MANIFOLDS MA3H5: COURSE NOTES, 2019

BRIAN H. BOWDITCH

0. Introduction

These are informal notes to accompany the course Manifolds MA3H5. They run parallel to the course, but are not necessarily identical. Where there is any difference, it is the material that is presented in lectures which will determine the examinable content of the course.

Sources :

In preparing the course, I have made use of the lecture notes and exercise sheets by David Mond, mentioned below, as well as hand-written notes by G. R. Allan (University of Cambridge).

Other notes :

The "Lecture Notes for MA455 Manifolds" by David Mond, Warwick, 2008, are available here:
http://homepages.warwick.ac.uk/∼masbm/Manifolds/manifolds.ps

They date from a time when "Manifolds" was a 4th-year module, but there is quite a bit of overlap with our course. Most of what we do will be covered by these notes (though maybe somewhat differently), and they contain lots more examples and pictures. We won't however be doing much on "transversality", which was a substantial part of the fourth-year course. Conversely, we plan to say a bit about vector bundles, which were not covered in the earlier course.

Books :

L. W. Tu, "An Introduction to Manifolds", Universitext Springer-Verlag (2010). QA613.T8.

J. M. Lee, "Introduction to Smooth Manifolds", Graduate Texts in Mathematics, Springer (2013). QA613.L3.

Date: September 2019.

F. Warner, "Foundations of differentiable manifolds and Lie groups", Graduate Texts in Mathematics, Springer (2010). QA614.3.W2.

W. Boothby, "An introduction to differentiable manifolds and Riemannian geometry", Academic Press (2003). QA614.3.B6.

Web pages :

The course has an official web page:
http://www2.warwick.ac.uk/fac/sci/maths/undergrad/ughandbook/year3/ma3h5/
and an unofficial one:
http://homepages.warwick.ac.uk/∼masgak/manifolds/course.html
(which can be found via the link on my homepage).

Summary :

This course is about the theory of smooth manifolds, that is those equipped with a differentiable structure. We will begin by discussing manifolds in euclidean space, $\mathbb{R}^n$, and will move on to an account of smooth manifolds in a more abstract setting. We aim to cover topics such as differential forms and integration on manifolds. We will briefly mention riemannian manifolds and Lie groups, which form a major component of certain fourth-year courses.

Content :

1. Background material.
2. What is a manifold?
3. Definition of a manifold in euclidean space.
4. Immersions and submersions.
5. Tangents, normals, orientations.
6. Abstract manifolds.
7. Vector bundles.
8. Extending smooth functions.
9. Manifolds with boundary.
10. Differential forms and integration.
11. Exterior derivatives and Stokes's Theorem.

1. Background material

In this section, we summarise a few facts we will be using. These should have been covered in the prerequisite 2nd year units.

Linear algebra

Unless otherwise stated, all vector spaces will be assumed real (defined over $\mathbb{R}$) and finite-dimensional.

We assume basic notions: bases, dimension, subspaces, linear maps, multilinear maps etc.

Given vector spaces, E, F, we write $E \leq F$ to mean that E is a subspace of F. Note that $\dim E \leq \dim F$.

Given two vector spaces, E, F, the direct product $E \times F$ is also a vector space, with addition and scalar multiplication defined pointwise. In fact, E and F are often respectively identified with the subspaces $E \times \{0\}$ and $\{0\} \times F$, of $E \times F$. Note that $\dim(E \times F) = \dim E + \dim F$.

We say that a vector space V is a *direct sum* of subspaces, E and F, if each element of V can be written uniquely in the form $x + y$ with $x \in E$ and $y \in F$. (That is, E, F are *complementary* subspaces of V.) In this case, one often writes $V = E \oplus F$. Note that the map $[(x, y) \mapsto x + y] : E \times F \longrightarrow E \oplus F$ is an isomorphism.

In the context of vector spaces, the terms "direct product" and "direct sum" are often regarded as synonymous. Here we will generally use the former notation and terminology for the domain of bilinear maps. We use the latter notation when we want to think a vectors space together with two complementary subspaces.

We write $L(E, F)$ for the vector space of all linear maps from E to F. Choosing bases for E and F, we can identify $L(E, F)$ with the space of $m \times n$ matrices, where $m = \dim E$, $n = \dim F$. Note that $\dim L(E, F) = mn$.

Write $E^* = L(E, \mathbb{R})$ for the dual space to E. Given $x \in E$ and $f \in E^*$ write $\langle x, f \rangle = f(x)$. (This is a bilinear form on $E \times E^*$.) Given a basis $e_1, \ldots, e_m$, for E, there is a *dual basis* $f_1, \ldots, f_n$ such that $\langle e_i, f_j \rangle = \delta_{ij}$, where $\delta_{ij} = 1$ if $i = j$ and 0 if $i \neq j$. Note that there is a natural injective linear map, α, from E to E^{**} given by $\langle f, \alpha(x) \rangle = \langle x, f \rangle$ for all $f \in E^*$. Since $\dim E^{**} = \dim E^* = \dim E$, it follows that α is bijective, so we can naturally identify E with E^{**} (hence the term "dual"). (Note this is *not* true in general for infinite dimensional vector spaces.) Of course, E and E^* are also isomorphic as vector spaces, since they have the same dimension. However, there is no natural isomorphism — it involves an arbitrary choice. Whereas, the isomorphism between E and E^{**} is canonical: the map described above did not involve any arbitrary choices. Elements of the dual space are sometimes referred to as *linear functionals*.

If E and F are vector spaces, then a multilinear map $T : E^p \longrightarrow F$ is *symmetric* if for all $x_i \in E_i$, $T(x_{\pi(1)}, \ldots, x_{\pi(p)}) = T(x_1, \ldots, x_p)$ for any permutation π of the arguments. It is *alternating* if $p \geq 1$ and $T(x_{\pi(1)}, \ldots, x_{\pi(p)}) = \mathrm{sig}(\pi) T(x_1, \ldots, x_p)$ where $\mathrm{sig}(\pi) \in \{-1, 1\}$ is the signature of π. If $F = \mathbb{R}$, we refer to symmetric and alternating *forms*.

Note that the dot product on $E = \mathbb{R}^n$ is an example of symmetric bilinear form. Given $V \subseteq \mathbb{R}^n$, we write $V^\perp = \{y \in \mathbb{R}^n \mid (\forall x \in V)(x.y = 0)\}$ for the orthogonal subspace. Note that $V^{\perp\perp} = V$ and that $\mathbb{R}^n \cong V \oplus V^\perp$.

More generally, an *inner product* on V is a symmetric bilinear form, $[(x, y) \mapsto \langle x, y\rangle]$ such that $\langle x, x\rangle > 0$ for all $x \neq 0$.

We will have more to say about alternating forms in Section 9.

By the *standard basis* of $\mathbb{R}^n$, we mean $\{e_1, e_2, \dots, e_n\}$, where $e_1 = (1, 0, \dots, 0)$, $e_2 = (0, 1, \dots, 0)$ etc.

On the few occasions we talk about infinite dimensional vector spaces, they will not be assumed to have any additional structure. (In particular, they are not considered to be topological vector spaces, which you may have encountered elsewhere).

Metric and topological spaces.

We will assume basic facts about metric and topological spaces: open and closed sets, connectedness, open covers, compactness, local compactness, continuous maps, homeomorphisms, proper maps, second countable spaces.

We recall that a continuous proper injective map from a compact space to hausdorff space is a homeomorphism onto its range. From this, one can see also that a continuous proper map from a locally compact space to a hausdorff space is a homeomorphism onto its range.

Given a metric space (X, d), $x \in X$ and $r \geq 0$, we will write $N(x; r) = \{y \in X \mid d(x, y) < r\}$ and $B(x; r) = \{y \in X \mid d(x, y) \leq r\}$ respectively for the open and closed r-balls about x.

Differentiation

We say that a map $f : \mathbb{R}^m \longrightarrow \mathbb{R}^n$ is *smooth* if it has derivatives of all orders.

If $x \in \mathbb{R}^m$, the derivative of f at x is a linear map $d_x f : \mathbb{R}^m \longrightarrow \mathbb{R}^n$. It is represented by the $m \times n$ "jacobian" matrix $(\partial f_i/\partial x_j)_{ij}$ of partial derivatives. We will denote the jacobian by $J_f(x)$.

Theorem 1.1. *If $d_x f$ has rank m, then f is locally injective at p. (That is, there is an open set, $U \subseteq \mathbb{R}^n$ with $p \in U$ such that $f|U$ is injective.)*

(Clearly, the hypotheses imply that $m \leq n$.)

Theorem 1.2. *(Inverse function theorem) If $d_x f : \mathbb{R}^n \longrightarrow \mathbb{R}^n$ has rank n (i.e. is invertible) then there are open subsets $U, V \subseteq \mathbb{R}^n$ with $p \in U$ such that $f|U$ is bijective, and $(f|U)^{-1} : V \longrightarrow U$ is also smooth.*

A *diffeomorphism* between two open sets, $U, V \subseteq \mathbb{R}^n$ is a smooth bijection $f : U \longrightarrow V$, whose inverse, $f^{-1} : V \longrightarrow U$ is also smooth.

A proof of Theorem 1.2 can be found in MA225. With some work, Theorem 1.1 can be deduced from Theorem 1.2.

We will also need some basic facts about integration in $\mathbb{R}^m$. Let $f : \mathbb{R}^m \longrightarrow \mathbb{R}$ be a smooth function. (In fact, any continuous function would do for the present discussion.) The *support* of f, denoted $\operatorname{supp} f$, is the closure of $\{x \in \mathbb{R}^m \mid f(x) \neq 0\}$. We say that f is *compactly supported* if $\operatorname{supp} f$ is compact. In this case, we can define the integral $\int_{\mathbb{R}^m} f(x)\,dx$. If $\phi : U \longrightarrow V$ is an orientation-preserving diffeomorphism between open sets, $U, V \subseteq \mathbb{R}^m$, with $\operatorname{supp} f \subseteq V$, then $f \circ \phi$ is also compactly supported, and $\int_{\mathbb{R}^m} f \circ \phi(x)\,dx = \int_{\mathbb{R}^m} f(x) J_\phi(x)\,dx$.

2. What is a manifold?

We first discuss informally what we mean by a manifold, and give some examples of things which are and are not manifolds. We will give formal definitions in Section 3. One can then retrospectively verify the statements made here.

Basically an "m-manifold" (or "manifold of dimension m") is something which looks locally like m-dimensional space $\mathbb{R}^m$. That is, one has a local system of real coordinates — one for each of the m dimensions. This course will be almost entirely about "smooth manifolds", where they come with some differentiable structure. (Though one can also talk more generally about "topological manifolds", where there is no such structure a-priori.)

For the first few Sections (3 to 5) we will deal only with manifolds which are subsets of $\mathbb{R}^n$ for some n, though we later deal with manifolds defined in a more abstract way (Section 6 onwards).

Here are some examples, and non-examples, which will motivate the formal definition we give in Section 3:

(E1) $\mathbb{R}^n$ is an n-manifold. So is any non-empty open subset of $\mathbb{R}^n$. Note that, we do not in general assume that a manifold is connected. It may be debatable whether or not the empty set is a manifold, though we will adopt the convention that manifolds are non-empty.

(E2) We allow $n = 0$, so a point is a 0-manifold. Indeed any non-empty finite subset of any $\mathbb{R}^n$ is a 0-manifold.

(E3) Any m-dimensional (vector) subspace of $\mathbb{R}^n$ is an m-manifold. Note that necessarily $m \leq n$. In fact (we will see) any manifold in $\mathbb{R}^n$ has dimension at most n.

(E4) The unit circle, $S^1 = \{(x_1, x_2) \in \mathbb{R}^2 \mid x_1^2 + x_2^2 = 1\}$ is a 1-manifold. We have nice smooth "local coordinates" $\theta \mapsto (\cos\theta, \sin\theta)$, where $\theta \in \mathbb{R}$, which means that it locally "looks like" $\mathbb{R} = \mathbb{R}^1$.

(E5) The unit sphere, $S^2 = \{(x_1, x_2, x_3) \in \mathbb{R}^2 \mid x_1^2 + x_2^2 + x_3^2 = 1\}$ is a 2-manifold. (A 2-manifold is often called a "surface".)

There are many ways to put local coordinates on S^2, One is *stereographic projection*. Let $p_N = (0, 0, 1)$ and $p_S = (0, 0, -1)$ be the "north" and "south poles". Define $\phi_N : S^2 \setminus \{p_N\} \longrightarrow \mathbb{R}^3$ by the formula:

$$\phi_N(x_1, x_2, x_3) = \left(\frac{x_1}{1 - x_3}, \frac{x_2}{1 - x_3} \right).$$

This is stereographic projection from the north pole to the equatorial plane, i.e. $\mathbb{R}^2$ is identified with $\mathbb{R}^2 \times \{0\} \subseteq \mathbb{R}^3$. Similarly, define $\phi_S : S^2 \setminus \{p_S\} \longrightarrow \mathbb{R}^2$ by:

$$\phi_S(x_1, x_2, x_3) = \left(\frac{x_1}{1 + x_3}, \frac{x_2}{1 + x_3} \right).$$

Together these two maps give local smooth coordinates on the whole of S^2.

There are lots of other ways of doing this. It's what cartography is all about, and you can read about the many different types of projections used in atlases.

(E6) The torus $T_{a,b}$ in $\mathbb{R}^3$ is a 2-manifold. Let $a > b > 0$. Let $C_a = \{(x_1, x_2, 0) \in \mathbb{R}^3 \mid x_1^2 + x_2^2 = a\}$. (This is a 1-manifold in $\mathbb{R}^3$.) Let $T_{a,b}$ be the set of points euclidean distance b, from C_a. In other words, $T_{a,b}$ can be defined as the set of $(x_1, x_2, x_3) \in \mathbb{R}^3$ satisfying:

$$\left(x_1 - \frac{a x_1}{\sqrt{x_1^2 + x_2^2}} \right)^2 + \left(x_2 - \frac{a x_2}{\sqrt{x_1^2 + x_2^2}} \right)^2 + x_3^2 = b^2.$$

Note that $T_{a,b}$ has local coordinates $(\theta, \phi) \mapsto ((a + b\cos\phi)\cos\theta, (a + b\cos\phi)\sin\theta, b\sin\phi)$. As such it is a 2-manifold.

Note that we can alter a and b (subject to $a > b > 0$). We get different tori in $\mathbb{R}^2$. But they all look similar intrinsically. It's not hard to see that they are all homeomorphic (exercise). In fact, we will see that they all have essentially the same smooth structure (i.e. they are "diffeomorphic"). This will be explained in Section 3.

Of course, these homeomorphisms do not respect distance. The metric is not intrinsically part of the structure of a manifold. It is an additional structure, which may, or may not, be relevant, depending on what one wants to use manifolds for.

(E7) Let $T = \{(x_1, x_2, x_3, x_4) \in \mathbb{R}^4 \mid x_1^2 + x_2^2 = x_3^2 + x_4^2 = 1\}$. This is a 2-manifold. In fact, it is also a "torus": it is "diffeomorphic" to $T_{a,b}$ in example (E6). We can think of it as a direct product of two circles, $S^1 \times S^1 \subseteq \mathbb{R}^2 \times \mathbb{R}^2 \equiv \mathbb{R}^4$ from example (E3). (In general, direct products of manifolds are manifolds.) Although we can embed the torus in $\mathbb{R}^3$ (that is, describe it as a subset of $\mathbb{R}^3$), it seems to be happier in $\mathbb{R}^4$.

(E8) Given $n \in \mathbb{Z}$, let $B_n = \{(\cos\theta, \sin\theta, r\cos(n\theta/2), r\sin(n\theta/2)) \mid \theta, r \in \mathbb{R}\}$. Then $B_n \subseteq \mathbb{R}^4$ is a 2-manifold. Note that the projection $[(x_1, x_2, x_3, x_4) \mapsto (x_1, x_2)]$ maps B_n onto a circle in $\mathbb{R}^2$. The preimage of any point in B_n is a straight line in $\mathbb{R}^4$. As we go once around this circle, this line spins around $n/2$ times. So, for example, B_0 is a cylinder, and B_1 is "Möbius band".

(E9) Now for some non-examples. The union of the coordinate axes $X = \{(x_1, x_2) \in \mathbb{R}^2 \mid x_1 = 0 \text{ or } x_2 = 0\}$ is *not* a manifold (of any dimension). Things go wrong at the origin $(0, 0)$. Even though it is in some sense "1-dimensional" it does not have local coordinates there.

However, $X \setminus \{(0, 0)\}$ is a (disconnected) 1-manifold.

(E10) Similarly, the cone $C = \{(x_1, x_2, x_3) \in \mathbb{R}^3 \mid x_1^2 + x_2^2 - x_3^2 = 0\}$ is not a manifold: it again goes awry at the origin.

However, $C \setminus \{(0, 0, 0)\}$ is a 2-manifold.

(E11) Of course, we don't really have to search so hard for non-examples. If you pick some "random" subset of $\mathbb{R}^n$, then it is unlikely to be a manifold.

(E12) What about the square: $Q = (\{-1, 1\} \times [-1, 1]) \cup ([-1, 1] \times \{-1, 1\}) \subseteq \mathbb{R}^2$?

This is homeomorphic to S^1: by radial projection from the origin (exercise).

It is a topological 1-manifold, but not a smooth manifold (so not a "manifold" for us). One cannot put smooth local coordinates near the corners.

(E13) Things get worse, or more interesting: the Koch snowflake curve in $\mathbb{R}^2$ is also homeomorphic to S^1, hence also a topological 1-manifold, but does not have smooth coordinates anywhere.

(E14) Alexander's horned sphere in $\mathbb{R}^3$: homeomorphic to S^2, but not smooth.

3. Definition of a manifold in euclidean space

Our aim in this section is to give a formal definition of a manifold, M, as a subset of $\mathbb{R}^n$. In this context, $\mathbb{R}^n$ is referred to as the "ambient space". It is worth taking note of when the ambient space features in the discussions of the next few sections (3 to 5) — many of the arguments make little or no reference to it. For many purposes, it can be dispensed with altogether, as we will see in Section 6.

We begin by generalising the notions of smooth maps and diffeomorphisms to arbitrary subsets of $\mathbb{R}^n$. Note that these are already defined for open subsets.

Definition. Let $X \subseteq \mathbb{R}^m$. A map $f : X \longrightarrow \mathbb{R}^n$ is *smooth* if, for all $x \in X$, there is an open set, $U \ni x$, in $\mathbb{R}^m$ and a smooth map (in the original sense), $F : U \longrightarrow \mathbb{R}^n$ such that $F|(U \cap X) = f|(U \cap X)$. We also speak of f as a smooth map from X to $f(X) \subseteq \mathbb{R}^n$. If $X \subseteq \mathbb{R}^m$ and $Y \subseteq \mathbb{R}^n$, then a *diffeomorphism* between X and Y is a smooth map from X to Y which has a smooth inverse.

It is easily checked that, if X is open, then this notion coincides with the original, so there will be no ambiguity. The following is also an easy exercise (given that it holds in the original sense).

Lemma 3.1. *The composition of smooth maps is smooth.*

Clearly smooth maps are continuous, and diffeomorphims are also homeomorphisms.

Examples.

(1) In the case of stereographic projection of S^2 described in Example (E5) in Section 2, the maps ϕ_N and ϕ_S are both smooth: one can use the same formulae to extend them to smooth maps defined on $\mathbb{R}^3 \setminus \{p_N\}$ and $\mathbb{R}^3 \setminus \{p_S\}$ respectively. In fact, they are both diffeomorphisms to $\mathbb{R}^2$: one can check the inverse maps are smooth by writing down explicit formulae for them (exercise).

(2) There is a diffeomorphism from the torus $T \subseteq \mathbb{R}^4$ defined by Example (E7), and $T_{a,b} \subseteq \mathbb{R}^3$ as defined in Example (E6). One can check that the map $(x_1, x_2, x_3, x_4) \mapsto ((a + bx_3)x_1, (a + bx_3)x_2, bx_4)$ from T to $T_{a,b}$ is a diffeomorphism: it extends to a smooth map defined on $\mathbb{R}^4$ by the same formula. There is also a smooth inverse which can be extended to an open set contain $T_{a,b}$ (exercise: write down a formula for this).

Definition. Let $M \subseteq \mathbb{R}^n$. A *smooth chart* on M consists a subset $U \subseteq M$, open in M, an open set, $V \subseteq \mathbb{R}^m$ and a diffeomorphism $\phi : U \longrightarrow V$.
A *smooth atlas* for M is a collection of charts whose union covers M.
We say that M is a *smooth manifold of dimension* m (or a *smooth* m*-manifold*) if it admits a smooth atlas with charts mapping to $\mathbb{R}^m$. (In other words we can find a smooth chart defined on a neighbourhood of any point of M.)

Note that we don't in general assume U, V to be connected (though they will be in the examples described below).

It is easily seen that $m \leq n$. (See also the discussion of tangent spaces below: Lemma 3.3). Also (exercise) the dimension, m, is determined by M.

Remark: The notation M^m is sometimes used to denote an m-manifold — to remind us that it has dimension m. Of course, this is *not* the direct product

$M \times \cdots \times M$. For example, we have already used S^2 to denote the 2-sphere, which is quite standard notation.

One can go through the examples described in Section 2, and verify that they are indeed smooth manifolds.

For example:

(1) In Example (E5), S^2 is a manifold: there is an atlas consisting of just two charts: $\phi_N : S^2 \setminus \{p_N\} \longrightarrow \mathbb{R}^2$ and $\phi_S : S^2 \setminus \{p_S\} \longrightarrow \mathbb{R}^2$.

(2) Indeed one can define stereographic projection of the n-sphere, $S^n = \{x \in \mathbb{R}^{n+1} \mid ||x|| = 1\}$ by the same geometric construction. It is given by similar formulae. Again we have two charts from the complements of the north and south poles (the points with final cordinate equal to 1 and -1). This shows that S^n is an n-manifold. This gives one way of seeing that $S^1 \subseteq \mathbb{R}^2$ is indeed a 1-manifold.

(3) If $M \subseteq \mathbb{R}^p$ and $N \subseteq \mathbb{R}^q$ are respectively an m-manifold and an n-manifold, then $M \times N \subseteq \mathbb{R}^{p+q}$ is an $(m + n)$-manifold (exercise). In particular, $T \subseteq \mathbb{R}^4$ (Example (E7)) is a 2-manifold.

(4) If $U \subseteq \mathbb{R}^m$ is open and $f : U \longrightarrow \mathbb{R}^n$ is smooth, then the graph of f, that is $M = \{(x, f(x)) \mid x \in U\} \subseteq \mathbb{R}^{m+n}$ is an m-manifold. There is an atlas consisting a single chart which is projection to U.

(5) We can identify the set, $\mathcal{M}(n, \mathbb{R})$, of all $n \times n$ real matrices with $\mathbb{R}^{n^2}$, just taking the entries as coordinates. The set, $GL(n, \mathbb{R})$, be set of invertible $n \times n$ matrices is then an open subset of $\mathbb{R}^{n^2}$ (since the determinant map is continuous). Note that $GL(n, \mathbb{R})$ is a group under multiplication of matrices. Moreover, the inversion map $[A \mapsto A^{-1}] : GL(n, \mathbb{R}) \longrightarrow GL(n, \mathbb{R})$ and the product map $[(A, B) \mapsto AB] : GL(n, \mathbb{R}) \times GL(n, \mathbb{R}) \longrightarrow GL(n, \mathbb{R})$ are both smooth. With this structure $GL(n, \mathbb{R})$ is an example of a "Lie group", about which we will say more later. A similar example is the group of upper triangular matrices with 1s on the diagonal (an open subset of $\mathbb{R}^{(n^2-n)/2}$). Another, is the group of invertible complex matrices, $GL(n, \mathbb{C})$, identified as an open subset of $\mathbb{C}^{n^2} \equiv \mathbb{R}^{2n^2}$. We will see more later.

Exercise : If $M \subseteq \mathbb{R}^n$ is an m-manifold, and $N \subseteq M$ is open in M, then N is also an m-manifold.

Definition. We say two manifolds are *diffeomorphic* if there is a diffeomorphism between them.

One sees easily that this is an equivalence relation on the class of manifolds.

Examples :

(1) For example, the torus $T \subseteq \mathbb{R}^4$ and $T_{a,b} \subseteq \mathbb{R}^3$ are diffeomorphic, as we stated informally in Section 2.

(2) In example (E9) of Section 2, the manifolds B_n and B_m are diffeomorphic if $m - n$ is even. (The converse is also true, as we will see later.)

