

Introduction to Manifolds.*

Kevin McGerty[†]

May 20, 2021

Contents

| | | |
|----------|---|-----------|
| 1 | Course Outline | 2 |
| 2 | Notation | 3 |
| 3 | Linear maps and continuity | 5 |
| 3.1 | Normed vector spaces | 5 |
| 3.2 | Bounded linear maps | 7 |
| 3.3 | Finite dimensional normed vector spaces | 9 |
| 4 | The derivative in higher dimensions | 10 |
| 4.1 | The definition | 10 |
| 4.2 | Partial derivatives and the total derivative | 14 |
| 4.3 | The Chain Rule | 18 |
| 4.4 | Real-valued functions of many variables on an inner product space | 20 |
| 4.5 | Mean Value Theorems | 20 |
| 4.6 | *Higher order derivatives | 23 |
| 4.6.1 | Symmetry of mixed partial derivatives | 24 |
| 5 | The Inverse and Implicit Function Theorems | 27 |
| 5.1 | The Inverse Function Theorem | 27 |
| 5.2 | *Variants of the Inverse Function Theorem. | 29 |
| 5.3 | The Implicit Function Theorem | 30 |
| 6 | Submanifolds of a normed vector space | 37 |
| 6.1 | Definition and basic properties | 37 |
| 6.2 | Tangent spaces and normal vectors | 41 |
| 6.3 | *Abstract Manifolds | 45 |

*See the end of the Course Outline for a description of the edits in this version.

[†]Please contact me at mcgerty@maths.ox.ac.uk if you find any ambiguities or errors in these notes.

| | | |
|----------|--|-----------|
| 7 | Lagrange multipliers | 45 |
| 8 | *Appendix | 50 |
| 8.1 | Proof of the Inverse Function Theorem | 50 |
| 8.2 | An alternative definition of a submanifold | 52 |
| 8.3 | Multilinear maps and higher derivatives | 53 |
| 8.3.1 | Symmetries of higher derivatives | 54 |
| 8.4 | Finite dimensional normed vector spaces | 56 |
| 8.4.1 | Subspaces and quotient spaces | 56 |

Sections, proofs, or individual Remarks which are marked with a (*) are non-examinable.

1 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]

Updates in this version of the notes:

1. Some more details are given in §4.5 on mean-value theorems and inequalities.
2. Section 4.6 discusses in more detail the case of the second derivative of real-valued function on an open subset of an n -dimensional normed vector space, and proves a slightly stronger version of the result on the symmetry of mixed partial derivatives.
3. Section 6.1 has been slightly edited with an additional remark added.
4. Section 6.2 has been updated to match the presentation in the lecture videos.
5. The Appendix, in Section 8.2, now gives a proof of the equivalence of the definition of a submanifold of \mathbb{R}^n given in the (non-examinable) Remark 6.5, which in turn has been clarified. (The precise statement required is more subtle than one might guess reading the previous phrasing.)
6. The Appendix in Section 8.3 describes the higher derivatives as multi-linear maps, and proves they are symmetric multi-linear maps when they exist (refining the result on the symmetry of mixed partial derivatives).

2 Notation

- $B(v, r) = \{w \in V : \|v - w\| < r\}$ denotes the open ball of radius r centred at $v \in V$, where V is a normed vector space.
- $\bar{B}(v, r) = \{w \in V : \|v - w\| \leq r\}$ denotes the closed ball of radius r centred at $v \in V$, where V is a normed vector space. If $r > 0$ then the closed ball $\bar{B}(v, r)$ is the closure of $B(v, r)$ (this is not necessarily true in a general metric space.)
- $\mathcal{L}(V, W)$ denotes the space of linear maps from V to W , where V and W are vector spaces.
- $\mathcal{B}(V, W)$ denotes the space of bounded linear maps from V to W , where V and W are normed vector spaces.
- I_V denotes the identity element of $\mathcal{B}(V, V)$. If $V = \mathbb{R}^n$ then we write I_n instead of $I_{\mathbb{R}^n}$. We also write I_n for the $n \times n$ identity matrix.
- $\|T\|_\infty$ denotes the operator norm of $T \in \mathcal{B}(V, W)$ (see Example 3.7).
- 0_V denotes the zero vector (or origin) of a normed vector space V . (Where there is no possibility for confusion, we may simply write 0 .)
- 0_n denotes the zero vector (or origin) of \mathbb{R}^n . (Where there is no possibility for confusion, we may simply write 0 .)
- $\text{Mat}_{m,n}(\mathbb{R})$ denotes the space of $m \times n$ matrices (*i.e.* matrices with m rows and n columns) with entries in \mathbb{R} . When $n = m$ we sometimes write $\text{Mat}_n(\mathbb{R})$ instead of $\text{Mat}_{n,n}(\mathbb{R})$.
- $0_{n,m}$ denotes the origin in $\text{Mat}_{n,m}(\mathbb{R})$.
- I_n denotes the identity map from \mathbb{R}^n to itself, (and the identity matrix in $\text{Mat}_n(\mathbb{R})$).
- $O(\|x\|)$ denotes any function $f: U \rightarrow W$ where U is a neighbourhood of the origin 0_V in a normed vector space V , taking values in a normed vector space W , such that $\|(f(x)/\|x\|)\|$ is bounded as $\|x\| \rightarrow 0$.
- $o(\|x\|)$ denotes any function $f: U \rightarrow W$ where U is a neighbourhood of the origin 0_V in a normed vector space V , taking values in a normed vector space W , such that $\|(f(x)/\|x\|)\| \rightarrow 0$ as $\|x\| \rightarrow 0$.
- $C^0(V, W)$ denotes the space of continuous functions on an open subset U of a normed vector space V taking values in a normed vector space W . See Definition 4.6.
- $C^1(U, W)$ denotes the space of continuously differentiable functions on an open subset U of a normed vector space V taking values in a normed vector space W . See Definition 4.15.

- $C^k(U, W)$ denotes the space of k -times continuously differentiable functions on an open subset U of a normed vector space V taking values in a normed vector space W . See Definition [4.26](#).
- a *diffeomorphism* is a bijective continuously differentiable function whose inverse is also continuously differentiable. See Remark [5.4](#).
- $f: (U, a) \rightarrow (V, b)$ denotes a function $f: U \rightarrow V$ satisfying $f(a) = b$ (and hence $a \in U$ and $b \in V$). See the sentence before Definition [5.10](#).

3 Linear maps and continuity

3.1 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. Almost everything¹ in this section was already treated in the lectures and problem sets for the Metric Spaces part of A2 in Michaelmas.

Definition 3.1. A normed vector space $(V, \|\cdot\|)$ is a pair consisting of a real² vector space V and a function $\|\cdot\|: V \rightarrow \mathbb{R}$ which satisfies

1. $\|v\| \geq 0$ with equality if and only if $v = 0$. (*Positivity.*)
2. For $\lambda \in \mathbb{R}$ and $v \in V$ we have $\|\lambda.v\| = |\lambda|.\|v\|$. (*Homogeneity.*)
3. $\|v + w\| \leq \|v\| + \|w\|$ for all $v, w \in V$. (*Triangle inequality.*)

Note that (2) implies that $\|0\| = 0$ and thus by (3) we must have

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

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 $\|v_1 - v_2\|$.

We will normally write $\|\cdot\|$ 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 ambiguity³, such as when we consider more than one norm on the same vector space, we will decorate the norm with a subscript, e.g. $\|\cdot\|_V$ or $\|\cdot\|_1$.

Recall that if V is a normed vector space and $v \in V$ we say that a subset $U \subseteq V$ is a *neighbourhood* of v 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 $B(x, r_x) \subseteq U$.

Example 3.2. If V is one-dimensional, it is easy to understand all possible norms on V . Indeed if we pick $e_1 \in V \setminus \{0\}$, then for any $v \in V$ there is a unique $\lambda \in \mathbb{R}$ such that $v = \lambda.e_1$. Now if $f: V \rightarrow \mathbb{R}_{\geq 0}$ 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 V provided f is not identically zero. Since any norm on V necessarily satisfies the homogeneity condition, it

¹The operator norm $\|T\|_\infty$ in Definition 3.7 and Corollary 3.15 are the exceptions I believe.

²In fact one just needs a field with a sensible notion of “absolute value” – for example the complex numbers equipped with the modulus function.

³If you find an ambiguity I have missed, please let me know.

follows that any norm $\|\cdot\|$ on V has the form $\|v\| = c \cdot |\lambda|$ for $c > 0$ a positive real number (where, as above, $v = \lambda \cdot e_1$).

If $\dim(V) > 1$ – indeed even for $\dim(V) = 2$ – one cannot give such an explicit classification of all possible norms⁴, 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 3.3. Let $V = \mathbb{R}^n$. 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

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

These norms are all *equivalent* in the following sense:

Definition 3.4. If $\|\cdot\|_a$ and $\|\cdot\|_b$ are two norms on a vector space V , we say that they are *equivalent* if there exist constants $C_1, C_2 > 0$ such that $C_1\|v\|_a \leq \|v\|_b \leq C_2\|v\|_a$.

Note that if $\|\cdot\|_a$ and $\|\cdot\|_b$ are equivalent norms on a vector space V , then not only do they give the same notions of convergence and continuity, but also of completeness (*i.e.* V is complete as a metric space for the norm $\|\cdot\|_a$ if and only if it is complete as a metric space for the norm $\|\cdot\|_b$).

Example 3.5. Consider the norms $\|\cdot\|_1$ and $\|\cdot\|_2$ on \mathbb{R}^n defined above. We claim that they are equivalent. Indeed if $x = (x_1, \dots, x_n)$, then clearly

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

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

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

⁴Giving a norm $\|\cdot\|$ on \mathbb{R}^n is equivalent to giving the set $B_{\|\cdot\|} = \{v \in V : \|v\| \leq 1\}$ of vectors in its closed unit ball. Such a set $B_{\|\cdot\|}$ 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.

3.2 Bounded linear maps

Definition 3.6. If V and W are vector spaces, we write $\mathcal{L}(V, W)$ for the vector space of all linear maps from V to W .

Definition 3.7. If $(V, \|\cdot\|_V)$ and $(W, \|\cdot\|_W)$ are normed vector spaces, we say a linear map $T: V \rightarrow W$ is *bounded* if it takes bounded subsets⁵ of V to bounded subsets of W . In a normed vector space V , a subset $X \subseteq V$ is bounded if and only if there is an $R > 0$ such that $X \subseteq \bar{B}(0_V, R)$. Thus T will take bounded sets to bounded sets precisely if, for each $R > 0$ there is some $C_R > 0$ such that $T(\bar{B}(0_V, R)) \subseteq \bar{B}(0_W, C_R)$, i.e. $\|T(v)\| \leq C_R$ whenever $\|v\| \leq R$. But since $\|T(v/R)\| = R^{-1}\|T(v)\|$, it suffices to check this for $R = 1$.

Thus⁶ T is bounded if there is a constant $C > 0$ such that $\|T(v)\| \leq C$ for all $v \in \bar{B}(0, 1)$. If T is bounded we set

$$\|T\|_\infty = \sup\{\|T(v)\| : v \in V, \|v\| \leq 1\}.$$

Notice that the homogeneity of the norm shows that the condition that $\|T(v)\| \leq C$ for all $v \in \bar{B}(0, 1)$ is equivalent to the condition that $\|T(v)\| \leq C\|v\|$ for all $v \in V$, since if $v \neq 0$ then $v/\|v\| \in \bar{B}(0, 1)$.

We will write $\mathcal{B}(V, W)$ for the subspace of $\mathcal{L}(V, W)$ consisting of bounded linear maps from V to W . (Check you see this is indeed a linear subspace.) It is a normed vector space, with the norm, known as the *operator norm* given by $T \mapsto \|T\|_\infty$. Using standard facts about suprema, you can check that this norm is *submultiplicative*, in the sense that if U, V and W are normed vector spaces, $S: U \rightarrow V$ and, as above $T: V \rightarrow W$, then $\|T \circ S\|_\infty \leq \|T\|_\infty \|S\|_\infty$.

Remark 3.8. In Metric Spaces, you studied the space $B(X)$ of real-valued *bounded functions* on a 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 $\alpha: V \rightarrow W$ between normed vector spaces is never bounded, thus the usages are not, at first sight, consistent.

This apparent inconsistency is not, however, impossible to resolve⁷: Since it is compatible with scaling, a linear map α is completely determined by its values on $B_V = \bar{B}(0_V, 1)$, indeed if $v \neq 0$ then $u = v/\|v\| \in \bar{B}(0, 1)$ and $\alpha(v) = \|v\|\alpha(u)$. Thus we get an injective map $r: \mathcal{B}(V, W) \rightarrow C(B_V, W)$, from $\mathcal{B}(V, W)$ to the space of continuous functions on B_V taking values in W . Here $r(\alpha)$ is just the restriction of α to the closed ball B_V . By definition, it gives an isometric embedding of $\mathcal{B}(V, W)$, equipped with the operator norm, into $C_b(B_V, W)$, where the latter space is equipped with the usual supremum norm: $\|f\|_\infty = \sup\{\|f(x)\| : x \in B_V\}$.

Lemma 3.9. A linear map $T: V \rightarrow W$ between normed vector spaces is bounded if and only if it is continuous.

⁵A subset Y of a metric space (X, d) is bounded if $\sup\{d(x, y) : x, y \in Y\}$ is finite. Equivalently, Y is contained in some closed ball of X .

⁶This is the standard definition you will see in most textbooks.

⁷It, of course, is perfectly acceptable to just remember the apparent inconsistency in usage.

Proof. If T is bounded, so that we may find a $C > 0$ such that $\|T(v)\| \leq C\|v\|$, then

$$\|T(v_1) - T(v_2)\| = \|T(v_1 - v_2)\| \leq C\|v_1 - v_2\|,$$

so that T is in fact Lipschitz continuous. For the converse, if T is continuous, it is certainly continuous at 0, hence there is a $\delta > 0$ such that $\|T(u)\| < 1$ for all $u \in V$ with $\|u\| < \delta$. But if $v \in V \setminus \{0\}$ then $v_1 = \frac{\delta}{2\|v\|} \cdot v \in B(0, \delta)$ and hence $\|T(v_1)\| \leq 1$, so that $\|T(v)\| \leq 2\delta^{-1}\|v\|$. Since this inequality also clearly holds for $v = 0$ it follows T is bounded as required. \square

Definition 3.10. We say that $\alpha \in \mathcal{B}(V, W)$ is a *topological isomorphism* of normed vector spaces if it has a bounded linear inverse (The linearity of the inverse is automatic, but the boundedness is not.) More precisely, $\alpha \in \mathcal{B}(V, W)$ is a topological isomorphism if there is a $\beta \in \mathcal{B}(W, V)$ such that $\alpha \circ \beta = I_W$ and $\beta \circ \alpha = I_V$. Two normed vector spaces V and W are said to be *topologically isomorphic* if there is an isomorphism of normed vector spaces $\alpha: V \rightarrow W$ from V to W , that is, a linear isomorphism which is also a homeomorphism between V and W as metric spaces.

Remark 3.11. If V is a vector space with two norms $\|\cdot\|_a$ and $\|\cdot\|_b$, then $\|\cdot\|_a$ and $\|\cdot\|_b$ are equivalent if and only if the identity map is a topological isomorphism from $(V, \|\cdot\|_a)$ to $(V, \|\cdot\|_b)$.

***Remark 3.12.** Let $V = C([0, 1])$ be the space of continuous functions on the interval $[0, 1]$ and let $W = C_0^1([0, 1])$ 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 V and W as normed vector spaces using the supremum norm. Then we have a linear map $T: V \rightarrow W$, where if $f \in C([0, 1])$,

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

The fundamental theorem of calculus shows that $T(f)$ is indeed in $C_0^1([0, 1])$ 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 $T \in \mathcal{B}(V, W)$. While T is invertible with inverse $D: W \rightarrow V$, where $D(g) = g'$ for all $g \in W$, 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.

Example 3.13. For the vector space $\mathcal{L}(\mathbb{R}^n, \mathbb{R}^m)$ it can be useful to have a more explicit norm than the operator norm. Using the standard basis, we may identify $\mathcal{L}(\mathbb{R}^n, \mathbb{R}^m)$ with $\text{Mat}_{m,n}(\mathbb{R})$, and hence with \mathbb{R}^{mn} equipped with the $\|\cdot\|_2$ -norm. Then if $v \in \mathbb{R}^n$, and $A = (a_{ij})$,

we have (using the $\|\cdot\|_2$ also for \mathbb{R}^n and \mathbb{R}^m)

$$\begin{aligned}\|A(v)\|_2^2 &= \sum_{i=1}^m \left(\sum_{j=1}^n a_{ij} v_j \right)^2 \\ &\leq \sum_{i=1}^m \left(\sum_{j=1}^n a_{ij}^2 \right) \left(\sum_{j=1}^n v_j^2 \right) \\ &= \left(\sum_{i,j} a_{ij}^2 \right) \left(\sum_{j=1}^n v_j^2 \right) \\ &= \|A\|_2^2 \cdot \|v\|_2^2.\end{aligned}$$

where in the second line we use the Cauchy-Schwarz inequality. Thus we see that the norm $\|\cdot\|_2$ on $\mathcal{L}(\mathbb{R}^n, \mathbb{R}^m)$ gives an upper bound for the operator norm. One can also check, again using Cauchy-Schwarz for example, that the norm $\|\cdot\|_2$ is submultiplicative. This norm on the space of linear maps between inner product spaces is sometimes known as the *Hilbert-Schmidt* norm. The associated inner product is $\langle A, B \rangle = \text{tr}(A \cdot B^t)$.

