Manifolds

Ed Segal
Autumn 2016

Contents

1 Introduction

A manifold is a particular kind of mathematical space, which encodes an idea of 'smoothness'. They're the most general kind of space on which we can easily do calculus - differentiation and integration. This makes them very important, and they're fundamental objects in geometry, topology, and analysis, as well as having lots of uses in applied maths and theoretical physics.

The simplest example of a manifold is the real vector space $\mathbb{R}^n$, for any n. More generally, a manifold is a space that 'locally looks like $\mathbb{R}^n$', so if you zoom in close enough, you can't tell that you're not in $\mathbb{R}^n$.

Example 1.1. The surface of the Earth is approximately a 2-dimensional sphere, a space that we denote S^2. There's a myth that people used to think the Earth was flat - the myth is obviously false, in fact the ancient Greeks had a decent estimate of the radius of the Earth! But the story has a grain of plausibility, because 'close up' the Earth does look flat, and we could imagine that we're living on the surface of the plane $\mathbb{R}^2$. Hence the sphere S^2 is an example of a 2-dimensional manifold.

Example 1.2. The surface of a ring doughnut is a space we call a (2-dimensional) torus, and denote T^2 (see Figure 1). If you were a very small creature sitting on the doughnut, it wouldn't be immediately obvious that you weren't sitting on $\mathbb{R}^2$. So T^2 is another example of a 2-dimensional manifold.

Let's go down a dimension:

Example 1.3. A circle, sometimes denoted S^1, is an example of a 1-dimensional manifold. A small piece of a circle looks just like a small piece of the real line $\mathbb{R}$.

Here's an example of a different flavour:

Example 1.4. Let $\mathrm{Mat}_{2\times 2}(\mathbb{R})$ be the set of all 2×2 real matrices, this is a 4-dimensional real vector space so it's isomorphic to $\mathbb{R}^4$. Now let

$$GL_2(\mathbb{R}) \subset \mathrm{Mat}_{2\times 2}(\mathbb{R})$$

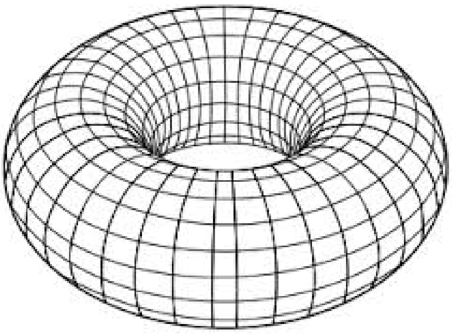

Figure 1: A torus.

denote the subset of *invertible* matrices. If M is an invertible matrix, and N is any matrix whose entries are sufficiently small numbers, then the matrix $M + N$ will still be invertible. So every matrix 'nearby' M also lies in $GL_2(\mathbb{R})$ (i.e. $GL_2(\mathbb{R})$ is an open subset). This means that a small neighbourhood of M looks exactly like a small neighbourhood of the origin in $\mathrm{Mat}_{2\times 2}(\mathbb{R}) \cong \mathbb{R}^4$. Hence $GL_2(\mathbb{R})$ is an example of a 4-dimensional manifold.

This is an example of a *Lie group*, a group that is also a manifold. Lie groups are very important, but they won't really be covered in this course.

We often picture manifolds as being subsets of some larger vector space, e.g. we think of S^2 or T^2 as smooth surfaces sitting inside $\mathbb{R}^3$. This is very helpful for our intuition, but the theory becomes much more powerful when we can talk about manifolds *abstractly*, without reference to any ambient vector space. A lot of the hard work in this course will involve developing the necessary machinery so that we can do this.

2 Topological manifolds and smooth manifolds

2.1 Topological manifolds

We now begin formalizing the concept of a manifold. The full definition is rather complicated, so we begin with a simpler version, called a *topological manifold*.

Definition 2.1. Let X be a topological space. A **co-ordinate chart** on X is the data of:

- An open set $U \subset X$.
- An open set $\tilde{U} \subset \mathbb{R}^n$, for some n.
- A homeomorphism
$$f : U \xrightarrow{\sim} \tilde{U}$$

When we want to specify a co-ordinate chart we always need to specify this triple $(U, \tilde{U}, f)$, but often we'll be lazy and just write (U, f), leaving the $\tilde{U}$ implicit.

The key distinguishing property of manifolds is that co-ordinate charts exist!

Definition 2.2. Let X be a topological space, and fix a natural number $n \in \mathbb{N}$. We say that X is an n-**dimensional topological manifold** iff for any point $x \in X$ we can find a co-ordinate chart

$$f : U \xrightarrow{\sim} \tilde{U} \subset \mathbb{R}^n$$

with $x \in U$.

In words, this says that at any point in X we can find an open neighbourhood which is homeomorphic to some open set in $\mathbb{R}^n$. A concise way to say this is that X is 'locally homeomorphic' to $\mathbb{R}^n$ (some people use the term 'locally Euclidean').

It's possible to prove that an open set in $\mathbb{R}^n$ cannot be homeomorphic to an open set in $\mathbb{R}^m$ unless $n = m$, so the dimension of a topological manifold is unambiguous. The proof of this fact is not difficult, but it uses some algebraic topology that isn't in this course.

Remark 2.3. There are two more conditions that are usually part of the definition of a topological manifold, namely that the space X should be:

- Hausdorff, and
- second-countable.

These are technical conditions used to rule out certain 'pathological' examples (see Appendix A). Every space we see in this course will be Hausdorff and second-countable, and we're going to avoid mentioning these conditions as far as possible.

Example 2.4. The circle S^1 is a 1-dimensional topological manifold. Let's prove this carefully. Firstly, let's define S^1 to be the subset

$$S^1 = \left\{(x, y);\ x^2 + y^2 = 1\right\} \subset \mathbb{R}^2$$

and equip it with the subspace topology. Next we need to find some co-ordinate charts, we'll do this using *stereographic projection.*

Let (x, y) be a point in S^1, not equal to $(0, -1)$. Draw a straight line through (x, y) and the point $(0, -1)$, and let $\tilde{x} \in \mathbb{R}$ be the point where this line crosses the x-axis, so:

$$\tilde{x} = \frac{x}{1 + y}$$

This sets up a bijection between points in S^1 (apart from $(0, -1)$) and points in the x-axis. So let's set

$$U_1 = S^1 \setminus (0, -1)$$

and note that this is an open set, since it's the intersection of S^1 with the open set $\{y \neq -1\} \subset \mathbb{R}^2$. Now set $\tilde{U}_1 = \mathbb{R}$, and

$$\begin{aligned} f_1 : U_1 &\to \tilde{U}_1 \\ (x, y) &\mapsto \tilde{x} = \frac{x}{1 + y} \end{aligned}$$

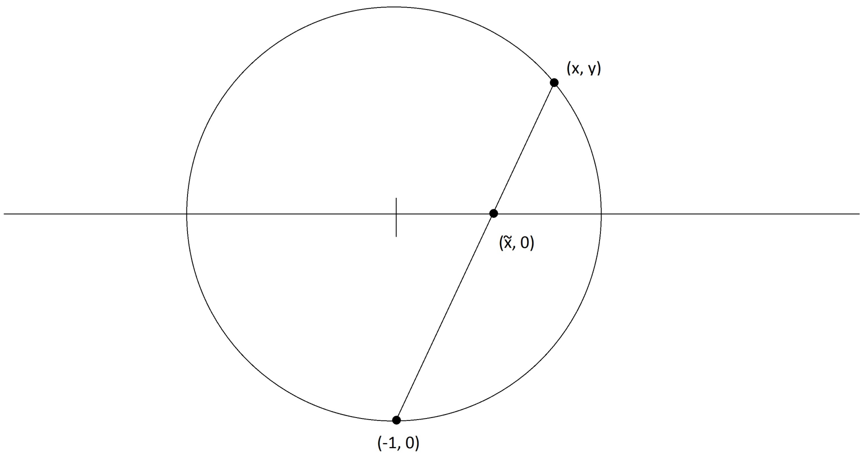

Figure 2: Stereographic projection.

Then f_1 is continuous, since it's the restriction to U_1 of a continuous function defined on $\{y \neq -1\} \subset \mathbb{R}^2$. To show that f_1 is a bijection we write down the inverse function:

$$\begin{aligned} f_1^{-1} : \tilde{U}_1 &\to U_1 \\ \tilde{x} &\mapsto \left(\frac{2\tilde{x}}{1+\tilde{x}^2}, \frac{1-\tilde{x}^2}{1+\tilde{x}^2} \right) \end{aligned}$$

An elementary calculation shows that $f_1^{-1}(\tilde{x})$ really does lie in U_1 for any $\tilde{x} \in \mathbb{R}$, and that f_1 and f_1^{-1} really are inverse to each other. Also f_1^{-1} is continuous (since it's evidently continuous when viewed as a function to $\mathbb{R}^2$), so we conclude that f_1 is a homeomorphism. The triple $(U_1, \tilde{U}_1, f_1)$ defines our first co-ordinate chart.

For our second co-ordinate chart, we use the same trick but we project from the point $(0, 1)$ instead. So we define $U_2 = S^1 \setminus (0, 1)$ and $\tilde{U}_2 = \mathbb{R}$, and:

$$\begin{aligned} f_2 : U_2 &\xrightarrow{\sim} \tilde{U}_2 \\ (x, y) &\mapsto \frac{x}{1-y} \end{aligned}$$

We repeat the previous arguments to check that this is also a co-ordinate chart. Now any point in S^1 lies in either U_1 or U_2 (most points lie in both) so we have proved that S^1 is a 1-dimensional topological manifold.

Now let's do the same thing for the n-dimensional sphere S^n.

Example 2.5. Let

$$S^n = \left\{ (x_0, ..., x_n);\ \sum x_i^2 = 1 \right\} \subset \mathbb{R}^{n+1}$$

with the subspace topology. We get our first co-ordinate chart using sterero-graphic projection from the point $(0, ..., 0, -1)$ (the "south pole"). So we define

$$U_1 = S^n \setminus (0, ..., 0, -1)$$
$$\tilde{U}_1 = \mathbb{R}^n$$

and

$$f_1 : U_1 \to \tilde{U}_1$$
$$(x_0, ..., x_n) \mapsto \left(\frac{x_0}{1 + x_n}, ..., \frac{x_{n-1}}{1 + x_n} \right)$$

We can prove that f_1 is a homeomorphism using the arguments from the previous example, in particular the inverse to f_1 is the function

$$f_1^{-1} : (\tilde{x}_0, ..., \tilde{x}_{n-1}) \mapsto \left(\frac{2\tilde{x}_0}{1 + \sum \tilde{x}_i^2}, \; ... \; , \frac{2\tilde{x}_{n-1}}{1 + \sum \tilde{x}_i^2}, \frac{1 - \sum \tilde{x}_i^2}{1 + \sum \tilde{x}_i^2} \right)$$

For our second co-ordinate chart we project from the point $(0, ..., 0, 1)$ (the "north pole"), i.e. we set

$$U_2 = S^n \setminus (0, ..., 0, 1)$$
$$\tilde{U}_2 = \mathbb{R}^n$$

and:

$$f_2 : U_2 \to \tilde{U}_2$$
$$(x_0, ..., x_n) \mapsto \left(\frac{x_0}{1 - x_n}, ..., \frac{x_{n-1}}{1 - x_n} \right)$$

Since every point in S^n lies in at least one of U_1 and U_2, this proves that S^n is a topological manifold, of dimension n.

2.2 Smooth atlases

Let's go back to S^1 again. Pick a point $(x, y) \in S^1$ which isn't $(0, -1)$ or $(0, 1)$. We've found two co-ordinate charts that we could use near this point; if we use U_1 (and f_1) then our point has co-ordinate $\frac{x}{1+y}$, but if we use U_2 (and f_2) then our point has co-ordinate $\frac{x}{1-y}$. How can we switch between the two co-ordinate systems?

The intersection of our two co-ordinate charts is

$$U_1 \cap U_2 = S^1 \setminus \{(0, -1), (0, 1)\}$$

In this locus, both the functions f_1 and f_2 are defined. Now notice that in the co-ordinate chart U_1, the point $(0, 1)$ gets mapped to the origin in $\mathbb{R}$. So the function f_1 defines a homeomorphism:

$$f_1 : U_1 \cap U_2 \xrightarrow{\sim} \mathbb{R} \setminus 0$$

Similarly, the point $(0, -1)$ lies in the co-ordinate chart U_2, and it gets mapped by f_2 to the origin in $\mathbb{R}$. So the function f_2 also defines a homeomorphism:

$$f_2 : U_1 \cap U_2 \xrightarrow{\sim} \mathbb{R} \setminus 0$$

To change between co-ordinates we must consider the composition:

$$\phi_{21} = f_2 \circ f_1^{-1} : \mathbb{R} \setminus 0 \xrightarrow{\sim} \mathbb{R} \setminus 0$$

This sends $\tilde{x} \in \mathbb{R} \setminus 0$ to:

$$\phi_{21}(\tilde{x}) = f_2\left(\frac{2\tilde{x}}{1+\tilde{x}^2}, \frac{1-\tilde{x}^2}{1+\tilde{x}^2} \right) = \frac{1}{\tilde{x}}$$

So if our point has co-ordinate $\tilde{x}$ under our first co-ordinate chart, then it has co-ordinate $1/\tilde{x}$ in our second chart. The function ϕ_{21} is called a *transition function*.

