Multidimensional Analysis & Geometry.
Lecture Notes

Kevin McGerty

Trinity term, 2023.

Contents

Sections, proofs, or individual Remarks which are marked with an asterisk (∗) are non-examinable.
Please contact me at mcgerty@maths.ox.ac.uk if you find any ambiguities or errors in these notes.

Index of Notation

$B(a,r)$	the open ball of radius f centred at a.
$\bar{B}(a,r)$	the closed ball of radius r centred at a.
$\mathcal{B}(X,Y)$	the space of bounded linear maps $\beta\colon X \to Y$ between normed vector spaces X and Y.
$\mathbf{B}_X$	the closed ball $\bar{B}(0_X,1)$ of radius 1 centred at 0_X in a normed vector space X.
$C^k(U,Y)$	for k a non-negative integer this is the space of continuous functions $f\colon U \to Y$ defined on an open subset U of a normed vector space X taking values in a normed vector space Y which are k times continuously differentiable.
$C^\infty(U,Y)$	the space of infinitely differentiable functions on an open subset U of a normed vector space X taking values in a normed vector space Y.
$\mathcal{L}(V,W)$	the space of linear maps $\alpha\colon V \to W$ between vector spaces V and W.
$\mathrm{Mat}_{m,n}(\mathbb{R})$	the space of $n \times m$ matrices with entries in $\mathbb{R}$.
$\mathrm{Mat}_n(\mathbb{R})$	the space of $n \times n$ matrices with entries in $\mathbb{R}$.
0_X	the zero vector in a vector space X. If $V = \mathbb{R}_n$ we write 0_n in place of $0_{\mathbb{R}^n}$, and if the vector space in question is clear from the context we suppress the subscript and write 0 rather than 0_X.
$O_Y(\lVert x\rVert)$	the space of functions f defined on a neighbourhood of 0_X in a normed vector space X taking values in a normed vector space Y with the property that there exist constants $C, r > 0$ such that $\frac{\lVert f(x)\rVert}{\lVert x\rVert} \leq C$ for all $x \in B(0_X, r)$.
$o_Y(\lVert x\rVert)$	the space of functions f defined on a neighbourhood of 0_X in a normed vector space X taking values in a normed vector space Y with the property that $\lim_{x\to 0} \frac{\lVert f(x)\rVert}{\lVert x\rVert} = 0$.
(U,a)	a pointed set, *i.e.* U is a set and $a \in U$ is an element of U.

Course Outline

- Definition of a derivative of a function from $\mathbb{R}^n$ to $\mathbb{R}^m$; examples; elementary properties; partial derivatives; the chain rule; the gradient of a function from $\mathbb{R}^n$ to $\mathbb{R}$; Jacobian. Continuous partial derivatives imply differentiability. Mean Value Theorems. [3 lectures]
- The Inverse Function Theorem and the Implicit Function Theorem (proofs are non-examinable). [2 lectures]
- The definition of a submanifold of $\mathbb{R}^n$. Its tangent and normal space at a point, examples, including two-dimensional surfaces in $\mathbb{R}^3$. [2 lectures]
- Lagrange multipliers. [1 lecture]

1 Review from A1: Linear maps and continuity

Everything in sections §1.1 and §1.2 apart from Definition 1.10 is covered in the Metric Spaces part of the A.1 core course. The only significant new result is proved in section §1.3: Theorem 1.17 shows that a linear map between normed vector spaces whose domain is finite-dimensional is automatically continuous.

1.1 Normed vector spaces

Before discussing the notion of differentiability for functions of many (real) variables, we begin by reviewing the relationship between the conditions of continuity and linearity for functions, in the natural context where both notions are defined, namely that of normed vector spaces.

Definition 1.1. A normed vector space $(X, \|.\|)$ is a pair consisting of a real[1] vector space X and a function $\|.\|\colon X \to \mathbb{R}$ which satisfies, for all $v, w \in X$ and $\lambda \in \mathbb{R}$:

1. $\|v\| \geq 0$ with equality if and only if $v = 0$. (*Positivity*.)
2. $\|\lambda.v\| = |\lambda|.\|v\|$. (*Homogeneity*.)
3. $\|v + w\| \leq \|v\| + \|w\|$. (*Triangle inequality*.)

We write 0_X for the zero vector in X (or simply 0 if there is no possibility for confusion). Taking $\lambda = 0$ in (2) we see that $\|0_X\| = 0$ and thus by (2) and (3) we must have

$$0 = \|0_X\| \leq \|v\| + \|-v\| = 2\|v\|.$$

Hence (2) and (3) in fact imply the inequality in (1), however the implication $\|v\| = 0 \implies v = 0$ does *not* follow from (2) and (3). A normed vector space is automatically a metric space, where the distance between $v_1, v_2 \in V$ is defined to be $\|v_1 - v_2\|$.

Remark 1.2. We will normally write $\|.\|$ for the norm on an arbitrary vector space, as it will be clear from context which vector space is in question. When there might be ambiguity[2], such as when we consider more than one norm on the same vector space, we will decorate the norm with a subscript, *e.g.* $\|.\|_X$ or $\|.\|_1$.

We will largely follow the notational conventions of the Metric Spaces and Complex Analysis course, and write, for example, for $a \in X$ and $r \geq 0$

$$B(a, r) = \{x \in X : \|x - a\| < r\}, \quad \bar{B}(a, r) = \{x \in X : \|x - a\| \leq r\},$$

for the open and closed balls respectively about a of radius r. Note that in a normed vector space, unlike in a general metric space, if $r > 0$ then the closed ball $\bar{B}(a, r)$ is always the closure $\overline{B(a, r)}$ of $B(a, r)$. When $V = \mathbb{R}^n$ we will write 0_n in place of $0_{\mathbb{R}^n}$.

We will also write $\mathbf{B}_X$ for the closed ball $\bar{B}(0_X, 1)$ and $S_X = \{v \in X : \|v\| = 1\}$ for its boundary, the unit sphere centred at 0_X.

Recall that if X is a normed vector space and $a \in X$ we say that a subset $U \subseteq X$ is a *neighbourhood* of a if there is some $r > 0$ such that the open ball $B(a, r)$ of radius r centred at a is contained in U. We say U is *open* if it is a neighbourhood of each of its points, that is, for every $x \in U$ there is some $r_x > 0$ such that $Bxr_x \subseteq U$.

Example 1.3. If X is one-dimensional, it is easy to understand all possible norms on X. Indeed if we pick $e_1 \in X\backslash\{0\}$, then for any $v \in X$ there is a unique $\lambda \in \mathbb{R}$ such that $v = \lambda.e_1$. Now if $f\colon X \to \mathbb{R}_{\geq 0}$ is homogeneous, so that $f(t.v) = |t|.f(v)$ for all $t \in \mathbb{R}$, then $f(v) = |\lambda|.f(e_1)$. Since it is easy to check that the absolute-value function $t \mapsto |t|$ on $\mathbb{R}$ is a norm, it follows from the formula $f(v) = |\lambda|f(e_1)$ that f is a norm on X provided f is not identically zero. Since any norm on X necessarily satisfies the homogeneity condition, it follows that any norm $\|.\|$ on X has the form $\|v\| = c.|\lambda|$ for $c > 0$ a positive real number (where, as above, $v = \lambda.e_1$).

[1] In fact one just needs a field with a sensible notion of "absolute value" – for example the complex numbers equipped with the modulus function.

[2] If you find an ambiguity I have missed, please let me know.

If $\dim(X) > 1$ – indeed even for $\dim(X) = 2$ – one cannot give such an explicit classification of all possible norms[3], but we will shortly see that, for finite dimensional vector spaces, all norms are *equivalent* in a sense which immediately implies they all yield the same notion of convergence, continuity, and uniform continuity.

Example 1.4. Let $X = \mathbb{R}^n$. Then there are many norms which are natural to consider. Perhaps the three most commonly used ones are the following: For $v = (v_1, \dots, v_n) \in \mathbb{R}^n$, we set

$$\begin{aligned} \|v\|_\infty &= \max_{1\leq i\leq n} |x_i|, \\ \|v\|_1 &= \sum_{i=1}^{n} |x_i| \\ \|v\|_2 &= \left(\sum_{i=1}^{n} x_i^2\right)^{1/2} \end{aligned}$$

Where it is important to emphasize which norm we are using on $\mathbb{R}^n$, we will write $\ell^n_\dagger$ for the normed vector space $(\mathbb{R}^n, \|.\|_\dagger)$ (where $\dagger \in \{1, 2, \infty\}$).

Example 1.5. The normed vector space ℓ^n_2 is an example of an *inner product space*, meaning that the norm comes from a positive definite symmetric bilinear form (or inner product): if $x, y \in \mathbb{R}^n$, then the pairing $\langle x, y\rangle = \sum_{i=1}^n x_i y_i$ (the standard "dot product") is such a form and $\|x\| = \langle x, x\rangle^{1/2}$. Inner product spaces have both a notion of distance and angle.

If X and Y are are finite-dimensional inner product spaces, and we write $\langle v_1, v_2\rangle_X$ denote the inner product on X and $\langle w_1, w_2\rangle_Y$ the inner product on Y, then, as in A0 Linear Algebra, for any $T \in \mathcal{L}(X, Y)$, there is a unique $T^* \in \mathcal{L}(Y, X)$ such that

$$\langle T(v), w\rangle_Y = \langle v, T^*(w)\rangle_X, \quad \forall v \in X, w \in Y. \tag{$\dagger$}$$

Indeed if one picks orthonormal bases B_X and B_Y for X and Y respectively, then applying ($\dagger$) to the elements of B_X and B_Y shows that if T has matrix A with respect to these bases then T^* must have matrix A^t. On the other hand it is easy to see using bilinearity ("multiplying out") that if T^* satisfies ($\dagger$) for $v \in B_X$ and $w \in B_Y$ then it satisfies ($\dagger$) for all $v \in X$ and $w \in Y$, thus T^* is just the linear map corresponding to the matrix A^t and the bases B_X, B_Y. Notice that this also shows $\mathrm{tr}(T) = \mathrm{tr}(T^*)$ since the trace of a matrix is equal to that of its transpose.

When X and Y are inner product spaces, we can make $\mathcal{L} := \mathcal{L}(X, Y))$ into an inner product spaces by setting

$$\langle S_1, S_2\rangle_{\mathcal{L}} = \mathrm{tr}_X(S_1^* S_2) = \mathrm{tr}_Y(S_2^* S_1), \quad \forall S_1, S_2 \in \mathcal{L}(X, Y)$$

where the second equality holds because $(S_1^* S_2)^* = S_2^*(S_1^*)^* = S_2^* S_1$ and since, as noted above, for any $T \in \mathcal{L}(X, Y)$ we have $\mathrm{tr}(T^*) = \mathrm{tr}(T)$, this is a symmetric bilinear form.

If we pick orthonormal bases $B_X = \{b_1, \dots, b_n\}$ and $B_Y = \{c_1, \dots c_m\}$ of X and Y respectively, then if $A = (a_{ij}) = {}_{B_Y}[S]_{B_X}$ is the matrix of S with respect to these bases, we have $a_{ij} = \langle c_i, S(b_j)\rangle_Y$, and hence

$$\langle S, S\rangle_{\mathcal{L}} = \mathrm{tr}(A^t A) = \sum_{\substack{1\leq k\leq n\\ 1\leq j\leq m}} a^t_{kj} a_{jk} = \sum_{\substack{1\leq k\leq n\\ 1\leq j\leq m}} a^2_{jk}$$

hence $\langle S, T\rangle_{\mathcal{L}}$ is positive definite – indeed it follows that $\mathcal{L}$ has an orthonormal basis consisting of the linear maps corresponding to the elemenary matrices $\{E_{ij}\}_{1\leq i,j\leq n}$. The associated norm on $\mathcal{L}(X, Y)$ is called the *Hilbert-Schmidt* norm, $\|S\|_{HS} = \langle S, S\rangle^{1/2}_{\mathcal{L}}$.

[3]Giving a norm $\|.\|$ on $\mathbb{R}^n$ is equivalent to giving the set $B_{\|.\|} = \{v \in V : \|v\| \leq 1\}$ of vectors in its closed unit ball. Such a set $B_{\|.\|}$ must be closed and bounded (both with respect to the Euclidean metric), convex, and preserved by the map $x \mapsto -x$, but otherwise can be arbitrary.

1.2 Bounded linear maps

Definition 1.6. If X and Y are vector spaces, we write $\mathcal{L}(X, Y)$ for the vector space of all linear maps from X to Y. If $X = Y$ then we write I_X for the identity map from X to itself. (In the case where $X = \mathbb{R}^n$ we will usually write I_n rather than $I_{\mathbb{R}^n}$.)

If we pick bases $B_X = \{e_1, \dots, e_n\}$ of X and $B_Y = \{f_1, \dots, f_m\}$ of Y respectively, then we can identify $\mathcal{L}(X, Y)$ with $\mathrm{Mat}_{m,n}(\mathbb{R})$ the space of n-by-m matrices where if $\alpha \in \mathcal{L}(X, Y)$ the $\alpha \mapsto A = (a_{ij})$ with $\alpha(e_j) = \sum_{i=1}^m a_{ij} f_i$. If $\dim(X) = \dim(Y) = n$, then we write $\mathrm{Mat}_n(\mathbb{R})$ instead of $\mathrm{Mat}_{n,n}(\mathbb{R})$.

Definition 1.7. A linear map $T\colon X \to Y$ is said to be *bounded* if there is some constant $C > 0$ such that

$$\|T(x)\| \leq C.\|x\|, \quad \forall x \in X.$$

We will write $\mathcal{B}(X, Y)$ for the set of bounded linear maps from X to Y. Note that, for $x \neq 0$, this condition is equivalent to $\|T(\frac{x}{\|x\|})\| \leq C$, thus T is bounded if and only if $\|T(x)\|$ is bounded on $\bar{B}(0_X, 1)$.

Exercise 1.8. In Problem Sheet 1, you are asked to show that a linear map $T \in \mathcal{L}(X, Y)$ is bounded if and only if it takes bounded subsets of X to bounded subsets of Y.

Bounded linear maps are clearly continuous, indeed Lipschitz continuous: if C is an upper bound for $T\colon X \to Y$ on $\bar{B}(0_X, 1)$ then if $x_1, x_2 \in X$ then $\|T(x_1) - T(x_2)\| = \|T(x_1 - x_2)\| \leq C.\|x_1 - x_2\|$, so that T is Lipschitz continuous with Lipschitz constant C. The following Lemma refines this observation slightly, using the notational conventions described in §5.1 of the Appendix.

Lemma 1.9. *Let X and Y be normed vector spaces. Then if $C^0(X, Y)$ denotes the space of continuous functions from X to Y we have*

$$\mathcal{B}(X, Y) = O_Y(\|v\|) \cap \mathcal{L}(X, Y) = C^0(X, Y) \cap \mathcal{L}(X, Y) = \mathcal{N}_0(X, Y) \cap \mathcal{L}(X, Y)$$

In particular, $\mathcal{B}(X, Y)$ is a vector space.

Proof. If $T\colon X \to Y$ is bounded then it is clear from the definition that it lies in $O_Y(\|.\|)$, and we have already seen above that it must be continuous. Since continuity implies continuity at 0_X, to complete the proof it suffices to show that if T is continuous at 0_X, then it is bounded. But if T is continuous at 0_X, then there is a $\delta > 0$ such that $\|T(v)\| < 1$ for all $v \in B(0_X, \delta)$. But then for any $v \in X$ with $\|v\| \leq 1$, we have $(1/2\delta).v \in B(0_X, \delta)$ so that $\|T((\delta/2).v)\| \leq 1$, and hence for all $v \in V$ with $\|v\| \leq 1$ we have $\|T(v)\| \leq 2/\delta$, that is, T is bounded. □

Definition 1.10. The space of bounded linear maps $\mathcal{B}(X, Y)$ is a normed vector space, with the norm, known as the *operator norm* given by $T \mapsto \|T\|_\infty$, where $\|T\|_\infty$ is defined as above. Using standard facts about suprema, you can check that this norm is *submultiplicative*, in the sense that if X, Y and Z are normed vector spaces, $S\colon X \to Y$ and, as above $T\colon Y \to Z$, then $\|T \circ S\|_\infty \leq \|T\|_\infty.\|S\|_\infty$.

Remark 1.11. In Metric Spaces, you studied the space $B(X)$ of real-valued *bounded functions* on an arbitrary set X and, for a metric space X, the space of bounded, real-valued, continuous functions $C_b(X)$. In that setting, a function is said to be bounded if its image is a bounded set. The image of a non-zero linear map $\alpha\colon X \to Y$ between normed vector spaces is never bounded, thus the usages are not, at first sight, consistent.

This apparent inconsistency is not, however, impossible to resolve[4]: Since it is compatible with scaling, a linear map α is completely determined by its values on $\mathbf{B}_X = \bar{B}(0_X, 1)$, indeed if $v \neq 0$ then $u = v/\|v\| \in \mathbf{B}_X$ and $\alpha(v) = \|v\|\alpha(u)$. Thus we get an injective map $r\colon \mathcal{B}(X, Y) \to C(\mathbf{B}_X, Y)$, from $\mathcal{B}(X, Y)$ to the space of continuous functions on $\mathbf{B}_X$ taking values in Y. Here $r(\alpha)$ is just the restriction of α to the closed ball $\mathbf{B}_X$. By definition, it gives an isometric embedding of $\mathcal{B}(X, Y)$, equipped with the operator norm, into $C_b(\mathbf{B}_X, Y)$, where the latter space is equipped with the usual supremum norm: $\|f\|_\infty = \sup\{\|f(x)\| : x \in \mathbf{B}_X\}$.

Definition 1.12. If X and Y are normed vector spaces, we say that $\alpha \in \mathcal{B}(X, Y)$ is a *topological isomorphism* if it has a bounded linear inverse. More precisely, $\alpha \in \mathcal{B}(X, Y)$ is a topologial isomorphism if there is a $\beta \in \mathcal{B}(Y, X)$ such that $\alpha \circ \beta = I_Y$ and $\beta \circ \alpha = I_X$. By Lemma 1.9, this is equivalent to the condition that α has a continuous linear inverse. When such an isomorphism exists, we say that X and Y are *topologically isomorphic*.

[4] It, of course, is perfectly acceptable to just remember the apparent inconsistency in usage.

Note that because a linear map is continuous if and only if it is uniformly continuous, and indeed Lipschitz continuous, if X and Y are normed vector spaces and X is a complete, then if Y is topologically isomorphic to X, it must also be complete, since uniformly continuous maps preserve Cauchy sequences.

Definition 1.13. If X is a vector space with two norms $\|.\|_a$ and $\|.\|_b$, then $\|.\|_a$ and $\|.\|_b$ are *equivalent* if the identity map is a topological isomorphism from $(X, \|.\|_a)$ to $(X, \|.\|_b)$.

To make this explicit, let $\iota\colon (X, \|.\|_a) \to (X, \|.\|_b)$ be the identity map viewed as a map between two different normed vector spaces $(X, \|.\|_a)$ and $(X, \|.\|_b)$. The fact that ι is bounded is equivalent to the existence of a constant $C_1 > 0$ such that, for all $v \in X$ we have $\|v\|_b = \|\iota(v)\|_b \leq C_1.\|v\|_a$. On the other hand, the fact that ι^{-1} is bounded is equivalent to the existence of a constant $C_2 > 0$ such that $\|v\|_a = \|\iota^{-1}(v)\|_a \leq C_2.\|v\|_b$. Setting $c = C_1^{-1}$ and $C = C_2$, this is equivalent to the existence of constants $c, C > 0$ such that

$$c.\|v\|_b \leq \|v\|_a \leq C.\|v\|_b \quad \forall v \in X. \tag{1.1}$$

If $\|.\|_a$ and $\|.\|_b$ are equivalent, then they yield the same notions of continuity, convergence, and uniform continuity and a function f is $o(\|x\|_a)$ if and only if it is $o(\|x\|_b)$.

Example 1.14. Consider the norms $\|.\|_1$ and $\|.\|_2$ on $\mathbb{R}^n$ defined above. We claim that they are equivalent. Indeed if $x = (x_1, \ldots, x_n)$, then clearly

$$\|x\|_2^2 = \sum_{i=1}^n |x_i|^2 \leq \sum_{i=1}^n |x_i|^2 + 2\sum_{i<j} |x_i|.|x_j| = (\sum_{i=1}^n |x_i|)^2 = \|x\|_1^2.$$

so that $\|x\|_2 \leq \|x\|_1$. On the other hand, applying Cauchy-Schwarz to the vectors $u_1 = (1, 1 \ldots, 1)$ and $u_2 = (|x_1|, \ldots, |x_n|)$, we see that

$$\|x\|_1 = \sum_{i=1}^n |x_i| = \sum_{i=1}^n 1.|x_i| \leq n^{1/2}.\|x\|_2,$$

Remark 1.15. Let $X = C([0, 1])$ be the space of continuous functions on the interval $[0, 1]$ and let $Y = C_0^1([0, 1])$ be the space of continuously differentiable functions on the same interval (with one-sided derivatives at the end-points) which vanish at the origin. View both X and Y as normed vector spaces using the supremum norm. Then we have a linear map $T\colon X \to Y$, where if $f \in X$,

$$T(f)(x) = \int_0^x f(t)dt.$$

The fundamental theorem of calculus shows that $T(f)$ is indeed in $Y = C_0^1([0, 1])$ if $f \in C([0, 1])$, and the triangle equality for integrals shows that $\|T(f)\| \leq \int_0^1 |f(t)|dt \leq \|f\|_\infty$, so that $T \in \mathcal{B}(X, Y)$. While T is invertible with inverse $D\colon Y \to X$, where $D(g) = g'$ for all $g \in Y$, it is easy to see that D is unbounded. Thus while T is a linear isomorphism, it is not a topological isomorphism.

This difference between integration and differentiation is closely related to the ideas discussed in Picard's Theorem in Differential Equations 1.

1.3 Finite dimensional normed vector spaces

Lemma 1.16. *Let X be a normed vector space and let $T\colon \ell_1^n \to X$ be a linear map, (where $\ell_1^n = (\mathbb{R}^n, \|.\|_1)$). Then T is automatically bounded, and moreover, if T is bijective, then it is a topological isomorphism.*

Proof. Let $\{e_1, \ldots, e_n\}$ be the standard basis of $\mathbb{R}^n$, and set $M_1 = \max\{\|T(e_i)\| : 1 \leq i \leq n\}$. Now any $x \in \mathbb{R}^n$ can be written as $x = \sum_{i=1}^n \lambda_i e_i$, and hence

$$\|T(x)\| = \|\sum_{i=1}^n \lambda_i T(e_i)\| \leq \sum_{i=1}^n |\lambda_i|.\|T(e_i)\| \leq M_1.\|x\|_1,$$

and so T is bounded.

Now suppose that T is bijective. Its set-theoretic inverse is automatically linear, and to show it is continuous, *i.e.* bounded, we must show there is some $M_2 > 0$ such that $\|T^{-1}(v)\|_1 \leq M_2\|v\|$, for all $v \in X$, or equivalently (setting $x = T^{-1}(v)$ and $C = M_2^{-1}$) some $C > 0$ such that

$$C.\|x\|_1 \leq \|T(x)\| \iff C \leq \left\| T\left(\frac{x}{\|x\|_1}\right)\right\|.$$

Now if $S_1 = \{x \in \mathbb{R}^n : \|x\|_1 = 1\}$ (the "sphere" of unit radius in the $\|.\|_1$-norm) then, by Bolzano-Weierstrass, S_1 is compact, and $x \mapsto \|T(x)\|$ is continuous, its image is closed and bounded in $\mathbb{R}$. Now since $\|T(x)\| > 0$ for all $x \in S_1$ (since $\|.\|$ is a norm) $m = \min\{\|T(x)\| : x \in S_1\} > 0$, and hence we may take $C = m$. □

Theorem 1.17. *Let X and Y be normed vector spaces. If X is finite-dimensional then $\mathcal{L}(X, Y) = \mathcal{B}(X, Y)$, that is, every linear map from X to Y is automatically continuous. In particular, any two norms on X are equivalent.*

Proof. Let $n = \dim(X)$ and suppose $T : X \to Y$ is a linear map. Picking a basis $B = \{v_1, \dots, v_n\}$ of X induces an bijective linear map $\phi_B : \mathbb{R}^n \to X$ given by $\phi_B(\lambda_1, \dots, \lambda_n)^t = \sum_{i=1}^n \lambda_i v_i$. Then by the previous Lemma we see that ϕ_B is a topological isomorphism, and also that the composition $T \circ \phi_B : \mathbb{R}^n \to Y$ is continuous. But then $T = (T \circ \phi_B) \circ \phi_B^{-1}$ is a composition of continuous functions and hence is continuous as required.

For the final sentence, let $\|.\|_a$ and $\|.\|_b$ be two norms on X, By the first part of the Lemma, the identity map, viewed as a map from $(X, \|.\|_a)$ to $(X, \|.\|_b)$ is continuous, as is its inverse, which is the identity map viewed as a map from $(X, \|.\|_b)$ to $(X, \|.\|_a)$, which precisely says that $\|.\|_a$ and $\|.\|_b$ are equivalent. □

Corollary 1.18. *Let X be a normed vector space and let F be a finite dimensional subspace. Then F is a closed subset of X.*

Proof. If $\dim(F) = k$, then Theorem 1.17 show that a linear isomorphism $\phi : \mathbb{R}^k \to F$ is automatically continuous (viewing $\mathbb{R}^k$ as a normed vector space with the $\|.\|_1$-norm). Since a continuous linear map is automatically Lipschitz continuous, and $\mathbb{R}^k$ is complete, so is F. As a complete subspace of a metric space it must be closed (see the proof of Lemma 6.2.1 in **[?]** – a closed subset of a complete metric space is complete, but a complete subspace of a metric space is always closed whether or not the the ambient space is complete). □

Remark 1.19. The upshot of the previous discussion is that, for the purposes of this course, we do not lose any generality by assuming our normed vector spaces are of the form $\mathbb{R}^n$ equipped with the $\|.\|_2$ norm associated to the standard dot product (and thus the spaces of linear maps between them can also be viewed as an inner product space using the Hilbert-Schmidt norm, or as a normed vector space using the operator norm). However, the results of this section shows that we are free to use whichever norm is convenient (*e.g.* in the proof of the previous corollary, the $\|.\|_1$ norm is the simplest to consider) and that, even if we state results for $(\mathbb{R}^n, \|.\|_2)$, they hold for any finite-dimensional normed vector space.

Indeed part of our goal in this course is to show the advantages of being able to choose good "local" coordinates when studying differentiable functions, by analogy with the way in which we study linear maps by finding a basis with respect to which they are as simple as possible (*e.g.* diagonalisable) we will take care however to point out when the concepts we study require a choice of basis for our vector space or not.

2 The derivative in higher dimensions

Suppose that U is an open subset of $\mathbb{R}^n$ and $f\colon U \to \mathbb{R}^m$ is an $\mathbb{R}^m$-valued function. We would like to extend the one-variable notion of the differentiability to functions of this kind, which have both higher-dimensional input and output. First however, it is important to note that we must equip $\mathbb{R}^n$ and $\mathbb{R}^m$ with metrics in order for the notion of a limit to make sense, and if such a metric obeys some natural compatibilities with vector addition and scalar multiplication, it is induced by a norm. Thus a more invariant (or "coordinate free") way to phrase our goal, is the following: Given (finite-dimensional) normed vector spaces X and Y and an open subset U of X, what is a sensible definition of the derivative of a function $f\colon U \to Y$?