Exercise :
Note that the definition of *diffeomorphic* could also be applied to any subsets of euclidean space (though it is rarely applied in this generality). Suppose that $M \subseteq \mathbb{R}^p$ and $N \subseteq \mathbb{R}^q$ are diffeomorphic, then M is an m-manifold if and only if N is an m-manifold.
Together with earlier remarks, this implies that the torus, $T_{a,b} \subseteq \mathbb{R}^3$, is a 2-manifold in $\mathbb{R}^3$. Exercise: give an explicit description of an atlas.

Remark : We can define the notion of a topological manifold similarly by using a "topological atlas" consisting of "topological charts". A "topological chart" can be defined by replacing the word "diffeomorphism" by "homeomorphism", so that there is no reference at all to differentiation.

We saw some examples of topological manifolds in Section 2 ((E12)–(E14)). In fact, each of those topological manifolds was homeomorphic to a smooth manifold (circle or sphere in those cases). It turns out that if $n \leq 3$, then any topological n-manifold is homeomorphic to a smooth manifold, but this is not in general true if $n \geq 4$. Furthermore, if two smooth manifolds of dimension $n \leq 3$ are homeomorphic, they are also diffeomorphic. Again, this fails in higher dimensions. We will say a bit more about this in Section 6.

In the meantime, all our manifolds will be assumed smooth, and so we generally omit the adjective "smooth", as this is implicitly understood.)

Suppose that $\phi_\alpha : U_\alpha \longrightarrow V_\alpha$ and $\phi_\beta : U_\beta \longrightarrow V_\beta$ are charts. (We will often use α, β as indices for charts in an atlas.) We get a map

$$\phi_\beta \circ \phi_\alpha^{-1} : \phi_\alpha(U_\alpha \cap U_\beta) \longrightarrow \phi_\beta(U_\alpha \cap U_\beta).$$

This is necessarily a diffeomorphism. A map of this sort is called a *transition map.* (Note that $U_\alpha \cap U_\beta$ might be empty, but the empty map is a diffeomorphism, so we are happy.)

As an example, consider again S^2 with the stereographic projections. In this case, we have a transition map $\mathbb{R}^2 \setminus \{(0,0)\} \longrightarrow \mathbb{R}^2 \setminus \{(0,0)\}$ which one can check is inversion in the unit circle: that is, it is given by the formula:

$$(x_1, x_2) \mapsto \left(\frac{x_1}{x_1^2 + x_2^2}, \frac{x_2}{x_1^2 + x_2^2} \right)$$

which is clearly a diffeomorphism. (Here, the transition map is the same as its inverse, though that is special to this case.)

Recall that if $f : U \longrightarrow \mathbb{R}^n$ is a smooth map defined on an open set $U \subseteq \mathbb{R}^m$, and $x \in U$, we have the derivative $d_x f : \mathbb{R}^m \longrightarrow \mathbb{R}^n$, which is a linear map. Under composition, this obeys the "chain rule": $d_x(g \circ f) = d_{f(x)} g \circ d_x f$.

If $I \subseteq \mathbb{R}$ is an open interval, we define a smooth path to be a map $\gamma : I \longrightarrow \mathbb{R}^n$. We say that that γ is a *smooth path in* $X \subseteq \mathbb{R}^n$ if $\gamma(I) \subseteq X$. If $t \in I$, we write $\gamma'(t) = (d_t \gamma)(1)$. This is the usual "tangent vector" to γ. In fact, we are now ready to define tangents more generally.

Definition. Let $M \subseteq \mathbb{R}^n$ be an m-manifold, and let $x \in M$. The *tangent space* $T_x M$, of M at x is the space $d_a(\phi^{-1})(\mathbb{R}^m)$, where $\phi : U \longrightarrow \mathbb{R}^m$ is a chart defined on an open neighbourhood, U, of x in M, and $a = \phi(x)$.

Certainly such a chart exists, by hypothesis. Also, since $d_a(\phi^{-1})$ is a linear map, $T_x M$ is a subspace of $\mathbb{R}^m$. However, we need to check that it doesn't depend on the choice of the chart ϕ.

For this, we will use the following equivalent, more intuitive, formulation of tangent space.

Lemma 3.2. *$T_x M$ is the set of vectors, $v \in \mathbb{R}^n$, of the form $v = \gamma'(0)$ for some smooth path $\gamma : I \longrightarrow M$ with $0 \in I$ and with $x = \gamma(0)$.*

Here, $T_x M$ is interpreted, for the moment, as being defined with respect to a given chart, ϕ, we have chosen.

Proof. Suppose $v = \gamma'(0)$ has this form. Restricting the domain, we can assume that $\gamma(I) \subseteq U$. Let $\sigma = \phi \circ \gamma : I \longrightarrow \mathbb{R}^m$. This is curve in $V \subseteq \mathbb{R}^m$. By the chain rule, we get $v = \gamma'(0) = (\phi^{-1} \circ \sigma)'(0) = d_a(\phi^{-1})(\sigma'(0)) \in T_x M$.

Conversely, if $v \in T_x M$. Write $v = d_a(\phi^{-1})(w)$ where $w \in \mathbb{R}^m$. Let $\sigma : I \longrightarrow V$ be a curve with $\sigma(0) = a = \phi(x)$ and with $\sigma'(0) = w$ (e.g. $[t \mapsto a + tw]$). Let $\gamma = \phi^{-1} \circ \sigma : I \longrightarrow U \subseteq M$. Now $v = d_a(\phi^{-1})(w) = d_a(\phi^{-1})(\sigma'(0)) = (\phi^{-1} \circ \sigma)'(0) = \gamma'(0)$. $\square$

Note that this formulation makes no reference to charts, and so we see retrospectively, that it is independent of ϕ. Hence, our original definition did not depend on ϕ.

(Of course, we could instead have defined the tangent space using curves in M. But then, it would not be immediately clear that it has to be a subspace of $\mathbb{R}^n$.)

In fact we have:

Lemma 3.3. $\dim T_x M = m$.

Proof. Let Φ be any smooth extension of ϕ in a neighbourhood of x in $\mathbb{R}^m$. Then $\Phi \circ \phi^{-1}$ is the identity on a neighbourhood of a in $\mathbb{R}^m$. Thus, $d_x\Phi \circ d_a(\phi^{-1})$ is the identity map on $\mathbb{R}^m$. It follows that $d_a(\phi^{-1})$ is injective, so its image has dimension m. $\square$

Note, in particular, that if $M \subseteq \mathbb{R}^n$, then $m \leq n$.

Examples

(1) The tangent space of any vector subspace of $\mathbb{R}^n$ is the space itself. In particular, $T_x(\mathbb{R}^n) = \mathbb{R}^n$ for all $x \in \mathbb{R}^n$.

(2) The tangent space to $x \in S^n \subseteq \mathbb{R}^{n+1}$ is $\{y \in \mathbb{R}^{n+1} \mid x.y = 0\}$.

(3) Suppose that $M = f^{-1}(0)$, where $f : \mathbb{R}^n \longrightarrow \mathbb{R}$ is smooth. Suppose that M is a manifold, and that $x \in M$. Suppose that $d_xf : \mathbb{R}^n \longrightarrow \mathbb{R}$ has rank 1. Then $T_xM = \ker d_xf$.

Suppose now that $f : M \longrightarrow N$ is a smooth map between manifolds, $M \subseteq \mathbb{R}^p$ and $N \subseteq \mathbb{R}^q$. We aim to define a linear map $d_xf : T_xM \longrightarrow T_{f(x)}N$.

One way to do this would be to choose a smooth extension, F, of f in a neighbourhood of x in $\mathbb{R}^p$, and set $d_xf = d_xF|T_xM$. Alternatively, given $v \in T_xM$, choose any smooth curve γ in M with $\gamma(0) = v$, $\gamma'(0) = v$, and then set $d_xf(v) = (f \circ \gamma)'(0)$. Of course, both involve making a choice.

In fact, these two definitions coincide, since $d_xF(\gamma'(0)) = (F \circ \gamma)'(0) = (f \circ \gamma)'(0)$. It follows retrospectively that the definition is independent of the choice of extension F (in the first formulation) and also independent of the choice of γ (in the second formulation). In other words, we have a well defined map.

Definition. The *derivative* of f at x is the linear map, $d_xf : T_xM \longrightarrow T_{f(x)}N$ defined above.

Example.

If $V \subseteq \mathbb{R}^p$ and $W \subseteq \mathbb{R}^q$ are vector subspaces and $L : V \longrightarrow W$ is a linear map, then $d_xL = L$ for all $x \in V$.

The following chain rule is now a simple exercise given the usual chain rule in $\mathbb{R}^n$:

Lemma 3.4. *If M, N, P are smooth manifolds, and $f : M \longrightarrow N$ and $g : N \longrightarrow P$ are smooth maps, then $d_x(g \circ f) = d_{f(x)}g \circ d_xf$.*

Note that we can now view a chart $\phi : U \longrightarrow \mathbb{R}^m$ as a diffeomorphism between manifolds $U \subseteq M$ and $V \subseteq \mathbb{R}^m$. If $x \in U$, then $T_xU = T_xM$ and we get a derivative map $d_x\phi : T_xM \longrightarrow T_x(\mathbb{R}^m) = \mathbb{R}^m$, as we have just defined for manifolds.

This is just the inverse of the map $d_{\phi(x)}(\phi^{-1})$ used in defining T_xM in the first place. In particular, $d_x\phi$ is an isomorphism from T_xM to $\mathbb{R}^m$. The above is just a matter of formally checking the definitions.

Germs.

The following, somewhat formal, discussion will be relevant to defining the tangent space to an abstract manifold is Section 6.

One can think of a tangent vector as giving a means of differentiating smooth functions on M. This can be expressed as follows.

Let $C^\infty(M)$ be the set of all smooth functions, $f : M \longrightarrow \mathbb{R}$. Note that this is naturally an (infinite dimensional) vector space over $\mathbb{R}$ with addition and scalar multiplication defined pointwise in the obvious way.

Given $x \in M$, write $C^\infty_x(M)$ for the set of "local" smooth functions at x, i.e. smooth functions, $f : U \longrightarrow \mathbb{R}$, defined on some open set U containing x. In other words, $C^\infty_x(M)$ is the union of all $C^\infty(U)$ as U ranges over all open sets containing x. (We think of U as being "small", since we can always restrict any such function to an even smaller open set.)

Unfortunately, $C^\infty_x(M)$ is not a vector space: it does not make sense to add two functions defined on different domains. To fix this, we define a relation, $\sim$, on $C^\infty_x(M)$ by deeming $f : U \longrightarrow \mathbb{R}$ to be related to $g : V \longrightarrow \mathbb{R}$ if there is an open set, $W \subseteq U \cap V$, with $x \in W$, such that $f|W = g|W$. It is readily checked that $\sim$ is an equivalence relation on $C^\infty_x(M)$. We write $\mathcal{G}_x(M) = C^\infty_x(M)/\sim$.

Definition. A *(smooth) germ* at x is an element of $\mathcal{G}_x(M)$.

We can think of a germ as a smooth function defined on an arbitrarily small neighbourhood of x. We can add two germs: just take representatives in $C^\infty_x(M)$, and add them on the intersection of their domains. It is easily checked that this is well defined. We can also define scalar multiplication by a fixed real number in the obvious way. This gives $\mathcal{G}_x(M)$ the structure of a vector space, which is what we were aiming for.

Exercise : Write these definitions more formally, and check that $\mathcal{G}_x(M)$ is indeed a vector space with respect to these operations.

Note that any smooth function on M determines a germ at any $x \in M$. We will see later (see Corollary 8.3) that every germ at x arises in this way (from a smooth function defined on all of M), but we don't need to worry about that for now.

Now, given any $v \in T_xM$ and any $f \in \mathcal{G}_x(M)$, choose any smooth curve $\gamma : I \longrightarrow M$ with $\gamma(0) = x$ and $\gamma'(0) = v$. Write $v.f = (f \circ \gamma)'(0)$. More precisely, take a representative $f : U \longrightarrow \mathbb{R}$, of the germ in $C^\infty_x(M)$, and take the derivative in sense already defined. We have seen that this is well defined independently of the choice of γ. It is also independent of the choice of representative of the germ. The map $[f \mapsto v.f]$ is a linear functional on $\mathcal{G}_x(M)$ (that is, a linear map from $\mathcal{G}_x(M)$ to $\mathbb{R}$, in other words, an element of the dual space). Note that it also satisfies the product (or "Leibniz") rule:

(L): If $f, g \in \mathcal{G}_x(M)$, then $v.(fg) = f(x)(v.g) + g(x)(v.f)$.

Here, of course, $f(x)$ and $g(x)$ are just real numbers.

(One can check that this is equivalent to the usual product rule for $\mathbb{R}^m$, on taking co-ordinates.)

We will see later that this gives a way of making sense of tangents in "abstract manifolds" in Section 6.

Notation : The following informal notation is often used to denote tangent vectors to a manifold. Suppose that $U \longrightarrow V \subseteq \mathbb{R}^m$ is a chart of M. Given $x \in U$, we have an isomorphism, $d_x\phi : T_xM \longrightarrow \mathbb{R}^m$. Let $e_1, \ldots, e_m$ be the standard basis of m. We can denote the tangent vector $(d_x\phi)^{-1}e_i$, by $\frac{\partial}{\partial x_i}\Big|_x \in T_xM$, or simply by $\frac{\partial}{\partial x_i}$, when the point $x \in M$ is understood. Thus, a general tangent vector at x is a linear combination, $\sum_i \lambda_i \frac{\partial}{\partial x_i}$, where $\lambda_i \in \mathbb{R}$. In this notation, the map ϕ is implicitly understood. Note that $\frac{\partial}{\partial x_i}$ can be thought of as differentiating functions in the x_i coordinate direction. That is, if $f : U \longrightarrow \mathbb{R}$ is a smooth (locally defined) function, then

$$\frac{\partial}{\partial x_i}.f = \frac{\partial(f \circ \phi^{-1})}{\partial x_i}$$

in the usual sense. This is often abbreviated to $\frac{\partial f}{\partial x_i}$, suppressing mention of ϕ. In other words, although we are really in M, we are pretending we are in $\mathbb{R}^m$.

If we have an atlas, $\{\phi_\alpha : U_\alpha \longrightarrow V_\alpha\}_{\alpha \in \mathcal{A}}$ for M, then the notation, $\frac{\partial}{\partial x_i^\alpha}$ is often used for the vector $\frac{\partial}{\partial x_i}$, with respect to the chart, ϕ_α (that is, formally $\frac{\partial}{\partial x_i^\alpha} = (d_x\phi)^{-1}e_i$). In this notation, if $x \in U_\alpha \cap U_\beta$, then we have

$$\frac{\partial}{\partial x_i^\alpha} = \sum_{j=1}^m \frac{\partial x_j^\beta}{\partial x_i^\alpha} \frac{\partial}{\partial x_j^\beta}.$$

In the expression $\frac{\partial x_j^\beta}{\partial x_j^\alpha}$, x_j^β denotes the j'th coordinate of the image under ϕ_β, which we can think of as a function on U.

More formally,

$$\frac{\partial x_j^\beta}{\partial x_i^\alpha} = \frac{\partial}{\partial x_i}(\phi_\beta \circ \phi_\alpha^{-1}).$$

In other words $\left(\frac{\partial x_j^\beta}{\partial x_i^\alpha}\right)_{ji}$ is the jacobian of the transition function, $\phi_\beta \circ \phi_\alpha^{-1}$ evaluated at $\phi_\beta(x)$. This can be seen as an expression of the chain rule. Applied to a smooth function, f, on M we get the usual formula

$$\frac{\partial f}{\partial x_i^\alpha} = \sum_j \frac{\partial x_j^\beta}{\partial x_i^\alpha} \frac{\partial f}{\partial x_j^\beta}.$$

Example : The map, $[(r,\theta) \mapsto (r\cos\theta, r\sin\theta)] : (0,\infty) \times (0,2\pi) \longrightarrow \mathbb{R}^2$ is a diffeomorphism from $(0,\infty) \times (0,2\pi)$ to an open subset, $U \subseteq \mathbb{R}$. The inverse, $\psi : U \longrightarrow \mathbb{R}^2$ is a chart. The coordinate functions are r, θ ("polar coordinates). In the above notation (suppressing mention of ψ), we have vectors $\frac{\partial}{\partial r}$ and $\frac{\partial}{\partial \theta}$ on U, given by

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

and

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

We finish this section with an observation which will be useful later.

Note that it is an almost immediate consequence of the definition that any manifold is locally path connected. That is, every point, x, of M has a neighbourhood W, such that if $y \in M$ there is a path in W from x to y, (i.e. a continuous map $\gamma : [0,1] \longrightarrow W$, with $\gamma(0) = x$ and $\gamma(1) = y$). (Let $\phi : U \longrightarrow V$ be a chart with $x \in U$. Choose any $r > 0$ so that $N(\phi x; r) \subseteq V$, and set $W = \phi^{-1}(N(\phi x; r) \subseteq U$.) From this we get:

Lemma 3.5. *Any connected manifold is path connected.*

Proof. Fix any $x \in M$ and let $P \subseteq M$ be the set of points connected to x by a path in M. Check that P is both open and closed, so $P = M$. □

In fact, one can show that any to points of M can be connected by a smooth path. This follows by a similar argument, however there is a technical point. We need to know that if we concatenate two smooth paths, we can smooth them out near the join to make a single smooth path. We leave this as a somewhat technical exercise for those interested. (One way is to use "bump functions" which we discuss in Section 8.)

4. Immersions and submersions

We begin with a formulation of the Inverse Function Theorem which works for manifolds. This amounts mostly to reinterpreting what we already knew in terms of our new definitions.

Let $M \subseteq \mathbb{R}^p$ and $N \subseteq \mathbb{R}^q$ be manifolds. Suppose $f : M \longrightarrow N$ is smooth, and $x \in M$. Write $y = f(x)$. Let $\phi_1 : U_1 \longrightarrow V_1$ and $\phi_2 : U_2 \longrightarrow V_2$ be charts in M and N respectively, with $x \in U_1$ and $y \in U_2$. After replacing U_1 with $U_1 \cap f^{-1}U_2$ and restricting ϕ_1 to the new domain, we can assume that $fU_1 \subseteq U_2$. We now have a map

$$\psi = \phi_2 \circ f \circ \phi_1^{-1} : V_1 \longrightarrow V_2.$$

Let $a = \phi_1 x$. By the chain rule, we get $d_a\psi = P \circ d_x f \circ Q^{-1}$, where $Q = d_x\phi_1$ and $P = d_y\phi_2$ are both linear isomorphisms. In other words, $d_a\psi$ and $d_x f$ are similar, so in particular have the same rank.

Suppose, for the moment, that $\dim M = \dim N = n$ and that $d_x f$ has rank n: that is, it is invertible. The same is true of $d_a\psi$, and so by the usual Inverse Function Theorem (as given in Section 1), there are open sets $V_1' \subseteq V_1$ and $V_2' \subseteq V_2$ with $a \in V_1'$, such so that $\psi|V_1'$ is a diffeomorphism onto V_2'. Now let $U_1' = \phi_1^{-1}V_1'$ and $U_2' = \phi_2^{-1}V_2'$. We see that $f|U_1' : U_1' \longrightarrow U_2'$ is a diffeomorphism.

Replacing U_1 and U_2 by these smaller open sets, we have now shown:

Theorem 4.1. *Suppose M, N are n-manifolds, and that $f : M \longrightarrow N$ is smooth. Suppose $x \in M$ and that $d_x f$ is invertible. Then there are open sets $U_1 \subseteq M$ and $U_2 \subseteq N$ with $x \in U_1$ such that $f|U_1$ is a diffeomorphism to U_2.*

Note that we can assume that U_1 and U_2 are the domains of charts $\phi_1 : U_1 \longrightarrow V_1$ and $\phi_2 : U_2 \longrightarrow V_2$. We have seen that $\psi : V_1 \longrightarrow V_2$ is a diffeomorphism, and it follows that $\psi \circ \phi_1 : U_1 \longrightarrow V_2$ is also a chart. Replacing ϕ_1 by this new chart, we get the following addendum to Theorem 4.1.

Theorem 4.2. *With the same hypotheses as Theorem 4.1, we can find charts $\phi_1 : U_1 \longrightarrow V \subseteq \mathbb{R}^n$ and $\phi_2 : U_2 \longrightarrow V \subseteq \mathbb{R}^n$ around x and $f(x)$ respectively, such that $\phi_1 = \phi_2 \circ f$.*

In other words, with respect to suitable local coordinates, f corresponds to the identity on $\mathbb{R}^n$.

Example. The map $[t \mapsto (\cos t, \sin t)] : \mathbb{R} \longrightarrow S^1$, or indeed any covering space (for people who know about covering spaces).

We want to generalise the above to the case where M and N have different dimensions.

First we need some linear algebra.

Let V, W be vector spaces, and set $m = \dim V$, $n = \dim W$.

Definition. We say that a linear map, $L : V \longrightarrow W$, has *maximal rank* if $\operatorname{rank} L = \min\{m, n\}$.

There are two cases. If $m \leq n$, then L is injective, and if $m \geq n$, then L is surjective. (Of course, these cases overlap if $n = m$ and L is invertible.)

The simplest examples are respectively the *standard immersion*, ι, of $\mathbb{R}^m$ in $\mathbb{R}^n$, given by

$$\iota(x_1, \ldots, x_m) = (x_1, \ldots, x_m, 0, \ldots, 0),$$

and the *standard submersion*, σ, of $\mathbb{R}^m$ to $\mathbb{R}^n$ given by

$$\sigma(x_1, \ldots, x_m) = (x_1, \ldots, x_n).$$

(We will explain the terms "immersion" and "submersion" later.) While these two cases are qualitatively quite different, we will deal with them in parallel, since many of the results have similar formulations.

The following is a simple exercise in linear algebra:

Lemma 4.3. *Let $L : V \longrightarrow W$ be a linear map of maximal rank. Then there are invertible linear maps $P : V \longrightarrow \mathbb{R}^m$ and $Q : W \longrightarrow \mathbb{R}^n$ such that $Q \circ L \circ P^{-1}$ is a standard immersion or submersion.*

One can prove this using bases. Alternatively, it can be reinterpreted as a statement about matrices, which can be proven using either column or row operations. From the latter point of view, the two cases can be viewed as the same result — swapping rows and columns.

We want a variation on this for smooth maps. We state it first for $\mathbb{R}^m$ and $\mathbb{R}^n$, which can be viewed as an elaboration on the Inverse Function Theorem. We will user 0^p to denote the origin in $\mathbb{R}^p$.

Proposition 4.4. *Suppose that $U \subseteq \mathbb{R}^m$ is open and that $f : U \longrightarrow \mathbb{R}^n$ is a smooth map. Suppose that $c \in U$ and that $d_c f$ has maximal rank (that is, $\min\{m, n\}$). Then there are open sets $U_1, V_1 \subseteq \mathbb{R}^m$ and $U_2, V_2 \subseteq \mathbb{R}^n$, with $c \in U_1 \subseteq U$, and $f(U_1) \subseteq U_2$, together with diffeomorphisms, $\theta_1 : U_1 \longrightarrow V_1$ and $\theta_2 : U_2 \longrightarrow V_2$, such that $\theta_1 c = 0^m$, and $\theta_2 \circ f \circ \theta_1^{-1} : V_1 \longrightarrow V_2$ is the restriction of the standard immersion or submersion of $\mathbb{R}^m$ to $\mathbb{R}^n$.*

Of course, this is really two results depending on whether $m \leq n$ or $m \geq n$. They intersect in the case where $m = n$, where Proposition 4.4 just becomes a formulation the Inverse Function Theorem. We will deal with the two cases separately. (In fact, we will see that θ_1 or θ_2 can be taken to be a linear map in the respective cases.)

To simplify the argument, we note that there is no loss in assuming that $d_c f$ is just the standard immersion or submersion. This follows since by Lemma 4.3, we can find linear isomorphisms, P, Q of $\mathbb{R}^n$ and $\mathbb{R}^n$ respectively, so that $P \circ d_c f \circ Q^{-1}$ is standard. Using the chain rule, we can then just replace f by $P \circ f \circ Q^{-1}$. After postcomposing by a translation of $\mathbb{R}^m$, we can also assume that $c = 0^m$ and $f(c) = 0^n$.