3.3 Finite dimensional normed vector spaces

Lemma 3.14. *Any two norms on a finite-dimensional vector space are equivalent.*

Proof. If $(V, \|\cdot\|_V)$ is any finite-dimensional normed vector space, then if $B = \{v_1, \dots, v_n\}$ is a basis of V , it gives a linear isomorphism $\alpha: \mathbb{R}^n \rightarrow V$, where if $x = (\lambda_1, \dots, \lambda_n)$ we define $\alpha(x) = \sum_{i=1}^n \lambda_i v_i$. If we set $\|x\| = \|\alpha(x)\|_V$ then α is a linear isometry from $(\mathbb{R}^n, \|\cdot\|)$ to $(V, \|\cdot\|_V)$. Thus we may assume that $V = \mathbb{R}^n$. Moreover since equivalence of norms is an equivalence relation, it suffices to show that if $\|\cdot\|$ is any norm on \mathbb{R}^n , then $\|\cdot\|$ is equivalent to $\|\cdot\|_1$, where as usual if $x = (\lambda_1, \dots, \lambda_n)$ then $\|x\|_1 = \sum_{i=1}^n |\lambda_i|$.

Now let $\{e_1, \dots, e_n\}$ be the standard basis of \mathbb{R}^n , and set $M_1 = \max\{\|e_i\| : 1 \leq i \leq n\}$. Then if $x \in \mathbb{R}^n$ and we write $x = \sum_{i=1}^n \lambda_i e_i$, we have

$$\|x\| = \left\| \sum_{i=1}^n \lambda_i e_i \right\| \leq \sum_{i=1}^n |\lambda_i| \cdot \|e_i\| \leq M_1 \cdot \|x\|_1$$

Hence to show that $\|\cdot\|_1$ and $\|\cdot\|$ are equivalent, it remains to show that there is some $M_2 > 0$ such that $M_2 \|x\|_1 \leq \|x\|$ for all $x \in \mathbb{R}^n$. For this, first note that $x \mapsto \|x\|$ is Lipschitz continuous on $(\mathbb{R}^n, \|\cdot\|_1)$: indeed by the reverse triangle inequality

$$\left| \|x\| - \|y\| \right| \leq \|x - y\| \leq M \|x - y\|_1.$$

Now for any $x \neq 0$, if $M_2 \|x\|_1 \leq \|x\|$ then $M_2 \leq \|(x/\|x\|_1)\|$, and clearly $\|(x/\|x\|_1)\|_1 = 1$. Thus to show that M_2 exists it suffices to show that $\|\cdot\|$ is bounded away from 0 on $S = \{x \in \mathbb{R}^n : \|x\|_1 = 1\}$. But S is a closed bounded subset of $(\mathbb{R}^n, \|\cdot\|_1)$, and so is compact. Since $\|\cdot\|$ is continuous, it follows that $\|\cdot\|$ attains its minimum value on S . Thus we may pick $v \in S$ such that $\|v\| \leq \|x\|$ for all $x \in S$. But since $\|\cdot\|_1$ is a norm, $0 \notin S$, and hence, as $\|\cdot\|$ is a norm, $\|v\| > 0$. It follows we may take $M_2 = \|v\|$. \square

Corollary 3.15. *Suppose that V and W are normed vector spaces and that V is finite-dimensional. Then $\mathcal{B}(V, W) = \mathcal{L}(V, W)$, that is, every linear map $\alpha: V \rightarrow W$ is bounded.*

Proof. By the proof of the previous Lemma, we see that we can assume $V = (\mathbb{R}^n, \|\cdot\|_1)$. Then if $\{e_1, \dots, e_n\}$ denotes the standard basis of \mathbb{R}^n , and we set $M = \max\{\|\alpha(e_i)\| : 1 \leq i \leq n\}$, then if $x = \sum_{i=1}^n x_i e_i \in \mathbb{R}^n$, we have

$$\|\alpha(x)\| = \left\| \sum_{i=1}^n x_i \alpha(e_i) \right\| \leq \sum_{i=1}^n |x_i| \cdot \|\alpha(e_i)\| \leq M \|x\|_1,$$

and hence $\|\alpha\|_\infty \leq M$, that is, α is bounded as required. \square

Corollary 3.16. *Let V be a normed vector space and let U be a finite dimensional subspace. Then U is a closed subset of V .*

Proof. If $\dim(U) = k$, then Corollary 3.15 show that U is topologically isomorphic to $(\mathbb{R}^k, \|\cdot\|_1)$. But by the Bolzano-Weierstrass theorem $(\mathbb{R}^k, \|\cdot\|_1)$ is complete, hence so is U . As a complete subspace of a metric space it must be closed (see the proof of Lemma 6.2.1 in [B] – 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 ambient space is complete). \square

Remark 3.17. 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 $\|\cdot\|_2$ norm associated to the standard dot product (and thus the spaces of linear maps between them can be viewed as matrices equipped with either the operator norm or the Hilbert-Schmidt 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 $\|\cdot\|_1$ norm is the simplest to consider) and that, even if we state results for $(\mathbb{R}^n, \|\cdot\|_2)$, 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.

4 The derivative in higher dimensions

4.1 The definition

We now consider how the notion of differentiability can be extended to \mathbb{R}^m -valued functions on open subsets of \mathbb{R}^n . When $n = 1$ this is straight-forward: the ordinary definition still makes sense, since if $U \subset \mathbb{R}$ is open and $f: U \rightarrow \mathbb{R}^m$, we can define

$$Df(x) = f'(x) = \lim_{t \rightarrow x} \frac{f(t) - f(x)}{t - x}, \quad (4.1)$$

when, of course, this limit exists. Note that if we write f in terms of its components, $f = (f_1, \dots, f_m)$, then $f'(x) = (f'_1(x), \dots, f'_m(x))$, i.e. by taking components we can reduce to the one-variable case.⁸

To extend the notion of differentiability to the case where $n > 1$, it is useful to review the heuristics which motivate the definition in the one-variable case: There are a number of 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 $(a, f(a))$. Since our conception of time is resolutely one-dimensional, it is this latter geometric interpretation of the derivative which leads more readily to a notion of differentiability in many dimension. To see how this works, consider the following rephrasing of the standard definition of the derivative:

Lemma 4.1. *Let $U \subseteq \mathbb{R}$ be an open set and suppose that $f: U \rightarrow \mathbb{R}$ is a function. If $a \in U$, so that for some $r > 0$ we have $(a - r, a + r) \subseteq U$, then f is differentiable at a (in the sense of (4.1)) if and only if there is a real number α such that*

$$f(x) = f(a) + \alpha(x - a) + |x - a|\epsilon(x), \quad (4.2)$$

where $\epsilon(x) \rightarrow 0 = \epsilon(a)$ as $x \rightarrow a$. If it exists, α is unique, and $\alpha = f'(a) = Df(a)$.

Proof. If $x \neq a$, then $|x - a| \neq 0$ and hence, for such x , (4.2) determines $\epsilon(x)$ uniquely. Moreover, rearranging, we find that, for such x , we have

$$|\epsilon(x)| = \left| \frac{f(x) - f(a)}{x - a} - \alpha \right|,$$

thus the condition that $\epsilon(x) \rightarrow 0$ as $x \rightarrow a$ is equivalent to $\lim_{x \rightarrow a} \frac{f(x) - f(a)}{x - a} = \alpha$, proving the Lemma. \square

Remark 4.2.

1. Notice that, phrased this way, the *definition* of the derivative is to require Taylor's theorem to hold to first order: the statement of Lemma 4.1 is a rigorous formulation of assertion that the function $L_a(x) = f(a) + f'(a)(x - a)$ is, by definition, the “best linear approximation” to $f(x)$ at the point a , or, in more geometric terms, its graph gives the tangent line to the graph of f at $(a, f(a))$. If $L_1(x) = f(a) + \alpha_1(x - a)$ and $L_2(x) = f(a) + \alpha_2(x - a)$ are two linear functions⁹ passing through the point $(a, f(a))$, then $|L_1(x) - L_2(x)| = |\alpha_1 - \alpha_2| \cdot |x - a|$. Thus if $|f(x) - L_1(x)| = |\epsilon(x)| \cdot |x - a|$, where $|\epsilon(x)| \rightarrow 0$ as $x \rightarrow a$ then $L_1(x)$ approximates $f(x)$ near a better than any other linear function: if $L_2 \neq L_1$ then $C = |\alpha_1 - \alpha_2| > 0$ and $|f(x) - L_2(x)|$, for x close to a , will be approximately $C \cdot |x - a|$.

⁸On the other hand, the expression in (4.1) makes sense for any function $f: \mathbb{R} \rightarrow V$ taking values in a normed vector space V , whether or not it is finite dimensional – ultimately at least, we should seek a notion of differentiability which is coordinate-free.

⁹“Affine-linear” might be a more suitable term – their graphs are lines not necessarily passing through the origin.

2. This formulation of the derivative is also the easiest to use when proving the chain rule. We will see this later when discussing the chain rule in the multivariable case, but a quick check of the standard one-variable proofs of the chain rule should confirm that they all (more or less explicitly depending on the author) harnesses the formulation of Definition 4.1.
3. Finally, and for us, most importantly, this definition has the advantage that it immediately generalizes to many variables: The standard one-variable definition of the derivative of a function f at a point a considers the ratio of the difference $f(x) - f(a)$ with $x - a$. Apart from the case of \mathbb{R}^2 viewed as \mathbb{C} , however, we cannot consider the ratio of two vectors. On the other hand, as soon as we equip a vector space with a norm, we can compare different approximations to a given function.

We can now give the definition of the derivative in higher dimensions: Motivated by Lemma 4.1, we require that, if $f: U \rightarrow \mathbb{R}^m$ is a function defined on an open subset U of \mathbb{R}^n , then for f to be differentiable at $x = a$, it should have a “best first-order approximation” near a , that is, there should be a linear map $T: \mathbb{R}^n \rightarrow \mathbb{R}^m$ such that $f(a) + T(x - a)$ estimates $f(x)$ better than any other such affine-linear function. Formally, the definition is as follows:

Definition 4.3. Suppose that m and n are positive integers. If $U \subseteq \mathbb{R}^n$ is an open set, and $f: U \rightarrow \mathbb{R}^m$. Then we say that f is differentiable at $a \in U$ if there is a linear map $T: \mathbb{R}^n \rightarrow \mathbb{R}^m$ such that

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

where $\epsilon(x) \rightarrow \epsilon(a) = 0$ as $x \rightarrow a$. If such a map T exists it is unique, and we denote it as $Df(a)$ or Df_a : since $Df(a)$ is a linear map, we will often apply it to a vector $v \in \mathbb{R}^n$, and the notation $Df(a)(v)$ is less compact than $Df_a(v)$. It is known as the *total derivative*¹⁰ of f at a . We say that f is differentiable on U if it is differentiable at every $a \in U$. In that case, we obtain a function $Df: U \rightarrow \mathcal{L}(\mathbb{R}^n, \mathbb{R}^m)$.

Remark 4.4.

- 1 The definition of the derivative is sometimes written using the “little o ” notation in terms of the vector $h = x - a$, that is, as $f(a + h) = f(a) + T(h) + o(\|h\|)$.
- 2 One can prove the uniqueness of the linear map Df_a directly, and the problem set asks you to do this. The uniqueness of Df_a can also be deduced, however, as we will shortly see, by understanding its relationship to the notion of partial derivatives which you have already met in multivariable calculus.
- 3 If U is an open subset of \mathbb{R}^n and $f: U \rightarrow \mathbb{R}^m$, then if $f = (f_1, \dots, f_m)$, you can check that f is differentiable at $a \in U$ if and only if each f_i is, and $Df_a = \sum_{i=1}^m Df_{i,a} \cdot e_i$, that is, if $v \in \mathbb{R}^n$, we have $Df_a(v) = \sum_{i=1}^m Df_{i,a}(v) \cdot e_i$. This can be checked directly, and is in essence a very special case of the multi-variable version of the *Chain Rule*.

¹⁰As opposed to the partial derivatives.

- 4 The notion of differentiability is independent of the norms used on \mathbb{R}^n and \mathbb{R}^k , since, by Lemma 3.14, all norms are equivalent, and in particular give the same notion of convergence. In fact the definition of differentiability makes sense for any function $f: V \rightarrow W$ provided V and W are normed vector spaces.

[*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 3.15, is automatic in the finite-dimensional setting.]

- 5 If $f: U \rightarrow \mathbb{R}^m$ is differentiable on U , then it defines a function $Df: U \rightarrow \mathcal{L}(\mathbb{R}^n, \mathbb{R}^m)$. Viewed as a function “taking values in (linear) functions” it appears to be a more complicated object than the original function f . However, $\mathcal{L}(\mathbb{R}^n, \mathbb{R}^m)$ is just a $n \cdot m$ -dimensional normed vector space, and hence by picking a basis for it, we can view Df as a map from U to $\mathbb{R}^{n \cdot m}$. Thus, at least in principle, Df is no more complicated an object than f . We discuss this in more detail in Section 4.6.

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

Lemma 4.5. *Let U be an open subset of \mathbb{R}^n and $f: U \rightarrow \mathbb{R}^m$ be a function. Suppose that f is differentiable at $a \in U$. Then there are constants $C, r > 0$ so that, for all $x \in B(a, r)$,*

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

In particular, f is continuous at a .

Proof. Since f is differentiable at a , there is a function $\epsilon: U \rightarrow \mathbb{R}^m$ such that

$$f(x) = f(a) + Df_a(x - a) + \|x - a\|\epsilon(x), \quad \forall x \in U, \quad (4.3)$$

where $\epsilon(x) \rightarrow 0 = \epsilon(a)$ as $x \rightarrow a$. In particular, we have

$$\begin{aligned} \|f(x) - f(a)\| &= \|Df_a(x - a) + \|x - a\|\epsilon(x)\| \\ &\leq (\|Df_a\|_\infty + \|\epsilon(x)\|) \cdot \|x - a\|, \end{aligned} \quad (4.4)$$

where we write $\|Df_a\|_\infty$ for the operator norm of Df_a as in Example 3.7. Since $\epsilon(x) \rightarrow 0$ as $x \rightarrow a$ there is some $s > 0$ such that for x with $\|x - a\| < s$ we have $\|\epsilon(x)\| < 1$. It follows that we may take $r = s$ and $C = \|Df_a\| + 1$. \square

Definition 4.6. If V and W are normed vector spaces and U is an open subset of V , then we write $C^0(U, W)$ for the vector space of continuous functions on U taking values in W . The previous Lemma thus shows that if $f: U \rightarrow W$ is differentiable on all of U then $f \in C^0(U, W)$.

Example 4.7. Constant functions $c: \mathbb{R}^n \rightarrow \mathbb{R}^k$ are clearly differentiable, with derivative 0, since if c is constant $c(x) = c(a)$. If $T: \mathbb{R}^n \rightarrow \mathbb{R}^k$ is linear, that is $T \in \mathcal{L}(\mathbb{R}^n, \mathbb{R}^k)$, then, for any $a \in \mathbb{R}^n$ we have $Df_a = T$, since

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

(and thus the error term $\epsilon(x) \cdot \|x\|$ is identically zero). Thus if $f = T$ is linear, $Df: \mathbb{R}^n \rightarrow \mathcal{L}(\mathbb{R}^n, \mathbb{R}^k)$ is the constant function $x \mapsto T$.

If U is an open subset of \mathbb{R}^n and $f, g: U \rightarrow \mathbb{R}^k$ 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 $T \in \mathcal{L}(\mathbb{R}^n, \mathbb{R}^k)$ and $b \in \mathbb{R}^k$, then f is differentiable with $Df_a = T$ for all $a \in \mathbb{R}^n$.

Example 4.8. If $\|\cdot\|$ is a norm on \mathbb{R}^n , we may view it as a function $\|\cdot\|: \mathbb{R}^n \rightarrow \mathbb{R}$. 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 $\epsilon(h) \rightarrow 0$ as $\|h\| \rightarrow 0$ 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 $\|\cdot\|_1$, $\|\cdot\|_2$ and $\|\cdot\|_\infty$.

4.2 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:

Definition 4.9. If U is an open subset of \mathbb{R}^n and $f: U \rightarrow \mathbb{R}^m$ then, for any $v \in \mathbb{R}^n$ consider the function $f_{a,v}(t) = f(a + tv)$, which, since U is open, is defined for all sufficiently small real numbers t . Define the *directional derivative* of f in the direction v at $a \in U$ is defined to be

$$\partial_v f(a) = \frac{d}{dt}(f_{a,v}(t))_{t=0} = \lim_{t \rightarrow 0} \frac{f(a + tv) - f(a)}{t},$$

if this limit exists. Note that if $v_1 = s \cdot v$ for some $s \in \mathbb{R}_{>0}$, then letting $t_1 = s \cdot t$, we have

$$\partial_{v_1} f(a) = \partial_{s \cdot v} f(a) = \lim_{t \rightarrow 0} \frac{f(a + t \cdot s \cdot v) - f(a)}{t} = s \cdot \lim_{t_1 \rightarrow 0} \frac{f(a + t_1 \cdot v) - f(a)}{t_1} = s \cdot \partial_v f(a).$$

Thus if we scale the direction vector v , the directional derivative scales correspondingly. It follows that directional derivatives are completely determined by direction vectors v with unit length.

For each $j \in \{1, 2, \dots, n\}$ we define the *j-th partial derivative* of f , which we denote by $\partial_j f(a)$ or $\partial f / \partial x_j(a)$, to be the directional derivative of f for the direction vector $v = e_j$, so that

$$\partial_j f(a) = \frac{\partial f}{\partial x_j}(a) := \partial_{e_j}(f)(a) = \lim_{t \rightarrow 0} \frac{f(a_1, \dots, a_j + t, \dots, a_n) - f(a_1, \dots, a_n)}{t},$$

(again, when this limit exists).