More generally, suppose X is any topological manifold, and let

$$f_1 : U_1 \xrightarrow{\sim} \tilde{U}_1 \subset \mathbb{R}^n$$

and

$$f_2 : U_2 \xrightarrow{\sim} \tilde{U}_2 \subset \mathbb{R}^n$$

be two co-ordinate charts on X. The intersection $U_1 \cap U_2$ is an open subset of U_1, and f_1 gives us a homeomorphism:

$$f_1 : U_1 \cap U_2 \xrightarrow{\sim} f_1(U_1 \cap U_2) \subset \tilde{U}_1$$

The image of this homeomorphism is some open subset of $\mathbb{R}^n$, contained in $\tilde{U}_1$. Similarly, f_2 gives us a homeomorphism

$$f_2 : U_1 \cap U_2 \xrightarrow{\sim} f_2(U_1 \cap U_2) \subset \tilde{U}_2$$

onto some other open subset of $\mathbb{R}^n$, which is contained in $\tilde{U}_2$.

Definition 2.6. Let X be a topological manifold, and let (U_1, f_1) and (U_2, f_2) be two co-ordinate charts on X. The **transition function** between these two co-ordinate charts is the function:

$$\phi_{21} = f_2 \circ f_1^{-1} : f_1(U_1 \cap U_2) \xrightarrow{\sim} f_2(U_1 \cap U_2)$$

The transition function is automatically a homeomorphism between these two open subsets of $\mathbb{R}^n$, since it's a composition of two homeomorphisms. Notice that it's possible that the intersection of U_1 and U_2 is empty, but then the transition function isn't very interesting!

Also notice that ϕ_{21} depends on the ordering of the two co-ordinate charts; it's the transition function *from* the chart U_1 *to* the chart U_2. If we reverse the order then we get the transition function

$$\phi_{12} = f_1 \circ f_2^{-1} : f_2(U_1 \cap U_2) \xrightarrow{\sim} f_1(U_1 \cap U_2)$$

but these two functions are inverse to each other:

$$\phi_{12} = (\phi_{21})^{-1}$$

Example 2.7. Let $X = S^n$, and consider the two co-ordinate charts that we found in Example 2.5. We have

$$U_1 \cap U_2 = S^n \setminus \{(0, ..., 0, -1), (0, ..., 0, 1)\}$$

$$f_1(U_1 \cap U_2) = \mathbb{R}^n \setminus 0$$
$$f_2(U_1 \cap U_2) = \mathbb{R}^n \setminus 0$$

(in this example both $f_1(U_1 \cap U_2)$ and $f_2(U_1 \cap U_2)$ happen to be the *same* subset of $\mathbb{R}^n$, but this is a coincidence). The transition function between these two charts is the function:

$$\phi_{21} : \mathbb{R}^n \setminus 0 \xrightarrow{\sim} \mathbb{R}^n \setminus 0$$
$$(\tilde{x}_0, ..., \tilde{x}_{n-1}) \mapsto \left(\frac{\tilde{x}_0}{\sum \tilde{x}_i^2}, \; ..., \; \frac{\tilde{x}_{n-1}}{\sum \tilde{x}_i^2} \right)$$

A transition function tells us how to change co-ordinates between two different charts in some region of our manifold, so it tells us each new co-ordinate as some continuous function of the old co-ordinates. However, we don't want our change-of-co-ordinate functions to be merely continuous, we really want them to be smooth.

Recall that a function
$$F : \mathbb{R}^n \to \mathbb{R}^m$$
is called *smooth* (or C^∞) if we can take partial derivatives of F to any order, in any direction. This definition also makes sense if F is only defined on some open subset of $\mathbb{R}^n$. Since a transition function is a function from an open set in $\mathbb{R}^n$ to some other open set in $\mathbb{R}^n$, it makes sense to ask if the transition function is smooth.

Note that if we just have a single chart (U, f) then (in general) it makes no sense to ask if f is smooth! This is because U is just an open set in an abstract topological space X, so we have no way of defining differentiation. It's only when we are *comparing* two charts that we can ask about smoothness.

Definition 2.8. Let X be a topological manifold. An **atlas** for X is a collection of co-ordinate charts on X

$$f_i : U_i \to \tilde{U}_i \subset \mathbb{R}^n, \qquad i \in I$$

indexed by some (possibly-infinite) set I, such that

$$\bigcup_{i \in I} U_i = X$$

So an atlas is a set of co-ordinate charts that collectively cover the whole of X. By the definition of a topological manifold, an atlas always exists. The next definition is more important:

Definition 2.9. An atlas for a topological manifold X is called **smooth** iff for any two charts in the atlas, the transition function

$$\phi_{ij} : f_j(U_i \cap U_j) \xrightarrow{\sim} f_i(U_i \cap U_j), \quad i, j \in I$$

is a smooth function.

So a smooth atlas is a collection of co-ordinate charts that cover the whole of X, and such that whenever we change co-ordinates the new co-ordinates depend smoothly on the old co-ordinates.

Recall that a bijection ϕ between two open subsets of $\mathbb{R}^n$ is called a *diffeomorphism* if both ϕ and ϕ^{-1} are smooth. If we have a smooth atlas then all the transition functions have to be diffeomorphisms, because the definition requires that both ϕ_{ij} and the inverse transition function $\phi_{ji} = \phi_{ij}^{-1}$ are smooth.

Example 2.10. Let $X = S^n$, and consider the atlas consisting of the two coordinate charts that we found in Example 2.5. In Example 2.7 we wrote down the transition function ϕ_{21} between the two charts, and it is clearly a smooth function. By symmetry, the transition function ϕ_{12} in the other direction is also a smooth function (in fact it's easy to check that $\phi_{12} = \phi_{21}$ in this example). Therefore this is a smooth atlas.

Note that it was important to check that both ϕ_{12} and ϕ_{21} were smooth, because there do exist smooth bijections whose inverses are not smooth!

The next example shows another way that we might approach the circle.

Example 2.11. Let the group $\mathbb{Z}$ act on the real numbers $\mathbb{R}$ by translations, so the orbit of a real number x is the set:

$$[x] = \{x + n;\ n \in \mathbb{Z}\}$$

Let $T^1 = \mathbb{R}/\mathbb{Z}$ be the set of orbits, i.e. T^1 is the quotient set of $\mathbb{R}$ by the equivalence relation $x \sim y \iff x - y \in \mathbb{Z}$. The notation T^1 here means '1-dimensional torus'. Let

$$q : \mathbb{R} \to T^1$$

be the quotient map, which sends x to $[x]$. We give T^1 the quotient topology, so a set $U \subset T^1$ is open iff its preimage $q^{-1}(U)$ is open, this means that q is automatically continuous. Notice that q is also an open mapping, i.e. it sends open sets to open sets. This is because if $W \subset \mathbb{R}$ is any open set then $q^{-1}(q(W))$ is the union of all translates of W, hence it is an open set, hence $q(W)$ is an open set in T^1.

Every equivalence class apart from $[0]$ has a unique representative in the interval $[0, 1]$, so:

$$T^1 = [0, 1]/(0 \sim 1)$$

i.e. we form T^1 by taking an interval and then gluing the two ends together. This gives us a circle! In fact T^1 is homeomorphic to S^1, via the map:

$$x \mapsto \big(\cos(2\pi x), \sin(2\pi x)\big)$$

However, thinking of the circle as T^1 gives a different way to find a smooth atlas. Let $\tilde{U}_1$ be the open interval

$$\tilde{U}_1 = (0, 1) \subset \mathbb{R}$$

and let:

$$U_1 = q(\tilde{U}_1) \subset T^1$$

Then the quotient map $q : \tilde{U}_1 \to U_1$ is a bijection. We let

$$f_1 : U_1 \to \tilde{U}_1 = (0, 1) \subset \mathbb{R}$$

be the inverse function to $q : \tilde{U}_1 \to U_1$, so $f_1([x])$ is the unique representative of the orbit $[x]$ which lies in the interval $(0, 1)$. Since q is an open mapping,

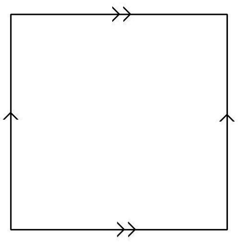

Figure 3: The 2-dimensional torus $T^2 = \mathbb{R}^2/\mathbb{Z}^2$.

the function f_1 is continuous, and thus a homeomorphism. Hence (U_1, f_1) is a co-ordinate chart on T^1.

Now let $\tilde{U}_2 = (-\frac{1}{2}, \frac{1}{2}) \subset \mathbb{R}$, so we get a second co-ordinate chart by defining $U_2 = q(\tilde{U}_2)$ and defining f_2 to be the inverse of the quotient map $q : \tilde{U}_2 \to U_2$. These two charts cover the whole of T^1, now let's look at the transition functions. We have

$$U_1 \cap U_2 = T^1 \setminus \left\{[0], [\tfrac{1}{2}]\right\}, \quad f_1(U_1 \cap U_2) = (0, \tfrac{1}{2}) \sqcup (\tfrac{1}{2}, 1), \quad f_2(U_1 \cap U_2) = (-\tfrac{1}{2}, 0) \sqcup (0, \tfrac{1}{2})$$

and the transition function is:

$$\begin{aligned} \phi_{21} : (0, \tfrac{1}{2}) \sqcup (\tfrac{1}{2}, 1) &\longrightarrow (-\tfrac{1}{2}, 0) \sqcup (0, \tfrac{1}{2}) \\ x &\mapsto \begin{cases} x, & \text{for } x < \frac{1}{2} \\ x - 1, & \text{for } x > \frac{1}{2} \end{cases} \end{aligned}$$

This function is a diffeomorphism; both ϕ_{21} and its inverse ϕ_{12} are smooth. So this is a smooth atlas on T^1.

Now let's do the 2-dimensional version of the previous example:

Example 2.12. Let the group $\mathbb{Z}^2$ act on $\mathbb{R}^2$ by translations, so the orbits are:

$$[(x, y)] = \{(x + n, y + m);\ n, m \in \mathbb{Z}\}$$

Let $T^2 = \mathbb{R}^2/\mathbb{Z}^2$ be the quotient space. We can picture this as a square with opposite sides glued together:

$$T^2 = [0, 1] \times [0, 1] \;/\; (x, 0) \sim (x, 1) \text{ and } (0, y) \sim (1, y)$$

(see Figure 3). Hopefully it's clear that this produces a 2-dimensional torus.

Let q denote the quotient map $q : \mathbb{R}^2 \to T^2$. Now consider the following four open subsets of $\mathbb{R}^2$:

$$\begin{aligned} \tilde{U}_1 &= (0, 1) \times (0, 1), & \tilde{U}_2 &= (-\tfrac{1}{2}, \tfrac{1}{2}) \times (0, 1) \\ \tilde{U}_3 &= (0, 1) \times (-\tfrac{1}{2}, \tfrac{1}{2}), & \tilde{U}_4 &= (-\tfrac{1}{2}, \tfrac{1}{2}) \times (-\tfrac{1}{2}, \tfrac{1}{2}) \end{aligned}$$

For each i we set $U_i = q(\tilde{U}_i) \subset T^2$, and then since $q : \tilde{U}_i \to U_i$ is a bijection we can define $f_i : U_i \to \tilde{U}_i$ to be the inverse to q. Using the same arguments as in Example 2.11 we can show that each (U_i, f_i) is a co-ordinate chart, and it's easy to check that this is a smooth atlas.

We can generalize this to $T^n = \mathbb{R}^n/\mathbb{Z}^n$ for any n; this gives us an n-dimensional torus.

Here's a trivial but important example:

Example 2.13. Let $X = \mathbb{R}^n$. Now let $U_0 = X$, let $\tilde{U}_0 = \mathbb{R}^n$, and let

$$f_0 : U_0 \xrightarrow{\sim} \tilde{U}_0$$

be the identity function. This is a co-ordinate chart, and since it covers all of X it is in fact an atlas (consisting of a single chart). Hence $\mathbb{R}^n$ is a topological manifold, of dimension n. Furthermore this is a smooth atlas, because there are no non-trivial transition functions!

More generally we can let X be any open set inside $\mathbb{R}^n$, then the same procedure provides a smooth atlas on X (with a single chart).

2.3 Smooth structures

If we're given a smooth atlas for a topological manifold X then we have a collection of co-ordinate charts which cover all of X. However, in practice we might wish to also use other co-ordinate charts, since different co-ordinates are useful for different problems.

Definition 2.14. Let X be a topological manifold, and let

$$\mathcal{A} = \{(U_i, f_i);\ i \in I\}$$

be a smooth atlas for X. Let (U, f) be any co-ordinate chart on X. We say that (U, f) is **compatible** with the atlas $\mathcal{A}$ iff the transition functions between (U, f) and any chart in $\mathcal{A}$ are diffeomorphisms.

In other words, the new chart (U, f) is compatible with the atlas $\mathcal{A}$ iff the union $\mathcal{A} \cup \{(U, f)\}$ is still a smooth atlas.

Once we've fixed a smooth atlas $\mathcal{A}$, it's not very important to know which co-ordinate charts are actually in $\mathcal{A}$: the important thing is to know which charts are *compatible* with $\mathcal{A}$. These are the charts that we are 'allowed to use', we should disregard all the charts that are not compatible with $\mathcal{A}$.

There are two obvious ways to produce new charts which are compatible with $\mathcal{A}$.

Lemma 2.15. *Let $\mathcal{A}$ be a smooth atlas on an n-dimensional topological manifold, and let*

$$f : U \xrightarrow{\sim} \tilde{U}$$

be a chart in $\mathcal{A}$.

(i) Let V be an open subset of U. Then $f|_V : V \to f(V)$ is a co-ordinate chart on X which is compatible with $\mathcal{A}$.

(ii) Let $\tilde{V} \subset \mathbb{R}^n$ be an open set and let $g : \tilde{U} \to \tilde{V}$ be a diffeomorphism. Then $g \circ f : U \to \tilde{V}$ is a co-ordinate chart on X which is compatible with $\mathcal{A}$.