Proof. We split into the two cases:
(1) $m \leq n$. Identify $\mathbb{R}^n \equiv \mathbb{R}^m \times \mathbb{R}^{n-m}$. Define a map $F : U \times \mathbb{R}^{n-m} \longrightarrow \mathbb{R}^n$ by $F(x) = f(x_1, \ldots, x_m) + (0, \ldots, 0, x_{m+1}, \ldots, x_n))$. Note that $d_{0^n} F$ is now the

identity on $\mathbb{R}^n$. Therefore, by the Inverse Function Theorem, there are open sets, $U_2, V_2 \subseteq \mathbb{R}^m$ both containing the origin, and a diffeomorphism $\theta_2 : U_2 \longrightarrow V_2$, which is the inverse of $F|V_2$. Let $U_1 = U \cap V_2$. Now if $y \in U_1$, we have $\theta_2(f(y)) = \theta_2(f(y) + 0^n) = \theta_2(F(y, 0^{n-m})) = (y, 0^{n-m}))$. The result now follows by setting θ_1 to be the identity restricted to U_1.
(2) $m \geq n$. Identify $\mathbb{R}^m \equiv \mathbb{R}^n \times \mathbb{R}^{m-n}$. Define a map $G : U \longrightarrow \mathbb{R}^m$ by $G(x) = (f(x), \pi(x))$, where $\pi(x_1, \ldots, x_n) = (x_{n+1}, \ldots, x_m)$. Note that $d_0 G$ is the identity on $\mathbb{R}^m$. By the Inverse Function Theorem, we have open sets $U_1, V_1 \subseteq \mathbb{R}^m$, both containing 0, such that $G|U_1 : U_1 \longrightarrow U_2$ is a diffeomorphism. We now set $\theta_1 = G|U_1$, $U_2 = fU_1$, and θ_2 to be the identity on U_2. $\square$

We can translate this back into a statement about manifolds:

Theorem 4.5. *Suppose that M, N are manifolds, and that $f : M \longrightarrow N$ is smooth. Suppose that $c \in M$, and that $d_c f$ has maximal rank. Then there are charts, $\phi_1 : U_1 \longrightarrow V_1$ and $\phi_2 : U_2 \longrightarrow V_2$ around c and $f(c)$ respectively, with $\phi_1 x = 0$, such that $\phi_2 \circ f \circ \phi_1^{-1}|V_1 : V_1 \longrightarrow V_2$ is the restriction of a standard immersion or submersion.*

This uses Proposition 4.4 in the same way that Theorem 4.2 used the Inverse Function Theorem. The argument is essentially the same, so we leave it as an exercise.

Definition. We say that a map $f : M \longrightarrow N$ is an *immersion* if for all $x \in M$, $d_x f$ is injective. We say that it is a *submersion* if for all $x \in M$, $d_x f$ is surjective.

Clearly, these imply, respectively, that $\dim M \leq \dim N$, and $\dim M \geq \dim N$. Note that the composition of immersions is an immersion and a composition of submersions is a submersion.

Note that an immediate consequence of Theorem 4.5 is that an immersion is locally injective: that is, for all $x \in M$, there is an open set, U, containing x such that $f|U$ is injective. Similarly, a submersion is open, that is fU is open for any open $U \subseteq M$.

Examples. Lots of familiar curves in the plane, such as the lemniscate ("figure of eight"), are examples of immersions. Provided they don't have cusps — such as the cuspidal cubic.

The familiar picture of the Klein bottle drawn in 3-space is an example of a 2-manifold immersed in $\mathbb{R}^3$. The "Boy surface" in an immersion of the projective plane into $\mathbb{R}^3$.

Definition. A smooth map, $f : M \longrightarrow N$ between manifolds is an *embedding* if it is a diffeomorphsim onto its range. We refer to the range, $f(M)$, of such a map as a *submanifold* of N.

Clearly any embedding is an injective immersion, though the converse need not be true. A counterexample is the injective map of $[0, 1)$ to the plane whose image is a "figure of six".

Note that if $M \subseteq \mathbb{R}^p$ is a manifold in $\mathbb{R}^p$ (according to our original definition of such), then M is a submanifold of $\mathbb{R}^p$, according to the definition we have just given.

Recall that a map (between topological spaces) is *proper* if the preimage of any compact set is compact.

Lemma 4.6. *Any proper injective immersion is an embedding.*

Proof. This is an immediate consequence of the fact (mentioned in Section 1, that a continuous proper injective map between locally compact spaces is a homeomorphism onto its range. □

We refer to the image of a proper embedding as a *proper submanifold*.

Note this is the same as being closed in $\mathbb{R}^p$ (Exercise).

Exercise : Suppose that N, N' are manifolds, and that $M \subseteq N$ and $M' \subseteq N'$ are submanifolds. Suppose that $f : N \longrightarrow N'$ is smooth, with $f(M) \subseteq M'$. Then $f|M : M \longrightarrow M'$ is smooth (with respect to the intrinsic smooth structures). If f is a diffeomorphism, and $f(M) = M'$, then $f|M'$ is a diffeomorphism from M to M'.

Definition. Let $f : M \longrightarrow N$ be a smooth map between manifolds. We say that $x \in M$ is a *regular point* if $d_x f$ is surjective, and a *critical point* otherwise. A point $y \in N$ is a *regular value* if each point of $f^{-1}(y)$ is a regular point, otherwise it is a *critical value*.

In other words, a critical value is the image of a critical point. Note that any point of $N \setminus f(M)$ has empty preimage, and is therefore a regular value.

Theorem 4.7. *If $f : M \longrightarrow N$ is smooth, and $y \in f(M) \subseteq N$ is a regular value, then $f^{-1}(y)$ is a proper submanifold of M of dimension $m - n$.*

Since f is continuous, $f^{-1}(y)$, is closed, so it is a proper submanifold.

The fact that we insist that manifolds are non-empty, explains why we disallow $y \in M \setminus f(N)$, in this statement. Note that we can assume that $\dim M \leq \dim N$ (otherwise, there is no regular point, and the statement becomes vacuous).

Proof. Let $x \in f^{-1}(y)$. Choose charts as given by Theorem 4.5. Then $f^{-1}(y) \cap U_1$ is the set of points in U_1 with first m coordinates all equal to 0. Let $\pi : \mathbb{R}^n \longrightarrow \mathbb{R}^{m-n}$ be the projection to the final $m - n$ coordinates. Let $U = f^{-1}(y) \cap U_1$, and set $\phi = \pi \circ \phi_1 : U \longrightarrow \mathbb{R}^{m-n}$. The $x \in U$, and $\phi : U \longrightarrow \mathbb{R}^{m-n}$ is a chart for $f^{-1}(y)$. □

We also note:

Lemma 4.8. *With the hypotheses of Theorem 4.7, we have* $T_x(f^{-1}(y)) = \ker d_x f$.

Proof. By Lemma 3.2 any tangent vector in $T_x(f^{-1}(y))$ has the form $\gamma'(0)$ for a smooth curve, $\gamma : I \longrightarrow f^{-1}(y)$, where $\gamma(0) = x$. Now, $f \circ \gamma$ is constant, so $d_x f(\gamma'(0)) = (f \circ \gamma)'(0) = 0$, so $\gamma'(0) \in \ker d_x f$. This shows that $T_x(f^{-1}(y)) \subseteq \ker d_x f$. But $\dim T_x(f^{-1}(y)) = m - n = \dim(\ker d_x f)$, so $T_x(f^{-1}(y)) = \ker d_x f$. $\square$

Examples

(1) Define $f : \mathbb{R}^3 \longrightarrow \mathbb{R}$ by $f(x_1, x_2, x_3) = x_1^2 + x_2^2 + x_3^2$. The only critical point is $(0, 0, 0)$, so the only critical value is 0. Also, $f^{-1}(t)$ is non-empty precisely if $t \geq 0$. Thus, $f^{-1}(t)$ is a manifold if $t > 0$ (the sphere of radius $\sqrt{t}$).

(2) Similarly, Define $f : \mathbb{R}^3 \longrightarrow \mathbb{R}$ by $f(x_1, x_2, x_3) = x_1^2 + x_2^2 - x_3^2$. Again, the only critcal point is $(0, 0, 0)$, so the only critical value is 0. In this case, $f^{-1}(t)$ is always non-empty. Thus, $f^{-1}(t)$ is a manifold if $t \neq 0$: it is a hyperboloid. Note that $f^{-1}(0)$ is a cone — not a manifold.

(3) Let $T_{a,b} \subseteq \mathbb{R}^3$ be the torus, described in Example (6). Define $f : T_{a,b} \longrightarrow \mathbb{R}$ by $f(x_1, x_2, x_3) = x_1$. In this case, $f^{-1}(t)$ is a manifold for $t \in (-a - b, -a + b) \cup (-a + b, a - b) \cup (a - b, a + b)$.

(4) We noted in Section 3 that $GL(n, \mathbb{R})$ is an n^2-manifold, with the group operations smooth. From this, we can derive other similar examples. For example, consider the map $\Delta : GL(n, \mathbb{R}) \longrightarrow \mathbb{R}$, given by $\Delta(A) = \det A$. Considered as a map from $\mathbb{R}^{n^2}$ to $\mathbb{R}$, it is smooth (in fact, polynomial). We claim that if $t \neq 0$, then t is a regular value. For suppose $\Delta(A) = t \neq 0$. Consider the path $\gamma : \mathbb{R} \longrightarrow GL(n, \mathbb{R})$ given by $\gamma(u) = (1 + u)A$. Now $\Delta \circ \gamma(u) = (1 + u)^n t$, so we get $d_{\gamma(0)}\Delta(\gamma'(0)) = (\Delta \circ \gamma)'(0) = nt \neq 0$, and so $d_A \Delta$ has rank 1. It follows that $\Delta^{-1}(t)$ is a manifold. Let $SL(n, \mathbb{R}) = \Delta^{-1}(1) = \{A \in GL(n, \mathbb{R}) \mid \det A = 1\}$. This is a group. Moreover (using an earlier exercise) the group operations are smooth (since they are restrictions of the smooth group operations on $GL(n, \mathbb{R})$).

A similar argument shows that $O(n, \mathbb{R}) = \{A \in GL(n, \mathbb{R}) \mid A^T A = I\}$ is a manifold, with smooth group operations.

(5) Some "real world" examples arise from mechanical linkages. These have been used by engineers since ancient times. One can think of a linkage as a collection of rods connected at pivots, where they are allowed to flex. In mathematical terms, one can think of the pivots as a set of n points in the plane (for a planar linkage), and so described by a point in $\mathbb{R}^{2n}$. The rods determine a number of constraints, such as the fact that certain pivots lie on a straight line, or that the distances between them is fixed. There are given by smooth equations, that is a map from

$\mathbb{R}^{2n}$ to some $\mathbb{R}^d$. The configuration of the linkage is then constrained to lie in the preimage of a point. If this is a regular point (as one would expect generically, or as one might hope, if the linkage is to function smoothly), then this is a manifold.

(6) Analogous situations arise in many contexts in physics, where one has a number of invariants in a system. For example, we may have a set of particles, or planets or whatever, whizzing around in space. The state of the system at a given moment may describe a finite number of coordinates of position and velocity, say. The energy of the system is constant, so they are constrained to live in some subset of the coordinate space — typically a manifold.

5. Tangents, normals, orientations

Let $M \subseteq \mathbb{R}^n$ be an m-manifold.

Definition. The *tangent bundle*, TM, to M is defined by:

$$TM = \{(x, v) \in M \times \mathbb{R}^n \mid v \in T_x M\}.$$

We write $p : TM \longrightarrow M$ for the projection, $p(x, v) = x$. Thus, $p^{-1}x = \{x\} \times T_x M$, which we can identify with $T_x M$.

Proposition 5.1. *TM is a manifold of dimension $2m$, and $p : TM \longrightarrow M$ is a submersion.*

Proof. Let $\phi : U \longrightarrow V$ be a chart. Define a map $\psi : p^{-1}U \longrightarrow V \times \mathbb{R}^m$ by $\psi(x, v) = (\phi x, d_x\phi(v))$. This is smooth, and has smooth inverse defined by $[(y, w) \mapsto (\phi^{-1}(y), d_y\phi^{-1}(w))]$. We see that the collection of such maps form an atlas for TM.

Now $p = \phi^{-1} \circ \sigma \circ \psi$, where $\sigma : V \times \mathbb{R} \longrightarrow V$ is projection to the first coordinate. Thus, p is a composition of submersions, hence a submersion. $\square$

Examples TS^1 is a cylinder. The map $[(x_1, x_2, t) \mapsto (x_1, x_2, -tx_2, tx_1)$ gives a diffeomorphism from $S^1 \times \mathbb{R} \subseteq \mathbb{R}^3$ to $TS^1 \subseteq \mathbb{R}^2 \times \mathbb{R}^2 \equiv \mathbb{R}^4$.

Suppose $f : M \longrightarrow N$ is a smooth map between manifolds. We get an induced map, $f_* : TM \longrightarrow TN$ defined by $f_*(v) = d_x f(v) \in T_{\phi x} M$ where $x = pv$.

Exercise : $f_* : TM \longrightarrow TN$ is smooth.

The following is a special case of a more general definition we give later, in Section 7.

Definition. A *section* of TM is a smooth map, $s : M \longrightarrow TM$ such that $p \circ s$ is the identity on M.

Note that a section, s, is necessarily an immersion. In fact, a proper embedding of M in TM.

In other words, we are assigning to each $x \in M$ a tangent vector, $s(x) \in T_x M$, in a smooth manner. In more familiar terms, s is just a smooth "vector field", on M.

Definition. A *frame field* on M is a family, of m smooth vector fields, $v_1, \ldots, v_m$, such that $v_1(x), \ldots, v_m(x)$ forms a basis for $T_x M$ for all $x \in M$.

Note that a frame field need not always exist on all of M. (Will see an example in Section 7.)

However, one can always find frame field locally. In fact, suppose $\phi : U \longrightarrow \mathbb{R}^m$ is a chart for M. Let $e_1, \ldots, e_m$ be the standard basis for $\mathbb{R}^m$ (or indeed any fixed basis). If $x \in U$, set $v_i(x) = (d_x\phi)^{-1} e_i$. Then $v_1, \ldots, v_n$ is a frame field on U (which we can think of an open submanifold of M).

Notation : Again, the notation $\frac{\partial}{\partial x_i}$ is often used informally for $v_i(x)$ defined locally. In this context, we can think of $\frac{\partial}{\partial x_i}$ as denoting a vector field on $U \subseteq M$. If $\lambda_i : U \longrightarrow \mathbb{R}$ are smooth functions, then $[x \mapsto \sum_{i=1}^m \lambda_i(x) v_i(x)]$ is also a vector field, often informally denoted as $\sum_i \lambda_i \frac{\partial}{\partial x_i}$.

We say that a smooth frame field is *orthonormal* if $(v_i(x))_i$ is a orthonormal basis with respect to the dot product induced by the embedding of $T_x M$ in $\mathbb{R}^n$. That is, $v_i(x).v_j(x) = \delta_{ij}$ where $\delta_{ij} = 0$ if $i \neq j$ and $\delta_{ij} = 1$ if $i = j$.

We note:

Lemma 5.2. *If M admits a (global) frame field, then it admits a global orthonormal frame field.*

Proof. Recall the Gram-Schmidt process for producing an orthonormal basis from a given basis. Start with any frame field $v_i(x)$. Set $v_1'(x) = v_1(x)$. Then set

$$v_2'(x) = v_2(x) - \frac{v_1(x).v_2(x)}{v_1(x).v_1(x)} v_1(x).$$

Then set

$$v_3'(x) = v_3(x) - \frac{v_1(x).v_3(x)}{v_1(x).v_1(x)} v_1(x) - \frac{v_2(x).v_3(x)}{v_2(x).v_2(x)} v_2(x),$$

etc. Finally, set

$$w_i(x) = v_i'(x)/||v_i'(x)||.$$

Then $(w_i(x))_i$ is orthonormal.

It suffices to observe that all the above operations are smooth (they are given by nice simple formulae). Therefore the resulting maps $[x \mapsto w_i(x)]$ are smooth. $\square$

In particular, we see that there is always a locally defined orthonormal frame field for a manifold embedded in euclidean space.

We now move on to consider normal vectors.

Recall that $(T_xM)^\perp$ is the orthogonal complement to T_xM in $\mathbb{R}^n$; that is the space of "normal vectors" to M at x, with respect to the dot product on $\mathbb{R}^n$.

Definition. The *normal bundle* to M in $\mathbb{R}^n$ is defined by:

$$\nu(M, \mathbb{R}^n) = \{(x, v) \in M \times \mathbb{R}^n \mid v \in (T_xM)^\perp\}.$$

We write $p : \nu(M, \mathbb{R}^n) \longrightarrow M$ for the projection map.

Proposition 5.3. *$\nu(M, \mathbb{R}^n)$ is an n-manifold, and $p : \nu(M, \mathbb{R}^n) \longrightarrow M$ is a submersion.*

(Here, we should think of the dimension, n, as $n = m + (n - m)$. Note that m of the coordinates come from M, and the remaining $n - m$ from the normal space.)

Proof. Let $x_0 \in M$. Let $\phi : U \longrightarrow \mathbb{R}^m$ be a chart, with $x_0 \in U$, and let $v_1(x), \dots, v_m(x)$, be a frame field defined on U (as discussed above). In partcular, $v_1(x_0), \dots, v_m(x_0)$ is a basis for $T_{x_0}(M) \subseteq \mathbb{R}^m$, and we extend this arbirarily to a basis $v_1(x_0) \dots, v_m(x_0), v_{m+1}, \dots, v_n$ of $\mathbb{R}^n$. By continuity, we see easily that $v_1(x) \dots, v_m(x), v_{m+1}, \dots, v_n$ is a basis of $\mathbb{R}^n$ for all x in some open neighbourhood, say U', of x in $U \subseteq M$. Now for $x \in U'$ set $v_i(x) = v_i$ for $i \in \{m+1, \dots, n\}$. We apply the Gram-Schmidt process to this to give us an orthonormal frame, $w_1(x), \dots, w_n(x)$, for $\mathbb{R}^n$, which varies smoothly in x. (This is just an extension of the construction used in proving Lemma 5.2.) Now $w_i(x) \in T_xM$ for $i \leq m$, and $w_i(x) \in (T_xM)^\perp$ for $i > m$. Set $e_i(x) = w_{m+i}(x)$ for $i > m$. Then $e_1(x), \dots, e_{n-m}(x)$ is an orthonormal basis for $(T_xM)^\perp$. (We will no longer need $w_1(x), \dots, w_m(x)$.)

We now define a map $\psi : p^{-1}U' \longrightarrow V \times \mathbb{R}^{n-m}$ by setting $\psi(x, v) = (\phi(x), \lambda(x))$, where $\lambda(x) = (v.e_1(x), \dots, v.e_{n-m}(x)) \in \mathbb{R}^{n-m}$. This is a smooth map, and has smooth inverse given by $(y, \lambda) \mapsto (\phi^{-1}y, \sum_{i=1}^{n-m} \lambda_i(x)e_i(x))$, where $\lambda(x) = (\lambda_1(x), \dots, \lambda_{n-m}(x))$.

We finally note that p is a submersion for a similar reason as for the tangent bundle. $\square$

Examples. $\nu(S^m, \mathbb{R}^{m+1})$ is diffeomorphic to $S^m \times \mathbb{R}$.

We can define a *section* of $\nu(M, \mathbb{R}^n)$ as a smooth map, $\kappa : M \longrightarrow \nu(M, \mathbb{R}^m)$ with $p \circ \kappa$ the identity on M. Such a section is commonly called a *normal field* to M in $\mathbb{R}^n$.

We now consider orientations on a manifold. We begin with some general linear algebra.

Let V be a vector space of dimension $m > 0$. Let $I(V)$ be the set of linear isomorphisms from V to $\mathbb{R}^m$. Given $\rho, \sigma \in I(V)$, we get a linear automorphism, $\sigma \circ \rho^{-1} : \mathbb{R}^m \longrightarrow \mathbb{R}^m$. Note that $\det(\sigma \circ \rho^{-1}) \neq 0$. We write $\sigma \sim \rho$ if $\det(\sigma \circ \rho^{-1}) > 0$. This is easily seen to define an equivalence relation on $I(V)$, and we write $\mathrm{Or}(V) = I(V)/\sim$. Thus $|\mathrm{Or}(V)| = 2$. In the case where $\dim V = 0$, we set $\mathrm{Or}(V) = \{-1, +1\}$.

Definition. An *orientation* on V is an element of $\mathrm{Or}(V)$. An *oriented vector space* is a (finite dimensional) vector space equipped with an orientation.

This can be thought of, perhaps more intuitively, in terms of bases. Let V be an oriented vector space of positive dimension. Note that a basis, $v_1, \ldots, v_m$, of V determines an element of $I(V)$ sending the basis to the standard basis of $\mathbb{R}^m$. We refer to $(v_i)_i$ as *positively* oriented if this map lies in the class of the orientation, and *negatively oriented* otherwise.

Note that $\mathbb{R}^m$ itself comes with a natural "standard" orientation, namely the class of the identity map. In other words, the standard basis is deemed to be positively oriented.

If V, W are oriented vector spaces of the same dimension, we say that a linear isomorphism $L : V \longrightarrow W$, is *orientation preserving* if it sends some (hence any) positively oriented basis to a positively oriented basis. Otherwise, it is *orientation reversing*. (Of course, this can also be expressed directly in terms of the orientation classes of isomorphism to $\mathbb{R}^m$.)

Note that an automorphism, L, of $\mathbb{R}^m$ with the standard orientation is orientation preserving if $\det L > 0$, and orientation reversing if $\det L < 0$.

Suppose that V, W are vector spaces and $E = V \oplus W$. Then orientations on V and W determine an orientation on E. For example, choose positively oriented bases for V and W. Their union is a basis for E, which we deem to be positively oriented in E. Exercise: Check this is well defined, independently of the bases we choose. Conversely, an orientation on V and an orientation on E determine an orientation on W.

We can now define an orientation on an m-manifold, M, in $\mathbb{R}^n$ when $m > 0$.

Definition. An *orientation* on M is an assignment of an orientation to each tangent space T_xM such that there is an atlas of charts, $\phi_\alpha : U_\alpha \longrightarrow \mathbb{R}^m$, indexed by some set $\mathcal{A}$, such that for all $\alpha \in \mathcal{A}$ and all $x \in U_\alpha$, the map $d_x\phi_\alpha : T_xM \longrightarrow \mathbb{R}^m$ orientation preserving (i.e. in the orientation class of the orientation of T_xM).

We refer to M as an *oriented manifold*.

Definition. We say that M is *orientable* if it admits an orientation.

Given an orientation on M we have an *opposite* orientation obtained by a reversing the orientation on each tangent space. To see that this is indeed or orientation, take the atlas, and postcompose every chart with an orientation reversing linear automorphism of $\mathbb{R}^m$ (e.g. $[(x_1, \dots, x_m) \mapsto (-x_1, \dots, x_m)]$), to give another smooth atlas with all orientations reversed.

Not every manifold is orientable (as we will see). However, every manifold is "locally orientable" in the sense that every point is contained in an open set, U, which is orientable. In fact, if $\phi : U \longrightarrow V$ is chart, then we just use the maps $d_x\phi : T_xM \longrightarrow \mathbb{R}^m$ to define the orientation on the tangent space $T_xU = T_xM$ at x.

It is easily seen that orientability is invariant under diffeomorphism.

Examples :

(1): Clearly $\mathbb{R}^n$ is orientable.

(2): $S^n \subseteq \mathbb{R}^{n+1}$ is orientable. If we take the atlas given by stereographic projections from the north and south poles, then the transition function reverses orientation. (Recall that it is inversion in an $(n-1)$-sphere.) To fix this, we just postcompose one of the charts with any orientation reversing diffeomorphism of $\mathbb{R}^n$, so as to give an oriented atlas.

(3): The direct product of any two orientable manifolds is orientable. (So for example, the torus is orientable. So is the cylinder $S^1 \times \mathbb{R}$.)

(4) The Möbius band is not orientable. (In particular, it is not diffeomorphic to the cylinder.)

We have noted that if M is orientable, then it has at least two orientations. It might have many, since orientations on different connected components of M are independent of each other. However, if M is connected, it has precisely two.

To see this we begin with the following observation.