Now suppose that f is differentiable at a with total derivative $T = Df_a$. It follows from the definitions that

$$\frac{f(a + t.v) - f(a)}{t} = T(v) \pm \epsilon(t.v)\|v\| \rightarrow T(v), \quad \text{as } t \rightarrow 0.$$

Thus the directional derivatives of f at a all exist and are equal to $T(v)$. In particular, if $x = (x_1, \dots, x_n)$ and we write $f(x) = (f_1(x), \dots, f_m(x))$ (so that the f_i are the components of f) then

$$T(e_j) = \partial_{e_j} f(a) = \partial_j f(a) = \begin{pmatrix} \partial_j f_1(a) \\ \partial_j f_2(a) \\ \vdots \\ \partial_j f_m(a) \end{pmatrix}$$

Thus we see that, if $T = Df_a$ exists, then its matrix with respect to the standard bases of \mathbb{R}^n and \mathbb{R}^m has columns given by $\partial_j f(a)$, and hence is

$$\left(\frac{\partial f_i}{\partial x_j} \right)_{x=a} = \begin{pmatrix} Df_1(a) \\ Df_2(a) \\ \vdots \\ Df_m(a) \end{pmatrix}.$$

Definition 4.10. As in multi-variable calculus, the above matrix is called the *Jacobian matrix* of the partial derivatives of f at a . It follows therefore immediately from the fact that the partial derivatives are unique (if they exist) that the total derivative, if it exists, is also unique. For later use, we note that the determinant $\det(Df_a) = \det(\partial_j f_i(a))$, is called the *Jacobian determinant*. It is often denoted $J_f(a)$.

The condition that the total derivative actually exists, however, is not equivalent to the existence of all the partial derivatives: The following example shows that a function need not be continuous at a point where all of its partial derivatives exist, whereas, as we have seen, a function is automatically continuous at a point where the total derivative exists.

Example 4.11. Let Ω be the open subset $\{(x_1, x_2) \in \mathbb{R}^2 : x_1^4 < x_2 < x_1^2\}$ and let $\chi = \mathbb{1}_\Omega$ be the indicator function of Ω , so that $\chi(x_1, x_2) = 1$ if $(x_1, x_2) \in \Omega$ and $\chi(x_1, x_2) = 0$ otherwise. The problem sheet asks you check that if $v \in \mathbb{R}^2 \setminus \{0\}$, the directional derivative $\partial_v \chi(0)$ exists and is equal to 0. Thus all the directional derivatives of χ exist at the origin, but χ fails to be continuous at the origin.

If the previous function χ feels artificial, it is worth noting that the function $f: \mathbb{R}^2 \rightarrow \mathbb{R}$ in Figure 1 given by

$$f(x_1, x_2) = \begin{cases} (x_1 x_2^2)/(x_1^2 + x_2^4), & (x_1, x_2) \neq (0, 0), \\ 0, & (x_1, x_2) = (0, 0) \end{cases}$$

is such that all of the directional derivatives $\partial_v f(0)$ exists for all non-zero $v \in \mathbb{R}^2$, but the total derivative Df_0 does not exist.




Figure 1: Graph of $f(x, y) = xy^2/(x^2 + y^4)$. All its directional derivatives exist at 0_2 but it is not differentiable there.

In two variables, complex analysis gives us many examples of functions $f: U \rightarrow \mathbb{R}^2$ on open subsets U of \mathbb{R}^2 whose total derivative exists on all of U .

Example 4.12. If U is an open subset of \mathbb{C} and $f: U \rightarrow \mathbb{C}$ 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

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

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.

Example 4.11 shows that the existence of all the partial derivatives for the function $\chi: \mathbb{R}^2 \rightarrow \mathbb{R}$ at the origin 0 is not sufficient to ensure that χ is continuous at that point. Since Lemma 4.5 shows that the existence of the total derivative at a point implies continuity at that point, χ cannot be differentiable at the origin. The function $f: \mathbb{R}^2 \rightarrow \mathbb{R}$ 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).

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.

Theorem 4.13. Suppose $U \subseteq \mathbb{R}^n$ is an open subset and that $f: U \rightarrow \mathbb{R}^m$ is a function whose partial derivatives exist near $a \in U$ and are continuous at a . Then the total derivative Df_a exists. With respect to the standard bases, the matrix of Df_a is therefore given by the Jacobian matrix of partial derivative.

Proof. Working component by component, it is enough to prove the result for f taking values in \mathbb{R} . Thus we suppose that $f: U \rightarrow \mathbb{R}$ and that there is an $r > 0$ with $B(a, r) \subseteq U$, such that for each $i \in \{1, \dots, n\}$ and $x \in B(a, r)$, the partial derivatives $\partial_i f(x)$ of f exist, and moreover each $\partial_i f(x)$ is continuous at $x = a$. Now if it exists, we know that the matrix of the total derivative Df_a with respect to the standard basis $\{e_1, \dots, e_n\}$ must be the row vector $(\partial_1 f(a), \dots, \partial_n f(a))$, viewed as an element of the dual space $(\mathbb{R}^n)^*$.

We write $a = (a_1, \dots, a_n)$. Let $\epsilon: B(0, r) \rightarrow \mathbb{R}^n$ be given by $\epsilon(0) = 0$ and, for $h = (h_1, \dots, h_n) \neq 0$, by

$$\epsilon(h) = \|h\|^{-1}(f(a+h) - f(a) - \partial_1 f(a).h_1 - \dots - \partial_n f(a).h_n). \quad (4.5)$$

Thus by definition, f is differentiable at a if we can show that $\epsilon(h) \rightarrow 0$ as $h \rightarrow 0$.

To show this, let $a^0(h) = a$ and, for $1 \leq k \leq n$, let $a^k(h) = a^{k-1}(h) + h_k e_k$. Then we have a telescoping sum

$$f(a+h) - f(a) = \sum_{j=1}^n (f(a^j(h)) - f(a^{j-1}(h))). \quad (4.6)$$

For $k = 1$, using the formulation of the derivative of Lemma 4.1, we see that $f(a^1) - f(a) = \partial_1 f(a).h_1 + \epsilon_1(h_1)$, where $\epsilon_1(h_1) \rightarrow 0 = \epsilon_1(0)$ as $h_1 \rightarrow 0$. Now suppose that $k \geq 2$. For $t \in \mathbb{R}$ sufficiently small, we can set $g_k(t) = f(a^{k-1}(h) + t e_k)$. Then g_k is differentiable with derivative $\partial_k f(a^{k-1}(h) + t e_k)$ and so the by the single-variable mean-value theorem, there is some $\theta_k \in (0, 1)$ such that

$$f(a^k(h)) - f(a^{k-1}(h)) = g_k(h_k) - g_k(0) = g'_k(\theta_k.h_k).h_k = \partial_k f(a^{k-1}(h) + \theta_k.h_k).h_k.$$

and hence (4.6) becomes

$$f(a+h) - f(a) = \partial_1 f(a).h_1 + \epsilon_1(h_1).h_1 + \sum_{j=2}^n \partial_k f(a^{k-1}(h) + \theta_k.h_k).h_k$$

and hence

$$\epsilon(h) = \epsilon_1(h_1) \cdot \frac{|h_1|}{\|h\|} + \sum_{j=2}^n \left(\partial_k f(a^{k-1}(h) + \theta_k.h_k) - \partial_k f(a) \right) \cdot \frac{h_k}{\|h\|}$$

But now if $\{\delta_k : 1 \leq k \leq n\}$ denotes the basis of $(\mathbb{R}^n)^*$ dual to the standard basis, $h_k = \delta_k(h)$, so that $|h_k| = |\delta_k(h)| \leq \|\delta_k\|_\infty \cdot \|h\|$, and hence, for any $k \in \{1, \dots, n\}$ the ratio $|h_k|/\|h\|$ is bounded as $h \rightarrow 0$. Similarly

$$\|a^{k-1}(h) + \theta_k h_k e_k - a\| \leq \left(\sum_{i=1}^k |h_i| \|e_i\| \right) \leq \left(\sum_{i=1}^k \|\delta_i\|_\infty \|e_i\| \right) \|h\|,$$

hence $a^{k-1}(h) + \theta_k h_k e_k \rightarrow a$ as $h \rightarrow 0$. It follows by the continuity of the partial derivatives $\partial_k f$ at a for each k (and the fact that $\epsilon_1(h_1) \rightarrow 0$ as $h \rightarrow 0$) that $\epsilon(h) \rightarrow 0$ as $h \rightarrow 0$ as required. \square

Remark 4.14. Note that in fact the proof didn't use 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 4.15. If U is an open subset of \mathbb{R}^n and $f: U \rightarrow \mathbb{R}^m$, then we say that f is *continuously differentiable* if¹¹ $Df: U \rightarrow \mathcal{L}(\mathbb{R}^n, \mathbb{R}^m)$ 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 $1 \leq j \leq n, 1 \leq i \leq m$. Let $C^1(U, \mathbb{R}^m)$ for the vector space of continuously differentiable functions on U taking values in \mathbb{R}^m .

***Remark 4.16.** If $f: U \rightarrow \mathbb{R}^k$ and $a \in U$, we say that f is *strongly differentiable* at a if there is a linear map $T \in \mathcal{L}(\mathbb{R}^n, \mathbb{R}^k)$ and an $r > 0$ such that

$$f(x) - f(y) = T(x - y) + o(\|x - y\|), \quad \forall x, y \in B(a, r).$$

If all of the partial derivatives of $f: U \rightarrow \mathbb{R}^k$ exist in a neighbourhood of $a \in U$ and are continuous at a , the technique of Theorem 4.13 shows that f is strongly differentiable at a : Let $h = y - x = \sum_{i=1}^n h_i e_i$, and in place of the a^k consider $x^k = x^{k-1} + h_k e_k$, where $x^0 = x$. Obviously, taking $y = a$ shows that, if it exists, the linear map T must be Df_a , but in general, a function which is differentiable at a point need not be strongly differentiable at that point – see Remark 5.4.

4.3 The Chain Rule

One of the fundamental properties of the differentiability 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 4.17. Let $f: U \rightarrow \mathbb{R}^m$ and $g: V \rightarrow \mathbb{R}^p$ be functions on an open subset $U \subseteq \mathbb{R}^n$ and an open subset $V \subseteq \mathbb{R}^m$ respectively. If f is differentiable at $a \in U$, and g is differentiable at $b = f(a) \in V$, then the composition $h = g \circ f$ is differentiable at a and its derivative is given by:

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

¹¹Since $\mathcal{L}(\mathbb{R}^n, \mathbb{R}^m)$ is a normed vector space, it makes sense to ask if Df is continuous.

Proof. Since f is differentiable at a , there is a function $\epsilon_1 : U \rightarrow \mathbb{R}^m$ such that

$$f(x) = f(a) + Df_a(x - a) + \|x - a\|\epsilon_1(x), \quad \forall x \in U, \quad (4.7)$$

where $\epsilon_1(x) \rightarrow 0 = \epsilon_1(a)$ as $x \rightarrow a$. In particular, by Lemma 4.5, there are constants $r, C > 0$ such that $\|f(x) - f(a)\| < C\|x - a\|$ for all $x \in B(a, r)$. Now since $b = f(a) \in V$ and V is open, we may find some $\delta > 0$ such that $B(b, \delta) \subseteq V$, and hence $g \circ f$ is defined on $B(a, r_1)$ where $r_1 = \min\{r, \delta/C\}$.

Now as g is differentiable at $b = f(a)$, there is a function $\epsilon_2 : V \rightarrow \mathbb{R}^p$ such that

$$g(y) = g(b) + Dg_b(y - b) + \|y - b\|\epsilon_2(y), \quad \forall y \in V, \quad (4.8)$$

and $\epsilon_2(y) \rightarrow 0 = \epsilon_2(b)$ as $y \rightarrow b = f(a)$. It follows that applying g to (4.7) and using (4.8) we have

$$\begin{aligned} h(x) &= g \circ f(x) \\ &= g(f(a) + Df_a(x - a) + \|x - a\|\epsilon_1(x)) \\ &= h(a) + Dg_b \circ Df_a(x - a) + Dg_b(\|x - a\|\epsilon_1(x)) + \|f(x) - f(a)\|\epsilon_2(f(x)) \\ &= h(a) + Dg_b \circ Df_a(x - a) + \|x - a\|\eta(x), \end{aligned}$$

where the final equality defines $\eta(x - a)$ for all $x \neq a$ and we set $\eta(a) = 0$. In order to show that $Dg_b \circ Df_a$ is the derivative of h at a , it therefore suffices to show that $\eta(x) \rightarrow 0$ as $x \rightarrow a$. Now for $x \neq a$, we have

$$\eta(x) = Dg_b(\epsilon_1(x)) + \frac{\|f(x) - f(a)\|}{\|x - a\|} \epsilon_2(f(x)).$$

But if r_1 is as above, for any $x \in B(a, r_1)$ we have $\|f(x) - f(a)\| \leq C\|x - a\|$, and hence

$$\begin{aligned} \|\eta(x)\| &= \|Dg_b(\epsilon_1(x)) + \frac{\|f(x) - f(a)\|}{\|x - a\|} \epsilon_2(f(x))\| \\ &\leq \|Dg_b(\epsilon_1(x))\| + C\|\epsilon_2(f(x))\|. \end{aligned}$$

Hence, since Dg_b and f are continuous at 0 and a respectively, it follows directly from the definitions that $\|\eta(x)\| \rightarrow 0$ as $x \rightarrow a$ as required. □

Remark 4.18. 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 $f'(a)$.

4.4 Real-valued functions of many variables 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 $f: U \rightarrow \mathbb{R}$ 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 $\delta: E \rightarrow E^*$, where $\delta(a)(v) = a \cdot v$ for all $a, v \in E$.

Definition 4.19. If $f: U \rightarrow \mathbb{R}$ is differentiable on U then we define $\nabla f: U \rightarrow E$ 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 4.20. 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, $|\nabla f(a) \cdot v| \leq \|\nabla f(a)\| \cdot \|v\| = \|\nabla f(a)\|$, 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 $\gamma: (-1, 1) \rightarrow \mathbb{R}^n$ 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 4.17:

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

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

4.5 Mean Value Theorems

For functions of a single variable, the Mean Value Theorem asserts that, if $f: U \rightarrow \mathbb{R}$ 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 $f(b) - f(a) = f'(c) \cdot (b - a)$. 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 dimensions greater than 1: 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 4.21. Let $f: \mathbb{R}^1 \rightarrow \mathbb{R}^2$ be given by $f(t) = (\cos(t), \sin(t))$. Then the derivative of f is $f'(t) = (-\sin(t), \cos(t))$, which is non-zero for all t . But if we take $a = 0$ and $b = 2\pi$ then $f(b) - f(a) = 0$, while $f'(t) \cdot (2\pi - 0) \neq 0$ for all $t \in \mathbb{R}$.

The way we will address the above example is rather pragmatics. Given normed vector spaces V and W and a differentiable function $f: U \rightarrow W$ defined on an open subset U of V , we will try and obtain information about $f(b) - f(a)$, for $a, b \in U$, by applying the single-variable Mean Value Theorem to compositions $f_{\gamma, \delta} = \delta \circ f \circ \gamma$, where $\gamma: I \rightarrow V$ is a “curve” joining a and b – that is, a function from an open interval $I \subseteq \mathbb{R}$ with $\gamma(0) = a$ and $\gamma(1) = b$ – and $\delta: W \rightarrow \mathbb{R}$ is an appropriate projection map:

$$\mathbb{R} \supseteq I \xrightarrow{\gamma} V \xrightarrow{f} W \xrightarrow{\delta} \mathbb{R}.$$

If $\gamma([0, 1]) \subseteq U$, so that $f_{\gamma, \delta}$ is defined in some open subset of \mathbb{R} containing $[0, 1]$, then the single-variable mean value theorem then gives an expression for $\delta(f(b)) - \delta(f(a))$, which the chain rule allows us to translate into a statement involving the derivative Df , but also involving the auxiliary data of $D\gamma$ and $D\delta$. In order for us to then obtain conclusions which only involve Df , we will require γ and δ to be linear, since then the derivative of $f_{\gamma, \delta}$ at $t = t_0$ is $Df_{\gamma, \delta, t_0} = \delta \circ Df_{\gamma(t_0)} \circ \gamma$. This requirement, in the case of γ , will impose a restriction on U the domain of definition of f , while in the case of δ it will mean we will obtain only a “Mean Value Inequality” rather than equality.

We first consider the case where $W = \mathbb{R}$, so that we do not need to consider the additional complication of an auxiliary function $\delta: W \rightarrow \mathbb{R}$.

Definition 4.22. Let V be a vector space. If $a, b \in V$ we write $\llbracket a, b \rrbracket$ for the line segment joining them, that is, $\llbracket a, b \rrbracket = \{ta + (1-t)b : 0 \leq t \leq 1\}$. The function $\gamma_{a,b}(t) = (1-t)a + tb$ is a path from a to b with image $\llbracket a, b \rrbracket$. A concatenation of finitely many such paths is called a *piecewise-linear* path. (See the Metric Spaces notes for more details.)

Proposition 4.23. Let $U \subset \mathbb{R}^n$ be an open subset of \mathbb{R}^n and let $f: U \rightarrow \mathbb{R}$ be a differentiable function. If $a, b \in U$, and the segment $\llbracket a, b \rrbracket \subseteq U$, then there exists an $\xi \in \llbracket a, b \rrbracket$, $\xi \notin \{a, b\}$, such that

$$f(b) - f(a) = Df_{\xi}(b - a).$$

Proof. Let $\gamma_{a,b}: [0, 1] \rightarrow U$ be defined by $\gamma_{a,b}(t) = (1-t)a + tb$ as above, and set $g(t) = f(\gamma_{a,b}(t))$. By the one-variable mean-value theorem, there is some $s \in (0, 1)$ such that $g(1) - g(0) = g'(s)$. Now by the chain rule,

$$\gamma'(s) = Df_{\gamma_{a,b}(s)} \gamma'_{a,b}(s),$$

and since $\gamma'_{a,b}(t) = b - a$ while $g(1) = f(b)$ and $g(0) = f(a)$, setting $\xi = \gamma_{a,b}(s)$ the result follows immediately. \square

Any easy application of this result is the following:

Proposition 4.24. *Suppose that U is a connected open subset of \mathbb{R}^n and $f: U \rightarrow \mathbb{R}^m$. 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 $\llbracket a, b \rrbracket \subseteq U$ is a line segment, Proposition 4.23 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. \square

If $f: U \rightarrow W$ is differentiable but takes values in a vector space of dimension greater than 1, we must take the components of f with respect to a basis of W and analyse them separately, or make do with an inequality rather than an equality. In practice the inequality is often the more useful of these choices, and we give such a result next. Before we state it, recall that a subset $U \subseteq \mathbb{R}^n$ is said to be *convex* if, for each $a, b \in U$ the line segment $\llbracket a, b \rrbracket \subseteq U$.