Proof. Exercise. □

Example 2.16. We saw in Example 2.13 that there is a 'trivial' smooth atlas on the topological manifold $X = \mathbb{R}$. Let $\mathcal{A}$ denote this atlas; it contains a single chart (U_0, f_0) with $U_0 = \mathbb{R}$ and f_0 the identity function.

Here are two more co-ordinate charts on X:

$$U_1 = \mathbb{R}_{>0}, \quad \tilde{U}_1 = \mathbb{R}_{>0}, \quad f_1(x) = \frac{1}{x}$$

$$U_2 = \mathbb{R}_{<1}, \quad \tilde{U}_2 = \mathbb{R}_{>0}, \quad f_2(x) = \frac{1}{1-x}$$

Since f_0 is the identity, the transition function from (U_0, f_0) to (U_1, f_1) is simply:

$$\phi_{10} = f_1 : \mathbb{R}_{>0} \xrightarrow{\sim} \mathbb{R}_{>0}$$

This is a diffeomorphism, so the chart (U_1, f_1) is compatible with the atlas $\mathcal{A}$. Similarly the transition function from (U_0, f_0) to (U_2, f_2) is the diffeomorphism

$$\phi_{20} = f_2 : \mathbb{R}_{<1} \xrightarrow{\sim} \mathbb{R}_{>0}$$

so (U_2, f_2) is also compatible with $\mathcal{A}$.

For completeness let's write down the transition functions between (U_1, f_1) and (U_2, f_2). We have

$$U_1 \cap U_2 = (0, 1), \quad f_1(U_1 \cap U_2) = \mathbb{R}_{>1}, \quad f_2(U_1 \cap U_2) = \mathbb{R}_{>1}$$

and:

$$\phi_{21}(\tilde{x}) = \frac{\tilde{x}}{\tilde{x} - 1} = \phi_{12}(\tilde{x})$$

Notice that in this example the two charts (U_1, f_1) and (U_2, f_2) cover the whole of $\mathbb{R}$, so they form an atlas; in fact they form a smooth atlas, because ϕ_{21} is a diffeomorphism. In an important sense, this atlas is 'equivalent' to the atlas $\mathcal{A}$.

Definition 2.17. Let X be a topological manifold, and let $\mathcal{A}$ and $\mathcal{B}$ be two smooth atlases for X. We say that $\mathcal{A}$ and $\mathcal{B}$ are **compatible** iff every chart in $\mathcal{B}$ is compatible with the atlas $\mathcal{A}$.

Equivalently, we could say that $\mathcal{A}$ and $\mathcal{B}$ are compatible iff every chart in $\mathcal{A}$ is compatible with the atlas $\mathcal{B}$, or iff the union $\mathcal{A} \cup \mathcal{B}$ is a smooth atlas.

Example 2.18. In Example 2.16 we found two smooth atlases on $\mathbb{R}$, one with a single chart, $\mathcal{A} = \{(U_0, f_0)\}$, and one with two charts, $\mathcal{B} = \{(U_1, f_1), (U_2, f_2)\}$. These two atlases are compatible.

As we've said, knowing exactly which charts are in our atlas $\mathcal{A}$ is not so important, all we really care about is the set of charts that are compatible with $\mathcal{A}$. The next lemma says that if we replace $\mathcal{A}$ by a compatible atlas $\mathcal{B}$ then this information does not change: the set of compatible charts remains the same.

Lemma 2.19. *Let X be a topological manifold, and let*

$$\mathcal{A} = \{(U_i, f_i);\ i \in I\} \quad \textit{and} \quad \mathcal{B} = \{(U_j, f_j);\ j \in J\}$$

be two compatible smooth atlases for X. Let (U_0, f_0) be a co-ordinate chart on X which is compatible with the atlas $\mathcal{A}$. Then (U_0, f_0) is also compatible with $\mathcal{B}$.

Proof. Pick any chart (U_j, f_j) in $\mathcal{B}$ and consider the transition function:

$$\phi_{j0} : f_0(U_0 \cap U_j) \xrightarrow{\sim} f_j(U_0 \cap U_j)$$

We need to show that both ϕ_{j0} and ϕ_{j0}^{-1} are smooth functions. Pick any point $U_0 \cap U_j$. We're going to show that ϕ_{j0} is smooth at the point $f_0(x)$, and that ϕ_{j0}^{-1} is smooth at the point $f_j(x) = \phi_{j0}(x)$. If we can do this for any point x, then we'll have shown that both functions are smooth, and proved the lemma.

Since $\mathcal{A}$ is an atlas, there exists some chart $(U_i, f_i) \in \mathcal{A}$ with $x \in U_i$. Set $W = U_0 \cap U_j \cap U_i$, this is an open neighbourhood of x. We have homeomorphisms:

$$\begin{tikzcd} & W \arrow[dl, "f_0"'] \arrow[d, "f_i"] \arrow[dr, "f_j"] & \\ f_0(W) & f_i(W) & f_j(W) \end{tikzcd}$$

The set $f_0(W)$ is an open set inside $f_0(U_0 \cap U_j)$, and the composition

$$f_j \circ f_0^{-1} : f_0(W) \xrightarrow{\sim} f_j(W)$$

is just the restriction of the transition function ϕ_{j0}. Similar statements apply when we move between $f_i(W)$ and either of the other two charts, so we have that:

$$\phi_{j0}|_{f_0(W)} = \phi_{ji}|_{f_i(W)} \circ \phi_{i0}|_{f_0(W)}$$

Since ϕ_{ji} and ϕ_{i0} are smooth by assumption, it follows that ϕ_{j0} is smooth within the open set $f_0(W)$, and in particular it is smooth at the point $f_0(x)$. Since ϕ_{ij}^{-1} and ϕ_{i0}^{-1} are also smooth, ϕ_{j0}^{-1} is smooth at the point $f_j(x)$. □

Corollary 2.20. *Compatibility is an equivalence relation on smooth atlases.*

Proof. Exercise. □

Finally we can define the objects we really care about!

Definition 2.21. A **smooth manifold** is a topological manifold X together with an equivalence class $[\mathcal{A}]$ of compatible smooth atlases on X. We call the equivalence classes of atlases a **smooth structure** on X.

If we want to specify a smooth structure on X then we have to give a specific smooth atlas $\mathcal{A}$, but once we've done that then we are free to change $\mathcal{A}$ to any other compatible atlas. Notice that it makes sense to say that a co-ordinate chart is compatible with a smooth structure, since the set of compatible charts is independent of the specific choice of atlas.

Example 2.22. Let $X = \mathbb{R}^n$ and let $\mathcal{A}$ be the 'trivial' atlas from Example 2.13. Then $(\mathbb{R}^n, [\mathcal{A}])$ is a smooth manifold. This is called the *standard* smooth structure on $\mathbb{R}^n$.

Example 2.23. Let $X = S^n$ and let $\mathcal{A}$ be the 'stereographic projection' atlas from Example 2.5. Since this is a smooth atlas (Example 2.10), the pair $(S^n, [\mathcal{A}])$ defines a smooth manifold.

Example 2.24. Let $X = T^1$, and let $\mathcal{B}$ be the smooth atlas from Example 2.11. Then $(T^1, [\mathcal{B}])$ is a smooth manifold. Similarly $(T^n, [\mathcal{B}])$ is a smooth manifold, where $\mathcal{B}$ is the smooth atlas on the n-dimensional torus described in Example 2.12.

As we shall see later, our two versions of the circle, $(S^1, [\mathcal{A}])$ and $(T^1, [\mathcal{B}])$, really are 'the same' smooth manifold.

Now let's see an example of two atlases which are *not* compatible.

Example 2.25. Let $X = \mathbb{R}$ again, and consider the function:

$$\begin{aligned} g &: \mathbb{R} \to \mathbb{R} \\ x &\mapsto \begin{cases} x, & \text{for } x \leq 0 \\ 2x, & \text{for } x > 0 \end{cases} \end{aligned}$$

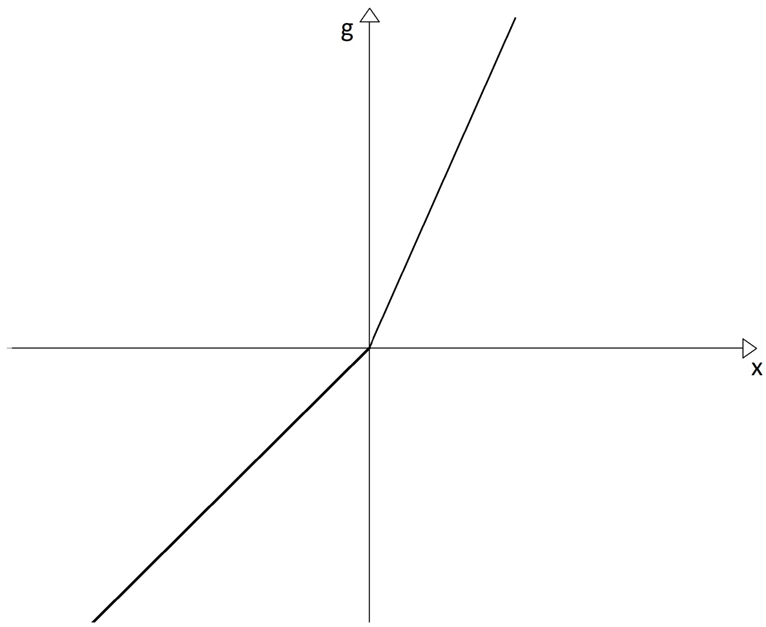

Figure 4: The function g from Example 2.25.

This is a homeomorphism, so $(\mathbb{R}, g)$ is a co-ordinate chart on X. Furthermore the chart covers the whole of X, so it gives us an atlas

$$\mathcal{C} = \{(\mathbb{R}, g)\}$$

which is automatically smooth since there are no transition functions to check. Hence $(\mathbb{R}, [\mathcal{C}])$ defines a smooth manifold.

However, because the function g is not smooth (it fails to have a derivative at the point $x = 0$), this co-ordinate chart is not compatible with the 'trivial' atlas $\mathcal{A}$ from Example 2.13. So $[\mathcal{C}]$ is a different smooth structure from the standard one $[\mathcal{A}]$.

Nevertheless, it will turn out that the two smooth manifolds $(\mathbb{R}, [\mathcal{A}])$ and $(\mathbb{R}, [\mathcal{C}])$ are still 'the same' smooth manifold (in the same sense that S^1 and T^1 are 'the same').

2.4 Some more examples

We're now going to show that a smooth atlas doesn't just determine the smooth structure on X, it also determines the underlying topology. This is quite useful for constructing examples of smooth manifolds, as we shall see.

Suppose that X is just a set, not a topological space. Now suppose we have a subset $U \subset X$, and a bijection

$$f : U \xrightarrow{\sim} \tilde{U}$$

where $\tilde{U}$ is an open subset of $\mathbb{R}^n$. This is quite like a co-ordinate chart, but note that it makes no sense to ask if f is a homeomorphism. Let's refer to this data (U, f) as a **pseudo-chart** on X. Similarly, let's call a collection $\mathcal{A} = \{(U_i, f_i)\}$ of pseudo-charts on X a **pseudo-atlas** if the union of all the U_i is the whole of X (*warning*: this is not standard terminology, I made up these words).

It still makes sense to talk about 'transition functions' for pseudo-charts. If (U_1, f_1) and (U_2, f_2) are two pseudo-charts on X, then $f_1(U_1 \cap U_2)$ is some (not necessarily open) subset of $\tilde{U}_1$, and $f_2(U_1 \cap U_2)$ is a subset of $\tilde{U}_2$, and we get a bijection:

$$\phi_{21} = f_2 \circ f_1^{-1} : f_1(U_1 \cap U_2) \to f_2(U_1 \cap U_2)$$

In general there's no reason for ϕ_{21} to be continuous. However, in many natural examples it will be continuous, and we have the following result:

Proposition 2.26. *Let X be a set, and let $\mathcal{A} = \{(U_i, f_i)\}$ be a pseudo-atlas for X. Suppose that for any pair $(U_1, f_1), (U_2, f_2)$ of pseudo-charts in $\mathcal{A}$, we have that:*

(a) both $f_1(U_1 \cap U_2) \subset \tilde{U}_1$ and $f_2(U_1 \cap U_2) \subset \tilde{U}_2$ are open subsets, and

(b) the transition function ϕ_{21} is continuous.

Then there is a topology on X such that each (U_i, f_i) is a co-ordinate chart. This topology is unique.

So a pseudo-atlas with properties (a) and (b) gives X the structure of a topological manifold.

Proof. Suppose that we can find such a topology on X. Then each U_i must be open, by definition. If $V \subset X$ is any other open set, then it must satisfy:

$$f_i(V \cap U_i) \text{ is open in } \tilde{U}_i, \quad \text{for all } (U_i, f_i) \in \mathcal{A} \tag{2.27}$$

Conversely, if $V \subset X$ is a subset that satisfies this property, then each $V \cap U_i$ is open in X, and since

$$V = \bigcup_{(U_i, f_i) \in \mathcal{A}} (V \cap U_i)$$

the set V must also be open. So the open sets are precisely the subsets that satisfy property (2.27). This proves uniqueness.

To prove existence, we can *define* the topology on X by declaring that the open sets are the subsets that satisfy (2.27). We leave it as an exercise to prove that this does define a topology, and that for this topology each (U_i, f_i) is a co-ordinate chart. $\square$

Corollary 2.28. *If we make the further assumption that each ϕ_{21} is smooth, then $[\mathcal{A}]$ is a smooth structure on the topological manifold X.*

Proof. By definition! $\square$

Now let's use this trick to define a manifold called 'real projective space'.