To extend the notion of differentiability to the case where $n > 1$, it is useful to recall some of the natural interpretations of the (one-variable) derivative: In dynamics, the derivative arises from the notion of instantaneous speed or velocity, while in geometry, the derivative at a point a gives the slope of the tangent line to the graph of f at the point $(a, f(a))$.

2.1 The one-dimensional case

Let us first consider the case of a function $f\colon X \to Y$, where $\dim(X) = \dim(Y) = 1$. Recall that, for a function $g\colon \mathbb{R} \to \mathbb{R}$, the derivative of g at a point $a \in \mathbb{R}$ is defined to be

$$\begin{aligned} Dg(a) = g'(a) &:= \lim_{x\to a} \frac{g(x) - g(a)}{x - a} \\ &= \lim_{h\to 0} \frac{g(a+h) - g(a)}{h} \end{aligned} \tag{2.1}$$

But now if we are given a function $f\colon X \to Y$ between two 1-dimensional different vector spaces, the if $x \neq a$ are vectors in X, the difference $f(x) - f(a)$ is a vector in Y, while $x - a$ is a vector in X, so it seems meaningless to consider their quotient. The obvious response to this problem is to pick coordinates so that we can identify both X and Y with $\mathbb{R}$, and then apply the standard definition. Thus let us pick a basis vector $e_1 \in X$ and a basis vector $e_2 \in Y$, and let us identify X with $\mathbb{R}$ via $t \mapsto i_1(t) = a + te_1$, and similarly we identify Y with $\mathbb{R}$ via $s \mapsto i_2(s) = f(a) + se_2$, that is, we centre our coordinates at a and $f(a)$ respectively.

Using these identifications, we obtain a scalar function $F_{e_1,e_2}\colon \mathbb{R} \to \mathbb{R}$, which is given by the equation

$$f(a) + F_{e_1,e_2}(t).e_2 = f(a + te_1).$$

One can view this equation as the requirement that, in the diagram:

$$\begin{tikzcd} 0_X \in X \arrow[r, "f"] & Y \ni 0_Y \\ 0 \in \mathbb{R} \arrow[u, "i_1"] \arrow[r, "F_{e_1,e_2}"'] & \mathbb{R} \ni 0 \arrow[u, "i_2"'] \end{tikzcd}$$

if one goes from the bottom left to top right by either of the possible compositions, one gets the same answer, that is $f \circ i_1 = i_2 \circ F_{e_1,e_2}$. Note that $F_{e_1,e_2}(0) = 0$, and, as a function from $\mathbb{R}$ to itself we can ask if F_{e_1,e_2} is differentiable at $t = 0$, that is, as $F_{e_1,e_2}(0) = 0$, if

$$\lim_{t\to 0} \frac{F_{e_1,e_2}(t)}{t}$$

exists. If it does, we denote it by $D_{e_1,e_2}f(a) = F'_{e_1,e_2}(0)$.

If $D_{e_1,e_2}f(a)$ was actually independent of the choice of bases $\{e_1\}, \{e_2\}$, then it would give a natural defintion of the derivative of f at a. However, if we choose different basis vectors $e'_1 = \lambda.e_1$ and $e'_2 = \mu.e_2$, then the associated scalar function $F_{e'_1,e'_2}$ is given by $F_{e'_1,e'_2}(t) = \mu^{-1}.F_{e_1,e_2}(\lambda.t)$, and hence $F'_{e'_1,e'_2}(0) = (\lambda/\mu).F'_{e_1,e_2}(0)$. In other words $D_{e'_1,e'_2}f(a) = (\lambda/\mu)D_{e_1,e_2}f(a)$.

Remark 2.1. One conclusion we might draw from the calculations above is that this is not the correct definition. With a bit more thought, however, it turns out that the correct conclusion to take from them is that the derivative $Df(a)$ is not in fact a scalar! It is instead an object to which we can associate a scalar once we choose bases of X

and Y respectively. Moreover, if we know this scalar for one choice of bases $\{e_1\}, \{e_2\}$, we can determine the scalar associated to any other choice of bases provided we can express those bases in terms of the bases $\{e_1\}, \{e_2\}$.

If this sounds esoteric, it is worth noticing that in fact we already knew this from physics: Recall that if a particle moves in space so that its position $x(t)$ is a function of the time t, then the derivative $\frac{dx}{dt}(t)$ is the velocity of the particle at time t. But velocity is not a dimensionless scalar, it has (S.I.) units ms^{-1}, and the factor λ/μ we found above matches those units: the choice of e_1 provided our "units", or scale, for the domain of f (which in the case of $x(t)$ is time, which is measured in seconds) and the choice of e_2 provides "units" for the codomain of f, which for $x(t)$ is space, and distance is measured in metres. Viewing a change of the choice of bases from $\{e_1\}$ and $\{e_2\}$ to $\{e'_1\}$ and $\{e'_2\}$ as a change of units, for example, changing the unit of time to hours, so that $h = 3600s$, and the unit of distance to kilometres, so that $km = 1000m$, then if the velocity is $v(t) = \frac{dx}{d}(t)$ in ms^{-1}, it becomes $3.6 = 3600/1000$ times $v(t)$ in $km.h^{-1}$, which is precisely the factor (λ/μ) which we just observed above.

The previous remark hopefully confirms that $Df(a)$ has to be something other than a scalar, but perhaps it does not quite tell us how what kind of object we should expect $Df(a)$ to be. We can gain some insight into this simply by considering more carefully where we are forced to take coordinates (rather than just picking coordinates wherever we can). Noticing that in a vector space we can of course divide by any nonzero scalar, we see that it makes sense to ask if the limit

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

exists – that is, the standard formula for the derivative becomes syntactically coherent as soon as we chose a basis $\{e_1\}$ of X, so we did not need to pick a basis for Y. For $e_1 \in X$ non-zero, we may therefore define

$$D_{e_1} f(a) := \lim_{t\to 0} \frac{f(a+t.e_1) - f(a)}{t} \tag{2.2}$$

wherever this limit exits. Note that $D_{e_1} f_a$ is now an element of Y, rather than a scalar. However, as

$$\frac{f(a+te_1) - f(a)}{t} = \frac{F_{e_1,e_2}(t)}{t}.e_2,$$

it follows easily that that $D_{e_1} f(a) = D_{e_1,e_2} f(a).e_2$. Thus simply by replacing $D_{e_1,e_2} f(a)$ by the corresponding multiple of e_2 we remove the dependence on the choice of a basis is Y. Now consider (2.2) when $e_1 \in X$ is arbitrary:

(i) If we take $e_1 = 0_X$ in (2.2), then $f(a + t.0_X) = f(a)$ and hence the limit on the right-hand side exists, and is equal to 0_Y.

(ii) It follows that if the limit in (2.2) exists for some non-zero vector in X, say a vector e_0 with $\|e_0\| = 1$. Then (2.2) defines, for *any* $v \in X$, a vector $D_v f(a)$ in Y where if $v = \lambda.e_0$ then $D_v f(a) = \lambda.D_{e_0} f(a)$. Since $\dim(X) = 1$, this shows that $v \mapsto D_v f(a)$ is a linear map from X to Y.

Thus we have finally have a natural description of what $Df(a)$ is: it is a linear map from X to Y sending $v \in X$ to $D_v f(a) \in Y$.

Remark 2.2. Of course, in addition to velocity and speed, the classic interpretation of the derivative of a function f at a point a is as the "slope of the tangent line" to the graph of f at $(a, f(a))$. Indeed the tangent line is just the graph of the function $f(t) = f(a) + f'(a)(t-a)$. Here again we can see that viewing the derivative, or slope, as a scalar is adequate if one is considering functions from $\mathbb{R}$ to itself, but as soon as we consider functions $f\colon X \to Y$ between two arbitrary one-dimensional vector spaces, we see that the tangent line must be the graph of a function of the form $t \mapsto f(a) + \alpha(t-a)$, where $\alpha \in \mathcal{L}(X, Y)$ is linear. Thus we are also led to consider $Df(a)$ as a linear function from X to Y by the "slope" interpretation of the derivative.

Notice that when $X = Y$, the scalar multiplication action of $\mathbb{R}$ on X gives a natural isomorphism $\mathbb{R} \to \mathcal{L}(X, X)$. Thus when $X = Y = \mathbb{R}$ the linear map really is just the scalar which gives its slope.

Remark 2.3. The considerations above for the one-dimensional case also really only used the fact that $\dim(X) = 1$ – the dimension of Y was not important. Thus we have in fact obtained a definition of the derivative for functions from an open subset of a one-dimensional vector space to a vector space of arbitrary dimension.

Definition 2.4. (*The 1-dimensional case.*) Let X and Y be normed vector spaces and suppose that $\dim(X) = 1$. Let $U \subseteq X$ be an open set and suppose $f : U \to Y$ is a function. If $a \in U$ then we define the derivative of f at a to be the linear map $Df_a \in \mathcal{L}(X, Y)$ given by

$$Df_a(v) = \lim_{t \to 0} \frac{f(a + t.v) - f(a)}{t},$$

where this limit exists. As noted above, the limit is compatible with scalar multplication, so that $Df_a(\lambda.v) = \lambda.Df_a(v)$ for any $\lambda \in \mathbb{R}$ and $v \in X$, and as X is 1-dimensional, this implies Df_a is a linear map. Indeed this also shows that if we know $Df_a(v)$ exists for a single non-zero vector $v_0 \in X$, then it exists for any $v \in X$.

2.2 The general case

Our consideration of the one-dimensional case gives some indication of what we should seek in the higher dimensional context: If X and Y are arbitrary finite-dimensional vector spaces, and $f : U \to Y$ is a function defined on an open subset U of X, then for $a \in U$, given our examination of the one-dimensional case, it is natural to demand that the derivative[5] Df_a of f at a is an element of $\mathcal{L}(X, Y)$.

Moreover, our definition in the one-dimensional case also yields a sensible notion in higher dimensions:

Definition 2.5. Let $f : U \to Y$ be as above and suppose $a \in U$ and $v \in X$. The *directional derivative* of f at $a \in U$ in the direction v is defined to be

$$\partial_v f(a) = \lim_{t \to 0} \frac{f(a + t.v) - f(a)}{t},$$

where this limit exists. Assuming it exists, it is an easy exercise to check that, for any $s \in \mathbb{R}$, we have $\partial_{s.v} f(a) = s.\partial_v f(a)$. That is, the directional derivative is *homogeneous* in v. For this reason, when taking a directional derivative we normally assume the direction vector v has unit length, *i.e.* $\|v\| = 1$. Note also that, if $\dim(X) = 1$, then we have $Df_a(v) = \partial_v f(a)$.

The above definition and its relation to the derivative in the one-dimensional case suggests that either of following might be reasonable:

Provisional Definitions: If $f : U \to Y$ is a function defined on an open subset U of a normed vector space X taking values in a normed vector space Y, then:

1. Proposal 1: f is differentiable at a if all the directional derivatives at a exist, and we define its derivative[6] at a to be the function $P_1 f_a(v) = \partial_v f(a)$.
2. Proposal 2: f is differentiable at a if there is a linear map $T \in \mathcal{L}(X, Y)$ such that for all $v \in X$, we have $T(v) = \partial_v f(a)$. This linear map T, if it exists, is certainly unique, and will be denoted $P_2 f_a$. Clearly, when it exists $P_2 f_a = P_1 f_a$.

The following examples show that these proposals are genuinely different:

Example 2.6.

(i) Let $f : \mathbb{R}^2 \to \mathbb{R}$ in Figure 1 given by

$$f_1(x_1, x_2) = \begin{cases} x_1 x_2(x_1 + x_2)/(x_1^2 + x_2^2), & (x_1, x_2) \neq (0, 0), \\ 0, & (x_1, x_2) = (0, 0) \end{cases}$$

Consider the directional derivative of f_1 in the direction $v = (v_{1,2})$.

$$\partial_v f(0) = \lim_{t \to 0} \frac{f_1(tx_1, tx_2)}{t} = \lim_{t \to 0} \frac{t^3 v_1 v_2(v_1 + v_2)}{t(t^2 v_1^2 + t^2 v_2^2)} = \frac{v_1.v_2(v_1 + v_2)}{v_1^2 + v_2^2} = f(v)$$

Thus all the directional derivatives exist, and so using Proposal 1, f_1 is differentiable at 0_2 with $P_1 f_{0_2} = f_1$, that is, f_1 is its own derivative at 0_2! On the other hand, since f_1 is clearly not a linear function, f_1 is not differentiable in the sense of Proposal 2.

[5]We write Df_a rather than $Df(a)$ because $Df_a \in \mathcal{L}(X, Y)$ so it is a function itself, and $Df_a(v)$ is more compact to read than $Df(a)(v)$.

[6]The use of the letter "P" is to indicate "provisional".

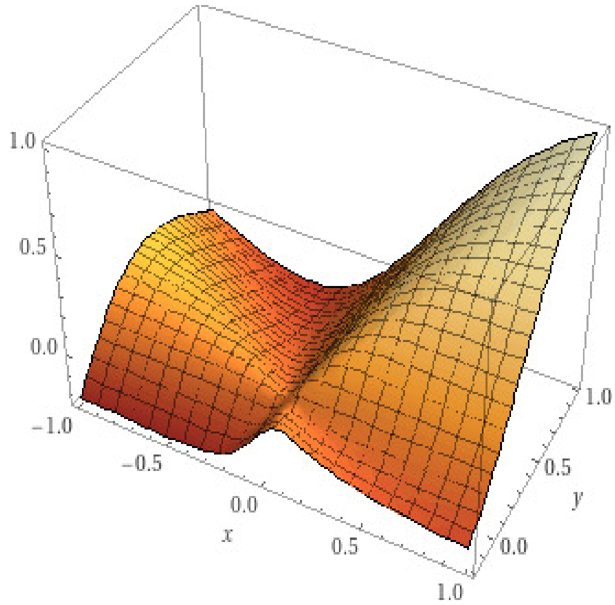

Figure 1: Graph of $f(x_1, x_2) = x_1 x_2(x_1 + x_2)/(x_1^2 + x_2^2)$. All its directional derivatives exist at 0_2 but it is not differentiable there.

(ii) Let Ω be the open subset $\{(x_1, x_2) \in \mathbb{R}^2 : 0 < x_1, 0 < x_2 < x_1^2\}$ and let $f_2 = \mathbb{1}_\Omega$ be the indicator function of Ω, so that $f_2(x_1, x_2) = 1$ if $(x_1, x_2) \in \Omega$ and $f_2(x_1, x_2) = 0$ otherwise. To calculate the directional derivatives of f_2 at 0_2, suppose that $v = (v_1, v_2) \in S_{\mathbb{R}^2}$. Clearly, since $f_2(t.(v_1, v_2)) = 0$ whenever $v_1.v_2 \leq 0$, $\partial_v f_2(0_2) = 0$ unless $v_1.v_2 > 0$. But if $v_1.v_2 > 0$, then if $|t| < |v_2|/v_1^2$, $t.(v_1, v_2) \notin \Omega$, hence $\lim_{t\to 0} f_2(t.(v_1, v_2))/t = \lim_{t\to 0} 0/t = 0$. Hence all of the directional derivative $\partial_v f_2(0)$ exists and equal 0_2. It follows that f_2 is differentiable in the sense of both proposals, with it derivative $P_2 f_{0_2}$ being the zero linear map.

The function f_1 above shows the difficulty with Proposal 1: this notion of differentiability will only be useful if we first develop a theory of homogeneous functions, as Df_a will only be homogeneous, *i.e.* be compatible with scalar multiplication, rather than linear. If you note that a homogeneous function is determined by its values on the unit sphere S_X, and that any continuous function $f\colon S_X \to Y$ from the unit sphere on X to a normed vector space Y extends to a homogeneous function from X to Y provided $f(-x) = -f(x)$ for all $x \in S_X$, it is clear that the space of continuous homogeneous functions from X to Y is a much more complicated one that the space of linear maps from X to Y, so any such theory will be much harder than linear algebra. Indeed the function f_1 in Example 2.6 is differentiable at 0_2 according to suggestion 1, but by the provisional definition P_1 the derivative is $Df_{1,0_2}(v) = f_1(v)$, so that passing to Df_1 does not provide a simpler object to study.

On the other hand, the function f_2 shows that simply demanding that the directional derivatives yield a linear function may not be the correct condition: If we recall the idea that the derivative at a point a should provide the tangent plane to the function at a, then the plane $T = \{(x_1, x_2, 0) : x_1, x_2 \in \mathbb{R}\}$ does not seem like a reasonable candidate for the tangent plane to the graph of f_2 at 0_2.

Moreover, f_2 is not even continuous at the origin. Indeed if we consider the curve $c(t) = (t, t^3)$ for $t \in \mathbb{R}$, then since for $t \in (0, 1)$ we have $0 < t^3 < t^2$, we see that $\lim_{t\downarrow 0} f_2(c(t)) = 1$, while $\lim_{t\downarrow 0} f_2(t.v) = 0$ for all $v \in \mathbb{R}^2, v \neq 0_2$. This example suggests one way in which our considerations so far might be deficient: In one dimension there are only two ways to approach a point (from the left or the right), however, even in two dimensions, there are infinitely many different curves through which one can approach a point, and moreover many more than simply by travelling along a straight line – focusing on directional derivatives therefore does an injustice to the geometry of linear spaces of dimension greater than 1.

This issue can be resolved easily however, in that it was already addressed in the Metric Spaces material of A0: if $f\colon X \to \mathbb{R}$ is a real-valued function on a metric space, then for $f(x)$ to tend to a limit α as $x \to a \in X$, the values of f must be close to α for all x sufficiently close to a. There is simply no need to specify a curve on which x lies as it tends to a. In order to be able to use this idea however, we need to rewrite the expression we have for a directional derivative in a way which only uses the norm functions. Let us do this first in the one-dimensional case: the condition that $Df_a(v)$ is given by the directional derivative as

$$\lim_{t\to 0}\frac{f(a+tv)-f(a)-Df_a(tv)}{t}=0_Y \iff \lim_{t\to 0}\frac{1}{|t|}\|f(a+t.v)-f(a)-Df_a(t.v)\|=0 \quad \forall v\in X, v\neq 0.$$

Notice that this formulation does not utilise the norm on X. This is however a relic of the Prelims definition we started with: by the homogeneity of directional derivatives, we may assume $\|v\|=1$, and then if we let $x=a+t.v\in X$, then $\|x-a\|=|t|$, and the above condition becomes

$$\lim_{x\to a}\frac{\|f(x)-f(a)-Df_a(x-a)\|}{\|x-a\|}\to 0 \tag{2.3}$$

But it makes sense to ask for the same limit to hold for any $f\colon U\to Y$ defined on an open subset $U\subseteq X$ taking values in Y, where X and Y are normed vector spaces, and this (finally!) gives us the definition of the derivative in higher dimensional that we will use:

Definition 2.7. Let X and Y be finite-dimensional normed vector spaces and let $U\subseteq X$ be an open subset of X. If $f\colon U\to Y$ is a function and $a\in U$, we say that f is *differentiable at a* if there is a linear map $T\in\mathcal{L}(X,Y)$ such that if the function $\epsilon\colon U\to Y$ given by $\epsilon(a)=0$ and, for $x\in U\backslash\{a\}$ by the equation

$$f(x)=f(a)+T(x-a)+\|x-a\|.\epsilon(x),$$

then ϵ is continuous at a, that is $\lim_{x\to a}\epsilon(x)=0_Y=\epsilon(a)$. If such a map T exists, it is unique and we denote it by Df_a.[7]

Remark 2.8. This definition takes some time to absorb!

1. Note that for $x\neq a$,

   $$\epsilon(x)=\frac{f(x)-f(a)-T(x-a)}{\|x-a\|}$$

   so that the continuity of ϵ at a is precisely the condition of Equation (2.3).

2. The function f_2 from Example 2.6 is not differentiable at $a=0_2$ in the above sense. Indeed because all of the directional derivatives of f_2 exist and equal 0, the only candidate for $Df_{2,a}$ is the zero linear map. But since 0_2 lies in the closure of Ω, we have $|f_2(x)-f_2(0_2)|=1$ for x arbitrarily close to 0_2, and so $|f(x)-f(0_2)|/\|x\|$ is unbounded near 0_2, hence the zero linear map fails to satisfy the requirement of Definition 2.7. In particular, it is important to note that Definition 2.7 requires more than the existence of all directional derivatives.

3. As the previous point notes, the linear map Df_a is unique if it exists, because its values are given by the directional derivatives, which are certainly unique (again, assuming they exist). One can also prove the uniqueness of the linear map Df_a directly, and the problem set asks you to do this.

4. One can write the condition required of the linear map Df_a using the little o notation, that is, as $f(a+h)=f(a)+Df_a(h)+o(\|h\|)$, where $h=x-a$.

5. If U is an open subset of $\mathbb{R}^n$ and $f\colon U\to\mathbb{R}^m$, then if $f=(f_1,\dots,f_m)$, then, as promised in the discussion of the definition of differentiability, f is differentiable at $a\in U$ if and only if each f_i is, and $Df_a=\sum_{i=1}^m Df_{i,a}.e_i$, that is, if $v\in\mathbb{R}^n$, we have $Df_a(v)=\sum_{i=1}^m Df_{i,a}(v).e_i$. This can be checked directly, and is in essence a very special case of the multi-variable version of the *Chain Rule*, which we will prove shortly.

6. It is straight-forward to check that equivalent norms yield the same condition for a function to be differentiable, since they give the same notion of convergence. Since all norms on finite-dimensional vector space are equivalent, it follows that the definition of the derivative is independent of the choice of norms on X and Y when both X and Y are finite-dimensional.

   [**Non-examinable: Since norms on an infinite-dimensional space need not be equivalent however, in the infinite-dimensional setting, the notion of differentiability may depend on the norm. Moreover, in the infinite-dimensional setting, the total derivative Df_a is required to be a bounded linear map, a condition which, by Corollary 1.17, is automatic in the finite-dimensional setting.*]

[7] The total derivative in this sense is sometimes called the *Fréchet derivative*.

7. If $f: U \to Y$ is differentiable on U, then it defines a function $Df: U \to \mathcal{L}(X, Y)$. Viewed as a function "taking values in (linear) functions" it appears to be a more complicated object than the original function f. However, $\mathcal{L}(X, Y)$ is just a $\dim(X).\dim(Y)$-dimensional normed vector space – using the operator norm $\|.\|_\infty$ – and if we pick a basis of X and Y then we can identify it with $\mathrm{Mat}_{m,n}(\mathbb{R})$. Thus, at least in principle, Df is no more complicated an object than f. We discuss this in more detail in Section 2.8.

As in the one-variable case, if f is differentiable at a point a, then it is continuous there:

Lemma 2.9. *Let X and Y be normed vector spaces and let U be an open subset of X. If $f: U \to Y$ is a function which is differentiable at $a \in U$, then there are constants $C, r > 0$ such that for all $x \in B(a, r)$,*

$$\|f(x) - f(a)\| \leq C.\|x - a\|.$$

In particular, f is continuous at a.

Proof. Replacing $f(x)$ with the function $f(x-a)-f(a)$ we may assume that $a = 0_X$ and $f(a) = 0_Y$. The statement of the Lemma is then simply that if f is differentiable at 0_X then $f \in O_Y(\|x\|)$. But $f(x) = Df_{0_X}(x) + o_Y(\|x\|)$, and since Df_{0_X} is a bounded linear map it lies in $O_Y(\|x\|)$, while $o_Y(\|x\|)$ is a subspace of $O_Y(\|x\|)$, hence $f(x) \in O_Y(\|x\|)$ as required. □

Definition 2.10. If X and Y are normed vector spaces and U is an open subset of X, then we write $C^0(U, Y)$ for the vector space of continuous functions on U taking values in Y. The previous Lemma thus shows that if $f: U \to Y$ is differentiable on all of U then $f \in C^0(U, Y)$.

Example 2.11. Constant functions $c: X \to Y$ are clearly differentiable, with derivative 0, since if c is constant $c(x) = c(a)$. If $T: X \to Y$ is linear, that is $T \in \mathcal{L}(X, Y)$, then, for any $a \in X$ we have $Df_a = T$, since

$$T(x) = T(a) + T(x - a),$$

(and thus the error term $\epsilon(x).\|x\|$ is identically zero). Thus if $f = T$ is linear, $Df: X \to \mathcal{L}(X, Y)$ is the constant function $x \mapsto T$, for all $x \in U$.

If U is an open subset of X and $f, g: U \to Y$ are differentiable at a point $a \in U$ then it is easy to see that $f + g$ is also, and $D(f + g)_a = Df_a + Dg_a$. In particular, if $f(x) = T(x) + b$, where $T \in \mathcal{L}(X, Y)$ and $b \in Y$, then f is differentiable with $Df_a = T$ for all $a \in U$.

Example 2.12. If $\|.\|$ is a norm on $\mathbb{R}^n$, we may view it as a function $\|.\|: \mathbb{R}^n \to \mathbb{R}$. This function is *not* differentiable at the origin: Indeed suppose that T is a linear map. Then $\epsilon(h) = \|h\|^{-1}(\|h\| - T(h)) = 1 - T(h/\|h\|)$, and since $T(h/\|h\|)$ is independent of $\|h\|$, if $\epsilon(h) \to 0$ as $\|h\| \to 0$ we must have $T(h/\|h\|) = 1$. But since $T(-h/\| - h\|) = -T(h/\|h\|)$ this is impossible.

The question of whether a norm is differentiable at other points in $\mathbb{R}^n$ may depend on the norm – consider for example the norms $\|.\|_1$, $\|.\|_2$ and $\|.\|_\infty$.

2.3 Partial derivatives and the total derivative

We now relate the notion of the total derivative to the notion of partial derivatives which were introduced in Prelims multivariable calculus:

In fact we work in slightly greater generality, as it clarifies the idea and reduces the notational clutter.

Definition 2.13. Suppose that X and Y are normed vector spaces and $U \subseteq X$ is an open subset with $f: U \to Y$ a function defined on U. If we are further given a subspace Z of X, then we can consider the function $f_{a,Z}: Z \to Y$ given by $f_{a,Z}(x) = f(a + z)$, and we set $\partial_Z f(a) = Df_{a,Z}(0_Z)$, so that $\partial_Z f(a)$ satisfies

$$\frac{\|f(a + z) - f(a) - \partial_Z f(a)(z)\|}{\|z\|} \to 0, \text{ as } z \to 0, \quad (z \in Z).$$

It is immediate from the definitions that, if the total derivative $Df(a)$ exists, then $Df(a)_{|Z} = \partial_Z f(a)$. Similarly, the values of the partial derivative $\partial_Z f(a) \in \mathcal{L}(Z, Y)$, like the total derivative, are given by the corresponding directional derivatives of f, so it is unique if it exists.

