbb{R}^k$ and let $a \in A$ be an interior point. A function $f : A \to \mathbb{R}^m$ is said to be **differentiable** (דיפרנציאבילית) at a is there exists a linear transformation $T \in \mathrm{Hom}(\mathbb{R}^k, \mathbb{R}^m)$ such that*

$$\lim_{x\to 0} \frac{f(a+x) - f(a) - Tx}{\|x\|_k} = 0. \tag{3.1}$$

*We will denote the linear transformation T by $(Df)_a$ (the **derivative** of f at the point a). We denote by $(Df)_a(x)$ the differential of f at a operating on a vector $x \in \mathbb{R}^k$, resulting in a vector in $\mathbb{R}^m$.*

Comment: Recall that the limit in (3.1) amounts to a component-wise limit, i.e., for all $j = 1, \dots, m$,

$$\lim_{x\to 0} \frac{f_j(a+x) - f_j(a) - (Tx)_j}{\|x\|_k} = 0.$$

Comment: The **derivative function** (פונקצית הנגזרת) Df is a (generally nonlinear) function,

$$Df : \mathbb{R}^k \to \mathrm{Hom}(\mathbb{R}^k, \mathbb{R}^m).$$

Its evaluation at a point a is denoted $(Df)_a$ rather than $Df(a)$, since it is a linear operator acting on vectors $x \in \mathbb{R}^k$, and the notation $Df(a)(x)$ would be somewhat confusing. Note again the different role played by a and x; the first is an element of a set A, where the second is an element in a vector space.

We first need to check that the derivative $(Df)_a$ is well defined, i.e., that if a linear transformation with the required properties exists, then it is unique.

Proposition 3.2 If f is differentiable at a then the linear transformation T in the definition (3.1) *is unique.*

Proof: Suppose that $T, S \in \mathrm{Hom}(\mathbb{R}^k, \mathbb{R}^m)$ both satisfy

$$\lim_{x\to 0} \frac{f(a+x) - f(a) - Tx}{\|x\|_k} = 0 \qquad \text{and} \qquad \lim_{x\to 0} \frac{f(a+x) - f(a) - Sx}{\|x\|_k} = 0.$$

For every finite displacement $x \in \mathbb{R}^n$ for which $a + x \in A$,

$$(S - T)x = (f(a + x) - f(a) - Tx) - (f(a + x) - f(a) - Sx).$$

By the triangle inequality,

$$\lim_{x\to 0} \frac{\|(S - T)x\|_m}{\|x\|_k} \le \lim_{x\to 0} \frac{\|f(a + x) - f(a) - Tx\|_m}{\|x\|_k} + \lim_{x\to 0} \frac{\|f(a + x) - f(a) - Sx\|_m}{\|x\|_k} = 0.$$

Using the homogeneity of the norm,

$$\lim_{x\to 0} \left\|(S - T)\left(\frac{x}{\|x\|_k}\right)\right\|_m = 0.$$

This limit holds for every sequence $x \to 0$. Taking $x = te_j$ and letting $t \to 0$, we deduce that for every j,

$$(S - T)(e_j) = 0,$$

which implies that $S - T = 0$. ■

At this stage, the meaning of $(Df)_a$ may seem intriguing; in particular, what is the significance of $(Df)_a$ acting on a vector x? The following proposition provides an answer.

Proposition 3.3 *Let $A \subset \mathbb{R}^k$ and let $f : A \to \mathbb{R}^m$ be differentiable at an interior point $a \in A$. Then, for every $x \in \mathbb{R}^k$,*

$$(Df)_a(x) = \lim_{t\to 0} \frac{f(a + tx) - f(a)}{t}. \tag{3.2}$$

That is, $(Df)_a(x)$ is the rate of change of f when moving from the point a along the direction x.

Proof: By definition, if f is differentiable at a, then for every $x \neq 0$,

$$\lim_{t\to 0} \frac{f(a + tx) - f(a) - (Df)_a(tx)}{\|tx\|_k} = 0.$$

Using the linearity of $(Df)_a$ and the homogeneity of the norm, if f is differentiable at a, then

$$\lim_{t\to 0} \frac{f(a+tx) - f(a) - t\,(Df)_a(x)}{t} = 0,$$

which reduces to (3.2). ■

The vector $(Df)_a(x) \in \mathbb{R}^m$ it is the rate of change of f when we displace its argument from a in the x-direction; it is called the **directional derivative** (נגזרת כיוונית) of f at a in the x-direction (despite the name, note that x is a general vector, and not only a unit vector defining a direction). It is useful to note the following:

Corollary 3.4 Let $f : A \to \mathbb{R}$ and $a \in A$ be defined as above and let $x \in \mathbb{R}^k$. We Define the function $g : (-\varepsilon, -\varepsilon) \to \mathbb{R}$ where $\varepsilon > 0$ is sufficiently small such that

$$g(t) = f(a+tx).$$

Then,

$$g'(0) = (Df)_a(x).$$

Eq. (3.2) may be somewhat intriguing: think of it with respect to the vector $x \in \mathbb{R}^k$. On the left-hand side, there is linear operator $(Df)_a \in \mathrm{Hom}(\mathbb{R}^k, \mathbb{R}^m)$ acting on x. The right-hand side doesn't look linear in x. The fact that it is linear in x is due to f being differentiable. In other words, the following non-obvious relation holds,

$$\begin{aligned}\lim_{t\to 0} \frac{f(a+t(\alpha x+\beta y)) - f(a)}{t} &= \alpha \lim_{t\to 0} \frac{f(a+tx) - f(a)}{t} \\ &\quad + \beta \lim_{t\to 0} \frac{f(a+ty) - f(a)}{t}.\end{aligned}$$

✎ *Exercise 3.1* Find a continuous function f for which the limit (3.2) exists for all x, but it is nevertheless not a linear function of x.

Definition 3.5 Let $A \subset \mathbb{R}^k$ and let $f : A \to \mathbb{R}^m$ be differentiable at an interior point $a \in A$. For every $j = 1, \dots, k$ we define the ***partial derivative*** *(נגזרת חלקית) of f at a in the j-th direction,*

$$\partial_j f(a) = (Df)_a(e_j) = \lim_{t\to 0} \frac{f(a+te_j) - f(a)}{t}.$$

Such an equation can be written for every one of the m components of f, i.e., there are $m \times k$ partial derivatives,

$$\partial_j f_i(a) = \lim_{t \to 0} \frac{f_i(a + te_j) - f_i(a)}{t} \qquad i = 1, \dots, m \quad j = 1, \dots, k.$$

Comment: The more common notation for the partial derivative in the j-th direction is $\frac{\partial f}{\partial x_j}$. I prefer not to use it for the same reasons that I prefer f' over $\frac{df}{dx}$ in the univariate case.

Example: Consider the case where $A = \mathbb{R}^k$ and f is a linear transformation, $f(a) = Ta$, with $T \in \operatorname{Hom}(\mathbb{R}^k, \mathbb{R}^m)$. Such a mapping is always differentiable as

$$\frac{f(a+x) - f(a) - Tx}{\|x\|_k} = 0.$$

By definition, $(Df)_a = T$ for all $a \in \mathbb{R}^k$. That is

$$f(a) = Ta \qquad \text{implies} \qquad (Df)_a(x) = Tx \quad \forall a \in A.$$

This generalizes the identity $(cx)' = c$ for $f : \mathbb{R} \to \mathbb{R}$, $f(x) = cx$. ▲▲▲

Example: If f is a constant function, $f = c \in \mathbb{R}^m$, then $(Df)_a$ is the zero transformation, $(Df)_a : \mathbb{R}^k \ni x \mapsto 0 \in \mathbb{R}^m$. ▲▲▲

As we have already pointed out, if $f : \mathbb{R}^k \supset A \to \mathbb{R}^m$ is differentiable, then for every $a \in A$, $(Df)_a \in \operatorname{Hom}(\mathbb{R}^k, \mathbb{R}^m)$, i.e.,

$$Df : A \to \operatorname{Hom}(\mathbb{R}^k, \mathbb{R}^m).$$

Since $\operatorname{Hom}(\mathbb{R}^k, \mathbb{R}^m)$ is a normed space, we have a notion of continuity of Df.

*Definition 3.6 Let $A \subset \mathbb{R}^k$ be an open set and let $f : A \to \mathbb{R}^m$ be differentiable in A. f is said to be **continuously-differentiable** (גזירה ברציפות) at a if the mapping*

$$Df : A \to \operatorname{Hom}(\mathbb{R}^k, \mathbb{R}^m)$$

is continuous at a.

That is, Df is continuous at $a \in A$ if for every converging sequence in A, $a_n \to a$,

$$\lim_{n\to\infty} (Df)_{a_n} = (Df)_a \qquad \text{in } \mathrm{Hom}(\mathbb{R}^k, \mathbb{R}^m),$$

which amounts to

$$\lim_{n\to\infty} \|(Df)_{a_n} - (Df)_a\|_{k,m} = 0.$$

This further means that for all $\mathbb{R}^k \ni x \neq 0$,

$$\begin{aligned}\lim_{n\to\infty} \|(Df)_{a_n}(x) - (Df)_a(x)\|_m &= \lim_{n\to\infty} \|((Df)_{a_n} - (Df)_a)(x)\|_m \\ &\leq \lim_{n\to\infty} \|(Df)_{a_n} - (Df)_a\|_{k,m}\|x\|_k = 0.\end{aligned}$$

Proposition 3.7 Let $f : A \subset \mathbb{R}^k \to \mathbb{R}^m$ be differentiable at a point $a \in A$. Then, all the $m \times k$ partial derivatives of f at a exists and the operator $(Df)_a$ has the matrix representation

$$(Df)_a = \begin{pmatrix} \partial_1 f_1(a) & \dots & \partial_k f_1(a) \\ \vdots & \ddots & \vdots \\ \partial_1 f_m(a) & \dots & \partial_k f_m(a) \end{pmatrix}.$$

Proof: By definition,

$$(Df)_a(e_j) = \lim_{t\to 0} \frac{f(a + te_j) - f(a)}{t}.$$

Both sides are vectors in $\mathbb{R}^m$. Taking the inner product with $e_i \in \mathbb{R}^m$, and noting that limits commute with inner products, we obtain

$$\begin{aligned}(e_i, (Df)_a(e_j)) &= \left(e_i, \lim_{t\to 0} \frac{f(a + te_j) - f(a)}{t}\right) \\ &= \lim_{t\to 0} \frac{(e_i, f(a + te_j) - f(a))}{t} \\ &= \lim_{t\to 0} \frac{f_i(a + te_j) - f_i(a)}{t} \\ &= \partial_j f_i(a).\end{aligned}$$

■

Example: Let $f : \mathbb{R}^2 \to \mathbb{R}^2$ be given by

$$f(a) = \begin{pmatrix} a_1^2 + a_2 \\ a_1 + a_2^2 \end{pmatrix}.$$

First note that

$$\begin{aligned} \partial_1 f(a) &= \lim_{t\to 0} \frac{f\begin{pmatrix} a_1 + t \\ a_2 \end{pmatrix} - f\begin{pmatrix} a_1 \\ a_2 \end{pmatrix}}{t} = \lim_{t\to 0} \frac{\begin{pmatrix} (a_1+t)^2 + a_2 \\ a_1 + t + a_2^2 \end{pmatrix} - \begin{pmatrix} a_1^2 + a_2 \\ a_1 + a_2^2 \end{pmatrix}}{t} \\ &= \lim_{t\to 0} \frac{\begin{pmatrix} 2ta_1 \\ t \end{pmatrix}}{t} = \begin{pmatrix} 2a_1 \\ 1 \end{pmatrix}, \end{aligned}$$

and

$$\begin{aligned} \partial_2 f(a) &= \lim_{t\to 0} \frac{f\begin{pmatrix} a_1 \\ a_2 + t \end{pmatrix} - f\begin{pmatrix} a_1 \\ a_2 \end{pmatrix}}{t} = \lim_{t\to 0} \frac{\begin{pmatrix} a_1^2 + a_2 + t \\ a_1 + (a_2+t)^2 \end{pmatrix} - \begin{pmatrix} a_1^2 + a_2 \\ a_1 + a_2^2 \end{pmatrix}}{t} \\ &= \lim_{t\to 0} \frac{\begin{pmatrix} t \\ 2ta_2 \end{pmatrix}}{t} = \begin{pmatrix} 1 \\ 2a_2 \end{pmatrix}. \end{aligned}$$

If f is differentiable, then its derivative at $a = (a_1, a_2)^T$ is represented by the 2-by-2 matrix

$$(Df)_a = \begin{pmatrix} 2a_1 & 1 \\ 1 & 2a_2 \end{pmatrix}.$$

To verify whether it is differentiable, we need to check whether

$$\lim_{x\to 0} \frac{1}{\sqrt{x_1^2 + x_2^2}} \left(\begin{pmatrix} (a_1+x_1)^2 + (a_2+x_2) \\ (a_1+x_1) + (a_2+x_2)^2 \end{pmatrix} - \begin{pmatrix} a_1^2 + a_2 \\ a_1 + a_2^2 \end{pmatrix} - \begin{pmatrix} 2a_1 & 1 \\ 1 & 2a_2 \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \end{pmatrix} \right) = 0.$$

That is

$$\lim_{x\to 0} \frac{1}{\sqrt{x_1^2 + x_2^2}} \begin{pmatrix} x_1^2 \\ x_2^2 \end{pmatrix} = 0,$$

which is indeed the case.