Theorem 4.25. (*Mean Value Inequality.*) *Let $U \subseteq \mathbb{R}^n$ be an open convex subset of \mathbb{R}^n and suppose that $f: U \rightarrow \mathbb{R}^m$ is differentiable. Suppose that $a, b \in U$ are such that $\llbracket a, b \rrbracket \subset U$. Then there is some $c \in \llbracket a, b \rrbracket \setminus \{a, b\}$ such that*

$$\|f(b) - f(a)\| \leq \|Df_c(b - a)\|.$$

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 Lipschitz continuous with constant K .

Proof. Fix $a, b \in U$, and define $g: U \rightarrow \mathbb{R}$ by $g(x) = (f(x) - f(a)) \cdot (f(b) - f(a))$, for any $x \in U$. Then $g(b) = \|f(b) - f(a)\|^2$ and $g(a) = 0$. Since the dot product is bilinear, the map $x \mapsto (f(b) - f(a)) \cdot x$ is linear, hence the chain rule shows that $Dg_x(v) = (f(b) - f(a)) \cdot Df_x(v)$, (for any $v \in \mathbb{R}^n$).

Now by Proposition 4.23 we have $\|f(b) - f(a)\|^2 = g(b) - g(a) = Dg_c(b - a)$ for some $c \in \llbracket a, b \rrbracket \setminus \{a, b\}$, and hence using the Cauchy-Schwarz inequality we see that

$$\|f(b) - f(a)\|^2 = \|(f(b) - f(a)) \cdot Df_c(b - a)\| \leq \|f(b) - f(a)\| \cdot \|Df_c(b - a)\|.$$

Thus we see that $\|f(b) - f(a)\| \leq \|Df_c(b - a)\|$ as required. For the final part, note that if U is convex then the above applies to all $a, b \in U$ and by the definition of the operator norm,

$$\|Df_c(b - a)\| \leq \|Df_c\|_\infty \|b - a\| \leq K \|b - a\|,$$

hence the result follows. \square

4.6 *Higher order derivatives

We briefly wish to discuss the notion of higher derivatives for multivariate functions. There are two ways of thinking about these, the first of which works concretely with partial derivatives, while the second works with the total derivative.

Suppose that V and W are finite-dimensional normed vector spaces and $f: U \rightarrow W$ is a differentiable function on an open subset U of V taking values in W . If we fix bases $\{v_1, \dots, v_n\}$ of V and $\{w_1, \dots, w_m\}$ of W respectively, let $\partial_j f_i$ denote the directional derivative of the i -th components of f in the direction v_j so that $(\partial_j f_i(a))$ is the Jacobian matrix of Df_a for $a \in U$.

Now the partial derivatives $\partial_j f_i$ are real-valued functions on U . We can thus consider all of their partial derivatives $\partial_{j_1} \partial_{j_2} f_i$, where $j_1, j_2 \in \{1, \dots, n\}$ and $i \in \{1, \dots, m\}$. If these all exist and are continuous, we say that f is twice continuously differentiable. Inductively we then have:

Definition 4.26. If $f: U \rightarrow W$ is as above and $f = \sum_{i=1}^m f_i \cdot w_i$ 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 $\beta = (j_1, j_2, \dots, j_{k-1}) \in \{1, 2, \dots, n\}^{k-1}$. The k -th partial derivatives of f are indexed by pairs (α, i) where $\alpha \in \{1, 2, \dots, n\}^k$ and $i \in \{1, 2, \dots, m\}$, where if $\alpha = (j_1, j_2, \dots, j_n)$ then setting $\beta = (j_2, \dots, j_n) \in \{1, 2, \dots, n\}^{k-1}$ we define

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

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

Remark 4.27. Theorem 4.13 shows that the $f \in C^1(U, \mathbb{R}^m)$ if and only if the total derivative exists and is continuous. Now the total derivative $Df: U \rightarrow \mathcal{L}(\mathbb{R}^n, \mathbb{R}^m)$, is a function on U taking values in the finite-dimensional vector space $\mathcal{L}(\mathbb{R}^n, \mathbb{R}^m)$. Using the standard bases of \mathbb{R}^n and \mathbb{R}^m respectively, we can identify $\mathcal{L}(\mathbb{R}^n, \mathbb{R}^m)$ with $\text{Mat}_{m,n}(\mathbb{R})$, the space of $m \times n$ matrices¹². The space $\text{Mat}_{m,n}(\mathbb{R})$ can then be identified with \mathbb{R}^{mn} , and the components of Df with respect to this identification are precisely the (first) partial derivatives of f . Theorem 4.13 thus shows that Df is continuously differentiable if and only if all the second partial derivatives exist and are continuous. Our definition of the spaces $C^k(U, \mathbb{R}^m)$ can thus be reformulated in terms of total derivatives rather than partial derivatives.

¹²If 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.

Example 4.28. The crucial case to understand in the above discussion is that of a twice-differentiable real-valued function on a n -dimensional vector space V . Let $B = \{e_1, \dots, e_n\}$ be a basis for V , and let $D = \{\delta_1, \dots, \delta_n\} \subset V^*$ be the dual basis. Suppose that U is an open subset of V and $f: U \rightarrow \mathbb{R}$ is twice differentiable on U . Then the derivative of f is a function $Df: U \rightarrow \mathcal{L}(V, \mathbb{R}) = V^*$. The components of Df with respect to the basis D of V^* are just the partial derivatives $\partial_i f$ of f , since if $Df_a = \sum_{j=1}^n Df_{a,j} \delta_j$ then

$$\partial_i f(a) = Df_a(e_i) = \sum_{j=1}^n Df_{a,j} \delta_j(e_i) = Df_{a,i}.$$

But now to calculate $D^2 f_a$ in terms of the second partial derivatives, we simply apply the same reasoning to each component $Df_{a,i}$ of Df : Indeed since the derivative is linear, we have

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

In the second equality we use the fact that if $w \in V^*$ and $g: U \rightarrow \mathbb{R}$, 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 $\mathcal{L}^2(V, \mathbb{R}) = \mathcal{L}(V, \mathcal{L}(V, \mathbb{R}))$ induced by our choice of basis $\{v_1, \dots, v_n\}$ of V is $\{\delta_j \cdot \delta_i : 1 \leq i, j \leq n\}$, where $\delta_j \cdot \delta_i$, viewed as an element of $\mathcal{L}^2(V, \mathbb{R})$, is given by $(\delta_j \cdot \delta_i)(v) = \delta_i(v) \cdot \delta_j \in V^*$ for any $v \in V$.

But the products $\delta_j \cdot \delta_i$ can also be thought of as *bilinear* maps on V taking values in \mathbb{R} , that is $(\delta_j \cdot \delta_i)(v_1, v_2) = \delta_i(v_1) \cdot \delta_j(v_2)$, and, since we have calculated that it is just a linear combination of the products $(\delta_j \cdot \delta_i)$, we can similarly view $D^2 f_a$ as a bilinear map on V : Explicitly, if we let $H_a = (\partial_{ji} f(a))$ be the *Hessian* matrix of $D^2 f_a$, and write $v_1 = \sum_{i=1}^n x_i e_i$ and $v_2 = \sum_{j=1}^n y_j e_j$, then

$$\begin{aligned} D^2 f_a(v_1)(v_2) &= D^2 f_a(v_1) \left(\sum_{j=1}^n y_j e_j \right) = \sum_{j=1}^n y_j D^2 f_a(v_1)(e_j) \\ &= \sum_{i,j} x_i y_j D^2 f_a(e_i)(e_j) = \sum_{i,j} x_i y_j \partial_{ji} f(a) \\ &= y^t H x. \end{aligned}$$

4.6.1 Symmetry of mixed partial derivatives

We now wish to show that, provided they result in a continuous function, the order in which the partial derivatives are taken does not matter, that is, if $f: U \rightarrow \mathbb{R}^m$, and $a \in U$, then for any $j_1, j_2 \in \{1, \dots, n\}$ we have

$$\partial_{j_1} \partial_{j_2} f(a) = \partial_{j_2} \partial_{j_1} f(a)$$

provided both second partial derivatives are continuous at a . To prove this we need the following:

Definition 4.29. Let $f: U \rightarrow \mathbb{R}$ be a function defined on an open set $U \subset \mathbb{R}^n$. Then if $s \in \mathbb{R} \setminus \{0\}$ and $j \in \{1, \dots, n\}$ let $\Delta_j^s(f)$ be the function given by

$$\Delta_j^s(f)(x_1, \dots, x_n) = \frac{f(x + s.e_j) - f(x)}{s},$$

Note that if f is differentiable at x then $\partial_j f(x) = \lim_{s \rightarrow 0} \Delta_j^s(f)(x)$.

It is straight-forward to check that, for any $s, t \in \mathbb{R} \setminus \{0\}$, and any $j_1, j_2 \in \{1, 2, \dots, n\}$ we have

$$\Delta_{j_1}^s \circ \Delta_{j_2}^t(f)(x) = \Delta_{j_2}^t(\Delta_{j_1}^s(f))(x).$$

Indeed a routine calculation shows that both sides are equal to

$$\frac{f(x + s.e_{j_1} + t.e_{j_2}) - f(x + s.e_{j_1}) - f(x + t.e_{j_2}) + f(x)}{st}.$$

Thus the difference operators $f \mapsto \Delta_{j_1}^s(f)$ and $f \mapsto \Delta_{j_2}^t(f)$ commute with each other. Moreover, since they are linear, the chain rule shows that they commute with partial differentiation: For all $j_1, j_2 \in \{1, \dots, n\}$ we have

$$\partial_{j_2} \Delta_{j_1}^s(f)(x) = \Delta_{j_1}^s(\partial_{j_2} f)(x). \quad (4.9)$$

We wish to use this fact to deduce that the corresponding partial differential operators also commute, but because of the limits involved, this will not be automatic, and we will need to impose the additional hypotheses that the relevant second partial derivatives of f are continuous functions.

Proposition 4.30. *Suppose that $f: U \rightarrow \mathbb{R}^m$ is such that all its second partial derivatives exist in a neighbourhood of $a \in U$. Then for any $i, j \in \{1, \dots, n\}$ we have*

$$\partial_i \partial_j f(a) = \partial_j \partial_i f(a)$$

whenever one of $\partial_i \partial_j f$ or $\partial_j \partial_i f$ is continuous at a .

Proof. Taking components we may immediately reduce to the case $m = 1$. Fix $a \in U$. By assumption we may pick $\epsilon, \delta > 0$ such that $\Delta_i^s(f)$ and $\Delta_j^t(f)$ are defined on $B(a, \epsilon)$ for all s, t with $|s|, |t| < \delta$, and we assume that $\partial_j \partial_i f$ is continuous at a (an identical argument, interchanging the roles of i and j establishes the result when $\partial_i \partial_j f$ is continuous at a).

Now using (4.9) and the fact that the difference operator Δ_i^s becomes ∂_i in the limit as $s \rightarrow 0$ we see that

$$\begin{aligned} \partial_i \partial_j f(a) &= \lim_{s \rightarrow 0} \Delta_i^s(\partial_j f)(a) = \lim_{s \rightarrow 0} \partial_j(\Delta_i^s(f)(a)) \\ &= \lim_{s \rightarrow 0} \lim_{t \rightarrow 0} (\Delta_j^t \circ \Delta_i^s(f)(a)) \end{aligned}$$

To establish the Proposition, we use the single-variable Mean Value Theorem to express $\Delta_j^t \circ \Delta_i^s(f)(a)$ in terms of $\partial_j \partial_i f$, and then use the continuity assumption to see that the limit

is $\partial_j \partial_i f(a)$. First we apply the one-variable mean value theorem to the function $g_i(y) = f(a + y.e_i)$ to obtain $\Delta_i^s(f) = \partial_i f(a + s_1.e_i)$ for some $s_1 \in (0, s)$, and hence

$$\Delta_j^t \circ \Delta_i^s(f)(a) = \Delta_j^t(\partial_i f(a + s_1.e_i))$$

Applying the single-variable Mean Value Theorem again, but now to the function $g_j(x) = \partial_i f(a + s_1.e_i + x.e_j)$, we see that $\Delta_j^t \partial_i f(a + s_1.e_i) = \partial_j \partial_i f(a + s_1.e_i + t_1.e_j)$, for some $t_1 \in (0, t)$, and hence

$$\Delta_j^t \circ \Delta_i^s(f)(a) = \partial_j \partial_i f(a + s_1.e_i + t_1.e_j)$$

But then as $(s, t) \rightarrow (0, 0)$ clearly (s_1, t_1) also tend to $(0, 0)$ and so by the continuity of $\partial_j \partial_i f$ at a , it follows

$$\partial_i \partial_j f(a) = \lim_{s \rightarrow 0} \lim_{t \rightarrow 0} (\Delta_j^t \circ \Delta_i^s(f)(a)) = \lim_{s \rightarrow 0} \lim_{t \rightarrow 0} \partial_j \partial_i f(a + s_1.e_i + t_1.e_j) = \partial_j \partial_i f(a)$$

by the continuity¹³ of $\partial_j \partial_i f$ at a . □

Example 4.31. The requirement that the second partial derivatives are continuous cannot simply be omitted. Indeed if we let $f: \mathbb{R}^2 \rightarrow \mathbb{R}$ be given by

$$f(x_1, x_2) = \begin{cases} x_1 x_2 \cdot \frac{x_1^2 - x_2^2}{x_1^2 + x_2^2} = x_1 x_2 (x_1 - x_2)(x_1 + x_2) / (x_1^2 + x_2^2) & (x_1, x_2) \neq (0, 0) \\ 0, & (x_1, x_2) = (0, 0). \end{cases}$$

The you can check that $\partial_1 \partial_2 f(0_2) = -1$ while $\partial_2 \partial_1 f(0_2) = +1$.

Remark 4.32. Returning to Example 4.28, we see now that, viewing $D^2 f(a)$ as a bilinear form, it has matrix $(\partial_{ij} f)_{ij}$, which, by Theorem 4.30, is a symmetric. Thus $D^2 f(a)$ can be viewed as a symmetric bilinear form.

Remark 4.33. Using induction and the previous Proposition, it follows that if $f: U \rightarrow \mathbb{R}^m$ is k -times continuously differentiable, then if $\alpha = (j_1, \dots, j_k) \in \{1, \dots, n\}^k$, and $\beta = (j_{\sigma(1)}, \dots, j_{\sigma(k)})$ is any re-ordering of the terms of α (so $\sigma \in S_k$ the symmetric group) the $\partial_\alpha f = \partial_\beta f$.

Example 4.34. Let $\square = \partial_1^2 - \partial_2^2$ be the (one-dimensional) wave operator. Provided we are only interested in acting on twice-continuously differentiable functions $u = u(x_1, x_2)$ so that $\partial_1 \partial_2(u) = \partial_2 \partial_1(u)$, we can factorize \square as

$$\square = (\partial_1 - \partial_2)(\partial_2 + \partial_1).$$

This leads to the classical D'Alembert solution of the one-dimensional wave equation.

¹³If you assume that $\partial_j \partial_i f$ is continuous in a neighbourhood of a this is easy, as for fixed s the limit as $t \rightarrow 0$ is just $\partial_j \partial_i f(a + s_1.e_i)$, and then taking the limit as $s \rightarrow 0$ gives the result. The iterated limit still equals $\partial_j \partial_i f(a)$ even if you only assume continuity at a , but the argument to show this is a little more elaborate.

5 The Inverse and Implicit Function Theorems

We begin with a result on the set of invertible linear maps.

Lemma 5.1. *Let $\Omega \subset \mathcal{L}(\mathbb{R}^n, \mathbb{R}^n)$ be the set of invertible linear maps from \mathbb{R}^n to itself. The we have*

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

Proof. Using the standard basis, we may identify $\mathcal{L}(\mathbb{R}^n, \mathbb{R}^n)$ with $\text{Mat}_n(\mathbb{R})$ the space of $n \times n$ matrices over \mathbb{R} . A linear map α is invertible if and only if its matrix A satisfies $\det(A) \neq 0$. The function \det is a polynomial function of the entries of A , hence it is continuous, thus the set of matrices corresponding to invertible linear maps is open as required.

If A is a matrix we may use Cramer's rule to calculate the matrix of its inverse: Recall that if $A = (a_{ij})$, then we write C_{ij} for the $(n-1) \times (n-1)$ given by removing the i -th row and j -th column from A . If we write $\text{adj}(A) = (c_{ij})$ where $c_{ij} = \det(C_{ji})$, then $A \cdot \text{adj}(A) = \det(A) \cdot \text{Id}$, and hence $A^{-1} = (\det(A))^{-1} \text{adj}(A)$. Since the determinant of a matrix is a polynomial function of its entries, the entries of $\text{adj}(A)$ are polynomials in the entries of A , it follows that the map $A \mapsto A^{-1}$ is continuous as required. \square

Remark 5.2. The first problem sheet gives another approach to the previous Lemma.

5.1 The Inverse Function Theorem

The following theorem is known as the Inverse Function Theorem. A complete proof is given in the Appendices.

Theorem 5.3. *Let $E \subset \mathbb{R}^n$ be an open set, $f: E \rightarrow \mathbb{R}^n$ a differentiable function, and let $a \in E$. Suppose that Df_a is invertible and that Df is continuous at $x = a$. Then there are open neighbourhoods U and V of a and $b = f(a)$ respectively, such that f is a bijection between U and V . Moreover if $g: V \rightarrow U$ is the inverse of f then g is differentiable with*

$$Dg_y = (Df_{g(y)})^{-1},$$

so that Dg is continuous at y whenever Df is continuous at $x = g(y)$. In particular, Dg is continuous at $f(a)$.