If we have a decomposition of X into a direct sum $X = X_1 \oplus X_2$, then the partial derivatives $\partial_{X_1} f(a)$ and $\partial_{X_2}(f)(a)$ determine $Df(a)$: if $\pi_1 \colon X \to X_1$ and $\pi_2 \colon X \to X_2$ denote the projection maps from X to X_1 and X_2 respectively, and $\iota_1 \colon X_1 \to X, \iota_2 \colon X_2 \to X$ denote the inclusion maps, then

$$I_X = \iota_1 \circ \pi_1 + \iota_2 \circ \pi_2,$$

where I_X denotes the identity map from X to itself. Thus, noting that $Df(a)_{|X_j} = Df(a) \circ \iota_j$ ($j \in \{1,2\}$), we have

$$\begin{aligned} Df(a) &= Df(a) \circ I_X = Df(a) \circ (\iota_1\pi_1 + \iota_2\pi_2) = (Df(a)\iota_1) \circ \pi_1 + (Df(a)\iota_2) \circ \pi_2 \\ &= Df(a)_{|X_1} \circ \pi_1 + Df(a)_{|X_2} \circ \pi_2 \end{aligned}$$

and hence

$$Df(a) = \partial_{X_1} f(a) \circ \pi_1 + \partial_{X_2} f(a) \circ \pi_2 \tag{2.4}$$

Remark 2.14. Obviously, in the same way, if we have any direct sum decomposition $X = X_1 \oplus X_2 \oplus \ldots \oplus X_k$ of X, the partial derivatives $\partial_{X_j} f(a)$ determine $Df(a)$, where if $\pi_j \colon X \to X_j$ denotes the projection map to the j-th summand X_j,

$$Df(a) = \sum_{j=1}^{k} \partial_{X_j} f(a) \circ \pi_j. \tag{2.5}$$

Motivated by matrix notation, we will sometimes write $Df_a = (\partial_{X_1} f(a) \mid \partial_{X_2} f(a))$ to express the decomposition of $Df(a)$ according to the direct sum decomposition $\mathcal{L}(X, Y) = \mathcal{L}(X_1, Y) \oplus \mathcal{L}(X_2, Y)$.d

2.3.1 Partial derivatives in multivariable calculus

In multivariable calculus, the term "partial derivative" usually refers to the directional derivatives of a function in the directions given by a choice of basis of X. This is essentially a special case of the above setting, as we now explain: Let $B_X = \{v_1, v_2, \ldots, v_n\}$ be a basis of X, and let $X_j = \mathbb{R}.v_j$ denote the line spanned by v_j ($1 \leq j \leq n$). We thus obtain a direct sum decomposition $X = X_1 \oplus \ldots X_n$ of X into n lines, *i.e.*, 1-dimensional subspaces.

Applying (2.5) to this decomposition, we see that $Df(a) = \sum_{j=1}^{n} \partial_{X_j} f(a) \circ \pi_j$. But if $B_X^* = \{x_1, \ldots, x_n\}$, so that if $u \in X$, we have $u = \sum_{j=1}^{n} x_j(u).v_j$, and hence $\pi_j(u) = x_j(u).v_j$. Thus

$$\partial_{X_j} f(a)\pi_j(u) = \partial_{X_j} f(a)(x_j(u)v_j) = x_j(u).\partial_{X_i} f(a)(v_j) = x_j(u)Df_a(v_j) = x_j(u)\partial_{v_j} f(a).$$

Thus the directional derivative $\partial_{v_j} f(a)$ completely determines $\partial_{X_j} f(a) \in \mathcal{L}(X_j, Y)$.

Definition 2.15. If we are given a basis $B = \{v_1, \ldots, v_n\}$ of X, then we will write

$$\partial_j f(a) = \frac{\partial f}{\partial x_j}(a) := \partial_{v_j} f(a) = \lim_{t \to 0} \frac{f(a + tv_j) - f(a)}{t} \in Y$$

The fractional notation $\frac{\partial f}{\partial x_j}$ is commonplace, but becomes cumbersome when considering higher-order partial derivatives. We will normally prefer to write $\partial_j f$.

Using this notation, the expression for the total derivative becomes

$$Df(a) = \sum_{j=1}^{n} x_j.\partial_j f(a) = \sum_{j=1}^{n} x_j.\frac{\partial f}{\partial x_j}(a). \tag{2.6}$$

We may refine this further if we pick a basis $B_Y = \{w_1, \ldots, w_m\}$ of Y: Using B_Y we may write $f(x) = \sum_{i=1}^{m} f_i(x).w_i$ where $f_i \colon U \to \mathbb{R}$, and hence we have $Df(a) = \sum_{i=1}^{m} Df_i(a).w_i$. Applying (2.6) to each Df_i and summing we obtain

$$Df(a) = \sum_{i=1}^{m} \left(\sum_{j=1}^{n} \partial_j f_i(a).x_j \right) w_i = \sum_{i=1}^{m} \left(\sum_{j=1}^{n} \frac{\partial f_i}{\partial x_j}(a).x_j \right) w_i \tag{2.7}$$

Notice that this last equation shows that the matrix of $Df(a)$ with respect to the bases B_X of X and B_Y of Y is just

$$ {}_{B_Y}[Df_a]_{B_X} = \begin{pmatrix} \partial_1 f_1(a) & \dots & \partial_n f_1(a) \\ \vdots & \ddots & \vdots \\ \partial_1 f_m(a) & \dots & \partial_n f_m(a) \end{pmatrix} $$

Thus, if we know the derivative exists, then we can compute its matrix with respect to a choice of bases of X and Y by computing the directional derivatives of the components of f along the directions given by the basis in X.

Definition 2.16. As in multi-variable calculus, the above matrix $\big(\partial_j f_i(a)\big)$ is called the *Jacobian matrix* of the partial derivatives of f at a. Note that the determinant $\det(Df_a) = \det(\partial_j f_i(a))$, is also often called the Jacobian. We will refer to it as the *Jacobian determinant*. It is often denoted $J_f(a)$.

Remark 2.17. In a similar way, if $X = X_1 \oplus X_2$, the partial derivative $\partial_{X_j} f(a)$ are given by block submatrices of the Jacobian matrix, and if you like, you can think of them as essentially just a notational shorthand for such submatrices. Indeed as we already noted above, if Df_a exists then $\partial_{X_j} f(a)$ is just the restriction of Df_a to X_1. But if our basis $B_X = \{v_1, \dots, v_n\}$ is adapted to this direct sum decomposition, so that for some k, $1 \leq k \leq n$, the subsets $B_1 = \{v_1, \dots, v_k\}$ and $B_2 = \{v_{k+1}, \dots, v_n\}$ are bases of X_1 and X_2 respectively, then

$$ {}_{B_Y}[Df_a]_{B_X} = \begin{pmatrix} {}_{B_Y}[\partial_{X_1} f(a)]_{B_1} & {}_{B_Y}[\partial_{X_2} f(a)]_{B_2} \end{pmatrix} $$

Example 2.18. If U is an open subset of $\mathbb{C}$ and $f\colon U \to \mathbb{C}$ is holomorphic, then, identifying $\mathbb{C}$ with $\mathbb{R}^2$ via $z \mapsto (\Re(z), \Im(z))$, we may view f as a function from $\mathbb{R}^2$ to itself, which, for clarity, we write as F. Since complex multiplication is $\mathbb{R}$-linear, F is differentiable in the real sense: explicitly, if $f'(z) = a + ib$ then the total derivative of F at z is the $\mathbb{R}$-linear map given by multiplication by $f'(z)$, and hence its matrix is

$$ DF_{z=(x,y)} = \begin{pmatrix} a & -b \\ b & a \end{pmatrix} $$

The Cauchy-Riemann equations follow immediately from this – they express the fact that the linear map given by the derivative is complex-linear rather than just real-linear, and so is given by multiplication by a complex number.

Remark 2.19. Example 2.6 shows that the existence of all the partial derivatives for the function $f_2\colon \mathbb{R}^2 \to \mathbb{R}$ at the origin 0 is not sufficient to ensure that f_2 is continuous at that point. Since Lemma 2.9 shows that the existence of the total derivative at a point implies continuity at that point, this gives another way of seeing that f_2 is not differentiable at the origin. The function $f_1\colon \mathbb{R}^2 \to \mathbb{R}$ in the same Example is continuous at the origin, but nevertheless, even though all of its directional derivatives exist at the origin, it is not differentiable there. (The first problem sheet asks you to check this).

We will see shortly, however, that if the partial derivatives exist and are continuous, then this is sufficient to show that the total derivative exists.

2.4 The Chain Rule

One of the fundamental properties of the differentiablity is that it is preserved under composition, just like continuity. The single variable version of this result is both a basic computational tool, and also the key to one version of the Fundamental Theorem of Calculus. We now establish its higher-dimensional analogue.

Theorem 2.20. *Let X, Y and Z be normed vector spaces, let $f\colon U_1 \to Y$ be a function defined on an open subset U_1 of X, and let $g\colon U_2 \to Z$ be a function defined on an open subset U_2 of Y. Suppose that $a \in U_1$ and $f(a) = b \in U_2$, then if f is differentiable at a and g is differentiable at b, their composition $h = g \circ f\colon f^{-1}(U_2) \to Z$ is differentiable at a and its derivative is given by*

$$ Dh_a = Dg_{f(a)} \circ Df_a. $$

Proof. Note that since f is differentiable at a, it is continuous there, and hence $f^{-1}(U_2)$ is a neighbourhood of a, hence it makes sense to ask if h is differentiable at a. By translating if necessary, we may assume that $a = 0_X$ and $f(a) = b = 0_Y$. To avoid cluttered notation, we will write 0 for the zero vector in all vector spaces in the rest of this proof.

Since f is differentiable at 0 we see that $f(x) = Df_0(x) + \epsilon_1(x)$ where $\epsilon_1(x) \in o_Y(\|x\|)$. Similarly since g is differentiable at $f(0) = 0$, we have $g(y) = Dg_0(y) + \epsilon_2(y)$, where $\epsilon_2(y) \in o_Z(\|y\|)$. It follows that

$$g \circ f(x) = Dg_0(Df_0(x)) + Dg_0(\epsilon_1(x)) + \epsilon_2(f(x)).$$

Thus to complete the proof, we must show that $Dg_0(\epsilon_1(x)) + \epsilon_2(f(x)) \in o_Z(\|x\|)$, which certainly follows if each summand lies in $o_Z(\|x\|)$. But since the linear map Dg_0 is bounded and $\epsilon_1(x) \in o_Y(\|x\|)$,

$$\frac{\|Dg_0(\epsilon_1(x))\|}{\|x\|} \leq \|Dg_0\|.\frac{\|\epsilon_1(x)\|}{\|x\|} \to 0, \quad \text{as } x \to 0.$$

hence $Dg_0(\epsilon_1(x)) \in o_Z(\|x\|)$. For the second term, recall that we may write $\epsilon_2(y) = \|y\|.\eta(y)$ where $\eta(y) \to 0 = \eta(0)$ as $y \to 0$. Then

$$\frac{\|\epsilon_2(f(x))\|}{\|x\|} = \frac{\|f(x)\|}{\|x\|}.\|\eta(f(x))\|$$

But now since f is differentiable at 0, we have $f \in O(\|x\|)$, hence the ratio $\|f(x)\|/\|x\|$ is bounded as $x \to 0$, hence it suffices to show that $\eta(f(x)) \to 0$ as $x \to 0$. But by definition $\eta(y) \to 0$ as $y \to 0$, thus we need only check $f(x) \to 0 = f(0)$ as $x \to 0$, but this again follows from $f \in O(\|x\|)$ (see Lemma 2.9) and so we are done.

□

Remark 2.21. It is worth noticing that this is almost word-for-word the proof in the single-variable case. The only difference lies in the fact that in higher dimensions we can only bound the ratio of norms $\|f(x)-f(a)\|/\|x-a\|$, whereas in the single-variable case, the ratio $(f(x) - f(a))/(x - a)$ of course converges to $f'(a)$.

2.5 The Mean Value Inequality

For functions of a single variable, the Mean Value Theorem asserts that, if $f\colon U \to \mathbb{R}$ is differentiable on an open subset U of $\mathbb{R}$ and $[a, b] \subset U$, then $(f(b) - f(a))/(b - a)$, the slope of the chord between $(a, f(a))$ and $(b, f(b))$, is equal to $f'(c)$ for some $c \in (a, b)$. In higher dimensions, as we have noted before, we can only divide by scalars, and so to obtain a statement which at least is syntactically correct, we can rewrite this as $f(b)-f(a) = f'(c).(b-a)$. There is however a more fundamental issue here: Namely the condition that c lies "between a and b", that is, $c \in (a, b)$, is not a meaningful one in higher dimensions: two points in an open subset U of $\mathbb{R}^n$ do not bound any region in U. One consequence of this is that the most naive attempt to generalize the Mean Value Theorem to arbitrary dimensions is simply false:

Example 2.22. Let $f\colon \mathbb{R}^1 \to \mathbb{R}^2$ be given by $f(t) = (\cos(2\pi t), \sin(2\pi t))$. Then the derivative of f is $f'(t) = 2\pi(-\sin(2\pi t), \cos(2\pi t))$, which is non-zero for all t. But if we take $a = 0$ and $b = 1$ then $f(b) - f(a) = 0$, while for any $t_0 \in [0, 1]$ we have $(2\pi - 0)f'(t_0) = 4\pi^2(-\sin(2\pi t), \cos(2\pi t)) \neq 0$.

Example 2.22 also suggests what the reason for the failure of the naive attempt at a generalisation of the Mean Value Theorem: Notice that $f'(t) = 2\pi(-\sin(2\pi t), \cos(2\pi t))$, and so by the Fundamental Theorem of Calculus[8] we have

$$f(1) - f(0) = \int_0^1 f'(t)dt = 2\pi\Big(\int_0^1 -\sin(2\pi t)dt, \int_0^1 \cos(2\pi t)dt\ \Big) = (0, 0).$$

Thus it is still true that $f(1)-f(0)$ is the *average* value of $f'(t)$ over the interval $[0, 1]$, it is just that this average value is not the value of $f'(t)$ for any $t \in [0, 1]$. This suggests that it should be possible to bound $\|f(b) - f(a)\|$ relative to $|b - a|$ by bounding $\|Df_t\|_\infty$, that is, we will prove a Mean Value *Inequality* rather than an equality.

[8] One can define the integral of a function $f\colon [0, 1] \to X$ where X is a finite-dimensional normed vector space by picking a basis and integrating componentwise. The resulting integral does not depend on the choice of basis made.

Definition 2.23. If X is a normed vector space and $a, b \in X$ we write $\gamma_{a,b}\colon [0,1] \to X$ for the line-segment path $\gamma_{a,b}(t) = (1-t)a + tb$, and write $[\gamma_{a,b}]$ for its image, that is $[\gamma_{a,b}] = \{\gamma_{a,b}(t) : t \in [0,1]\}$.

Recall that a subset C of X is *convex* if, for any $a, b \in C$ we have $[\gamma_{a,b}] \subseteq C$.

Theorem 2.24. *(Mean Value Inequality.) Let X and Y be finite-dimensional normed vector spaces and let $U \subset X$ be an open subset. Suppose that $f\colon U \to Y$ is differentiable, and $z_1, z_2 \in U$ are such that the image of γ_{z_1,z_2} lies entirely in U. Then there is some $c \in [\gamma_{z_1,z_2}]$ such that*

$$\|f(z_2) - f(z_1)\| \leq \|Df_c(z_2 - z_1)\|.$$

In particular, if U is convex and $\|Df_x\|_\infty \leq K$ for all $x \in U$ then $\|f(x) - f(y)\| \leq K.\|x - y\|$ for all $x, y \in U$, that is, f is Lipchitz continuous with constant K.

Proof. We give a proof only in the case where X is an inner product space. If $f(z_1) = f(z_2)$, we may choose c arbitrarily, so we may assume $f(z_1) \neq f(z_2)$. Let $e = \|f(z_2) - f(z_1)\|^{-1}.(f(z_2) - f(z_1))$. Define

$$g(x) = \langle e, f(x) - f(z_1)\rangle,$$

so that $g(z_1) = 0$ and $g(z_2) = \|f(z_2) - f(z_1)\|$. Now if we let $G(t) = g(\gamma_{z_1,z_2}(t))$, we see that $G\colon [0,1] \to \mathbb{R}$ is a real-valued function on $[0,1]$, satisfying $G(1) - G(0) = \|f(z_1) - f(z_1)\|$. Applying the Mean Value Theorem for a single variable shows that there is some $\xi \in (0,1)$ such that

$$\|f(z_2) - f(z_1)\| = G(1) - G(0) = G'(\xi) = Dg_{\gamma_{z_1,z_2}(\xi)}(\gamma'_{z_1,z_2}(\xi))$$

But $Dg_z(v) = \langle e, Df_z(v)\rangle$, and $\gamma'_{z_1,z_2}(t) = (z_2 - z_1)$, hence if we let $c = \gamma_{z_1,z_2}(\xi)$, the right-hand side of the previous equality is just $\langle e, Df_c(z_2 - z_2)\rangle$, and so by the Cauchy-Schwarz inequality we see

$$\|f(z_2) - f(z_1)\| \leq \|e\|.\|Df_c(z_2 - z_1)\| = \|Df_c(z_2 - z_1)\|,$$

as required.

For the final part, since $\|Df_c(z_2 - z_1)\| \leq \|Df_c\|_\infty.\|z_2 - z_1\|$, if U is convex and $\|Df_x\|_\infty \leq K$ for all $x \in U$, we may apply the first part of the Theorem to any $x, y \in U$ and the above inequalities to see that f is Lipschitz with constant K on U. □

Remark 2.25. The reason the above proof needs to assume X is an inner product space is so that we can identify $(X^*, \|.\|_\infty)$ as a normed vector space with $(X, \|.\|)$ via the map $v \mapsto [x \mapsto \langle x, v\rangle]$. If X is an arbitrary normed vector space, given $v = f(z_2) - f(z_1)$, one needs to find $\delta \in X^*$ such that $\delta(f(z_2) - f(z_1)) = \|f(z_2) - f(z_1)\|$ and $\|\delta\|_\infty = 1$. Clearly, if e is as in the proof above, we want $\delta(e) = 1$, but then one needs to show that this functional on $\mathbb{R}.e$ can be extended to all of X without increasing its operator norm. This is possible, and the required result is proved in Appendix 5.5.

Any easy application of this result is the following:

Proposition 2.26. *Suppose that U is a connected open subset of $\mathbb{R}^n$ and $f\colon U \to \mathbb{R}^m$. Then if $Df_x = 0$ for all $x \in U$ the function f is constant.*

Proof. Since U is open and connected in $\mathbb{R}^n$, it is path connected, and in fact any two points can be joined by piecewise-linear path. But if $\gamma_{a,b}$ is a line-segment path whose image lies in U then Proposition 2.24 and the hypothesis $Df = 0$ on U shows that $f(b) = f(a)$. It follows immediately that f must be constant on U as required. □

2.6 Continuity of partial derivatives and the existence of the total derivative

The next result shows that however that the existence *and continuity* of the partial derivatives give a sufficient condition for the total derivative to exist.

First note that, if $X = X_1 \oplus X_2$ and $\pi_1\colon X \to X_1$ and $\pi_2\colon X \to X_2$ denote the corresponding projections, then

$$\|x\|_d := \|\pi_1(x)\| + \|\pi_2(x)\|,$$

is a norm on X. Indeed the triangle inequality follows from the linearity of the projection maps and the triangle inequality for $\|.\|$, the original norm on X, and the positivity follows in the same way. Now since for any $x \in X$ we have $x = \pi_1(x) + \pi_2(x)$, the triangle inequality shows that $\|x\| \leq \|x\|_d$. On the other hand, we have $\|x\|_d = \|\pi_1(x)\| + \|\pi_2(x)\| \leq (\|\pi_1\|_\infty + \|\pi_2\|_\infty).\|x\|$, so that $\|.\|$ and $\|.\|_d$ are equivalent norms on X. (Of course, all norms on a finite dimensional vector space are equivalent, but this discussion gives more precise information on the relationship between the two norms.)

The above discussion shows that, if we are given a decomposition of X into a direct sum $X = X_1 \oplus X_2$, then, replacing the norm $\|.\|$ on X by the equivalent norm $\|.\|_d$, we may assume that if $x = (x_1, x_2)$ where $x_i \in X_i$, then $\|x\| = \|x_1\| + \|x_2\|$.

Example 2.27. If X is an inner product space, then if X_1 is a subspace, it has a natural complement given by $X_2 = X_1^\perp = \{v \in X : \langle v, x\rangle = 0, \quad \forall x \in X_1\}$. If π_1, π_2 denote the projection maps to X_1 and X_2 respectively, then

$$\begin{aligned}\|x\|^2 &= \langle x, x\rangle = \langle \pi_1(x) + \pi_2(x), \pi_1(x) + \pi_2(x)\rangle \\ &= \langle \pi_1(x), \pi_1(x)\rangle + \langle \pi_2(x), \pi_2(x)\rangle = \|\pi_1(x)\|^2 + \|\pi_2(x)\|^2,\end{aligned}$$

hence in this case (*c.f.* Example 1.14), $\|\pi_1(x)\| + \|\pi_2(x)\| \leq \sqrt{2}\|x\|$.

Theorem 2.28. *Let X and Y be finite-dimensional normed vector spaces and suppose that $f\colon U \to Y$ is a function defined on an open subset of X. Suppose that $X = X_1 \oplus X_2$, and that the partial derivatives $\partial_{X_1} f(x)$, $\partial_{X_2} f(x)$ both exist for all $x \in U$. Then if for some $a \in U$ both $\partial_{X_1} f(x)$ and $\partial_{X_2} f(x)$ are contiuous at a, then the total derivative of f exists, where necessarily $Df_a = (\partial_{X_1} f(a) \mid \partial_{X_2} f(a))$ and hence Df_a is also continuous at a.*

Proof. Let π_1, π_2 be the projections to X_1 and X_2 respectively. As the statement of the theorem notes, if Df_a exists, it must be given by $\partial_{X_1} f(0) \circ \pi_1 + \partial_{X_2} f(0) \circ \pi_2$, hence replacing $f(x)$ by

$$f_1(x) = f(a + x) - f(a) - \partial_{X_1} f(a) \circ \pi_1 - \partial_{X_2} f(a) \circ \pi_2,$$

we need only consider the case where $a = 0_X$, $f(0_X) = 0_Y$ and $\partial_{X_1} f(0) = 0$ and $\partial_{X_2} f(0) = 0$. Moreover, since the theorem is local, we may replace U by a sufficiently small ball centred at 0_X, and hence we may assume that U is convex.

Given these assumption, to prove the theorem, we must show that $f(x) \in o_Y(\|x\|)$. Now since $\partial_{X_1} f(0) = 0$ it follows that $f(x_1, 0) \in o_Y(\|x_1\|)$, so that, if $\epsilon > 0$ is given, there is some $\delta_1 > 0$ such that if $\|x_1\| < \delta$, then $\|f(x_1, 0)\| < \epsilon.\|x_1\|$.

Moreover, the partial derivative $\partial_{X_2} f(x)$ is continuous at $x = 0$, and $\partial_{X_2} f(0) = 0$, hence there is a $\delta_2 > 0$ such that, for $\|x\| < \delta_2$ we have $\|\partial_{X_2} f(x)\|_\infty < \epsilon$. Thus applying Theorem 2.24, $\|f(x_1, x_2) - f(x_1, 0)\| \leq \epsilon.\|x_2\|$, provided $\|x\| = \|x_1\| + \|x_2\| < \delta_2$ (since then $\|(x_1, t.x_2)\| = \|x_1\| + t.\|x_2\| < \delta_2$ for all $t \in [0, 1]$).

It follows that if $\delta = \min\{\delta_1, \delta_2\}$ and $\|x\| < \delta$, then

$$\begin{aligned}\|f(x_1, x_2)\| = \|f(x_1, 0) + (f(x_1, x_2) - f(x_1, 0)\| &\leq \|f(x_1, 0)\| + \|f(x_1, x_2) - f(x_1, 0)\| \\ &\leq \epsilon\|x_1\| + \epsilon.\|x_2\| = \epsilon.\|x\|\end{aligned}$$

so that $\|f(x_1, x_2)\| = o_Y(\|x\|)$ as required. □

Corollary 2.29. *If $f\colon U \to Y$ is as in the previous theorem, and $B_X = \{v_1, \ldots, v_n\}$ and $B_Y = \{w_1, \ldots, w_m\}$ are bases of X and Y respectively, then if the partial derivatives $\partial_j f_i(x)$ exist on U and are continuous at $a \in U$, the total derivative Df_a exists and is given by the matrix $(\partial_j f_i(a))$ and therefore it is also continuous.*

Proof. Use induction on $\dim(X) = |B_X|$ and the previous Theorem. In more detail, for $n = 1$ the result is trivial. If $\dim(X) > 1$, then write $X = X_1 \oplus X_2$, where $X_1 = \mathrm{Span}\{v_1, \ldots, v_{n-1}\}$ and $X_2 = \mathbb{R}.v_n$. Then since $\dim(X_1) = n - 1 < \dim(X)$, by induction we know that $\partial_{X_1} f(a)$ exists and is continuous at a (recall that the matrix of $\partial_{X_1} f(a)$ is just the submatrix of Df_a given by the first $n - 1$ columns of the Jacobian matrix for Df_a), and since $\partial_{X_2} f(a)$ is given by the final column vector $\partial_n f_i(x)$ it is also continuous at a. We may thus apply the previous theorem to conclude that Df_a exists and has matrix given by the Jacobian matrix of partial derivatives $(\partial_j f_i(a))$ as required. □

Remark 2.30. Note that in fact the proof of Corollary 2.29 doesn't in fact need the full strength of the hypothesis of the theorem – we assumed the existence and continuity of all of the partial derivatives of f at a, but it sufficed to know the continuity for all but one of them to conclude that f is real-differentiable at a (as one might suspect considering the case $n = 1$ of course!) In practice however, this weaker hypothesis is rarely useful.

Definition 2.31. If X and Y are finite dimensional normed vector spaces and U is an open subset of X then if $f\colon U \to Y$, we say that f is *continuously differentiable* if[9] $Df\colon U \to \mathcal{L}(X,Y)$ is continuous. This is equivalent to requiring the continuity of all of the partial derivatives $\partial_j f_i$, where $f = (f_1, \dots, f_m)$ and $1 \leq j \leq n$, $1 \leq i \leq m$. We will write $C^1(U,Y)$ for the vector space of continuously differentiable functions on U taking values in Y.

***Remark 2.32.** If $f\colon U \to Y$ and $a \in U$, we say that f is *strongly differentiable* at a if there is a linear map $T \in \mathcal{L}(X,Y)$ such that, for any $\epsilon > 0$ there is a $\delta > 0$

$$\|f(x) - f(y) - T(x-y)\| \leq \epsilon\|x-y\|, \quad \forall x, y \in B(a,\delta).$$

Equivalently, $\lim_{x,y\to a} \|f(x) - f(y) - T(x-y)\|/\|x-y\| = 0$. The linear map T is then the *strong total derivative* of f at a. Taking $y = a$ one sees immediately that if the strong total derivative exists, then f is differentiable and the total derivative is equal to T. On the other hand, a function which is differentiable at a point need not be strongly differentiable there.

Modifying the proof of Theorem 2.29 by applying the same technique used for $\partial_{X_2} f$ to $\partial_{X_1} f$ as well, one can show that if X and Y are finite-dimensional and the partial derivatives of $f\colon U \to Y$ exist in a neighbourhood of $a \in U$ and are continuous at a, then f is strongly differentiable at a.

2.7 Real-valued functions on an inner product space

Let E be a normed finite-dimensional vector space. (If you prefer you can take E to be $\mathbb{R}^n$, the reason we do not do that here is to try and make clearer what structures are being used where).