Now, what would be the directional derivative of f at a in the direction $(2,5)^T$? By definition it is the vector

$$\lim_{t\to 0} \frac{f\begin{pmatrix} a_1+2t \\ a_2+5t \end{pmatrix} - f\begin{pmatrix} a_1 \\ a_2 \end{pmatrix}}{h} = \lim_{t\to 0} \frac{\begin{pmatrix} (a_1+2t)^2+(a_2+5t) \\ (a_1+2t)+(a_2+5t)^2 \end{pmatrix} - \begin{pmatrix} a_1^2+a_2 \\ a_1+a_2^2 \end{pmatrix}}{t}.$$

However, we don't need to recalculate a limit, Since f is differentiable, this directional derivative equals

$$(Df)_a\begin{pmatrix} 2 \\ 5 \end{pmatrix} = 2\,(Df)_a\begin{pmatrix} 1 \\ 0 \end{pmatrix} + 5\,(Df)_a\begin{pmatrix} 0 \\ 1 \end{pmatrix} = 2\,\partial_1 f(a) + 5\,\partial_2 f(a) = \begin{pmatrix} 4a_1 & 5 \\ 2 & 10a_2 \end{pmatrix}.$$

▲▲▲

Example: Consider the function $f : \mathbb{R}^2 \to \mathbb{R}$,

$$f(a) = \sqrt{a_1^2 + a_2^2}.$$

This function is not differentiable at 0 because its partial derivatives do not exist. For example,

$$\lim_{t\to 0} \frac{f(0+te_1) - f(0)}{t} = \lim_{t\to 0} \frac{|t|}{t}$$

does not exist. ▲▲▲

Example: For $k = 1$, f can be thought of as a collection of m univariate functions. That is, there is only one "direction" along which f can be differentiated. For $f : \mathbb{R} \supset A \to \mathbb{R}^m$ we have $(Df)_a \in \mathrm{Hom}(\mathbb{R}, \mathbb{R}^m)$; a linear map $\mathbb{R} \to \mathbb{R}^m$ acts on numbers by multiplication by a vector in $\mathbb{R}^m$. In other words, $\mathrm{Hom}(\mathbb{R}, \mathbb{R}^m)$ is *naturally* isomorphic to $\mathbb{R}^m$, by its action on the "vector" $1 \in \mathbb{R}$.

By definition, for $a \in \mathbb{R}$,

$$(Df)_a(1) = \lim_{t\to 0} \frac{f(a+t) - f(a)}{t},$$

and by linearity, for $x \in \mathbb{R}$,

$$(Df)_a(x) = (Df)_a(1)\,x.$$

That is, the derivative of f at a is represented by a column vector whose entries are

$$(Df)_a(1) = \lim_{t\to 0}\frac{1}{h}\left(\begin{pmatrix} f_1(a+t)\\ \vdots \\ f_m(a+t)\end{pmatrix} - \begin{pmatrix} f_1(a)\\ \vdots \\ f_m(a)\end{pmatrix}\right) = \begin{pmatrix} f_1'(a)\\ \vdots \\ f_m'(a)\end{pmatrix}.$$

For functions defined on subsets of $\mathbb{R}$ we often use the notation $f'(a)$ for their derivative at $a \in \mathbb{R}$. ▲▲▲

Example: We next consider the case where the range of the function is one-dimensional. For $f : \mathbb{R}^k \supset A \to \mathbb{R}$ and $a \in A$,

$$(Df)_a \in \operatorname{Hom}(\mathbb{R}^k, \mathbb{R}),$$

i.e., the derivative of f at a is represented by a row vector whose entries are

$$(\partial_1 f(a), \dots, \partial_k f(a)).$$

The space $\operatorname{Hom}(\mathbb{R}^k, \mathbb{R})$ is k-dimensional; as such, it is isomorphic to $\mathbb{R}^k$, however, there is no *natural* isomorphism unless some structure is incorporated beyond the vector space structure. If $\mathbb{R}^k$ is endowed with an inner-product, then one obtains an isomorphism between $\mathbb{R}^k$ to $\operatorname{Hom}(\mathbb{R}^k, \mathbb{R})$. Define

$$\iota : \mathbb{R}^k \to \operatorname{Hom}(\mathbb{R}^k, \mathbb{R})$$

via,

$$\iota(v)(x) = (v, x).$$

Clear, for every $v \in \mathbb{R}^k$, $\iota(v)$ is an element of $\operatorname{Hom}(\mathbb{R}^k, \mathbb{R})$. To show that it is an isomorphism, it suffice to show that it is one-to-one, i.e., that it has a trivial kernel. Indeed, if there exists a v such that $\iota(v)$ is the zero element of $\operatorname{Hom}(\mathbb{R}^k, \mathbb{R})$, then for every $x \in \mathbb{R}^k$,

$$0 = \iota(v)(x) = (v, x),$$

and in particular $(v, v) = 0$, from which we deduce that $v = 0$. In fact, the mapping ι is even an isometry: for every $v \in \mathbb{R}^k$,

$$\|\iota(v)\|_{k,1} = \max_{\|x\|_k = 1} |\iota(v)(x)|_1 = \max_{\|x\|_k = 1} |(v, x)|.$$

By the Cauchy-Schwarz inequality,

$$\|\iota(v)\|_{k,1} \le \max_{\|x\|_k = 1} \|v\|_k \|x\|_k = \|v\|_k.$$

The maximum is attained by taking $x = v/\|v\|_k$.

For the particular case where the inner-product is Euclidean,

$$\iota(v)(x) = \sum_{i=1}^k v_i x_i,$$

i..e, $\iota(v)$ is a row vector whose elements are the elements of the column vector v. While the isomorphism between column vectors and row vectors seems the "most natural" it is not; it reflects a particular choice of an inner-product. In the case where the norm on $\mathbb{R}^k$ is obtained from an inner-product, then there is a natural inner-product, however recall that not every norm is induced by an inner-product. Our choice has been to work with the Euclidean norm, hence we use the Euclidean inner-product to obtain an isomorphism from $\mathbb{R}^k$ to $\mathrm{Hom}(\mathbb{R}^k, \mathbb{R})$.

Definition 3.8 Let $f : \mathbb{R}^k \supset A \to \mathbb{R}$ and $a \in A$. The ***gradient*** *(גרדיאנט) of f at a is a vector*

$$\nabla f(a) \in \mathbb{R}^k,$$

defined by

$$(\nabla f(a), x) = (Df)_a(x) \qquad \forall x \in \mathbb{R}^k.$$

That is,

$$\nabla f(a) = \begin{pmatrix} \partial_1 f(a) \\ \vdots \\ \partial_k f(a) \end{pmatrix}.$$

The gradient of f at a, $\nabla f(a) \in \mathbb{R}^k$, has a geometric interpretation. Consider $f : A \to \mathbb{R}$ in the vicinity of $a \in A$. One may ask in which direction is the rate of change of f maximal? That is, which among all unit vectors $\hat{x} \in \mathbb{R}^k$ maximizes the directional derivative.

$$\lim_{t \to 0} \frac{f(a + t\hat{x}) - f(a)}{t}.$$

By (3.2), we are looking for a maximizer of

$$(Df)_a(\hat{x}) = (\nabla f(a), \hat{x}).$$

The unit vector along which f changes the fastest is parallel to the gradient,

$$\hat{x} = \frac{\nabla f(a)}{\|\nabla f(a)\|_k}.$$

Thus, the gradient of f is a vector pointing in the direction of maximal growth rate of f, and whose magnitude is the rate of change of f along that direction.

Moreover, consider the space

$$M_f^{\perp}(a) \stackrel{\text{def}}{=} \left\{x \in \mathbb{R}^k : \ (Df)_a(x) = 0\right\} = \left\{x \in \mathbb{R}^k : \ x \perp \nabla f(a) = 0\right\}.$$

This is a $(k-1)$-dimensional subspace of $\mathbb{R}^k$ that spans all the directions along which f does not change, to first order in the displacement. This space spans the plane tangent to the **level sets** of f at the point a. ▲▲▲

We next address the following question: does the existence of all the $m \times k$ partial derivatives ensure the differentiability of a function? We approach this question in two steps. We first show that differentiability implies the differentiability of all the components (which shouldn't come as a surprise since limits in $\mathbb{R}^m$ coincide with component-wise limits).

Proposition 3.9 Let $A \subset \mathbb{R}^k$ be an open set and $f : A \to \mathbb{R}^m$; we denote its components by $f_1, \dots, f_m$, which are all functions $A \to \mathbb{R}$. Then f is differentiable at $a \in A$ if and only if all its components $f_j : A \to \mathbb{R}$ are differentiable at a.

Proof: The function f is differentiable at a if (and only if) there exists a $T \in \mathrm{Hom}(\mathbb{R}^k, \mathbb{R}^m)$, such that

$$\lim_{x \to 0} \frac{f(a+x) - f(a) - T(x)}{\|x\|_k} = 0.$$

Since convergence of the norm amounts to the convergence of each component, f is differentiable at a if and only if for every $j = 1, \dots, m$,

$$\lim_{x \to 0} \frac{f_j(a+x) - f_j(a) - (T(x))_j}{\|x\|_k} = 0.$$

Define $S_j \in \mathrm{Hom}(\mathbb{R}^k, \mathbb{R})$ by

$$S_j(x) = \sum_{i=1}^{k} T_{ji} x_i = (T(x))_j.$$

We conclude that f is differentiable at a if and only if each of each components f_j is differentiable at a. ■

Having reduced the question of differentiability to that of the differentiability of scalar-valued functions, we prove the following:

Theorem 3.10 *Let* $A \subset \mathbb{R}^k$ *be an open set and let* $f : A \to \mathbb{R}$*. If all the k partial derivatives of* f *exist in a neighborhood of* $a \in \mathbb{R}^k$ *and are continuous at a, then* f *is differentiable at a.*

Proof: We need to show that there exists a linear operator $T \in \mathrm{Hom}(\mathbb{R}^k, \mathbb{R})$ such that

$$\lim_{x\to 0} \frac{f(a+x) - f(a) - Tx}{\|x\|_k} = 0.$$

We know what T, <u>if</u> it exists, should be; it is the row vector

$$T = (\partial_1 f(a), \dots, \partial_k f(a)).$$

That is, we need to show that

$$\lim_{x\to 0} \frac{1}{\|x\|_k} \left(f(a+x) - f(a) - \sum_{j=1}^k \partial_j f(a) x_j \right) = 0.$$

Since f has continuous partial derivatives we can invoke the mean-value theorem along each component of its arguments. We first write the displacement of f in the form of a telescoping sum,

$$\begin{aligned} f(a+x) - f(a) &= f(a + x_1 e_1) - f(a) \\ &+ f(a + x_1 e_1 + x_2 e_2) - f(a + x_1 e_1) \\ &+ \dots \\ &+ f(a+x) - f(a + x_1 e_1 + x_2 e_2 + \cdots + x_{k-1} e_{k-1}). \end{aligned}$$

That is, —33h(2019)—

$$f(a+x) - f(a) = \sum_{j=1}^k \left(f\left(a + \sum_{i=1}^j x_i e_i\right) - f\left(a + \sum_{i=1}^{j-1} x_i e_i\right)\right).$$

Each term in this sum consists of a variation of f upon the displacement of its argument along a single axis. Consider the j-th term; define the function $g : [0, x_k] \to \mathbb{R}$,

$$g(t) = f\left(a + \sum_{i=1}^{j-1} x_i e_i + t e_j\right),$$

so that the j-th term equals $g(x_k) - g(0)$. Furthermore,

$$g'(t) = \lim_{h\to 0} \frac{g(t+h)-g(t)}{h} = \lim_{h\to 0} \frac{f\left(a+\sum_{i=1}^{j-1} x_i e_i + t e_j + h e_j\right) - f\left(a+\sum_{i=1}^{j-1} x_i e_i + t e_j\right)}{h}$$

By the definition of the directional derivative,

$$g'(t) = \partial_j f\left(a + \sum_{i=1}^{j-1} x_i e_i + t e_j\right).$$

By the (one-dimensional!) mean-value theorem, there exists a number $\theta_j \in (0,1)$, such that

$$g(x_j) - g(0) = g'(\theta_j x_k) x_k.$$

Putting it all together,

$$f(a+x) - f(a) = \sum_{j=1}^{k} \partial_j f\left(a + \sum_{i=1}^{j-1} x_i e_i + \theta_j x_j e_j\right).$$

Thus,

$$\frac{1}{\|x\|_k}\left(f(a+x) - f(a) - \sum_{j=1}^{k} \partial_j f(a) x_j\right) = \sum_{j=1}^{k}\left[\partial_j f\left(a + \sum_{i=1}^{j-1} x_i e_i + \theta_j x_j e_j\right) - \partial_j f(a)\right]\frac{x_j}{\|x\|_k}.$$

Letting $x \to 0$, the right-hand side vanishes by the continuity of the partial derivatives. ■

Example: A "counter example" is when f has all its $m\times k$ derivatives at a point, but it is nevertheless not differentiable. This can happen only if the partial derivatives are not continuous. Let $f : \mathbb{R}^2 \to \mathbb{R}$ be given by

$$f(a) = \begin{cases} \dfrac{a_1^3}{a_1^2 + a_2^2} & a \neq 0 \\ 0 & a = 0. \end{cases}.$$

This function has partial derivatives at zero,

$$\partial_1 f(0) = \lim_{t\to 0} \frac{f(0+te_1) - f(0)}{t} = \lim_{t\to 0} \frac{t^3/t^2}{t} = 1$$

$$\partial_2 f(0) = \lim_{t\to 0} \frac{f(0+te_2) - f(0)}{t} = 0.$$

On the other hand, this function is not differentiable at 0. If it were, we would have

$$\lim_{x\to 0} \frac{1}{\sqrt{x_1^2+x_2^2}} \left(\frac{x_1^3}{x_1^2+x_2^2} - \begin{pmatrix} 1 & 0 \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \end{pmatrix} \right) = 0,$$

i.e.,

$$\lim_{x\to 0} \frac{-x_1 x_2^2}{(x_1^2+x_2^2)^{3/2}} = 0.$$

This is not the case as seen when taking $x_1 = x_2 \to 0$.