Example 2.29. Let $\mathbb{RP}^1$ denote the set of lines through the origin (i.e. 1-dimensional subspaces) in $\mathbb{R}^2$. Any non-zero vector $(x, y) \in \mathbb{R}^2$ lies in a unique line, which we denote by $x : y$. Two vectors lie in the same line iff they're proportional, so $x : y$ and $\lambda x : \lambda y$ are the same line for any non-zero $\lambda \in \mathbb{R}$. Let's show that we can make $\mathbb{RP}^1$ into a smooth 1-dimensional manifold, using Corollary 2.28.

The y-axis is the line $0 : 1$. Any other line $x : y$ has a well-defined gradient $y/x \in \mathbb{R}$, and this gives a bijection:

$$f_1 : U_1 = \mathbb{RP}^1 \setminus 0:1 \xrightarrow{\sim} \mathbb{R}$$

This is a pseudo-chart, and we can get a second one by setting $U_2 = \mathbb{RP}^1 \setminus 1:0$ and considering:

$$\begin{aligned} f_2 : U_2 &\xrightarrow{\sim} \mathbb{R} \\ x:y &\mapsto x/y \end{aligned}$$

So we have pseudoatlas. Both $f_1(U_1 \cap U_2)$ and $f_2(U_1 \cap U_2)$ are the subset $\mathbb{R} \setminus 0 \subset \mathbb{R}$, so condition (a) from Proposition 2.26 holds. The inverse to f_1 is the function

$$f_1^{-1} : \tilde{y} \mapsto 1:\tilde{y}$$

so the transition function ϕ_{21} is:

$$\begin{aligned} \phi_{21} : \mathbb{R} \setminus 0 &\xrightarrow{\sim} \mathbb{R} \setminus 0 \\ \tilde{y} &\mapsto 1/\tilde{y} \end{aligned}$$

Both ϕ_{21} and its inverse ϕ_{12} are continuous, and smooth. So this determines a topology on $\mathbb{RP}^1$ making it a topological manifold, and this a smooth atlas.

In fact $\mathbb{RP}^1$ is just the circle again. To see this observe that there's a well-defined map:

$$\begin{aligned} \mathbb{R}/\mathbb{Z} &\to \mathbb{RP}^1 \\ x &\mapsto \cos \pi x : \sin \pi x \end{aligned}$$

It should be clear that this is a bijection between T^1 and $\mathbb{RP}^1$, and it's not too hard to prove that it's a homeomorphism. It will also be straightforward to show that $\mathbb{RP}^1$ is the same smooth manifold as T^1 (and hence also the same as S^1), once we've learnt how to say that precisely.

However, in higher dimensions we get some genuinely new examples:

Example 2.30. Let $\mathbb{RP}^n$ denote the set of lines in $\mathbb{R}^{n+1}$, and for a non-zero vector $(x_0, ..., x_n)$ we denote the corresponding line by $x_0 : ... : x_n$. For each $i \in [0, n]$ we can get a pseudo-chart on $\mathbb{RP}^n$ by setting $U_i = \{x_0 : ... : x_n,\ x_i \neq 0\}$ and:

$$\begin{aligned} f_i : U_i &\to \mathbb{R}^n \\ x_0 : ... : x_n &\mapsto \left(\frac{x_0}{x_i},\ ...,\ \frac{x_{i-1}}{x_i}, \frac{x_{i+1}}{x_i},\ ...,\ \frac{x_n}{x_i} \right) \end{aligned}$$

This is a bijection, its inverse is:

$$f_i^{-1} : (\tilde{x}_1, ..., \tilde{x}_n) \mapsto \tilde{x}_1 : ... : \tilde{x}_i : 1 : \tilde{x}_{i+1} : ... : \tilde{x}_n$$

If we take a pair of these charts (U_i, f_i) and (U_j, f_j) (with $i < j$ say) then the transition function is

$$\begin{aligned} \phi_{ji} : \mathbb{R}^n \setminus \{\tilde{x}_j = 0\} &\xrightarrow{\sim} \mathbb{R}^n \setminus \{\tilde{y}_{i+1} = 0\} \\ (\tilde{x}_1, ..., \tilde{x}_n) &\mapsto \frac{1}{\tilde{x}_j} \left(\tilde{x}_1, ..., \tilde{x}_i, 1, \tilde{x}_{i+1}, ..., \hat{j}, ..., \tilde{x}_n \right) \end{aligned}$$

(where the $\hat{j}$ indicates that we skip the component $\tilde{x}_j / \tilde{x}_j$). This is a smooth function between open subsets of $\mathbb{R}^n$. By Corollary 2.28 this gives $\mathbb{RP}^n$ the structure of a smooth manifold.

If we generalize this example by considering k-dimensional subspaces of $\mathbb{R}^n$, rather than 1-dimensional subspaces, we get 'Grassmannian' manifolds. This gets a bit more fiddly than our previous examples, and we won't go into all the details.

Example 2.31. For any $k \in [0, n]$, we define $\mathrm{Gr}(k, n)$ to be the set of k-dimensional subspaces of $\mathbb{R}^n$ (so $\mathbb{RP}^n = \mathrm{Gr}(1, n+1)$). We claim that we can give $\mathrm{Gr}(k, n)$ the structure of a smooth manifold, with dimension $k(n-k)$.

Suppose $S \subset \mathbb{R}^n$ is a k-dimensional subspace. If we fix a basis for S then this determines a rank k matrix $M \in \mathrm{Mat}_{k \times n}(\mathbb{R})$ whose rows are the basis vectors. Changing basis in S corresponds exactly to multiplying M by an invertible $k \times k$ matrix, so $\mathrm{Gr}(k, n)$ is the quotient set:

$$\{M \in \mathrm{Mat}_{n \times k}(\mathbb{R}),\ \mathrm{rank}(M) = k\} \ / \ \mathrm{GL}_k(\mathbb{R})$$

Now take one of these matrices M, and split it into two blocks

$$M = \left(M' \,|\, M'' \right)$$

where M' is a $k \times k$ matrix and M'' is a $k \times (n-k)$ matrix. Suppose M' is invertible, so we can multiply M by $(M')^{-1}$. Then M is equivalent in $\mathrm{Gr}(k, n)$ to a matrix of the form

$$\left(I_k \,|\, N \right)$$

where I_k is the $k \times k$ identity matrix and $N \in \mathrm{Mat}_{k \times (n-k)}(\mathbb{R})$. Furthermore there is exactly one point in the orbit of M that has this form. So this gives us a bijection between $\mathrm{Mat}_{k \times (n-k)}(\mathbb{R})$ and a subset of $\mathrm{Gr}(k, n)$, namely the set of orbits $[M]$ such that M' is invertible.

This gives us one pseudo-chart on $\mathrm{Gr}(k,n)$. To get another pseudo-chart, fix a subset $J \subset \{1,...,n\}$ of size k, and let M'_J be the $k \times k$ matrix obtained by taking the corresponding set of columns from M (so in our first pseudo-chart were setting $J = \{1,...,k\}$). If M'_J is invertible then $(M'_J)^{-1}M$ is a matrix of a special form: if you extract the J columns from it then you get the identity matrix I_k. For example, if $n = 5, k = 2$, and $J = \{1,3\}$ then

$$(M'_J)^{-1}M = \begin{pmatrix} 1 & a & 0 & b & c \\ 0 & d & 1 & e & f \end{pmatrix}$$

for some 2×3 matrix $\left(\begin{smallmatrix} a & b & c \\ d & e & f \end{smallmatrix}\right)$. So we get a bijection between $\mathrm{Mat}_{k\times(n-k)}(\mathbb{R})$ and the subset:

$$\{[M],\ \det(M'_J) \neq 0\}\ \subset \mathrm{Gr}(k,n)$$

To see that these pseudo-charts cover the whole of $\mathrm{Gr}(k,n)$ we need to quote a result from linear algebra: if every $k \times k$ minor of a matrix M vanishes (*i.e.* if $\det(M'_J) = 0$ for each J) then M has rank less than k.

Finally, by staring at this definition you should be able to convince yourself that the transition functions between these pseudo-charts are all given by rational functions, so they are smooth. So this defines the structure of a smooth manifold on $\mathrm{Gr}(k,n)$.

Now that we've laid down some technical foundations, we can get on with actually studying some manifolds. From this point on we're going to assume that 'everything in sight is smooth', i.e. we're going to say

- 'Manifold' when we mean 'smooth manifold'.
- 'Atlas' when we mean 'smooth atlas'.
- 'Co-ordinate chart' when we mean 'compatible co-ordinate chart'.

3 Submanifolds

3.1 Definition of a submanifold

Consider the following two subsets of $\mathbb{R}^2$:

$$Z_1 = \{(x, \sin x)\}$$
$$Z_2 = \{(x, x),\ x \le 0\} \cup \{(x, 2x),\ x \ge 0\}$$

So Z_1 is the graph of the function $x \mapsto \sin x$, and Z_2 is the graph of the function g from Example 2.25 and Figure 4. The subset Z_1 is nice and smooth, it looks like a (1-dimensional) manifold. But Z_2 doesn't look like a manifold, it has a sharp corner at the origin.

However, both of these subsets are homeomorphic to $\mathbb{R}$, in fact the graph of *any* continuous function $\mathbb{R} \to \mathbb{R}$ is always homeomorphic to $\mathbb{R}$. So Z_2 is certainly a topological manifold, and we could equip it with a smooth atlas if we wanted to. So why doesn't it *look* like a smooth manifold?

The problem is (of course) the way in which Z_2 is sitting inside the ambient space $\mathbb{R}^2$. To understand what's happening precisely, we need to introduce the concept of a *submanifold*.

Recall that an *affine* subspace of $\mathbb{R}^n$ is a translation of a linear subspace, *i.e.* a subset of the form

$$A = \{x + v;\ x \in W\} \subset \mathbb{R}^n$$

for some vector $v \in \mathbb{R}^n$ and some linear subspace $W \subset \mathbb{R}^n$. Linear subspaces are special case of affine subspaces.

Roughly speaking: a manifold is a space X that locally looks like $\mathbb{R}^n$, and a submanifold is a subset of X which locally looks like an affine subspace in $\mathbb{R}^n$.

Example 3.1. Consider the subset $Z_1 = \{(x, \sin x)\} \subset \mathbb{R}^2$ again. Now let $U = \mathbb{R}^2$, $\tilde{U} = \mathbb{R}^2$, and:

$$f : U \to \tilde{U}$$
$$(x, y) \mapsto (x, y - \sin x)$$

It's easy to check that f is a diffeomorphism, so (U, f) defines a co-ordinate chart on $\mathbb{R}^2$ (compatible with the standard smooth structure). Futhermore, $f(Z_1)$ is the subset:

$$\{(x, 0)\} \subset \tilde{U}$$

So in these co-ordinates, Z_1 is just the linear subspace $\mathbb{R} \times \{0\} \subset \mathbb{R}^2$.

Now we write down the formal definition.

Definition 3.2. Let X be a (smooth) manifold. Let Z be any subset of X. We say that Z is an m**-dimensional submanifold** of X iff, for any point $z \in Z$, there exists a (compatible) co-ordinate chart on X

$$f : U \xrightarrow{\sim} \tilde{U} \subset \mathbb{R}^n$$

with $z \in U$, and an m-dimensional affine subspace $A \subset \mathbb{R}^n$, such that:

$$f(U \cap Z) = \tilde{U} \cap A$$

In Example 3.1 we proved that the subset Z_1 inside $\mathbb{R}^2$ is a 1-dimensional submanifold (when $\mathbb{R}^2$ is equipped with the standard smooth structure). The subset Z_2 is not a submanifold, but we shall not prove this fact now.

Obviously any affine subspace $A \subset \mathbb{R}^n$ is a submanifold of $\mathbb{R}^n$ (with its standard smooth structure), just use the trivial chart on $\mathbb{R}^n$.

Example 3.3. Recall our definition of S^1:

$$S^1 = \left\{x^2 + y^2 = 1\right\} \subset \mathbb{R}^2$$

Let's prove that this subset is a (1-dimensional) submanifold of $\mathbb{R}^2$. We just need to use polar co-ordinates! Set

$$U_1 = \mathbb{R}^2 \setminus \{(x, 0),\ x \leq 0\}, \quad \tilde{U}_1 = \mathbb{R}_{>0} \times (-\pi, \pi)$$

and:

$$\begin{aligned} f_1^{-1} : \tilde{U}_1 &\to U_1 \\ (r, \theta) &\mapsto (r\cos\theta, r\sin\theta) \end{aligned}$$

Then it's clear that f_1^{-1} is a smooth bijection. It's slightly less obvious that the inverse function $f_1 : U_1 \to \tilde{U}_1$ is also smooth, but this can be shown using the Inverse Function Theorem (more on this shortly). Hence (U_1, f_1) is a co-ordinate chart on $\mathbb{R}^2$, compatible with the standard atlas. We have

$$f(S^1 \cap U_1) = \{(1, \theta),\ \theta \in (-\pi, \pi)\} \subset \tilde{U}_1$$

which is the intersection of $\tilde{U}_1$ with the 1-dimensional affine subspace $\{1\} \times \mathbb{R} \subset \mathbb{R}^2$. Every point on S^1 except for $(-1, 0)$ lies in U_1, so this co-ordinate chart demonstrates that S^1 satisfies the submanifold condition at every point except for $(-1, 0)$. To deal with this final point we can take a second co-ordinate chart (U_2, f_2) using polar co-ordinates with $\theta \in (0, 2\pi)$ (so we delete the positive x-axis).

If $Z \subset X$ is an m-dimensional submanifold, then Z is automatically a topological space because we can give it the subspace topology. We claim that in fact Z itself is a topological manifold of dimension m (as the name suggests!), and moreover it has a natural smooth structure induced from the smooth structure on X. To prove this, we first need to understand how to get co-ordinate charts on a submanifold.

Let's introduce a little more terminology. For any $m \in [0, n]$, let us define the **standard affine subspace** (of dimension m) to be the subspace

$$\{(x_1, ..., x_m, 0, ..., 0)\} \subset \mathbb{R}^n$$

and we'll just write $\mathbb{R}^m \subset \mathbb{R}^n$ when we mean this standard subspace.