If $U \subseteq E$ is an open subset and $f\colon E \to \mathbb{R}$ is differentiable on U, then its derivative Df takes values in $E^* = \mathcal{L}(E,\mathbb{R})$. If the norm on E comes from an inner product $(v,w) \mapsto v \cdot w$ however, we can use it to identify E and E^* via the map $\delta\colon E \to E^*$, where $\delta(a)(v) = a \cdot v$ for all $a, v \in E$.

Definition 2.33. If $f\colon U \to \mathbb{R}$ is differentiable on U then we define $\nabla f\colon U \to E$ to be the *gradient vector field* of f, where $\nabla f(a) = \delta^{-1}(Df_a)$. Thus $\nabla f(a)$ is characterized by the property that

$$Df_a(v) = \nabla f(a) \cdot v, \quad \forall v \in E.$$

Example 2.34. If we take $E = \mathbb{R}^n$, with the standard dot product, then we may view Df_a as a row vector, with entries $\partial_i f(a)$. The vector field $\nabla f(a)$ is then just the corresponding column vector.

$\nabla f(a)$ points in the direction of greatest change for f. More precisely, if $v \in E$ is a direction vector with norm 1, the directional derivative at a of f in the direction v is

$$\partial_v f(a) = Df_a(v) = \nabla f(a) \cdot v.$$

By the Cauchy-Schwarz inequality, $|\nabla f(a) \cdot v| \leq \|\nabla f(a)\|.\|v\| = \|\nabla f(a)\|$, with equality if and only if v and $\nabla f(a)$ are in the same direction. Thus the magnitude of the directional derivative of f at a is maximized when v is in the direction of $\nabla f(a)$.

Another important observation about the gradient vector field is that it is a *normal vector* to the level sets of f, that is, in a suitable sense, it is perpendicular to the level sets of f: If $\gamma\colon (-1,1) \to \mathbb{R}^n$ is a curve such that $f(\gamma(t)) = c$ for some constant $c \in \mathbb{R}$, and $p = \gamma(0)$, the gradient ∇f_p is perpendicular to $\gamma'(0)$, the "velocity vector" of γ at p, because, for all $t \in (-1,1)$ we have $g(t) = f(\gamma(t)) = c$, hence by Theorem 2.20:

$$0 = \frac{dg}{dt}_{t=0} = Df_{\gamma(0)}(\gamma'(0)) = \nabla f(p).\gamma'(0) = 0.$$

We will explore this in more detail when we discuss tangent spaces.

[9] Since, as X is finite-dimensional, $\mathcal{L}(X,Y) = \mathcal{B}(X,Y)$ is a normed vector space, it makes sense to ask if $Df\colon U \to \mathcal{L}(X,Y)$ is continuous.

2.8 *Higher order derivatives

We briefly wish to discuss the notion of higher derivatives for functions $f: U \to Y$, where as before, the domain of f is an open subset U of a normed vector space X and its codomain is a normed vector space Y. There are two ways of thinking about these, the first of which takes bases and works concretely with partial derivatives, while the second works with the total derivative in a coordinate-free manner.

Given bases $\{v_1, \ldots, v_n\}$ of X and $\{w_1, \ldots, w_m\}$ of Y, we obtain the components f_i of f as $f(x) = \sum_{i=1}^{m} f_j(x).w_j$, and then the directional derivatives in the direction of the v_js give the partial derivatives $\partial_j f_i$. But these are just real-valued functions on U, and hence we can consider all of their partial derivatives $\partial_{j_1}\partial_{j_2} f_i$, where $j_1, j_2 \in \{1, \ldots, n\}$ and $i \in \{1, \ldots, m\}$. If these all exist and are continuous, we say that f is twice continuously differentiable. Indeed we can proceed inductively and define:

Definition 2.35. If $f: U \to Y$ is as above and $f = \sum_{i=1}^{m} f_i.w_i$ so that the f_i are the components of f, we define that higher partial derivatives of f inductively as follows: If $k = 1$ these are just the partial derivatives $\partial_j f_i$, $(1 \leq j \leq n, 1 \leq i \leq m)$. For $k > 1$, we suppose that by induction we have defined the partial derivatives of order $k-1$, and write them as $\partial_\beta f_i$ where $\beta = (j_1, j_2, \ldots, j_{k-1}) \in \{1, 2, \ldots, n\}^{k-1}$. The k-th partial derivatives of f are indexed by pairs (α, i) where $\alpha \in \{1, 2, \ldots, n\}^k$ and $i \in \{1, 2, \ldots, m\}$, where if $\alpha = (j_1, j_2, \ldots, j_n)$ then setting $\beta = (j_2, \ldots, j_n) \in \{1, 2, \ldots, n\}^{k-1}$ we define

$$\begin{aligned} \partial_\alpha f_i &:= \partial_{j_1}(\partial_\beta f_i) \\ &= \partial_{j_1}\partial_{j_2} \ldots \partial_{j_k} f_i. \end{aligned}$$

We say that f is k-times continuously differentiable, and write $f \in C^k(U, Y)$, if the partial derivatives $\partial_\alpha f_i$ exist and are continuous for all $\alpha \in \{1, \ldots, n\}^k$ and $i \in \{1, \ldots, m\}$. We say that f is *smooth* or *infinitely differentiable* if the partial derivatives of all orders $k \geq 1$ exist, and write $C^\infty(U, Y)$ for the space of smooth functions on U taking values in Y.

Remark 2.36. One unsatisfactory aspect of this approach to the higher derivatives is that we do not get any sense for how to think about the second derivative $D(D(f))$ of f. In the case of the first derivative, the total derivative gives us the description of Df_a as the "best linear approximation" to f near a. In the same way, we gain a more conceptual understanding of the higher derivatives by considering the higher *total derivative* $D(Df)$ of Df. Theorem 2.29 shows that $f \in C^1(U, Y)$ if and only if the total derivative exists and is continuous. The latter condition makes sense because the total derivative Df is a function from U to $\mathcal{L}(X, Y)$, and $\mathcal{L}(X, Y)$ is a normed vector space when equipped with the operator norm $\|.\|_\infty$. By the same token, our definition of the derivative makes sense, and we can ask if $Df: U \to \mathcal{L}(X, Y)$ is (continuously) differentiable! This leads to an alternative definition of $C^2(U, Y)$, namely

$$C^2(U, Y) = \{f: U \to Y : D(Df): U \to \mathcal{L}(X, \mathcal{L}(X, Y)) \text{ exists and is continuous}\}.$$

To see how this relates to our definition using partial derivatives, notice that our choice of bases for X and Y allows us to identify $\mathcal{L}(X, Y)$ with $\mathrm{Mat}_{m,n}(\mathbb{R})$, the space of $m \times n$ matrices[10]. The space $\mathrm{Mat}_{m,n}(\mathbb{R})$ can then be identified with $\mathbb{R}^{mn}$, and the components of Df with respect to this identification are the (first) partial derivatives of f.[11] Theorem 2.29 thus shows that Df is continuously differentiable if and only if all the second partial derivatives exist and are continuous. In this way you can show by induction that the condition the k-th total derivative of f exists and is continuous is equivalent to the condition that all the k-th partial derivatives exist and are continuous.

We still, however, have not given a satisfactory answer to the question of how one should think of the second derivative. with the total derivative approach we see that $D^2 f_a \in \mathcal{L}(X, \mathcal{L}(X, Y))$, that is $D^2 f_a$ is a linear map from X to the space of linear maps from X to Y. Which is a mouthful.

The standard way to deal with this issue is to notice that $\mathcal{L}(X, \mathcal{L}(X, Y))$ can be less painfully thought of as the space of *bilinear maps* from $X \times X$ to Y! The details of this identification are in the Appendices, and we content ourselves here to trying to understand, explicitly, how one sees this for real-valued functions on an open subset of a normed vector space X.

[10] If we associate a matrix to the linear map given by left-multiplication on column vectors, $\mathcal{L}(\mathbb{R}^n, \mathbb{R}^m)$ is identified with the space of matrices with m rows and n columns.

[11] Here we are identifying the directional derivatives $\partial_{E_{ij}}(Df)$ with the partial derivative associated to the subspace $\mathbb{R}.E_{ij}$.

Example 2.37. Let X be an n-dimensional normed vector space, and let $B = \{e_1, \dots, e_n\}$ be a basis for X. Write $B^* = \{x_1, x_2, \dots, x_n\} \subset X^*$ for the corresponding dual basis.

Suppose that U is an open subset of X and $f\colon U \to \mathbb{R}$ is twice differentiable on U. The derivative of f is a function $Df\colon U \to \mathcal{L}(X, \mathbb{R}) = X^*$. Its components with respect to the basis B^* of X^* are just the partial derivatives $\partial_i f$ of f, since if $Df_a = \sum_{j=1}^n c_j(a).x_j$, where $c_j(a) \in \mathbb{R}$, then

$$c_j(a) = Df_a(e_j) = \partial_{e_j} f(a) = \partial_j f(a).$$

and so $Df = \sum_{j=1}^n (\partial_j f) x_j$. But now, as we already noted, the derivative D is a linear map, hence to calculate $D^2 f$ in terms of the second partial derivatives, we simply apply the same reasoning to each component $\partial_i f\colon U \to \mathbb{R}$ of Df: Indeed since the derivative is linear, we have

$$D(Df) = D\left(\sum_{i=1}^n \partial_i f.x_i\right) = \sum_{i=1}^n D(\partial_i f)x_i = \sum_{i=1}^n \left(\sum_{j=1}^n \partial_j(\partial_i f).x_j\right) x_i = \sum_{1\leq i,j\leq n} (\partial_{ji} f).(x_j x_i).$$

In the second equality we use the fact that if $w \in X^*$ and $g\colon U \to \mathbb{R}$, then $D(g.w) = (Dg).w$, which follows, for example, by the chain rule applied to the composition of g with the map $t \mapsto t.w$ (for $t \in \mathbb{R}$). Thus we see that the basis for $\mathcal{L}^2(X, \mathbb{R}) = \mathcal{L}(X, \mathcal{L}(X, \mathbb{R}))$ induced by our choice of basis $\{v_1, \dots, v_n\}$ of V is the set $\{x_j x_i : 1 \leq i, j \leq n\}$, of pairwise products of the dual basis vectors.

It is useful to explicitly describe $x_j.x_i$ as an element of $\mathcal{L}^2(X, \mathbb{R})$: if $v_1 \in X$ then $(x_j.x_i)(v_1)$ should be an element of X^*, and we may obtain one simply by applying x_j to v_1 to obtain $x_j(v_1).x_i$. Explicitly, it is the functional which assigns to a vector $v_2 \in X$ the scalar $x_j(v_1)x_i(v_2)$.

But it is equally reasonable, however, to think of $x_j.x_i$ as a real-valued function of a pair of vectors $(v_1, v_2) \in X \times X$, namely the function $(v_1, v_2) \mapsto x_j(v_1).x_i(v_2)$. From this point of view it is easy to check that $\{x_j.x_i : 1 \leq i, j \leq n\}$ is a basis of the space $\mathcal{M}^2(X, \mathbb{R})$ of *bilinear* maps from $X \times X$ to $\mathbb{R}$, and hence, since it is just a linear combination of the $x_j x_i' s$ we may view $D^2 f_a$ as a bilinear form on $X \times X$ taking values in $\mathbb{R}$. To see this more concretely, if we let $H = (\partial_{ji} f)$ be the *Hessian* matrix of $D^2 f$, and noting that if $u \in X$ then $u = \sum_{i=1}^n x_i(u).e_i$, we see that for any $v, w \in X$

$$D^2 f_a(v)(w) = \sum_{1\leq i,j\leq n} (\partial_{ji} f).[(x_j x_i)(v)](w) \sum_{i,j=1}^n x_j(v)(\partial_{ji} f).x_i(w) = \mathbf{x}(v)^t.H.\mathbf{x}(w)$$

where we write $\mathbf{x}(v)$ for the column vector $(x_1(v), x_2(v), \dots, x_n(v))^t$. Thus we see that the second derivative is just the symmetric bilinear form given by the Hessian (where the symmetry is a consequence of the symmetry of mixed partial derivatives – Appendix 5.2 gives more details on this which are however non-examinable).

3 The Inverse and Implicit Function Theorems

In this chapter we will discuss the theorems which lie at the heart of all the main results of this course.

Lemma 3.1. *Let $\Omega \subset \mathcal{L}(X, Y)$ be the set of invertible linear maps from X to Y. The we have*

1. *The set Ω is open.*
2. *The inverse map $\iota\colon \Omega \to \Omega$ given by $\iota(\alpha) = \alpha^{-1}$ is continuous.*

Proof. The first problem sheet asks you to establish this carefully. If X and Y have different dimensions, then Ω is empty and there is nothing to prove. If they have the same dimension, then there is an isomorphism $\gamma\colon Y \to X$ and it induces a linear map $\gamma_*\colon \mathcal{L}(X, Y) \to \mathcal{L}(X, X)$ given by $\alpha \mapsto \gamma \circ \alpha$. Its inverse is $(\gamma^{-1})_*$ and since in the finite-dimensional setting all linear maps are continuous, it follows that γ_* is a topological isomorphism, so we may assume that $X = Y$. But then Ω forms a group under composition, which acts on itself by left multiplication. Since $\|\alpha_1 \circ \alpha_2\|_\infty \leq \|\alpha_1\|_\infty.\|\alpha_2\|_\infty$, this action is by homeomorphisms, hence it follows that to show that Ω is open, it is enough to check that it is a neighbourhood of I_X. In fact we have $B(I_X, 1) \subseteq \Omega$.

To see this, note that any element of $B(I_X, 1)$ can be written as $I_X - H$ where $\|H\|_\infty < 1$. Now let $s_n(H) = \sum_{k=0}^n H^k$. Then $s_n(H)(I_X - H) = I_X - H^{n+1}$, and since $\|H^{n+1}\|_\infty \leq \|H\|_\infty^{n+1} \to 0$, it follows that, if we can show $s_n(H)$ converges, then its limit $s(H)$ is $(I_X - H)^{-1}$, and so in particular $I_X - H \in \Omega$ as claimed.

But $\mathcal{L}(X, X)$ is complete (since it is finite dimensional) hence it suffices to show that $(s_n(H))_{n\geq 0}$ is a Cauchy sequence. But if $\|H\|_\infty = r < 1$ then for $m < n$ we have

$$\|s_n(H) - s_m(H)\|_\infty = \|\sum_{k=m}^{n-1} H^k\|_\infty \leq \sum_{k=m+1}^{n} \|H^k\|_\infty \leq \frac{r^{m+1}}{1-r},$$

and so since $r^m/(1 - r) \to 0$ as $m \to \infty$ we see that $(s_n(H))_{n\geq 0}$ is Cauchy as required.

Finally, to see that the inversion map ι is continuous on Ω, the left action of Ω on itself can again be used to show that it suffices to check that ι is continuous at I_X. But $\iota(I_X) = I_X$, hence

$$\|\iota(I_X) - \iota(I_X - H)\| = \lim_{n\to\infty} \|s_0(H) - s_n(H)\|_\infty,$$

but we saw above that $\|s_0(H) - s_n(H)\| \leq \|H\|_\infty/(1 - \|H\|_\infty) \to 0$ as $\|H\|_\infty \to 0$, hence ι is continuous at I_X. □

3.1 The Inverse Function Theorem

Theorem 3.2. *Suppose that X and Y are finite-dimensional normed vector spaces, $U \subseteq X$ an open subset, and $f\colon U \to Y$ is a differentiable function. If $a \in U$ is such that Df_a is invertible and Df is continuous at a, then there is an open neighbourhood $U_1 \subseteq U$ of a such that $f_{|U_1}$ is a homeomorphism from U_1 to $V_1 = f(U_1)$ an open neighbourhood of $b = f(a)$. Moreover if $g\colon V_1 \to U_1$ denotes the inverse of f, then g is differentiable with*

$$Dg_y = (Df_{g(y)})^{-1}, \quad \forall y \in V_1.$$

Thus by the Lemma 3.1, Dg is continuous at y whenever Df is continuous at $x = g(y)$. In particular, Dg is continuous at $f(a)$.

Strategy of proof: Since linear maps are their own derivatives, one can replace f with $(Df_a)^{-1} \circ f$ and hence assume $f\colon X \to X$ and $Df_a = I_X$. Moreover, we can further replace f by $f(x + a) - f(a)$ and hence assume $a = f(a) = 0$.

We then write $f(x) = x + \varphi(x)$, so that $\varphi(x)$ measures the difference between f and the identity map. The intuition is then that a function which is a "small perturbation" of the identity should remain invertible. The insight is then that a "small perturbation" should be rigorously interpreted as a contraction mapping! Using the Mean Value Inequality and the continuity of Df at 0_X, one can show that, in $B(0_X, r)$ for small enough r, φ is Lipschitz with a Lipschitz constant less than 1. This ensures f is injective on $B(0_X, r)$ and, by an application of the contraction mapping theorem, that $f(B(0_X, r))$ is a neighbourhood of $0_X = f(0_X)$. It then follows that there is an open set V_1 containing 0_X such that $f_{|V_1}$ is a homoeomorphism and moreover both f and its inverse g are Lipschitz continuous. It is then easy to check that the inverse function g is differentiable.

Remark 3.3. A few comments about the theorem:

- Checking the condition that Df_a is invertible is straight-forward: It is equivalent to the non-vanishing of the determinant $J_f(a) = \det(Df_a)$ of the Jacobian matrix of Df_a.
- Let $U \subseteq X$ and $V \subseteq Y$ be open subsets of normed vector spaces X and Y respectively. We say that a continuously differentiable function $f\colon U \to Y$ is a *diffeomorphism* from U to V if it is injective with image $f(U) = V$, and its inverse $g\colon V \to U$ is continuously differentiable. The inverse function theorem can then be stated as follows: Let $f\colon D \to Y$ be a continuously differentiable function on an open subset $D \subseteq X$ taking values in a normed vector space Y. If Df_a is invertible, then there is an open neighbourhood $U \subseteq D$ of a on which f restricts to a diffeomorphism between U and its image $f(U) \subseteq Y$.

  [*Warning*: some references may only require f and g to be differentiable, while others may require that f and g are infinitely differentiable. To avoid ambiguity, one can also say C^1-diffeomorphism.]
- The formula for the derivative of g is forced on us by the chain rule – if g is differentiable, the chain rule applied to the composite $I_Y = f \circ g$, shows that $I_Y = DI_Y = Df(g(y)) \circ Dg(y)$ and so $Dg(y) = Df(g(y))^{-1}$.
- It is not sufficient, even if just wanted f to have a continuous inverse, for the function f to be differentiable with $f'(a)$ invertible: Consider the example $f\colon \mathbb{R} \to \mathbb{R}$, where $f(x) = x + 2x^2 \sin(1/x)$, which is extended by continuity to $x = 0$, so $f(0) = 0$. Then computing directly from the definition, we find $f'(0) = 1$ (which is invertible), but f is not injective in any neighborhood of 0.

  [**For those who read Remark 2.32, the function f is differentiable but not strongly differentiable at $x = 0$.*]
- The hypotheses of the theorem are also not necessary for f to have a *continuous* inverse – the function $f\colon \mathbb{R} \to \mathbb{R}$ given by $f(x) = x^3$ is continuous and has a continuous inverse $x \mapsto x^{1/3}$, however $f'(0) = 0$ so the inverse function theorem does not apply (and indeed the inverse function is not differentiable at 0).
- If $f\colon U \to \mathbb{R}^n$ is continuously differentiable with Df_x invertible for all $x \in U$, then although $f(U)$ is open in $\mathbb{R}^n$ (as we shall see below) f need not give a diffeomorphism between U and $f(U)$. Indeed f need not be injective. This happens already in two dimensions: Suppose that $U = \mathbb{R}^2\backslash\{0\}$ and $f\colon U \to \mathbb{R}^2$ is given by $f(x_1, x_2) = (x_1^2 - x_2^2, 2x_1x_2)$. Then $f(U) = U$, and we have

$$Df_{(x_1,x_2)} = \begin{pmatrix} 2x_1 & -2x_2 \\ 2x_2 & 2x_1 \end{pmatrix}.$$

  Since $\det(Df_{(x_1,x_2)}) = 4(x_1^2 + x_2^2)$ we see that $Df_{(x_1,x_2)}$ is invertible on all of $\mathbb{R}^2\backslash\{0\}$. But clearly $f(x_1, x_2) = f(-x_1, -x_2)$, so that f is not injective on U. If however we assume in addition that $f\colon U \to \mathbb{R}^n$ is injective, then it is indeed a diffeomorphism from U to $f(U)$ – see below.

3.2 *Proof of the Inverse Function Theorem

As noted above, by replacing f with $Df_a^{-1}(f(x + a) - f(a))$ we may assume that $Y = X$ and $Df_a = I_X$, and that $a = f(a) = 0_X$.

The heart of the proof is the following Proposition, which establishes a rigorous version of the idea that a small perturbation of the identity map should still be invertible, that is $I_X + \varphi$ should be invertible is φ is sufficiently small" compared to I_X. In the case of the space of linear maps $\mathcal{L}(X, X)$, our proof of Lemma 3.1 shows that $B(I_X, 1)$ consists of invertible elements, so in this case a "small perturbation" can be taken to mean a linear map map of operator norm strictly less than 1. But a linear map α has $\|\alpha\|_\infty < 1$ exactly when it is a contraction (that is, a Lipschitz map with a Lipschitz factor less than 1), and thus a natural candidate for a "small perturbation" is a contraction map *i.e.* a Lipschitz map with Lipschitz constant less than 1. (Note this is consistent with the requirement in the linear case at least!)

The next Proposition shows that using this notion of a small perturbation for functions defined on a closed ball, the contraction mapping theorem does indeed provide the tools to show that such a perturbation has a continuous (in fact Lipschitz continuous) inverse, at least if we shrink the domain of f to a ball of smaller radius.

Proposition 3.4. *Let X be a finite-dimensional normed vector space. Suppose that for some $r > 0, C \in (0,1)$ we are given a function $\varphi\colon \bar{B}(0_X, r) \to X$ satisfying $\varphi(0_X) = 0_X$ and*

$$\|\varphi(x) - \varphi(y)\| \leq C.\|x - y\| \quad \forall x, y \in \bar{B}(0, r).$$

Then if $f\colon \bar{B}(0_X, r) \to X$ is given by $f(x) = x + \varphi(x)$, and $y \in \bar{B}(0, 1 - C).r)$, there is a unique $x \in \bar{B}(0, r)$ such that $f(x) = y$. Moreover, the function $g\colon \bar{B}(0, (1 - C).r) \to \bar{B}(0, r)$ defined by $f(g(y)) = y$ is Lipschitz continuous with Lipschitz constant $(1 - C)^{-1}$.

Proof. Given $y \in \bar{B}((,0)(1 - C).r$, let $\varphi_y(x) = y - \varphi(x)$. Then we have

$$\|\varphi_y(x)\| = \|y - \varphi(x)\| \leq \|y\| + \|\varphi(x)\| \leq (1 - C).r + C.r = r,$$

so that φ_y maps $\bar{B}(0, r)$ to itself. Since $\bar{B}(0, r) \subset X$ is closed and X is complete, $\bar{B}(0, r)$ itself is complete and non-empty (since $0_X \in \bar{B}(0, r)$). Moreover,

$$\|\varphi_y(x) - \varphi_y(x')\| = \|\varphi(x') - \varphi(x)\| \leq C.\|x - x'\|, \quad \forall x, x' \in \bar{B}(0, r),$$

thus φ_y is a contraction on $\bar{B}(0, r)$. The Contraction Mapping Theorem thus implies that there is a unique point x_y with $\varphi_y(x_y) = x_y$, that is, $f(x_y) = x_y + \varphi(x_y) = y$. Let $g\colon \bar{B}(0, r/2) \to \bar{B}(0, r)$ be given by $g(y) = x_y$.

To see that g is continuous, let $y_1, y_2 \in \bar{B}(0, r)$. Then if $x_1 = g(y_1), x_2 = g(y_2)$ we have

$$\begin{aligned}\|f(x_1) - f(x_2\| &= \|(x_1 - x_2) + (\varphi(x_1) - \varphi(x_2)\| \geq \|x_1 - x_2\| - \|\varphi(x_1) - \varphi(x_2)\| \\ &\geq \|x_1 - x_2\| - C.\|x_1 - x_2\| = (1 - C).\|x_1 - x_2\|,\end{aligned}$$

thus $\|y_1 - y_2\| \leq (1 - C)^{-1}.\|g(y_1) - g(y_2)\|$ and hence g is Lipschitz continuous on $\bar{B}(0, (1 - C).r)$. □

The proof the Inverse Function Theorem for differentiable functions follows from this Proposition and two additional facts:

i) If $Df_{0_X} = I_X$ and Df_x is continuous at 0_X, then f is a "small" perturbation of I_X in $\bar{B}(0_X, r)$ for sufficiently small $r > 0$, so that we can apply the above Proposition.

ii) The inverse function g given by the Proposition is differentiable at $y = f(x)$ provided f is differentiable at x.

The first of these is an easy consequence of the Mean Value Inequality. Indeed we can even choose which value of C we prefer, for example we may take $C = 1/2$.

Lemma 3.5. *Suppose that X is a finite-dimensional normed vector space, $U \subset X$ is an open neighbourhood of 0_X, and let $f\colon U \to X$ be a differentiable function on U. If Df is continuous at 0_X and $Df_{0_X} = I_X$, then if $\varphi\colon U \to X$ is given by $\varphi(x) = f(x) - x$, there is an $r > 0$ such that for all $x, y \in \bar{B}(0_X, r) \subset U$,*

$$\|\varphi(x) - \varphi(y)\| \leq \frac{1}{2}.\|x - y\|.$$

Proof. By definition, since f is differentiable at $x \in U$, so is φ. Indeed for all $x \in U$ we have $D\varphi_x = Df_x - I_n$. In particular, $D\varphi_{0_X} = 0_{\mathcal{L}(X,X)}$. Since $D\varphi$ is continuous at a, there is an $r_1 > 0$ such that $\|D\varphi_x\|_\infty \leq 1/2$ for all $x \in B(0_X, r_1)$. But then by the Mean Value Inequality (Theorem 2.24), we have $\|\varphi(x) - \varphi(y)\| \leq \frac{1}{2}\|x - y\|$ for all $x, y \in B(0, r_1)$ hence on $\bar{B}(0, r)$ for any $r \in (0, r_1)$. □

The final part of the proof, checking where the inverse function is differentiable, is also straight-forward:

Lemma 3.6. *Suppose that X is a finite-dimensional normed vector space, U is an open subset of X, and $f\colon U \to X$ a injective function whose image $f(U)$ contains an open subset V. If $g\colon V \to U$ is the inverse of the restriction of f to $f^{-1}(V)$ and g is continuous at $b = f(a) \in V$, where Df_a is invertible, then g is differentiable at b and $Dg_b = (Df_a)^{-1}$.*

Proof. By replacing f by $x \mapsto Df_a^{-1}(f(a+x) - f(a))$ we may assume that $a = f(a) = 0_X$, and $Df_{0_X} = I_X$, so that

$$f(x) = x + \epsilon(x)\|x\| \tag{3.1}$$

where $\epsilon(x)$ is continuous at $x = 0_X$ and $\epsilon(0_X) = 0_X$. In order to show that $g = f^{-1}$ is differentiable at 0_X with derivative equal to $I_X^{-1} = I_X$, we must show that $g(y) = y + o_X(\|y\|)$.