Proof. (Outline – proof not examinable.) By replacing f by $x \mapsto Df(a)^{-1}(f(x+a) - f(a))$, we may assume that $a = f(a) = 0$ and $Df(a) = I_n$. Let $\varphi: E \rightarrow \mathbb{R}^n$ be given by $\varphi(x) = f(x) - x$. Then clearly φ is differentiable and $D\varphi$ is continuous at 0_n with $D\varphi(0) = 0_{n,n}$. Thus there is an $r > 0$ such that $\|D\varphi(x)\| < 1/2$ for all $x \in \bar{B}(0, r)$. Now since $f(x) = x + \varphi(x)$, and hence $f(x) = y$ if and only if $\phi_y(x) = x$, where $\phi_y(x) = y - \varphi(x)$. But using the Mean Value Inequality, one can show that for $y \in B(0, r/2)$ the function ϕ_y is a contraction on

$\bar{B}(0, r)$, and hence there is a unique $x \in \bar{B}(0, r)$ with $f(x) = y$ for any $y \in \bar{B}(0, r/2)$. Define $g: \bar{B}(0, r/2) \rightarrow \bar{B}(0, r)$ by setting $g(y)$ to be the unique point in $\bar{B}(0, r)$ with $f(g(y)) = y$.

Again using the Mean Value Inequality one can check that g is actually continuous, and, restricting the domain appropriately, that g is in fact continuously differentiable: By the chain rule we know that if $y = f(x)$ then $Dg(y) = Df(g(y))^{-1}$, and then right-hand side we know to be a continuous function of y . What one needs to show, therefore, is that the candidate $Df_{g(y)}^{-1}$ is indeed the total derivative of g at y – once one checks this, the continuity of the Dg follows from the continuity of Df and the inversion map.

□

Remark 5.4. 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, V be open subsets of \mathbb{R}^n . We say that a continuously differentiable function $f: U \rightarrow V$ is a *diffeomorphism* if it is bijective, and its inverse $g: V \rightarrow U$ is continuously differentiable. (Warning: other references may only require f and g be differentiable, still others that f be infinitely differentiable). The inverse function theorem can then be stated as follows: Let $f: E \rightarrow \mathbb{R}^n$ be a continuously differentiable function on an open subset $E \subset \mathbb{R}^n$. If $Df(a)$ is invertible, then there is an open neighbourhood $U \subset E$ of a on which f is a 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 $Id = f \circ g$, shows that $Id = Df(g(y)) \circ Dg(y)$ and so $Dg(y) = Df(g(y))^{-1}$.
- It is not sufficient, even if just wanted f to have a continuous inverse, for the f to be differentiable with $f'(a)$ invertible: Consider the example $f: \mathbb{R} \rightarrow \mathbb{R}$, where $f(x) = x + 2x^2 \sin(1/x)$ (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 4.16, the function f is differentiable but not strongly differentiable at $x = 0$.]
- If $f: U \rightarrow \mathbb{R}^n$ 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 $U = \mathbb{R}^2 \setminus \{0\}$ and $f: U \rightarrow \mathbb{R}^2$ 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 $\mathbb{R}^2 \setminus \{0\}$. 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 $f: U \rightarrow \mathbb{R}^n$ is injective, then it is indeed a diffeomorphism from U to $f(U)$ – see below.

- The hypotheses of the theorem are not necessary for f to have a *continuous* inverse – the function $f: \mathbb{R} \rightarrow \mathbb{R}$ 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).

Finally we want to note a consequence of the theorem which is global.

Definition 5.5. Let (X, d) and (Y, ρ) be metric spaces. A continuous function $g: X \rightarrow Y$ 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 5.6. Let $U \subset \mathbb{R}^n$ be an open set, and $f: U \rightarrow \mathbb{R}^n$ 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 $a \in V$, and the inverse function theorem applies to $f|_V: V \rightarrow \mathbb{R}^n$ 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. \square

Remark 5.7. 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 5.8. Let $E \subset \mathbb{R}^n$ be an open subset and let $f: E \rightarrow \mathbb{R}^n$ 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 $h: f(E) \rightarrow E$ by setting $h(y)$ to be this point x . But then h is continuously differentiable by the inverse function theorem, since at any point $y \in f(E)$, if $x = h(y)$ there are open sets U, V containing x and y respectively, such that $f|_U: U \rightarrow V$ is a diffeomorphism. But then $g|_V$ is continuously differentiable, and so h is continuously differentiable at $y \in V$. \square

5.2 *Variants of the Inverse Function Theorem.

One can in fact weaken the hypotheses of the Inverse Function Theorem somewhat in a number of ways: if U is an open subset of \mathbb{R}^n and $f: U \rightarrow \mathbb{R}^n$ 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 $g = f|_{U_1}^{-1}: U_2 \rightarrow U_1$, 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 .

One can also prove a local result imposing a condition on f only at the point $a \in U$: namely that f is strongly differentiable at a (see Remark 4.16), that is for any $\epsilon > 0$ there is some open neighbourhood N of a such that for all $x, y \in N$ we have $f(x) - f(y) = Df_a(x - y) + \eta(x, y)$ where $\|\eta(x, y)\| \leq \epsilon\|x - y\|$. If one assumes only that $f: U \rightarrow \mathbb{R}^n$ is strongly differentiable at a , then there are open neighbourhoods U and V of a and $f(a)$ such that f restricts to give a homeomorphism between U to V , and moreover its inverse g is (strongly) differentiable at $y \in V$ precisely when f is (strongly) differentiable at $g(y)$.

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 3). 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 then its inverse is automatically continuous.

5.3 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 V is an n -dimensional vector space, then picking a basis $B = \{v_1, \dots, v_n\}$ of V gives us coordinates for the vectors in V : for any vector $v \in V$ we assign to it the coordinates $(c_1, \dots, c_n) \in \mathbb{R}^n$ where $v = \sum_{i=1}^n c_i v_i$. Equivalently, the basis defines an invertible linear map $\theta: V \rightarrow \mathbb{R}^n$ given by sending B to the standard basis of \mathbb{R}^n . Thus giving such a map is equivalent to giving a (linear) coordinate systems on V . In the setting of differentiable functions, diffeomorphisms play the same role: if U is an open subset of V and $f: U \rightarrow \mathbb{R}^n$ is a diffeomorphism, then we can use the components of f to parametrize the points in U .

Example 5.9. Suppose that V is 2-dimensional with basis $\{v_1, v_2\}$. The function $g: \mathbb{R}^2 \rightarrow E$ given by $g: (r, s) \mapsto r \cos(s).v_1 + r \sin(s).v_2$ has Jacobian determinant $J_g = r$, so that if $U = E \setminus \{t.v_1 : t \geq 0\}$ and $f: U \rightarrow (0, \infty) \times (0, 2\pi)$ is given by $g(f(x, y)) = (x, y)$, then f is a diffeomorphism on U , and the components of f give polar coordinates on U .

The domain U is chosen to ensure that f is injective, so that we do indeed get a diffeomorphism between an open subset of E and \mathbb{R}^2 . In fact, because $J_g \neq 0$ on all of $\mathbb{R}^2 \setminus \{0\}$, one can modify the definitions suitably so as to allow f to have range $(0, \infty) \times [0, 2\pi)$. At the origin however, we do not have C^1 coordinates, as Df_0 is singular so that f is not a

diffeomorphism at 0. Notice that polar coordinates fail to reflect whether a function is continuously differentiable at 0: the function $v \mapsto \|v\|$ is not differentiable at $(x, y) = (0, 0) \in E$ (see Example 4.8) while the function on \mathbb{R}^2 given by $(r, s) \mapsto r$ certainly is.

The Inverse Function Theorem ensures that if U is open in V and $f: U \rightarrow \mathbb{R}^n$ 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 V , then, the components of f provide a (non-linear) parameterization of neighbourhood of p . In the next definition when we write $f: (U, a) \rightarrow (V, b)$ we mean that $f: U \rightarrow W$ is a function from U to W and $f(a) = b$ where $a \in U$ and $b \in V$.

Definition 5.10. Suppose that V is a normed vector space and $p \in V$. A *system of local coordinates* at p is a diffeomorphism $\psi: (U, p) \rightarrow (\Omega, 0_n)$ from a connected open neighbourhood U of p to a connected open neighbourhood¹⁴ Ω of $0_n \in \mathbb{R}^n$ such that $\psi(p) = 0_n$. The coordinates (x_1, \dots, x_n) of \mathbb{R}^n at 0 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 $f: U \rightarrow \mathbb{R}^k$ is any function, then by the chain rule, $f \circ \psi^{-1}$ is continuously differentiable if f is, and similarly, if function $g: \Omega \rightarrow \mathbb{R}^k$ is continuously differentiable, then so is $g \circ \psi$, since, as ψ is a diffeomorphism, both ψ and ψ^{-1} are continuously differentiable.

In fact this shows that the map $\psi^*: C^1(U, \mathbb{R}^k) \rightarrow C^1(\Omega, \mathbb{R}^k)$ 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 $f: U \rightarrow \mathbb{R}^k$ 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 . Recall that one strategy in the study of linear maps is to try and find the “simplest” form of a matrix representing a given linear map. Though it is not always phrased that way, this is one way of stating the rank-nullity theorem: if V and W are vector spaces, and $\alpha: V \rightarrow W$ is a linear map, then one can find bases of V and W with respect to which α has matrix A where if $0_{r,s}$ denotes the 0-matrix of size $r \times s$, and I_k the $k \times k$ identity matrix, then

$$A = \left(\begin{array}{c|c} 0_{k,n} & I_k \\ \hline 0_{m-k,n} & 0_{m-k,k} \end{array} \right) \quad \dim(V) = n + k, \quad \dim(W) = m.$$

In the case where α is surjective, this becomes $A = (0_{k,n} | I_k)$, and we obtain two nice consequences:

1. Using the coordinates given by our choice of bases, the map α takes a particularly simple form: it just projects along the first n coordinates, that is, it is given by $(x_1, \dots, x_{n+k}) \mapsto (x_{n+1}, \dots, x_{n+k})$.

¹⁴Most of the time, it will not be important for U to be connected, but since U is chosen to be a neighbourhood of a , we may always replace U by its connected component which contains a , so the connectedness assumption imposes no additional constraint.

2. The kernel $\ker(\alpha)$ thus has basis given by the first n vectors in our basis, so that our basis also gives us a coordinate system, or parametrization, of the subspace $\ker(\alpha)$.

The Implicit Function Theorem is a non-linear version of this result, though the price for extending to the differentiable setting is that it will only hold locally. Before we state and prove it, however, it is instructive to see how the strategy we will use in our proof works in the linear case:

Lemma 5.11. *Suppose that E is a finite-dimensional vector space and $\alpha: E \rightarrow \mathbb{R}^k$ is a surjective linear map. Then there is a basis B of E such that the matrix of α with respect to B and the standard basis of \mathbb{R}^k is $(0_{n-k,k}, I_k)$*

Proof. As usual we write $\{e_1, \dots, e_k\}$ for the standard basis of \mathbb{R}^k . Pick a basis $\{v_1, \dots, v_n\}$ of $\ker(\alpha)$, and extend it to a basis $B_1 = \{v_1, \dots, v_{n+k}\}$ of E . Let $\{x_1, \dots, x_{n+k}\}$ be the dual basis of E^* associated to B_1 , so that if $v \in E$ then $v = \sum_{i=1}^{n+k} x_i(v) \cdot v_i$. Define $\beta: E \rightarrow E$ by $\beta(v) = \sum_{i=1}^n x_i(v) \cdot v_i + \sum_{i=1}^k \alpha_i(v) \cdot v_{n+i}$, where $\alpha(v) = \sum_{i=1}^k \alpha_i(v) \cdot e_i$. We claim that β is an isomorphism.

To see that β is injective, suppose that $\beta(v) = 0$. Then clearly $\alpha_i(v) = 0$ for each i , $1 \leq i \leq k$, hence $v \in \ker(\alpha)$. But then by definition, $v = \sum_{i=1}^n x_i(v) \cdot v_i$, and so $\beta(v) = v$ and hence $v = 0$. To see that β is surjective, note that if $v = \sum_{i=1}^{n+k} \lambda_i v_i$, then, since α is surjective, there is some $v_1 \in V$ with $\alpha(v_1) = (\lambda_{n+1}, \dots, \lambda_{n+k})$. But now we may write $v_1 = \sum_{i=1}^{n+k} a_i v_i$, and letting $v_2 = \sum_{i=1}^n (\lambda_i - a_i) v_i \in \ker(\alpha)$ it follows that $\beta(v_1 + v_2) = v$.

Thus we can apply β^{-1} to our basis B_1 to obtain a new basis $B = \{w_1, \dots, w_{n+k}\}$, where $w_i = \beta^{-1}(v_i)$. Since $\beta(v_i) = v_i$ for $1 \leq i \leq n$, clearly $w_i = v_i$ for $1 \leq i \leq n$, and so $\{w_1, \dots, w_n\}$ is a basis of $\ker(\alpha)$. Now consider $\{w_{n+1}, \dots, w_{n+k}\}$. By definition

$$\beta(w_{n+i}) = \sum_{j=1}^n x_j(w_{n+i}) v_j + \sum_{s=1}^k \alpha_s(w_{n+i}) v_{n+s} = v_{n+i},$$

hence $\alpha_s(w_{n+i}) = \delta_{s,i}$ (that is, equals 1 when $s = i$ and 0 otherwise) and hence $\alpha(w_{n+i}) = e_i$. It follows that α has matrix $(0_{k,n-k} | I_k)$ with respect to the basis B as required. \square

We are now ready to state and prove the Implicit Function Theorem:

Theorem 5.12. *(Implicit Function Theorem) Let E be a finite-dimensional normed vector space and let U be an open subset of E . If $f: U \rightarrow \mathbb{R}^k$ lies in $C^1(U, \mathbb{R}^k)$, and $p \in U \cap f^{-1}(0_k)$ such that $Df_p: V \rightarrow \mathbb{R}^k$ is surjective, then there is an open neighbourhood V of 0_E , a diffeomorphism $\psi: V \rightarrow V_1$ where V_1 is an open subset of U , and a basis $B = \{v_1, \dots, v_{n+k}\}$ of V*

1. $\psi(0_E) = p$ and $f \circ \psi(x_1, \dots, x_{n+k}) = (x_{n+1}, \dots, x_{n+k})$ where (x_1, \dots, x_{n+k}) are the coordinates on V given by B .
2. The level-set $f^{-1}(0) \cap V_1$ is precisely the image of ψ restricted to $E_1 \cap V$ where $E_1 = \text{span}\{e_1, \dots, e_n\}$, that is ψ gives a parametrization of $f^{-1}(0) \cap V_1$ so that if $q \in f^{-1}(0) \cap V_1$ then $q = \psi(x_1, \dots, x_n, 0_k)$ for some $(x_1, \dots, x_n, 0_k) \in E_1 \cap V$.

Proof. (Non-examinable.) Let $\alpha = Df_p: E \rightarrow \mathbb{R}^k$. By the previous Lemma we may pick a basis B of E with respect to which the matrix of α is just $(0_{k,n}, I_k)$. Let (x_1, \dots, x_{n+k}) denote the coordinates of E with respect to the basis B , so that if $v \in E$ then $v = \sum_{i=1}^{n+k} x_i(v).v_i$, and consider $G: U \rightarrow V$ given by

$$G(v) = \sum_{i=1}^n (x_i(v) - x_i(p)).v_i + \sum_{i=1}^k f_i(v).v_{n+i}, (v \in U)$$

where $f(v) = \sum_{i=1}^k f_i(v).e_i$, that is, the f_i are the components of $f: U \rightarrow \mathbb{R}^k$. Thus $G(p) = 0_E$, and the Jacobian matrix of DG_p with respect to the basis B is

$$DG_p = \begin{pmatrix} I_n & 0_n \\ 0_{k,n} & I_k \end{pmatrix} = I_{n+k}.$$

But now we may use the Inverse Function Theorem to see that G restricts to a diffeomorphism on some open neighbourhood $V_1 \subseteq U$ of p . Let $V = G(V_1)$, so that V is an open neighbourhood of $G(p) = 0_E$. Then setting $\psi = (G|_{V_1})^{-1}$ it follows $\psi: V \rightarrow V_1$ is a diffeomorphism. But now $v = \sum_{i=1}^{n+k} x_i(v).v_i \in V$, then

$$\sum_{i=1}^{n+k} x_i(v).v_i = G(\psi(v)) = \sum_{i=1}^n (x_i(\psi(v)) - x_i(p)).v_i + \sum_{i=1}^k f_i(\psi(v)).v_{n+i}$$

from which (1) and (2) follow immediately. \square

Remark 5.13. Note that the proof shows a little more than in the statement of the theorem: Let $E_n = \text{span}\{v_1, \dots, v_n\}$ and $E_k = \text{span}\{v_{n+1}, \dots, v_{n+k}\}$. Then $E = E_n \oplus E_k$ and we can write any $v \in E$ uniquely as $v = v^n + v^k$, where $v^n \in E_n$ and $v^k \in E^k$.

Then it is clear from the definition of G that if $\phi = \psi - p$ (so that $\phi(0) = 0$) then if $u \in E_n$ we have $\phi(u) = u + \phi(u)^k$, that is, we may view ϕ as the graph of the function $u \mapsto \phi(u)^k$ from a neighbourhood of $0 \in E_n$ to E_k . Hence the parametrization of $f^{-1}(0)$ near p actually exhibits $f^{-1}(0)$ as the graph of a function.

Remark 5.14. Another way to think of this result is as a differentiable analogue of the following linear algebra fact: If $\{v_1, \dots, v_k\}$ are linearly independent vectors in a finite-dimensional vector space E , then we may extend it to a basis $\{v_1, \dots, v_k, \dots, v_n\}$ of E . If $U \subset E$ is an open set with $f: U \rightarrow \mathbb{R}^k$, then if Df_p has maximal rank, i.e. rank k , then its components $\{f_1, f_2, \dots, f_k\}$ can be extended to a system of local coordinates near p .

In practice the vector space E may be identified with \mathbb{R}^{n+k} , i.e. it may already have a preferred choice of basis/coordinates. If $\alpha: E \rightarrow \mathbb{R}^k$ is a linear map, then identifying E with \mathbb{R}^{n+k} using our preferred choice of coordinates, let A the matrix of α with respect to the standard bases. At least in small examples, it is often easier to produce a set of k columns of A which are linearly independent (showing A has rank k) than it is to produce a basis for E with respect to which α takes the form $(0_{n,k}|I_k)$, i.e. producing an invertible matrix P such that $PAP^{-1} = (0_{n,k}|I_k)$.