▲▲▲

The existence of partial derivatives is a necessary by not sufficient condition for differentiability. The continuity of the partial derivatives ensures differentiability. In fact, it ensures that the derivative is also continuous. Conversely, the continuity of the derivative implies the continuity of the partial derivatives. This leads us to the following corollary:

Corollary 3.11 Let $A \in \mathbb{R}^k$ be an open set. Then, $f : A \to \mathbb{R}^m$ is continuously differentiable at $a \in A$ if and only if its $m \times k$ partial derivatives exist and are continuous at a.

Definition 3.12 Let $A \in \mathbb{R}^k$ be an open set. We denote by $C^1(A;\mathbb{R}^m)$ the set of continuously-differentiable functions from A into $\mathbb{R}^m$.

✎ *Exercise 3.2* Let $f : \mathbb{R}^k \to \mathbb{R}$ such that $(Df)_x = 0$ on an open connected set $U \subset \mathbb{R}^k$. Prove that f is constant on U.

✎ *Exercise 3.3* Let $A \subset \mathbb{R}^k$ be an open set and let $f : A \to \mathbb{R}^m$. Show that if f is differentiable in A then it is also continuous, in fact, Lipschitz continuous.

We have thus defined a notion of differentiability that generalizes the "old" concept of differentiability for functions $\mathbb{R} \to \mathbb{R}$. Moreover, we have a formula for calculating the derivative of a function $\mathbb{R}^k \to \mathbb{R}^m$, which involves the calculation of $m \times k$ partial derivatives. The next stage is to do "useful" things with this new concept, and in particular, generalize results known for real-valued functions.

Proposition 3.13 (The chain rule) Let $A \subseteq \mathbb{R}^k$ and $B \subseteq \mathbb{R}^m$ be open domains and let $f : A \to B$ and $g : B \to \mathbb{R}^n$. Assume that f is differentiable at $a \in A$ and that g is differentiable at $f(a) \in B$. Then, $g \circ f : A \to \mathbb{R}^n$ is differentiable at a and its derivative is given by the chain rule:

$$\underbrace{(D(g \circ f))_a}_{\mathbb{R}^k \to \mathbb{R}^n} = \underbrace{(Dg)_{f(a)}}_{\mathbb{R}^m \to \mathbb{R}^n} \circ \underbrace{(Df)_a}_{\mathbb{R}^k \to \mathbb{R}^m}.$$

Proof: For sufficiently small displacements $x \in \mathbb{R}^n$ and $y \in \mathbb{R}^m$ we define the functions, $r : \mathbb{R}^k \to \mathbb{R}^m$ and $s : \mathbb{R}^m \to \mathbb{R}^n$ by

$$\begin{aligned} r(x) &= f(a+x) - f(a) - (Df)_a(x) \\ s(y) &= g(f(a)+y) - g(f(a)) - (Dg)_{f(a)}(y). \end{aligned}$$

By the definition of $(Df)_a$ and $(Dg)_{f(a)}$,

$$\lim_{x \to 0} \frac{\|r(x)\|_m}{\|x\|_k} = 0 \qquad \text{and} \qquad \lim_{y \to 0} \frac{\|s(y)\|_n}{\|y\|_m} = 0.$$

By the definition of r,

$$f(a+x) = f(a) + (Df)_a(x) + r(x),$$

and by the definition of s, substituting $y = (Df)_a(x) + r(x)$,

$$\begin{aligned} s((Df)_a(x) + r(x)) = {} & g(f(a) + (Df)_a(x) + r(x)) \\ & - g(f(a)) \\ & - (Dg)_{f(a)}((Df)_a(x) + r(x)), \end{aligned}$$

which we can re-organize as follows,

$$\begin{aligned} g(f(a+x)) - g(f(a)) = {} & (Dg)_{f(a)}((Df)_a(x)) \\ & + (Dg)_{f(a)}(r(x)) \\ & + s((Df)_a(x) + r(x)). \end{aligned}$$

Moving $(Dg)_{f(a)}((Df)_a(x))$ to the right hand side, dividing by $\|x\|_k$, and writing expressions such as $g(f(a))$ as compositions, we obtain

$$\frac{g \circ f(a+x) - g \circ f(a) - (Dg)_{f(a)} \circ (Df)_a(x)}{\|x\|_k} = (Dg)_{f(a)} \frac{r(x)}{\|x\|_k} + \frac{s((Df)_a(x) + r(x))}{\|x\|_k}.$$

Taking norms, and using the triangle inequality,

$$\frac{\|g \circ f(a+x) - g \circ f(a) - (Dg)_{f(a)} \circ (Df)_a(x)\|_n}{\|x\|_k} \le \|(Dg)_{f(a)}\|_{m,n} \frac{\|r(x)\|_m}{\|x\|_k} + \frac{\|s((Df)_a(x) + r(x))\|_n}{\|x\|_k}.$$

We need to prove that the right-hand side tends to zero as $x \to 0$. The first term tends to zero by the defining property of $r(x)$. It remains to prove that

$$\lim_{x \to 0} \underbrace{\frac{\|s((Df)_a(x) + r(x))\|_n}{\|x\|_k}}_{\stackrel{\text{def}}{=} \psi(x)} = 0.$$

Then,

$$\psi(x) = \begin{cases} \dfrac{\|s((Df)_a(x) + r(x))\|_n}{\|(Df)_a(x) + r(x)\|_m} \dfrac{\|(Df)_a(x) + r(x)\|_m}{\|x\|_k} & (Df)_a(x) + r(x) \neq 0 \\ 0 & \text{otherwise.} \end{cases}$$

By the triangle inequality,

$$\psi(x) \le \begin{cases} \dfrac{\|s((Df)_a(x) + r(x))\|_n}{\|(Df)_a(x) + r(x)\|_m} \left(\|(Df)_a\|_{k,m} + \dfrac{\|r(x)\|_m}{\|x\|_k} \right) & (Df)_a(x) + r(x) \neq 0 \\ 0 & \text{otherwise.} \end{cases}$$

Let now $x \to 0$. Then, $(Df)_a(x) + r(x) \to 0$. By the definition of s, the first factor tends to zero whereas the second term is bounded, hence

$$\lim_{x \to 0} \psi(x) = 0.$$

■

Example: Consider $f : \mathbb{R}^2 \to \mathbb{R}^3$ given by

$$f(a) = \begin{pmatrix} \cos a_1 \cos a_2 \\ \cos a_1 \sin a_2 \\ \sin a_1 \end{pmatrix},$$

and let $g : \mathbb{R}^3 \to \mathbb{R}$ be given by

$$g(b) = b_1^2 + 2b_2^2 + 3b_3^2.$$

The composite function is

$$(g \circ f)(a) = \cos^2 a_1 \cos^2 a_2 + 2\cos a_1^2 \sin a_2^2 + 3\sin^2 a_1.$$

The derivative of f at a is given by

$$(Df)_a = \begin{pmatrix} -\sin a_1 \cos a_2 & -\cos a_1 \sin a_2 \\ -\sin a_1 \sin a_2 & \cos a_1 \cos a_2 \\ \cos a_1 & 0 \end{pmatrix},$$

whereas the derivative of g at b is given by

$$(Dg)_b = (2b_1, 4b_2, 6b_3).$$

The composition of the derivatives gives

$$(Dg)_{f(a)} \circ (Df)_a = (2\cos a_1 \cos a_2, 4\cos a_1 \sin a_2, 6\sin a_1) \begin{pmatrix} -\sin a_1 \cos a_2 & -\cos a_1 \sin a_2 \\ -\sin a_1 \sin a_2 & \cos a_1 \cos a_2 \\ \cos a_1 & 0 \end{pmatrix}.$$

On the other hand,

$$(D(g \circ f))_a = \left(\sin 2a_1 \left(-\cos^2 a_2 - 2\sin a_2^2 + 3\right) \quad -\cos^2 a_1 \sin 2a_2 + 2\cos a_1^2 \sin 2a_2\right).$$

It is easily checked that the two are equal. ▲▲▲

Definition *3.14 Let $f : \mathbb{R}^k \supset A \to \mathbb{R}^k$ be differentiable at a. The **Jacobian** (יעקוביאן) of f at a is defined as*

$$Jf(a) = \det(Df)_a.$$

Comment: A non-zero Jacobian means that the derivative is an invertible transformation. The Jacobian also plays an important role in integration theory, since it quantifies how volumes are transformed under mappings. Specifically,

$$Jf(a) = \lim_{r\to 0} \frac{\mathrm{Vol}(f(\mathcal{B}_r(a)))}{\mathrm{Vol}(\mathcal{B}_r(a))}$$

Example: Consider the transformation

$$f : a \mapsto \begin{pmatrix} a_1 \cos a_2 \\ a_1 \sin a_2 \end{pmatrix}.$$

Its Jacobian is $Jf(a) = a_1$. ▲▲▲

Example: The Jacobian of the identity is 1. ▲▲▲

Proposition 3.15 Let $f, g : \mathbb{R}^k \to \mathbb{R}^k$. Than,

$$J(g \circ f)(a) = Jg(f(a))\, Jf(a).$$

Proof: Immediate from the properties of the determinant, as

$$J(g \circ f)(a) = \det(D(g \circ f))_a = \det\big((Dg)_{f(a)}(Df)_a\big) = \det(Dg)_{f(a)} \det(Df)_a.$$

■

For real-valued function, a "small" derivative implies that the function changes by "little" relative to the displacement of its argument. We expect a similar result to hold for functions between multi-dimensional Euclidean spaces, except that now the magnitude of the derivative should be the *norm* of its derivative.

Lemma 3.16 Let $\gamma : [0,1] \to \mathbb{R}^m$ be a continuously-differentiable path in $\mathbb{R}^m$ such that

$$\sup_{0\le t\le 1} \|(D\gamma)_t\|_{1,m} = M.$$

Then,

$$\|\gamma(1) - \gamma(0)\|_m \le M.$$

(Recall that the derivative $(D\gamma)_t = \gamma'(t)$ is a column vector whose entries are $\gamma_j'(t)$.)