But now $g(y) = x$ and $f(x) = y$, hence in terms of g, Equation (3.1) becomes $g(y) = y - \|g(y)\|\epsilon(g(y))$, and so we must show that $\|g(y)\|.\epsilon(g(y)) \in o_X(\|y\|)$, that is, we must show

$$\frac{\|g(y)\|}{\|y\|}.\epsilon(g(y)) \to 0 \text{ as } \|y\| \to 0.$$

But ϵ and g are continuous at 0_X and $\epsilon(0_X) = g(0_X) = 0_X$, and hence $\epsilon(g(y)) \to \epsilon(g(0_X)) = 0_X$ as $y \to 0_X$. Thus it suffices to show that $\|g(y)\|/\|y\|$ is bounded for $\|y\|$ small. But by the continuity of $\epsilon(g(y))$, there is a $\delta > 0$ such that if $\|y\| < \delta$ then $\|\epsilon(g(y))\| < 1/2$. Thus if $\|y\| < \delta$, since $y = g(y) + \epsilon(g(y).\|g(y)\|$, we have $\|y\| \geq \|g(y)\| - (1/2).\|g(y)\| = (1/2).\|g(y)\|$, and hence $\|g(y)\|/\|y\| \leq 2$ as required. □

Remark 3.7. It is worth comparing the proof of the Inverse Function Theorem above to the proof of the single-variable theorem. In that case, the differentiable inverse function theorem is also deduced from a continuous inverse function theorem. This is often misleadingly[12] presented as follows: Each $y \in V$ has $y = g(x)$ for a unique $x \in U$, or equivalently $f(x) = y$, hence

$$\lim_{y\to y_0} \frac{g(y) - g(y_0)}{y - y_0} = \lim_{y\to y_0} \frac{g(f(x)) - g(f(x_0))}{f(x) - f(x_0)} = \lim_{y\to y_0} \frac{x - x_0}{f(x) - f(x_0)} = \lim_{x\to x_0} \frac{x - x_0}{f(x) - f(x_0)} = 1/f'(x_0)$$

The algebraic manipulation is of course straight-forward, however the real content in the deduction is the justification for the second-last equality, that is, showing that one can switch from taking $\lim_{y\to y_0}$ to taking $\lim_{x\to x_0}$. It is here that the continuity of the inverse function is essential, since if $g = f^{-1}$ is continuous at y_0 then and hence if $y \to y_0$ then $g(y) \to g(y_0)$, that is $x \to x_0$, and thus the change of limit is indeed legitimate.

Remark 3.8. The continuous inverse function theorem in the single-variable case has a rather different proof to the many-variable case. This is because it is usually stated for functions on a closed interval, $f\colon [a,b] \to \mathbb{R}$. In this case, if f is injective, you can show it must be strictly increasing or decreasing, and replacing f with $(-f)$ if necessary we can assume it is increasing. It is then easy to see that the inverse, $f^{-1}\colon f([a,b]) \to [a,b]$ is also increasing, and by the Intermediate Value Theorem, $f([a,b])$ is the interval $[f(a), f(b)]$. But an increasing function can only have "jump" discontinuities, *i.e.*, the one-sided limits $f(x_0)^+ = \lim_{x\to x_0^+} f(x)$ and $f(x_0)^- = \lim_{x\to x_0^-} f(x)$ both exist, and $f(x_0)^- \leq f(x) \leq f(x_0)^+$, but the inequalities may all be strict. Since the image of f^{-1} is, by assumption, the interval [a,b], there can be no such discontinuities in the case of f^{-1}, and so it is continuous.

Thus, rather bizarrely, the continuity of the inverse in the one-dimensional theorem proved in Prelims is deduced from a criterion for continuity for increasing functions on an interval – namely that it is necessary and sufficient for its image to be an interval. In higher dimensions there is no reasonable notion of an increasing or decreasing function, so this argument does not generalise.

Remark 3.9. If, instead of assuming that $f\colon U \to \mathbb{R}^n$ is differentiable on U with Df continuous at $a = 0$, we assume only that it is strongly differentiable at a (see Remark 2.32), then one can modify the proof of Lemma 2.9 to show that Proposition 3.4 still holds on $\bar{B}(0,r)$ for small enough r. Similarly, Lemma 3.6 can be adapted to show that the inverse g is (strongly) differentiable at y if f is (strongly) differentiable at $x = g(y)$.

****Remark 3.10.** One can in fact somewhat weaken the hypotheses of the Inverse Function Theorem in a number of ways: if U is an open subset of $\mathbb{R}^n$ and $f\colon U \to \mathbb{R}^n$ has Df_x invertible for all $x \in U$, then f is locally invertible with differentiable inverse: More explicitly, for any $a \in U$ there are open sets U_1, V_1 with $a \in U_1 \subseteq U$ and $f(a) \in U_2$ such that f restricts to a bijection from U_1 to U_2 and if $g = f_{|U_1}^{-1}\colon U_2 \to U_1$, then g is differentiable with derivative $Df_{g(y)}^{-1}$ for all $y \in U_2$. Indeed by the chain rule, it follows that invertibility of Df_x for all $x \in U$ is equivalent to the local invertibility of f.

[12] In that it hides the key point in a subscript.

More importantly, especially for applications in the study of partial differential equations, the inverse function theorem holds for continuously differentiable functions on open subsets of any complete normed vector space, whether or not it is finite dimensional. In this context, the derivative must be a continuous linear map (that is, a bounded linear map – see Section 1). Thus the condition that the derivative at a point be invertible has to demand instead that the inverse linear map exists and is bounded, but then the whole theorem (and its proof) go through just as above. In fact, it is the case (though we do not quite have the tools to show it) that in a *complete* normed vector space (the ones in which the inverse function theorem holds) if a linear map is invertible (*i.e.* has a linear inverse) then its inverse is automatically continuous.

3.3 Applications of the Inverse Function Theorem

Definition 3.11. Let (X, d) and (Y, ρ) be metric spaces. A continuous function $g\colon X \to Y$ is said to be an *open mapping* if, for any open set $U \subset X$, its image $g(U)$ is open in Y. Notice that a continuous bijection is a homeomorphism precisely if it is an open mapping.

Corollary 3.12. *Let $U \subset \mathbb{R}^n$ be an open set, and $f\colon U \to \mathbb{R}^n$ be a continuously differentiable function such that Df_x is invertible for every $x \in U$. Then f is an open mapping.*

Proof. Let V be an open subset of $\mathbb{R}^n$ contained in E. We want to show that $f(V)$ is open. Pick $b \in f(V)$. We need to show that $f(V)$ contains some open ball centered at b. Now $b = f(a)$ for some $a \in O$, and the inverse function theorem applies to $f_{|V}\colon V \to \mathbb{R}^n$ and $a \in V$. Hence there are open sets V_1, V_2 with $a \in V_1 \subset V$ and $f(a) = b \in V_2$ such that f is a bijection between V_1 and V_2. But then there is a $\delta > 0$ such that $B(b, \delta) \subset V_2 = f(V_1) \subset f(V)$, and we are done. □

Remark 3.13. In fact the proof of this theorem used only the first part of the inverse function theorem – the fact that the inverse of f on U is continuously differentiable was not needed.

Another consequence of the inverse function theorem is the following:

Corollary 3.14. *Let $E \subset \mathbb{R}^n$ be an open subset and let $f\colon E \to \mathbb{R}^n$ be continuously differentiable, such that f is injective and Df_x is invertible for all $x \in E$. Then f is a diffeomorphism between E and $f(E)$.*

Proof. By assumption, given $y \in f(E)$ there is a unique $x \in E$ with $f(x) = y$, so that we can define $h\colon f(E) \to E$ by setting $h(y)$ to be this point x. But then g is continuously differentiable by the inverse function theorem, since at any point $y \in f(E)$, if $x = g(y)$ there are open sets U, V containing x and y respectively, such that $f_{|U}\colon U \to V$ is a diffeomorphism. But then $g_{|V}$ is continuously differentiable, and so g is continuously differentiable at $y \in V$. □

3.4 Systems of local coordinates and the Implicit Function Theorem.

The goal of our study of differentiable functions is to try to extend to such functions, in as much as this makes sense, results from linear algebra. To try and make this analogy between results in the linear and non-linear setting a little more concrete, consider the notion of coordinates on a vector space: If X is an n-dimensional vector space, then picking a basis $B_X = \{v_1, \ldots, v_n\}$ of X gives us coordinates for the vectors in V: for any vector $v \in X$ we assign to it the coordinates $(c_1, \ldots, c_n) \in \mathbb{R}^n$ where $v = \sum_{i=1}^n c_i v_i$. Equivalently, the basis defines an invertible linear map $\theta\colon X \to \mathbb{R}^n$ given by sending B_X to the standard basis of $\mathbb{R}^n$. Thus giving such a map is equivalent to giving a (linear) coordinate systems on X. In the setting of differentiable functions, diffeomorphisms play the same role: if U is an open subset of X and $f\colon U \to \mathbb{R}^n$ is a diffeomorphism onto its image $f(U) \subseteq \mathbb{R}^n$, then we can use the components of f to parameterise the points in U.

This gives one way of thinking of the Inverse Function Theorem, namely, it ensures that if U is open in X and $f\colon U \to \mathbb{R}^n$ is continuously differentiable, then if Df_p is invertible, at least near p, f is a diffeomorphism. In other words, if the derivative Df_p gives (linear) coordinates on X, then, the components of f provide a (non-linear) parameterization of neighbourhood of p.

Example 3.15. Suppose that X is 2-dimensional with basis $\{v_1, v_2\}$. The function $g\colon \mathbb{R}^2 \to X$ given by $g\colon (r, s) \mapsto r\cos(s).v_1 + r\sin(s).v_2$ has Jacobian determinant $J_g = r$, thus if we let $V = (0, \infty) \times (0, 2\pi)$, then $g\colon V \to U$, where $U = X \backslash \{t.v_1 : t \geq 0\}$, and $J_g \neq 0$ on all of V, so the inverse function theorem ensures that g has an inverse

$f\colon U \to V = (0,\infty)\times(0,2\pi)$. Since $g(f(v)) = v$, the function f simply assigns to $v \in V$ its "polar coordinates" (r,θ).

Note that U, the domain of f, is not all of X. If we wanted to enlarge the domain of definition of f, we would need to extend g to some $\tilde{U} \supseteq U$ to make it bijective, but it we try and do this, two problems present themselves: Firstly, if (r,s) has s close to 2π and s' close to 0, then $g(r,s)$ and $g(r,s')$ will both be close to rv_1, indeed $\lim_{s\to 2\pi} g(r,s) = \lim_{s'\to 0} g(r,s') = rv_1$. This forces the inverse of g to have a discontinuity at rv_1 – the limits $\lim_{t\downarrow 0} g(rv_1+tv_2) = (r,0)$ while $\lim_{t\uparrow 0} g(rv_1+tv_2) = (r,2\pi)$. Worse still, for 0_X to lie in the image of g, we must add to U an element of $(0,s)$, say $(0,s_0)$ but for any $s_1 \in \mathbb{R}$ we have $\lim_{r\to 0} g(r,s_1) = 0_X$, so that any choice of s_0 will for f to be discontinuous at 0_X. This latter problem is a consequence of the fact that, although g is defined on all of $\mathbb{R}^2$, its derivative is only nonsingular when $r \neq 0$. The former problem is an example of the local nature of the inverse function theorem – a continuously differentiable inverse is only guaranteed to exist sufficiently close to the point you apply it to. This is often less problematic – for example with polar coordinates, although any choice will have a discontinuity along any path which encircles the origin, we can control where this appears: for example we can chose $U' = (0,\infty)\times(\alpha,\alpha+2\pi)$ for the domain of g so that f is discontinuous on the ray $t(\cos(\alpha)v_1 + \sin(\alpha)v_2)$.

Definition 3.16. A *pointed set* is a pair (X,a) consisting of a set X and an element a of X. If (X,a) and (Y,b) are pointed sets, then we will write $f\colon (X,a)\to(Y,b)$ to indicate that f is a function from X to Y which maps a to b, that is, $f(a) = b$.for a function $f\colon X\to Y$ which satisfies $f(a) = b$, and refer to it as a map (or function) of pointed sets.

Remark 3.17. Many algebraic objects are naturally pointed – a vector space X has a zero vector, any group has an identity element *etc.*

Definition 3.18. Suppose that X is a normed vector space and $p \in X$. A *system of local coordinates* at p is a diffeomorphism $\psi\colon (U,p)\to(\Omega,0_n)$ from a connected[13] open neighbourhood U of the origin p in X to a connected open neighbourhood Ω of $0_n \in \mathbb{R}^n$. The standard coordinates $(x_1,\dots,x_n)$ of $\mathbb{R}^n$ at 0_n then give a system of coordinates $(t_1,\dots,t_n)$ at p, where, for $y \in U$, we set $t_i(y) = x_i \circ \psi(y)$, for $i \in \{1,\dots,n\}$.

If $f\colon U \to \mathbb{R}^k$ is any function, then by the chain rule, $f\circ\psi^{-1}$ is continuously differentiable when f is, and similarly, if a function $g\colon \Omega\to\mathbb{R}^k$ is continuously differentiable, then so is $g\circ\psi$, since, as ψ is a diffeomorphism, both ψ and ψ^{-1} are continuously differentiable. Thus the map $\psi^*\colon C^1(\Omega,\mathbb{R}^k)\to C^1(U,\mathbb{R}^k)$ given by $\psi^*(f) = f\circ\psi$ is an isomorphism of vector spaces, with inverse $(\psi^{-1})^*$ where $(\psi^{-1})^*(g) = g\circ\psi^{-1}$. More prosaically, this just says that if we wish to check if a function $f\colon U\to\mathbb{R}^k$ is continuously differentiable, we just need to check that it is continuously differentiable when viewed as a function of the coordinates $(t_1,\dots,t_n)$ given by the diffeomorphism ψ.

In this section we will use the Inverse Function Theorem to show that, for functions $f \in C^1(U,\mathbb{R}^k)$, structural information about the linear map Df_p at a point $p \in U$ can often be extended to give information about the behaviour of f near p.

Our main example of this is the Implicit Function Theorem. The linear algebra toy model for this theorem is the description of a surjective linear map $\alpha\colon X\to Y$. If $\{v_1,\dots,v_l\}$ is a basis for $\ker(\alpha)$, then we may extend it to a basis $\{v_1,\dots,v_{k+l}\}$ of X. The images the additional vectors yield a basis of Y, and in the coordinates these provide for X and Y the map α takes the form $\alpha(t_1,\dots,t_{k+l}) = (t_{k+1},\dots,t_{k+l})$. Similarly, the basis $\{v_1,\dots,v_l\}$ provide a set of coordinates $\psi\colon\ \ker(\alpha)\to\mathbb{R}^l$ for $\ker(\alpha)$.

From a computational point of view however, there is still the question of how one constructs the basis $\{v_1,\dots,v_l\}$ of $\ker(\alpha)$. In practice if $\alpha\colon X\to Y$ is a surjective linear map, it is likely to be given via its matrix with respect to some bases B_X, B_Y of X and Y respectively which have no particular compatibility with α.[14]. In such cases, it may be easier to find a subspace $X_2 \subseteq X$ such that $\alpha_{|X_2}\colon X_2\to Y$ is actually bijective. Then, if we pick any complementary subspace X_1, we may decompose α accordingly as $\alpha = (\alpha_1+\alpha_2)$, where $\alpha_i = \alpha\circ\pi_i$ for $i = 1,2$ and the maps π_i are the natural projection operators with images X_1 and X_2 respectively. For example,

[13]the assumption that Ω is connected is not necessary, but it is easy to ensure – if V is an arbitrary open neighborhood of 0_X then if C is the connected component of V containing 0_X, it is again an open neighbourhood of 0_X which is, of course, connected.

[14]In the context of experimental science or economics, for example, the bases B_X and B_Y are likely to be constructed in a way that reflects those qualities we can most readily measure.

if we have chosen bases B_X and B_Y for X and Y and $A = {}_{B_Y}[\alpha]_{B_X} \in \mathrm{Mat}_{k,n}(\mathbb{R})$ is the matrix of α with respect to these bases, then in low rank cases it is often not hard to find a $k \times k$ submatrix of A which has full rank. This partitioning of the columns of A into two sets of size k and $n-k$ then yield a correspond to a decomposition of X into a direct sum $X_1 \oplus X_2$. The following Lemma then shows that one can obtain a concrete description of the kernel of α using this decomposition. (This is really the description one obtains from the reduced row eschelon form of a matrix as in Prelims Linear algebra.)

Lemma 3.19. *Let $\alpha\colon X \to Y$ be a surjective linear map and suppose $B = B_1 \sqcup B_2$ is a basis of X such that if $X_i = Span(B_i)$, for $i = 1, 2$, then $\alpha_{|X_2}\colon X_2 \to Y$ is an isomorphism. Then there is a linear map $\theta\colon X_1 \to X_2$ such that $ker(\alpha) = \{(x, \theta(x)) : x \in X_1\}$.*

Moreover, if $B_Y = \alpha(B_2)$ is a basis of Y, and is we set $B_1^\theta = \{(b, \theta(b)) : b \in B_1\}$, then B_1^θ is a basis of $\ker(\alpha)$ and if $B^\theta = B_1^\theta \cup B_2$, then B^θ is a basis of X and the matrix ${}_{B_Y}[\alpha]_{B^\theta}$ with respect to these bases is in the canonical form $(0|I_k)$.

Proof. Let π_1, π_2 be the projection maps from X to X_1 and X_2 respectively (so $\ker(\pi_1) = X_2$ and $\ker(\pi_2) = X_1$). If $\gamma\colon Y \to X_2$ is the inverse of $\alpha_{|X_2}$, we may use it to identify Y with X_2, that is, we replace α with $\beta = \gamma \circ \alpha$, so that we may view α as a linear map from X to X_2 where if $\beta_i = \pi_i \circ \beta$ then $\beta = \beta_1 + \beta_2$ and $\beta_2(x) = \pi_2 \circ \gamma \circ \alpha = \pi_2(x)$.

Now let $T\colon X \to X$ be given by $T(x) = \pi_1(x) + \beta(x) = x + \beta_1(x)$, so that in terms of the decomposition $X = X_1 \oplus X_2$ we have $T = \begin{pmatrix} \pi_1 & 0 \\ \beta_1 & \pi_2 \end{pmatrix}$. Then T has inverse $T^{-1}(x) = x - \beta_1(x)$. It follows that $\beta(x) = 0$ if and only if $\pi_2(x) = \beta_1(x)$, so that $\ker(\alpha) = \{(x, \beta_1(x)) : x \in X_1\}$ as required. The final sentence then follows immediately from the above. $\square$

We now state the Implicit Function Theorem: Its formulation is almost identical to the linear algebra result given above: we take a differentiable function $f\colon U \to Y$ in place of the linear map α, but then, for a point $p \in U$ where the hypothesis of the previous Lemma are satisfied by the derivative Df_p of our function at p, just as in the case of the Inverse Function Theorem, we obtain a "local" consequence for the function f, that is, a statement about the nature of our function in a neighbourhood of the point in question.

Definition 3.20. If X and Y are normed vector spaces and $f \in C^1(U, Y)$, and $p \in U$ is such that $Df_p\colon X \to Y$ is surjective, the set $U_{\max} = \{x \in U : Df_x \text{ is surjective}\}$ is an open neighbourhood of p and we say that the restriction of f to $U_{\max}$ is a *submersion*.

Exercise 3.21. Check that you see why $U_{\max}$ is open – compare with Lemma 3.1.

Theorem 3.22. *(The Implicit Function Theorem.) Suppose that X and Y are normed vector spaces and we are given a direct sum decomposition $X = X_1 \oplus X_2$, with π_1, π_2 the corresponding projections to X_1 and X_2 respectively. Let U be an open subset of X, and let $f\colon U \to Y$ be a differentiable function. If $p = (x_0, y_0) \in U$ is such that $f(x_0, y_0) = 0$ and, for $i = 1, 2$, we write $\partial_i f(q)$ for the partial derivative $\partial_{X_i} f(q)$ of f with respect to X_i at $q \in U$, so that we have the decomposition*

$$Df = \partial_1 f(q) \circ \pi_1 + \partial_2 f(q) \circ \pi_2, \quad \forall q \in U.$$

If $\partial_2 f(q)$ is continuous at p and $\partial_2 f(p)$ is invertible, then there are open neighbourhoods V_1, W_1 of the zero vectors $0_1 = 0_{X_1}$ and $0_2 = 0_{X_2}$ respectively, and a differentiable function $\theta\colon V_1 \times W_1 \to \Omega$, where $\Omega \subseteq U$ is an open neighbourhood of p, such that if $\theta(x, y) = (\theta_1(x, y), \theta_2(x, y))$ then $\theta_1(x, y) = x + x_0$, and if $(x, y) \in \Omega$, then $f(x, y) = 0$ if and only if $(x, y) = (x, \theta_2(x - x_0, 0))$.

$$\begin{tikzcd}
X_1 \oplus X_2 \arrow[d, "\pi_1"] & V_1 \times W_1 \arrow[l, hook'] \arrow[r, "\theta"] & \Omega \arrow[d, "f"] \\
X_1 \arrow[rr, "Dg_x"] & & Y
\end{tikzcd}$$

Equivalently, if $g(x) = \psi_2(x - x_0, 0)$, then g is continuously differentiable, and if $(x, y) \in \Omega$ then $f(x, y) = 0$ if and only if $y = g(x)$. That is, within Ω, the set $f(x, y) = 0$ can be described as the graph of $g\colon V_1 \to X_2$. Moreover, the derivative of g is given by

$$Dg_x = -\partial_2 f(x, g(x))^{-1} \circ \partial_1 f(x, g(x))$$

Proof. (*Non-examinable*:) Let $\beta\colon Y \to X_2$ be the inverse of $\partial_2 f(p)$. By replacing f with $\beta \circ f$, we may assume that $f\colon X \to X_2$ and that $\partial_2 f(p) = \pi_2$. Similarly, by replacing f by $f(x_0 + x, y_0 + y) - f(x_0, y_0)$, we may assume that $p = 0_X = f(0_X)$. Define $G\colon U \to X$ be given by

$$G(x,y) = (x, f(x,y)), \quad x \in X_1, y \in X_2.$$

so that $G(0_X) = G(x_0, y_0) = 0_X$. Then, for any $q = (x,y) \in U$ decomposing $Df_q = \partial_1 f(q) \circ \pi_1 + \partial_2 f(q) \circ \pi_2$ according to the direct sum decomposition of X, we have

$$DG_q = \left(\begin{array}{c|c} I_{X_1} & 0 \\ \hline \partial_1 f(q) & \partial_2 f(q) \end{array}\right) \quad \forall q \in U$$

Thus G is differentiable, and continuously differentiable wherever f is. Moreover, since $\partial_2 f(p) = I_{X_2}$ is invertible, it follows that DG_p is invertible. It follows from the Inverse Function Theorem that there is an open set $\Omega \subseteq U$ with $0_X \in \Omega$ such that $G_{|\Omega}\colon \Omega \to V = G(\Omega)$ is a diffeomorphism from Ω to an open set V which contains $0_X = G(0_X)$. It follows that we may find open neighbourhoods V_1 and W_1 of 0_{X_1} and 0_{X_2} respectively such that $V_1 \times W_1 \subseteq V$, and if we let $\theta = (G_{|\Omega})^{-1}$, then $\theta(V_1 \times W_1) \subseteq \Omega$ is an open subset of X containing 0_X, so that by replace Ω with $\theta(V_1 \times W_1)$ we may assume V is a product of the form $V_1 \times W_1$.

Now if, for $(s,t) \in V_1 \times W_1$ we set $\theta(s,t) = (\theta_1(s,t), \theta_2(s,t)) \in X_1 \oplus X_2$, then if $(x,y) \in \Omega$, we have $G(x,y) = (x, f(x,y))$, hence $(x,y) = \theta \circ G(x,y)) = \theta(x, f(x,y))$. In particular, since G is surjective $\theta_1(s,t) = s$. Moreover, $f(x,y) = 0$ if and only if $(x,y) = \theta_2(x, 0_{X_2})$.

It follows that if we let $N(f) = \{(x,y) \in U : f(x,y) = 0\}$ and $g\colon V_1 \to X_2$ be given by $g(x) = \theta_2(x, 0_{X_2})$, then $N(f) \cap \Omega = \{(x, g(x)) : x \in V_1\}$.

Thus the theorem is proved except for the expression for the derivative of $g(x) = \theta_2(x, 0)$. But this follows by invertible the matrix of DG_q above, or by noting $0 = f(x, g(x))$, which implies by the chain rule that

$$0 = \left(\begin{array}{c|c} \partial_1 f(x, g(x)) & \partial_2 f(x, g(x)) \end{array}\right)\left(\begin{array}{c} I_{X_1} \\ \hline Dg_x \end{array}\right).$$

and hence $\partial_1 f(x, g(x)) + \partial_2 f(x, g(x)) \circ Dg_x = 0$, so that $Dg_x = -\partial_2 f(x, g(x))^{-1} \partial_1 f(x, g(x))$. □

Remark 3.23. This result is called the "Implicit Function Theorem" because one can view it as saying that, if we pick a basis for Y and consider the corresponding real-valued functions f_i given by the components of f with respect to this basis, then provided the linear map $\partial_2 f(x_0, y_0)$ is invertible, the system non-linear of equations $f_i(x,y) = 0$ for $i = 1, 2, \ldots, k$, can be solved, in the sense that the equations *implicitly* make the y-variables functions of the x-variables, at least locally near (x_0, y_0), as the existence of the function g demonstrates.

In this sense, the theorem gives a rigorous justification for the calculus technique of "implicit differentiation" – compare that technique to the calculation of Dg at the end of the above proof.

Corollary 3.24. *(Local normal form for a submersion): We can also formulate the Implicit Function theorem in terms of local systems of coordinates: Pick a basis B_1 of X_1 and B_Y of Y, so that if $B_2 = \partial_2 f(p)^{-1}(B_Y)$, then B_2 is a basis of X_2 and $B_X = B_1 \sqcup B_2$ is a basis of X. Then using θ we obtain a system of local coordinates, $(t_1, \ldots, t_n)$ say, for (Ω, p), where $(x,y) \in \Omega$ has coordinates $(t_1, \ldots, t_n)$ if $\theta_1(t_1, \ldots, t_n) = x$ and $\theta_2(t_1, \ldots, t_n) = y$. With respect to this system of local coordinates, and the linear coordinates on Y given by B_Y, it follows immediately from the definition of G that, if $n = \dim(X)$ and $k = \dim(Y)$, then the map f takes the form $(t_1, \ldots, t_n) \mapsto (t_{n-k+1}, \ldots, t_n)$.*

Proof. This follows immediately from the discussion above. Note that in this formulation, the theorem shows that the components of f can be extended to a local system of coordinates for X near p provided f is continuously differentiable and Df_p has full rank (*i.e.* there is a subspace X_2 of X for which the restriction $Df_{p|X_2}\colon X_2 \to Y$ is an isomorphism). □

Example 3.25. In this example, we will write z for a general vector in $\mathbb{R}^4$ and write $z = (x,y)$ where $x \in \mathbb{R}^2$, $y \in \mathbb{R}^2$. Let $f\colon \mathbb{R}^4 \to \mathbb{R}^2$ be given by

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

and consider the level set $M = f^{-1}\{(1,2)\}$ of f, so that

$$M = \left\{ z = (x_1, x_2, y_1, y_2) \in \mathbb{R}^4 : \begin{array}{ll} x_1^2 - x_2^2 + y_1^2 + 2y_2^2 & = 1 \\ x_1^2 + x_2^2 - y_1^2 - y_2^2 & = 2 \end{array} \right\}.$$

The total derivative Df_z has Jacobian matrix

$$Df_z = (Df_{1,x}|Df_{2,y}) = \begin{pmatrix} 2x_1 & -2x_2 & 2y_1 & 4y_2 \\ 2x_1 & 2x_2 & -2y_1 & -2y_2 \end{pmatrix}, \tag{3.2}$$

Thus considering 2×2 submatrices, we see that Df has rank 0 only at $z = 04$, and rank 1 if z lies on the coordinate axes (*i.e.* all but one of x_1, x_2, y_1, y_2 equal to zero), or if $x_1 = y_2 = 0$. Everywhere else Df_z has maximal rank. Now if $x \in M$ we have $2x_1^2 + y_2^2 = 3$, hence M does not intersect the plane $\{z \in \mathbb{R}^4 : x_1 = y_2 = 0\}$. Similarly it is easy to see that M does not intersect the coordinate axes, and hence Df has maximal rank on all of M. (In the terminology of the next section, this means that M is a 2-dimensional submanifold of $\mathbb{R}^4$.)