In that context it can be useful to use the following variant of the Implicit Function Theorem. To state it we need some more notation: By reordering, we suppose that the last k columns of A are linearly independent. Now we may view $\mathbb{R}^{n+k} = \mathbb{R}^n \oplus \mathbb{R}_n^k$, where $\mathbb{R}_n^k = \text{span}\{e_{n+1}, \dots, e_{n+k}\}$, and we will write (x, y) for a vector in \mathbb{R}^{n+k} where $x \in \mathbb{R}^n$ and $y \in \mathbb{R}_n^k$. If $\alpha: \mathbb{R}^{n+k} \rightarrow \mathbb{R}^k$, then we may also decompose into $\alpha = \alpha_n \oplus \alpha_k$ where $\alpha_n: \mathbb{R}^n \rightarrow \mathbb{R}^k$ and $\alpha_k: \mathbb{R}_n^k \rightarrow \mathbb{R}^k$ are given by $\alpha_n(x) = \alpha(x, 0)$ and $\alpha_k(y) = \alpha(0, y)$, so that $\alpha(x, y) = \alpha_n(x) + \alpha_k(y)$. In terms of matrices, α_n has matrix A_n given by the first n columns of A , the matrix of α and α_k has matrix A_k given by the last k columns of A , that is, $A = (A_n | A_k)$. The analogue of Lemma 5.11 is then:

Lemma 5.15. *Let $\alpha: \mathbb{R}^{n+k} \rightarrow \mathbb{R}^k$ be a linear map, and let $\alpha = \alpha_n \oplus \alpha_k$ be its decomposition according to the direct sum $\mathbb{R}^{n+k} = \mathbb{R}^n \oplus \mathbb{R}_n^k$ as above. Then if α_k is invertible, there is a linear map $\theta: \mathbb{R}^n \rightarrow \mathbb{R}_n^k$ such that $\ker(\alpha) = \{(x, \theta(x)) : x \in \mathbb{R}^n\}$.*

Proof. If α has matrix A with respect to the standard bases, then $A = (A_n | A_k)$ where α_n has matrix A_n and α_k has matrix A_k . We use essentially the same argument as in Lemma 5.11: Let $\beta: \mathbb{R}^{n+k} \rightarrow \mathbb{R}^{n+k}$ be the linear map with matrix (with respect to the standard basis)

$$B = \left(\begin{array}{c|c} I_n & 0 \\ \hline A_n & A_k \end{array} \right)$$

In other words, $\beta(x, y) = (x, \alpha_n(x) + \alpha_k(y))$. Now since the diagonal blocks I_n and A_k are invertible, it follows B (and hence β) is invertible, indeed it is easy to calculate its inverse explicitly:

$$B^{-1} = \left(\begin{array}{c|c} I_n & 0 \\ \hline -A_k^{-1}A_n & A_k^{-1} \end{array} \right)$$

Thus if $\theta: \mathbb{R}^n \rightarrow \mathbb{R}_n^k$ is the linear map $-\alpha_k \circ \alpha_n$, we have $\beta^{-1}(x, y) = (x, \theta(x) + \alpha_k^{-1}(y))$. Then $(x, y) \in \ker(\alpha)$ if and only if $\beta(x, y) = (x, 0)$, if and only if

$$(x, y) = \beta^{-1}(x, 0) = (x, \theta(x) + \alpha_k^{-1}(0)) = (x, \theta(x)),$$

so that $\ker(\alpha) = \{(x, \theta(x)) : x \in \mathbb{R}^n\} = \text{graph}(\theta)$. □

Theorem 5.16. *Suppose that U is an open subset of \mathbb{R}^{n+k} , and $f: U \rightarrow \mathbb{R}^k$ is continuously differentiable. If $p = (x_0, y_0) \in U$ is such that $f(x_0, y_0) = 0$ and $Df_{p,k}$ is invertible, where $Df_p = Df_{p,n} \oplus Df_{p,k}$ is the decomposition of Df_p as above, then there is a continuously differentiable function $\psi: V_1 \rightarrow V$, where $V \subseteq U$ is an open neighbourhood of p , and V_1 is an open neighbourhood of 0_{n+k} , such that if $\psi(x, y) = (\psi_n(x, y), \psi_k(x, y))$ then $\psi_n(x, y) = x + x_0$, and if $(x, y) \in V$, then $f(x, y) = 0$ if and only if $(x, y) = (x, \psi_k(x - x_0, 0))$. Equivalently, if $g(x) = \psi_k(x - x_0, 0)$, then g is continuously differentiable, and if $(x, y) \in V$ then $f(x, y) = 0$ if and only if $y = g(x)$. Moreover, the derivative of g is given by*

$$Dg_x = -Df_{(x, g(x)), k}^{-1} \circ Df_{(x, g(x)), n}.$$

Proof. (Non-examinable:) The proof follows the same strategy as the proof of Theorem 5.12. Let $G: U \rightarrow \mathbb{R}^{n+k}$ be given by

$$G(x, y) = (x - x_0, f(x, y)) = (x - x_0, f_n(x, y) + f_k(x, y))$$

(thus strictly speaking we should write $G(x, y) = (x - x_0, i_n(f(x, y)))$ where $i_n: \mathbb{R}^k \rightarrow \mathbb{R}^{n+k}$ is the inclusion identifying \mathbb{R}^k with \mathbb{R}_n^k) so that $G(p) = G(x_0, y_0) = 0_{n+k}$. Then, for any $q = (x, y) \in U$ decomposing $Df_q = Df_{q,n} \oplus Df_{q,k}$ according to the direct sum decomposition $\mathbb{R}^{n+k} = \mathbb{R}^n \oplus \mathbb{R}_n^k$, we have

$$DG_q = \left(\begin{array}{c|c} I_n & 0 \\ \hline Df_{q,n} & Df_{q,k} \end{array} \right)$$

Thus $G \in C^1(U, \mathbb{R}^{n+k})$, since $f \in C^1(U, \mathbb{R}^k)$. Moreover, since $Df_{p,k}$ is invertible, it follows that DG_p is invertible. It follows from the Inverse Function Theorem that there is an open set $V_1 \subseteq U$ with $p \in V_1$ such that $G|_{V_1}: V_1 \rightarrow V = G(V_1)$ is a diffeomorphism. It follows that if $\psi: V_1 \rightarrow V$ is given by $(G|_{V_1})^{-1}$, the ψ is continuously differentiable, and, from the definition of G we must have $\psi(x, y) = (\psi_n(x, y), \psi_k(x, y)) = (x + x_0, \psi_k(x, y))$.

Finally, if $(x, y) \in V_1$, then $f(x, y) = 0$ if and only if $G(x, y) = (x, 0)$, hence

$$(x, y) = \psi \circ G(x, y) = \psi(x - x_0, 0) = (x, \psi_k(x - x_0, 0)).$$

Thus the theorem is proved except for the expression for the derivative of $g(x) = \psi_k(x - x_0, 0)$. 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

$$0 = \left(Df_{(x,g(x)),n} \mid Df_{(x,g(x)),k} \right) \left(\begin{array}{c} I_n \\ Dg_x \end{array} \right).$$

and hence $Df_{(x,g(x)),n} + Df_{(x,g(x)),k} Dg_x = 0$, so that $Dg_x = -Df_{(x,g(x)),k}^{-1} Df_{(x,g(x)),n}$. \square

Remark 5.17. This version of the Implicit Function Theorem explains the reason for the name: One can view it as saying that, provided the submatrix $Df_{(x_0,y_0),k}$ is invertible, the system non-linear of equations $f_i(x, y) = 0$ for $i = 1, 2, \dots, k$, can be solved. In other words, system $(f_i(x, y)) = 0_k$ *implicitly* makes the y -variables functions of the x -variables by these equations, and the theorem shows that, at least near (x_0, y_0) the function $g(x)$ makes this *explicit*.

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.

Example 5.18. 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 $f: \mathbb{R}^4 \rightarrow \mathbb{R}^2$ 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

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

The total derivative Df_z has Jacobian matrix

$$Df_z = (Df_{1,x} | Df_{2,y}) = \begin{pmatrix} 2x_1 & -2x_2 & 2y_1 & 4y_2 \\ 2x_1 & 2x_2 & -2y_1 & -2y_2 \end{pmatrix}, \quad (5.1)$$

Thus considering 2×2 submatrices, we see that Df has rank 0 only at $z = 0_4$, 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 5.16, 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 $f(x_1, x_2, y_1, y_2) = (1, 2)$, *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 (5.1) we write $Df_z = (Df_{1,x} | Df_{2,y})$. Explicitly this becomes:

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

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 $g_2 \cdot \partial_1 g_2 = -2x_1$ 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

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

***Remark 5.19.** 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.

6 Submanifolds of a normed vector space

6.1 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 \mathbb{R}^n .

Example 6.1. 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 $\phi: [-1, 1] \times [-\pi, \pi] \rightarrow \mathbb{R}^3$ 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 $S \in \text{graph}(h) = \{(x_1, x_2, h(x_1, x_2)) : (x_1, x_2) \in V\}$, where $V = \{(x_1, x_2) \in \mathbb{R}^2 : x_1^2 + x_2^2 < 1\}$.

Definition 6.2. Let $M \subseteq \mathbb{R}^n$ be a closed subset. We say that M is a k -dimensional *submanifold* of \mathbb{R}^n if, for every point $p \in M$, there is an open subset U of \mathbb{R}^n containing p and a smooth¹⁵ function $f: U \rightarrow \mathbb{R}^{n-k}$ 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 $\text{rank}(Df_p) = n - k$.

We say that M is C^k if we can choose $f \in C^k(U, \mathbb{R}^{n-k})$ 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 $\{e_1, \dots, e_k\}$.

Theorem 6.3. Let M be a k -dimensional submanifold of \mathbb{R}^n , and let $p \in M$. Then there is an open subset V of \mathbb{R}^n containing p , and a diffeomorphism $\psi: U \rightarrow V$ such that $M \cap V = \psi(U \cap \mathbb{R}^k)$. In particular, $\psi|_{U \cap \mathbb{R}^k}: U \cap \mathbb{R}^k \rightarrow M \cap V$ gives a parametrization of $M \cap V$.

Proof. By definition, there is an open set V_1 containing p and a function $f: V \rightarrow \mathbb{R}^{n-k}$ such that $V_1 \cap M = \{x \in V : f(x) = 0_{n-k}\}$, and $\text{rank}(Df_x) = n - k$ for all $x \in V_1$. But then Theorem 5.12 shows that there is a diffeomorphism $\psi: U \rightarrow V \subseteq V_1$, 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, \dots, t_n) given by $t_i = x_i \circ \psi^{-1}$, the function f is given by (t_{k+1}, \dots, t_n) (that is, for $v \in V_1$, we have $f(v) = (t_{k+1}(v), \dots, t_n(v))$). Moreover, the functions (t_1, \dots, t_k) parameterise the submanifold M on the open subset $M \cap V$ of M : if $(t_1, \dots, t_k, 0, \dots, 0) \in \mathbb{R}^k \cap U$, and we set $\phi(t_1, \dots, t_k) = \psi(t_1, \dots, t_k, 0, \dots, 0)$ then $\phi(t_1, \dots, t_k) \in M \cap V$ and if $u \in M \cap V$ then $u = \phi(t_1, \dots, t_k)$ for $t_i = x_i \circ \psi^{-1}$. □

Remark 6.4. As discussed in Remark 5.13, our proof of Implicit Function Theorem in fact shows that, at least locally, a submanifold M can be viewed as the graph of a C^1 function, or in terms of our more explicit version of the Implicit Function Theorem. if we pick some $(n - k) \times (n - k)$ of the Jacobian matrix of Df_{x_0} which is invertible, then in a neighbourhood of x_0 the equation $f(x) = 0$ allows us to express the corresponding coordinates as functions of the remaining k coordinates.

To put it another way, let us define a k -dimensional *subgraphold*¹⁶ of \mathbb{R}^n to be a subset $M \subseteq \mathbb{R}^n$ such that, for any point $a \in M$, there is an open neighbourhood U of a together with a decomposition $\mathbb{R}^n = V_1 \oplus V_2$ with $\dim(V_1) = k$, and a function $\psi \in C^1(U \cap (a + V_1), V_2)$ 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 any k -submanifold of \mathbb{R}^n is a k -subgraphold. In fact the converse is also true: indeed, as we show in Lemma 6.7 below, if

¹⁵At least continuously differentiable, but many texts automatically assume infinitely differentiable.

¹⁶The term is completely non-standard, and therefore, to honest, deliberately chosen to be clunky.

$V = V_1 \oplus V_2$ 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 subgraphoid 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 6.5.** Our definition of a k -dimensional sub-manifold M is a subset of \mathbb{R}^n which is locally given as a level-set for a C^1 -function f taking values in \mathbb{R}^{n-k} for which Df_x has rank $n - k$. Theorem 6.3 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 of \mathbb{R}^k , where $D\psi$ has rank k . This is *not* the same as requiring M to be locally given as the graph of a function, since the domain \mathbb{R}^k is not identified with a subspace of \mathbb{R}^n in such a way that the image of ψ takes values in a complementary subspace.

Nevertheless, it turns out to be true that if $M \subseteq \mathbb{R}^n$ is locally given as the image of an injective C^1 -map from a suitable open subset of \mathbb{R}^k whose derivative has rank k , then M is a sub-manifold in the sense of Definition 6.3: More precisely, if $M \subseteq \mathbb{R}^n$ is a subset of \mathbb{R}^n such that, for each $a \in M$ there is an open neighbourhood U of a , an injective function $\phi \in C^1(B(0, R), \mathbb{R}^n)$ defined on the open ball $B(0_k, R)$ of $0_k \in \mathbb{R}^k$ with Df_x injective for all $x \in B(0_k, R)$, such that¹⁷ for some $0 < r < R$ we have $\phi(B(0_k, r), 0_k) = (M \cap U, a)$, then M is a k -dimensional submanifold of \mathbb{R}^n . This can be proved using the same ideas as we used in the proof of the implicit function theorem: to deduce the implicit function theorem from the inverse function theorem we effectively extend a linearly independent set in V^* to a basis, while in this case one needs to extend a linearly independent set in V to a basis. The details are given in Section 8.2 of the Appendix.

Example 6.6. Suppose that $g: \mathbb{R}^2 \rightarrow \mathbb{R}$ is given by $g(x_1, x_2) = x_1 x_2$. Then $Dg_{(x_1, x_2)} = (x_2, x_1)$ and hence $\text{rank}(Dg_{(x_1, x_2)}) = 1$ 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 6.7. *Let $V = V_1 \oplus V_2$ and suppose that $\psi \in C^1(\Omega, V_2)$ is a continuously differentiable function on an open subset Ω of V_1 taking values in V_2 . Then the graph $\Gamma(\psi) = \{(v, \psi(v)) : v \in \Omega\}$ is a submanifold of V .*

¹⁷That is, $\phi(\Omega) = M \cap U$ and $\phi(0_k) = a$.

Proof. But if we let $g: \Omega_1 \times V_2 \rightarrow V_2$ be given by $g(v_1, v_2) = v_2 - \phi(v_1)$, then clearly $g \in C^1(\Omega_1 \times V_2, V_2)$ and $(v_1, v_2) \in \Gamma(\phi)$ if and only if $g(v_1, v_2) = 0$. Moreover, if $a = a_1 + a_2 \in V_1 \oplus V_2$, then $Dg_{(a_1, a_2)}(v_1, v_2) = -D\phi_{a_1}(v_1) + v_2$. Thus for any $v_2 \in V_2$ and any $a \in \Omega_1 \times V_2$ we have $Dg_a(0, v_2) = v_2$, and hence the derivative $Dg_{(a_1, a_2)}$ is surjective for all $a \in \Omega_1 \times V_2$. Picking a basis for V_2 we identify it with \mathbb{R}^{n-k} , and thus $\Gamma(\phi)$ is a k -submanifold of \mathbb{R}^n . \square

Example 6.8. 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 differentiable 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 $\Gamma(f) = \{(x, y) \in U \times \mathbb{R} : g(x, y) = 0 \text{ where } g(x, y) = y - f(x)\}$. 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} .

Example 6.9. Suppose that $n \in \mathbb{R}^3$ is a unit vector and $C = \{x = (x_1, x_2, x_3) \in \mathbb{R}^3 : x_1^2 + x_2^2 - x_3^2 = 0, n \cdot x = d\}$. Then C is a level set of the function $f: \mathbb{R}^3 \rightarrow \mathbb{R}^2$, where $f(x_1, x_2, x_3) = (x_1^2 + x_2^2 - x_3^2, n_1 x_1 + n_2 x_2 + n_3 x_3)$: indeed $C = f^{-1}((0, d))$. We have

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

Then Df has rank 2 off the line $\mathbb{R} \cdot (n_1, n_2, -n_3)$. 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 $\mathbb{R} \cdot (n_1, n_2, -n_3)$ does not intersect the surface $f_1(x) = 0$, 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 5.16 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 $I_\phi = \mathbb{R} \setminus (-2, 2)$ 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 equivalently¹⁸ $(s_1, s_2) \mapsto \rho(s_1, s_2, s_1)$, 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 $\ell = d / \cos(\phi)$. We then find that C can be parameterized by $s \in \mathbb{R}$ via

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

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 5.1, which can be solved in this case by separation of variables.

6.2 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 \mathbb{R}^n .

Definition 6.10. Let X be a subset of a normed vector space V and let $p \in X$. A *path on X centred at p* is a function $\gamma \in C^1((-r, r), V)$, where $r > 0$, such that the image of γ lies in X and $\gamma(0) = p$. We write $\mathcal{P}(X, p)$ for the set of all paths on X centred at p . Let $T : \mathcal{P}(X, p) \rightarrow V$ be the map given by $T(\gamma) = \gamma'(0)$. The image of T is called the *tangent space* to X at p and is denoted $T_p X$.