Proof: Define the function, $f:[0,1]\to\mathbb{R}$, given by

$$f(t)=\|\gamma(t)-\gamma(0)\|_m=\left(\sum_{j=1}^k(\gamma_j(t)-\gamma_j(0))^2\right)^{1/2}.$$

Differentiating f^2 (this is a univariate function),

$$\begin{aligned}
f^{2\prime}(t)&=2\sum_{j=1}^k\gamma_j'(t)(\gamma_j(t)-\gamma_j(0))\\
&=2\,(\gamma'(t),\gamma(t)-\gamma(0))\\
&\le 2\|(D\gamma)_t\|_{1,m}\|\gamma(t)-\gamma(0)\|\\
&\le 2M\,f(t),
\end{aligned}$$

where in the passage to the third line we used the Cauchy-Schwarz inequality. It follows that

$$2f(t)f'(t)\le 2M\,f(t),$$

and since $f(t)$ is non-negative, $f'(t)\le M$ for all t. It then follows from the mean-value theorem that $f(1)\le M$, which is the desired result. ■

With this lemma we prove the following theorem:

Theorem 3.17 Let A be an open subset of $\mathbb{R}^k$, $g:A\to\mathbb{R}^m$ *a differentiable function, and* $a,b\in A$ *such that*

$$I=\{tb+(1-t)a:\ 0\le t\le 1\}\subset A$$

(the segment connecting a and b is in the domain). Assume furthermore that

$$\sup_{c\in I}\|(Dg)_c\|_{k,m}=M.$$

Then,

$$\|g(b)-g(a)\|_m\le M\,\|b-a\|_k.$$

Proof: We define the path $\gamma:[0,1]\to A$,

$$\gamma(t)=g(f(t)),\qquad\text{where}\qquad f(t)=tb+(1-t)a.$$

By the chain rule.

$$(D\gamma)_t = (Dg)_{f(t)} \circ (Df)_t = (Dg)_{f(t)}(b-a).$$

For every $t \in [0,1]$,

$$\|(D\gamma)_t\|_{1,m} \le \|(Dg)_{f(t)}\|_{k,m} \|b-a\|_k \le M\,\|b-a\|_k.$$

By Lemma 3.16, $\|\gamma(1)-\gamma(0)\|_m \le M\|b-a\|_k$, which is the desired result. ■ —35h(2019)—

3.2 Higher derivatives

The space $C^1(A;\mathbb{R}^m)$, $A \subset \mathbb{R}^k$ is the space of continuously-differentiable functions from A to $\mathbb{R}^m$, and it coincides with the space of functions that have continuous partial derivatives. For $f \in C^1(A;\mathbb{R}^m)$, its derivative is a mapping

$$Df : A \mapsto \mathrm{Hom}(\mathbb{R}^k, \mathbb{R}^m),$$

i.e., a (generally nonlinear!) mapping between two finite-dimensional normed spaces. Thus, we can define the derivative of Df at a: Df is differentiable at a if there exists a linear operator

$$(D^2 f)_a \in \mathrm{Hom}(\mathbb{R}^k, \mathrm{Hom}(\mathbb{R}^k, \mathbb{R}^m)),$$

such that

$$\lim_{x\to 0} \frac{(Df)_{a+x} - (Df)_a - (D^2 f)_a(x)}{\|x\|_k} = 0,$$

which is an equality between elements in $\mathrm{Hom}(\mathbb{R}^k, \mathbb{R}^m)$. Note that $(D^2 f)_a$ maps an ordered pair of vectors in $\mathbb{R}^k$ into a vector in $\mathbb{R}^m$. Specifically,

$$\begin{aligned}
&(D^2 f)_a \in \mathrm{Hom}(\mathbb{R}^k, \mathrm{Hom}(\mathbb{R}^k, \mathbb{R}^m))\\
&(D^2 f)_a(x) \in \mathrm{Hom}(\mathbb{R}^k, \mathbb{R}^m)\\
&(D^2 f)_a(x)(y) \in \mathbb{R}^m.
\end{aligned}$$

Proposition 3.18 *The second derivative of f at a is* ***bilinear*** *(בילינארית), namely,*

$$(D^2 f)_a(\alpha x_1 + \beta x_2)(y) = \alpha\,(D^2 f)_a(x_1)(y) + \beta\,(D^2 f)_a(x_2)(y),$$

and

$$(D^2 f)_a(x)(\alpha y_1 + \beta y_2) = \alpha\,(D^2 f)_a(x)(y_1) + \beta\,(D^2 f)_a(x)(y_2),$$

Proof: This is an immediate consequence of $x \mapsto (D^2 f)_a(x)$ being linear and $y \mapsto (D^2 f)_a(x)(y)$ being linear. ■

To understand how does the second derivative operate, we use the notion of the directional derivative (3.2) to get

$$(D^2 f)_a(x) = \lim_{t\to 0} \frac{(Df)_{a+tx} - (Df)_a}{t}.$$

Note that the right-hand side is a limit in $\text{Hom}(\mathbb{R}^k, \mathbb{R}^m)$.

Applying both sides on $y \in \mathbb{R}^k$, and using the fact that matrix-vector multiplication is continuous,

$$(D^2 f)_a(x)(y) = \lim_{t\to 0} \frac{(Df)_{a+tx}(y) - (Df)_a(y)}{t}.$$

In particular, for $x = e_i$ and $y = e_j$,

$$\begin{aligned}(D^2 f)_a(e_i)(e_j) &= \lim_{t\to 0} \frac{(Df)_{a+te_i}(e_j) - (Df)_a(e_j)}{t} \\ &= \lim_{t\to 0} \frac{\partial_j f(a+te_i) - \partial_j f(a)}{t} \\ &= \partial_i \partial_j f(a).\end{aligned}$$

Since $(D^2 f)_a$ is bilinear (Proposition 3.18), for

$$x = \sum_{i=1}^k x_i e_i \qquad \text{and} \qquad y = \sum_{j=1}^k y_j e_j,$$

$$(D^2 f)_a(x)(y) = \sum_{i=1}^k \sum_{j=1}^k \partial_i \partial_j f(a) x_i y_j.$$

That is, the second derivative of f at a is an $\mathbb{R}^m$-valued bilinear form represented by the $m \times k \times k$ numbers

$$\partial_i \partial_j f^\ell(a).$$

In a similar way, higher derivatives may be defined. The ℓ-th derivative of f is a map

$$(D^\ell f)_a \in \text{Hom}(\mathbb{R}^k, \text{Hom}(\mathbb{R}^k, \dots, \text{Hom}(\mathbb{R}^k, \mathbb{R}^m))),$$

which can be defined as a ℓ-**multilinear** map from $\mathbb{R}^k$ to $\mathbb{R}^m$. Moreover,

$$(D^\ell f)_a(e_{i_1}) \dots (e_{i_\ell}) = \partial_{i_1} \partial_{i_2} \dots \partial_{i_\ell} f(a).$$

Comment: For $T \in \mathrm{Hom}(\mathbb{R}^k, \mathrm{Hom}(\mathbb{R}^k, \mathbb{R}^m))$ we will write for convenience $T(x,y)$ rather than $T(x)(y)$.

The following lemma will come up handy:

Lemma 3.19 Let $f : \mathbb{R}^k \supset A \to \mathbb{R}^m$ be twice-differentiable at a. Define $g : A \to \mathbb{R}^m$ by

$$g(a) = (Df)_a(y),$$

where $y \in \mathbb{R}^k$ is a fixed vector. Then, g is differentiable at a, and

$$(Dg)_a(z) = (D^2 f)_a(z,y).$$

(Note that g is in general a nonlinear function.)

Proof: By definition,

$$\begin{aligned}(Dg)_a(z) &= \lim_{t\to 0} \frac{g(a+tz) - g(a)}{t} \\ &= \lim_{t\to 0} \frac{(Df)_{a+tz}(y) - (Df)_a(y)}{t} \\ &= \left(\lim_{t\to 0} \frac{(Df)_{a+tz} - (Df)_a}{t}\right)(y) \\ &= (D^2 f)(z)(y).\end{aligned}$$

■

3.3 Multivariate mean-value and Taylor theorems

Proposition 3.20 Let $A \subseteq \mathbb{R}^k$ be an open, path-connected set (i.e., for every $a,b \in A$ there exists a differentiable path connecting a and b). Then $f : A \to \mathbb{R}^m$ is constant if and only if $(Df)_a = 0$ (the zero operator) for all $a \in A$.

Proof: We have already seen that $f = \text{const}$ implies $Df = 0$. Suppose now that $Df = 0$ in A, and let $a,b \in A$. Since A is path-connected, there exists a (continuously-differentiable) path $\varphi : [0,1] \to A$ such that $\varphi(0) = a$ and $\varphi(1) = b$. Define the path $g : [0,1] \to \mathbb{R}^m$ by

$$g(t) = f(\varphi(t)).$$

By the chain rule,

$$g'(t) = (Dg)_t = (Df)_{\varphi(t)} \circ (D\varphi)_t = 0,$$

from which we conclude that all the components of g are constant, i.e., $f(a) = f(b)$. ■

Corollary 3.21 Let $A \subseteq \mathbb{R}^k$ be open and path connected and $f, g \in C^1(A; \mathbb{R}^m)$ such that

$$(Df)_a = (Dg)_a \qquad \text{for all } a \in A.$$

Then there exists a constant $c \in \mathbb{R}^m$ such that

$$f = g + c.$$

Proof: Apply the previous proposition to the function $h(a) = f(a) - g(a)$. ■

We now consider a generalization of the **mean-value theorem** (משפט הערך הממוצע). Recall that for $f \in C^1([a,b]; \mathbb{R})$, there exists a $\theta \in (0,1)$ such that

$$f(b) - f(a) = f'(a + \theta(b-a))(b-a).$$

The question is whether this theorem holds for functions $f : \mathbb{R}^k \to \mathbb{R}^m$, i.e., is it true that for every $a, b \in \mathbb{R}^k$ there exists a $\theta \in (0,1)$ such that

$$f(b) - f(a) = (Df)_{b+\theta(b-a)}(b-a)?$$

In general, this is false. It only holds for real-valued functions.

Example: The following example show that the mean-value theorem does not generally hold when the range is $\mathbb{R}^m$. Consider the function $f : \mathbb{R} \to \mathbb{R}^3$:

$$f : t \mapsto \begin{pmatrix} \cos t \\ \sin t \\ t \end{pmatrix}.$$

This function is differentiable with

$$(Df)_t = \begin{pmatrix} -\sin t \\ \cos t \\ 1 \end{pmatrix}.$$

Now,

$$f(2\pi) - f(0) = \begin{pmatrix} 0 \\ 0 \\ 2\pi \end{pmatrix} \qquad \text{whereas} \qquad (Df)_{2\pi\theta}(2\pi - 0) = 2\pi \begin{pmatrix} -\sin 2\pi\theta \\ \cos 2\pi\theta \\ 1 \end{pmatrix},$$

and there is no value of θ for which the two are equal. ▲▲▲

Theorem 3.22 (mean-value theorem) Let $A \subseteq \mathbb{R}^k$ be open, $a, b \in A$, such that the segment connecting them is in A, and $f \in C^1(A; \mathbb{R})$. Then there exists a $\theta \in (0,1)$ such that

$$f(b) - f(a) = (Df)_{a+\theta(b-a)}(b-a).$$

Proof: The idea is to use the mean-value theorem for univariate function. Consider the function $\varphi : [0,1] \to A$ given by

$$\varphi(t) = a + t(b-a),$$

and the function $g : I \to \mathbb{R}$ given by

$$g(t) = f(\varphi(t)).$$

The function g is differentiable with

$$g'(t) = (Df)_{\varphi(t)} \circ (D\varphi)_t = (Df)_{\varphi(t)}(b-a).$$

By the univariate mean-value theorem there exists a $\theta \in (0,1)$, such that

$$g(1) - g(0) = g'(\theta) = (Df)_{a+\theta(b-a)}(b-a),$$

which concludes the proof. ■

Comment: Since a function $f : \mathbb{R}^k \to \mathbb{R}^m$ is a row of m real-valued functions, each component satisfies the mean-value theorem, but with a different θ, that is there exist $\theta_1, \dots, \theta_m \in (0,1)$ such that

$$f_j(b) - f_j(a) = (Df_j)_{a+\theta_j(b-a)}(b-a).$$

—37h(2019)—

We proceed to show that if $f : \mathbb{R}^k \to \mathbb{R}^m$ has continuous second partial derivatives, i.e., $f \in C^2(\mathbb{R}^k;\mathbb{R}^m)$, then its second derivative is symmetric. In particular, for every $i, j = 1, \dots, k$,

$$\partial_i \partial_j f = \partial_j \partial_i f.$$

Clearly, it is enough to consider real-valued functions.