We now consider how to parametrize M. Using Theorem 3.22, and noting that the final two columns form an invertible matrix provided $y_1 y_2 \neq 0$, we see that in a neighbourhood of a point $p = (a, b, c, d) \in M$ for which $c.d \neq 0$, the condition that $f(x_1, x_2, y_1, y_1) = (1, 2)$, *i.e.* implicitly defines a function g in a neighbourhood of (a, b) such that

$$f(x_1, x_2, y_1, y_2) = (1,2) \iff (y_1, y_2) = g(x_1, x_2),$$

that is, locally near p, the level set M is the graph of a function.

The theorem however does not produce the parameterizing function $g = (g_1, g_2)$. However, it does allow us to calculate the derivative Dg_x: If $z = (x, g(x))$ we have $Dg_x = -Df_{2,g(x)}^{-1} Df_{1,x}$, where, as in (3.2) we write $Df_z = (Df_{1,x}|Df_{2,y})$. Explicitly this becomes:

$$\begin{aligned} Dg_x = \begin{pmatrix} \partial_1 g_1 & \partial_2 g_1 \\ \partial_1 g_2 & \partial_2 g_2 \end{pmatrix} &= -(4g_1g_2)^{-1} \begin{pmatrix} -2g_2 & -4g_2 \\ 2g_1 & 2g_1 \end{pmatrix} . \begin{pmatrix} 2x_1 & -2x_2 \\ 2x_1 & 2x_2 \end{pmatrix} \\ &= (4g_1g_2)^{-1} \begin{pmatrix} 12x_1g_2 & 4x_2g_2 \\ -8x_1g_1 & 0 \end{pmatrix}. \\ &= \begin{pmatrix} 3x_1/g_1 & x_2/g_1 \\ -2x_1/g_2 & 0 \end{pmatrix}. \end{aligned}$$

Indeed one can view the Implicit Function Theorem (or indeed the Inverse Function Theorem) as asserting the unique solution to a system of differential equations. Of course in general we may not be able to readily solve these equations explicitly, but this example is simple enough that we can:

To start, note that $\partial_2 g_2 = 0$, so g_2 is independent of x_2, while $g_2.\partial_1 g_2 = -2x_1$ so that the only equation governing g_2 is $\partial_1 g_2 = 2x_1/g_2$. Indeed we already noted that on M, $2x_1^2 + y_2^2 = 3$, that is, $2x_1^2 + g_2^2 = 3$, hence $g_2(x_1, x_2) = \pm\sqrt{3 - 2x_1^2}$, where the sign will be determined by the sign of d, the corresponding coefficient of p. Note that we have $\partial_1(\sqrt{3 - 2x_1^2}) = -2x_1/\sqrt{3 - 2x_1^2}$ as expected. Having determined g_2, it is not so difficult to determine g_1, using, for example, the first component of f:

$$g_1(x_1, x_2) = \pm\sqrt{1 - x_1^2 + x_2^2 - 2.(3 - 2x_1^2)} = \pm\sqrt{3x_1^2 + x_2^2 - 5},$$

where again, the sign is determined by that of the corresponding coefficient of p (which is c in this case). Note again that $\partial_1 g_1 = 3x_1/g_1$ and $\partial_2 g_1 = x_2/g_1$. Thus we have

$$(g_1(x), g_2(x)) = \left(\pm\sqrt{3x_1^2 + x_2^2 - 5},\ \pm\sqrt{3 - 2x_1^2}\right)$$

Example 3.26. A more abstract application of the Implicit Function Theorem is a "smooth" version of the problem of extracting the roots of a polynomial equation. It is a famous result of Abel and Ruffini[15] that for equations

[15] This predates Galois, who developed a complete theory in which the Abel-Ruffini theorem sits as a special case.

of degree $n = 5$ and higher, one cannot express the roots of a polynomial equation $p(t) = \sum_{k=0}^n a_k t^k$ "in radicals" – that is, using only the ordinary algebraic operations along with taking k-th roots for $k \leq n$. One can still however, consider how a root of p varies as we continuously vary the coefficients $\mathbf{a} = (a_k) \in \mathbb{C}^{n+1}$. It seems intuitively clear that a root will move continuously with the coefficients, and the Implicit Function Theorem allows us to make this precise:

Suppose that $c \in \mathbb{C}$ is a simple root of $p(t)$ - so $(t-c)$ divides p but $(t-c)^2$ does not. Equivalently $p(c) = 0$ but $p'(c) \neq 0$. Let $f\colon \mathbb{C}^{n+2} \to \mathbb{C}$ be the function $f(a_0, \ldots, a_n, t) = \sum_{k=0}^n a_k t^k$, that is, f is the function obtained from p by viewing it as a function of t and of all of its coefficients. Then $\partial_t f(\mathbf{a}, c) = p'(c) \neq 0$, so that if we decompose $\mathbb{C}^{n+2} = \mathbb{C}^{n+1} \oplus \mathbb{C}$, the implicit function theorem shows that there is an open neighbourhood V of $(\mathbf{a}, c)$ in which $f(\mathbf{x}, t) = 0$ if and only if $t = g(\mathbf{x})$, where $g(\mathbf{a}) = c$.

Since a polynomial is smooth (*i.e.* infinitely differentiable) we can conclude that $g(\mathbf{x})$ is also smooth. Thus the roots of a polynomial (at least when they are simple) are smooth functions of the coefficients, even if they cannot be written in the form of radicals as the mathematicians of the 17th century had wished.

***Remark 3.27.** In the setting of infinite dimensional complete normed vector spaces, the Inverse Function Theorem can be used to prove a version of the Implicit Function Theorem. Such a result can be used to prove a version of Picard's Theorem on existence and uniqueness of solutions to differential equations. See [R] for more details.

3.5 Lagrange multipliers

Suppose first that X is a normed vector space and U is an open set in X with $f\colon U \to \mathbb{R}$ a differentiable function.

Lemma 3.28. *If $f\colon U \to \mathbb{R}$ has a local minimum at $a \in U$, so that for some $r > 0$ we have $g(a) \leq g(x)$ for all $x \in B(a, r)$, then $Dg_a = 0$.*

Proof. Suppose for the sake of contradiction that $Dg_a \neq 0$. Then we may find $v \in X$ such that $Dg_a(v) > 0$ and $\|v\| = 1$. For $t \in \mathbb{R}$ let $\gamma(t) = a + t.v$, then $\gamma^{-1}(U)$ is an open set in $\mathbb{R}$ containing 0, hence for some $\delta > 0$, the function $g \circ \gamma$ is defined on $(-\delta, \delta)$. Now by definition we have

$$0 \leq g(x) - g(a) = Dg_a(x-a) + \|x-a\|\eta(x),$$

where $\eta(x) \to 0 = \eta(a)$ as $x \to a$. Thus for all $t \in (-\delta, \delta)$ we have

$$0 \leq g(\gamma(t)) - g(a) = t.[Dg_a(v) \pm \eta(a + t.v)].$$

But since $\eta(a + t.v) \to 0$ as $t \to 0$, and $Dg_a(v) > 0$, there is a $\delta_1 < \delta$ such that if $t \in (-\delta_1, \delta_1)$ then $Dg_a(v) \pm \eta(a + tv) > Dg_a(v)/2$. But then for all $t \in (-\delta_1, 0)$ the inequality above cannot hold, giving a contradiction. □

We now wish to study the problem of minimizing $g\colon U \to \mathbb{R}$ given constraints on $x \in U$. Before formulating the general result, consider the problem of trying to minimize a function $g\colon \mathbb{R}^3 \to \mathbb{R}$ on a surface $S = \{x \in \mathbb{R}^3 : f(x) = 0\}$. In the unconstrained setting, as we just saw, if a point $a \in \mathbb{R}^3$ is a local minimum for g we must have $\nabla g(a) = 0$: This need not be the case in the constrained setting.

Example 3.29. Let $f(x) = x_1^2 + x_2^2 + x_3^2 - 1$, and let $S = \{x \in \mathbb{R}^3 : f(x) = 0\}$. Suppose that we wish to mimimize $g(x) = x_3$ on S. Clearly $Dg_x = (0, 0, 1)$ never vanishes, but it is easy to check that $p = (0, 0, -1)$ minimizes g on S. Notice that, since $Df_x = 2(x_1, x_2, x_3)$, so that at p we have $2Dg_p + Df_p = (0, 0, 2) + (0, 0, -2) = 0$.

This dependence is not a coincidence: In the proof of Lemma 3.28, when $Dg_p \neq 0$ we can find a direction to move in where the linear approximation to g, given by Dg_p increases in value (and so decreases in the opposite direction) and that the approximation has an error of magnitude $o_Y(\|x-a\|)$ suffices to show that the failure of the linearized problem to have a local minimum forces the same to be true of the original nonlinear problem. In the situation of the constrained minimum in this example, $Df_p(x) = 0$ can be seen as the linear approximation to the non-linear constraint $f(x) = 0$ near p. If Dg_p and Df_p are multiples of each other, then Dg_p actually vanishes on the locus given by the linearized constraint $Df_p(x) = 0$. This suggests the replacement for the condition $Dg_p = 0$ in the unconstrained problem should be that Dg_p vanishes on the linearization at p of the constraint $f(x) = 0$, that is, we should have $\ker(Df_p) \subseteq \ker(Dg_p)$.

To make this observation into a theorem, we need to show that the linearised problem is a good enough approximation to the original non-linear constrained optimization problem for the linear condition we just obtained to remain necessary in the original problem. But this is exactly what the Implicit Function Theorem does for us!

Theorem 3.30. *Suppose that U is an open subset of a finite-dimensional normed vector space X and $g\colon U \to \mathbb{R}$ is continuously differentiable. Let $f\colon U \to \mathbb{R}^k$ be constraint function, and consider the optimization problem given by seeking to minimize $g(x)$ subject to $x \in S = \{x \in U : f(x) = 0\}$.*

If z is a local minimum for g on S, then if Df_{x_0} has rank k, there exist scalars $\lambda_0, \lambda_1, \dots, \lambda_k \in \mathbb{R}$ such that

$$\lambda_0 Dg_z + \sum_{i=1}^{k} \lambda_i Df_{i,z} = 0,$$

where $f(x) = \sum_{i=1}^{k} f_i(x).e_i$, with $\{e_i : 1 \leq i \leq k\}$ the standard basis of $\mathbb{R}^k$.

Proof. The hypothesis of the theorem ensures that we can apply the Implicit Function Theorem: Df_z has rank k, hence there is a subspace $X_2 \leq X$ on which Df_z restricts to give an isomorphism from X_2 to $\mathbb{R}^k$. If we pick any complementary subspace X_1, then the Implicit Function Theorem shows that there is an open neighbourhood Ω of z in which there is a system of local coordinates $(t_1, \dots, t_n)$ for which $f = (t_{n-k+1}, \dots, t_n)$. Thus restricting g to $f(x) = 0$ simply sets $t_{n-k+1} = \dots t_n = 0$, and hence from the previous Lemma we must have

$$Dg_z = (\partial_1 g(z) | \partial_2 g(z)) = (0 | \partial_2 g(z)),$$

and hence Dg_z lies in the span of $\{Dt_i : i \geq n-k+1\}$, or equivalently the span of $\{Df_i : 1 \leq i \leq k\}$, which is equivalent to the existence of the linear dependence in the statement of the theorem. □

Remark 3.31. Since the hypothesis of the Theorem assumes that Df_z has rank k, and the Jacobian matrix of Df_z has rows given by the derivatives of the components $Df_{i,z}$, these are linearly independent, so that the scalar λ_0 must be non-zero. It follows that one can rescale the λ_i to ensure $\lambda_0 = 1$, and some texts will state the result this way. (In practice, in some situations the calculations are tidier setting $\lambda_0 = 1$ and in others it can be easier not to distinguish λ_0 in this way.)

Example 3.32. Consider the problem of finding the extrema of the function $g\colon \mathbb{R}^3 \to \mathbb{R}$ given by

$$g(x_1, x_2, x_3) = x_1 + x_2 + 3x_3,$$

subject to the constraints that $x = (x_1, x_2, x_3)$ must satisfy $(f_1(x), f_2(x)) = (2, 1)$ where

$$f_1(x) = x_1^2 + x_2^2, \quad f_2(x) = x_1 + x_2 + x_3.$$

That is, x lies on the cylinder of radius $\sqrt{2}$ centred along the x_3-axis and on the plane perpendicular to $(1,1,1)$ passing through $\frac{1}{3}(1,1,1)$. Let $C = \{x \in \mathbb{R}^3 : f_1(x) = 2, f_2(x) = 1\}$ denote this locus, a level-set of $f\colon \mathbb{R}^3 \to \mathbb{R}^2$, where $f = (f_1, f_2)$.

It is easy to check that C is bounded, and hence as any level-set is closed, it is compact. It follows g attains a maximum and minimum on C. By the Lagrange multiplier theorem, at such an extremum $c = (c_1, c_2, c_3)$ there must exist scalars $\lambda_1, \lambda_2 \in \mathbb{R}$ such that

$$Dg_c = \lambda_1 Df_{1,c} + \lambda_2 Df_{2,c},$$

and hence

$$(1, 1, 3) = \lambda_1(2c_1, 2c_2, 0) + \lambda_2(1, 1, 1).$$

Thus $\lambda_2 = 3$, and hence $2\lambda_1 c_1 = 2\lambda_1 c_2 = -2$. It follows that $c = (-\lambda_1^{-1}, -(\lambda_1)^{-1}, c_3)$. The constraint $f_1(c) = 2$ then implies $\lambda_1 = \pm 1$ so that since $f_2(c) = 1$ we see that if we set $c_\pm = (\pm 1, \pm 1, 1 \mp 2)$, the points $c_\pm$ are the only possibilities for extrema of g on C, and since we know g attains a maximum and minimum value, we see that $-1 = g(c_+) \leq g(x) \leq g(c_-) = 7$ for all $x \in C$.

Example 3.33. Let us prove the Cauchy-Schwarz inequality using Lagrange multipliers. Thus we wish to show that, for any two vectors $a, b \in \mathbb{R}^n$ we have $|a \cdot b| \leq \|a\|.\|b\|$. This is trivially true if either a or b is zero, so we may assume both are non-zero. But then we may rewrite the inequality as $(a/\|a\|) \cdot (b/\|b\|) \leq 1$. Since $a/\|a\|$ and $b/\|b\|$ are unit vectors, we are thus reduced to the following:

Problem: Maximize $x \cdot y$ for $x, y \in \mathbb{R}^n$ subject to the contraints that $\|x\| = \|y\| = 1$.

Let us formulate this in the language of Theorem 3.30. Let $g \colon \mathbb{R}^{2n} = X_1 \oplus X_2$ (the span of the first n and last n standard basis vectors respectively) be given by $g(x, y) = x \cdot y$ (thus we use the same notational conventions as in Theorem 3.22) and let $f \colon \mathbb{R}^{2n} \to \mathbb{R}^2$ be given by $f(x.y) = (x \cdot x, y \cdot y)$. We wish to maximize g subject to the condition that $(x, y) \in S = \{(x, y) \in \mathbb{R}^{2n} : f(x.y) = (1, 1)\}$.

Now S is clearly compact (as it is closed and bounded) hence g attains a maximum value on S. Now for any $(x, y) \in S$ we have $Df_{1,(x,y)} = 2(x, 0)$ and $Df_{2,(x,y)} = 2(0, y)$, and hence $\mathrm{rank}(Df_{(x_0,y_0)}) = 2$, so that S is a $2n - 2$-dimensional submanifold of $\mathbb{R}^{2n}$. Hence, by Theorem 3.30, if $p = (x_0, y_0)$ is a local maximum for g on S, there must exist scalars $\lambda_1, \lambda_2 \in \mathbb{R}$, not all zero, such that

$$Dg_{(x_0,y_0)} = \lambda_1 Df_{1,(x_0,y_0)} + \lambda_2 Df_{2,(x_0,y_0)}.$$

Now it is easy to see that $Dg_{(x_0,y_0)} = (y_0, x_0)$, hence the previous equation becomes

$$(y_0, x_0) = (2\lambda_1.x_0, 2\lambda_2.y_0),$$

so that, taking components in $\mathbb{R}^n$ and $\mathbb{R}^n_n$ we must have

$$y_0 = 2\lambda_1.x_0, \quad x_0 = 2\lambda_2.y_0.$$

But then we must have $y_0 = \lambda_1.x_0$ and $x_0 = \lambda_2.y_0$, so that $\lambda_1\lambda_2 = 1$, and since $\|x_0\| = \|y_0\| = 1$, we must have $|\lambda_1| = |\lambda_2| = 1$ and hence either $x_0 = y_0$ or $x_0 = -y_0$. Since $g(x_0, x_0) = \|x_0\| = 1$ and $g(x_0, -x_0) = -\|x_0\| = -1$, it follows immediately that $-1 \leq g(x, y) \leq 1$ on S and we obtain the equalities $g(x, y) = \pm 1$ if and only if $x = \pm y$.

4 Submanifolds of a normed vector space

4.1 Definition and basic properties

The goal of this section is to apply the inverse and implicit function theorems to geometry. The theorems allow us to show the equivalence of two natural definitions of a smooth surface in $\mathbb{R}^3$, and, more generally, define the notion of a *submanifold* of a normed vector space X.

Example 4.1. Let $S = \{x \in \mathbb{R}^3 : x_1^2 + x_2^2 + x_3^2 = 1\}$ is the standard unit sphere. It is smooth (in a sense that we have yet to make precise) and we can describe the points which lie on it in (at least) two ways. The first is implicit in the definition – a point $p = (x_1.x_2.x_3)$ lies in S if the function $f(x_1, x_2, x_3) = x_1^2 + x_2^2 + x_3^2$ evaluates to 1 on p, that is, S is a *level set* of the function f.

The second way to describe points on S is via a *parametrization*: for example, the map $\phi\colon [-1, 1] \times [-\pi, \pi) \to \mathbb{R}^3$ given by $(t, \theta) \mapsto (\cos(\theta).\sqrt{1-t^2}, \sin(\theta).\sqrt{1-t^2}, t)$ has S as its image, thus we can use the parameters (t, θ) to study S. Note that our parametrizing map ϕ is *not* injective, though it is on much of its domain. In general we will usually only be able to obtain parametrizations of a surface locally, that is, given a point p on our surface S, we will show that there is a diffeomorphism from an open subset U of $\mathbb{R}^2$ to an open subset V of our surface containing p.

On the other hand, if we only wish to obtain parametrizations for open subsets of a surface, we can often use the Implicit Function Theorem to turn the condition $f(x_1, x_2, x_3) = 0$ into an equation for one of the variables in terms of the others. For example, if $H_3 = \{x \in \mathbb{R}^3 : x_3 > 0\}$, then on $H_3 \cap S$ we may write S as the graph of $h(x_1, x_2) = \sqrt{1 - x_1^2 - x_2^2}$, that is, in H_3 we have $x \in S$ if and only if $S \in \operatorname{graph}(h) = \{(x_1, x_2, h(x_1, x_2)) : (x_1, x_2) \in V\}$, where $V = \{(x_1, x_2) \in \mathbb{R}^2 : x_1^2 + x_2^2 < 1\}$.

Definition 4.2. Let $M \subseteq X$ be a closed subset of an n-dimensional normed vector space X. We say that M is a k-dimensional *submanifold* of X if, for every point $p \in M$, there is an open subset U of X containing p and a smooth[16] function $f\colon U \to Y$, where Y is an $(n-k)$-dimensional normed vector space, such that $M \cap U = f^{-1}(0)$, and at each $p \in M \cap U$ the derivative Df_p has maximal rank, that is $\operatorname{rank}(Df_p) = n - k$.

We say that M is C^k if we can choose $f \in C^k(U, Y)$ where $k \in \mathbb{N} \cup \{\infty\}$. If $k = \infty$ we say M is a *smooth* submanifold of $\mathbb{R}^n$.

Informally, this definition says that, locally (*i.e.* near any given point of M) the submanifold is given as the level-set of $n - k$ smooth functions (the components of f) which are not "tangent to each other" – this last requirement being captured by the rank condition.

The Implicit Function Theorem allows us to relate this definition to the second method of understanding surfaces discussed above, namely, via parametrizations. In the next theorem, for $k \leq n$ we view $\mathbb{R}^k$ as a subspace of $\mathbb{R}^n$ spanned by $\{e_1, \ldots, e_k\}$.

Theorem 4.3. *Let M be a k-dimensional submanifold of an n-dimensional normed vector space X, and let $p \in M$. Then there is a direct sum decomposition $X = X_1 \oplus X_2$ where $\dim(X_1) = k$, $\dim(X_2) = n - k$, and open neighbourhoods V and $U_1 \times U_2$ of p and 0_X respectively, where for $i = 1, 2$, U_i is an open subset of X_i, and a diffeomorphism $\psi\colon U_1 \times U_2 \to V$ such that $M \cap V = \psi(U_1 \times \{0_{X_2}\})$. In particular, $\psi_{|U_1 \times 0_{X_2}}\colon U_1 \to M \cap V$ gives a parametrization of $M \cap V$.*

Proof. By definition, there is an open set V_1 containing p and a function $f\colon V \to \mathbb{R}^{n-k}$ such that $V_1 \cap M = \{x \in V : f(x) = 0_{n-k}\}$, and $\operatorname{rank}(Df_x) = n - k$ for all $x \in V_1$. But then Theorem 3.22 shows that there is a diffeomorphism $\psi\colon U \to V \subseteq V_1$, where U an open neighbourhood of 0_n and $V_1 \subseteq V$ is an open neighbourhood of p, such that in the coordinate system $(t_1, \ldots, t_n)$ given by $t_i = x_i \circ \psi^{-1}$, the function f is given by $(t_{k+1}, \ldots, t_n)$ (that is, for $v \in V_1$, we have $f(v) = (t_{k+1}(v), \ldots, t_n(v))$). Moreover, the functions $(t_1, \ldots, t_k)$ parameterise the submanifold M on the open subset $M \cap V$ of M: if $(t_1, \ldots, t_k, 0, \ldots, 0) \in \mathbb{R}^k \cap U$, and we set $\phi(t_1, \ldots, t_k) = \psi(t_1, \ldots, t_k, 0, \ldots, 0)$ then $\phi(t_1, \ldots, t_k) \in M \cap V$ and if $u \in M \cap V$ then $u = \phi(t_1, \ldots, t_k)$ for $t_i = x_i \circ \psi^{-1}$.

□

[16] At least continuously differentiable, but many texts automatically assume infinitely differentiable.

Remark 4.4. The Implicit Function Theorem shows that, at least locally, a submanifold M can be viewed as the graph of a C^1 function. To put it another way, let us define a k-dimensional *subgrapholdt*[17] of a normed vector space X to be a subset $M \subseteq X$ such that, for any point $a \in M$, there is an open neighbourhood U of a together with a decomposition $X = X_1 \oplus X_2$ with $\dim(X_1) = k$, and a function $\psi \in C^1(U \cap (a + X_1), V_2)$ such that $M \cap U = \Gamma(\psi)$, where $\Gamma(\psi) = \{(v, \psi(v)) : v \in U \cap (a + V_1)\}$ is the graph of ψ. In this terminology, the previous discussion shows that that any k-submanifold of X is a k-subgraphold. In fact the converse is also true: indeed, as we show in Lemma 4.7 below, if $V = V_1 \oplus V_2$ and $\phi \in C^1(\Omega_1, V_2)$ for some open subset $\Omega_1 \subseteq V_1$ of V_1, then $\Gamma(\phi)$, the graph of ϕ, is always a submanifold of V.

Thus the two notions – that of submanifold and subgraphold are equivalent, and we can use either local description to study submanifolds. One advantage of the definition in terms of level-sets is that it does not require introducing an auxiliary decomposition of $\mathbb{R}^n$ into a direct sum.

***Remark 4.5.** Our definition of a k-dimensional sub-manifold M is a subset of a normed vector space X which is locally given as a level-set for a C^1-function f taking values in an $(n-k)$-dimensional vector space Y for which Df_x has rank $n-k$. Theorem 4.3 shows that, if M is a submanifold, then M is locally given as the image of a C^1-map ψ from an open subset V of a k-dimensional normed vector space Z, where $D\psi$ has rank k. This is, *a priori* strictly weaker, since the domain V is not identified with an open subset of a subspace X_1 of X in such a way that the image of ψ takes values in a complementary subspace.

Nevertheless, it turns out to be true that if $M \subseteq X$ is locally given as the image of an injective C^1-map from a suitable open subset V of a k-dimensional normed vector space Z whose derivative has rank k at each point of V, then M is a sub-manifold in the sense of Definition 4.3: More precisely, if $V \subseteq \mathbb{R}^k$ is an open subset of $\mathbb{R}^k$ and $\psi \in C^1(V, \mathbb{R}^n)$ we say that ψ is an *immersion* if $\operatorname{rank}(D\psi_p) = k$ for all $p \in V$. The *immersion criterion* states that a subset $M \subseteq \mathbb{R}^n$ is a k-submanifold in the sense of Definition 4.2 if, for every $a \in M$ there is a neighbourhood U_a of a, and an immersion $\psi \in C^1(B(0_k, r), \mathbb{R}^n)$ from an open ball of radius $r > 0$ centred at $0_k \in \mathbb{R}^k$ such that $\psi(0_k) = a$ and $M \cap U_a = \operatorname{im}(\psi)$. For more details on this see Appendix 5.4.

Example 4.6. Suppose that $g\colon \mathbb{R}^2 \to \mathbb{R}$ is given by $g(x_1, x_2) = x_1x_2$. Then $Dg_{(x_1,x_2)} = (x_2, x_1)$ and hence $\operatorname{rank}(Dg_{(x_1,x_2)}) = 1$ unless $(x_1, x_2) = (0, 0)$. Then for all $c \neq 0$, the level-sets $L_c = g^{-1}(c)$ are smooth 1-submanifolds of $\mathbb{R}^2$, but $L_0 = g^{-1}(0) = \{(x, 0) : x \in \mathbb{R}\} \cup \{(0, y) : y \in \mathbb{R}\}$, which is not smooth at the origin $(0, 0)$, exactly the point where Dg fails to have maximal rank.