Lemma 5.4. *Suppose that we have two orientations on M. Let $A \subseteq M$ be the set of $x \in M$ such that the two orientations agree on T_xM. Then A is open and closed in M.*

Proof. Let $x \in M$. Let $\phi_\alpha : U_\alpha \longrightarrow \mathbb{R}^m$ and $\phi_\beta : U_\beta \longrightarrow \mathbb{R}^m$ be charts for the respective orientations, with $x \in U_\alpha \cap U_\beta$. For each $y \in U_\alpha \cap U_\beta$, we have a linear automorphism $(d_y\phi_\beta) \circ (d_y\phi_\alpha)^{-1} : \mathbb{R}^m \longrightarrow \mathbb{R}^m$. Let $\Delta(y) = \det((d_y\phi_\beta \circ (d_y\phi_\alpha)^{-1})$. Then $\Delta(y) \neq 0$, and $\Delta : U_\alpha \cap U_\beta \longrightarrow \mathbb{R}^m$ is smooth, so in particular continuous. Therefore, if $x \in A$, then $\Delta(x) = 1$, so $\Delta(y) = 1$ for all y in a neighbourhood of x, so $y \in A$. This shows that A is open. Similarly, if $\Delta = -1$, then $\Delta(y) = -1$ on a neighbourhood of x, and so A is also closed. $\square$

Lemma 5.5. *A connected orientable manifold has precisely two orientations.*

Proof. Choose any $x_0 \in M$. After reversing one of the orientations if necessary, we can assume that they both agree on $T_{x_0}M$. So, in Lemma 5.4 $A \neq \emptyset$, and so $A = M$. $\square$

Suppose that $M \subseteq \mathbb{R}^n$ is a smooth manifold. Given $x \in M$, then $\mathbb{R}^n \cong T_x\mathbb{R}^n = T_xM \oplus (T_xM)^\perp$. Now $\mathbb{R}^n$ comes with a standard orientation. Thus, as discussed before, an orientation on T_xM determines an orientation on $(T_xM)^\perp$ and conversely.

Consider the case when $n = m + 1$. Then $(T_xM)^\perp \cong \mathbb{R}$, and so an orientation on T_xM, hence on $(T_xM)^\perp$ gives us a unique $\kappa(x) \in (T_xM)^\perp$ with $||\kappa(x)|| = 1$ and with $\{\kappa(x)\}$ a positive basis for $(T_xM)^\perp$. We can think of $\kappa(x)$ as the "outward" unit normal vector. If M is oriented, then $\kappa(x)$ is defined everywhere. Also, the map $[x \mapsto \kappa(x)]$ is smooth (Exercise: see the proof of Proposition 5.3.) Thus, $\kappa(x)$ is a normal field.

Conversely, if κ is a (global) nowhere vanishing normal field on M, then κ gives rise to an orientation on T_xM for all x. From this, it is not hard to construct an orientable atlas.

From this one can deduce:

Theorem 5.6. *Let $M \subseteq \mathbb{R}^{m+1}$ be an m-manifold. The following are equivalent:*
(1) M is orientable.
(2) M admits a nowhere-vanishing normal field.
(3) M admits a unit normal field.

Definition. Suppose that M and N are oriented manifolds, and that $f : M \longrightarrow N$ is a diffeomorphism. We say that f *preserves orientation* if $d_xf : T_xM \longrightarrow T_{f(x)}N$ respects the given orientations for all $x \in M$. We say that it *reverses orientation* if it reverses orientation for all $x \in M$.

Exercises :

(1) If M, N are connected and oriented, then every diffeomorphism either preserves or reverses orientation.

(2) Let $f : S^n \longrightarrow S^n$ be the antipodal map on the n-sphere $S^n \subseteq \mathbb{R}^{n+1}$. (That is, $f(x) = -x$.) Show that f is orientation preserving if n is odd, and orientation reversing if n is even.

6. Abstract manifolds

Up until now, all our manifolds have come with embeddings into some euclidean space, $\mathbb{R}^n$: the "ambient space". They had to, because that formed part of our definition of a "manifold". However, many of the constructions only really make essential reference to an atlas of charts. For example, one can talk about smooth functions, immersions and submersions of one manifold in another, etc, just using the atlas. Our definition of tangent space made reference to the embedding in euclidean space, though it seems that appropriately formulated, a "tangent vector" should in some sense live in the manifold itself, so we should be able to talk about a vector field "on a manifold" without worrying about the ambient space. The same should apply to orientability. The key point here is that these notions are respected by diffeomorphisms of one manifold to another.

Of course, a few of our constructions, such the normal bundle, do make essential reference to the ambient space. Similarly, the notion of an "orthonormal frame" uses the ambient space in order to make sense of the dot product. These will not in general be respected by diffeomorphism — they are not properties of general abstract manifolds.

In this section, we will start all over again, defining a manifold intrinsically in terms of an atlas. We can develop the theory quite quickly, since much of it is essentially just repetition of what we have already done for manifolds with the old definition. The trickiest bit will be to define what we mean by a "tangent space" in this context.

One reason for going to all this trouble is that many manifolds are easy to describe in the abstract (see for example the "quotient spaces" below), but it would be quite complicated and unnatural to describe them as subsets of euclidean space.

We move on to a more formal definition.

Recall that a topological space is *second countable* if it has a countable base of open sets.

Definition. A *topological m-manifold* is a non-empty second countable hausdorff topological space for which every point has an open neighbourhood which is homeomorphic to an open subset of $\mathbb{R}^m$.

Definition. Let M be a topological space. A *topological chart* for M is a homeomorphism, $\phi : U \longrightarrow V$, where $U \subseteq M$, is open, and where $V \subseteq \mathbb{R}^m$ is open.
An *topological atlas* for M is a collection, $\{\phi_\alpha : U_\alpha \longrightarrow V_\alpha\}_{\alpha\in\mathcal{A}}$, of charts, indexed by some set, $\mathcal{A}$, such that $M = \bigcup_{\alpha\in\mathcal{A}} U_\alpha$.

In other words, a topological manifold is a non-empty separable hausdorff topological space which admits an atlas.

Remark : The "second countable" assumption turns out to be necessary to develop the theory beyond a certain point (as we will see). Without this there are curious examples, even in dimension 1 (read about the "long line" if you are interested). Even in 2 dimensions describing such examples systematically becomes a pretty hopeless project. Moreover, there are few contexts in which such examples arise naturally. If one assumes non-empty hausdorff, and the existence of an atlas, then it turns out that being second countable is equivalent to a range of other hypotheses: for example "separable" (that is, having a countable dense subset) or "metrisable" (that is, admitting a metric which induces the original topology), or indeed at least thirty others. One we will refer to later is "paracompact". We will say more about that in Section 8. We won't be using this assumption for the moment.

Let M be a topological manifold, with atlas $\{\phi_\alpha\}_{\alpha \in \mathcal{A}}$. If $\alpha, \beta \in \mathcal{A}$, we have a *transition map* $\phi_\beta \circ \phi_\alpha^{-1} : \phi_\alpha(U_\alpha \cap U_\beta) \longrightarrow \phi_\beta(U_\alpha \cap U_\beta)$.

Definition. An atlas is *smooth* if the transition map $\phi_\beta \circ \phi_\alpha^{-1}$ is smooth for all $\alpha, \beta \in \mathcal{A}$.

(Of course, this clause is trivially satisfied whenever $\alpha = \beta$, and is vacuously satisfied whenever $U_\alpha \cap U_\beta = \emptyset$.)

Examples.

(1) Every smooth m-manifold in $\mathbb{R}^n$, by the definition in Section 3 gives such an example. By hypothesis it has a (smooth) atlas, and the transition maps for such an atlas are necessarily smooth, as we have observed. It is also automatically second countable.

(2) The real projective space. Define an equivalence relation, $\sim$, on $\mathbb{R}^{n+1} \setminus \{0\}$ by writing $x \sim y$ if there is some $\lambda \in \mathbb{R} \setminus \{0\}$ with $y = \lambda x$. Note that this restricts to the antipodal equivalence relation on $S^n \subseteq \mathbb{R}^{n+1}$. One can check (exercise) that the inclusion $S^n/\!\sim \longrightarrow (\mathbb{R}^{n+1} \setminus \{0\})/\!\sim$ is a homeomorphism. We denote this by P^n (sometimes $\mathbb{R}P^n$) "(real) projective n space".

If $n = 2$, this is the projective plane. It is homeomorphic to a manifold in $\mathbb{R}^6$ (see exercise sheet). One can generalise this to show that P^n is a manifold (in some higher dimensional euclidean space, $\mathbb{R}^N$). However, this a rather complicated way of describing it. One can show that it is an abstract manifold much more simply (see the exercise sheets).

(3) Let Γ be a group acting by isometries on $\mathbb{R}^n$. We say that an open set $U \subseteq \mathbb{R}^n$ is *wandering* if $gU \cap U = \emptyset$ for all $g \in \Gamma \setminus \{1\}$. We say that Γ is a *deck* group if, given $x \in \mathbb{R}^n$, there is a wandering open set $U \subseteq \mathbb{R}^n$ containing x. (This is the same as acting "freely and properly discontinuously" in terminology from other courses.) In this case, the quotient space, $M = \mathbb{R}^n/\Gamma$ is hausdorff, and we write $p : \mathbb{R}^n \longrightarrow M$ for the quotient map. In fact, for any wandering open set U, the map $p|U : U \longrightarrow pU$ is a homeomorphism. We can define an atlas on M by taking the collection of maps $(p|U)^{-1} : pU \longrightarrow U \subseteq \mathbb{R}^m$, as U ranges over all wandering sets. It can be checked (exercise) that gives M the structure of a smooth manifold. This construction gives rise to many examples:

(3a) $\mathbb{Z}$ acts on $\mathbb{R}$ by $n.x = x + n$. The quotient $\mathbb{R}/\mathbb{Z}$ is diffeomorphic to S^1 (exercise).

(3b) $\mathbb{Z}$ acts on $\mathbb{R}^2$ by $n.(x, y) = (x+n, y)$. Then $\mathbb{R}^2/\mathbb{Z}$ is diffeomorphic a cylinder.

(3c) $\mathbb{Z}$ acts on $\mathbb{R}^2$ by $n.(x, y) = (x + n, (-1)^n y)$. Then $\mathbb{R}^2/\mathbb{Z}$ is diffeomorphic to a Möbius band.

(3d) $\mathbb{Z}^2$ acts on $\mathbb{R}^2$ by $(m, n).(x, y) = (x+m, y+n)$. Then $\mathbb{R}^2/\mathbb{Z}^2$ is diffeomorphic to the torus.

(3e) The group, Γ, generated by $[(x, y) \mapsto (x+1, -y)]$ and $[(x, y) \mapsto (x, y+1)]$ is a deck group. The quotient, $\mathbb{R}^2/\Gamma$ is a compact 2-manifold, whose diffeomorphism type is called the *Klein bottle*.
Exercise: Describe an embedding of the Klein bottle in $\mathbb{R}^4$. (One can show that it cannot be embedded in $\mathbb{R}^3$.)

(3f) There are many more examples in higher dimensions. For example, the are exactly 10 diffeomorphism types of compact 3-manifolds arising in this way as quotients of $\mathbb{R}^3$ (one obvious example is the 3-torus, $S^1 \times S^1 \times S^1 = \mathbb{R}^3/\mathbb{Z}^3$). It would be complicated and unnatural to attempt to describe each of these as a submanifold of euclidean space. (Read about "Bieberbach groups" or "crystallographic groups" if you want to learn more.)

(4) One can do similar constructions in other situations. For example, on can take quotients of hyperbolic space to construct (locally) hyperbolic manifolds.

Suppose that M is a manifold equipped with a smooth atlas. We say that a map, $f : M \longrightarrow \mathbb{R}^n$, is *smooth* if $f \circ \phi_\alpha^{-1} : V_\alpha \longrightarrow \mathbb{R}^n$ is smooth for all charts $\phi_\alpha : U_\alpha \longrightarrow V_\alpha$ in the given atlas. (We will generalise this definition shortly.)

Definition. Suppose we have two smooth atlases, $\{\phi_\alpha\}_{\alpha\in\mathcal{A}}$ and $\{\phi_\beta\}_{\beta\in\mathcal{B}}$ on the same topological manifold, M. We say that these two atlases are *equivalent* if their union $\{\phi_\alpha\}_{\alpha\in\mathcal{A}\cup\mathcal{B}}$ is also an atlas.

In other words, the transition maps between the charts of the different atlases are also smooth. Clearly, this is an equivalence relation on the set of all atlases for M.

Exercise : Two atlases are equivalent if and only if they determine the same set of smooth functions to $\mathbb{R}$. (By the second statement, we mean that $f : M \longrightarrow \mathbb{R}$ is smooth with respect to the first atlas if and only if it smooth with respect to the second.)

Definition. A *smooth structure* on a topological manifold, M, is an equivalence class of smooth atlases.

Note that the class of smooth function to $\mathbb{R}^n$ depends only on the smooth structure of M.

Remark : It is not hard to see that, in fact, any equivalence class of atlases contains a unique maximal atlas in the class. In other words any other atlas in the class is a subset of this maximal atlas. This gives another way of defining a "smooth structure": a smooth structure on M is essentially the same thing as a maximal atlas.

Definition. A *smooth manifold* is a topological manifold equipped with a smooth structure.

Definition. Let M, N be smooth manifolds. A map $f : M \longrightarrow N$ is *smooth* if $\psi_\beta \circ f \circ \phi_\alpha^{-1} : \phi_\alpha(U_\alpha \cap f^{-1}U_\beta) \longrightarrow \mathbb{R}^n$ is smooth for all charts $\phi_\alpha : U_\alpha \longrightarrow V_\alpha$ and all charts $\psi_\beta : U_\beta \longrightarrow V_\beta$ in the smooth atlases of M and N respectively.

Here, by "the" smooth atlas, we mean any atlas in the given smooth structure: it doesn't matter which one. (Alternatively, take the unique maximal atlas in the class.)

It is easily checked that a smooth map is necessarily continuous. (Or just add this the hypothesis, if you don't want to be bothered.)

Definition. A *diffeomorphism* between two smooth manifolds is a smooth map with a smooth inverse.

This agrees with the definitions given for smooth manifolds embedded in euclidean space.

Example. The identity map from $\mathbb{R}^n$ to $\mathbb{R}^n$ gives us a smooth atlas on $\mathbb{R}^n$ consisting of a single chart. This gives us a smooth structure on $\mathbb{R}^n$ as a n-manifold: the "standard" smooth structure.

If $n = 1$, this is the standard smooth structure on $M = \mathbb{R}$. But we could (somewhat perversely) take $N = \mathbb{R}$ with the map $[x \mapsto x^3]$. This also gives us a one-chart atlas, hence a smooth structure. But it is different smooth structure. The union of these two atlases (in other words the two charts together) do not form an atlas. This is because the transition map $[x \mapsto x^3]$ is not a diffeomorphism: its inverse is not smooth.

Nevertheless, M and N are diffeomorphic. The map $[x \mapsto x^{1/3}] : M \longrightarrow N$ is a diffeomorphism. This is just a question of formally checking the definitions. So N is really just a copy of the standard real line after all — it's just that we initially chose the "wrong" coordinates for it.

Unless otherwise stated, we will always assume $\mathbb{R}^n$ to be equipped with the standard smooth structure.

Exercises : Check the following:

(1) Any open subset of a smooth manifold has itself a natural structure as a smooth manifold (by restricting charts).

(2) If $\phi : U \longrightarrow V$ is smooth chart in the atlas of M, then ϕ is a diffeomorphisms from $U \subseteq M$ to $V \subseteq \mathbb{R}^m$, thought of as smooth manifolds in their own right (as in part (1)).

Remarks

(1) It is true (though not so easy to prove) that if a nonempty open subset of $\mathbb{R}^m$ is homeomorphic to an open subset of $\mathbb{R}^n$, then $m = n$. The weaker statement with "homeomorphic" replaced by "diffeomorphic", on the other hand, is simple linear algebra, given that tangent spaces are isomorphic as vector spaces. This means that the dimension of a (topological or smooth) manifold is uniquely determined.

(2) It is natural to ask how the theories of topological manifolds and smooth manifolds relate. As mentioned in the case of manifolds in euclidean space, the situation is simpler in low dimensions. If $m \leq 3$, then any (abstract) topological m-manifold admits a smooth structure. Moreover, if $n \leq 3$, if two smooth m-manifolds are homeomorphic, then they are also diffeomorphic (though of course, not every homeomorphism is a diffeomorphism, even in dimension 1). However, for any $m \geq 4$, there are compact topological m-manifolds which do not admit any smooth structure. There are also pairs of homeomorphic smooth n-manifolds which are not diffeomorphic. In fact, there are 28 pairwise non-diffeomorphic smooth structures on the topological 7-sphere, S^7 (Milnor). There is also a smooth structure on $\mathbb{R}^4$ which is not diffeomorphic to the standard one (Donaldson). Non-standard smooth structures of this type are ofter referred to as "exotic". It is an open problem as to whether the 4-sphere, S^4, admits an exotic smooth structure — the 4-dimensional smooth Poincaré conjecture.

Note that, in example (E4) of Section 2, the circle, S^1, comes equipped with a natural smooth structure. In examples (E12) and (E13), it does not — though of course, we already know that it does admit such a structure. This similarly applies to the 2-sphere respectively in examples (E5) and (E14).

(3) It's not hard to see that the circle is the only compact 1-manifold up to diffeomorphism. Once can also give a complete classification of compact 2-manifolds

up to diffeomorphism. This is not a feasible project in higher dimensions though.

(4) One can define different "categories" of manifolds, by replacing the assumption that transition maps are diffeomophisms, by the assumption that they are X, where "X" could be a range of different properties. Note that transition maps are always homeomorphisms, so if we take $X =$ "homeomorphism", we just recover the notion of a topological manifold. This is the weakest sensible assumption. But we could also take X to be "C^r" for some finite $r \in \mathbb{N}$, or to be "bilipschitz" or "conformal" etc. (if you know what this means). Identifying $\mathbb{R}^{2m}$ with $\mathbb{C}^m$, we could also take X to be "complex analytic" in even (real) dimension, giving rise to the notion of a complex manifold. The key point in the discussion is that X should be closed under inverses and composition. Each of these categories has its own theory associated to it.

(5) One can show that any abstract m-manifold can be embedded in some $\mathbb{R}^n$ (indeed for $n = 2m + 1$). This is the Whitney Embedding Theorem. (In other words, it is diffeomorphic to a manifold in euclidean space, as defined in Section 3). However, such embeddings are often not particularly natural.

We next want to define tangent spaces. For a manifold, M, in $\mathbb{R}^n$, this was defined to be subspace of $\mathbb{R}^n$. In other words, it made reference to the ambient space. However, we saw that a vector gave a means of differentiating smooth functions. That will be the basis of defining the tangent space for abstract manifolds.

So, let M be an (abstract) smooth m-manifold. Write $C^\infty(M)$ for the vector space of smooth functions on M. Write $C^\infty_x(M)$ for the set of local smooth functions at x, and $\mathcal{G}_x(M) = C^\infty_x(M)/\!\sim$ for the vector space of germs at x. These are defined as in Section 3, except now, of course, interpreting "smooth function" as we have defined it for an abstract manifold.

Let v be a linear functional on $\mathcal{G}_x(M)$. Write $v.f = v(f)$. Of course, linearity just means that $v.(f + g) = v.f + v.g$ and $v.(\lambda f) = \lambda v.f$ for all $f, g \in \mathcal{G}_x(M)$ and $\lambda \in \mathbb{R}$. The space of linear functionals is itself a vector space, with addition and scalar multiplication defined by $(v + w).f = v.f + w.f$ and $(\lambda v).f = \lambda v.f$. What we have just described is, of course, just the dual space, $(\mathcal{G}_x(M))^*$.

Definition. Let M be a smooth manifold and $x \in M$. The (new) *tangent space* to M at x is the space of all linear functionals v on $\mathcal{G}_x(M)$ which satisfy:
(L) $v.(fg) = f(x)(v.g) + g(x)(v.f)$ for all $f, g \in \mathcal{G}_x(M)$.

We said "space" because it is easily verified (exercise) that it is a subspace of the dual space, $(\mathcal{G}_x(M))^*$, hence intrinsically a vector space. (For the moment, we know nothing about its dimension.) We will denote it by $\hat{T}_x M$. (The hat is temporary. Once we have related it to our earlier definition, we will omit it.) We will refer to elements of the tangent space as *tangent vectors*.

Exercise: If $U \subseteq M$ is open and $x \in U$, then there is natural identification of $\hat{T}_x U$ with $\hat{T}_x M$. (Recall that U is a manifold in its own right.)

As an example, suppose that $\gamma : (-1, 1) \longrightarrow M$ is a smooth curve with $\gamma(0) = x$. Given $f \in \mathcal{G}_x(M)$, write $v.f = (f \circ \gamma)'(0)$. Thus, $v \in (\mathcal{G}_x(M))^*$. Clearly, $v.f$ only depends on f restricted to an arbitrarily small neighbourhood of x. It follows from the usual product rule (for smooth maps of $\mathbb{R}$) that v satisfies (L), and so $v \in \hat{T}_x M$. We write $\hat{\gamma}'(0) = v$ (again, we will drop the hat later). In fact, we will see that every tangent vector arises in this way (Theorem 6.3).

First, we show that the tangent space is what we expect in the case where our manifold is just $\mathbb{R}^m$. For this, we will need the following, known as "Hadamard's Lemma":

Lemma 6.1. *Let $a \in \mathbb{R}^m$. Then every $f \in C^\infty_a(\mathbb{R}^m)$ has the form*

$$f(x) = f(a) + \sum_{i=1}^{m} (x_i - a_i) g_i(x),$$

for local functions, $g_i \in C^\infty_a(\mathbb{R}^m)$.

Proof. The idea is to apply Taylor's theorem separately to the functions

$$f(a_1, a_2, \dots, a_{i-1}, x_i, x_{1+1}, \dots, x_m) - f(a_1, a_2, \dots, a_{i-1}, a_i, x_{i+1}, \dots, x_m),$$

and then take the sum over i. Details left as an exercise. $\square$

In fact, we perform the construction in a canonical way. Thus, we can choose the maps g_i so that if we apply the same construction to maps with domain U and V which agree on some open set $W \subseteq U \cap V$ containing x, then the corresponding g_i also agree on W.

Note that we necessarily have $g_i(a) = \frac{\partial f}{\partial x_i}(a)$.

Now, for each i, we have an element $d_i \in \hat{T}_a \mathbb{R}^m$ defined by $d_i.f = \frac{\partial f}{\partial x_i}(a)$ for all $f \in C^\infty_a(\mathbb{R}^m)$. This gives rise to a linear map from $\mathbb{R}^m$ to $\hat{T}_a \mathbb{R}^m$ sending the standard basis elements, e_i, to d_i.

Let $\pi_i : \mathbb{R}^m \longrightarrow \mathbb{R}$ be projection to the ith coordinate (that is, $\pi_i(x) = x_i$). Then $d_i.\pi_j = \delta_{ij}$. It follows easily that the d_i are linearly independent. In fact:

Lemma 6.2. *$\{d_1, \dots, d_m\}$ is a basis for $\hat{T}_a \mathbb{R}^m$.*

Proof. It remains to prove that these elements span. So let $v \in \hat{T}_a \mathbb{R}^m$. Write $\lambda_i = v.\pi_i$.

First consider the constant function, 1, on any open $U \subseteq \mathbb{R}^m$. By (L), we have $v.1 = v.(1 \times 1) = 1v.1 + 1v.1 = 2v.1$, so $v.1 = 0$. By linearity, it follows that v kills all (locally) constant functions on $\mathbb{R}^m$.

Now suppose $f \in C_a^\infty(\mathbb{R}^m)$. By Hadamard's Lemma (7.1), we can write

$$f(x) = f(a) + \sum_{i=1}^{m} (\pi_i(x) - \pi_i(a)) g_i(x),$$

and so by (L):

$$\begin{aligned} v.f &= 0 + \sum_{i=1}^{n} (\pi_i(a) - \pi_i(a))(v.g_i) + g_i(a)(v.\pi_i) \\ &= \sum_{i=1}^{m} \lambda_i g_i(a) = \sum_{i=1}^{m} \lambda_i \frac{\partial f}{\partial x_i}(a) = \sum_{i=1}^{m} \lambda_i d_i.f. \end{aligned}$$

Thus, $v = \sum_{i=1}^m \lambda_i d_i$. □

We can now identify $\hat{T}_a\mathbb{R}^m$ with $\mathbb{R}^m$ by sending each d_i to the standard basis element e_i.

To apply this to manifolds, we note, quite generally, that if M and N are smooth manifolds, and $x \in M$, and $\theta : M \longrightarrow N$ is a diffeomorphism, we get a linear map $d_x\theta : \hat{T}_x M \longrightarrow \hat{T}_{\theta(x)} N$ defined by

$$(d_x\theta)(v).f = v.(f \circ \theta)$$

for all $f \in \mathcal{G}_x(N)$. One needs to check (exercise) that this $d_x\theta$ satisfies (L), and so indeed lies in the tangent space to N. Note also, that if θ is a diffeomorphism, then $d_x\theta$ is a linear isomorphism.