Theorem 3.23 (Equality of mixed derivatives) Let $A \subset \mathbb{R}^k$ be an open set and let $f \in C^2(A;\mathbb{R})$. Then, $(D^2 f)_a \in \mathrm{Hom}(\mathbb{R}^k, \mathrm{Hom}(\mathbb{R}^k, \mathbb{R}))$ is a symmetric bilinear operator. That is, for every $x, y \in \mathbb{R}^k$,

$$(D^2 f)_a(x,y) = (D^2 f)_a(y,x).$$

Proof: Fixing $a \in A$ and $x, y \in \mathbb{R}^k$, consider the expression

$$\begin{aligned} I &= (f(a+tx+sy) - f(a+tx)) - (f(a+sy) - f(a)) \\ &= (f(a+tx+sy) - f(a+sy)) - (f(a+tx) - f(a)). \end{aligned}$$

For every s we define a function $g_s : A \to \mathbb{R}$,

$$g_s(z) = f(z+sy) - f(z).$$

Likewise, for every t we define a function $h_t : A \to \mathbb{R}$,

$$h_t(z) = f(z+tx) - f(z).$$

Then,

$$I = g_s(a+tx) - g_s(a) = h_t(a+sy) - h_t(a).$$

By the the mean-value theorem, there exist a family of constants $\theta_s \in (0,1)$ and a family of constants $\tau_t \in (0,1)$, such that

$$I = (Dg_s)_{a+\theta_s tx}(tx) = (Dh_t)_{a+\tau_t sy}(sy).$$

By the chain rule,

$$\begin{aligned}(Dg_s)_{a+\theta_s tx}(tx) &= t\left((Df)_{a+\theta_s tx+sy}(x) - (Df)_{a+\theta_s tx}(x)\right)\\ (Dh_t)_{a+\tau_t sy}(sy) &= s\left((Df)_{a+\tau_t sy+tx}(y) - (Df)_{a+\tau_t sy}(y)\right).\end{aligned}$$

We are going to apply the mean-value theorem a second time. Define the functions $p, r : \mathbb{R} \to \mathbb{R}$ (both depend on s and t),

$$\begin{aligned}p(u) &= (Df)_{a+\theta_s tx+uy}(x) - (Df)_{a+\theta_s tx}(x)\\ r(u) &= (Df)_{a+\tau_t sy+ux}(y) - (Df)_{a+\tau_t sy}(y).\end{aligned}$$

Then, there exist $\xi_{s,t}, \zeta_{s,t} \in (0,1)$, such that

$$\begin{aligned}(Dg_s)_{a+\theta_s tx}(tx) &= t\left(p(s) - p(0)\right) = ts\, p'(\xi_{s,t}s)\\ (Dh_t)_{a+\tau_t sy}(sy) &= s\left(r(t) - r(0)\right) = ts\, r'(\zeta_{s,t}t).\end{aligned}$$

However,

$$\begin{aligned}p'(\xi_{s,t}s) &= (D^2 f)_{a+\theta_s tx+\xi_{s,t}sy}(y,x)\\ r'(\zeta_{s,t}t) &= (D^2 f)_{a+\tau_t sy+\zeta_{s,t}tx}(x,y).\end{aligned}$$

Putting it all together,

$$\frac{I}{st} = (D^2 f)_{a+\theta_s tx+\xi_{s,t}sy}(y,x) = (D^2 f)_{a+\tau_t sy+\zeta_{s,t}tx}(x,y).$$

Set now $s = t$ and let $t \to 0$. By the continuity of D^f we obtain that

$$(D^2 f)(x,y) = (D^2 f)_a(y,x).$$

■

Finally, we prove a multivariate version of **Taylor's theorem** (משפט טיילור הרב ממדי). Here again, we consider real-valued functions. The result applies for $\mathbb{R}^m$-valued functions in a row-by-row fashion.

Lemma 3.24 Let $f : \mathbb{R}^k \supset A \to \mathbb{R}$ be p-times differentiable, let $a \in A$ and $x \in \mathbb{R}^k$ such that the segment connecting a and $a + x$ is in A, and let $g : [0,1] \to \mathbb{R}$ be given by

$$g(t) = f(a + tx).$$

Then,

$$\begin{aligned} g'(t) &= (Df)_{a+tx}(x) \\ g''(t) &= (D^2 f)_{a+tx}(x, x) \\ g'''(t) &= (D^3 f)_{a+tx}(x, x, x) \\ \vdots &= \vdots \\ g^{(p)}(t) &= (D^p f)_{a+tx}(x, x, \ldots, x). \end{aligned}$$

Proof: Define $\varphi : \mathbb{R} \to \mathbb{R}^k$ by

$$\varphi(t) = a + tx.$$

Then $g = f \circ \varphi$, and by the chain rule

$$g'(t) = (Df)_{\varphi(t)} \circ \varphi'(t) = (Df)_{a+xt}(x).$$

For the second derivative, we've already proven such a result, but the simplest would be to re-derive it,

$$\begin{aligned} g''(t) &= \lim_{h\to 0} \frac{(Df)_{a+xt+xh}(x) - (Df)_{a+xt}(x)}{h} \\ &= \left(\lim_{h\to 0} \frac{(Df)_{a+xt+xh} - (Df)_{a+xt}}{h} \right)(x) \\ &= (D^2 f)_{a+tx}(x)(x). \end{aligned}$$

Likewise,

$$\begin{aligned} g'''(t) &= \lim_{h\to 0} \frac{(D^2 f)_{a+xt+xh}(x, x) - (D^2 f)_{a+xt}(x, x)}{h} \\ &= \left(\lim_{h\to 0} \frac{(D^2 f)_{a+xt+xh} - (D^2 f)_{a+xt}}{h} \right)(x, x) \\ &= (D^3 f)_{a+tx}(x)(x, x), \end{aligned}$$

and so on. ■

Theorem 3.25 (Multivariate Taylor) *Let* $A \subset \mathbb{R}^k$ *be an open set,* $f \in C^{p+1}(A;\mathbb{R})$*,* $a \in A$*,* $a + x \in A$ *(as well as the segment connecting the two points). Then there exists a* $\theta \in (0,1)$ *such that*

$$f(a+x) = f(a) + (Df)_a(x) \; + \frac{1}{2}(D^2 f)_a(x,x) + \frac{1}{3!}(D^3 f)_a(x,x,x) + \cdots + R_p(x),$$

where

$$R_p(x) = \frac{1}{(k+1)!}(D^{p+1} f)_{a+\theta x}(x,x,\dots,x).$$

✎ *Exercise 3.4* Convince yourself that the multivariate Taylor theorem reduces to the univariate Taylor theorem for $k = 1$.

Proof: Here again, we base the proof on the univariate version of Taylor's theorem. Define $g : (-1,1) \to \mathbb{R}$ by

$$g(t) = f(a + tx),$$

and expand $g(1)$ about $t = 0$. By the previous lemma,

$$\begin{aligned} g'(t) &= (Df)_{a+tx}(x) \\ g''(t) &= (D^2 f)_{a+tx}(x,x) \\ g'''(t) &= (D^3 f)_{a+tx}(x,x,x) \\ \vdots &= \vdots \end{aligned}$$

Finally,

$$g(1) = g(0) + g'(0) + \frac{1}{2}g''(0) + \cdots + \frac{1}{(p+1)!}g^{(p+1)}(\theta).$$

■

Comment: For functions $f : \mathbb{R}^k \to \mathbb{R}^m$ each component f_j satisfies Taylor's theorem, but every component will have its own θ, as we have already seen for the mean-value theorem.

Example: Let $f : \mathbb{R}^2 \to \mathbb{R}$ by given by

$$f(a) = e^{a_1 + 2a_2}.$$

Then, $f(0) = 1$,

$$(Df)_0(x) = \partial_1 f(0) x_1 + \partial_2 f(0)\, x_2 = x_1 + 2x_2,$$

$$(D^2 f)_0(x,x) = \partial_1\partial_1 f(0) x_1^2 + 2\partial_1\partial_2 f(0)\, x_1 x_2 + \partial_2\partial_2 f(0) x_2^2 = x_1^2 + 2x_1 x_2 + 4x_2^2,$$

so that the second-order Taylor polynomial of f at zero is

$$P_2 f(x) = 1 + (x_1 + 2x_2) + \frac{1}{2}(x_1^2 + 4x_1 x_2 + 4x_2^2).$$

▲ ▲ ▲

Example: Taylor's theorem is above all an approximation method. It states that every (smooth) function can, to zeroth order, be approximated by a constant, to first order by a linear function, and so on. First-order approximation gives rise to the **multivariate Newton method**. Suppose we have a function $f : \mathbb{R}^k \to \mathbb{R}^k$ whose root(s) we want to compute. That is, we are looking for $r \in \mathbb{R}^k$ for which $\mathbb{R}^k \ni f(r) = 0$. Suppose we have an initial guess a_0 close enough to the desired r. By Taylor's theorem,

$$f(a) = f(a_0) + (Df)_{a_0}(a - a_0) + \text{remainder}.$$

Retaining only the linear approximation, and provided that $(Df)_{a_0}$ is not singular, an approximation a_1 for r is obtained by

$$f(a_0) + (Df)_{a_0}(a_1 - a_0) = 0 \qquad \Rightarrow \qquad a_1 = a_0 - [(Df)_{a_0}]^{-1} f(a_0).$$

This suggests the following **iterative method**:

$$a_{n+1} = a_n - [(Df)_{a_n}]^{-1} f(a_n),$$

known as the multivariate Newton method. The hope is that $a_n \to r$. This sequence does not always converge, but when it does, it does it extremely fast! ▲ ▲ ▲

✎ *Exercise 3.5* Assume that $f \in C^1(\mathbb{R}^k;\mathbb{R}^k)$, r a root of f and $Jf(r) \neq 0$. Show that there exists a neighborhood of r in which Newton's method converges to r.

3.4 Minima and maxima

We devote this short section to the identification of extrema of multivariate real-valued functions. Throughout this section we consider functions $f \in C^2(A;\mathbb{R})$, where $A \subseteq \mathbb{R}^k$ is open.

Definition 3.26 An interior point $a \in A$ is called a ***local maximum*** *(מקסימום מקומי) of f if there exists an open neighborhood of a, U, such that*

$$f(a) = \sup_{b \in U} f(b).$$

A local minimum is defined similarly.

The following proposition generalizes the well-known property of extremal points for univariate functions.

Proposition 3.27 If $a \in A$ is a local maximum of f then $(Df)_a = 0$.

Comment: A point where the derivative vanishes is called a **critical point** (נקודה קריטית).

Proof: Let U be a neighborhood of a in which $f(a)$ is maximal. By definition, there exists an open ball $\mathcal{B}_r(a) \subset U$ in which $f(a)$ is maximal. To show that $(Df)_a = 0$ we need to show that $(Df)_a(x) = 0$ for every vector x. Let $\hat{x}$ be an arbitrary unit vector and consider the function $g : (-r, r) \to \mathbb{R}$,

$$g(t) = f(a + t\hat{x}).$$

The function g is differentiable and reaches a local maximum at $t = 0$, hence

$$0 = g'(0) = (Df)_a(\hat{x}).$$

Since $(Df)_a(\hat{x}) = 0$ for every unit vector $\hat{x}$, $(Df)_a(x) = 0$ for every $x \in \mathbb{R}^k$. ■

In the univariate case, a vanishing first derivative does not guarantee a local extremum. A sufficient (but not necessary!) condition for a univariate functions f to reach a local maximum at a is that its first derivative vanishes and its second derivative is negative. We expect a similar condition to hold for multivariate functions, except that the second derivative is a symmetric bilinear form,

$$(D^2 f)_a(x, y) = (D^2 f)_a(y, x) = \sum_{i,j=1}^{n} \partial_j \partial_i f(a)\, x_i y_j.$$

Definition 3.28 A bilinear operator $T \in \mathrm{Hom}(\mathbb{R}^k, \mathrm{Hom}(\mathbb{R}^k, \mathbb{R}))$ is called ***positive-definite*** *(חיובית בהחלט) if*

$$T(x,x) \ge 0 \qquad \text{for all } x \in \mathbb{R}^k,$$

with equality only if $x = 0$. T is called ***positive-semidefinite*** *(חיובית למחצה) if equality may also hold for $x \ne 0$. Negative-definiteness and negative semi-definiteness are defined similarly.*

Note that $(D^2 f)_a$ is positive-definite if the matrix H whose entries are

$$h_{ij} = \partial_i \partial_j f(a)$$

satisfies

$$(x, Hx) > 0$$

for all $x \ne 0$.

Lemma 3.29 A bilinear form $T \in \mathrm{Hom}(\mathbb{R}^k, \mathrm{Hom}(\mathbb{R}^k, \mathbb{R}))$ is positive definite if and only if

$$T(\hat{x}, \hat{x}) > 0$$

for every unit vector $\hat{x} \in \mathbb{R}^k$.

Proof: The “only-if” part is trivial. The “if” part follows from the bilinearity of A. If

$$T(\hat{x}, \hat{x}) > 0$$

for every unit vector $\hat{x}$, then for every non-zero vector $x \in \mathbb{R}^k$,

$$T(x,x) = T\left(\frac{x}{\|x\|_k}, \frac{x}{\|x\|_k}\right) \|x\|_k^2 > 0.$$

■

Lemma 3.30 *Let $T \in \operatorname{Hom}(\mathbb{R}^k, \operatorname{Hom}(\mathbb{R}^k,\mathbb{R}))$ be positive definite. Then, there exists an $\alpha > 0$, such that*

$$T(\hat{x},\hat{x}) \ge \alpha$$

for every unit vector $\hat{x}$. Similarly, if T is negative-definite, then there exists an $\alpha > 0$, such that

$$T(\hat{x},\hat{x}) \le -\alpha$$

for every unit vector $\hat{x}$.