Lemma 3.4. *If X is an n-dimensional manifold, and $Z \subset X$ is an m-dimensional submanifold, then for any point $z \in Z$ there exists a chart (U, f) on X containing z such that $f(U \cap Z)$ is the intersection of $\tilde{U}$ with the standard affine subspace $\mathbb{R}^m \subset \mathbb{R}^n$.*

Proof. Pick $z \in Z$. By definition, there is a chart $f : U \to \tilde{U}$ around z such that $f(U \cap Z)$ is the intersection of $\tilde{U}$ with some affine subspace $A \subset \mathbb{R}^n$. It's elementary to see that there is an invertible affine map

$$\tau : \mathbb{R}^n \to \mathbb{R}^n$$

i.e. the composition of an invertible linear map and a translation, which maps A to the standard subspace $\mathbb{R}^m$ (just pick an appropriate basis). Since τ is a diffeomorphism we can form a new chart $(U, \tau \circ f)$, and this has the required property. $\square$

So around any point in a submanifold Z we can always choose co-ordinates which make Z look like the *standard* subspace. We can use this fact to produce co-ordinate charts on Z, as follows. Pick a chart

$$f : U \xrightarrow{\sim} \tilde{U}$$

on X as in the previous lemma. The intersection $V = Z \cap U$ is an open set in Z, and f induces a homeomorphism:

$$g : V \xrightarrow{\sim} \tilde{V} = \tilde{U} \cap \mathbb{R}^m$$

Since $\tilde{V}$ is an open set in $\mathbb{R}^m$, this is a co-ordinate chart on Z.

The procedure we've just performed is simple, but important. Make sure you understand it!

Example 3.5. In Example 3.3 we showed that S^1 is a submanifold of $\mathbb{R}^2$. Let (U_1, f_1) be the polar co-ordinate chart on $\mathbb{R}^2$ from that example. Then as we saw, f_1 induces a homeomorphism between

$$V_1 = S^1 \cap U_1 = U_1 \setminus (-1, 0)$$

and the intersection of $\tilde{U}_1$ with the affine subspace $\{r = 1\}$. As in Lemma 3.4 we can compose f_1 with the affine map $\tau : (r, \theta) \mapsto (\theta, r - 1)$ to get a new chart on $\mathbb{R}^2$:

$$\begin{aligned} \tau \circ f_1 : \; U_1 &\xrightarrow{\sim} (-\pi, \pi) \times (-1, \infty) \\ (r\cos\theta, r\sin\theta) &\mapsto (\theta, r - 1) \end{aligned}$$

Then the image of V_1 in these co-ordinates is the subset $(-\pi, \pi) \times \{0\}$, *i.e.* the intersection of the standard subspace $\mathbb{R} \subset \mathbb{R}^2$ with the codomain of the chart. So the induced chart on S^1 is:

$$\begin{aligned} g_1 : \; V_1 &\xrightarrow{\sim} \tilde{V}_1 = (-\pi, \pi) \quad \subset \mathbb{R} \\ (\cos\theta, \sin\theta) &\mapsto \theta \end{aligned}$$

We can also do this procedure starting from the polar co-ordinate chart (U_2, f_2) with domain $U_2 = \mathbb{R}^2 \setminus \{x \geq 0\}$ and codomain $\mathbb{R}_{>0} \times (0, 2\pi)$. Then the resulting chart on S^1 is:

$$\begin{aligned} g_2 : \; V_2 = S^1 \setminus (1, 0) &\xrightarrow{\sim} \tilde{V}_2 = (0, 2\pi) \quad \subset \mathbb{R} \\ (\cos\theta, \sin\theta) &\mapsto \theta \end{aligned}$$

Together, these two charts form an atlas for S^1. Let's check that they form a smooth atlas; we'll do it in a way that's a bit more complicated than necessary, because it's helpful for understanding the proof of the next lemma.

Consider the two charts $(U_1, \tau \circ f_1)$ and $(U_2, \tau \circ f_2)$ on $\mathbb{R}^2$, which both map S^1 to the standard subspace. The transition function between them is:

$$\begin{aligned}\phi_{21}: \big((-\pi,0)\cup(0,\pi)\big)\times\mathbb{R}_{>-1} &\xrightarrow{\sim} \big((0,\pi)\cup(\pi,2\pi)\big)\times\mathbb{R}_{>-1}\\ (\theta, r') &\mapsto \begin{cases} (\theta, r') & \text{for } \theta > 0\\ (\theta+2\pi, r') & \text{for } \theta < 0\end{cases}\end{aligned}$$

Notice that this function maps the subset $\big((-\pi,0)\cup(0,\pi)\big)\times\{0\}$ to the subset $\big((0,\pi)\cup(\pi,2\pi)\big)\times\{0\}$; this had to be true because these are subsets corresponding to S^1. So if we restrict ϕ_{21} to the subset $\{r'=0\}$ we get an induced function:

$$\begin{aligned}\psi_{21}: \big((-\pi,0)\cup(0,\pi)\big) &\xrightarrow{\sim} \big((0,\pi)\cup(\pi,2\pi)\big)\\ \theta &\mapsto \begin{cases} \theta & \text{for } \theta > 0\\ \theta+2\pi & \text{for } \theta < 0\end{cases}\end{aligned}$$

Now consider the induced charts (V_1, g_1) and (V_2, g_2) on S^1; the transition function between these two charts is exactly the function ψ_{21}. It's evidently a diffeomorphism, so this is indeed a smooth atlas for S^1.

It's a straight-foward exercise to check that this smooth atlas is compatible with the 'stereographic projection atlas' from Example 2.4; just compute all the transition functions. Hence it defines the same smooth structure as the one we already have.

Now we generalize this example.

Lemma 3.6. *Let X be an n-dimensional manifold, and let $Z \subset X$ be an m-dimensional submanifold of X. Let (U_1, f_1) and (U_2, f_2) be two charts on X each of which maps Z to the standard subspace. Let (V_1, g_1) and (V_2, g_2) be the two induced charts on Z, and let ψ_{21} be the transition function between them. Then ψ_{21} is smooth.*

Proof. To simplify our notation let's set $U = U_1 \cap U_2$, and $V = V_1 \cap V_2 = U \cap Z$. By assumption, we have:

$$f_1(V) = f_1(U) \cap \mathbb{R}^m \qquad \text{and} \qquad f_2(V) = f_2(U) \cap \mathbb{R}^m$$

We know we have a smooth transition function

$$\phi_{21}: f_1(U) \xrightarrow{\sim} f_2(U)$$

between the charts on X, and restricting this to the subspace $\mathbb{R}^m$ gives a smooth function:

$$\widehat{\psi}_{21} = \phi_{21}|_{\mathbb{R}^m} : f_1(V) \hookrightarrow f_2(U)$$

However, we also know that ϕ_{21} must map $f_1(V)$ to $f_2(V)$, so the image of $\widehat{\psi}_{21}$ must be the subset $f_2(U)\cap\mathbb{R}^m$. So if we think of $\widehat{\psi}_{21}$ as an n-tuple of smooth real functions defined on $f_1(V)$, then the last $n-m$ of these functions are identically zero. The first m of these functions define a smooth map

$$\psi_{21}: f_1(V) \xrightarrow{\sim} f_2(V)$$

but this is exactly the transition function from (V_1, g_1) to (V_2, g_2). $\square$

Proposition 3.7. *Let X be an n-dimensional manifold, and let $Z \subset X$ be an m-dimensional submanifold of X. Then Z is an m-dimensional topological manifold, and it carries a smooth structure induced from the smooth structure on X.*

Proof. Let $\mathcal{C}$ denote the set of all charts on X which map Z to the standard subspace $\mathbb{R}^m \subset \mathbb{R}$. By Lemma 3.4, for any point $z \in Z$ there is a chart in $\mathcal{C}$ containing z. This induces a co-ordinate chart (V, g) on Z with codomain in $\mathbb{R}^m$, and with z lying in V. This proves that Z is an m-dimensional topological manifold.

If we pick any set of charts $\mathcal{A} \subset \mathcal{C}$ whose domains cover Z, then the induced set of charts on Z form an atlas. By Lemma 3.6 this atlas is smooth. The resulting smooth structure is independent of our choices, because any two atlases for Z produced by this method will be compatible with each other (by Lemma 3.6 again). □

Example 3.8. Recall the definition of real projective space $\mathbb{RP}^n$ from Example 2.30, with its atlas $\{(U_i, f_i), i \in [0, n]\}$. The inclusion of the standard subspace $\mathbb{R}^{m+1} \subset \mathbb{R}^{n+1}$ induces a well-defined map

$$\begin{aligned} \mathbb{RP}^m &\hookrightarrow \mathbb{RP}^n \\ x_0 : \ldots : x_m &\mapsto x_0 : \ldots : x_m : 0 : \ldots : 0 \end{aligned}$$

which is clearly an injection. Let's show that the image of this map is a submanifold in $\mathbb{RP}^n$.

Let $Z \subset \mathbb{RP}^n$ denote the image. At any point in Z at least one of $x_0, \ldots, x_m$ must be non-zero, so we must be in the domain of at least one of the co-ordinate charts $(U_0, f_0), \ldots, (U_m, f_m)$. For any of these charts we have $\tilde{U}_i = \mathbb{R}^n$, and $f_i(U_i \cap Z)$ is precisely the standard subspace $\mathbb{R}^m \subset \mathbb{R}^n$. So Z is indeed an m-dimensional submanifold of $\mathbb{RP}^n$.

By Proposition 3.7 we get an induced smooth structure on $Z = \mathbb{RP}^m$, by considering co-ordinate charts on $\mathbb{RP}^n$ that make Z look like the standard affine subspace and using them to induce co-ordinate charts on Z. We've just seen that the charts $(U_0, f_0), \ldots, (U_m, f_m)$ are of this form, and the charts they induce on Z are exactly the charts we used to define the smooth structure on $\mathbb{RP}^m$. So this induced smooth structure is the same as the one we have already.

3.2 A short detour into real analysis

Before we can continue our study of submanifolds, we need to recall a few things from real analysis.

Suppose we have a smooth function:

$$h = (h_1, \ldots., h_k) : \mathbb{R}^n \to \mathbb{R}^k$$

If we fix a point $x = (x_1, \ldots, x_n) \in \mathbb{R}^n$, recall that the *derivative of h at x* is the linear map

$$Dh|_x : \mathbb{R}^n \to \mathbb{R}^k$$

given by the k-by-n matrix

$$Dh|_x = \left(\left. \frac{\partial h_i}{\partial x_j} \right|_x \right)$$

of all partial derivatives of the components of h, evaluated at the point x. This is also known as the *Jacobian* of h (at the point x). This definition also makes sense if h is only defined in some open neighbourhood of x.

We have the following fundamental result, which is usually proved in a first course on multi-variable real analysis.

Theorem 3.9 (Inverse Function Theorem)**.** *Let U be an open subset of $\mathbb{R}^n$, and let*

$$F : U \to \mathbb{R}^n$$

be a smooth function. Let $x \in U$ be a point such that the derivative of F at x

$$DF|_x : \mathbb{R}^n \to \mathbb{R}^n$$

is an isomorphism. Then there exists an open neighbourhood $V \subset U$ of x such that the function

$$F : V \to F(V) \quad \subset \mathbb{R}^n$$

is a diffeomorphism.

If we already know that our function F is a smooth bijection, we can use the Inverse Function Theorem to test if it's a diffeomorphism:

Corollary 3.10. *If $F : U \to W$ is a smooth bijection between two open subsets of $\mathbb{R}^n$, and the derivative $DF|_x$ is an isomorphism for all points $x \in U$, then the inverse function $F^{-1} : W \to U$ is also smooth.*

Proof. This follows immediately. □

Example 3.11. In Example 3.3 we used the function:

$$\begin{aligned} F : \mathbb{R}_{>0} \times (-\pi, \pi) &\to \mathbb{R}^2 \setminus \{x < 0\} \\ (r, \theta) &\mapsto (r\cos\theta, r\sin\theta) \end{aligned}$$

This is evidently a smooth bijection, and Corollary 3.10 can be used to prove that it is a diffeomorphism. Hence $f = F^{-1}$ is a chart on $\mathbb{R}^2$ compatible with the standard smooth structure.

3.3 Level sets in $\mathbb{R}^n$

In this section we're going to introduce a very general trick for producing submanifolds of $\mathbb{R}^n$.

Let's think again about our two subsets

$$Z_1 = \{y - \sin x = 0\} \subset \mathbb{R}^2$$

which is a submanifold (Example 3.1), and

$$Z_2 = \{y - g(x) = 0\} \subset \mathbb{R}^2$$

(where g is the function defined in Example 2.25), which is not a submanifold. We also know that the subset

$$S^1 = \left\{y^2 + x^2 = 1\right\} \subset \mathbb{R}^2$$

is a submanifold (Example 3.3). All of these subsets are of the form

$$\{h(x,y) = \alpha\}$$

for some function $h : \mathbb{R}^2 \to \mathbb{R}$ and some real number $\alpha \in \mathbb{R}$. So we should ask: when is such a subset in fact a submanifold?

More generally, if h is a function

$$h : \mathbb{R}^n \to \mathbb{R}^k$$

and $\alpha \in \mathbb{R}^k$, when is the subset $\{h(x) = \alpha\} \subset \mathbb{R}^n$ a submanifold? This is an important question, which we will explore in some detail.