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

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

Remark 6.11. Note that while the normal space N_pX is by definition a linear subspace of V , the tangent space need not in general be a linear subspace (see Example 6.16). Indeed since $T_aX \subseteq (T_pX^\perp)^\perp = N_pX^\perp$ with equality if and only if T_pX is itself a linear subspace of V . Thus N_aX^\perp is the smallest subspace of V containing T_aX , that is, N_pX is the linear span of T_pX . We will shortly see that $T_pX = N_pX^\perp$ when X is a submanifold.

Remark 6.12. Let V be a normed vector space and $X \subseteq Y \subseteq V$ be subsets. For any $p \in X$ clearly $\mathcal{P}(X, p) \subseteq \mathcal{P}(Y, p)$ and hence $T_pX \subseteq T_pY$.

Slightly less trivially, if $p \in X$ and U is an open subset containing p , then $T_p(U \cap X) = T_pX$. Since $X \cap U \subseteq X$, by the above we see that $T_p(U \cap X) \subseteq T_pX$. For the reverse inclusion, note that if $v \in T_pX$ then we may pick a path $\gamma \in \mathcal{P}(X, a)$ with $T(\gamma) = a$. Then γ is continuous, so $\gamma^{-1}(U)$ is an open neighbourhood of 0 (since $\gamma(0) = a$) and so contains an open interval of the form $(-s, s)$. Let $\gamma_s = \gamma|_{(-s, s)}$. Then $\gamma_s \in \mathcal{P}(X \cap U, a)$, 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 $v \in T_p(U \cap X)$.

Thus the tangent space T_aX of X at a is only sensitive to the nature of X near a . 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 6.13. Let V be a normed vector space and let X be any subset of X . If $\psi \in C^1(U, W)$, and $p \in U \cap X$, then the derivative of ψ at p induces a map

$$D\psi_p: T_aX \rightarrow T_qY,$$

for any $Y \supseteq \psi(U \cap X)$, and $q = \psi(p)$.

Proof. Let $v \in T_pX$. By Remark 6.12, we may assume that $v = T(\gamma)$ for $\gamma \in \mathcal{P}(X \cap U, a)$. But then $\psi \circ \gamma \in \mathcal{P}(\psi(U \cap X), b) \subseteq \mathcal{P}(Y, b)$, so that $T(\psi \circ \gamma) \in T_bY$. But by the Chain Rule,

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

so that $D\psi_a(v) \in T_bY$ as required. \square

Corollary 6.14. Let V and W be normed vector spaces, U an open subset of V , and X any subset of V . Suppose that $\psi \in C^1(U, W)$ and $p \in U \cap X$. Then we have the following:

1. If $D\psi_p$ is an invertible linear map, then $D\psi_p$ gives a bijection between T_pX and T_qY , where $q = \psi(p)$ and $Y = \psi(U \cap X)$.
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 , an open subset of W containing $q = \psi(p)$. But then if $\theta: \Omega \rightarrow U_1$ is the inverse of ψ , by Lemma 6.13 applied to ψ and θ , we have $D\psi_p: T_pX \rightarrow T_qY$ and $D\theta_q: T_qY \rightarrow T_pX$, and $D\psi_p$ and $D\theta_q$ are inverse, the result follows.

For the second part, Lemma 6.13 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 $T_pX \subseteq \ker(D\psi_p)$. \square

Example 6.15. If M is a k -submanifold of V , 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 6.12 and Corollary 6.14 part (2), we see that

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

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 $T_p X \subseteq \ker(Df_p)$ 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 $T_p M = \ker(Df_p)$.

Example 6.16. Consider Example 6.6 again, that is, let $g: \mathbb{R}^2 \rightarrow \mathbb{R}$ the continuously differentiable function given by $g(x_1, x_2) = x_1 \cdot x_2$, and, for $c \in \mathbb{R}$ let $L_c = \{(x_1, x_2) \in \mathbb{R}^2 : x_1 \cdot x_2 = c\}$. 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 6.14 shows that $T_a(L_c) \subseteq \ker(Dg_a) = \{(x_1, x_2) : a_2 x_1 + a_1 x_2 = 0\}$, while at $a = 0$ we only get the trivial bound $T_0 L_0 \subseteq \ker(Dg_0) = \mathbb{R}^2$. In fact you can check that $T_a L_c = \ker Dg_a$ for all $a \neq 0$, while at $a = 0$, $T_0 L_0 = L_0$, giving an example where the tangent space of a level-set is not a linear subspace.

Example 6.17. 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 $f(x) = (x_{k+1}, \dots, x_n)$. Then it is clear that Df_0 has kernel given by $\text{span}_{\mathbb{R}}\{e_1, \dots, e_k\}$. On the other hand, if $v = (v_1, \dots, v_k, 0, \dots, 0)$, then $\gamma(t) = t \cdot v$ lies in M , and $\gamma'(0) = v$, hence we see that $v \in T_0 M$ if and only if $Df_0(v) = 0$.

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

Proposition 6.18. *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 $f: U \rightarrow \mathbb{R}^{n-k}$ is continuously differentiable with Df_x of maximal rank for all $x \in U$. Then we have*

$$T_p M = \ker(Df_p).$$

In particular, $T_p M$ is a k -dimensional vector subspace.

Proof. We have already shown the containment $T_p M \subseteq \ker(Df_p)$ in Corollary 6.14, so it remains to establish the reverse inclusion. In the case where $f = (x_{k+1}, \dots, x_n)$ 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 $\psi: V \rightarrow U$ from an open neighbourhood V of 0_n to an open neighbourhood U of p taking $N \cap V$ to $M \cap U$ where $N = \{x \in U : (x_{k+1}, \dots, x_n) = 0_{n-k}\}$. The result then follows from Lemma 6.13. \square

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

Proposition 6.19. 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 $f: U \rightarrow \mathbb{R}^{n-k}$ is a continuously differentiable function, then if $f = (f_1, \dots, f_{n-k})$ we have

$$T_p M^\perp = \text{span}_{\mathbb{R}}\{\nabla f_1(p), \dots, \nabla f_{n-k}(p)\}.$$

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

Proof. By Proposition 6.18, the tangent space $T_p M = \ker(Df_p)$ is a k -dimensional subspace of \mathbb{R}^n . Let $f = (f_1, \dots, f_{n-k})$ and let $N = \text{span}_{\mathbb{R}}\{\nabla f_1(p), \dots, \nabla f_{n-k}(p)\}$, 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 $v \in T_p M$ if and only if $v \in N^\perp$. Thus $T_p M = N^\perp$ and hence $N = T_p M^\perp$ as required (since, for any subspace W of an inner product space V we have $(W^\perp)^\perp = W$). \square

Example 6.20. 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

$$\begin{aligned} T_a S &= \{v = (v_1, v_2, v_3) \in \mathbb{R}^3 : 2a_1 \cdot v_1 + 4a_2 \cdot v_2 - 14a_3 \cdot v_3 = 0\}, \\ T_p S^\perp &= \{\lambda \cdot (2a_1, 4a_2, -14a_3) : \lambda \in \mathbb{R}\} \end{aligned}$$

Example 6.21. Let $O_n(\mathbb{R}) = \{X \in \text{Mat}_n(\mathbb{R}) : X \cdot X^T = I_n\}$ be the orthogonal group, the group of linear isometries of \mathbb{R}^n (equipped with the $\|\cdot\|_2$ -norm). We claim this is a smooth submanifold of $\text{Mat}_n(\mathbb{R})$ of dimension $n(n - 1)/2$.

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

$$q(X + H) = (X + H) \cdot (X + H)^T = q(X) + H \cdot X^T + X \cdot H^T + H \cdot H^T,$$

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

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

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

so that Dq is surjective as required.

The group $O_n(\mathbb{R})$ 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 $\ker(Dq_{I_n}) = \{H \in \text{Mat}_n(\mathbb{R}) : H + H^T = 0\}$. 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 derivative of the group structure on $O_n(\mathbb{R})$. This is studied in detail in courses in Part C.

6.3 *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 $\psi: B(0, r) \rightarrow U$ such that $\psi(V \cap (\mathbb{R}^k \oplus 0_{n-k})) = M \cap U$. If we write $\psi^{-1}(x) = (t_1, \dots, t_n)$, then if $f: M \cap U \rightarrow \mathbb{R}$ is any function, we may define $\tilde{f}: U \rightarrow \mathbb{R}$ by

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

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, \dots, 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” $\{\phi_i: U_i \rightarrow V_i : i \in I\}$, 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 $f: M \rightarrow \mathbb{R}$ 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 $\psi_j^{-1} \circ \psi_i: U_i \cap U_j \rightarrow U_i \cap U_j$ is diffeomorphism, and this is exactly the compatibility condition which is imposed. Abstract differentiable manifolds are studied in the Part C course “Differentiable Manifolds”.

7 Lagrange multipliers

Suppose U is an open subset of \mathbb{R}^n and $g: U \rightarrow \mathbb{R}$ is differentiable. We wish to study local extrema of g – unconstrained at first, and then on submanifold $M \cap U \subset \mathbb{R}^n$.

Lemma 7.1. *Suppose that $p \in U$ is a local minimum of g , so that there is an $r > 0$ such that if $x \in B(p, r) \subseteq U$, then $g(x) \geq g(p)$. Then $Dg_p = 0$.*

Proof. It is convenient to use the gradient vector field ∇g . Suppose for the sake of contradiction that $Dg_p \neq 0$. Then $\nabla g(p) \neq 0$ and we may set $u = \nabla g(p)/\|\nabla g(p)\|$, so that u is the vector of unit length in the direction of $\nabla g(p)$. Then the curve $\gamma: (-r, r) \rightarrow \mathbb{R}^n$ given by $\gamma(t) = p + t.u$ has image in U , so that we may define $G(t) = g(\gamma(t))$. By the chain rule, we have

$$G'(0) = Dg(\gamma(0))(\gamma'(0)) = Dg_p(u) = \nabla g(p) \cdot u = \|\nabla g(p)\| > 0.$$

Thus, by Lemma 4.1, we have

$$G(t) = G(0) + G'(0).t + \epsilon(t).|t| = G(0) + t.(G'(0) \pm \epsilon(t)),$$

where the first equality defines $\epsilon(t)$ for $t \neq 0$ and $\epsilon(t) \rightarrow 0 = \epsilon(0)$ as $t \rightarrow 0$. Thus since $G'(0) > 0$ it follows that for all sufficiently small negative t we must have $G(t) < G(0) = g(p)$, contradicting the fact that $g(p)$ is a local minimum. \square

We now wish to study the problem of minimizing $g: U \rightarrow \mathbb{R}$ given constraints on $x \in U$. Before formulating the general result, consider the problem of trying to minimize a function $g: \mathbb{R}^3 \rightarrow \mathbb{R}$ 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 7.2. 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 minimize $g(x) = x_3$ 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 $Df_x = 2(x_1, x_2, x_3)$, we have $2\nabla g(p) + \nabla f(p) = 0$, so that $\nabla g(p)$ is normal to the surface S at the extreme point $p \in S$.

This example is not a coincidence: if we consider the proof of the previous lemma, the strategy relies on the fact that perturbing p in the direction $-\nabla g(p)$ should decrease the value of g . Now if we are to stay on S , then we cannot necessarily move along a curve in the direction of $\nabla(g)$. However, unless $\nabla g(p)$ is perpendicular to $T_p S$, we may write $\nabla g(p) = v + n$ where $v \in T_p S$ and $n \in T_p S^\perp$, with $v \neq 0$. We can check that if we perturb in the direction of $-v$, then g decreases:

To make this precise, since $v \in T_p S$, we can find a curve $\gamma: (-r, r) \rightarrow S$ such that γ is continuously differentiable, $\gamma(0) = p$, and $\gamma'(0) = v$. Then if $G(t) = g(\gamma(t))$, we have

$$G'(t) = Dg(\gamma(t))(\gamma'(t)) = Dg(p)(v) = \nabla g(p) \cdot v = \|v\| > 0,$$

hence, as in the proof of Lemma 7.1, for all sufficiently small negative t , we have $g(\gamma(t)) < g(p)$, and so p cannot be a local minimum of g on S . It follows that at such a local minimum, we must have $\nabla g(p) \in T_p S^\perp$. But we know that $T_p S^\perp$ is spanned by $\nabla f(p)$ where $S = \{x \in \mathbb{R}^3 : f(x) = 0\}$, so that the condition that $\nabla g(p) \in T_p S^\perp$ is equivalent to the existence of scalars $\lambda_0, \lambda_1 \in \mathbb{R}$ such that

$$\lambda_0 Dg(p) + \lambda_1 Df(p) = 0.$$

The scalars λ_0, λ_1 are known as *Lagrange multipliers*. The following theorem shows that this analysis works in general:

Theorem 7.3. Let U be an open subset of \mathbb{R}^n , and let $g: U \rightarrow \mathbb{R}$ be continuously differentiable. If M is a k -submanifold of \mathbb{R}^n , and $x_0 \in U$ is a local minimum on M , then $\nabla g(x_0) \in T_{x_0}M^\perp$. Equivalently, if V is an open neighbourhood of x_0 and $f: V \rightarrow \mathbb{R}^{n-k}$ is continuously differentiable with $M \cap V = f^{-1}(0)$, and components $f = (f_1, \dots, f_{n-k})$, there are scalars $\lambda_0, \dots, \lambda_{n-k} \in \mathbb{R}$, with $\lambda_0 \neq 0$ such that

$$\lambda_0 Dg(x_0) + \sum_{i=1}^{n-k} \lambda_i Df_i(x_0) = 0. \quad (7.1)$$

Proof. As above, if $g(x_0)$ is a local minimum on M , and M is given by the vanishing of f near x_0 , then by Proposition 6.19, $T_{x_0}M^\perp$ has basis $\{\nabla f_1(x_0), \dots, \nabla f_{n-k}(x_0)\}$, so that $\nabla g(x_0) \in T_{x_0}M^\perp$ if and only if (7.1) holds for some scalars λ_i , $0 \leq i \leq n-k$ with $\lambda_0 \neq 0$.

Thus it remains to check that at a local minimum on M , we have $\nabla g(x_0) \in T_{x_0}M^\perp$. Suppose for the sake of a contradiction this is not the case. Then we may write $\nabla g(x_0) = v + n$, where $v \in T_{x_0}M$ and $n \in T_{x_0}M^\perp$, where $v \neq 0$. Then we may find a continuously differentiable curve $\gamma: (-r, r) \rightarrow \mathbb{R}^n$ whose image is in M , with $\gamma(0) = x_0$ and $\gamma'(0) = v$.

Then let $G(t) = g(\gamma(t))$. By the chain rule,

$$G'(0) = Dg_{x_0}(v) = \|v\| > 0,$$

and hence for all sufficiently small negative t , we have $G(t) < G(0) = g(x_0)$, contradicting the minimality of $g(x_0)$. □

Remark 7.4. Since the theorem asserts that $\lambda_0 \neq 0$, one can, by rescaling the linear dependence, assume that $\lambda_0 = 1$.

Remark 7.5. If we wish to optimize g over a set $N = \{x \in \mathbb{R}^n : f(x) = 0\}$, where $f: V \rightarrow \mathbb{R}^{n-k}$ is a C^1 -function on an open subset V of \mathbb{R}^n , but where we do not know the rank of Df . If x_0 is a local extremum, . The proof still shows, however, that at a local extremum $x_0 \in N$, $\nabla g(x_0)$ must lie in the normal space $T_{x_0}N^\perp$. However, without the maximal rank condition, this gives less information: the gradient vector fields $\nabla f_i(x_0)$ of the components of f no longer have to be linearly independent, and the containment $T_{x_0}N \subseteq \ker(Df_{x_0})$ need not be an equality.

Remark 7.6. Despite the above remark, the Lagrange Multiplier theorem gives a remarkably powerful general technique for finding extrema in constrained optimization problems.

Example 7.7. Consider the problem of finding the extrema of the function $g: \mathbb{R}^3 \rightarrow \mathbb{R}$ 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 both $f_1(x) = x_1^2 + x_2^2 = 2$ and $f_2(x) = x_1 + x_2 + x_3 = 1$, 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 $f: \mathbb{R}^3 \rightarrow \mathbb{R}^2$, 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

$$\nabla g(c) = \lambda_1 \nabla f_1(c) + \lambda_2 \nabla f_2(c),$$

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 7.8. 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 $|a \cdot b| \leq \|a\| \cdot \|b\|$. 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 constraints that $\|x\| = \|y\| = 1$.

Let us formulate this in the language of Theorem 7.3. Let $g: \mathbb{R}^{2n} = \mathbb{R}^n \oplus \mathbb{R}_n^n$ be given by $g(x, y) = x \cdot y$ (thus we use the same notational conventions as in Theorem 5.16) and let $f: \mathbb{R}^{2n} \rightarrow \mathbb{R}^2$ 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 $\text{rank}(Df_{(x_0,y_0)}) = 2$, so that S is a $2n - 2$ -dimensional submanifold of \mathbb{R}^{2n} . Hence, by Theorem 7.3, 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

$$\nabla g(x_0, y_0) = \lambda_1 \nabla f_1(x_0, y_0) + \lambda_2 \nabla f_2(x_0, y_0)$$

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

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

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

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

But then we must have $y_0 = \lambda_1 \cdot x_0$ and $x_0 = \lambda_2 \cdot y_0$, 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 $g(x_0, x_0) = \|x_0\|^2 = 1$ and $g(x_0, -x_0) = -\|x_0\|^2 = -1$, 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 $x = \pm y$.

References

- [B] B. Green, *Metric spaces*, lecture notes for A2 “Metric spaces and complex analysis”, MT 2020, available at <https://courses.maths.ox.ac.uk/node/50681/materials>
- [R] J. Robbin, *On the existence theorem for differential equations*, Proc. Amer. Math. Soc. 19 (1968), 1005–1006.

8 *Appendix

8.1 Proof of the Inverse Function Theorem

The following result is key to establishing the Inverse Function Theorem. It should be interpreted as saying that if we perturb the identity map $\text{id}: U \rightarrow U$ by a small enough function ϕ then the resulting function is still invertible.