Proof: Consider the function $g : \mathbb{R}^k \to \mathbb{R}$,

$$g(x) = T(x,x).$$

We will show that g is continuous. Indeed,

$$\begin{aligned}
|T(x_n,x_n) - T(x,x)| &= |T(x_n, x_n - x) + T(x_n - x, x)| \\
&\le \left| T\left(\frac{x_n}{\|x_n\|_k}, \frac{x_n - x}{\|x_n - x\|_k} \right) \right| \|x_n\|_k \|x_n - x\|_k + \left| \frac{x_n - x}{\|x_n - x\|_k}, \frac{x}{\|x\|_k} \right| \|x\|_k \|x_n - x\|_k \\
&\le \|T\|_{\text{op}} \|x_n\|_k \|x_n - x\|_k + \|A\|_{\text{op}} \|x\|_k \|x_n - x\|_k.
\end{aligned}$$

Letting $x_n \to x$ we obtain that $T(x_n,x_n) \to T(x,x)$. Since the unit sphere is compact, g assumes a minimum on the unit sphere; that is, there exists a unit vector $\hat{z}$, such that for all unit vectors $\hat{x}$,

$$T(\hat{x},\hat{x}) \ge T(\hat{z},\hat{z}) \stackrel{\text{def}}{=} \alpha > 0.$$

$\blacksquare$ —39h(2019)—

Lemma 3.31 *Let $A \subset \mathbb{R}^k$ be an open domain, and let $f \in C^2(A;\mathbb{R})$. Suppose that $(D^2 f)_a$ is negative-definite. Then, there exists an open ball $\mathcal{B}_r(a)$, such that $(D^2 f)_b$ is negative-definite for all $b \in \mathcal{B}_r(a)$.*

Proof: In essence, this lemma asserts that if the second derivative is continuous, then, its "definiteness" is continuous as well. By Lemma 3.30, there exists an $\alpha > 0$, such that for every unit vector $\hat{y}$,

$$(D^2 f)_a(\hat{y},\hat{y}) \le -\alpha.$$

We need to show that there exists an open ball $\mathcal{B}_r(a)$, such that for all $b \in \mathcal{B}_r(a)$ and all unit vectors $\hat{y}$,

$$(D^2 f)_b(\hat{y}, \hat{y}) < 0.$$

Since $D^2 f$ is continuous, there exists an r, such that for all $b \in \mathcal{B}_r(a)$,

$$\|(D^2 f)_b - (D^2 f)_a\|_{\text{op}} < \frac{\alpha}{2}.$$

Recall the definition of the operator norm of an element of $T \in \text{Hom}(\mathbb{R}^k, \text{Hom}(\mathbb{R}^k, \mathbb{R}))$,

$$\|T\|_{\text{op}} = \sup_{\hat{y}, \hat{z}} |T(\hat{y}, \hat{z})|.$$

Thus, for every $b \in \mathcal{B}_r(a)$ and every unit vector $\hat{y}$,

$$\begin{aligned}(D^2 f)_b(\hat{y}, \hat{y}) &= (D^2 f)_a(\hat{y}, \hat{y}) + ((D^2 f)_b(\hat{y}, \hat{y}) - (D^2 f)_a(\hat{y}, \hat{y})) \\ &\le -\alpha + |((D^2 f)_b - (D^2 f)_a)(\hat{y}, \hat{y})| \\ &\le -\alpha + \|(D^2 f)_b - (D^2 f)_a\|_{\text{op}} \\ &\le -\alpha + \frac{\alpha}{2} < 0.\end{aligned}$$

■

Theorem 3.32 *Let $f \in C^2(A; \mathbb{R})$, where $A \subset \mathbb{R}^k$ is open, and let $a \in A$ be a critical point.*

1. *If $(D^2 f)_a$ is negative-definite then a is a local maximum.*
2. *If a is a local maximum then $(D^2 f)_a$ is negative-semidefinite.*

Proof:

1. Suppose that $(Df)_a = 0$ and $(D^2 f)_a(x, x) < 0$ for all $x \neq 0$. By Taylor's theorem, there exists for every sufficiently small x a $\theta \in (0, 1)$ such that

$$f(a + x) = f(a) + (D^2 f)_{a+\theta x}(x, x),$$

where we used the fact that a is a critical point, i.e., $(Df)_a(x) = 0$. Since $(D^2 f)_a$ is negative-definite, it follows from Lemma 3.31 that there exists an open ball $\mathcal{B}_r(a)$, such that for every $a + x \in \mathcal{B}_r(a)$, $(D^2 f)_{a+x}$ is negative-definite. Thus, for every $x \in \mathcal{B}_r(a)$,

$$f(a+x) = f(a) + (D^2 f)_{a+\theta x}(x,x) \le f(a),$$

which proves that a is a local maximum.

2. Suppose that a is a local maximum of f. Here we may use the univariate condition of a maximum. Let $\mathcal{B}_r(a)$ be an open ball in which f reaches its maximum at a, and let $\hat{x}$ be a unit vector. Consider the path $g : (-r, r) \to \mathbb{R}$,

$$g(t) = f(a + t\hat{x}).$$

Then,

$$g'(t) = (Df)_{a+t\hat{x}}(\hat{x}),$$

and

$$g''(t) = (D^2 f)_{a+t\hat{x}}(\hat{x}, \hat{x}).$$

Since g reaches a local maximum at $t = 0$,

$$g'(0) = (Df)_a(\hat{x}) = 0 \qquad \text{and} \qquad g''(0) = (D^2 f)_a(\hat{x}, \hat{x}) \le 0,$$

i.e., $(D^2 f)_a$ is negative semi-definite.

■

3.5 The inverse function theorem

In this section we consider under what conditions does a function $f : \mathbb{R}^k \supset A \to B \subset \mathbb{R}^k$ have an inverse.

Example: Consider the function $f : A \to B$, where

$$A = (0, \infty) \times (0, \pi) \to \mathbb{R} \times (0, \infty)$$

given by

$$f(a) = \begin{pmatrix} a_1 \cos a_2 \\ a_1 \sin a_2 \end{pmatrix}.$$

Then f is invertible. Indeed, suppose that $f(a) = b$, i.e.,

$$b_1 = a_1 \cos a_2 \qquad \text{and} \qquad b_2 = a_1 \sin a_2.$$

Then,

$$b_1^2 + b_2^2 = a_1^2 \qquad \text{and} \qquad \frac{b_2}{b_1} = \tan_2 .$$

Thus,

$$f^{-1}(b) = \begin{pmatrix} \sqrt{b_1^2 + b_2^2} \\ \tan^{-1} \frac{b_2}{b_1} \end{pmatrix}.$$

▲▲▲

First of all, note that the domain and the range have the same dimension. Indeed, we know that in the particular case of linear functions this is a mandatory condition. In fact, we know the following:

Proposition 3.33 Let $T \in \mathrm{Hom}(\mathbb{R}^k, \mathbb{R}^k)$ and let $f : \mathbb{R}^k \to \mathbb{R}^k$ be the linear function $f(a) = Ta$. Then, f is invertible if and only if at some point $a \in \mathbb{R}^k$,

$$(Df)_a \in \mathrm{Hom}(\mathbb{R}^k, \mathbb{R}^k)$$

is invertible, which is the case if and only if $\det(Df)_a > 0$.

Proof: For every $a \in \mathbb{R}^k$, $(Df)_a = T$. The function inverse to f is $b \mapsto T^{-1}b$, and it exists if and only if T is invertible, which as we learned in linear algebra, occurs if and only if $\det T > 0$. ■

If f is nonlinear, then its derivative is not a constant matrix. If Df is invertible at a point, this is not sufficient to ensure that f is invertible, but as we will see, it ensures that f restricted to some neighborhood of a is invertible (on its image).

Lemma 3.34 Let $A \subset \mathbb{R}^k$ be an open set and $f \in C^1(A; \mathbb{R}^m)$. Let $a, b, c \in A$ such that the segment I connecting b and c is in A. Then,

$$\|f(c) - f(b) - (Df)_a(c - b)\|_m \le \|c - b\|_k \sup_{w \in I} \|(Df)_a - (Df)_w\|_{k,m}.$$

Proof: Define the function $g : A \to \mathbb{R}^m$,

$$g(w) = f(w) - (Df)_a(w - a).$$

The derivative of the linear operator $(Df)_a$ is constant and equals $(Df)_a$, hence,

$$(Dg)_w = (Df)_w - (Df)_a.$$

By Theorem 3.17,

$$\|g(c) - g(b)\|_m \le \|c - b\|_n \sup_{w \in I} \|(Dg)_w\|_{k,m},$$

which is the desired result. ■

—41h(2019)—

Theorem 3.35 (Inverse function theorem (משפט הפונקציה ההופכית)) Let $A \subset \mathbb{R}^k$ be an open set and let $f \in C^1(A; R^k)$. Suppose that $\det(Df)_a \neq 0$ for some point $a \in A$. Then there exists an open neighborhood U of a, $a \in U \subset A$, such that $V = f(U)$ is also open, and $f|_U$ is one-to-one and onto V. This defines an inverse function $f^{-1} : V \to U$, which is also continuously differentiable.

Comment: For real-valued functions, the inverse function theorem states that if $f : \mathbb{R} \supset A \to \mathbb{R}$ is continuously-differentiable and $f'(a) \neq 0$ at a point a, then there exists a neighborhood U of a in which $f|_U : U \to f(U)$ has an inverse, which is also continuously-differentiable.

Comment: If the theorem holds, then we can derive a formula for the derivative of f^{-1}. Since

$$f^{-1} \circ f = \mathrm{Id},$$

it follows from the chain rule that

$$(Df^{-1})_{f(a)} \circ (Df)_a = \mathrm{Id}$$

(the derivative of the linear operator Id at every point is Id), or,

$$(Df^{-1})_{f(a)} = (Df)_a^{-1}.$$

Proof: Since $\det(Df)_a \neq 0$, the linear transformation $(Df)_a$ has an inverse $(Df)_a^{-1}$. Moreover, since Df is continuous, there exists a closed ball $\hat{\mathcal{B}}_r(a) \subset A$ such that

$$\|(Df)_b - (Df)_a\|_{k,k} \le \frac{1}{2\,\|(Df)_a^{-1}\|_{k,k}}$$

for all $b \in \hat{\mathcal{B}}_r(a)$.

Step 1: Show that f is one-to-one in $\hat{\mathcal{B}}_r(a)$: By Lemma 3.34 and by our choice of r, for every two points $b, c \in \hat{\mathcal{B}}_r(a)$

$$\begin{aligned}\|f(c) - f(a) - (Df)_a(c-b)\|_k &\le \|c-b\|_k \sup_{w \in \hat{\mathcal{B}}_r(a)} \|(Df)_w - (Df)_a\|_{k,k} \\ &\le \frac{\|c-b\|_k}{2\,\|(Df)_a^{-1}\|_{k,k}}.\end{aligned}$$

Applying the (reverse) triangle inequality,

$$\begin{aligned}\|f(c) - f(b)\|_k &\ge \|(Df)_a(c-b)\|_k - \|f(c) - f(b) - (Df)_a(c-b)\|_k \\ &\ge \|(Df)_a(c-b)\|_k - \frac{\|c-b\|_k}{2\,\|(Df)_a^{-1}\|_{k,k}} \\ &= \frac{1}{\|(Df)_a^{-1}\|_{k,k}}\left[\|(Df)_a^{-1}\|_{k,k}\,\|(Df)_a(c-b)\|_k - \frac{1}{2}\|c-b\|_k\right] \\ &\ge \frac{1}{\|(Df)_a^{-1}\|_{k,k}}\left[\|(Df)_a^{-1}(Df)_a(c-b)\|_k - \frac{1}{2}\|c-b\|_k\right],\end{aligned}$$

where in the last step we used the fact that $\|T\|_{k,k}\|x\|_k \ge \|Tx\|_k$. It follows that.,

$$\|f(c) - f(b)\|_k \ge \frac{\|c-b\|_k}{2\,\|(Df)_a^{-1}\|_{k,k}}. \tag{3.3}$$

This implies that f is one-to-one in $\hat{\mathcal{B}}_r(a)$.

Step 2: We show that Df is invertible in $\hat{\mathcal{B}}_r(a)$: It suffices to show that its kernel is trivial,

$$\forall b \in \hat{\mathcal{B}}_r(a) \quad \text{and} \quad \forall 0 \neq y \in \mathbb{R}^k \qquad (Df)_b(y) \neq 0.$$

Let $b \in \hat{\mathcal{B}}_r(a)$ and let $0 \neq y \in \mathbb{R}^k$; by the continuity of the norm and (3.3),

$$\begin{aligned}\|(Df)_b(y)\|_k &= \left\|\lim_{t\to 0} \frac{f(b+ty) - f(b)}{t}\right\|_k = \lim_{t\to 0} \frac{1}{|t|}\|f(b+ty) - f(b)\|_k \\ &\ge \lim_{t\to 0} \frac{1}{|t|} \frac{\|ty\|_k}{2\,\|(Df)_a^{-1}\|_{k,k}} = \frac{\|y\|_k}{2\,\|(Df)_a^{-1}\|_{k,k}} > 0,\end{aligned}$$

which proves that $(Df)_b$ is invertible.