Suppose that M is a smooth manifold, and $\phi : U \longrightarrow V$ is a chart of the atlas. We have observed that ϕ is a diffeomorphism between the intrinsic manifolds $U \subseteq M$ and $V \subseteq \mathbb{R}^m$. Moreover, we have natural identifications $\hat{T}_x(U) \equiv \hat{T}_x M$ and $\hat{T}_{\phi(x)} V \equiv \hat{T}_{\phi(x)}(\mathbb{R}^m)$. Furthermore, by Lemma 6.2, there is a natural isomorphism from $\hat{T}_{\phi(x)}(\mathbb{R}^m)$ to $\mathbb{R}^m$.

From this, we can draw some immediate conclusions.

Theorem 6.3. *Let M be a smooth m-manifold and $x \in M$. Then.*
(1) $\dim \hat{T}_x M = m$.
(2) If $v \in \hat{T}_x M$, there is a smooth curve, $\gamma : I \longrightarrow M$ with $\gamma(0) = x$ and $\hat{\gamma}'(0) = v$.
(3) If M is a smooth submanifold of $\mathbb{R}^n$, then there is a natural identification of $\hat{T}_x M$ with $T_x M$ (as defined earlier) so that if $\gamma : I \longrightarrow M$ is a smooth curve with $\gamma(0) = x$, then $\hat{\gamma}'(0)$ gets identified with $\gamma'(0)$ (the usual tangent vector to a curve in $\mathbb{R}^n$).

Proof.
(1) Directly from Lemma 6.2.
(2) Let $\phi : U \longrightarrow V \subseteq \mathbb{R}^m$ be a chart with $x \in U$, so that $d_x\phi : \hat{T}_x M \equiv \hat{T}_x U \longrightarrow \hat{T}_x \mathbb{R}^m \equiv \mathbb{R}^m$ is a linear isomorphism. Let $\delta : I \longrightarrow V$ be a curve with $\delta'(0) = d_x\phi(v)$, and set $\gamma = \phi^{-1} \circ \delta$.
(3) Recall that $T_x M$ was defined via the isomorphism to $T_{\phi x} M$ induced by ϕ. □

In view part (3), we will henceforth drop the hat from the notation for tangent space, T_xM, and from the notation $\gamma'(0)$.

Also note that (in view of (2)), given a smooth map, $\theta : M \longrightarrow N$, one can describe the map $d_x\theta : T_xM \longrightarrow T_{\theta(x)}N$ by saying that it sends $v \in T_xM$ to $(\theta \circ \gamma)'(0)$ for any curve γ in M with $\gamma'(0) = v$.

The tangent bundle

We have defined individual tangent spaces at points of x. In fact, all these tangent spaces can be "bundled" together to form a manifold, TM, called the "tangent bundle". In the case of manifolds in euclidean space, this came for free (Proposition 5.1). Here we have to construct it explicitly.

As a set, TM, is just the formal disjoint union, $TM = \bigsqcup_{x\in M} T_xM$. (More formally, this means that $TM = \bigcup_{x\in M}(\{x\} \times T_xM)$, though as before we will simplify things by identfying $\{x\} \times T_xM$ with T_xM.) We have a surjective map, $p : TM \longrightarrow M$, with $p(T_xM) = \{x\}$ for all $x \in M$. Note that if $U \subseteq M$ is open, then $TU = p^{-1}U = \bigsqcup_{x\in U} T_xM$. We need to define a topology on TM, as well as an atlas for TM. We do the latter first.

Given a chart, $\phi : U \longrightarrow V \subseteq \mathbb{R}^m$ of M, we have a map $\psi : TU \longrightarrow V \times \mathbb{R}^m$ given by $\psi(v) = (\phi x, d_x\phi(v))$, where $x = pv$. In this way, an atlas $\{\phi_\alpha : U_\alpha \longrightarrow V_\alpha\}_{\alpha\in\mathcal{A}}$ for M gives rise to a collection of maps, $\psi_\alpha : TU_\alpha \longrightarrow V_\alpha \times \mathbb{R}^m \subseteq \mathbb{R}^{2m}$. We deem a set, $O \subseteq TM$, to be "open" if $\psi_\alpha(O \cap TU_\alpha)$ is open in $V_\alpha \cap \mathbb{R}^m$ for all $\alpha \in \mathcal{A}$. One can check that this does indeed define a topology on TM. Moreover, the collection of maps $\{\psi_\alpha\}_{\alpha\in\mathcal{A}}$, is a smooth atlas. (The transition maps for TM can be defined in terms of the transition maps on M, which by definition are smooth.) One can also check that the resulting smooth structure does not depend on the atlas for M that we chose. (We could have taken the maximal atlas for M.) It is also easily checked that the map $p : TM \longrightarrow M$ is a smooth submersion.

Definition. The manifold, TM, together with the map $p : TM \longrightarrow M$ is called the *tangent bundle* of M.

This was all rather formal. The main point we want to take from it is that it allows us to speak of a "smooth vector field" on M, that is a section of the tangent bundle, defined in the same way as in Section 5. It can be understood in less abstract way, by the following exercise (with the same notion as in Section 5).

Exercises :

(1) If $\phi : U \longrightarrow M$ is chart, and $\lambda_i : U \longrightarrow \mathbb{R}$ are smooth functions, then $\sum_{i=1}^m \lambda_i \frac{\partial}{\partial x_i}$ is a smooth vector field on U. Conversely, every smooth vector field on U can be uniquely expressed in this form.

(2) If X is a smooth vector field on M, and $f \in C^\infty(M)$, then $Xf \in C^\infty(M)$, where Xf denotes the function given by $Xf(x) = X(x).[f]$, where $[f] \in \mathcal{G}_x(M)$ is the germ of f at x.

The point of these exercises is just to check smoothness in terms of the formal definition.

For future reference, we also note that if $x \in U_\alpha$, then $\psi_\alpha | T_x M : T_x M \longrightarrow \{x\} \times \mathbb{R}^m \equiv \mathbb{R}^m$ is a linear isomorphism.

If $f : M \longrightarrow N$ is a smooth map, we get a smooth map, $f_* : TM \longrightarrow TN$.

The discussion of immersions and submersions in Section 4 now goes through more or less verbatim — except with the new interpretation. In particular, Theorem 4.5 holds as stated. (We conveniently forgot to mention the ambient euclidean space in its statement.)

One can also define embeddings etc. One can check that a manifold, M, in $\mathbb{R}^m$ (as defined in Section 3) is the same as an abstract manifold, M, together with an embedding into $\mathbb{R}^m$ (as defined in this section.)

One can now to on to discuss other things, such as vector fields and orientations, though we postpone this until the next section, since they are more conveniently set in a broader context of "vector bundles".

Lie Groups.

We finish this section with a (very) brief discussion of Lie groups.

Definition. A *Lie group* is a manifold, G, together with smooth maps, $[x \mapsto x^{-1}] : G \longrightarrow G$ and $[(x, y) \mapsto xy] : G \times G \longrightarrow G$, which give G the structure of a group.

We have already seen several examples: $GL(n, \mathbb{R})$, $GL(n, \mathbb{C})$, $SL(n, \mathbb{R})$, $SO(n, \mathbb{R})$, etc.

7. Vector bundles

We will define the notion of a "vector bundle" over a manifold. The idea is to associate to each point in the manifold a vector space of some given dimension (perhaps different from that of the manifold) and assemble these together in a nice smooth way. We will see, for example, that the tangent and normal bundles we have already encountered are examples of such.

Suppose that M is an (abstract) m-manifold. By a *family of vector spaces* over M, we mean a manifold E, together with a smooth map, $p : E \longrightarrow M$, such that for all $x \in M$, $p^{-1}x$ comes equipped with the structure of a vector space. We write $E_x = p^{-1}x$, and refer to E_x as a *fibre* of E. An example of such is $M \times \mathbb{R}^q$, with $p : M \times \mathbb{R}^q \longrightarrow M$ just projection to the first coordinate, so that each fibre is just a copy of $\mathbb{R}^q$. If $U \subseteq M$ is open, write $E|U = p^{-1}U$. Then $p|U : E|U \longrightarrow U$ is a family of vector spaces over U.

Suppose $p' : E' \longrightarrow M$ is another family over M. An *isomorphism* from E to E' is a diffeomorphism, $\psi : E \longrightarrow E'$ with $p' \circ \psi = p$ and such that $\psi|E_x \longrightarrow E'_x$ is a linear isomorphism for all $x \in M$. If such exists, we say that E and E' are *isomorphic*. We say that a family E over M is *trivial* if it is isomorphic to $M \times \mathbb{R}^q$ for some q.

Definition. A *vector bundle* over M is a family of vector spaces, $p : E \longrightarrow M$, such that for all $x \in M$, there is some open $U \subseteq M$, with $x \in U$, and with $E|U$ trivial.

In other words, we can find an atlas, $\{\phi_\alpha : U_\alpha \longrightarrow V_\alpha\}_{\alpha \in \mathcal{A}}$, for M, such that for each α, we have diffeomorphism, $\psi_\alpha : E|U_\alpha \longrightarrow V_\alpha \times \mathbb{R}^q \subseteq \mathbb{R}^{m+q}$, which is a linear isomorphism on each fibre. Note that the maps $\{\psi_\alpha\}_{\alpha \in \mathcal{A}}$ form an atlas for E. We refer to it as a *locally trivialising atlas* for E. It is easily seen from the local structure, that $p : E \longrightarrow M$ is a surjective smooth submersion. Also, each fibre, E_x, is an embedded submanifold, and the vector space operations are smooth. The notion of vector bundle is preserved under isomorphism, so we can also talk about isomorphisms of vector bundles (with the same definition).

Henceforth, the only families of vector spaces we will encounter will be vector bundles (so we can forget about the more general notion).

Examples :

(1) Clearly, $M \times \mathbb{R}^q$ is vector bundle over M.

(2) If M is any smooth manifold, then the tangent bundle, TM, is a vector bundle over M. For a manifold in euclidean space, this follows directly from the proof that TM is a manifold (Proposition 5.1). For an abstract manifold, it follows from the discussion at the end of Section 6. In fact, the two constructions agree, or more precisely, they are isomorphic as vector bundles. So there is no ambiguity in referring to the "tangent bundle" of M.

(3) If $M \subseteq \mathbb{R}^n$ is manifold in $\mathbb{R}^n$, then the normal bundle, $\nu(M, \mathbb{R}^n)$ is a vector bundle. In fact, the proof that $\nu(M, \mathbb{R}^n)$ is a manifold (Propsition 5.3) effectively shows this — the map $\psi : p^{-1}U' \longrightarrow V \times \mathbb{R}^{n-m}$ constructed in the proof restricts to a linear isomorphism on each fibre.

(4) The space B_n, (example (E8) of Section 2) together with the projection to the circle is a fibre bundle. In particular, the Möbius band, B_1 is a fibre bundle over the circle. Note that B_m is isomorphic to B_n if $|m-n|$ is even (exercise). However B_0 cannot be isomorphic to B_1, since they are not diffeomorphic. In particular,

the Möbius band is non-trivial.

(5) TS^1 is trivial. There are a number of ways to see this.

For example, if we take S^1 to be the unit circle in $\mathbb{R}^2$, and define the tangent bundle as in Section 3. Recall that in Section 3, we observed that the map, $f : S^1 \times \mathbb{R} \longrightarrow TS^1$ given by explicitly by $f((x_1, x_2), t) = ((x_1, x_2, -tx_2, tx_1)$ for $x = (x_1, x_2) \in S^1$, is a diffeomorphism. In fact, we see that it is a bundle isomorphism.

It can also be seen more simply, by taking S^1 to be the abstract manifold, $S^1 = \mathbb{R}/\mathbb{Z}$. In this case, each tangent space is canonically identified with $\mathbb{R}$, so we get an identification of TS^1 with $S^1 \times \mathbb{R}$.

Let $p : E \longrightarrow M$ be a vector bundle.

Definition. A *section* of E is a smooth map $s : M \longrightarrow E$ with $p \circ s$ the identity on M.

This is a generalisation of a section of a tangent space, as defined in Section 5. A section of the tangent space is usually referred to as a *vector field* on M. Note that (as for a tangent space) any section will be an embedding of M into E.

For example, any smooth map $f : M \longrightarrow \mathbb{R}^q$, determines a section of $M \times \mathbb{R}^q$, given by $s(x) = (x, f(x))$. The image of the section is the graph of f in $M \times \mathbb{R}^q$.

We say that a section is *non-vanishing* (or *nowhere vanishing*) if $s(x) \neq 0$ for all $x \in M$.

The Möbius band, B, has no non-vanishing section. (Exercise, using the Intermediate Value Theorem.) This gives another (simpler) proof that the Möbius band is a non-trivial bundle.

There is no non-trivial section to TS^2. This is just a statement of the famous "Hairy Ball Theorem". Put another way:

Theorem 7.1. *Every vector field on the 2-sphere, S^2, has at least one zero.*

Proof. (Sketch) We can give a sketch of a proof of this as follows. We use polar coordinates, (θ, ϕ), for S^2. Here $\theta \in \mathbb{R}/2\pi\mathbb{Z}$ is the longitude, and $\phi \in [-\pi/2, \pi/2]$ is the latitude. The points $\phi = \pi/2$ and $\phi = -\pi/2$ correspond respectively to the north and south poles. We identify the circle, S^1, with $\mathbb{R}/2\pi\mathbb{Z}$.

Suppose, for contradiction that v is a nowhere-vanishing vector field on S^2. Given $\phi \in (-\pi/2, \pi/2)$ we define a continuous (in fact smooth) map $h_\phi : S^1 \longrightarrow S^1$ by setting $h_\phi(\theta)$ to be the angle which the tangent vector, $\partial/\partial\theta$, to the the ϕ-latitude at (θ, ϕ), makes with the vector $v(\theta, \phi)$ (taking account of the orientation of S^2). Now, the map h_ϕ has a "degree" $n(\phi)$, associated to it. Informally, this counts the number of times h_ϕ wraps the circle around itself, taking account of orientation. If we imagine walking once around the latitude, it is the number of times we spin around relative to the vector field v, on going once around this latitude. Now very near the north pole (that is, ϕ close to $\pi/2$), $n(\phi)$ must be equal to 1, since v is approximately constant there, and we spin around once in the positive direction. Similarly, near the south pole, $n(\phi)$ will be -1, since we spin around in the negative direction. But now, h_ϕ, varies continuously (in fact smoothly) in ϕ. The degree does not change under continuous deformation (that is, what topologists call a "homotopy"). (This should be intuitively clear, and is not hard to prove, though we won't make this precise here.) Thus, $[\phi \mapsto n(\phi)]$ is constant on all of $(-\pi/2, \pi/2)$ giving a contradiction. □

Remarks : It turns out any odd-dimensional manifold admits a non-vanishing vector field, though we won't prove that here.

An even-dimensional manifold may or may not. For example, the torus does (exercise). However it is the only compact connected orientable 2-manifold which does.

Back to general vector bundles, $p : E \longrightarrow M$. We note:

Lemma 7.2. *A vector bundle, E, is trivial if and only if there are sections $s_1, s_2, \dots, s_q$ such that $\{s_1(x), \dots, s_q(x)\}$ is a basis for E_x for all $x \in M$.*

Proof. If $f : M \times \mathbb{R}^q \longrightarrow E$ is an isomorphism, set $s_i(x) = f(x, e_i)$, where $e_1, \dots, e_q$ is any fixed basis for $\mathbb{R}^q$.

Conversely, if we have such sections, s_i, define a map $f : M \times \mathbb{R}^q \longrightarrow E$ by $f(x, (\lambda_1, \dots, \lambda_q)) = \sum_{i=1}^q \lambda_i s_i(x)$. This is an isomorphism. □

Example : Let $S^n \subseteq \mathbb{R}^{n+1}$ be the unit sphere in $\mathbb{R}^n$. Then the normal bundle, $\nu(S^n, \mathbb{R}^{n+1})$ is trivial: the outward pointing unit normal is a nowhere vanishing section.

Definition. A manifold, M, is *parallelisable* if the tangent bundle TM is trivial.

In view of Lemma 7.2, this is the same as saying that M admits a global frame field. (Recall, from Section 5 that a "frame field" on M is a family of vector fields, v_i, on M, such that $v_1(x), \dots, v_n(x)$ is a basis for T_xM for all $x \in M$. We observed there that frame fields exist locally, but not necessarily globally.)

One can check that any parallelisable manifold is orientable.

Examples :

(1) We have already observed that the circle is parallelisable. (Taking the unit circle in $\mathbb{R}^2$, we can take the tangent at (x_1, x_2) to be $(-x_2, x_1)$.)

(2) By the Hairy Ball Theorem, S^2 is not parallelisable (even though S^2 is orientable).

(3) S^3 is parallelisable. Take $S^3 \subseteq \mathbb{R}^4$ to be the unit sphere. If $x = (x_1, x_2, x_3, x_4) \in S^3$, then $(-x_2, x_1, x_4, -x_3)$, $(-x_3, -x_4, x_1, x_2)$, $(-x_4, x_3, -x_2, x_1)$ defines an orthonormal frame in $T_x S^3$. This follows from the fact these four vectors together form an orthonormal basis for the ambient space $\mathbb{R}^4$, and that the first, (x_1, x_2, x_3, x_4) is normal to S^3, so the last three must lie in the tangent space.
(In fact, it turns out that any orientable 3-manifold is parallelisable, though we won't prove that here.)

(4) S^7 is parallelisable. Take $S^7 \subseteq \mathbb{R}^8$ to be the unit sphere.
If $x = (x_1, x_2, x_3, x_4, x_5, x_6, x_7, x_8) \in S^7$, then
$(-x_2, +x_1, +x_4, -x_3, +x_6, -x_5, -x_8, +x_7)$,
$(-x_3, -x_4, +x_1, +x_2, +x_7, +x_8, -x_5, -x_6)$,
$(-x_4, +x_3, -x_2, +x_1, +x_8, -x_7, +x_6, -x_5)$,
$(-x_5, -x_6, -x_7, -x_8, +x_1, +x_2, +x_3, +x_4)$,
$(-x_6, +x_5, -x_8, +x_7, -x_2, +x_1, -x_4, +x_3)$,
$(-x_7, +x_8, +x_5, -x_6, -x_3, +x_4, +x_1, -x_2)$,
$(-x_8, -x_7, +x_6, +x_5, -x_4, -x_3, +x_2, +x_1)$
defines an orthonormal frame in $T_x S^7$. This follows as for S^3.
(In fact, it turns out that if S^n is parallelisable, then $n \in \{0, 1, 3, 7\}$. So, for example, S^5 is an example of a closed orientable odd-dimensional manifold that is not parallelisable.)

(5) Any Lie group is parallelisable. To see this, let G be a Lie group. Given $x \in G$, let $L_x : G \longrightarrow G$, be left multiplication: $L_x(y) = xy$. This a diffeomorphism of G to itself. It induces a smooth map $(L_x)_* : TG \longrightarrow TG$. In particular, $d_{L_x} : T_1 G \longrightarrow T_x G$ is a linear isomorphism. If $v \in T_1 G$, then setting $v(x) = d_{L_x}(v)$, we get a vector field $[x \mapsto v(x)]$ on G. Indeed, if $v_1, \ldots, v_m$ is any basis for T_G, then the maps $[x \mapsto v_i(x)]$ give us a frame field on G.

Remark : The above examples are related. If we identify $\mathbb{R}^2$ with the complex plane, $\mathbb{C}$, then the vector field $(-x_2, x_1)$ is obtained by multiplying by i: $(x_1 + x_2 i)i = -x_2 + x_1 i$. Similarly, if we identify $\mathbb{R}^4$ with the quaternions, $\mathbb{H}$ (if you know what these are), then the frame on S^3 is obtained by multiplying on the right by i, j, k: $(x_1 + x_2 i + x_3 j + x_4 k)i = -x_2 + x_1 i + x_4 j - x_3 k$ etc. For S^7 the frame field is similarly obtained by identifying $\mathbb{R}^8$ with the Cayley numbers (or "octonians"), $\mathbb{O}$. Thus the fact that these spheres are parallelisable is related to the existence of these number systems, which only exist in dimensions 1, 2, 4, and 8. In fact, the numbers $\mathbb{C}$, $\mathbb{H}$ and $\mathbb{O}$ can be used respectively to give S^1, S^3 and S^7 each the structure of a Lie group. (The fact that algebraic structures of this sort only exist in these dimensions can be proven using topological arguments: it is application of topology to a purely algebraic question.)

Vector fields.

Let X be a vector field on M (i.e. a section of the tangent bundle). Suppose that $f \in C^\infty(M)$; i.e. $f : M \longrightarrow \mathbb{R}$ is a smooth function. We can define a function, $Xf : M \longrightarrow \mathbb{R}$, by setting $Xf(x) = X.[f]$, where $[f] \in \mathcal{G}_x(M)$ is the germ of f at $x \in M$. In fact, we claim Xf is smooth. We can see this using local coordinates. To see that Xf is smooth at x, let x_i be local coordinates in a neighbourhood, U, of x in M. Writing $X = \sum_i \lambda_i \frac{\partial}{\partial x_i}$, where $\lambda_i : U \longrightarrow \mathbb{R}$ is smooth, we have $Xf = \sum_i \lambda_i \frac{\partial f}{\partial x_i}$ (exactly as described in Section 3), and so in particular it is smooth at x. Since x was abritrary, it follows that $Xf \in C^\infty(M)$, as claimed.

In fact, we see easily that $[f \mapsto Xf] : C^\infty(M) \longrightarrow C^\infty(M)$ is a linear map from $C^\infty(M)$ to itself. Moreover, directly from our definition of tangent space, we see that $X(fg) = fXg + gXf$ for all $f, g \in C^\infty(M)$. In fact, we have a converse:

Proposition 7.3. *Suppose that $L : C^\infty(M) \longrightarrow C^\infty(M)$ is a linear map satisfying $L(fg) = fL(g) + gL(f)$ for all $f, g \in C^\infty(M)$. Then there is a unique vector field, X, on M such that $L(f) = Xf$ for all $f \in C^\infty(M)$.*

Proof. For this, we need the fact that, given any $x \in M$, and any germ $h \in \mathcal{G}_x(M)$, there is a function, $f \in C^\infty(M)$ with $h = [f]$. This will be proven in Section 8, see Corollary 8.3. (Of course, by definition of germ, there must be such an f defined on a neighbourhood of x in M. But we want it defined on all of M, which calls for a result about extending smooth functions.)

Suppose that $x \in M$. Given the above fact, L determines a vector $X(x) \in T_xM$, so that $L(f)(x) = X(x).[f]$ where $[f] \in \mathcal{G}_x(M)$ is the germ of f. This gives us a map $X = [x \mapsto X(x)]$ which assigns a tangent vector to each point of M. We need to check that X is smooth.

To this end, we again use local coordinates x_i, on a neighbourhood, U, of x. We write $X = \sum_i \lambda_i \frac{\partial}{\partial x_i}$, where $\lambda_i : U \longrightarrow \mathbb{R}$, and we want to check that the λ_i are smooth. Now, let $f_i \in C^\infty(M)$ be a function that agrees with the coordinate projection, π_i, on a neighbourhood of x (as given by the result stated in the first paragraph). In this neighbourhood, we have $\frac{\partial f_i}{\partial x_i} \equiv 1$, and so $\lambda_i = \lambda_i \frac{\partial f_i}{\partial x_i} = X.[f_i] = L(f_i)$. But, by hypothesis, $L(f_i) \in C^\infty(M)$ is smooth, and so λ_i is smooth in a neighbourhood of x, and so also is X. Since x was arbitrary, it follows that X is smooth everywhere in M. Thus, X is a vector field, as required.

We still need to check uniqueness of X. By linearity, it's enough to show that if $Xf \equiv 0$ for all $f \in C^\infty(M)$, then $f \equiv 0$. Again, use local coordinates. In the notation of the previous paragraph, taking $f = f_i$, we have $\lambda_i = \lambda_i \frac{\partial f_i}{\partial x_i} = Xf_i = 0$. In particular, X is identically 0 in a neighbourhood of x, hence everywhere. $\square$

As an example of this, suppose X, Y are vector fields on M. One checks (exercise) that $[f \mapsto X(Yf) - Y(Xf)]$ satisfies the hypotheses of Proposition 7.3. It follows that there is a unique vector field, denoted $[X, Y]$, on M such that $[X, Y]f = X(Yf) - Y(Xf)$ for all $f \in C^\infty(M)$.