The subsets $\{h(x) = \alpha\} \subset \mathbb{R}^n$ are called the *level sets* of h. Firstly, note that functions

$$h_1(x,y) = y - \sin x, \quad \text{and} \quad h_3(x,y) = x^2 + y^2$$

are both smooth functions from $\mathbb{R}^2$ to $\mathbb{R}$, whereas the function

$$h_2(x,y) = y - g(x)$$

is not a smooth function (since g is not smooth). One might reasonably guess that the level sets of h are submanifolds provided that h is a smooth function. Unfortunately this is not enough, as the next example shows.

Example 3.12. Consider the smooth function:

$$\begin{aligned} h : \mathbb{R}^2 &\to \mathbb{R} \\ (x,y) &\to xy \end{aligned}$$

For any $\alpha \in \mathbb{R}$ let's denote the level set of h by

$$Z_\alpha = \{h(x,y) = \alpha\} \subset \mathbb{R}^2$$

(see Figure 5).

If $\alpha \neq 0$, then Z_α is the set $\{(x, \alpha/x); x \in \mathbb{R}_{\neq 0}\}$, it's both branches of a hyperbola. This a (1-dimensional) submanifold of $\mathbb{R}^2$: consider the co-ordinate chart with

$$U = \tilde{U} = \{x \neq 0\} \subset \mathbb{R}^2$$

and:

$$\begin{aligned} f : U &\to \tilde{U} \\ (x,y) &\mapsto (x, y - \alpha/x) \end{aligned}$$

Note that this really is a co-ordinate chart (f is a diffeomorphism), and that:

$$f(Z_\alpha \cap U) = \{(x,0)\} \subset \tilde{U}$$

Since Z_α is entirely contained in U, this demonstrates that Z_α is a submanifold for $\alpha \neq 0$.

However, for $\alpha = 0$, the level set

$$Z_0 = \{xy = 0\} = \{x = 0\} \cup \{y = 0\}$$

consists of both co-ordinate axes. This certainly doesn't look like a submanifold, because of the 'singularity' at the point where the two axes cross. In fact one can prove that Z_0 is not even a topological manifold.

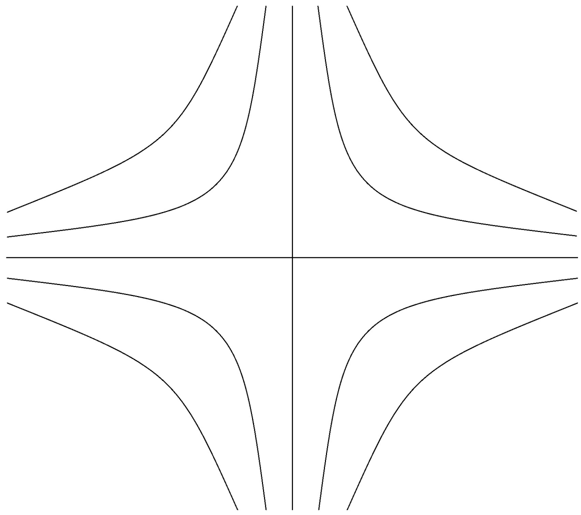

Figure 5: Level sets of $h = xy$.

So we need to find an additional condition to guarantee that a level set of h is a submanifold. In fact we must also look at the derivative of h.

Definition 3.13. Let $h : \mathbb{R}^n \to \mathbb{R}^k$ be a smooth function. A point $x \in \mathbb{R}^n$ is called a **regular point** of h iff the derivative

$$Dh|_x : \mathbb{R}^n \to \mathbb{R}^k$$

of h at x is a surjection. If x is not a regular point of h then it's called a **critical point**.

Definition 3.14. A point $\alpha \in \mathbb{R}^k$ is called a **regular value** of h iff every point in the level set

$$h^{-1}(\alpha) \subset \mathbb{R}^n$$

is a regular point of h. If α is not a regular value of h then it's called a **critical value**.

Notice that we can't have any regular points unless $k \leq n$. Also note that all these definitions work perfectly well if h is only defined on an open subset of $\mathbb{R}^n$.

Example 3.15. Let h be the function $h(x, y) = xy$ we considered in Example 3.12. If we fix a point $(x, y) \in \mathbb{R}^2$, then the derivative of h at this point is the 1-by-2-matrix (or linear map $\mathbb{R}^2 \to \mathbb{R}$):

$$Dh|_{(x,y)} = (y, x)$$

This is a surjection provided that at least one of x or y is not zero, so the origin $(x, y) = (0, 0)$ is a critical point of h and all other points are regular points. Hence the only critical value of h is $\alpha = 0$; if $\alpha \in \mathbb{R}$ is not zero then it is a regular value.

Comparing this with Example 3.12, we see that the level sets Z_α are submanifolds provided that α is a regular value of h (i.e. for $\alpha \neq 0$). However for the critical value $\alpha = 0$, the level set Z_0 contains a critical point $(0, 0)$, and Z_0 fails to be a submanifold near this point.

Here is the general result:

Proposition 3.16. *Let $h : \mathbb{R}^n \to \mathbb{R}^k$ be a smooth function. If $\alpha \in \mathbb{R}^k$ is a regular value of h then the level set*

$$Z_\alpha = h^{-1}(\alpha) \subset \mathbb{R}^n$$

is a submanifold of $\mathbb{R}^n$, of dimension $n - k$.

If X is any n-dimensional manifold, and $Z \subset X$ is an m-dimensional submanifold, then the difference

$$n - m$$

is called the *codimension* of Z. So the proposition says that the level set $Z_\alpha \subset \mathbb{R}^n$ is a submanifold of codimension k (provided that α is a regular value). These numbers should make intuitive sense: we start with a space having n degrees-of-freedom, then we impose k equations, so we have $n - k$ degrees-of-freedom left.

The proof of Proposition 3.16 will follow easily from another corollary of the Inverse Function Theorem, called the *Implicit Function Theorem*. For $k \leq n$, let's define the **standard projection** to be the linear map

$$\begin{aligned} \pi : \mathbb{R}^n &\to \mathbb{R}^k \\ (x_1, ..., x_n) &\mapsto (x_{n-k+1}, ..., x_n) \end{aligned}$$

which projects onto the last k co-ordinates (so the kernel of π is the standard subspace $\mathbb{R}^{n-k} \subset \mathbb{R}^n$). Since π is linear, for any point $x \in \mathbb{R}^n$ we have $D\pi|_x = \pi$, and since this is a surjection x is a regular point of π. The Implicit Function Theorem says that near a regular point any smooth function can be made to look like the standard projection π, by choosing the right co-ordinates.

Theorem 3.17 (Implicit Function Theorem)**.** *Let $U \subset \mathbb{R}^n$ be an open subset, and let*

$$h : U \to \mathbb{R}^k$$

be a smooth function, where $k \leq n$. Let $z \in U$ be a regular point of h. Then there exists an open neighbourhood $V \subset U$ of z, a second open subset $\tilde{V} \subset \mathbb{R}^n$, and a diffeomorphism $f : V \xrightarrow{\sim} \tilde{V}$, such that

$$h \circ f^{-1} = \pi : \tilde{V} \to \mathbb{R}^k$$

is the standard projection (restricted to $\tilde{V}$).

Before we give the proof let's look at what this is saying. Let's write $(h_1, ..., h_k)$ for the components of h. Then we're looking for a diffeomorphism $f = (f_1, ..., f_n)$ such that $\pi \circ f = h$, so $f_{n-k+i} = h_i$ for each $i \leq k$. In other words, we're looking for a co-ordinate system around z where $h_1, ..., h_k$ are the last k co-ordinates. The theorem says that if z is regular then this is always possible.

Proof. Let $x_1, ..., x_n$ be the standard co-ordinates on $\mathbb{R}^n$. The Jacobian $Dh|_z$ is a k-by-n real matrix, and we're assuming it has rank k. So the columns of this matrix span $\mathbb{R}^k$, and therefore some subset of the columns must form a basis for $\mathbb{R}^k$. After re-ordering the co-ordinates x_j if necessary, we may assume that the last k columns form a basis, *i.e.* the k-by-k matrix

$$M = \begin{pmatrix} \left.\frac{\partial h_1}{\partial x_{n-k+1}}\right|_z & \cdots & \left.\frac{\partial h_1}{\partial x_n}\right|_z \\ \vdots & & \\ \left.\frac{\partial h_k}{\partial x_{n-k+1}}\right|_z & \cdots & \left.\frac{\partial h_k}{\partial x_n}\right|_z \end{pmatrix} : \mathbb{R}^k \to \mathbb{R}^k \tag{3.18}$$

is invertible. Now consider the function:

$$\begin{aligned} f &: U \to \mathbb{R}^n \\ &x \mapsto (x_1, ..., x_{n-k}, h_1(x), ..., h_k(x)) \end{aligned}$$

The derivative of f at our point z is an n-by-n matrix of the form

$$Df|_z = \left(\begin{array}{c|c} I_{n-k} & 0 \\ \hline \star & M \end{array} \right)$$

where M is the matrix (3.18), and I_{n-k} is the identity matrix. Hence $\det Df|_z = \det M$ which is non-zero by assumption. Applying the Inverse Function Theorem (Theorem 3.9), we see that there is an open set $V \subset U$ containing z such that the function

$$f : V \to f(V) \subset \mathbb{R}^n$$

is a diffeomorphism. Then $\pi \circ f = h$ by construction. □

Notice that this proof tells us exactly how to construct the required co-ordinate system: for the last k co-ordinates we use $h_1, ..., h_k$, and for the first $n-k$ co-ordinates we use an appropriate subset of x_i's. Remember that we can't necessarily use $x_1, ..., x_{n-k}$ because of the re-ordering step in the proof.

Now we can prove that level sets at regular values are always submanifolds.

Proof of Proposition 3.16. Pick a regular value $\alpha \in \mathbb{R}^k$ of the function $h : \mathbb{R}^n \to \mathbb{R}^k$. Pick any point $z \in Z_\alpha$, so z is a regular point. By the Implicit Function Theorem (Theorem 3.17) there is a chart $f : V \xrightarrow{\sim} \tilde{V}$ around the point z in which h is just the standard projection function $\pi : \mathbb{R}^n \to \mathbb{R}^k$ (restricted to $\tilde{V}$). Then

$$f(Z_\alpha \cap V) = \tilde{V} \cap \pi^{-1}(\alpha)$$

and $\pi^{-1}(\alpha)$ is an affine subspace of $\mathbb{R}^n$ of dimension $n-k$. So the submanifold condition holds at all points $z \in Z_\alpha$. □

Example 3.19. Let's prove that S^n is a submanifold of $\mathbb{R}^{n+1}$. Consider the smooth function:

$$\begin{aligned} &h : \mathbb{R}^{n+1} \to \mathbb{R} \\ &h : (x_0, ..., x_n) \mapsto x_0^2 + ... + x_n^2 \end{aligned}$$

The derivative of h at a point $(x_0, ..., x_n)$ is the 1-by-$(n+1)$ matrix

$$(2x_0, ..., 2x_n) : \mathbb{R}^{n+1} \to \mathbb{R}$$

which is a surjection if at least one of the x_i is not zero. Hence the origin is the only critical point of h, and $0 \in \mathbb{R}$ is the only critical value. For $\alpha > 0$, the level set $Z_\alpha = h^{-1}(\alpha)$ is a n-dimensional sphere of radius $\sqrt{\alpha}$. So this sphere is a codimension-1 submanifold of $\mathbb{R}^{n+1}$.

If $\alpha < 0$ then $Z_\alpha = \phi$ (the empty set), technically this is also a codimension-1 submanifold of $\mathbb{R}^n$, but it's less interesting!

We can generalize slightly by considering level sets of functions which are only defined in an open set of $\mathbb{R}^n$. Take an open set $X \subset \mathbb{R}^n$, and consider a smooth function

$$h : X \to \mathbb{R}^k$$

(and recall from Example 2.13 that X is an n-dimensional manifold). Because the proof of Proposition 3.16 was entirely local, it shows immediately that the level sets

$$h^{-1}(\alpha) \subset X$$

are submanifolds of X, provided that α is a regular value.

Example 3.20. Consider the smooth function:

$$\begin{aligned} r : \mathbb{R}^2 \setminus (0,0) &\to \mathbb{R} \\ (x,y) &\mapsto \sqrt{x^2+y^2} \end{aligned}$$

Now let X be the open set

$$X = \{(x,y,z);\ (x,y) \neq (0,0)\} \subset \mathbb{R}^3$$

and let h be the smooth function:

$$\begin{aligned} h : X &\to \mathbb{R} \\ (x,y,z) &\mapsto (r-2)^2 + z^2 \end{aligned}$$

The derivative of h at (x,y,z) is the 1-by-3 matrix:

$$(2x(r-2)/r,\ 2y(r-2)/r,\ 2z) : \mathbb{R}^3 \to \mathbb{R}$$

This only fails to be a surjection if $z = 0$ and $r = 2$, meaning that the only critical value of h is $\alpha = 0$. For any other value of α, the level set Z_α is a 2-dimensional manifold.

If α lies in the interval $(0,4)$ then Z_α is the surface-of-revolution of the graph drawn in Figure 6, so it's a 2-dimensional torus.

Proposition 3.16 gives us an easy way to find new manifolds: just pick any smooth function from $\mathbb{R}^n$ (or an open set in $\mathbb{R}^n$) to $\mathbb{R}^k$ and then look at one of the level sets Z_α. Provided that α is a regular value, the level set Z_α is a submanifold of $\mathbb{R}^n$, and then by Proposition 3.7 Z_α automatically gets the structure of a manifold.

How can we find some charts on Z_α? In Section 3.1 we discussed how to get charts on a general submanifold: we have to find a chart on the ambient manifold which maps the submanifold to the standard affine subspace. In the case of level sets in $\mathbb{R}^n$ this is easy, because the Implicit Function Theorem tells us an explicit way to do it.