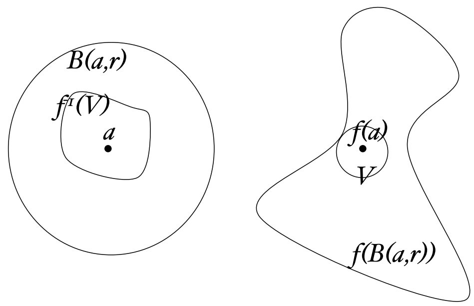

Step 3: construction of U **and** V: f is one-to-one in $\hat{\mathcal{B}}_r(a)$. It seems as if the *open* ball $\mathcal{B}_r(a)$ is a good candidate for the neighborhood U of a. The problem is that it is not necessarily true that $f(\mathcal{B}_r(a))$ is open, as continuous functions not necessarily map open sets into open sets. We rather show that $f(\mathcal{B}_r(a))$ contains an open neighborhood V of $f(a)$; since f is continuous, its pre-image intersected with $\mathcal{B}_r(a)$, which we denote by U, is open; f is then a one-to-one function from U onto V.

Let $b \in \partial\mathcal{B}_r(a)$; since f is one-to-one $f(b) \neq f(a)$. The boundary $\partial\mathcal{B}_r(a)$ is closed and bounded, hence compact, which implies that $f(\partial\mathcal{B}_r(a))$ is compact and does not include $f(a)$. We proved in the past that it implies that there is a positive distance, ε, between $f(a)$ and $f(\partial\mathcal{B}_r(a))$. Consider now the open ball $\mathcal{B}_{\varepsilon/2}(f(a))$. We claim that

$$\mathcal{B}_{\varepsilon/2}(f(a)) \subset f(\mathcal{B}_r(a)),$$

i.e.,

$$\forall w \in \mathcal{B}_{\varepsilon/2}(f(a)) \quad \exists b \in \mathcal{B}_r(a) \quad \text{such that} \quad f(b) = w.$$

While this may seem obvious by geometric intuition, this requires a proof. Let w be an arbitrary point in $\mathcal{B}_{\varepsilon/2}(f(a))$. Consider the function $h : \hat{\mathcal{B}}_a(r) \to \mathbb{R}$ defined by

$$h(b) = \|f(b) - w\|_k^2.$$

This function is continuous and defined on a compact set, hence it reaches a minimum. Denote this minimum by b, that is, $b \in \hat{\mathcal{B}}_a(r)$ satisfies

$$\|f(b) - w\|_k \le \|f(c) - w\|_k \qquad \text{for all } c \in \hat{\mathcal{B}}_a(r).$$

The point b cannot be on the boundary, since $b \in \partial\hat{\mathcal{B}}_a(r)$ would imply

$$\frac{\varepsilon}{2} < \|f(b) - w\|_k \le \|f(a) - w\|_k < \frac{\varepsilon}{2},$$

which is a contradiction.

Thus, $b \in \mathcal{B}_r(a)$. Since b is the minimizer of the differentiable function h, its derivative vanishes at b, i.e.,

$$0 = (Dh)_b = 2(Df)_b(f(b) - w).$$

Since $(Df)_b$ is invertible, it follows that $f(b) = w$, i.e., that $w \in f(\mathcal{B}_r(a))$. Since w was chosen as an arbitrary point in $\mathcal{B}_{\varepsilon/2}(f(a))$ it follows that

$$\mathcal{B}_{\varepsilon/2}(f(a)) \subset f(\mathcal{B}_a(r)).$$

At this point we set $V = \mathcal{B}_{\varepsilon/2}(f(a))$ and $U = f^{-1}(V) \cap \mathcal{B}_r(a)$. The sets U and V are open and f is one-to-one and onto V, hence f^{-1} is defined on V.

Step 4: We show that f^{-1} is Lipschitz continuous: Let $v, w \in V$. By (3.3), substituting $f^{-1}(v)$ and $f^{-1}(w)$,

$$\|f(f^{-1}(w)) - f(f^{-1}(v))\|_k \ge \frac{\|f^{-1}(w) - f^{-1}((v)\|_k}{2\,\|(Df)_a^{-1}\|_{k,k}},$$

i.e.,

$$\|f^{-1}(w) - f^{-1}(v)\|_k \le 2\,\|(Df)_a^{-1}\|_{k,k}\|w - v\|_k.$$

Step 5: We show that f^{-1} is continuously-differentiable: Let v be an arbitrary point in V, and $y \in \mathbb{R}^k$ sufficiently small such that $v + y \in V$. We define

$$b = f^{-1}(v) \qquad \text{and} \qquad x = f^{-1}(v + y) - f^{-1}(v),$$

and note that

$$v = f(b) \qquad \text{and} \qquad y = f(b + x) - f(b).$$

Since f^{-1} is Lipschitz continuous in V, there exists a constant M such that

$$\|x\|_k \le M\,\|y\|_k. \tag{3.4}$$

Then,

$$\begin{aligned} f^{-1}(v + y) - f^{-1}(v) - (Df)_b^{-1}(y) &= x - (Df)_b^{-1}(y) \\ &= (Df)_b^{-1}\left[(Df)_b(x) - y\right] \\ &= -(Df)_b^{-1}\left[f(b + x) - f(b) - (Df)_b(x)\right]. \end{aligned}$$

Taking norms, dividing by $\|y\|_k$ and using (3.4),

$$\begin{aligned}\frac{\|f^{-1}(v+y)-f^{-1}(v)-(Df)_b^{-1}(y)\|_k}{\|y\|_k} &\le \|(Df)_b^{-1}\|_{k,k}\frac{\|f(b+x)-f(b)-(Df)_b(x)\|_k}{\|y\|_k}\\ &\le M\|(Df)_b^{-1}\|_{k,k}\frac{\|f(b+x)-f(b)-(Df)_b(x)\|_k}{\|x\|_k}.\end{aligned}$$

Letting $y\to 0$, (3.4) implies that $x\to 0$, hence the right-hand side tends to zero by the definition of $(Df)_x$. It follows that f^{-1} is differentiable at y, and

$$(Df^{-1})_v = (Df)_b^{-1}.$$

Finally f^{-1} is continuously-differentiable due to the continuity of the matrix inverse. ■ —43h(2019)—

Before we proceed, some preliminaries. Recall that the **column rank** (**דרגת עמודות**) of a linear transformation $A\in\operatorname{Hom}(\mathbb{R}^k,\mathbb{R}^m)$ is the dimension of the vector space spanned by its columns (a subspace of $\mathbb{R}^m$), or the dimension of its image. Similarly, its **row rank** (**דרגת שורות**) is the dimension of the vector space spanned by its rows (a subspace of $\mathbb{R}^k$). The row rank and column rank are always equal, hence

$$\operatorname{rank}(A)\le\min(m,k).$$

A is said to have **full rank** (**דרגה מלאה**) if its rank of the largest possible given the dimensions. If A has full rank, then it is onto if and only if $m\le k$ (it is a transformation from a "large" space into a "small" space; A is a "fat-and-short" matrix).

Lemma 3.36 If $m\le k$ and $S\in\operatorname{Hom}(\mathbb{R}^k,\mathbb{R}^m)$ has full rank (i.e., it is onto), then there exists a matrix $T\in\operatorname{Hom}(\mathbb{R}^m,\mathbb{R}^k)$ such that $ST\in\operatorname{Hom}(\mathbb{R}^m,\mathbb{R}^m)$ is invertible.

Proof: Since S has rank m, its kernel is $(k-m)$-dimensional, i.e., $\mathbb{R}^k$ has an m-dimensional subspace which is orthogonal to $\ker S$—there exist m independent vectors $u_1,\dots,u_m\in\mathbb{R}^k$, such that

$$(\ker S)^\perp = \operatorname{Span}\{u_1,\dots,u_m\}.$$

Construct T such that its columns are the u_j's. Let $\mathbb{R}^m\ni x\neq 0$. Then,

$$Tx\in(\ker S)^\perp$$

(here we use the independence of the u_j's), hence $STx \neq 0$. This proves that ST is non-singular. ■

Consider momentarily functions $f : \mathbb{R} \to \mathbb{R}$. The derivative $f'(a)$ tells us that locally a displacement x from the point a will be mapped into a displacement $f'(a)x$ from the point $f(a)$. If $f'(a) \neq 0$, a small enough open neighborhood of a will be mapped into an open neighborhood of $f(a)$, that is, the mapping is locally an "open" one. Similarly, for functions $f : \mathbb{R}^k \to \mathbb{R}^m$ the derivative $(Df)_a$ tells us that a displacement x from a is mapped into a displacement $(Df)_a(x)$ of $f(a)$ (up to an $o(\|x\|_k)$ correction). If the rank of $(Df)_a$ is less than m, then there are directions in $\mathbb{R}^m$ which are not "covered" by local displacements. If, on the other hand, $(Df)_a$ has full rank, then we expect again that a small enough open neighborhood of a will be mapped into an open neighborhood of $f(a)$. The open mapping theorem formalizes these ideas.

Theorem 3.37 (Open mapping theorem (משפט ההעתקה הפתוחה)) Let $m \leq k$. Let $A \subset \mathbb{R}^k$ be an open set, and let $f \in C^1(A; \mathbb{R}^m)$ be such that $(Df)_a$ has full rank for all $a \in A$. Then f is an ***open mapping****: if $B \subset A$ is open in $\mathbb{R}^k$ then $f(B)$ is open in $\mathbb{R}^m$.*

Proof: We need to show that for every open set $B \subset A$ and for every $a \in B$, $f(a)$ is an interior point of $f(B)$, since it would imply that $f(B)$ only contains interior points, i.e., $f(B)$ is open. This follows from the next lemma. ■

Lemma 3.38 Let $m \leq k$. Let $A \subset \mathbb{R}^k$ be an open set and $f \in C^1(A; \mathbb{R}^m)$. Suppose that at a point $a \in A$, rank $(Df)_a = m$ *(i.e., the derivative has full rank), then $f(a)$ is an interior point of $f(A)$.*

Proof: Since $(Df)_a$ has full rank, there exists a linear mapping $T : \mathbb{R}^m \to \mathbb{R}^k$, such that $(Df)_a \circ T$ is invertible. Define the linear function $S : \mathbb{R}^m \to \mathbb{R}^k$,

$$S(x) = a + Tx,$$

and on $S^{-1}(A)$ we define the function

$$g(x) = f(S(x)),$$

which is a function from an open subset of $\mathbb{R}^m$ into $\mathbb{R}^m$ (since T is continuous). By the chain rule

$$(Dg)_0 = (Df)_{S(0)} \circ T = (Df)_a \circ T,$$

which is an invertible transformation. From the inverse mapping theorem, the point $0 \in \mathbb{R}^m$ has an open neighborhood which is mapped onto an open set that contains $g(0) = f(a)$. Since the image of g is a subset of the image of f, we conclude that $f(a)$ has a neighborhood included in $f(A)$, i.e., $f(a)$ is an interior point. ■

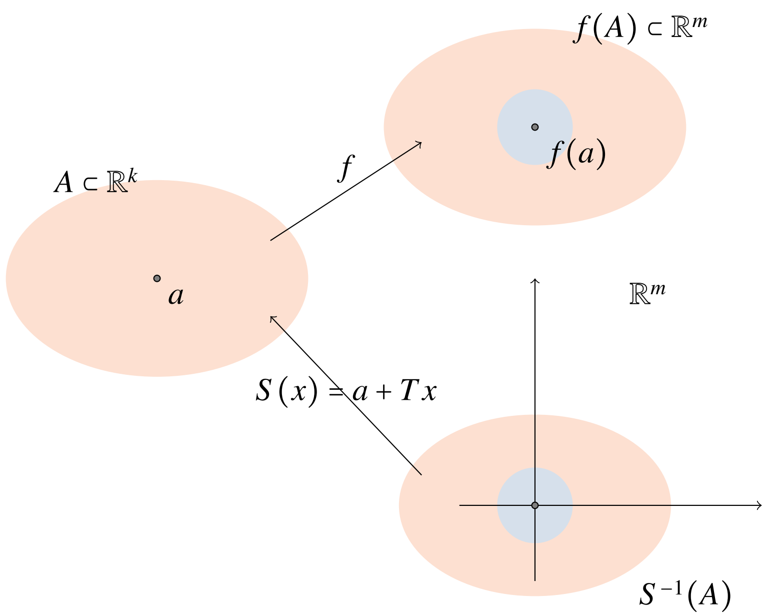

Comment: This theorem has a generalized version, known as the **Banach-Schauder theorem**: if X and Y are Banach spaces and $A : X \to Y$ is a surjective continuous linear transformation, then A is an "open mapping".