This is called the "Lie Bracket" of X and Y.

Exercise :

(1) If $X = \sum_i \lambda_i \frac{\partial}{\partial x_i}$ and $Y = \sum_i \mu_i \frac{\partial}{\partial x_i}$, in local coordinates, then $[X, Y] = \sum_i \nu_i \frac{\partial}{\partial x_i}$, where $\nu_i = \sum_j \left(\lambda_j \frac{\partial \mu_i}{\partial x_j} - \mu_j \frac{\partial \lambda_i}{\partial x_j} \right)$.

(This gives a way of seeing directly that $[X, Y]$ is indeed a smooth vector field.)

(2) Let G be a Lie group. Describe how one can define a "Lie Bracket", $[.,.]$, on the tangent space, $T_1 G$, at the identity. (Use the fact that any $v \in T_1 G$ determines a vector field, $[x \mapsto v(x)]$, on G as described in Example (5) above.) Check that $[(v, w) \mapsto [v, w]$ is bilinear. What other properties does it satisfy?

Operations on vector bundles.

Various canonical constructions on vector spaces can be applied also to vector bundles. The basic idea is that if we do constructions fibre by fibre, then provided that the operations are all smooth, we would expect them to fit together into a smooth bundle. In principle, this can all be set in a fairly general context, though here we will just deal with the constructions case by case.

The constructions are rather formal, though the idea is simple. If we have a vector space, or a collection of vectors spaces, depending smoothly on a parameter x, then any natural construction we perform to give us another vector space (such as dual, or direct sum, etc.) will also depend smoothly on x.

First, we consider direct sums, which in the context of bundles are called "Whitney sums".

Let $p : E \longrightarrow M$ and $p' : E' \longrightarrow M$ be vector bundles. Let $\hat{E} = \bigsqcup_{x \in M} (E_x \oplus E'_x)$ (as a set), and let $\hat{p} : \hat{E} \longrightarrow M$ be the obvious map. We want to give this the structure of a vector bundle.

Let $\{\phi : U_\alpha \longrightarrow V_\alpha\}_{\alpha \in \mathcal{A}}$ be an atlas for M, which gives rise to locally trivialising atlases, $\{\psi_\alpha : p^{-1} U_\alpha \longrightarrow V_\alpha \times \mathbb{R}^q\}_\alpha$ and $\{\psi'_\alpha : (p')^{-1} U_\alpha \longrightarrow V_\alpha \times \mathbb{R}^r\}_\alpha$ for E and E' respectively. (After intersecting domains of charts, we can assume that we can use the same atlas for M for both cases.)

Given $v \in E_x \oplus E'_x$, write $v = u + u'$, with $u \in E_x$ and $u' \in E'_x$. Let $w \in \mathbb{R}^q$ and $w' \in \mathbb{R}^r$ be respectively the second coordinates of $\psi_\alpha u$ and $\psi'_\alpha u'$. We define $\theta_\alpha : (\hat{p})^{-1} U_\alpha \longrightarrow V_\alpha \times \mathbb{R}^q \times \mathbb{R}^r$ by $\theta_\alpha(v) = (x, w, w')$, where $x = \phi v$.

We deem a set $O \subseteq \hat{E}$ to be open if $\theta_\alpha(O)$ is open in $V_\alpha \times \mathbb{R}^{q+r}$ for all $\alpha \in \mathcal{A}$. One can now check (similarly as in the construction of the tangent space of an abstract manifold) that this gives us a topology on $\hat{E}$, and endows it with the structure of a vector bundle. Moreover, the resulting structure is independent of the choice of atlases for M, E and E'.

The bundle $\hat{E}$ is called the *Whitney sum* of the bundles E and E', and is generally denoted $E \oplus E'$.

Exercise: There are natural embeddings of manifolds, $h : E \longrightarrow E \oplus E'$ and $h' : E' \longrightarrow E \oplus E'$, so that for all $x \in M$, $h(E_x) \oplus h'(E'_x) \cong E_x \oplus E'_x$.

Examples :

(1) The Whitney sum of trivial bundles is trivial.

(2) If $B \longrightarrow S^1$ is the Möbius band, then $B \oplus B$ is trivial (exercise).

(3) Let $M \subseteq \mathbb{R}^n$ be a manifold in $\mathbb{R}^n$. Then $TM \oplus \nu(M, \mathbb{R}^n)$ is trivial. In fact, the each fibre, $T_xM \oplus (T_xM)^\perp$ is just a copy of $\mathbb{R}^n$, and so $TM \oplus \nu(M, \mathbb{R}^n) \equiv M \times \mathbb{R}^n$. For example, we have $TS^2 \oplus \nu(S^2, \mathbb{R}^3) \cong S^2 \times \mathbb{R}^3$. We have seen that $\nu(S^2, \mathbb{R}^3) \cong S^2 \times \mathbb{R}$. This shows that the Whitney sum of a non-trivial bundle with a trivial bundle can be trivial.

We note that, up to isomorphism, the Whitney sum, $E \oplus E'$, can also be constructed as follows. Let $F = \{(v, w) \in E \times E' \mid pv = pw\}$. (Here, $E \times E'$ denotes the usual direct product of E and E'. As such, it is a $(2n + q + r)$+manifold.) There is an obvious map, $F \longrightarrow M$, so that the preimage of $x \in F$ is just a copy of $E_x \times E'_x$. We can naturally identify $E_x \oplus E'_x$ with $E_x \times E'_x$. This gives rise to a natural bijection from $E \oplus E'$ to $F \subseteq E \times E'$. One can check that this is, in fact, an embedding of the manifold $E \oplus E'$ into $E \times E'$. (Note that the notation is potentially confusing: while $E_x \oplus E'_x$ and $E_x \times E'_x$ are essentially the same thing viewed from different perspectives, the manifolds $E \oplus E'$ and $E \times E'$ are definitely different things.)

We can perform a similar construction for dual spaces.

Let $p : E \longrightarrow M$ be a vector bundle. Let $E^* = \bigsqcup_{x \in M} E^*_x$ (where E^*_x is the dual vector space to E_x) and let $p^* : E^* \longrightarrow M$ be the obvious map.

Now any linear isomorphism, $\theta : E_x \longrightarrow \mathbb{R}^q$, canonically gives rise to an isomorphism $\theta^* : E^*_x \longrightarrow \mathbb{R}^q$. This is defined by the condition that $(\theta^* v).(\theta w) = v(w)$ for all $w \in E_x$ and all $v \in E^*_x$. (Here "." denotes dot product on $\mathbb{R}^n$). Let $\{\phi_\alpha : U_\alpha \longrightarrow V_\alpha\}_\alpha$ be an atlas for M, giving rise to a locally trivialising atlas, $\{\psi_\alpha : (E|U_\alpha) \longrightarrow V_\alpha \times \mathbb{R}^q\}_\alpha$ for E. We define maps $\psi^*_\alpha : E^*|U_\alpha \longrightarrow V_\alpha \times \mathbb{R}^q$

by $\psi_\alpha^*(v) = (x, (\psi_\alpha|E_x)^*(v))$ where $x = \phi_\alpha v$, so $v \in E^*$. We use this to define a topology on E^*, similarly as before, and check that this gives it the structure of a vector bundle.

The construction was rather formal. All we really need to remember is that if we have a family of vector spaces depending smoothly on $x \in M$, then the dual spaces can also be assumed to vary smoothly in x.

Cotangent bundle.

The main case of interest to us is the dual to the tangent bundle, $(TM)^*$, more usually denoted T^*M. Its fibres have the form $(T_xM)^*$, again more commonly denoted T_x^*M. An element of T_x^*M called a *covector* at x. A section of T^*M is called a *covector field*, or more commonly a 1-*form*. (The latter terminology will be explained in Section 10.)

Suppose that $f \in C^\infty(M)$ is a smooth function on M. Then f determines an element of T_x^*M, given by the linear map $[v \mapsto v.f] : T_xM \longrightarrow \mathbb{R}$, where $v \in T_xM$. (This only requires f to be defined on a neighbourhood of x. Indeed, it only depends on the germ of f at x.) This element is denoted $df(x)$. The map $[x \mapsto df(x)] : M \longrightarrow T^*M$ is a section of T^*M is a section of T^*M, that is, a 1-form. It is denoted df. Note that we have the product rule: $d(fg) = gdf + fdg$.

If $\phi : U \longrightarrow \mathbb{R}^m$ is a chart, then we get 1-forms, $dx_1, \ldots, dx_m$, defined on U, where x_i is the ith coordinate of ϕ. In the notation for tangent vectors described in Section 3, we see that $dx_i(\frac{\partial}{\partial x_i}) = \delta_{ij}$. We see that, at each $x \in M$, $\{dx_1, \ldots, dx_m\}$ is a basis for T_x^*M. In fact, it is the dual basis to the basis, $\left\{\frac{\partial}{\partial x_1}, \ldots, \frac{\partial}{\partial x_m}\right\}$ for T_xM. Note that if f is as smooth function, we have $df = \frac{\partial f}{\partial x_i} dx_i$.

Note that if ω is a 1-form on X is a vector field, we get a map, denoted $\omega(X)$, from M to $\mathbb{R}$. This is defined pointwise, just applying $\omega(x) \in T_x^*$ to $X(x) \in T_xM$. One checks (from the definion of the smooth structure, that $\omega(X) : M \longrightarrow \mathbb{R}$ is again smooth. In other words $\omega(X) \in C^\infty(M)$.

Suppose that $\{\phi_\alpha : U_\alpha \longrightarrow V_\alpha\}_\alpha$ is an atlas for M. For each α, we have the 1-forms $dx_1^\alpha, \ldots, dx_m^\alpha$. On the overlap, $U_\alpha \cap U_\beta$, these transform according to the rule, $dx_i^\alpha = \sum_{j=1}^m \frac{\partial x_i^\alpha}{\partial x_j^\beta} dx_j^\beta$. Note that $J_{\alpha\beta}(x) = \left(\frac{\partial x_i^\alpha}{\partial x_j^\beta}\right)_{ij}$ is the Jacobian of the transition function $\phi_\alpha \circ \phi_\beta^{-1}$ (evaluated at $\phi_\beta x$). This is the inverse of the Jacobian of the transition function $\phi_\beta \circ \phi_\alpha^{-1}$ (evaluated at $\phi_\alpha x$). This ties in with the corresponding formula or tangent vectors, described in Section 3.

Example : For example, given polar coordinates in the (r, θ) in the euclidean xy-plane, we have 1-forms, dr and $d\theta$. These satisfy the familiar relations

$$dx = \cos\theta \, dr - r \sin\theta \, d\theta$$

and

$$dy = \sin\theta\, dr + r\cos\theta\, d\theta.$$

which we can now make formal sense of.

1-forms can be used to integrate along curves. Suppose that $\gamma : [a, b] \longrightarrow M$ is a smooth curve, and that ω is a 1-form on M. Then we get a smooth function $[t \mapsto (\omega(\gamma(t)))(\gamma'(t))]$, evaluating $\omega(\gamma(t)) \in T^*_{\gamma(t)}M$ at the vector $\gamma'(t) \in T_{\gamma(t)}M$. We write $\int_\gamma \omega = \int_a^b (\omega(\gamma(t)))(\gamma'(t))\, dt$.

Exercise : If f is a smooth function on M, then $\int_\gamma df = f(b) - f(a)$. (Here, "$df$" denotes the 1-form.)

Pull-backs :

Suppose M, N are smooth manifolds and $f : M \longrightarrow N$ is a smooth function. Given a 1-form, ω, on N, we can define a 1-form, η, on M as follows. Given $x \in M$, write $\eta(x)(v) = \omega(f(x))(f_* v)$, where $f_*(v) = (d_x f)(v)$. Thus $\eta(x) \in T^*_x(M)$. This gives a map $[x \longrightarrow \eta(x)] : M \longrightarrow T^*M$. One can check (exercise) that this is smooth, hence a section of T^*M. In other words, η is a 1-form on M.

This is commonly denoted $f^*\omega = \eta$, and called the *pull-back* of η to M.

As an example, suppose $I \subseteq M$ is an open interval, and $\gamma : I \longrightarrow M$ is smooth path. Given a 1-form, ω, on M, we get a 1-form $\gamma^*\omega$ on I, which can be written in the form $\gamma^*\omega = \lambda\, dt$, where t is the $\mathbb{R}$-coordinate, and $\lambda : I \longrightarrow \mathbb{R}$ is a smooth function. One checks (execise) that $\int_\gamma \omega = \int_I \lambda(t)\, dt$ (in the usual sense). In other words, the integral of a 1-form along a path can be equivalently defined by pulling back the 1-form the domain, and integrating in the usual way.

8. Extending smooth functions

This section we give some technical but useful results about extending smooth functions on manifolds. The basic building blocks are so-called "bump functions" which are smooth, and identically zero outside a compact set. For example:

Lemma 8.1. *There is a smooth function, $\theta_0 : \mathbb{R}^n \longrightarrow [0, 1] \subseteq \mathbb{R}$, with $\theta_0(x) = 1$ whenever $||x|| \leq 1$, and with $\theta_0(x) = 0$ whenever $||x|| \geq 2$.*

Proof. We construct θ_0 in stages. First, define $\theta_1 : \mathbb{R} \longrightarrow \mathbb{R}$ by $\theta_1(t) = e^{-1/t}$ for $t > 0$ and $\theta_1(t) = 0$ for $t < 0$. This is smooth (exercise). Now set $\theta_2(t) = \theta_1(t)/(\theta_1(t) + \theta_1(1 - t))$. Thus $\theta_2|(1, \infty) \equiv 1$ and $\theta_2|(-\infty, 0] \equiv 0$. Now set $\theta_0(x) = \theta_2(2 - ||x||)$ for $x \in \mathbb{R}^n$. This is also smooth. $\square$

We have the following corollary for manifolds.

Lemma 8.2. *Let M be an m-manifold, let $W \subseteq M$ be an open subset, and let $x \in W$. Then there is a smooth function $\theta : M \longrightarrow [0,1]$ such that $\theta|(M \setminus W) \equiv 0$ and such that θ is identically 1 on some neighbourhood of x.*

Proof. Let $\phi : U \longrightarrow V \subseteq \mathbb{R}^m$ be a chart with $x \in U$. After postcomposing with a translation of $\mathbb{R}^m$, we can suppose that $0 \in \phi(W \cap U)$, and after postcomposing again by a dilation, we can suppose that the 2-ball, $B(0;2)$, lies in $\phi(W \cap U)$. Now set $\theta(x) = \theta_0(\phi(x))$ for $x \in U$ and set $\theta(x) = 0$ for $x \in M \setminus U$. $\square$

Corollary 8.3. *Suppose that $f : W \longrightarrow \mathbb{R}$ is a smooth function defined on some open set $W \subseteq M$, and suppose that $x \in W$. Then there is a smooth function $g : M \longrightarrow \mathbb{R}$ which agrees with f on some neighbourhood of x in W.*

Proof. Let $\theta : M \longrightarrow [0,1]$ be the function given by Lemma 8.2, and set $g(x) = f(x)\theta(x)$ for $x \in W$ and $g(x) = 0$ for $x \notin W$. $\square$

Recall from Section 6 that $\mathcal{G}_x(M)$ is the space of germs at x. There is a natural linear map from $C^\infty(M)$ (the space of smooth real-valued function on M) to $\mathcal{G}_x(M)$. Corollary 8.3 now tells us that this is surjective. In other words, every germ arises from a global smooth function (defined on all of M).

Remark : Retrospectively, this means that we could have used $C^\infty(M)$ in place of $\mathcal{G}_x(M)$ in Section 6 to define the tangent space, T_xM, to M at x. However, this would be somewhat unnatural, in that it implicates the whole of M in what is really just a local construction.

Partitions of Unity.

We next need a brief digression into general topology.

Let X be a hausdorff topological space. Recall that an *open cover* of X is a collection, $\mathcal{U}$, of open subsets of X with $X = \bigcup \mathcal{U}$. We say that $\mathcal{U}$ is *locally finite* if, for all $x \in X$, there is an open set $O \subseteq X$, with $x \in O$ such that $\{U \in \mathcal{U} \mid O \cap U \neq \emptyset\}$ is finite. (If X is locally compact, this is equivalent to saying that any compact subset of X meets only finitely many elements of $\mathcal{U}$.)

Definition. A *refinement* of $\mathcal{U}$ is an open cover, $\mathcal{V}$, of X such that $(\forall V \in \mathcal{V})(\exists U \in \mathcal{U})(V \subseteq U)$.

Note that if X is locally compact, then any open cover has a refinement all of whose elements are relatively compact. (Recall that a subset $U \subseteq X$ is *relatively compact* if its closure $\bar{U}$ is compact.)

Definition. X is *paracompact* if every open cover has a locally finite refinement.

Clearly any (finite) subcover of an open cover is a (locally finite) refinement. Thus any compact space is paracompact. The following is a bit more challenging:

Exercise (for topologists) : A locally compact second countable hausdorff space is paracompact.

In particular, any manifold is paracompact.

If you don't want to be bothered with the exercise, you can just substitute "paracompact" for "second countable" in the definition of "manifold" in Section 6. (Since the converse is also holds, if one assumes that M is connected.) A lot of books define "manifold" in this way anyway.

The point behind this is that it allows us to pass to covers with nicer properties. For example, if we want, we can find an atlas for M which is locally finite, and such that the domain of each chart has compact closure.

We now get back to manifolds.

Definition. Let $f : M \longrightarrow [0, \infty) \subseteq \mathbb{R}$ be a smooth function. The *support* of f, denoted $\operatorname{supp} f$, is the closure of $\{x \in M \mid f(x) \neq 0\}$.

Let $\{U_\alpha\}_{\alpha\in\mathcal{A}}$ be a locally finite open cover of M.

Definition. A *partition of unity subordinate to* $\{U_\alpha\}_\alpha$ is a collection of smooth functions, $\rho_\alpha : M \longrightarrow [0, 1]$, indexed by $\alpha \in \mathcal{A}$, with $\operatorname{supp} \rho_\alpha \subseteq U_\alpha$ for all $\alpha \in \mathcal{A}$, and such that $\sum_{\alpha\in\mathcal{A}} \rho_\alpha(x) = 1$ for all $x \in \mathcal{M}$.

Note that this is just a finite sum, since $\rho_\alpha(x) = 0$ for all but finitely many $\alpha \in \mathcal{A}$.

We now have all the ingredients to prove the following:

Theorem 8.4. *Any locally finite open cover of M has a subordinate partition of unity.*

Proof. The basic idea is quite simple, though the details are a bit technical.

First note that it's enough to find such maps so that $\sum_{\alpha\in\mathcal{A}} \rho_\alpha(x) > 0$ for all $x \in M$ — since we can just renormalise, that is, replace ρ_α by $\rho_\alpha/(\sum_{\alpha\in\mathcal{A}} \rho_\alpha)$.

Second, note that we can assume that each $\bar{U}_\alpha$ is compact. (Since we can always find a locally finite subcover of, say $\{U'_\beta\}_{\beta\in\mathcal{B}}$, of $\{U_\alpha\}_{\alpha\in\mathcal{A}}$, by relatively compact sets U'_β.) Given any $\beta \in \mathcal{B}$, choose $\alpha(\beta) \in \mathcal{A}$ with $U'_\beta \subseteq U_{\alpha(\beta)}$. If we have a partition of unity, $\{\rho'_\beta\}_\beta$, subordinate to $\{U'_\beta\}_\beta$, then set $\rho_\alpha = \sum\{\rho'_\beta \mid \alpha(\beta) = \alpha\}$. After renormalising, this gives a partition of unity for the original cover, $\{U_\alpha\}_\alpha$.)

So let's suppose that each $\bar{U}_\alpha$ is compact. Given any $p \in M$, choose some $\alpha(p) \in \mathcal{A}$ with $p \in U_{\alpha(p)}$. By Lemma 8.2, we can find an open set $O_p \subseteq U_{\alpha(p)}$, with $p \in O_p$, together with a smooth function, $\theta_p : M \longrightarrow [0, 1]$ with $\sup \theta_p \subseteq U_{\alpha(p)}$ and with $\theta_p|O_p \equiv 1$. Now $\{O_p\}_{p\in M}$ is an open cover of M, so by paracompactness, there is a locally finite subcover, $\mathcal{V}$, say. For each $V \in \mathcal{V}$, choose $p(V) \in M$ with $V \subseteq O_{p(V)}$. Let $P = \{p(V) \mid V \in \mathcal{V}\}$. Given $\alpha \in \mathcal{A}$, let $P_\alpha = \{p \in P \mid \alpha(p) = \alpha\} \subseteq P \cap U_\alpha$. We claim that each P_α is finite. To see this, let $K = \bigcup\{\bar{U}_\beta \mid U_\alpha \cap U_\beta \neq \emptyset\}$. This is compact (since there only finitely many such β). Now, if $p = p(V) \in P_\alpha$, then $V \subseteq K$ (since $V \subseteq O_p \subseteq U_{\alpha(p)}$, and $p \in U_\alpha \cap U_{\alpha(p)}$, so $U_\alpha \cap U_{\alpha(p)} \neq \emptyset$ and so

$U_{\alpha(p)} \subseteq K$). Since $\mathcal{V}$ is a locally finite cover of M, hence of K, there are only finitely many such V, proving the claim. We now set $\rho_\alpha = \sum_{p \in P_\alpha} \theta_p$. We now have $\sum_{\alpha \in \mathcal{A}} \rho_\alpha(x) = \sum_{p \in P} \theta_p(x) > 0$ for all $x \in M$ (since there is some V with $x \in V \subseteq O_{p(V)}$, so $\theta_{p(V)}(x) = 1$). $\square$

For future reference, we note that if $\{U_\alpha\}_{\alpha \in \mathcal{A}}$ and $\{U'_\beta\}_{\beta \in \mathcal{B}}$ are open covers of M, then so is $\{U_\alpha \cap U'_\beta\}_{(\alpha,\beta) \in \mathcal{A} \times \mathcal{B}}$. If $\{U_\alpha\}_\alpha$ and $\{U'_\beta\}_\beta$ are both locally finite, then so is $\{U_\alpha \cap U'_\beta\}_{\alpha,\beta}$. If $\{\rho_\alpha\}_\alpha$ and $\{\rho'_\beta\}_\beta$ are respectively partitions of unity subordinate to $\{U_\alpha\}_\alpha$ and $\{U'_\beta\}_\beta$, then $\{\rho_\alpha \rho'_\beta\}_{\alpha,\beta}$ is a partition of unity subordinate to $\{U_\alpha \cap U'_\beta\}_{\alpha,\beta}$.

Riemannian manifolds.

As an example of an application of partitions of unity, we consider the existence of riemannian metrics (Theorem 8.5 below). (Some further applications will be described in Section 10.) First, we need to say what a "riemannian metric" is.

Let $p : TM \longrightarrow M$ be the tangent bundle to M. Let $P = \{(v,w) \in TM \times TM \mid pv = pw\}$. In other words, it consists of pairs of tangent vectors based at the same point of M. Clearly, there is a natural map $P \longrightarrow M$, and as discussed in Section 7, and we can identify P with the Whitney sum $TM \oplus TM$. In particular, P is a smooth $(3m)$-manifold, so it makes sense to speak of a function from P to $\mathbb{R}$ as being smooth.

Definition. A *riemannian metric* on M is a smooth map $f : P \longrightarrow \mathbb{R}$ such that for all $x \in M$, the restriction of f to the fibre $T_xM \times T_xM$, is an inner product on T_xM.

(Of course, this is not a "metric" in the standard metric space sense, though as we will note, it does give rise to one.)

Given $v, w \in T_xM$, we usually denote $f((v,w))$ as $\langle v, w \rangle$ — the usual notation for an inner product. If $v \in T_xM$, we write $||v|| = \sqrt{\langle v, v \rangle}$. This is the induced norm on T_xM. In particular, a riemannian metric gives us a way of measuring norms of tangent vectors in a nice smooth way.

If $M \subseteq \mathbb{R}^n$ is a submanifold of $\mathbb{R}^n$, then it already comes equipped with a riemannian metric. Recall that in this case, T_xM is identified with a subspace of $\mathbb{R}^n$, and we can simply take $\langle v, w \rangle = v.w$, to be the restriction of the dot product on $\mathbb{R}^n$.

Thus, $\mathbb{R}^n$ is itself a riemannian manifold. So is also S^n is a natural way. If $m = 2$ and $\mathbb{R}^3$, the riemannian metric is essentially the same thing as the "first fundamental form" on the surface $M \subseteq \mathbb{R}^3$.