On the other hand, if V_1 and V_2 are normed vector spaces and $\psi \in C^1(U, V_2)$ is a continuously differentiable function on an open subset U of V_1 taking values in V_2, then if we set

$$\Gamma(\psi) = \{(v, \psi(v) : v \in U\} \subset V = V_1 \oplus V_2,$$

then the following Lemma shows that $\Gamma(\psi)$ is always a submanifold of V.

Lemma 4.7. *Let X_1, X_2 be finite-dimensional normed vector spaces, and suppose that $\psi \in C^1(\Omega, X_2)$ is a continuously differentiable function on an open subset Ω of X_1 taking values in X_2. Then the graph $\Gamma(\psi) = \{(v, \psi(v)) : v \in \Omega\}$ is a submanifold of $X = X_1 \oplus X_2$.*

Proof. But if we let $g\colon \Omega_1 \times X_2 \to X_2$ be given by $g(v_1, v_2) = v_2 - \phi(v_1)$, then clearly $g \in C^1(\Omega_1 \times X_2, X_2)$ and $(v_1, v_2) \in \Gamma(\phi)$ if and only if $g(v_1, v_2) = 0$. Moreover, if $a = a_1 + a_2 \in X_1 \oplus X_2$, then $Dg_{(a_1,a_2)}(v_1, v_2) = -D\phi_{a_1}(v_1) + v_2$. Thus for any $v_2 \in X_2$ and any $a \in \Omega_1 \times X_2$ we have $Dg_a(0, v_2) = v_2$, and hence the derivative $Dg_{(a_1,a_2)}$ is surjective for all $a \in \Omega_1 \times X_2$. Thus $\Gamma(\phi)$ is a k-submanifold of $\mathbb{R}^n$, where $k = \dim(X_1)$. □

Example 4.8. The simplest case of the previous Lemma is when $V_1 = \mathbb{R}^n$ and $V_2 = \mathbb{R}$, so that $C^1(U, V_2) = C^1(U, \mathbb{R})$ is just the space of real-valued continuously differntiable functions on an open subset U of $\mathbb{R}^n$. If f is such a function, we can then view $\Gamma(f) = \{(x, f(x) : x \in U\}$ as a subset of $\mathbb{R}^{n+1} = \mathbb{R}^n \oplus \mathbb{R}$. Writing a point in $\mathbb{R}^{n+1}$ as (x, y) where $x \in \mathbb{R}^n$ and $y \in \mathbb{R}$, we see immediately that $\Gamma(f) = \{(x, y) \in U \times \mathbb{R} : g(x, y) = 0$ where $g(x, y) = y - f(x)$. Since $Dg_{(x,f(x))}$ has Jacobian matrix $(-\partial_1 f(x), \dots, -\partial_n f(x), 1)$, clearly $Dg_{(x,f(x))}$ always has rank 1, and so $\Gamma(f)$ is an n-submanifold of $\mathbb{R}^{n+1}$

[17]The term is completely non-standard, and therefore, to honest, deliberately chosen to be clunky.

Example 4.9. Suppose that $n \in \mathbb{R}^3$ is a unit vector and

$$C = \{x \in \mathbb{R}^3 : x_1^2 + x_2^2 - x_3^2 = 0, \langle n, x \rangle = d\}.$$

Then C is a level set of the function $f \colon \mathbb{R}^3 \to \mathbb{R}^2$, where f has components $f_1(x) = x_1^2 + x_2^2 - x_3^2$ and $f_2(x) = \langle n, x \rangle = n_1 x_1 + n_2 x_2 + n_3 x_3$: indeed $C = f^{-1}(\{(0, d)\})$. Now

$$Df_x = \begin{pmatrix} 2x_1 & 2x_2 & -2x_3 \\ n_1 & n_2 & n_3 \end{pmatrix}$$

hence Df has rank 2 on the complement of the line $\mathbb{R}.(n_1, n_2, -n_3)$. If $d = 0$ then clearly $0 \in C$ and Df_0 has rank 1, so we will suppose that $d \neq 0$. But then it is easy to check the line $\mathbb{R}.(n_1, n_2, -n_3)$ does not intersect the level set C, and hence Df has rank 2 at every point of C, and so C is a 1-dimensional submanifold of $\mathbb{R}^3$.

Suppose we wish to parameterize the curve C. The Implicit Function Theorem in the form of Theorem 3.22 shows that, at least locally we can write it as the graph of any one of our coordinates x_1, x_2, x_3. In fact, by rotating around the x_3-axis, we may assume that $n = (n_1, 0, n_3)$, and hence we may write $n = (\cos(\phi), 0, \sin(\phi))$ for some $\theta \in \mathbb{R}$. Then C is given by the system of equations:

$$\begin{aligned} x_2^2 = x_3^2 - x_1^2 &= (x_3 - x_1)(x_3 + x_1), \\ \cos(\phi)x_1 + \sin(\phi)x_3 &= d. \end{aligned}$$

If $\cos(\phi) = 0$, it is easy to see that C is just one of the circles $C_{\pm d} = \{(x_1, x_2, \pm d) : x_1^2 + x_2^2 = d^2\}$, so assume $\cos(\phi) \neq 0$. Moreover, if $\cos(\phi) = \sin(\phi)$ then C is clearly a parabola with parametrization $s \mapsto (d_1 + (s/2d_1)^2, s, d_1 - (s/2d_1)^2)$, where $d_1 = d/\sqrt{2}$. Otherwise, writing $\ell = d/\cos(\phi)$, we have $x_1 = \ell - \tan(\phi)x_3$, and hence our equations become

$$x_2^2 = ((1 + \tan(\phi))x_3 - \ell)((1 - \tan(\phi))x_3 + \ell) = (1 - \tan(\phi)^2)x_3^2 + 2\ell\tan(\phi).x_3 - \ell^2$$

Since $\ell = d/\cos(\phi) \neq 0$, then the quadratic on the right is non-negative on $I_\phi = \mathbb{R}\backslash(-2, 2)$ when $\tan(\phi) < 1$ and non-negative on $I_\phi = [2, 2]$ when $\tan(\phi) > 1$. and hence writing $t = \tan(\phi)$ we obtain a parameterization:

$$\begin{aligned} C &= \{(\ell - t.s, \pm\sqrt{(1 - t^2).s^2 + 2t\ell.s - \ell^2}, s) : s \in I_\phi\} \\ &= \{(1 - t.s, \pm\sqrt{(1 - t^2)s^2 + 2t.s - 1}, s) : s \in \ell.I_\phi\}. \end{aligned}$$

Thus we obtain ellipses or hyperbolas for $\tan(\phi) > 1$ and $\tan(\phi) < 1$ respectively. The signs which occur, as before, are determined, for example, by choosing a point $p \in C$ around which we wish to obtain a local parameterization.

Of course the Implicit Function Theorem can also be applied starting with different local coordinates at a point $p \in C$: Indeed it might, given the nature of f, be more sensible to start with the cylindrical polar coordinates $\rho(r, \theta, z) = (r\cos(\theta), r\sin(\theta), z)$: In these coordinates the level-set C becomes $\{p \in \mathbb{R}^3 : r^2 - z^2 = 0, r\cos(\theta)\cos(\phi) + z\sin(\phi) = d\}$, where $p = \rho(r, \theta, z) = (r(p), \theta(p), z(p))$.

Note that the derivative of $f = (f_1, f_2)$ with respect to these coordinates is

$$Df_{(r,\theta,z)} = \begin{pmatrix} 2r & 0 & -2z \\ \cos(\theta)\cos(\phi) & -r\sin(\theta)\cos(\phi) & \sin(\phi) \end{pmatrix}.$$

and so has rank 2 provided $r \neq 0$ and $\theta \neq n\pi$ (when $\cos(\phi) \neq 0$),

The level set $f_1(p) = 0$ is thus parameterized by $(s_1, s_2) \mapsto (s_1\cos(s_2), s_1\sin(s_2), s_1) \in \mathbb{R}^3$, or equivalently[18] $(s_1, s_2) \mapsto \rho(s_1, s_2, s_1)$, for $(s_1, s_2) \in \mathbb{R}^2$. Since the case $\cos(\phi) = 0$ is equally easy to handle in this setting, we assume $\cos(\phi) \neq 0$, and again set $\ell = d/\cos(\theta)$. We then find that C can be parameterized by $s \in \mathbb{R}$ via

$$s \mapsto \rho(r(s), \theta(s), z(s)) = \rho(\frac{\ell}{\tan(\phi) + \cos(s)}, s, \frac{\ell}{(\tan(\phi) + \cos(s))}).$$

Thus recovering the polar form for the equations of a parabola, ellipse or hyperbola. One can also determine the differential equation the function $g(s) = (r(s), z(s))$ must satisfy, as we did in Example 3.2, which can be solved in this case by separation of variables.

[18] If $z < 0$ then this shifts s_2 by π from the normal convention of $r > 0$.

4.2 Tangent spaces and normal vectors

We now wish to define the notion of tangent vectors and normal vectors at a point in a submanifold of a finite-dimensional inner product space E.

Definition 4.10. Let S be a subset of a normed vector space X and let $p \in S$. A *path on S centred at p* is a function $\gamma \in C^1((-r, r), X)$, where $r > 0$, such that the image of γ lies in S and $\gamma(0) = p$. We write $\mathcal{P}(S, p)$ for the set of all paths on S centred at p. Let $T : \mathcal{P}(S, p) \to X$ be the map given by $T(\gamma) = \gamma'(0)$. The image of T is called the *tangent space* to S at p and is denoted T_pS.

If V is an inner product space, we can also define $T_pS^\perp = \{n \in X : \langle n, v\rangle = 0, \forall v \in T_pS\}$, the *normal space* to S at p. This space is also sometimes denoted N_pS.

Remark 4.11. Note that while the normal space N_pX is by definition a linear subspace of X, the tangent space need not in general be a linear subspace (see Example 4.16). Indeed since $T_pS \subseteq (T_pS^\perp)^\perp = N_pS^\perp$ with equality if and only if T_pS is itself a linear subspace of X. Thus $N_pS^\perp$ is the smallest subspace of X containing T_pS, that is, $N_pS^\perp$ is the linear span of T_pS. We will shortly see that $T_pS = N_pS^\perp$ when S is a submanifold.

Remark 4.12. Let X be a normed vector space and $R \subseteq S \subseteq X$ be subsets. For any $p \in R$ clearly $\mathcal{P}(R, p) \subseteq \mathcal{P}(S, p)$ and hence $T_pR \subseteq T_pS$.

Slightly less trivially, if $p \in S$ and U is an open subset containing p, then $T_p(U \cap S) = T_pS$. Since $S \cap U \subseteq S$, by the above we see that $T_p(U \cap S) \subseteq T_pS$. For the reverse inclusion, note that if $v \in T_pS$ then we may pick a path $\gamma \in \mathcal{P}(S, p)$ with $T(\gamma) = v$. Then γ is continuous, so $\gamma^{-1}(U)$ is an open neighbourhood of 0 (since $\gamma(0) = p$) and so contains an open interval of the form $(-s, s)$. Let $\gamma_s = \gamma_{|(-s,s)}$. Then $\gamma_s \in \mathcal{P}(S \cap U, p)$, and, since it is the restriction of γ to an open set containing 0. $T(\gamma_s) = (\gamma_s)'(0) = \gamma'(0) = v$, and hence $v \in T_p(U \cap S)$.

Thus the tangent space T_pS of S at p is only sensitive to the nature of S near p. This simple observation, along with the Chain Rule, gives us the following Lemma, which although easy to prove, will be the key tool in calculating with tangent spaces.

Lemma 4.13. *Let X and Y be a normed vector spaces and let U be an open subset of X and let S be an arbitrary subset of X. If $\psi \in C^1(U, Y)$, and $p \in U \cap S$, then if $R \subseteq Y$ is such that $\psi(U \cap S) \subseteq R$, and $q = \psi(p)$, the derivative of ψ at p induces a map*

$$D\psi_p : T_pS \to T_qR$$

Proof. Let $v \in T_pS$. By Remark 4.12, we may assume that $v = T(\gamma)$ for $\gamma \in \mathcal{P}(X \cap U, p)$. But then $\psi \circ \gamma \in \mathcal{P}(\psi(U \cap S), b) \subseteq \mathcal{P}(R, q)$, so that $T(\psi \circ \gamma) \in T_qR$. But by the Chain Rule,

$$T(\psi \circ \gamma) = (\psi \circ \gamma)'(0) = D\psi_a(\gamma'(0)) = D\psi_a(v),$$

so that $D\psi_p(v) \in T_qR$ as required. □

Corollary 4.14. *Let X and Y be normed vector spaces, U an open subset of X, and S any subset of X. Suppose that $\psi \in C^1(U, Y)$ and $p \in U \cap S$. Then we have the following:*

1. *If $D\psi_p$ is an invertible linear map, then $D\psi_p$ gives a bijection between T_pS and T_qR, where $q = \psi(p)$ and $R = \psi(U \cap S)$.*
2. *If $\psi(X) = q$ then $T_pX \subseteq \ker(D\psi_p)$.*

Proof. Since $D\psi_p$ is invertible, the Inverse Function Theorem shows that ψ induces a diffeomorphism from a neighbourhood U_1 of p to Ω, and open subset of W containing $q = \psi(p)$. But then if $\theta : \Omega \to U_1$ is the inverse of ψ, by Lemma 4.13 applied to ψ and θ, we have $D\psi_p : T_pX \to T_qY$ and $D\theta_q : T_qY \to T_pX$, and $D\psi_p$ and $D\theta_q$ are inverse, the result follows.

For the second part, Lemma 4.13 shows that $D\psi_p(T_pX) \subseteq T_q\{q\}$. But clearly $\mathcal{P}(\{q\}, q)$ consists of the constant maps γ which take the value q, and hence have derivative 0. It follows that $T_q(\{q\}) = \{0\}$, and hence that $T_pX \subseteq \ker(D\psi_p)$. □

Example 4.15. If M is a k-submanifold of X, so that for any $a \in M$ we can find an open neighbourhood U of a such that $U \cap M = f^{-1}(0)$ for some $f \in C^1(U, \mathbb{R}^{n-k})$ for which Df_x has rank $n-k$ for all $x \in U$. Using Example 4.12 and Corollary 4.14 part (2), we see that

$$T_pM = T_p(U \cap M) = T_p(f^{-1}(0)) \subseteq \ker(Df_p).$$

If X is a subset of V and U is a neighbourhood of $a \in X$ such that $X \cap U = f^{-1}(0)$ for some $f \in C^1(U, \mathbb{R}^m)$, the containment $T_pX \subseteq \ker(Df_p)$ can, in general, be strict. However, when M is a submanifold of $\mathbb{R}^n$ locally defined by the vanishing of f, then we will shortly see that $T_pM = \ker(Df_p)$.

Example 4.16. Consider Example 4.6 again, that is, let $g\colon \mathbb{R}^2 \to \mathbb{R}$ the continuously differentiable function given by $g(x_1, x_2) = x_1.x_2$, and, for $c \in \mathbb{R}$ let $L_c = \{(x_1, x_2) \in \mathbb{R}^2 : x_1.x_2 = c\}$. Then $Dg_{(a_1,a_2)} = (a_2, a_1)$, which has maximal rank (*i.e.* rank 1) provided $a = (a_1, a_2) \neq 0$. Thus for any $a \neq 0$, if $g(a) = c$ Corollary 4.14 shows that $T_a(L_c) \subseteq \ker(Dg_a) = \{(x_1, x_2) : a_2x_1 + a_1x_2 = 0\}$, while at $a = 0$ we only get the trivial bound $T_0L_0 \subseteq \ker(Dg_0) = \mathbb{R}^2$. In fact you can check that $T_aL_c = \ker Dg_a$ for all $a \neq 0$, while at $a = 0$, $T_0L_0 = L_0$, giving an example where the tangent space of a level-set is not a linear subspace.

Example 4.17. Now case where $M = \{x \in \mathbb{R}^n : x_l = 0, \forall l > k\}$ and $p = 0_n$. Then M is defined by the vanishing of $f(x) = (x_{k+1}, \ldots, x_n\}$. Then it is clear that Df_0 has kernel given by $\operatorname{span}_{\mathbb{R}}\{e_1, \ldots, e_k\}$. On the other hand, if $v = (v_1, \ldots, v_k, 0, \ldots, 0)$, then $\gamma(t) = t.v$ lies in M, and $\gamma'(0) = v$, hence we see that $v \in T_0M$ if and only if $Df_0(v) = 0$.

The above example along with the Implicit Function Theorem shows the following:

Proposition 4.18. *Let M be a k-dimensional submanifold of $\mathbb{R}^n$ and let $p \in M$. Then if U is an open subset of $\mathbb{R}^n$ such that $M \cap U = f^{-1}(0)$, where $f\colon U \to \mathbb{R}^{n-k}$ is continuously differentiable with Df_x of maximal rank for all $x \in U$. Then we have*

$$T_pM = \ker(Df_p).$$

In particular, T_pM is a k-dimensional vector subspace.

Proof. We have already shown the containment $T_pM \subseteq \ker(Df_p)$ in Corollary 4.14, so it remains to establish the reverse inclusion. In the case where $f = (x_{k+1}, \ldots, x_n)$ this was shown in the previous Example, but the Implicit Function Theorem shows us that, for any point $p \in M$, we can find a diffeomorphism $\psi\colon V \to U$ from an open neighbourhood V of 0_n to an open neighbourhood U of p taking $N \cap V$ to $M \cap U$ where $N = \{x \in U : (x_{k+1}, \ldots, x_n) = 0_{n-k}\}$. The result then follows from Lemma 4.13. □

Using the notion of gradient vector fields, we can also describe the normal space $T_pM^\perp$ of a k-dimensional submanifold:

Proposition 4.19. *Suppose that M is a k-dimensional submanifold and $p \in M$. If U is an open neighbourhood of p such that $M \cap U$ is given by $f^{-1}(0)$ where $f\colon U \to \mathbb{R}^{n-k}$ is a continuously differentiable function, then if $f = (f_1, \ldots, f_{n-k})$ we have*

$$T_pM^\perp = \operatorname{span}_{\mathbb{R}}\{\nabla f_1(p), \ldots, \nabla f_{n-k}(p)\}.$$

In particular $T_pM^\perp$ is a vector space of dimension $n-k$.

Proof. By Proposition 4.18, the tangent space $T_pM = \ker(Df_p)$ is a k-dimensional subspace of $\mathbb{R}^n$. Let $f = (f_1, \ldots, f_{n-k})$ and let $N = \operatorname{span}_{\mathbb{R}}\{\nabla f_1(p), \ldots, \nabla f_{n-k}(p)\}$, an $(n-k)$-dimensional subspace. Now the rows of the Jacobian matrix of Df_p are given by $\nabla f_i(p)^T$, so that

$$Df_p(v) = \sum_{i=1}^{n-k} (\nabla f_i(p) \cdot v)e_i$$

It follows that $v \in T_pM$ if and only if $v \in N^\perp$. Thus $T_pM = N^\perp$ and hence $N = T_pM^\perp$ as required (since, for any subspace W of an inner product space V we have $(W^\perp)^\perp = W$). □

Example 4.20. Let $S = \{(x_1, x_2, x_3) \in \mathbb{R}^3 : x_1^2 + 2x_2^2 - 7x_3^2 = 1\}$. Then if $f(x) = x_1^2 + 2x_2^2 - 7x_3^2$, the surface S is a level-set of f. Since $\nabla f(x) = (2x_1, 4x_2, -14x_3)$, the function f has maximal rank (*i.e.* rank 1) everywhere except 0, and since $0 \notin S$, it follows that S is a 2-dimensional submanifold of $\mathbb{R}^3$. The tangent and normal spaces to S at a point $a = (a_1, a_2, a_3)$ is then

$$T_aS = \{v = (v_1, v_2, v_3) \in \mathbb{R}^3 : 2a_1.v_1 + 4a_2.v_2 - 14a_3.v_3 = 0\},$$
$$T_pS^{\perp} = \{\lambda.(2a_1, 4a_2, -14a_3) : \lambda \in \mathbb{R}\}$$

Example 4.21. Let $\mathrm{O}_n(\mathbb{R}) = \{X \in \mathrm{Mat}_n(\mathbb{R}) : X.X^T = I_n\}$ be the orthogonal group, the group of linear isometries of $\mathbb{R}^n$ (equipped with the $\|.\|_2$-norm). We claim this is a smooth submanifold of $\mathrm{Mat}_n(\mathbb{R})$ of dimension $n(n-1)/2$.

Now the definition of $\mathrm{O}_n(\mathbb{R})$ shows that it is a level-set of the function $q(X) = X.X^T$, which has entries which are degree two polynomials in the entries of X. Thus $q(X)$ is clearly continuously differentiable, and moreover $Dq_X(H) = X.H^T + H.X^T$, since

$$q(X + H) = (X + H).(X + H)^T = q(X) + H.X^T + X.H^T + H.H^T,$$

and $\|H.H^T\|_\infty \leq \|H\|_\infty.\|H^T\|_\infty$ so that $\|H\|_\infty^{-1} H.H^T \to 0$ as $H \to 0$ (since clearly $H^T \to 0$ as $H \to 0$).

Now $(X.X^T)^T = X.X^T$, so the image of q lies in the linear subspace $S(\mathbb{R}^n)$ of symmetric matrices in $\mathrm{Mat}_n(\mathbb{R})$, which is a subspace of dimension $n(n+1)/2$. Thus it will follows that $\mathrm{O}_n(\mathbb{R})$ is a submanifold of dimension $n(n-1)/2$ if we can show that Dq_X is a surjective linear map from $\mathrm{Mat}_n(\mathbb{R})$ to $S(\mathbb{R}^n)$. But if $C \in S$ then $(CX)^T = X^T.C = X^{-1}.C$, so that

$$Dq_X(\frac{1}{2}(C.X)) = \frac{1}{2}(C.X.X^T + X.(C.X)^T) = \frac{1}{2}(C.I_n + I_n.C) = C,$$

so that Dq is surjective as required.

The group $\mathrm{O}_n(\mathbb{R})$ is thus what is known as a *Lie group*. Its tangent space at the identity I_n is denoted by $\mathfrak{o}_n(\mathbb{R})$. Explicitly this is $\ker(Dq_{I_n}) = \{H \in \mathrm{Mat}_n(\mathbb{R}) : H + H^T = 0\}$. It carries a kind of non-associative product, called a *Lie bracket*: If $H_1, H_2 \in \mathfrak{o}_n(\mathbb{R})$ then you can check that $[H_1, H_2] = H_1H_2 - H_2H_1 \in \mathfrak{o}_n(\mathbb{R})$. The Lie algebra structure gives a kind of "infinitesimal" or deriviative of the group structure on $\mathrm{O}_n(\mathbb{R})$. This is studied in detail in courses in Part C.

Remark 4.22. Now that we have the language of tangent spaces and submanifolds, we can reinterpret the theory of Lagrange multipliers in more geometric terms: if U is an open subset of a normed vector space X and $f \in C^1(U, Y)$ is a constraint function and we seek to minimize $g(x)$ on the locus $C = \{x \in U : f(x) = 0\}$.

If $a \in C$ and ∇g_a has a non-trivial component in T_aC, then the same argument as the one used in Lemma 3.28 shows that a cannot be a local minimum (one must use a path γ centred at a lying on S which has $T(\gamma)$ equal to the projection of ∇g_a onto T_aC, but with this extra detail the same strategy works). It follows that a necessary condition for $a \in C$ to be a local minimum is that ∇g_a is normal to C at a. Provided that Df has maximal rank on C, if $f = \sum_{i=1}^k f_i.w_i$ for $\{w_1, \ldots, w_k\}$ some basis of Y, then Proposition 4.19 shows that this is equivalent to $\nabla g_a \in \mathrm{Span}\{\nabla f_i(a) : 1 \leq i \leq k\}$, and so we recover the theorem on Lagrange multipliers.

4.3 *Abstract Manifolds

Suppose that M is a k-dimensional submanifold of $\mathbb{R}^n$. If V is an open neighbourhood of a point $p \in M$, then there is an open subset of $\mathbb{R}^n$ with $V = M \cap U$. Shrinking V and U is necessary, we can find a diffeomorphism $\psi\colon B(0, r) \to U$ such that $\psi(V \cap (\mathbb{R}^k \oplus 0_{n-k})) = M \cap U$. If we write $\psi^{-1}(x) = (t_1, \ldots, t_n)$, then if $f\colon M \cap U \to \mathbb{R}$ is any function, we may define $\tilde{f}\colon U \to \mathbb{R}$ by

$$\tilde{f}(x) = f \circ (\psi(t_1, \ldots, t_k, 0, \ldots, 0)).$$

If $x \in M \cap U$ then $\tilde{f}(x) = f(x)$, so that $\tilde{f}$ extends f to a function on U an open subset of $\mathbb{R}^n$. We then say that f is C^1 at $x \in M \cap U$ if $\tilde{f}$ is. Using the chain rule, one can check that this definition is independent of the choice of diffeomorphism ψ. In effect, f is differentiable at $x \in M \cap U$ if it is differentiable as a function of the parameters $(t_1, \ldots, t_k)$. Thus the crucial fact is that we can equip M, at least locally, with "C^1-coordinates".

There is a notion of an abstract differentiable k-dimensional manifold: This is a topological space M, equipped with a collection of "charts" $\{\phi_i\colon U_i \to V_i : i \in I\}$, where the collection $\{V_i : i \in I\}$ forms an open cover of M (that is, $M = \bigcup_{i\in I} V_i$ and each V_i is an open subset of M) the U_i are open subsets of $\mathbb{R}^k$, and the ϕ_i are homeomorphisms. The charts allow us to say when a function $f\colon M \to \mathbb{R}$ is continuously differentiable: if $x \in M$, we say f is differentiable at $x \in M$ if $f \circ \psi_i$ is differentiable at $\psi_i^{-1}(x)$, where $i \in I$ is such that $x \in V_i$. In order for this definition to be consistent, the charts must satisfy a compatibility condition: if $x \in V_i \cap V_j$ lies in the image of two charts ψ_i and ψ_j we need $f \circ \psi_i$ to be differentiable at $\psi_i^{-1}(x)$ if and only if $f \circ \psi_j$ is C^1 at $\psi_j^{-1}(x)$. But by the chain rule, this follows if $\psi_j^{-1} \circ \psi_i\colon U_i \cap U_j \to U_i \cap U_j$ is diffeomorphism, and this is exactly the compatibility condition which is imposed. Abstract differentiable manifolds are studied in the Part C course "Differentiable Manifolds".

References

[G] B. Green, *Metric spaces*, lecture notes for A2 "Metric spaces and complex analysis", MT 2020, available here.

[S] M. J. Spivak, *Calculus on Manifolds*.

[R] J. Robbin, *On the existence theorem for differential equations*, Proc. Amer. Math. Soc. 19 (1968), 1005–1006. Available here.