Example: Here is a "counter example": the mapping $f : \mathbb{R} \to \mathbb{R}$,

$$f(x) = x^2,$$

has a point $x = 0$, where $(Df)_0 = 0$ has rank less than one. This mapping is not open since $f(\mathbb{R}) = [0, \infty)$. ▲▲▲

3.6 The implicit function theorem

Consider, for example, the equation,

$$G(x, y) = e^y \cos(y^3 + xy^2 + x^2) + e^x \sin(xy) - 1 = 0. \tag{3.5}$$

This equation defines a relation between x and y; for every x there may exist zero, one, or more values of y for which this equation is satisfied. Thus, such an equation defines a function $\mathbb{R} \to \mathbb{R}$, that may not be defined on the whole real line, and that may or may not be single-valued. If a (single-valued) function $f : \mathbb{R} \to \mathbb{R}$ satisfying

$$G(x, f(x)) = 0$$

does exists, in say, a domain $A \subset \mathbb{R}$, we say that (3.5) defines **implicitly** the function $f(x)$. The implicit function theorem, which is the main topic of this chapter, states, in a more general setting, conditions on g that ensure that the equation $g(x, y) = 0$ does indeed define a function $y = f(x)$.

Theorem 3.39 (Implicit function theorem) Let $A \subset \mathbb{R}^k$ and $B \subset \mathbb{R}^m$ be open sets, and

$$G \in C^1(B \times A; \mathbb{R}^m) \qquad \textit{satisfies} \qquad f(b, a) = 0$$

at a point $(b, a) \in B \times A$. Furthermore, suppose that the square matrix

$$\partial_i G_j(b, a), \qquad i, j = 1, \dots, m$$

is invertible. Then there exist an open neighborhood $U \subset B \times A$ of (b, a), and an open neighborhood $V \subset \mathbb{R}^k$ of a, and a function $f \in C^1(V; B)$ such that for all $(y, x) \in U$,

$$G(y, x) = 0 \qquad \textit{iff} \qquad y = f(x).$$

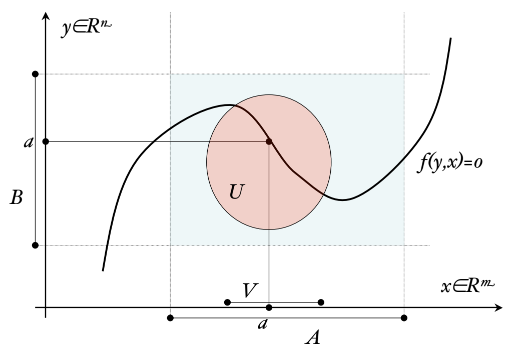

Comment: Just a variable count: the equation $G(y, x) = 0$ constitutes m equations for $m + k$ variables. Under the theorem's condition, if we fix k of those variables, we remain with m equations in m variables, and a unique solution exists. This implicitly defines a mapping between the k "fixed" variables and m resulting variables, i.e., a mapping $\mathbb{R}^k \to \mathbb{R}^m$.

Example: Consider the case $m = k = 1$ and the function

$$G(y, x) = x^2 y + xy^3 - 2.$$

At the point $(b, a) = (1, 1)$ we have $G(b, a) = 0$. Does there exist a neighborhood of $x = 1$ and a function $y = f(x)$ such that $G(y, x) = 0$ iff $y = g(x)$? According to the theorem, we only need to verify that

$$\partial_1 G(1, 1) = 1 + 3 = 4 \neq 0.$$

▲▲▲

Proof: Define a function $F : B \times A \to \mathbb{R}^{m+k}$ by

$$F : (y, x) \mapsto (G(y, x), x).$$

At the point $(y, x) = (b, a)$ we have

$$F(b, a) = (G(b, a), a) = (0, a),$$

and

$$(DF)_{(b,a)} = \begin{pmatrix} \partial_1 G_1 & \dots & \partial_m G_1 & \partial_{m+1} G_1 & \dots & \partial_{m+k} G_1 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ \partial_1 G_m & \dots & \partial_m G_m & \partial_{m+1} G_m & \dots & \partial_{m+k} G_m \\ 0 & \dots & 0 & 1 & \dots & 0 \\ 0 & \dots & 0 & 0 & \ddots & 0 \\ 0 & \dots & 0 & 0 & \ddots & 1 \end{pmatrix}$$

By assumption, the determinant of this matrix is non-zero. Thus, there exists, by the inverse function theorem, an open neighborhood U of (b,a) such that W is an open neighborhood of $(0,a)$, $F : U \to W$ is one-to-one and onto and $F^{-1} : W \to U$ is continuously differentiable. Note that F^{-1} is a mapping of the form

$$(z,x) \mapsto (h(z,x),x),$$

where $h : W \to B$ is continuously differentiable.

Define now

$$V = \{x \in \mathbb{R}^m : \ (0,x) \in W\}$$

and $f : V \to B$ by $f(x) = h(0,x)$. Since W is open so is V. And since h is continuously differentiable, so is f. Now,

$$\begin{aligned} (y,x) \in U,\ G(y,x) = 0 \quad &\Leftrightarrow \quad (y,x) \in U, \quad F(y,x) = (0,x) \\ &\Leftrightarrow \quad (0,x) \in W, \quad F^{-1}(0,x) = (y,x) \\ &\Leftrightarrow \quad x \in V, \quad h(0,x) = y \\ &\Leftrightarrow \quad x \in V, \quad f(x) = y. \end{aligned}$$

—45h(2019)— ■

Example: Consider the two equations:

$$\begin{aligned} 0 &= x^3 y t^3 + ytz + 3y^2 - 5 \\ 0 &= z^3 + xt^2 - 2. \end{aligned}$$

Does there exist a neighborhood of the point $x = y = z = t = 1$ where all the solutions are of the form

$$\begin{pmatrix} z \\ t \end{pmatrix} = f\begin{pmatrix} x \\ y \end{pmatrix},$$

where f is a continuously differentiable function $\mathbb{R}^2 \to \mathbb{R}^2$?

To comply to the structure of the theorem we rewrite this system as

$$\begin{aligned} G_1(y_1, y_2, x_1, x_2) &= x_1^3 x_2 y_2^3 + x_2 y_2 y_1 + 3x_2^2 - 5 = 0 \\ G_2(y_1, y_2, x_1, x_2) &= y_1^3 + x_1 y_2^2 - 2 = 0, \end{aligned}$$

where (y_1, y_2) are the old (z, t) and (x_1, x_2) are the old (x, y). Now

$$\begin{pmatrix} \partial_1 G_1 & \partial_2 G_1 \\ \partial_1 G_2 & \partial_2 G_2 \end{pmatrix} (1, 1, 1, 1) = \begin{pmatrix} 1 & 4 \\ 3 & 2 \end{pmatrix},$$

and the latter is invertible. ▲▲▲

3.7 Lagrange multipliers

In this section we consider the following problem: suppose we have a domain $B \subset \mathbb{R}^n$ and a function $f \in C^1(B; \mathbb{R})$, and we are looking for local minima or maxima of f. We have seen that a necessary condition is that Df (the gradient of f) vanishes as that point. Suppose however that we restrict ourselves to a subset of B determined by a set of k **constraints**: $g_i(x) = 0$, where $g_i \in C^1(B; \mathbb{R})$, $i = 1, \dots, k$. That is, we consider a restricted set:

$$A = \{x \in B : \ g_1(x) = \cdots = g_k(x) = 0\}.$$

(For A not to be a trivial set, we need $n \geq k + 1$.) The question is how to find a local extremum of f within the set A.

Example: A **probability distribution** on the set $\{1, \dots, n\}$ is a vector $p = (p_1, \dots, p_n)$ of non-negative entries that sum up to one. The **entropy** of the distribution is a function $H : \mathbb{R}^n \to \mathbb{R}$:

$$H(p) = -\sum_{i=1}^{n} p_i \log p_i.$$

We want to find the distribution that maximizes H under the constaint

$$g(p) = \sum_{i=1}^{n} p_i - 1 = 0.$$

▲▲▲

Before formalizing this problem, we can solve it by the following considerations. Recall that the gradient of f at a points along the direction normal to the $(n-1)$-dimensional hyperplane tangent to the level set of f. That is,

$$(\nabla f)_a \perp M_f^\perp(a) \stackrel{\text{def}}{=} \{x \in \mathbb{R}^n : (Df)_a(x) = 0\}.$$

For f to be a local extremum in the constrained set A, we need f not to (locally) vary along directions that are level sets of all the g_i at a. That is, we need

$$y \in M_{g_1}^\perp(a) \cap \cdots \cap M_{g_k}^\perp(a) \qquad \Rightarrow \qquad y \in M_f^\perp(a),$$

or.

$$M_{g_1}^\perp(a) \cap \cdots \cap M_{g_k}^\perp(a) \subseteq M_f^\perp(a).$$

This can be re-written as follows

$$\{x \in \mathbb{R}^n : (Df)_a(x) = 0\} \supseteq \left\{x \in \mathbb{R}^n : \sum_{j=1}^k \lambda_j (Dg_j)_a(x) = 0, \forall(\lambda_1, \ldots, \lambda_k)\right\}.$$

It follows that for a to be a local extremum of f in the hyperplane tangent to the level sets of all the g_i there must exist k numbers $(\lambda_1, \ldots, \lambda_k)$ such that

$$(Df)_a = \sum_{j=1}^k \lambda_j (Dg_j)_a.$$

This will be found to be a corollary of the following theorem:

Theorem 3.40 *Let $B \subset \mathbb{R}^n$, $k+1 \leq n$,*

$$f, g_1, \ldots, g_k \in C^1(B; \mathbb{R}),$$

$$A = \{x \in B : g_1(x) = \cdots = g_k(x) = 0\},$$

and $a \in A$ is a local extremum point of $f|_A$. Then the $(k+1)$-by-n matrix

$$\begin{pmatrix} \dfrac{\partial g_1}{\partial x_1} & \cdots & \dfrac{\partial g_1}{\partial x_n} \\ \vdots & \vdots & \vdots \\ \dfrac{\partial g_k}{\partial x_1} & \cdots & \dfrac{\partial g_k}{\partial x_n} \\ \partial_1 f & \cdots & \dfrac{\partial f}{\partial x_n} \end{pmatrix}$$

has rank less than $k+1$.

Proof: Suppose, wlog, that a is a local maximum point. That is, there exists an open neighborhood $U \subset B$ of a such that

$$f(a) = \max_{x \in A \cap U} f(x).$$

Consider now the function $F : U \to \mathbb{R}^{k+1}$ defined by

$$F(x) = \begin{pmatrix} g_1 \\ \vdots \\ g_k \\ f \end{pmatrix}.$$

We need to prove that $(DF)_a$ has rank less than $k+1$. By contradiction, if it had rank $k+1$, then by the open mapping theorem

$$F(a) = (0, \dots, 0, f(a))^T$$

would be an interior point of $F(U)$ in $\mathbb{R}^n$, which contradicts the fact that all the points of the form $(0, \dots, 0, t)$, $t > f(a)$ are not in $F(U)$. ■

Corollary 3.41 If $(Dg_1)_a, \dots, (Dg_k)_a$ are linearly independent, then there exist k numbers $\lambda_1, \dots, \lambda_k$, such that

$$(Df)_a = \sum_{i=1}^{k} \lambda_i (Dg_1)_a.$$

*The λ_i are called **Lagrange multipliers**. Thus, to find local extrema of f in A one needs to solve $n+k$ equations in $n+k$ unknowns:*

$$\begin{aligned} &g_1(x) = 0 \\ &\vdots \\ &g_k(x) = 0 \\ &(Df)_x = \sum_{i=1}^{k} \lambda_i (Dg_1)_x. \end{aligned}$$

Example: Let's return to the above example: to maximize

$$f(p) = -\sum_{i=1}^{n} p_i \log p_i,$$

subject to the constraint

$$g(p) = \sum_{i=1}^{n} p_i - 1$$

we solve the system in $n + 1$ variables,

$$\begin{aligned} 0 &= \sum_{i=1}^{n} p_i - 1 \\ -1 - \log p_j &= \lambda, \end{aligned}$$

from which we get that all the p_j are equal, i.e., equal to $1/n$. This means that the distribution that maximizes the entropy is the **uniform distribution**. ▲▲▲

Example: Find the non-negative vector in $\mathbb{R}^n$ that maximizes $f(x) = x_1 x_2 \dots x_n$ subject to the constraint that $x_1 + \cdots + x_n = n$. Here we solve the system

$$\begin{aligned} &x_1 + \cdots + x_n = n \\ &x_1 x_2 \dots x_{j-1} x_{j+1} \dots x_n = \lambda, \qquad j = 1, \dots, n. \end{aligned}$$

Here again, we deduce that all the x_j are equal to each other, and given by 1. This means that for all such x,

$$x_1 x_2 \dots x_n \le 1.$$

In particular, for every non-negative vector y the vector

$$x = \frac{ny}{y_1 + \cdots + y_n}$$

satisfies the normalization requirement, and therefore

$$\frac{n^n y_1 y_2 \dots y_n}{(y_1 + \cdots + y_n)^n} \le 1,$$

or,

$$(y_1 y_2 \dots y_n)^{1/n} \le \frac{1}{n}(y_1 + \cdots + y_n),$$

which is the well-known **arithmetic-mean-geometric-mean inequality**. ▲▲▲