In general, an abstract manifold does not come already equipped with a riemannian metric. However, we have:

Theorem 8.5. *Every manifold admits a riemannian metric.*

Proof. Let $\{\phi_\alpha : U_\alpha \longrightarrow V_\alpha\}_{\alpha\in\mathcal{A}}$ be an atlas for M, giving rise to a trivialising atlas, $\{\psi_\alpha\}_\alpha$ for TM. By paracompactness of M, we can assume $\{U_\alpha\}_\alpha$ to be locally finite. Let $\{\rho_\alpha\}_\alpha$ be a partition of unity on M subordinate to $\{U_\alpha\}_\alpha$, as given by Theorem 8.4.

Given $\alpha \in \mathcal{A}$, and $v, w \in T_xM$ with $x \in U_\alpha$, we set $\langle v, w\rangle_\alpha = (\psi_\alpha v).(\psi_\alpha w)$. We now set $\langle v, w\rangle = \sum_{\alpha\in\mathcal{A}} \rho_\alpha(x)\langle v, w\rangle_\alpha$. This is smooth, and its restriction to each tangent space is an inner product (since a positive linear combination of inner products is an inner product). □

In other words, we use the partition of unity to patch together inner products induced from the dot products given by charts. This is a typical application of this principle.

Finally, what are riemannian metrics good for?

They arise naturally in many different contexts. They can be also used to prove things about manifolds.

For example, if $\gamma : [a, b] \longrightarrow M$ is a smooth curve, we can its riemannian length, $\operatorname{length}\gamma$ as $\int_a^b ||\gamma'(t)||\, dt$, where $\gamma'(t) \in T_{\gamma(t)}M$ is its tangent, as defined in Section 6. (Note that if $M \subseteq \mathbb{R}^n$ is an embedded manifold, with the induced riemannian metric from $\mathbb{R}^n$, then this is the usual length in $\mathbb{R}^n$.)

One can go on to show that if M is connected, then any two points, $x, y \in M$, can be joined by a smooth path (a smooth version of Lemma 3.5). One can then define $d(x, y)$ to be the infimum of the lengths of all such paths. It's not to hard to see that d is a metric on M (in the traditional metric space sense), which induces the original topology. It then follows (by Theorem 8.5) that any (connected) riemannian manifold is metrisable. We won't however go into the details of that here.

9. Manifolds with boundary

This will be a brief section. We will summarise the basic theory of "manifolds" come which come with a boundary. Such spaces arise naturally, and are also an important tool in developing the theory of manifolds as we have so far described them.

Write $H^m = \{(x_1, \ldots, x_n) \in \mathbb{R}^m \mid x_n \geq 0\}$, and $\partial H^m = \mathbb{R}^{m-1} \times \{0\} \subseteq H^m$.

We now proceed to define a *manifold with boundary* exactly as for "manifold" in Section 6, except with H^m now replacing $\mathbb{R}^m$ as the model space. Thus, M is a second countable hausdorff topological space, which an atlas of charts, $\phi_\alpha : U_\alpha \longrightarrow V_\alpha$, indexed by $\alpha \in \mathcal{A}$, where V_α is now an open set in H^m. To get the smooth structure, we insist that the transition maps, $\phi_\beta \circ \phi_\alpha^{-1} : \phi_\alpha(U_\alpha \cap U_\beta) \longrightarrow \phi_\beta(U_\alpha \cap U_\beta)$, are all smooth in the sense defined in Section 3 for subsets of $\mathbb{R}^m$ (that is, that they extend to smooth functions defined on an open neighbourhood in $\mathbb{R}^m$).

The following is a basic observation:

Lemma 9.1. *Suppose that $x \in U_\alpha \cap U_\beta$. Then $\phi_\alpha x \in \partial H^m$ if and only if $\phi_\beta x \in \partial H^m$.*

This boils down to proving the following:

Lemma 9.2. *Let $\theta : U \longrightarrow V$ be a diffeomorphism between open subsets, $U, V \subseteq H^m$. Then $\theta(U \cap H^m) = V \cap \partial H^m$.*

We leave the proof as an exercise (using the inverse function theorem).
(The statement is also true for homeomorphisms, but much harder to prove.)

Definition. The boundary of M, denoted ∂M, is the set of $x \in M$ such that $\phi_\alpha x \in \partial H$ for some (hence every) chart $\phi_\alpha : U_\alpha \longrightarrow V_\alpha$, of M, with $x \in U_\alpha$.
The *interior* of M, denoted $\operatorname{int} M$, is $M \setminus \partial M$.

Note : This terminology is specific to manifolds, and in general differs from the standard terminology for topological spaces. (For example, the usual topological interior of M as a subset of M is of course just M itself!)

Exercise : $\operatorname{int} M$ is canonically an m-manifold (without boundary) and ∂M is canonically an $(m-1)$-manifold (provided it is non-empty).

We have allowed ∂M to be the empty set, which we decided should not be considered a manifold. Hence the somewhat artificial clause in the above statement.

We can now go on to define smooth maps and diffeomorphisms in the same way as before.

Examples :

(1) H^m is a manifold with boundary, $\partial H^m = \partial H^m$ (as already defined).

(2) Any interval in $\mathbb{R}$ is a 1-manifold with boundary. In fact, any 1-manfold with boundary is diffeomorphic to exactly one of S^1, $\mathbb{R}$, $[0, 1]$ or $[0, \infty)$.

(3) The unit ball, $B(0; 1)$, in $\mathbb{R}^n$ is a manifold with boundary, $\partial B(0; 1) = S^{n-1}$. (This is also the topological boundary in this case.)

One can now go on to discuss tangent spaces and such.

Note that if M is oriented, then we can get an orientation on ∂M, simply by restricting the charts in an oriented atlas to ∂M. In fact (for reasons that become clear in Section 11), we adopt the convention that the "orientation induced" on ∂M is given by such a retriction in the case where m is even, whereas it is the opposite orientation when m is odd.

10. Differential forms and integration

First, we need some linear algebra. For simplicity, we assume all vector spaces to be finite dimensional.

Given vector spaces, E, F, and $p \in \mathbb{N}$, $p \geq 1$, write $A(E^p, F)$ for the vector space of alternating p-linear maps from $E^p = E \times \cdots \times E \longrightarrow F$. (So that $A(E^1, F) = L(E, F)$ is just the space of linear maps.) We want to define an "exterior product" which allows us to convert any alternating multilinear map into linear one. It is based on the following formal observation:

Lemma 10.1. *Given E, p, there is a vector space, V, together with an surjective map $\mu \in A(E^p, V)$, such that if $\theta \in A(E^p, F)$, there is a linear map $\hat{\theta} \in L(V, F)$ such that $\theta = \hat{\theta} \circ \mu$.*

Proof. Let $V = (A(E^p, \mathbb{R}))^*$ — the dual space to $A(E^p, \mathbb{R})$. Define $\mu : E^p \longrightarrow V$ by $\mu(v)(\phi) = \phi(v)$, where $v \in E^p$ and $\phi \in A(E^p, \mathbb{R})$. Note that μ is alternating in the v_i, since ϕ is.

We first check that the conclusion holds when $F = \mathbb{R}$. To see this, let $\theta \in A(E^p, \mathbb{R})$. Define $\hat{\theta} : V \longrightarrow \mathbb{R}$ by $\hat{\theta}(f) = f(\theta)$, for all $f \in V$ (so $f : A(E^p, \mathbb{R}) \longrightarrow \mathbb{R}$). Now, if $v \in E^p$, we have $\hat{\theta}(\mu(v)) = (\mu(v))(\theta) = \theta(v)$, and so $\hat{\theta} \circ \mu = \theta$, as required.

Now a general finite dimensional vector space is isomorphic to $\mathbb{R}^d$ for some d, so we can just apply the construction of the previous paragraph to each coordinate separately.

One can also show that μ is surjective. But to save us the bother, we could just replace V by the image of $\mu(E^p) \subseteq V$, which will work just as well. $\square$

Note that the map, $\hat{\theta}$, obtained from θ is unique. Also, if $\mu' : E^p \longrightarrow V'$ is another such map satisfying the same property, then we get unique linear maps $\hat{\mu}' : V \longrightarrow V'$ and $\hat{\mu} : V' \longrightarrow V$, with $\mu' = \hat{\mu}' \circ \mu$ and $\mu = \hat{\mu} \circ \mu'$. These are inverse isomorphisms.

It follows that V is unique up to isomorphism (respecting μ). Henceforth, we write $V = \Lambda^p E$ and given $v_1, v_2, \ldots, v_p \in E$, we write

$$v_1 \wedge v_2 \wedge \cdots \wedge v_p = \mu(v_1, v_2, \ldots, v_p).$$

Definition. $\Lambda^p E$ is the *p'th exterior power* of E.

Note that $\Lambda^1 E$ is naturally identified with E, via the map μ which we can consider to be the identity. Also, $\Lambda^p E = \{0\}$ for $p > n$.

The above was all very formal, but there is a much more concrete way of understanding what $\Lambda^p E$ looks like.

Let $e_1, e_2, \ldots, e_m$ be a basis for E. Now since μ is surjective, $\Lambda^p E$ is spanned by elements of the form $e_{i_1} \wedge e_{i_2} \wedge \cdots \wedge e_{i_p}$, where $i_k \in I(m) = \{1, \ldots, m\}$ (since E^p is spanned by elements of the form $(e_{i_1}, \ldots, e_{i_p})$). Now, if two of the indices are equal, then this element is 0 (since the map μ is alternating). Thus, we can assume that all the i_k are distinct. Also, by swapping any two of the indices, the result changes by a factor of -1 (again, since μ is alternating). By performing a sequence of such transpositions, we can put all the indices in order, so we get $e_{i_1} \wedge \cdots \wedge e_{i_p} = \pm e_{j_1} \wedge \cdots \wedge e_{j_p}$, where $\{j_1, \ldots, j_p\} = \{i_1, \ldots, i_p\}$, and $j_1 < j_2 < \cdots < j_p$. Note that the sign is given by the signature of the permutation of the indices need to achieve this.

Now, write $\mathcal{I}(m,p)$ for the set of subsets of $I(m) = \{1, \ldots, m\}$ of cardinality p. Given $I \in \mathcal{I}(m,p)$, write $\mathbf{e}_I = e_{i_1} \wedge \cdots \wedge e_{i_p}$, where $I = \{i_1, \ldots, i_p\}$ and $i_1 < i_2 < \cdots < i_p$. We claim:

Lemma 10.2. *$\{\mathbf{e}_I \mid I \in \mathcal{I}(m,p)\}$ is a basis for $\Lambda^p E$.*

Proof. We have already seen that the $\mathbf{e}_I$ span $\Lambda^p E$. We need to show that they are linearly independent.

Now, given any $I \in \mathcal{I}(m,p)$, there is an alternating map, $\theta_I \in A(E^p, \mathbb{R})$, with $\theta_I(e_{i_1}, \ldots, e_{i_p}) = 1$, and with $\theta_I(e_{j_1}, \ldots, e_{j_p}) = 0$ if $\{j_1, \ldots, j_p\} \neq I$. (Given $v_1, \ldots, v_p \in E$, write $v_i = \sum_{j=1}^m a_{ij} e_j$ where $a_{ij} \in \mathbb{R}$, and set $\theta_I(v_1, \ldots, v_p)$ to be the determinant of the $p \times p$ matrix $(a_{ij})_{ij}$, where i ranges over $\{1, \ldots, p\}$, and j ranges over I.) Now θ_I gives rise to a map $\hat{\theta}_I : \Lambda^p E \longrightarrow \mathbb{R}$, with $\hat{\theta}(\mathbf{e}_I) = 1$ and $\hat{\theta}(\mathbf{e}_J) = 0$ for all $J \neq I$. It follows that the $\mathbf{e}_I$ are linearly independent, as claimed. $\square$

It follows that $\dim \Lambda^p E = |\mathcal{I}(m,p)| = \binom{m}{p}$.

Exercise : Show that $\dim A(E^p, \mathbb{R}) = \binom{m}{p}$. Given that, by our construction, $\Lambda^p E \subseteq (A(E^p, \mathbb{R}))^*$, deduce that, in fact, $\Lambda^p E = (A(E^p, \mathbb{R}))^*$.

A particular case is when $p = m$, so the dimension is 1. In this case, the single element $\mathbf{e}_{I(m)} = e_1 \wedge \cdots \wedge e_p$ serves as a basis.

What effect does changing the basis of E have on the basis of $\Lambda^p E$?

First, consider the case where $p = m$. Let $v_1, \ldots, v_m \in E$ be (for the moment) any set of m elements of E. Write $v_i = \sum_{j=1}^m a_{ij} e_j$, so $A = (a_{ij})_{ij}$ is an $m \times m$ matrix. We know that $v_{I(m)} = v_1 \wedge \cdots \wedge v_m$ must be some multiple of $\mathbf{e}_I$. To find this multiple, imagine expanding the expression $(\sum_j a_{1j} e_j) \wedge \cdots \wedge (\sum_j a_{mj} e_j)$. The only terms that are non-zero are those which select a complete set of e_i in some order. In other words, each is some multiple of $e_{\pi^{-1}(1)} \wedge \cdots \wedge e_{\pi^{-1}(m)}$ for some permutation, π, of $I(m)$. Putting these indices back in order, this expression is equal to $\mathrm{sig}(\pi)\mathbf{e}_I$, where $\mathrm{sig}(\pi)$ denotes the signature of π. The coefficient of $\mathbf{e}_I$ is thus $\mathrm{sig}(\pi) a_{1\pi(1)} \cdots a_{m\pi(m)}$. Each permutation occurs exactly once, so we now sum over all permutations and we get $\sum_\pi \mathrm{sig}(\pi) \prod_{i=1}^m a_{i\pi(i)}$, which is precisely the determinant, $\det(A)$. In other words, $v_{I(m)} = (\det(A))\mathbf{e}_{I(m)}$. Note that if v_i is a basis, then A is invertible, so the factor, $\det(A)$ is non-zero, as expected.

In $\mathbb{R}^m$, this has a nice geometric interpretation. Recall that $\det(A)$ is the signed volume of the parallelepiped with edges given by the vectors v_i (that is with vertices of the form $\sum_i \epsilon_i v_i$, where $\epsilon_i \in \{0,1\}$). By "signed volume" we mean that it is positive if $v_1, \dots, v_m$ is a positively oriented basis, and negative if it is a negatively oriented basis, and 0 if it does not span (so that the parallelepided is degenerate). Thus, we can think of $v_1 \wedge \cdots \wedge v_m$ as measuring the volume of this parallelepiped (relative to the unit cube, given by the standard basis, which is deemed to have volume 1).

In general, if $p \neq m$, then these things transform by a similar, though more complicated formula. It can be shown by a similar argument as in the case where $p = m$. Since we will not make explicit use of it, we leave the following as an exercise:

Exercise : Let $v_1, \dots, v_m \in E$, and define A as before. Given $I, J \in \mathcal{I}(m,p)$, let A_{IJ} be the $p \times p$ matrix $(a_{ij})_{ij}$ as i ranges over I and j ranges over J. Let $\Delta_{IJ} = \det A_{IJ}$. Then $v_I = \sum_{J \in \mathcal{I}(m,p)} \epsilon_{IJ} \Delta_{IJ} \mathbf{e}_J$, where ϵ_{IJ} is the minimal number of transpositions needed to shift all the indices I to J.

Given $p, q \in \mathbb{N}$ with $p, q \geq 1$, we can define a bilinear "wedge product", $[(\eta, \zeta) \mapsto \eta \wedge \zeta] : \Lambda^p E \times \Lambda^q E \longrightarrow \Lambda^{p+q} E$. To do this, choose a basis $e_1, \dots, e_m$, for E. Given indices $i_1 < \cdots < i_p$ and $j_1 < \cdots < j_q$ in $I(m)$, we set

$$(e_{i_1} \wedge \cdots \wedge e_{i_p}) \wedge (e_{j_1} \wedge \cdots \wedge e_{j_q}) = e_{i_1} \wedge \cdots \wedge e_{i_p} \wedge e_{j_1} \wedge \cdots \wedge e_{j_q}.$$

(Note the result is will be non-zero precisely when $I = \{i_1, \cdots, i_p\}$ and $J = \{j_1, \dots, j_q\}$ are disjoint, and then $\mathbf{e}_I \wedge \mathbf{e}_J = \pm \mathbf{e}_{I \cup J}$.) We need to check that the resulting map does not, in fact, depend on the basis $\{e_i\}_i$ we choose.

In fact, we claim that if $x_1, \dots, x_p, y_1, \dots, y_q$ are any elements of E, then

$$(x_1 \wedge \cdots \wedge x_p) \wedge (y_1 \wedge \cdots \wedge y_q) = x_1 \wedge \cdots \wedge x_p \wedge y_1 \wedge \cdots \wedge y_q.$$

To see this, first note that the formula defining $\wedge$ remains valid if we permute independently the indices $\{i_k\}_k$ and $\{j_l\}_l$ (since any transposition multiplies both sides by -1). Also, both sides are multilinear in the entries, so the statement follows by writing the x_i's and y_j's in terms of the e_k's. In particular, we see that the same formula holds if we choose a different basis.

Exercises :

(1) If $\eta \in \Lambda^p E$, $\zeta \in \Lambda^q E$ and $\omega \in \Lambda^r E$, then $(\eta \wedge \zeta) \wedge \omega = \eta \wedge (\zeta \wedge \omega)$.
This means that we can just denote the result by $\eta \wedge \zeta \wedge \omega$.

(2) If $\eta \in \Lambda^p E$, $\zeta \in \Lambda^q E$, then $\eta \wedge \zeta = (-1)^{pq} \zeta \wedge \eta$.
In particular, if $p = q$ is odd, then $\eta \wedge \eta = 0$.
Let $\eta = (e_1 \wedge e_2) + (e_3 \wedge e_4) \in \Lambda^2 \mathbb{R}^4$. Then $\eta \wedge \eta \neq 0$.

Note : This notation is consistent. Recall that we can identify E with $\Lambda^1 E$. In this case, if $v_i \in E$, the element $v_1 \wedge \cdots \wedge v_p \in \Lambda^p E$ as originally defined is the same as the wedge product of the v_i thought of as elements of $\Lambda^1 E$.

Example : The case of dimension 3 is special in that $\dim \Lambda^2 E = 3 = \dim E$.

Let e_1, e_2, e_3 be the standard basis for $\mathbb{R}^3$. This gives us basis, $e_1 \wedge e_2$, $e_2 \wedge e_3$, $e_3 \wedge e_1$, for $\Lambda^2 \mathbb{R}^3$. (We write $e_3 \wedge e_1$ instead of $e_1 \wedge e_3$ to simplify the sign conventions.) We have a linear isomorphism, $\theta : \Lambda^2 \mathbb{R}^3 \longrightarrow \mathbb{R}^3$, given by $\theta(e_1 \wedge e_2) = e_3$, $\theta(e_2 \wedge e_3) = e_1$ and $\theta(e_3 \wedge e_1) = e_2$. Note that in all cases, $\theta(e_i \wedge e_j) = e_i \times e_j$, where $\times$ denotes the usual cross product on $\mathbb{R}^3$. Since both sides are bilinear, it follows that $\theta(v \wedge w) = v \times w$ for all $v, w \in \mathbb{R}^3$. (Thus, if we identify $\Lambda^2 \mathbb{R}^3$ with $\mathbb{R}^3$ via θ, then the wedge product is the same as the cross product. Indeed, the cross product is often denoted by $\wedge$.)

There is also a connection with triple products. If $v, w, z \in \mathbb{R}^3$, then $v \wedge w \wedge z = \lambda e_1 \wedge e_2 \wedge e_3$, where λ is the triple product $(v \times w).z$.

The above identification of $\Lambda^2 \mathbb{R}^3$ with $\mathbb{R}^3$ only really requires the dot product and orientation on $\mathbb{R}^3$: one can check that any positively oriented orthonormal frame would have given rise to same identification. Indeed, the same construction works for any 3-dimensional oriented inner product space. However, for a general 3-dimensional vector space, E, there is no canonical identification of $\Lambda^2 E$ with E, so cross products do not make sense. Of course, it is also a phenomenon special to 3 dimensions.

Suppose that E and E' are finite dimensional vector spaces, and that $\phi : E \longrightarrow E'$ is a linear map. This gives rise to a multilinear map, $\phi^p : E^p \longrightarrow (E')^p$, just taking the direct products, and hence in turn to an alternating map, $\mu' \circ \phi^p : E^p \longrightarrow \Lambda^p E'$, where $\mu' : (E')^p \longrightarrow \Lambda^p E'$ is the exterior product of E'. By the defining property of $\Lambda^p E'$, we get a linear map, $\Lambda^p \phi : \Lambda^p E \longrightarrow \Lambda^p E'$, such that $(\Lambda^p \phi) \circ \mu = \mu' \circ \phi^p$.

Again, this is very formal, but it can be described more simply.

Note that if $x_1, \ldots, x_p \in E$, then

$$(\Lambda^p \phi)(x_1 \wedge \cdots \wedge x_p) = (\phi x_1) \wedge \cdots \wedge (\phi x_p).$$

In particular, if $e_1, \ldots, e_m$ is a basis for E, then we can describe $\Lambda^p \phi$ by setting

$$(\Lambda^p \phi)(e_I) = (\Lambda^p \phi)(e_{i_1} \wedge \cdots \wedge e_{i_p}) = (\phi e_{i_1}) \wedge \cdots \wedge (\phi e_{i_p}).$$

on the basis elements $(e_I)_I$ of $\Lambda^p E$, and extending linearly.

Exercise : Check explicitly that this is well defined.

We now finally get back to manifolds.

Let M be an m-manifold. Given $x \in M$, we can form the p'th exterior power, $\Lambda^p T_x^* M$ of the cotangent space $T_x^* M$. We can assemble these together into a vector bundle, denoted $\Lambda^p T^* M$. (In other words, $(\Lambda^p T^* M)_x = \Lambda^p T^*_x M$.) This can be achieved by a similar construction as for Whitney sums and dual bundles. We won't give details. We only really need to note that allows us talk about sections, that is smooth fields of exterior products. More precisely, they can be written simply in local coordinates, as we will see shortly. Note that $\Lambda^1 T^* M \equiv T^* M$.

Definition. A *p-form* on M is a section of the bundle $\Lambda^p T^* M$.

That is, for each $x \in M$, we have $\omega(x) \in \Lambda^p T^*_x N$, which varies in a nice smooth way. Note that a 1-form is the same a "covector field". (The latter term is rarely used.)

Given a p-form, η, and a q-form, ζ, we can form the $(p+q)$-form $\eta \wedge \zeta$. This is defined pointwise: $(\eta \wedge \zeta)(x) = \eta(x) \wedge \zeta(x)$. We have seen that $\eta \wedge \zeta = (-1)^{pq} \zeta \wedge \eta$, and that $(\eta \wedge \zeta) \wedge \omega = \eta \wedge (\zeta \wedge \omega)$ (so we can drop brackets from the notation).

Let $\phi : U \longrightarrow M$ be a chart, with local coordinates, $x_1, \ldots, x_m$. We have locally defined 1-forms, $dx_1, \ldots, dx_m$, which (evaluated at any given point of M) give a basis for the cotangent space. Given $I \in \mathcal{I}(m,p)$, we will write $\mathbf{d}x_I = dx_{i_1} \wedge \cdots \wedge dx_{i_p}$, where $I = \{i_1, \ldots, i_p\}$ and $i_1 < \cdots < i_p$. (This is not standard notation. We use bold font $\mathbf{d}$ in the above to avoid potential ambiguity of notation when we come to describe exterior derivatives in Section 11.) Thus $\{\mathbf{d}x_I\}_{I \in \mathcal{I}(m,p)}$ (evaluated at any point of M) gives us a basis for the exterior product. It follows (using Lemma 7.2) that any p-form, ω, on U can be uniquely written in the form $\omega = \sum_{I \in \mathcal{I}(m,p)} \lambda_I \mathbf{d}x_I$, where each $\lambda_I : U \longrightarrow \mathbb{R}$ is a smooth function. (This is all we really need from the bundle structure of $\Lambda^p T^* M$.) In particular, if $p = m$, any m-form locally looks like $\lambda dx_1 \wedge \cdots \wedge dx_m$, where λ is a smooth function.

Pull-backs.

We can generalise the notion of "pull-back" to p-forms.