5 Appendix

5.1 Notation: o and O

Definition 5.1. Let X and Y be normed vector spaces. Let $\mathcal{N}(X,Y)$ be the vector space of functions $f\colon D \to Y$ whose domain of definition $D \subseteq X$ is a neighbourhood of 0_X and let $\mathcal{N}_0(X,Y)$ be the subspace of $\mathcal{N}(X,Y)$ consisting of those functions $f \in \mathcal{N}(X,Y)$ which are continuous at 0_X and satisfy $f(0_X) = 0_Y$. Note that if $f\colon D_1 \to Y$ and $f_2\colon D_2 \to Y$, then their sum $f_1 + f_2$ is only defined on $D_1 \cap D_2$, but this is still a neighbourhood of 0_X, so that $\mathcal{N}(X,Y)$ is indeed a vector space. In fact, the same observation shows that if $c \in \mathcal{N}(X,\mathbb{R})$ and $f \in \mathcal{N}(X,Y)$ then $c.f \in \mathcal{N}(X,Y)$, and if $f \in \mathcal{N}_0(X,Y)$ so is $c.f$.

If g is a non-negative function in $\mathcal{N}(,\mathbb{R})$ then we will write $O_Y(g)$ for the subspace of $\mathcal{N}(X,Y)$ consisting of those functions $f\colon D \to Y$ for which there exists a constant $C > 0$ and an open ball $B(0_X,r) \subseteq D$ such that

$$\|f(x)\| \leq C.g(x), \quad \forall x \in B(0_X,r).$$

Note that if $g \in \mathcal{N}_0(X,\mathbb{R})$ it follows that $f \in \mathcal{N}_0(X,W)$ also, that is if $g \in \mathcal{N}_0(X,\mathbb{R})$ the $O_W(g) \subseteq \mathcal{N}_0(X,Y)$.

Similarly we write $o_Y(g)$ for the subspace of $\mathcal{N}(X,Y)$ consisting of those functions $f\colon D \to Y$ for which, given any $\epsilon > 0$, there is some $\delta > 0$ such that for all $x \in B(0_X,\delta)$ we have $\|f(x)\| \leq \epsilon.g(x)$. If g is non-vanishing in a neighbourhood of 0_X (except perhaps at 0_X itself) then this is equivalent to the condition that

$$\lim_{v \to 0_V} \frac{\|f(x)\|}{g(x)} = 0.$$

Notice that, again assuming g is non-vanishing on $B(0_X,r)\backslash\{0_X\}$ for some $r > 0$, if we set $f_1(x) = g(x)^{-1}.f(x)$ for $x \neq 0$ and $f_1(0_X) = 0_Y$, then by assumption f_1 defines an element of $\mathcal{N}_0(V,W)$, so that we may equivalently view $o_Y(g) = g.\mathcal{N}_0(V,W)$.

By a standard abuse of notation, we will write $f_1(x) = f_2(x) + o_Y(g)$ to mean $f_1(x) - f_2(x) \in o_Y(g)$, and similarly for $f_1(x) = f_2(x) + O(g)$. Note that if the target space Y is clear from the context, we will omit the subscript W and simply write $o(g)$ or $O(g)$.

Remark 5.2. Note that the functions in $O_Y(g)$ can, informally, be considered as those functions $f(x)$ for which $f(x) \to 0_Y$ as $x \to 0_X$ "at the same rate" as $g(x) \to 0$, while the functions in $o_W(g)$ tend to 0_Y "faster" than g tends to 0.

The easiest case to consider here is if g is continuous and $g(0) > 0$. Then, by continuity, $0 < g(0)/2 < g(x) < 3g(0)/2$ on some small ball $B(0_X,r)$ say, and hence $f \in O_Y(g)$ precisely if it is bounded near 0_X, while $f \in o_Y(g)$ precisely when $f(x) \to 0_Y$ as $x \to 0_X$.

5.2 *Multilinear maps and higher derivatives

In this section we describe how one can understand the higher derivatives of a function $f: U \to W$ without partial derivatives. The main point is to obtain a better understanding of the space in which $D^k f$ takes values when $k > 1$. Example 2.37 shows how the space $\mathcal{L}(V, \mathcal{L}(V, \mathbb{R}))$ is equivalent to the space $\mathrm{Bil}(V, \mathbb{R})$ of bilinear forms on V, that is functions $B: V \times V \to \mathbb{R}$ which are linear in each factor.

There is a similar way to describe the vector space of functions in which the higher derivatives $D^k f$ for $k \geq 2$ take values. The key point here is quite general:

Lemma 5.3. *Let X, Y and Z be sets, and write $F(X, Y)$ for the set of all functions from X to Y. Then there is a bijection $\theta: F(X, F(Y, Z)) \to F(X \times Y, Z)$ given by $\theta(f)(x, y) = f(x)(y)$, for all $x \in X, y \in Y$.*

Proof. This is trivial to check – the inverse map $\xi: F(X \times Y, Z) \to F(X, F(Y, Z)$ is given by $\xi(g)(x) = [y \mapsto g(x, y)]$, for all $x \in X, y \in Y$. □

Write $V^k = V \times \ldots \times V$ for the Cartesian product of V with itself k times, and let $\mathcal{M}^k(V, W)$ be the space of k-multilinear functions on V taking values in W:

$$\mathcal{M}^k(V, W) = \{f: V^k \to W : f(v_1, \ldots, v_k) \text{ is linear in each } v_i, 1 \leq i \leq k\}$$

Example 5.4. If $k = 1$ then $\mathcal{M}^1(V, W)$ is just the space of linear maps $\mathcal{L}(V, W)$. The space $\mathcal{M}^2(V, \mathbb{R})$ is just the space $\mathrm{Bil}(V, \mathbb{R})$ of bilinear forms on V. The determinant function, viewed as a function on the column vectors of an $n \times n$ matrix, is an element of $\mathcal{M}^n(\mathbb{R}^n, \mathbb{R})$.

Lemma 5.5. *Let V and W be finite dimensional normed vector spaces. For each $k \geq 1$ there is a natural isomorphism $\theta_k: \mathcal{L}(V, \mathcal{M}^{k-1}(V, W)) \to \mathcal{M}^k(V, W)$, and hence if $f: U \to W$ is a function on an open subset U of V which is k-times differentiable, we may view $D^k f$ as a function from U to $\mathcal{M}^k(V, W)$.*

Proof. Taking $X = V$, $Y = V^{k-1}$ and $Z = W$ in Lemma 5.3, you can check that the map θ in the proof of the Lemma restricts to give the required isomorphism θ_k. The final part of the Lemma then follows by induction on k. □

Thus we see that the higher derivatives $D^k f$ can be viewed as functions on U taking values in $\mathcal{M}^k(V, W)$, the space of k-multilinear functions on V taking values in W. Arguing essentially as we do in Example 2.37, it is possible to check that, if $\{w_1, \ldots, w_m\}$ is a basis of W, and we write $f = \sum_{i=1}^m f_i w_i$, so that the f_i are the components of f, and $\{e_1, \ldots, e_n\}$ is as before the basis of V, then

$$D^k f_i(e_{j_1}, \ldots, e_{j_k}) = \partial_\alpha f_i,$$

where $\alpha = (j_k, j_{k-1}, \ldots, j_1)$.

Proposition 5.6. *Let V, W be normed vector spaces, let U be an open subset of V, and let $f: U \to W$. Then $f \in C^k(V, W)$ if and only if the higher total derivative*

$$Df^k: U \to \mathcal{M}^k(V, W)$$

exists and is continuous. Moreover f is smooth if and only if all of the higher total derivatives Df^k exist.

5.3 *Symmetries of higher derivatives

The multivariable calculus result on the symmetry of the mixed partial derivatives is just the statement that the Hessian matrix of $D^2 f$ is symmetric which implies that $D^2 f_a$ is a symmetric bilinear form, thus the symmetry of mixed partial derivatives can be reinterpreted in a coordinate-free way, namely that $D^2 f_a(v_1, v_2) = D^2 f_a(v_2, v_1)$ for all $v_1, v_2 \in V$. An advantage of this formulation is that the famous "symmetry of mixed partial derivatives" obtains a natural invariant formulation, and moreover the symmetry holds as soon as the "total" second derivative exists, which is a weaker hypothesis than the classical one (which requires all second partial derivatives to exist and be continuous[19]).

We first need the following a simple Lemma. It is the analogue of the fact that, if $\alpha: V \to \mathbb{R}$ is a linear functional, and $\alpha = o(\|x\|)$ then $\alpha = 0$, as one readily sees by considering the operator norm of α.

[19] This is, unsurprisingly, reminiscent of the relationship between the total derivative and continuity of the partial derivatives.

Lemma 5.7. *Suppose that $\beta \colon V \times V \to \mathbb{R}$ is a bilinear map and suppose that $\beta(v,w) = o((\|v\| + \|w\|)^2)$. Then $\beta = 0$.*

Proof. Since β is bilinear, it suffices to show that $\beta(v_1, v_2) = 0$ for any $v_1, v_2 \in V$ with $\|v_1\| = \|v_2\| = 1$. Thus we fix unit vectors $v_1, v_2 \in V$. But now, for $s \in \mathbb{R}_{>0}$,

$$\frac{\beta(sv_1, sv_2)}{(\|sv_1\| + \|sv_2\|)^2} = \frac{s^2\beta(v_1, v_2)}{(2s)^2} = \frac{1}{4}\beta(v_1, v_2).$$

while $(\|sv_1\| + \|sv_2\|)^2 = 4s^2 \to 0$ as $s \to 0$. Thus if $\beta(v_1, v_2)$ is $o(\|v_1\|^2 + \|v_2\|^2)$ we must have $\beta(v_1, v_2) = 0$ as required. $\square$

The previous Lemma is the key to proving that $D^2 f_a$ is a symmetric bilinear form. (In examining the proof of the next result, it may be worth noting that the linear analogue of the previous Lemma is one way to see that the derivative Df_a is unique).

Proposition 5.8. *Let U be an open subset of a normed vector space V. If $f \colon U \to \mathbb{R}$ is twice differentiable at $a \in U$, then viewing $D^2 f$ as a bilinear form on V we have $D^2 f_a(v_1, v_2) = D^2 f_a(v_2, v_1)$.*

Proof. Note that, in order for $D^2 f$ to be defined, we must have f differentiable in a neighbourhood of a, and Df is continuous at a since it is differentiable at a.

Fix $r > 0$ such that $B = B(a,r) \subseteq U$ such that Df is defined for all $x \in B(a,r)$. Consider the function $A \colon B \times B \to \mathbb{R}$ given by

$$A(h,k) = f(a+h+k) - f(a+h) - f(a+k) + f(a).$$

Note that A has the virtue of being symmetric, that is $A(h,k) = A(k,h)$, but, unlike $D^2 f(h,k)$ it is not bilinear in h and k. The idea of the proof is to compare the two when $(h,k) \in V \oplus V$ is very small. Thus, fixing h for the moment, consider

$$J_1(k) = A(h,k) - D^2 f_a(h,k)$$

Now, noting $J_1(0) = 0$, and writing $i_h(D^2 f_a)$ for the linear functional $k \mapsto D^2 f_a(h,k)$, we can apply the Mean Value Inequality 2.24 to J to obtain

$$\|J_1(k)\| \leq \|k\|. \sup_{0 \leq t \leq 1} \|Df_{a+h+tk} - Df_{a+tk} - i_h(D^2 f_a)\|_\infty \tag{5.1}$$

Now as Df is differentiable at a, we may write

$$\begin{aligned} Df_{a+tk} &= Df_a + i_{tk}(D^2 f_a) + \|tk\|\epsilon_1(tk), \\ Df_{a+h+tk} &= Df_a + i_{h+tk}(D^2 f_a) + \|h+tk\|\epsilon_1(h+tk). \end{aligned}$$

where $\epsilon_1(x) \to 0$ as $x \to 0$. Hence we see that

$$Df_{a+h+tk} - Df_{a+tk} - i_h(D^2 f_a) = \|h+tk\|\epsilon_1(h+tk) - \|tk\|\epsilon_1(tk).$$

so that, in particular, if we let $\epsilon_2(h,k) = \sup\{\|\epsilon_1(s.h+t.k)\| : 0 \leq s,t \leq 1\}$, then $\epsilon_2(h,k) = \epsilon_2(k,h)$ and $\epsilon_2(h,k) \to 0$ as $(h,k) \to 0$ and

$$\|Df_{a+h+tk} - Df_{a+tk} - i_h(D^2 f_1)\| \leq (\|h\| + \|k\|).\epsilon_2(h,k)).$$

Thus returning to the inequality (5.1), we see that

$$\|J_1(k)\| = \|A(h,k) - D^2 f(h,k)\| \leq \|k\|.(\|h\| + \|k\|).\epsilon_2(h,k).$$

But carrying out the same analysis for $J_2(k) = A(k,h) - D^2 f(k,h)$ we see that $\|A(k,h) - D^2 f(k,h)\| \leq \|h\|(\|h\| + \|k\|).\epsilon_2(k,h)$, and hence if we let

$$\beta(h,k) = D^2 f_a(h,k) - D^2 f_a(k,h),$$

we see that β is a bilinear form which, by the symmetry of $A(h,k)$, satisfies:

$$\|\beta(h,k)\| \leq \|D^2 f_a(h,k) - A(h,k)\| + \|A(k,h) - D^2 f_a(k,h)\| \leq (\|h\| + \|k\|)^2 \epsilon_2(h,k). \tag{5.2}$$

But now Lemma 5.2 shows that $\beta = 0$ and hence $D^2 f_a$ is symmetric as required. $\square$

Remark 5.9. Using induction, it is straight-forward to use the previous Theorem to see that, whenever they exist, the higher derivatives $D^k f_a$ as symmetric k-multilinear forms.

5.4 *The immersion criterion for a submanifold

For completeness, we include here a proof of the equivalence of the definition of a submanifold given in Remark 4.5 with that given in Definition 4.2. In fact we prove something slightly stronger, giving a condition for the image of an injective immersion to yield a submanifold.

Proposition 5.10. *Let V be a n-dimensional normed vector space, and let 0_k denote the origin in $\mathbb{R}^k$. Suppose that $M \subseteq V$ is such that, for some $a \in M$, there exists*

- *an open neighbourhood U_a of a;*
- *an injective function $\psi \in C^1(B(0_k, R), V)$ whose derivative $D\psi_x$ is injective for every $x \in B(0_k, R)$;*
- *an $r \in (0, R)$ such that $\psi(B(0_k, r), 0_k) = (U \cap M, a)$.*

Then $M \cap U_a$ is a k-dimensional submanifold of V, and hence if the above conditions hold for all $a \in M$ then M is a submanifold of V.

Proof. If suffices to show that $\psi(B(0, r))$ is a k-submanifold of V. Suppose $p \in \psi(B(0_k, r))$. Then since ψ is injective, there is a unique $q \in B(0_k, r)$ such that $\psi(q) = p$. Let $V_1 = \mathrm{im}(D\psi_q)$, and pick a complementary subspace V_2 of V_1 in V, so that $V = V_1 \oplus V_2$. Let $i_2 \colon V_2 \to V$ denote the inclusion map. Let $\varphi \in C^1(B(0_k, r) \times V_2, V)$ be given by $\varphi(x, v) = \psi(x) + i_2(v)$. Since i_2 is a linear map, $D\varphi_{(q,0)} = D\phi_q + i_2$, and hence $D\varphi_{(q,0)}$ is an isomorphism. The inverse function theorem then shows that there is an open neighbourhood U_p of $\varphi(q, 0) = p$ and an open neighbourhood $\Omega_1 \times \Omega_2 \subseteq B(0, r) \times V_2$ of $(q, 0)$ such that φ restricts to a diffeomorphism from $(\Omega_1 \times \Omega_2, (q, 0))$ to (U_p, p). But now if $\theta \in C^1(U_p, \mathbb{R}^k \times V_2)$ is the inverse of $\varphi_{|\Omega_1 \times \Omega_2}$, and we write $\theta = \theta_1 \oplus \theta_2$ as the sum of its components in $\mathbb{R}^k$ and V_2 respectively, so that $\theta_2 \in C^1(U_p, V_2)$, it is easy to see $M \cap U_p = \theta_2^{-1}(0)$, and that that $D\theta_{2,p} = \pi_2$, where $\pi_2 \colon V \to V_2$ is the projection map with kernel V_1. It follows immediately that $D\theta_{2,p}$ has rank $\dim(V_2) = n - k$, and hence, since p was arbitrary, that $\psi(B(0_k, r))$ is a k-submanifold as required. □

5.5 *Normed vector spaces: duals and quotients

5.5.1 Bounded linear functionals

In Theorem 2.24, we assumed the differentiable function $f: U \to Y$ was a map between inner product spaces. In fact the proof only requires that Y is an inner products space: the goal of the theorem is to bound the length of a vector $y \in Y$ (where in the theorem $y = f(z_2) - f(z_1)$). The functional $\delta_y: Y \to \mathbb{R}$ given by $\delta_y(x) = \langle y, x\rangle$, *i.e.* taking the inner product with y, allows us to map our problem in Y to the real line in such a way that δ_y never increases the length of a vector (that is $|\delta_y(z)| \leq \|z\|$ is length preserving for vectors parallel to y, thus any bound we can calculate such as $\delta(v) \leq \delta(z)$ immediately implies that $\|v\| \leq \|z\|$.

Thus to use the same strategy of proof for an arbitrary normed vector space Y, one would need, for any vector $z \in Y$, a linear functional $\eta: Y \to \mathbb{R}$ with the property that $\|\eta\|_\infty = 1$ and $\eta(z) = \|z\|$. In fact, as we now show, one can prove that such functionals always exist for any normed vector space. Indeed if you have a functional $\eta: Z \to \mathbb{R}$ defined on a subspace Z of Y, then we say that a functional $\delta: Y \to \mathbb{R}$ is a norm-preserving extension of η if $\delta(z) = \eta(z)$ for all $z \in Z$ and $\|\delta\|_\infty = \|\eta\|_\infty$. If we take $Z = \mathbb{R}.z$ and η the linear functional defined by $\eta(z) = \|z\|$, then if δ is a norm preserving extension of η it has the properties we required above. The next Lemma shows that norm-preserving extensions always exist when Y is finite-dimensional[20]

Lemma 5.11. *Suppose that X is a finite-dimensional normed vector space and Z is a subspace of X. If $\eta_Z: Z \to \mathbb{R}$ is a linear functional on Z, then there is a functional $\delta: X \to \mathbb{R}$ which satisfies $\delta(z) = \eta(z)$ for all $z \in Z$. In other words η can be extended to a linear functional on X without increasing the operator norm.*

Proof. We use induction on $n = \dim(X)$. If $\dim(V) = 1$, then its only subspaces are $\{0\}$ and itself, and in each case the result is trivial. If $\dim(V) = n > 1$ and $Z \leq X$ is a subspace, then if $Z = X$ there is nothing to prove, while if $Z < X$, we may find a hyperplane H with $Z \leq H < V$, and by induction, there is a norm-preserving extension of δ to H, hence replacing Z with H if necessary, we may assume Z is codimension 1 in X.

Rescaling η if necessary, we may assume that $\|\eta\|_\infty = 1$. Pick $u \in X\backslash Z$, so that $X = \text{Span}\{Z, u\} = Z \oplus \mathbb{R}.u$. Any $\delta: X \to \mathbb{R}$ which restricts to η on Z is then determined by its value on u, say $\delta(u) = \lambda$, and the condition that $\|\delta\|_\infty = 1$ is

$$|\delta(z + t.u)| = |\eta(z) + t.\lambda| \leq \|z + t.u\|, \quad \forall t \in \mathbb{R}, z \in Z.$$

This is automatic if $t = 0$, while if $t \neq 0$, we may divide through by it to see that our condition is equivalent to $|\eta(z) + \lambda| \leq \|z + u\|$ for all $z \in Z$.

Rearranging, this becomes $\lambda \in I_z$ for every $z \in Z$, where $I_z = [-\|z + u\| - \eta(z), \|z + u\| - \eta(z)]$. Thus we need the intersection of the closed intervals I_z over all $z \in Z$ to be non-empty. But this follows precisely when, for any $z_1, z_2 \in Z$, the lower end-point of I_{z_1} is always at most the upper limit of I_{z_2}, that is, if and only if for all $z_1, z_2 \in Z$ we have

$$-\|z_1 + u\| - \eta(z_1) \leq \|z_2 + u\| - \eta(z_2)$$

But this is just $\delta(z_2 - z_1) \leq \|z_1 + u\| + \|z_2 + u\|$, and since η has norm 1 we have $|\eta(z_2 - z_1)| \leq \|z_2 - z_1\| \leq \|z_2 + u\| + \|z_1 + u\|$ as required. ◻

5.5.2 Quotients and normed vector spaces

If $(V, \|.\|)$ is a normed vector space, then any linear subspace F clearly inherits the structure of a normed vector space: the norm $\|.\|$ restricts to a norm on F. A somewhat more delicate question is whether the quotient vector space V/F inherits a norm. The first question is to decide what the notion of a norm on V/F should be? A natural suggestion is to consider how close the affine subspace $x + U$ comes to the origin in V. This leads to the definition of the function

$$x + F \mapsto \inf\{\|x + v\| : v \in F\}.$$

[20] The result (if you believe in the axiom of choice) holds for arbitrary normed vector spaces, and is called the *Hahn-Banach theorem*. It is important because it is a basic tool allowing one to build bounded linear functional having desirable properties.

Notice that while we might expect there to be a "closest point" on $x + F$ to the origin[21], it is not necessary to determine whether or not that is indeed the case in order to check this gives a norm on V/F, provided the subspace F is a closed subset of V.

Lemma 5.12. *Let X be a normed vector space and let F be a closed subspace, that is, a linear subspace which is also a closed subset of X. The the quotient vector space X/F inherits a norm:*

$$\|x + F\| := \inf\{\|x + u\| : u \in F\}.$$

Moreover, the quotient map $q\colon X \to X/F$ is bounded, with $\|q\|_\infty \leq 1$.

Proof. For any $x \in X$ we have $\|x + F\| = \inf_{u\in F} \|x - u\| = 0$ if and only if x is a limit point of F, thus since F is closed $\|x + F\| \geq 0$ for all x with equality if and only if $x + F = 0 + F$. Now suppose that $\lambda \in \mathbb{R}$. If $\lambda = 0$ then $\|\lambda.x + F\| = |\lambda|.\|x + F\| = 0$, while if $\lambda \neq 0$,

$$\|\lambda.x + F\| = \inf_{u\in F} \|\lambda.x + u\| = \inf_{u\in F} |\lambda|.\|x + \lambda^{-1}u\| = |\lambda| \inf_{u_1\in F} \|x + u_1\| = |\lambda|.\|x + F\|$$

For the triangle inequality, suppose $x + F, y + F \in V/F$. By the approximation property, for any $\epsilon > 0$, we may find $u_1, u_2 \in F$ such that $\|x + F\| \leq \|x + u_1\| < \|x + F\| + \epsilon$, and $\|y + F\| \leq \|y + u_2\| < \|y + F\| + \epsilon$. But then since $u_1 + u_2 \in F$, by definition we have

$$\begin{aligned}\|(x + y) + F\| &\leq \|(x + y) + (u_1 + u_2)\| = \|(x + u_1) + (y + u_2)\|\\ &\leq \|x + u_1\| + \|y + u_2\| < (\|x + F\| + \epsilon) + (\|y + F\| + \epsilon)\\ &= \|x + F\| + \|y + F\| + 2\epsilon,\end{aligned}$$

and since this holds for any $\epsilon > 0$, it follows that $\|(x + y) + F\| \leq \|x + F\| + \|y + F\|$, as required. Since $\|q(x)\| = \inf_{u\in F} \|x + u\| \leq \|x + 0\| = \|x\|$ we have $\|q\|_\infty \leq 1$, which completes the proof. □

The quotient construction for normed vector spaces in fact gives another approach to Theorem 1.17, as we now show: The key point is that, proving the statement by induction on dimension, it follows by the same argument used to prove Corollary 1.18 that subspaces of a finite-dimensional vector space are necessarily closed, hence any quotient is again a normed vector space.

Proposition 5.13. *Let V and W be normed vector spaces and suppose that* $\dim(V) < \infty$. *Then any linear map $\alpha\colon V \to W$ is automatically bounded, that is $\mathcal{B}(V, W) = \mathcal{L}(V, W)$.*

Proof. We use induction $\dim(V)$. In the case $\dim(V) = 1$, pick a vector $e \in V$ of norm 1. Then for any $v \in V$, we have $v = \pm\|v\|.e$ and hence $\|\alpha(v)\| = \|\alpha(e)\|.\|v\|$, so that $\|\alpha\|_\infty = \|\alpha(e)\|$, and α is bounded as required.

Next note that, for any given finite-dimensional vector space V, the statement of the proposition follows from the case $W = \mathbb{R}$, *i.e.* where $\alpha \in V^*$ is a linear functional. Indeed if $\dim(V) = n$ then $\dim(\alpha(V)) = m \leq n$, hence we can pick a basis $\{w_1, w_2, \ldots, w_m\}$ of $\alpha(V)$, and if, for $v \in V$ we define $\alpha_i\colon V \to \mathbb{R}$ by $\alpha(v) = \sum_{i=1}^m \alpha_i(v).w_i$, then the functions α_i are linear. and α is continuous if each α_i is. Indeed

$$\|\alpha(v)\| \leq \sum_{i=1}^{m} |\alpha_i(v)|.\|w_i\| \leq \left(\sum_{i=1}^{m} \|\alpha_i\|_\infty.\|w_i\|\right)\|v\|.$$

where the second inequality follows from the definition of the operator norm.

Now suppose that $n = \dim(V) > 1$, and that, by induction, we know any linear map whose domain is a normed vector space of dimension less than n must be bounded. Let $U < V$ be a subspace of V of dimension $k < n$. Picking a basis $\{u_1, \ldots, u_k\}$ of U defines a linear isomorphism $\phi\colon \mathbb{R}^k \to U$ where if $x = (x_1, \ldots, x_k) \in \mathbb{R}^k$ then $\phi(x) = \sum_{i=1}^k x_i u_i$. By our inductive hypothesis, ϕ is a topological isomorphism, and hence since $\mathbb{R}^k$ (viewed as a normed vector space using the $\|.\|_2$ norm) is complete, so is[22] U. It follows that U must therefore be closed in V.

[21]This is always true if F is finite-dimensional, but is in fact not necessarily the case when F is infinite-dimensional.

[22]Note that while completeness is not invariant under homeomorphism, continuous linear maps are Lipschitz continuous, and Lipschitz continuous functions preserve Cauchy sequences.

But now the fact that any linear functional $\alpha \in V^*$ is continuous follows from the 1-dimensional case together with Lemma 5.12: Indeed if α is zero, it is trivially continuous, and if $\alpha \neq 0$ then $H = \ker(\alpha)$ is $(n-1)$-dimensional subspace of V, and hence as noted above H is closed. But then by Lemma 5.12, the norm on V induces one on V/H and the quotient map $q\colon V \to V/H$ has operator norm $\|q\|_\infty \leq 1$. But the functional α can be written as the composition $\alpha = \bar{\alpha} \circ q$, where $\bar{\alpha}\colon V/H \to \mathbb{R}$ is the injective linear map induced by α on V/H. But since $\dim(V/H) = 1$ we know $\bar{\alpha}$ is bounded, and hence by the submultiplicativity of the operator norm, α is bounded as required. □

Remark 5.14. This proposition shows that the topology $\mathcal{T}$ induced by any norm on a finite dimensional vector space is independent of the choice of norm. In fact, with a bit more thought it follows that this topology is determined by the linear functionals on V: it is the topology generated by the condition that every linear functional on V is continuous.