Lemma 8.1. *Suppose U is an open subset of \mathbb{R}^n and that $f: U \rightarrow \mathbb{R}^n$ is differentiable on U . If $0 \in U$ and Df is continuous at $0 \in U$, with $Df_0 = I_n$, then if $\varphi: U \rightarrow \mathbb{R}^n$ is given by $\varphi(x) = f(x) - x$, there is an $r > 0$ such that for all $x, y \in \bar{B}(0, r) \subset U$,*

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

Proof. By definition, if f is differentiable at $x \in U$ then $D\varphi_x = Df_x - I_n$, so that $D\varphi_0 = 0$. Since $D\varphi$ is continuous at a , there is an $r > 0$ such that $\|D\varphi_x\|_\infty \leq 1/2$ for all $x \in B(0, r_1)$. But then by the Mean Value Inequality (Theorem 4.25), 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 $r \in (0, r_1)$. \square

The next Proposition is the key step in proving Inverse Function Theorem. Roughly speaking, it says that any function which is a small enough perturbation of the identity map should remain a bijection – c.f. Q.5 on Problem Sheet 1.

Proposition 8.2. *Let $U \subset \mathbb{R}^n$ be an open neighbourhood of 0, and let $\varphi: U \rightarrow \mathbb{R}^n$ be a function such that $\varphi(0) = 0$ and that, for some $r > 0$, we have*

$$\|\varphi(x) - \varphi(y)\| \leq \frac{1}{2} \|x - y\| \quad \forall x, y \in B(0, r).$$

Then if $f: U \rightarrow \mathbb{R}^n$ is given by $f(x) = x + \varphi(x)$, and $y \in \bar{B}(0, r/2)$, there is a unique $x \in \bar{B}(0, r)$ such that $f(x) = y$. Moreover, the function $g: \bar{B}(0, r/2) \rightarrow \bar{B}(0, r)$ defined by $f(g(y)) = y$ is continuous.

Proof. Given $y \in \bar{B}(0, r/2)$, let $\varphi_y(x) = y - \varphi(x)$. Then we have

$$\|\varphi_y(x)\| = \|y - \varphi(x)\| \leq \|y\| + \|\varphi(x)\| \leq r/2 + r/2 = r,$$

so that φ_y maps $\bar{B}(0, r)$ to itself. Since $\bar{B}(0, r) \subset \mathbb{R}^n$ is closed and \mathbb{R}^n is complete, $\bar{B}(0, r)$ itself is complete and non-empty. Moreover,

$$\|\varphi_y(x) - \varphi_y(x')\| = \|\varphi(x') - \varphi(x)\| \leq \frac{1}{2} \|x - x'\|, \quad \forall x, x' \in \bar{B}(0, r),$$

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 $g: \bar{B}(0, r/2) \rightarrow \bar{B}(0, r)$ 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

$$\begin{aligned}\|f(x_1) - f(x_2)\| &= \|(x_1 - x_2) + (\varphi(x_1) - \varphi(x_2))\| \geq \|x_1 - x_2\| - \|\varphi(x_1) - \varphi(x_2)\| \\ &\geq \|x_1 - x_2\| - \frac{1}{2}\|x_1 - x_2\| = \frac{1}{2}\|x_1 - x_2\|,\end{aligned}$$

thus $\|y_1 - y_2\| \leq 2\|g(y_1) - g(y_2)\|$ and hence g is Lipschitz continuous on $\bar{B}(0, r/2)$. \square

Theorem 8.3. (*Inverse Function Theorem.*) Let $U \subset \mathbb{R}^n$ be an open set, and let $f: U \rightarrow \mathbb{R}^n$ be a differentiable function. Suppose that $a \in U$ and f is continuously differentiable at a , with Df_a invertible. Then there are open neighbourhoods U and V of a and $b = f(a)$ respectively, such that f restricts to a bijection between U and V . Moreover if $g: V \rightarrow U$ is the inverse of f then g is differentiable with

$$Dg(y) = (Df(g(y)))^{-1}.$$

and hence Dg is continuous at x if and only if Df is continuous at $g(y)$.

Proof. By replacing f by $x \mapsto Df(a)^{-1}(f(x+a) - f(a))$, we may assume that $a = f(a) = 0$ and $Df(a) = I_n$. Then if we let $\varphi(x) = f(x) - x$, Lemma 8.1 shows that we may find an $r > 0$ such that $\|Df - I_n\|_\infty < 1/2$ on $B(0, r)$ and so $\|\varphi(x) - \varphi(y)\| \leq \frac{1}{2}\|x - y\|$. Thus it follows from Proposition 8.2 that there is a Lipschitz continuous function $g: \bar{B}(0, r/2) \rightarrow \bar{B}(0, r)$, with

$$f(g(y)) = y, \quad \|g(y_1) - g(y_2)\| \leq 2\|y_1 - y_2\|, \quad \forall y, y_1, y_2 \in \bar{B}(0, r/2).$$

Set $U = f^{-1}(B(0, r/2)) \cap B(0, r)$ which is open since f is continuous, and let $V = f(U)$. Then we have $V = g^{-1}(U)$ so that, as g is continuous, V is an open subset of $B(0, r/2)$, and f restricts to a homeomorphism between U and V . It remains to understand where g is differentiable.

Fix $y_0 \in V$ and let $x_0 = g(y_0)$, and $T = Df_{x_0}$. We wish to show that g is differentiable at y_0 if f is differentiable at x_0 . Write $g(y_0 + k) = x_0 + u(k)$, so that $\|u(k_1) - u(k_2)\| \leq 2\|k_1 - k_2\|$. If $T = Df_{x_0}$ then by definition we have

$$y_0 + k = f(x_0 + u(k)) = f(x_0) + T(u(k)) + \|u(k)\| \cdot \epsilon(x_0 + u(k)),$$

where $\epsilon(x) \rightarrow 0 = \epsilon(x_0)$. and hence $k = T(u(k)) + \|u(k)\| \cdot \epsilon(u(k))$. Now since $\|T - I_n\|_\infty < 1/2$, by the Q.6 on the first Problem Sheet, T is invertible, and hence we have

$$u(k) = T^{-1}(k) - \|u(k)\| \cdot T^{-1}(\epsilon(x_0 + u(k))).$$

Thus g is differentiable at y_0 with derivative T^{-1} provided

$$\eta(k) = -\|k\|^{-1}\|u(k)\| \cdot T^{-1}(\epsilon(u(k))) \rightarrow 0 \text{ as } k \rightarrow 0.$$

But since u is Lipschitz continuous, $u(k)/\|k\|$ is bounded – in fact $\|k\|^{-1}\|u(k)\| \leq 2$ – and as u, ϵ are continuous at $k = 0$, $x = x_0$ respectively, and T^{-1} is bounded, $\|\eta(k)\| \leq 2\|T^{-1}(\epsilon(u(k)))\| \rightarrow 0$ as $k \rightarrow 0$ as required.

Finally, note that we have shown that $Dg_{y_0} = (Df_{g(y_0)})^{-1} = \iota \circ Df \circ g$, where ι denotes the inversion map, which is continuous. It follows that Dg is continuous wherever Df is. \square

Remark 8.4. Notice also that the main difficulty in the proof is to show that f is locally a homeomorphism – once the Lipschitz continuity of the inverse function g is known, it is straight-forward to see where it must be differentiable. In fact it is possible to formulate an inverse function theorem for Lipschitz continuous functions (in the spirit of Proposition 8.2) and deduce the differentiable version from it. It is worth noticing that the continuity of the inverse mapping was also the difficult step in the case of a single variable: In Prelims Analysis you first prove a “continuous inverse function theorem” and then deduce the differentiable inverse function theorem, but the latter theorem is much easier given the continuous case.

Remark 8.5. If, instead of assuming that $f: U \rightarrow \mathbb{R}^n$ is differentiable on U with Df continuous at $a = 0$, we assume only that it is strongly differentiable at a (see Remark 4.16), then one can check that the hypotheses of Proposition 8.2 holds on $B(0, r)$ for small enough r . The arguments above then show that f is locally a homeomorphism at 0, and that its inverse g is (strongly) differentiable at y if f is (strongly) differentiable at $x = g(y)$.

8.2 An alternative definition of a submanifold

For completeness, we include here a proof of the equivalence of the definition of a submanifold given in Remark 6.5 with that given in Definition 6.2.

Proposition 8.6. *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 every $a \in M$, there exists*

- *an open neighbourhood U 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 is a k -dimensional submanifold of V .

Proof. It 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 $\phi(q) = p$. Let $V_1 = \text{im}(Df_q)$, and pick a complementary subspace V_2 of V_1 in V , so that $V = V_1 \oplus V_2$. Let $i_2: V_2 \rightarrow V$ 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, it is easy to see that $D\varphi_{q,0} = D\phi_q + i_2$, 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 $D\theta_{2,p} = \pi_2$, where $\pi_2: V \rightarrow V_2$ 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. \square

8.3 Multilinear maps and higher derivatives

In this section we describe how one can understand the higher derivatives of a function $f: U \rightarrow W$ 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 4.28 shows how the space $\mathcal{L}(V, \mathcal{L}(V, \mathbb{R}))$ is equivalent to the space $\text{Bil}(V, \mathbb{R})$ of bilinear forms on V , that is functions $B: V \times V \rightarrow \mathbb{R}$ 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 8.7. *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 $\theta: F(X, F(Y, Z)) \rightarrow F(X \times Y, Z)$ 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 $\xi: F(X \times Y, Z) \rightarrow F(X, F(Y, Z))$ is given by $\xi(g)(x) = [y \mapsto g(x, y)]$, for all $x \in X, y \in Y$. \square

Write $V^k = V \times \dots \times V$ 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 :

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

Example 8.8. 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 $\text{Bil}(V, \mathbb{R})$ 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 8.9. *Let V and W be finite dimensional normed vector spaces. For each $k \geq 1$ there is a natural isomorphism $\theta_k: \mathcal{L}(V, \mathcal{M}^{k-1}(V, W)) \rightarrow \mathcal{M}^k(V, W)$, and hence if $f: U \rightarrow W$ 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 8.7, 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 . \square

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 4.28, it is possible to check that, if $\{w_1, \dots, w_m\}$ 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 $\{e_1, \dots, e_n\}$ is as before the basis of V , then

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

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

Proposition 8.10. *Let V, W be normed vector spaces, let U be an open subset of V , and let $f: U \rightarrow W$. Then $f \in C^k(V, W)$ if and only if the higher total derivative*

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

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

8.3.1 Symmetries of higher derivatives

The continuity hypothesis imposed on the partial derivatives in Theorem 4.30 is of course reminiscent of the continuity condition we imposed on partial derivatives in order to ensure that the total derivative exists. Indeed, if, for all $i, j \in \{1, 2, \dots, n\}$, $\partial_{ij}f$ is continuous at a , then by Proposition 4.6, D^2f_a exists (and is continuous at a). But the symmetry of the mixed partial derivatives is just the statement that the Hessian matrix of D^2f is symmetric which implies that D^2f_a is a symmetric bilinear form, thus the symmetry of mixed partial derivatives can be reinterpreted in a coordinate-free way, namely that $D^2f_a(v_1, v_2) = D^2f_a(v_2, v_1)$ for all $v_1, v_2 \in V$. Expressed this way, we can slightly weaken the hypothesis of Theorem 4.30 to the assumption that D^2f exists at a , as we now shows.

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

Lemma 8.11. *Suppose that $\beta: V \times V \rightarrow \mathbb{R}$ 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 $(\|sv_1\| + \|sv_2\|)^2 = 4s^2 \rightarrow 0$ as $s \rightarrow 0$. 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^2f_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 8.12. *Let U be an open subset of a normed vector space V . If $f: U \rightarrow \mathbb{R}$ is twice differentiable at $a \in U$, then viewing D^2f as a bilinear form on V we have $D^2f_a(v_1, v_2) = D^2f_a(v_2, v_1)$.*

Proof. Note that, in order for D^2f 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 $A: B \times B \rightarrow \mathbb{R}$ 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 4.25 to J to obtain

$$\|J_1(k)\| \leq \|k\| \cdot \sup_{0 \leq t \leq 1} \|Df_{a+h+tk} - Df_{a+tk} - i_h(D^2 f_a)\|_\infty \quad (8.1)$$

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 $\epsilon_1(x) \rightarrow 0$ as $x \rightarrow 0$. 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 $\epsilon_2(h, k) \rightarrow 0$ as $(h, k) \rightarrow 0$ and

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

Thus returning to the inequality (8.1), we see that

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

But carrying out the same analysis for $J_2(k) = A(k, h) - D^2 f_a(k, h)$ we see that $\|A(k, h) - D^2 f_a(k, h)\| \leq \|h\|(\|h\| + \|k\|) \cdot \epsilon_2(k, h)$, 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:

$$\|\beta(h, k)\| \leq \|D^2 f_a(h, k) - A(h, k)\| + \|A(k, h) - D^2 f_a(k, h)\| \leq (\|h\| + \|k\|)^2 \epsilon_2(h, k). \quad (8.2)$$

But now Lemma 8.2 shows that $\beta = 0$ and hence $D^2 f_a$ is symmetric as required. \square

Remark 8.13. 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.

8.4 Finite dimensional normed vector spaces

Here we give an alternative approach to the fact that all norms on a finite dimensional vector space are equivalent.

8.4.1 Subspaces and quotient spaces

If $(V, \|\cdot\|)$ is a normed vector space, then any linear subspace U clearly inherits the structure of a normed vector space: the norm $\|\cdot\|$ restricts to a norm on U . A somewhat more delicate question is whether the quotient vector space V/U inherits a norm. The first question is to decide what the natural notion of a norm on V/U 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

$$x + U \mapsto \inf \{\|x + u\| : u \in U\}.$$

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

Lemma 8.14. *Let V be a normed vector space and let U be a closed subspace, that is, a linear subspace which is also a closed subset of V . Then the quotient vector space V/U inherits a norm:*

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

Moreover, the quotient map $q: V \rightarrow V/U$ is bounded, with $\|q\|_\infty \leq 1$.

Proof. The requirement that U be a closed linear subspace is what ensures the positivity of the norm $\|\cdot\|$ on V/U . Indeed since the norm on V is non-negative, certainly $\|x + U\| \geq 0$, but suppose now that $\|x + U\| = 0$. Then for any $\epsilon > 0$, there is some $u \in U$ with $\|x + u\| < \epsilon$. But then $-u \in B(x, \epsilon)$, and since $-u \in U$, it follows that x is a limit point of U (that is, in the notation of the metric spaces course, $x \in L(U)$). Since we are assuming U is closed, it follows $x \in U$, and hence $x + U = 0 + U$, so that $\|\cdot\|$ satisfies the positivity condition for a norm.

The homogeneity of the function $\|\cdot\|$ on V/U is straight-forward, so we will only check the triangle inequality here. Let $x + U, y + U \in V/U$. Then by the approximation property, for any $\epsilon > 0$, we may find $u_1, u_2 \in U$ such that $\|x + U\| \leq \|x + u_1\| < \|x + U\| + \epsilon$, and $\|y + U\| \leq \|y + u_2\| < \|y + U\| + \epsilon$. But then since $u_1 + u_2 \in U$, by definition we have

$$\|(x + y) + U\| \leq \|(x + y) + (u_1 + u_2)\| = \|(x + u_1) + (y + u_2)\|.$$

Thus, using the triangle inequality for the norm on V and our choice of u_1 and u_2 we have

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

Combining these two inequalities, it follows that $\|(x + y) + U\| < \|x + U\| + \|y + U\| + 2\epsilon$ for any $\epsilon > 0$, and hence $\|(x + y) + U\| \leq \|x + U\| + \|y + U\|$, as required. \square

Proposition 8.15. *Let V and W be normed vector spaces and suppose that $\dim(V) < \infty$. Then any linear map $\alpha: V \rightarrow W$ is automatically bounded, that is $\mathcal{B}(V, W) = \mathcal{L}(V, W)$.*

Proof. We prove this statement by induction on $n = \dim(V)$. First suppose $\dim(V) = 1$, and let $\alpha: V \rightarrow W$ be a linear map. Pick $e \in V$ a unit vector (so $\|e\| = 1$), so that if $v \in V$ is arbitrary, $v = \pm\|v\|e$ and hence $\|\alpha(v)\| = \|\alpha(e)\| \cdot \|v\|$, so that $\|\alpha\|_\infty = \|\alpha(e)\|$, that is, α is bounded as required.

Now suppose that $n = \dim(V) > 1$, and that 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 $\{u_1, \dots, u_k\}$ of U defines a linear isomorphism $\phi: \mathbb{R}^k \rightarrow U$ where if $x = (x_1, \dots, x_k) \in \mathbb{R}^k$ 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 $\|\cdot\|_2$ norm) is complete, so is U . It follows that U must therefore be closed in V .

Now suppose $\alpha: V \rightarrow W$. Since $\dim(\alpha(V)) \leq n$, we may pick a finite basis $\{w_1, \dots, w_m\}$ of the image of α in W . Let $\alpha_i: V \rightarrow \mathbb{R}$ be the components of α with respect to this basis, that is, the functionals α_i are defined by the equation $\alpha(v) = \sum_{i=1}^m \alpha_i(v) \cdot w_i$. Next note that α is bounded if all of the α_i are, since then

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

It thus suffices to show that any linear functional $\delta: V \rightarrow \mathbb{R}$ is bounded. This is clear if $\delta = 0$, so suppose $\delta \neq 0$. Then $H = \ker(\delta)$ is an $(n-1)$ -dimensional subspace of V , and hence, as noted above, it is closed. But then the quotient space V/H is a normed vector space and the quotient map $q: V \rightarrow V/H$ is bounded (with norm at most 1). But the functional δ can be written as the composition $\delta = \bar{\delta} \circ q$, where $\bar{\delta}: V/H \rightarrow \mathbb{R}$ is the injective linear map induced by δ on V/H . But since $\dim(V/H) = 1$ we know $\bar{\delta}$ is bounded, and hence its composition with q is also bounded, and so δ is bounded as required. \square

Remark 8.16. This theorem 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.