Suppose that M and N are manifolds (of dimension m and n) and that $f : M \longrightarrow N$ is a smooth function. Given a p-form, ω, on N, we can define a *pull-back* p-form, $f^*\omega$, on M as follows. Given $x \in M$, we have the derivative map $d_x f : T_x M \longrightarrow T_{fx} N$, hence a dual map $(d_x f)^* : T^*_{fx} N \longrightarrow T^*_x M$. This in turn gives rise to a linear map $\Lambda^p (d_x f)^* : \Lambda^p T^*_{fx} N \longrightarrow \Lambda^p T^*_x M$. We set $(f^*\omega)(x) = (\Lambda^p(d_x f)^*)(\omega(f(x)))$. One checks that this is smooth, and so $f^*\omega$ is a p-form on M. When $p = 1$, this reduces to the pull-back on 1-forms already defined.

In particular, we can pull back p-forms to any manifold embedded in a larger manifold, such as $\mathbb{R}^n$.

Again this is very formal. It's much easier to see with an example.

(1) Suppose S^1 is the unit circle in $\mathbb{R}^2$. Let θ be the angle co-ordinate, so that $x = \cos\theta$ and $y = \sin\theta$, where x, y are the usual coordinates on $\mathbb{R}^2$. Then the pull back of dx and dy are obtained by differentiating these formulae: $-\sin\theta \, d\theta$ and $\cos\theta \, d\theta$. From this, we can pull back an arbitrary 1-form, just by linear extension.

For example $2xy\,dx + xe^y\,dy$ pulls back to $(-2\cos\theta\,\sin^2\theta\ + \cos^2\theta e^{\sin\theta})\,d\theta$, etc.

(2) Suppose S^2 is the unit 2-sphere in $\mathbb{R}^3$. Consider spherical polar coordinates, θ, ϕ (away from the poles). The coordinates in $\mathbb{R}^3$ are given by $x = \sin\theta\,\cos\phi$, $y = \sin\theta\,\sin\phi$, $z = \sin\theta$. The pull-backs of dx, dy and dz are respectively: $\cos\theta\,\cos\phi\ d\theta - \sin\theta\,\sin\phi\ d\phi$, $\cos\theta\,\sin\phi\ d\theta + \sin\theta\,\cos\phi\ d\phi$ and $\cos\phi\,d\phi$. So the pull back of $dx \wedge dy$ is $\cos\theta\ \sin\theta\ d\theta \wedge d\phi$.

Exercise : Check that the above examples agree with the formal definition of pull-back.

Integration on manifolds.

For the rest of this section, we will focus on m-forms.

Let $\{\phi_\alpha : U_\alpha \longrightarrow V_\alpha\}_{\alpha\in\mathcal{A}}$ be an atlas for M. If $\alpha, \beta \in \mathcal{A}$, then on the overlap, $U_\alpha \cap U_\beta$, we have

$$dx_i^\alpha \wedge \cdots \wedge dx_m^\alpha = \Delta_{\alpha\beta}(x) dx_1^\beta \wedge \cdots \wedge dx_m^\beta$$

where $\Delta_{\alpha\beta}$ is the determinant of the Jacobian of the transition function, $\phi_\alpha \circ \phi_\beta^{-1}$. (Recall that 1-forms transform according to the Jacobian.) Note in particular, if we have an oriented atlas, then $\Delta_{\alpha\beta}(x) > 0$. This gives rise to:

Theorem 10.3. *An m-manifold is orientable if and only if it admits a nowhere vanishing m-form.*

Proof. Suppose M is orientable. Choose an oriented atlas. We can suppose (by paracompactness) that the cover, $\{U_\alpha\}_\alpha$, is locally finite. Let $\{\rho_\alpha\}_\alpha$ be a partition of unity subordinate to this cover (Theorem 8.4). Let $\omega = \sum_{\alpha\in\mathcal{A}} \rho_\alpha dx_1^\alpha \wedge \cdots \wedge dx_m^\alpha$. (We define $\omega(x)$ pointwise, so this a finite sum for any given x.) Now if $x \in U_\alpha$, we can express all of the contributions in terms of $dx_1^\alpha \wedge \cdots \wedge dx_m^\alpha$. By the above transformation rule, we see that all the coefficients are non-negative, and at least one must be positive. In particular, $\omega(x) \neq 0$.

Conversely, suppose that ω is a nowhere vanishing m-form. Let $\{\phi_\alpha : U_\alpha \longrightarrow V_\alpha\}_\alpha$ be an atlas for M. We can suppose that each U_α is connected. Writing $\omega = \lambda_\alpha dx_1^\alpha \wedge \cdots \wedge dx_m^\alpha$, we see that λ_α is either always positive or always negative on U_α. After postcomposing the chart with an orientation-reversing diffeomorphism of $\mathbb{R}^n$, if necessary, we can suppose that it is always positive. (We can take this diffeomorphism to be just $[(x_1, \ldots, x_m) \mapsto (-x_1, \ldots, x_m)]$). This now gives an oriented atlas, since if $x \in U_\alpha \cap U_\beta$, we have $\lambda_\beta(x) = \Delta_{\alpha\beta}(x)\lambda_\alpha(x)$, so $\Delta_{\alpha\beta}(x) > 0$. $\square$

In general, there is no preferred m-form on M. However, in some cases, if M has additional structure, there is a canonical choice.

Suppose, for example, that $M \subseteq \mathbb{R}^n$, is a manifold embedded in $\mathbb{R}^n$, equipped with an orientation. Given $x \in M$, let $v_1(x), \dots, v_m(x) \in T_x M$ be a positively oriented orthonormal basis for $T_x M$. Let $\eta_1(x), \dots, \eta_m(x) \in T_x^* M$ be the dual basis, and set $\omega(x) = \eta_1(x) \wedge \cdots \wedge \eta_m(x) \in \Lambda^p T_x^* M$. First note that $\omega(x)$ is, in fact, well defined independently of the choice of $v_1(x), \dots, v_m(x)$. This is because a different choice of basis would transform by an orthogonal matrix (in $SO(n, \mathbb{R})$), and the dual basis would also. But this has determinant 1, so we get the same element of the exterior power. Second, we need to note that ω varies smoothly in x. One way to see this is to recall, by Lemma 5.2 and subsequent discussion, that we choose the v_i locally to vary smoothly in x. Since everything is given by simple formulae, it then follows that the η_i and ω, vary smoothly also. Thus, ω is a canonically defined nowhere vanishing m-form on M. It is referred to as the *volume form*, for reasons that will become apparent below.

In fact, the same construction works for a riemannian manifold. (The proof of Lemma 5.2 works also in this case.) Thus, any oriented riemannian manifold also has a globally defined volume form.

Integration of m-forms.

Let M be an oriented manifold. Given an m-form, ω, on M, we define the *support* of ω, denoted $\operatorname{supp} \omega$, to be the closure of the set $\{x \in M \mid \omega(x) \neq 0\}$. In this section, we will suppose that $\operatorname{supp} \omega$ is compact. The aim is to define the integral, $\int_M \omega$, of ω over M.

To this end, let $\{\phi \mid U_\alpha \longrightarrow V_\alpha\}_{\alpha \in \mathcal{A}}$ be a locally finite oriented atlas.

Suppose, first that η is an m-form with $\operatorname{supp} \eta \subseteq U_\alpha$ for some $\alpha \in \mathcal{A}$. We write $\eta = \lambda_\alpha dx_1^\alpha \wedge \cdots \wedge dx_m^\alpha$, where $\lambda_\alpha : U_\alpha \longrightarrow \mathbb{R}$ is smooth and compactly supported, and set and set $I_\alpha(\eta) = \int_{V_\alpha} \lambda_\alpha \circ \phi_\alpha^{-1}(x)\, dx$ (in the usual sense of an integral of a compactly supported smooth function in $\mathbb{R}^m$).

Now, choose a partition of unity, $\{\rho_\alpha\}_\alpha$, subordinate to $\{U_\alpha\}_\alpha$, and set $\int_M \omega = \sum_{\alpha \in \mathcal{A}} I_\alpha(\rho_\alpha \omega)$. (Note this is a finite sum, since only finitely many U_α meet $\operatorname{supp} \omega$.)

We need to check that this is well defined, independently of the choice of atlas and partition of unity.

Let $\{\phi'_\beta : U'_\beta \longrightarrow V'_\beta\}_{\beta \in \mathcal{B}}$ be another locally finite oriented atlas with the same orientation.

Suppose first, that η is an m-form with $\operatorname{supp} \eta \subseteq U_\alpha \cap U'_\beta$, where $\alpha \in \mathcal{A}$ and $\beta \in \mathcal{B}$.

Lemma 10.4. *$I_\alpha(\eta) = I_\beta(\eta)$.*

Proof. Locally, we write $\eta = \lambda_\alpha \mathbf{d}x^\alpha_{I(m)} = \lambda_\beta \mathbf{d}x^\beta_{I(m)}$. Then, by the coordinate-change formula, we have $\mathbf{d}x^\alpha_{I(m)} = \Delta_{\alpha\beta}(x) \mathbf{d}x^\beta_{I(m)}$, so $\lambda_\beta = \Delta_{\alpha\beta}\lambda_\alpha$, where $\Delta_{\alpha\beta} : U_\alpha \cap U_\beta \longrightarrow \mathbb{R}$ is the determinant of the Jacobian of $\phi_\alpha \circ (\phi'_\beta)^{-1}$, evaluated at

$\phi'_\beta(x)$. We therefore get:

$$I_\beta(\eta) = \int_{V'_\beta} \lambda_\beta \circ (\phi'_\beta)^{-1}(x)\,dx = \int_{V'_\beta} (\lambda_\beta \circ (\phi_\alpha)^{-1}) \circ (\phi_\alpha \circ (\phi'_\beta)^{-1})(x)\,dx$$

$$= \int_{V'_\beta} \lambda_\beta \circ \phi_\alpha^{-1}(x)\,\Delta_{\alpha\beta}^{-1}(x)\,dx = \int_{V_\alpha} \lambda_\alpha \circ \phi_\alpha^{-1}(x)\,dx = I_\alpha(\eta).$$

Here the penultimate equality is the "change of variable" rule for integration in $\mathbb{R}^m$, as mentioned in Section 1. (This cancels out the Jacobian arising from the change of coordinates in the m-form, and explains why m-forms are the natural things to use for integation.) $\square$

We can now write this quantity unambiguously as $I_{\{\alpha,\beta\}}(\eta)$.

Let $\{\rho'_\beta\}_\beta$ be a partition of unity subordinate to $\{U'_\beta\}_\beta$. As observed in Section 8, $\{\rho_\alpha\rho'_\beta\}_{(\alpha,\beta)\in\mathcal{A}\times\mathcal{B}}$ is a partition of unity subordinate to $\{U_\alpha \cap U'_\beta\}_{(\alpha,\beta)\in\mathcal{A}\times\mathcal{B}}$. Since integration is additive in $\mathbb{R}^n$, if $\operatorname{supp}\eta \subseteq U_\alpha$, we have $I_\alpha(\eta) = \sum_{\beta\in\mathcal{B}} I_{\{\alpha,\beta\}}(\rho'_\beta\eta)$. Similarly, this holds interchanging the roles of α and β.

We now have

$$\sum_{\alpha\in\mathcal{A}} I_\alpha(\rho_\alpha\omega) = \sum_{\alpha\in\mathcal{A},\beta\in\beta} I_{\{\alpha,\beta\}}(\rho_\alpha\rho'_\beta\omega) = \sum_{\beta\in\mathcal{B}} I_\beta(\rho'_\beta\omega)$$

as required.

This shows that $\int_M \omega$ is well defined.

This now allows us to define the volume of a compact orientable riemannian manifold. Choose any orientation, and set ω to be the volume form. We define its *volume* to be $\operatorname{vol} M = \int_M \omega$. Note that this is necessarily positive, since ω is consistent with the orientation. If we reverse the orientation, we get the same answer (exercise).

In fact, if $f : M \longrightarrow \mathbb{R}$ is any smooth function, we can integrate f with respect to volume, that is, integrate the m-form $f\omega$. The result, $\int_M f\omega$, is often denoted informally as $\int_M f\,dV$ (or $\int_M f\,dA$ if $\dim M = 2$), where of course, "V" and "A" stand for "volume" and "area". However this notation is somewhat at odds with ours, and could lead to ambiguities with "exterior derivatives" described in the next section.

Question : How should one define the volume of a compact non-orientable riemannian manifold?

11. Exterior derivatives and Stokes's theorem

Let M be an m-manifold. We write $\Omega^p(M)$ for the space of p-forms, viewed as a real vector space. We identify $\Omega^0(M) \equiv C^\infty(M)$, that is, a "0-form" is just a smooth real-valued function on M. If $U \subseteq M$ is open, and $\omega \in \Omega^p(M)$, we write $\omega|U \in \Omega^p(U)$ for the restriction of ω to U.

We claim that we can consistently define a family of linear functions, $d : \Omega^p(M) \longrightarrow \Omega^{p+1}(M)$ (all traditionally denoted by the same symbol, d), so as to satisfy the following:

(D1) If $f \in \Omega^0(M)$, then df has its original meaning as a 1-form (as defined in Section 7).
(D2) If $f \in \Omega^0(M)$ and $\omega \in \Omega^p(M)$, then $d(f\omega) = (df) \wedge \omega + f d\omega$.
(D3) If $\omega \in \Omega^p(M)$, then $dd\omega = 0$.
(D4) If $U \subseteq M$ is open, then $d(\omega|U) = (d\omega)|U$.

Note that (D4) tells us that f is a local operation. If $x \in M$, then $d\omega(x)$ only depends on ω defined on an arbitrarily small neighbourhood of x.

In fact, we claim first, that such a system of maps (if it indeed exists) will be unique.

To see this, we begin with the observation that if $n \in \mathbb{N}$, and $f_1, \dots, f_n \in \Omega^0(M)$, then $d(df_1 \wedge \cdots \wedge df_n) = 0$. This can be seen by induction on n, as follows. Applying (D2) with $f = f_1$ and $\omega = df_2 \wedge \cdots \wedge df_n = 0$, we get

$$df_1 \wedge \cdots \wedge df_n = d(f_1 df_2 \wedge \cdots \wedge df_n) - f_1 d(df_2 \wedge \cdots \wedge df_n).$$

The second term of the right-hand side is 0 by induction. Applying d to both sides, (D3) now tells us that the result is 0.

Now, by (D4) it is enough to check uniqueness for forms defined on the domain, U, of a chart, $\phi : U \longrightarrow V$. As usual, we write $x_1, \dots, x_m$ for the coordinates. Now dx_i has its usual meaning (by (D1)). Any $\omega \in \Omega^p(U)$ has the form $\omega = \sum_{I \in \mathcal{I}(m,p)} \lambda_I \mathbf{d}x_I$, where $\lambda_I \in \Omega^0(U)$. By the previous paragraph, $d(\mathbf{d}x_I) = 0$, so applying (D2) and linearity, we get $d\omega = \sum_{I \in \mathcal{I}(m,p)} (d\lambda_I) \wedge \mathbf{d}x_I$. In other words, we have no choice in how we define $d\omega$. It is unique.

To show existence we first check that if we define $d\omega$ in the manner prescribed above, then it must satisfy (D1)–(D3).

So let $\phi : U \longrightarrow V$ be a chart as above. Given

$$\omega = \sum_{I \in \mathcal{I}(m,p)} \lambda_I \mathbf{d}x_I \in \Omega^p(U),$$

we now define

$$d\omega = \sum_{I \in \mathcal{I}(m,p)} (d\lambda_I) \wedge \mathbf{d}x_I,$$

where $d\lambda_I$ is defined as in Section 7. (If $\omega = \lambda \in \Omega^0(U)$, this is just interpreted as $d\lambda$ in the above sense.) Now (D1) holds, by definition. For (D2), we note that,

$$d((f\lambda_I)\mathbf{d}x_I) = d(f\lambda_I) \wedge \mathbf{d}x_I = (\lambda_I df) \wedge \mathbf{d}x_I + f d\lambda_I \wedge \mathbf{d}x_I = df \wedge (\lambda_I \mathbf{d}x_I) + f d\lambda_I \wedge \mathbf{d}x_I.$$

(Here we have used the product rule for $\Omega^0(U)$ as described in Section 7.) Thus, (D2) follows summing over all $I \in \mathcal{I}(m,p)$. For (D3), note that $d\lambda_I = \sum_{i=1}^m \frac{\partial \lambda_I}{\partial x_i} dx_i$. Thus,

$$dd(\lambda_I \mathbf{d}x_I) = d\left(\sum_j \frac{\partial \lambda_I}{\partial x_j} dx_j \wedge \mathbf{d}x_I\right) = \sum_{i,j} \frac{\partial^2 \lambda_I}{\partial x_i \partial x_j} dx_i \wedge dx_j \wedge \mathbf{d}x_I$$

$$= \sum_{i<j} \left(\frac{\partial^2 \lambda_I}{\partial x_i \partial x_j} - \frac{\partial^2 \lambda_I}{\partial x_j \partial x_i}\right) dx_i \wedge dx_j \wedge \mathbf{d}x_I = 0.$$

Note that the second equality above calls for the observation that if $f \in \Omega^0(U)$, then $d(f dx_{i_1} \wedge \cdots \wedge dx_{i_p}) = (df) \wedge dx_{i_1} \wedge \cdots \wedge dx_{i_p}$ for any indices, i_k, in any order. (This follows, since transposing any two indices in the defining formula has the same effect on both sides, and setting two indices equal sets both sides equal to 0.) Summing over I, now gives (D3).

Note, by the earlier argument, (D1)–(D3) are enough to show uniqueness of $d : \Omega^p(U) \longrightarrow \Omega^{p+1}(U)$, without reference to (D4).

Of course, we need to define $d : \Omega^p(W) \longrightarrow \Omega^{p+1}(W)$ for any open subset $W \subseteq M$, including M itself. But we can now do this pointwise. If $x \in W$, choose any chart $\phi : U \longrightarrow V$ with $x \in U \subseteq W$. Given any $\omega \in \Omega^p(U)$, the above gives $d(\omega|U) \in \Omega^{p+1}(U)$, and we set $d\omega(x) = d(\omega|U)(x)$. To check this is well defined, note that if $\phi' : U' \longrightarrow V'$ is another such chart, the operations satisfy (D1)–(D3) on the intersecton $U \cap U'$, so by uniqueness we must have $d(\omega|U)(x) = d(\omega|U')(x)$.

Now, by construction, (D4) holds. Since the properties are all local, we see that (D1)–(D3) hold also.

(Alternatively, one can check, using the change of basis formula left as an exercise in Section 10, that the definition of $d\omega$, is invariant under change of coordinates.)

In summary, we have now seen that there is a unique family of maps satisfying (D1)–(D4) above.

Definition. Given $\omega \in \Omega^p(M)$, $d\omega$ is called the *exterior derivative* of ω.

Exercise : If $\eta \in \Omega^p(M)$ and $\omega \in \Omega^q(M)$, then $d(\eta \wedge \omega) = (d\eta) \wedge \omega + (-1)^p \eta \wedge d\omega$.

Stokes's Theorem

Let M now be an m-manifold with boundary (as discussed in Section 9). Thus, ∂M is an $(m-1)$ manifold (if it is non-empty). We can define p-forms and exterior derivatives on M in essentially the same way, and we will use the same notation.

We first note that there is a natural linear map, $\iota : \Omega^p(M) \longrightarrow \Omega^p(\partial M)$, defined as follows. Suppose $\omega \in \Omega^p(M)$, and that $x_1, \ldots, x_m$ are local coordinates defined by some chart (to an open subset of H^m). We can write ω locally as $\sum_{I \in \mathcal{I}(m,p)} \lambda_I \mathbf{d}x_I$. We now define $\iota\omega$ locally by $\sum_{I \in \mathcal{I}(m-1,p)} \lambda_I \mathbf{d}x_I$, where we have identified $\mathcal{I}(m-1,p) \subseteq \mathcal{I}(m,p)$. In other words, we have simply thrown away all those terms that do not involve dx_m.

Exercise If $x \in \partial M$, then $\iota\omega(x)$ does not depend on the choice of chart containing x. (Note that any change of chart will send dx_m to a 1-form of the form λdx_m, where λ is a smooth function.)

Given this, we see that $\iota\omega$ is defined on all of ∂M.

In fact, the $\iota\omega$ is usually just denoted by ω, and is thought of as the *restriction* of ω to ∂M. This is the convention we adopt here.

We can finally state Stokes's theorem:

Theorem 11.1. *Let M be a compact oriented m-manifold with boundary. Let ω be an $(m-1)$-form on M. Then $\int_M d\omega = \int_{\partial M} \omega$.*

If $\partial M = \emptyset$, this says that $\int_M d\omega = 0$.

Proof. Suppose first that $\operatorname{supp} \omega \subseteq U$, where $\phi : U \longrightarrow V \subseteq H^m$ is a chart with local coordinates $x_1, \ldots, x_m$. We can write ω locally in the form $\omega = \sum_{i=1}^n \omega_i$ with $\omega_i = \lambda_i \mathbf{d}x_{I_i}$, where $I_i = I(m) \setminus \{i\}$. (In other words, $\mathbf{d}x_{I_i}$ can be written as "$dx_1 \wedge \cdots \wedge dx_m$" except with the "$dx_i$" term omitted.) Thus, ω restricted to ∂M (as defined above) is just ω_m.

Now

$$d\omega_i = \frac{\partial \lambda_i}{\partial x_i} dx_i \wedge \mathbf{d}x_{I_i} = (-1)^{i+1} \frac{\partial \lambda_i}{\partial x_i} \mathbf{d}x_{I(m)}.$$

Suppressing ϕ from the notation, we see (from the definition of integral) that

$$\int_M d\omega_i = \int_U d\omega_i = (-1)^{m+1} \int_{H^m} \frac{\partial \lambda_i}{\partial x_i} \, dx_1 \, dx_2 \, \ldots dx_m,$$

the last term being the usual integral in $H^m \subseteq \mathbb{R}^m$. Integrating the term in dx_i first (using Fubini's Theorem) we get that $\int_M d\omega_i = 0$ if $i \leq m$, and that

$$\int_M dw_m = (-1)^m \int_{\partial H^m} \lambda_m(x_1, \ldots, x_{m-1}, 0) \, dx_1 \, \ldots dx_{m-1} = \int_{\partial M} \omega_m = \int_{\partial M} \omega.$$

(We lose the $(-1)^m$ because of the conventions of orientation of ∂M defined at the end of Section 9.) Thus, summing over i, we get $\int_M d\omega = \int_{\partial M} \omega$.

For the general case, choose any finite oriented atlas, $\{\phi_\alpha : U_\alpha \longrightarrow V_\alpha\}_\alpha$, for M, and let $\{\rho_\alpha\}_\alpha$ be a subordinate partition of unity. Note that

$$\sum_\alpha d(\rho_\alpha \omega) = \sum_\alpha d\rho_\alpha \wedge \omega + \sum_\alpha \rho_\alpha d\omega = 0 + d\omega = d\omega,$$

since $\sum_\alpha d\rho_\alpha = d1 = 0$, and so

$$\int_M d\omega = \int_M \sum_\alpha d(\rho_\alpha \omega) = \sum_\alpha \int_M d(\rho_\alpha \omega) = \sum_\alpha \int_{\partial M} \rho_\alpha \omega = \int_{\partial M} \omega,$$

by the earlier case. $\square$

de Rham cohomology

We finish with a brief mention of "de Rham cohomology".

Let M be any manifold (without boundary).

Definition. We say that a p-form, ω, on M is *closed* if $d\omega = 0$. We say that ω is *exact* if there is a $(p-1)$-form, η, on M such that $\omega = d\eta$.

Since $dd = 0$, we see that every exact form is closed.

We write $\mathcal{B}^p(M) \subseteq \mathcal{Z}^p(M) = \Omega^p(M)$ for the subspaces of exact and closed p-forms. Write $\mathcal{H}^p(M) = \mathcal{Z}^p(M)/\mathcal{B}^p(M)$ for the quotient space.

Definition. $\mathcal{H}^p(M)$ is the p'th-*de Rham cohomology* group of M.

de Rham's theorem says that $\mathcal{H}^p$ is naturally isomorphic to the singular cohomology of M with real coefficients, generally denoted $H^p(M, \mathbb{R})$. In particular, it depends only on the topology of M (not on the smooth structure). It also follows that (if M is compact) $\mathcal{H}^p(M)$ is finite dimensional as a vector space over $\mathbb{R}$, even though $\mathcal{Z}^p(M)$ and $\mathcal{B}^p(M)$ are both infinite dimensional.

Mathematics Institute, University of Warwick, Coventry, CV4 7AL, Great Britain