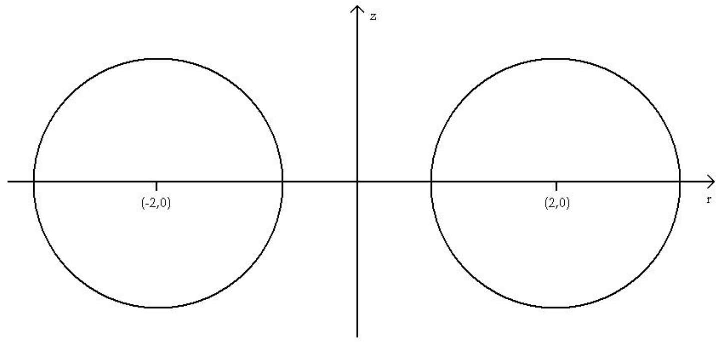

Figure 6: Level set of $(r-2)^2 + z^2$.

Example 3.21. Let's consider the submanifold $S^n = h^{-1}(1) \subset \mathbb{R}^{n+1}$, where h is the function from Example 3.19, and let's find some charts on S^n by inducing them from charts on $\mathbb{R}^{n+1}$. In Example 3.5 we did this in an *ad hoc* way for the case $n = 1$, now let's do it in a systematic way for all n by following the proof of the Implicit Function Theorem.

Let's start by finding a chart on S^n which includes the point $z = (0, ..., 0, 1)$. The derivative of h at this point z is the matrix $(0, ..., 2)$, and the final entry spans the 1-dimensional vector space $\mathbb{R}$. So if we set

$$f = (x_0, ..., x_{n-1}, h) : \mathbb{R}^{n+1} \to \mathbb{R}^{n+1}$$

then $\det(Df)|_x$ is not zero, and hence f must be a diffeomorphism on some open neighbourhood of z. In fact we can use the neighbourhood $U = \{x_n > 0\}$; it's clear that f is injective on U, and $\det(Df) = 2x_n$ is never zero in this set, so Corollary 3.10 applies. Hence (U, f) is a chart on $\mathbb{R}^{n+1}$.

This chart maps S^n to the affine subspace $\{\tilde{x}_n = 1\}$, but we can just compose it with the obvious translation τ to get a new chart $(U, \tau \circ f)$ which maps S^n to the standard subspace $\mathbb{R}^n \subset \mathbb{R}^{n+1}$. Then we get our induced chart on S^n, its domain is $V = S^n \cap \{x_n > 0\}$, and the co-ordinates are simply:

$$\begin{gathered} g : V \to \mathbb{R}^n \\ (x_0, ..., x_n) \mapsto (x_0, ..., x_{n-1}) \end{gathered}$$

The codomain $\tilde{V}$ of this chart is the open unit ball.

We could repeat this procedure with the obvious variations, and we end up with an atlas for S^n with $2(n+1)$ charts, having domains $\{\pm x_i > 0\} \cap S^n$. There is no need to check if this atlas is smooth because it follows automatically from Lemma 3.6. In the case $n = 1$, Lemma 3.6 also tells us that this atlas is compatible with the 'polar co-ordinates' atlas from Example 3.5.

However we should verify that this atlas is compatible with the stereographic projection atlas from Example 2.5, so they both define the same smooth structure. This is a simple exercise in computing transition functions.

Example 3.22. Let h be the function $h(x, y, z) = (r-2)^2 + z^2$ from Example 3.20, defined on the open set $X = \{(x, y) \neq (0, 0)\} \subset \mathbb{R}^3$. Let Z be the level set

$Z = h^{-1}(1)$. We saw that this was a submanifold, so let's find some charts on it.

First consider the function

$$f_1 = (x, y, h - 1) : \mathbb{R}^3 \to \mathbb{R}^3$$

(note that we're using $h - 1$ instead of h, *i.e.* we've done the translation step already). Then $\det(Df_1) = \partial_z h = 2z$. It's clear that f_1 is an injection on the set $U_1 = \{z > 0\}$, so it is a diffeomorphism on that subset. Hence (U_1, f_1) is a chart on X, and it maps Z to the standard subspace $\mathbb{R}^2 \subset \mathbb{R}^3$. The induced chart on Z is

$$\begin{aligned} g_1 : V_1 = Z \cap \{z > 0\} &\longrightarrow \mathbb{R}^2 \\ (x, y, z) &\mapsto (x, y) \end{aligned}$$

and this chart has codomain the open annulus $\tilde{V}_1 = \{1 < r < 3\}$.

Intuitively, critical points and critical values are rather rare. If we pick a point $x \in \mathbb{R}^n$ 'at random', then it is vanishingly unlikely that the derivative $Dh|_x$ is not surjective, since 'almost all' k-by-n matrices are surjective (provided that $k \leq n$). This suggests that 'almost all' level sets of a smooth function are submanifolds. This intuition is correct, and can be turned into a result known as Sard's Theorem. However it would take us on a significant detour to even state this result precisely.

4 Smooth functions

4.1 Definition of a smooth function

Suppose we have a manifold X, and a function:

$$h : X \to \mathbb{R}$$

If $X = \mathbb{R}^n$, or an open set in $\mathbb{R}^n$, then know what it means to say that h is a smooth function. If X is an arbitrary manifold, how should we decide if h is 'smooth' or not?

The answer is that we should 'look at h in co-ordinates'. If we pick a co-ordinate chart

$$f : U \xrightarrow{\sim} \tilde{U} \subset \mathbb{R}^n$$

on X then we can consider the function

$$\tilde{h} = h \circ f^{-1} : \tilde{U} \to \mathbb{R}$$

We should think of $\tilde{h}$ as the function h written in this choice of co-ordinates. If $\tilde{h}$ is a smooth function, then we should declare that h is also smooth, at least within the open set $U \subset X$. If we want to be more specific then we can choose a particular point $x \in U$, and declare that h is *smooth at x* iff the function $\tilde{h}$ is smooth at the point $f(x)$.

However, there might be a problem with this definition: it might depend on which co-ordinates we chose. Let's check that it doesn't. Suppose that (U_1, f_1) and (U_2, f_2) are two co-ordinate charts on X, both containing the point x. Then the functions

$$\tilde{h}_1 = h \circ f_1^{-1} \qquad \text{and} \qquad \tilde{h}_2 = h \circ f_2^{-1}$$

are related by the transition function ϕ_{12} between the two charts:

$$\tilde{h}_2 = \tilde{h}_1 \circ \phi_{12}$$

(note that this equality only makes sense on the open set $f_2(U_1 \cap U_2) \subset \tilde{U}_2$ where both sides are defined). Since ϕ_{12} is a diffeomorphism, the function $\tilde{h}_2$ is smooth at the point $f_2(x)$ iff the function $\tilde{h}_1$ is smooth at the point $f_1(x)$.

So if h looks smooth (at x) in one co-ordinate chart, then it will look smooth (at x) in any co-ordinate chart. Let's record this definition formally:

Definition 4.1. Let X be a manifold, let

$$h : X \to \mathbb{R}$$

be a function, and let x be a point in X. We say that h is **smooth at** x iff, for any co-ordinate chart

$$f : U \to \tilde{U} \subset \mathbb{R}^n$$

with $x \in U$, the function

$$h \circ f^{-1} : \tilde{U} \to \mathbb{R}$$

is smooth at the point $f(x)$. We say that h is **smooth everywhere**, or simply **smooth**, iff h is smooth at all points $x \in X$.

Notice that h is smooth everywhere iff, for any co-ordinate chart (U, f), the function $h \circ f^{-1}$ is smooth. However if want to check if h is smooth we don't need to check every co-ordinate chart, it's enough to pick an atlas $\{(U_i, f_i)\}$ for X, and verify that each function $h \circ f_i^{-1}$ is smooth. Also notice that if $X = \mathbb{R}^n$, or an open set in $\mathbb{R}^n$ (with the standard smooth structure), it's clear that this definition of a smooth function agrees with the ordinary definition.

Example 4.2. Let $X = S^1 = \left\{x^2 + y^2 = 1\right\} \subset \mathbb{R}^2$, and let:

$$\begin{aligned} h : S^1 &\to \mathbb{R} \\ (x, y) &\mapsto x^2 \end{aligned}$$

Let's verify that h is a smooth function. In Example 3.21 we constructed a convenient atlas on S^1, with four charts. The first chart has domain $U_1 = S^1 \cap \{y > 0\}$, codomain $\tilde{U}_1 = (-1, 1) \subset \mathbb{R}$, and co-ordinates $f_1 : (x, y) \mapsto x$. In this chart the function h becomes

$$\begin{aligned} \tilde{h}_1 = h \circ f_1^{-1} : (-1, 1) &\to \mathbb{R} \\ x &\mapsto x^2 \end{aligned}$$

which is smooth at all points. The second chart has domain $U_2 = S^1 \cap \{x > 0\}$, codomain $\tilde{U}_2 = (-1, 1)$ again, and co-ordinates $f_2 : (x, y) \mapsto y$. Since f_2^{-1} sends y to $(\sqrt{1 - y^2}, y)$, we see that in this chart h becomes

$$\tilde{h}_2 : y \mapsto 1 - y^2$$

which is also smooth at all points. The remaining two charts (which have domains $S^1 \cap \{y < 0\}$ and $S^1 \cap \{x < 0\}$) are similar.

Evidently we could replace h with any other smooth function of x and y in this example, and it would again define a smooth function on S^1.

It's obvious how to generalize Definition 4.1 to get a definition of when a function

$$h : X \to \mathbb{R}^k$$

is smooth. But why stop there? What we really need is a definition of a smooth function between *any* two manifolds.

Suppose X is a manifold of dimension n, and Y is a manifold of dimension k, and we have a function:

$$H : X \to Y$$

How should we decide if H is 'smooth' or not? Or more specifically, if we pick a point $x \in X$, how should we decide if H is 'smooth' at this point?

Again, we need to look at H in co-ordinates. So, pick a co-ordinate chart on X

$$f : U \xrightarrow{\sim} \tilde{U} \subset \mathbb{R}^n$$

with $x \in U$, and pick a co-ordinate chart on Y

$$g : V \xrightarrow{\sim} \tilde{V} \subset \mathbb{R}^k$$

with $H(x) \in V$. To study the function H in co-ordinates (near the point x) we should look at the composition

$$g \circ H \circ f^{-1} : \tilde{U} \to \tilde{V}$$

(note that we need to assume that $H(U)$ is contained in V for this to be defined). This is a function between an open set in $\mathbb{R}^n$ and an open set in $\mathbb{R}^k$, so it makes sense to ask if it is a smooth function.

Definition 4.3. Let X and Y be two manifolds, of dimension n and k, and let

$$H : X \to Y$$

be a continuous function. Fix a point $x \in X$. We say that H is **smooth at** x iff there exists a chart (U, f) on X containing x, and a chart (V, g) on Y containing all of $H(U)$, and the function

$$g \circ H \circ f^{-1} : \tilde{U} \to \tilde{V}$$

is smooth at the point $f(x)$.

We say that H is **smooth everywhere**, or just **smooth**, iff H is smooth at x for every point $x \in X$.

A smooth function is automatically continuous (exercise). Conversely, if you assume that H is continuous then finding charts such that $H(U) \subset V$ is easy: just pick any chart (U', f') on X, and any chart (V, g) on Y. Then $U = H^{-1}(V) \cap U'$ is an open set in U (since H is continuous), and restricting f' to this open subset gives a smaller chart with the required property.

If we set $Y = \mathbb{R}$ then it's easy to check that this definition agrees with Definition 4.1. Recall that in that case we observed that if a function 'looks smooth in one co-ordinate chart' then it 'looks smooth in all co-ordinate charts'. This is still true for our more general definition of a smooth function. Suppose (U_1, f_1) and (U_2, f_2) are two co-ordinate charts on X, and let ϕ_{21} be the transition function between them. Now let (V_1, g_1) and (V_2, g_2) be two co-ordinate charts on Y, and let ψ_{21} be the transition function between them. Assume that $H(U_1) \subset V_1$ and $H(U_2) \subset V_2$. Then we have an equality of functions

$$g_2 \circ H \circ f_2^{-1} = \psi_{21} \circ (g_1 \circ H \circ f_1^{-1}) \circ \phi_{12} \tag{4.4}$$

(on the open subset in $\tilde{U}_2$ where both sides are defined). Since all the transition functions are diffeomorphisms, the function $g_2 \circ H \circ f_2^{-1}$ will be smooth iff the function $g_1 \circ H \circ f_1^{-1}$ is smooth.

Example 4.5. Let $X = T^1$ (from Example 2.11) and $Y = S^1$ and

$$\begin{aligned} H : T^1 &\to S^1 \\ [t] &\mapsto (\cos 2\pi t, \sin 2\pi t) \end{aligned}$$

(note that this is well-defined). Let's show that H is smooth.

Start with the chart on T^1 having domain $U' = T^1 \setminus [0]$, codomain $\tilde{U}' = (0, 1) \subset \mathbb{R}$, and co-ordinates $f = q^{-1} : U' \xrightarrow{\sim} \tilde{U}'$. Now take a chart on S^1 with domain $V = S^1 \cap \{y > 0\}$, codomain $\tilde{V} = (-1, 1)$, and co-ordinates $g : (x, y) \mapsto x$ (as in Example 3.21). We have

$$H(U') = S^1 \setminus (1, 0)$$

which is not contained in V, but we can shrink the chart (U', f) to correct this: just set $\tilde{U} = (0, \frac{1}{2})$ and $U = q(\tilde{U})$, and restrict f to U. Then in these charts H

becomes:

$$\tilde{H} = g \circ H \circ f^{-1} \colon (0, \tfrac{1}{2}) \to (-1, 1)$$
$$t \mapsto \cos 2\pi t$$

This is smooth, which shows that H is smooth at all points in U. We could check that H is smooth at the remaining points by using variations on these charts.

The next result should not be surprising.

Lemma 4.6. *Let X, Y and Z be three manifolds, and let*

$$H : X \to Y \qquad \textit{and} \qquad G : Y \to Z$$

be smooth functions. Then $G \circ H$ is smooth.

Proof. Fix a point $x \in X$. Now pick a co-ordinate chart (U, f) on X containing the point x, a co-ordinate chart (V, g) on Y containing the point $H(x)$, and a co-ordinate chart (W, h) on Z containing the point $G(H(x))$. Since H is smooth, the function $g \circ H \circ f^{-1}$ is a smooth function, defined on some open neighbourhood of the point $f(x) \in \mathbb{R}^n$ (where n is the dimension of X). Similarly since G is smooth, the function $h \circ G \circ g^{-1}$ is a smooth function, defined on some open neighbourhood of the point $g(H(x)) \in \mathbb{R}^k$ (where k is the dimension of Y). To prove that $G \circ H$ is smooth at x, we need to know that the function

$$h \circ G \circ H \circ f^{-1}$$

is smooth at the point $f(x)$. But in a sufficiently small open neighbourhood of $f(x)$ we can factor this function as

$$h \circ G \circ H \circ f^{-1} = (h \circ G \circ g^{-1}) \circ (g \circ H \circ f^{-1})$$

and both factors are smooth. □

If you know what a *category* is, then this shows that there is a category whose objects are manifolds and whose arrows are smooth functions.

Lemma 4.7. *Let Z be a submanifold of X, and let*

$$\iota : Z \hookrightarrow X$$

be the inclusion function. Then ι is smooth.

Proof. Exercise. □

So if $H : X \to Y$ is smooth, and $Z \subset X$ is a submanifold, then the restriction $H|_Z : Z \to Y$ is automatically smooth (just combine Lemmas 4.7 and 4.6). For example, any smooth function $h : \mathbb{R}^2 \to \mathbb{R}$ restricts to give a smooth function $h : S^1 \to \mathbb{R}$, as we checked explicitly in Example 4.2.

Here is a another result that is often useful for proving that a function is smooth:

Lemma 4.8. *Let $H : X \to Y$ be a smooth function between two manifolds, and let $Z \subset Y$ be a submanifold of Y. Suppose that the image of H is contained in Z. Then H defines a smooth function from X to Z.*

Proof. Exercise. ☐

Example 4.9. Let's do Example 4.5 again using the previous result. Define a function:

$$\begin{aligned}\hat{H} &: T^1 \to \mathbb{R}^2 \\ [t] &\mapsto (\cos 2\pi t, \sin 2\pi t)\end{aligned}$$

We can check that $\hat{H}$ is smooth by looking at it in charts on T^1, then by Lemma 4.8 we know that $\hat{H}$ defines a smooth function $H : T^1 \to S^1$ (we don't need to use explicit charts on S^1).

4.2 The rank of a smooth function

We now begin to think about an extremely important concept: the *derivative* of a smooth function. It will take us a long time to really get to grips with this idea.

Suppose we have smooth functions

$$F : \mathbb{R}^n \to \mathbb{R}^k \qquad \text{and} \qquad G : \mathbb{R}^k \to \mathbb{R}^m$$

and we form their composition $G \circ F : \mathbb{R}^n \to \mathbb{R}^m$. If we pick a point $x \in \mathbb{R}^n$ then the derivative of $G \circ F$ at x is a linear map

$$D(G \circ F)|_x : \mathbb{R}^n \to \mathbb{R}^m$$

and you should recall that the formula

$$D(G \circ F)|_x = DG|_{F(x)} \circ DF|_x$$

holds. This is nothing but the chain rule for functions of more than one variable. Of course the formula still holds if F is only defined in some open neighbourhood of x, and G is only defined in some open neighbourhood of $F(x)$.

In particular, if $n = k = m$, and $G = F^{-1}$, we get that

$$D(F^{-1})|_{F(x)} = (DF|_x)^{-1}$$

since the derivative of the identity function $\mathbb{R}^n \to \mathbb{R}^n$ at any point is the identity linear map. So if F is a diffeomorphism then the derivative of F is an isomorphism at all points. This is the (much easier!) converse to the Inverse Function Theorem.

Now suppose we have two manifolds X and Y, of dimensions n and k respectively, and we have a smooth function:

$$F : X \to Y$$

Fix a point $x \in X$. Let's write F in co-ordinates near the point x, so we pick a co-ordinate chart (U_1, f_1) on X containing the point x, and a co-ordinate chart (V_1, g_1) on Y containing the point $F(x)$, and we consider the function:

$$\tilde{F}_1 = g_1 \circ F \circ f_1^{-1}$$

This is defined in some open neighbourhood of the point $f_1(x) \in \mathbb{R}^n$, and it lands in $\mathbb{R}^k$. This means we can take the derivative of this function at the point $f_1(x)$, it is some linear map:

$$D\tilde{F}_1|_{f_1(x)} : \mathbb{R}^n \to \mathbb{R}^k$$

What happens if we change co-ordinates? If we pick new charts (U_2, f_2) (containing x) and (V_2, f_2) (containing $F(x)$), then our function becomes:

$$\tilde{F}_2 = g_2 \circ F \circ f_2^{-1}$$

(which is is defined on some open neighbourhood of $f_2(x) \in \mathbb{R}^n$). The derivative of $\tilde{F}_2$ at the point $f_2(x)$ is also a linear map:

$$D\tilde{F}_2|_{f_2(x)} : \mathbb{R}^n \to \mathbb{R}^k$$

How are the two linear maps $D\tilde{F}_1|_{f_1(x)}$ and $D\tilde{F}_2|_{f_2(x)}$ related to each other?

We've already observed (4.4) that the functions $\tilde{F}_1$ and $\tilde{F}_2$ are related by the equation

$$\tilde{F}_2 = \psi_{21} \circ \tilde{F}_1 \circ \phi_{12} \tag{4.10}$$

where ϕ_{12} is the transition function between U_2 and U_1, and ψ_{21} is the transition function between V_1 and V_2 (we might have to restrict to a smaller open neighbourhood of $f_2(x)$ before this equation makes sense).

Now take the derivative of the equation (4.10) at the point $f_2(x)$. By the chain rule, we have:

$$D\tilde{F}_2|_{f_2(x)} = D\psi_{21}|_{g_1(F(x))} \circ D\tilde{F}_1|_{f_1(x)} \circ D\phi_{12}|_{f_2(x)} \tag{4.11}$$

So the linear maps $D\tilde{F}_1|_{f_1(x)}$ and $D\tilde{F}_2|_{f_2(x)}$ are not the same, but they are related by this formula. Now we can make an important observation: the linear maps $D\psi_{21}|_{g_1(H(x))}$ and $D\phi_{12}|_{f_2(x)}$ are isomorphisms, because the transition functions are diffeomorphisms. Therefore, the rank of $D\tilde{F}_2|_{f_2(x)}$ must be the same as the rank of $D\tilde{F}_1|_{f_1(x)}$.

This means we can make the following definition:

Definition 4.12. Let X and Y be manifolds (of dimensions n and k respectively) and let $F : X \to Y$ be a smooth function. Fix a point $x \in X$. Now pick a co-ordinate chart (U, f) containing x and a co-ordinate chart (V, g) containing $F(U)$, and consider the function:

$$\tilde{F} = g \circ F \circ f^{-1} : \tilde{U} \to \tilde{V}$$

We define the **rank** of F at x to be the rank of the derivative

$$D\tilde{F}|_{f(x)} : \mathbb{R}^n \to \mathbb{R}^k$$

of $\tilde{F}$ at $f(x)$.

This makes sense because of the formula (4.11); it doesn't matter which co-ordinate charts we choose, the rank of $D\tilde{F}|_{f(x)}$ will always be the same.

Now we can generalize Definitions 3.13 and 3.14.

Definition 4.13. Let $F : X \to Y$ be a smooth function between two manifolds, of dimensions n and k respectively. We say that a point $x \in X$ is a **regular point** of F if the rank of F at x is equal to k. If x is not a regular point then we call it a **critical point**.

We say that a point $y \in Y$ is a **regular value** of F if every point $x \in F^{-1}(y)$ is a regular point. If y is not a regular value then we call it a **critical value**.

So x is a regular point of F iff the derivative

$$D\tilde{F}|_{f(x)} : \mathbb{R}^n \to \mathbb{R}^k$$

is a surjection, where $\tilde{F}$ is F written in any co-ordinate charts. In other words x is a regular point of F iff $f(x)$ is a regular point of $\tilde{F}$, for any choice of co-ordinates. Clearly if we set X to be an open set in $\mathbb{R}^n$, and Y to be $\mathbb{R}^k$, then we recover our previous definitions.

We can also generalize Proposition 3.16 fairly easily:

Proposition 4.14. *Let $F : X \to Y$ be a smooth function between two manifolds, of dimensions n and k respectively. Let $y \in Y$ be a regular value of F. Then the level set*

$$Z_y = F^{-1}(y) \subset X$$

is a submanifold of X of codimension k.

Proof. Pick a point $x \in Z_y$, a co-ordinate chart (U, f) containing x, and a co-ordinate chart (V, g) containing $F(U)$, and consider the function:

$$\tilde{F} = g \circ F \circ f^{-1} : \tilde{U} \to \tilde{V}$$

Note that the level set $\tilde{F}^{-1}(g(y))$ is just $f(Z_y \cap U)$. Since y is a regular value of F, the point x must be a regular point of F, which means that $f(x)$ is a regular point of $\tilde{F}$. Now we can apply the Implicit Function Theorem and conclude that there is a chart (W, h) on $\mathbb{R}^n$, with $f(x) \in W \subset \tilde{U}$, such that:

$$h(\tilde{F}^{-1}(g(y)) = \mathbb{R}^{n-k} \cap \tilde{W}$$

Then we can use the co-ordinate chart $(f^{-1}(W), h \circ f)$ on X to demonstate that Z_y satisfies the submanifold condition at the point x. $\square$

Example 4.15. Consider the 2-sphere

$$S^2 = \left\{(x, y, z);\ x^2 + y^2 + z^2 = 1\right\} \subset \mathbb{R}^3$$

and let $F : S^2 \to \mathbb{R}$ be the function:

$$F : (x, y, z) \mapsto x$$

We have a chart on S^2 with domain $U_1 = S^2 \cap \{z > 0\}$ and co-ordinates $f_1 : (x, y, z) \mapsto (x, y)$, and in this chart F becomes:

$$\begin{aligned} \tilde{F}_1 : B(0,1) &\to \mathbb{R} \\ (x, y) &\mapsto x \end{aligned}$$

Then $D\tilde{F}_1 = (1,0)$ at all points, so the rank of F is 1 at all points in U. Similarly, it's easy to see that the rank of F is 1 at any point where $z < 0$, or any point where $y \neq 0$.

Now let's switch to the chart with domain $U_2 = S^2 \cap \{x > 0\}$ and co-ordinates $f_2 : (x, y, z) \mapsto (y, z)$. In this chart F becomes

$$\begin{aligned} \tilde{F}_2 : B(0,1) &\to \mathbb{R} \\ (y,z) &\mapsto \sqrt{1 - y^2 - z^2} \end{aligned}$$

and this has derivative:

$$D\tilde{F}_2|_{(y,z)} = \frac{-(y,z)}{\sqrt{1-y^2-z^2}}$$

This has rank 1 if $(y,z) \neq (0,0)$ (this had to be true from our calculations in other charts), but it has rank zero at $(0,0)$. So F has rank zero at the point $(1,0,0)$. A similar calculation shows that F also has rank zero at the point $(-1,0,0)$.

So the only critical points of F are $(\pm 1, 0, 0)$, and thus the only critical values of F are $\alpha = 1$ and $\alpha = -1$. If $|\alpha| < 1$ then the level set $F^{-1}(\alpha)$ is a circle, the intersection of the 2-sphere with the plane $\{x = \alpha\}$. Proposition 4.14 says that this is a 1-dimensional submanifold of S^2.

If $|\alpha| > 1$ then the level set $F^{-1}(\alpha)$ is empty. At the critical values $\alpha = \pm 1$ the level set consists of a single point, this is evidently not a 1-dimensional submanifold.

4.3 Some special kinds of smooth functions

We can use our definition of rank to single out some particularly important kinds of smooth functions.

Definition 4.16. A smooth function $F : X \to Y$ is called a **submersion** if the rank of F at any point is equal to the dimension of Y.

So F is a submersion iff the derivative at any point (in any co-ordinates) is a surjection, i.e. a submersion is exactly a smooth function that has no critical points. There is a dual notion to this:

Definition 4.17. A smooth function $F : X \to Y$ is called an **immersion** if the rank of F at any point $x \in X$ is equal to the dimension of X.

In other words, F is an immersion iff the derivative at any point (in any co-ordinates) is an injection.

Recall from Lemma 4.7 that the inclusion $\iota : Z \hookrightarrow X$ of a submanifold is a smooth function.

Lemma 4.18. *Let $Z \subset X$ be a submanifold of X. Then the inclusion map $\iota : Z \hookrightarrow X$ is an immersion.*

Proof. Exercise. □

This is the typical example of an immersion, however not every immersion is of this form.

Example 4.19. Let $X = \mathbb{R}$ and $Y = \mathbb{R}^2$, and let:

$$\begin{aligned} F : \mathbb{R} &\to \mathbb{R}^2 \\ t &\mapsto (t^2, t^3 - t) \end{aligned}$$

We can take the trivial co-ordinate charts on X and Y. The derivative of F at the point $t \in \mathbb{R}$ is the linear map

$$DF|_t = (2t, 3t^2 - 1) : \mathbb{R} \to \mathbb{R}^2$$

which is an injection for every t. Hence F is an immersion. The image of F is not a submanifold, the problem occurs at the 'intersection point' at $(1, 0) = F(\pm 1)$ (see Figure 7).

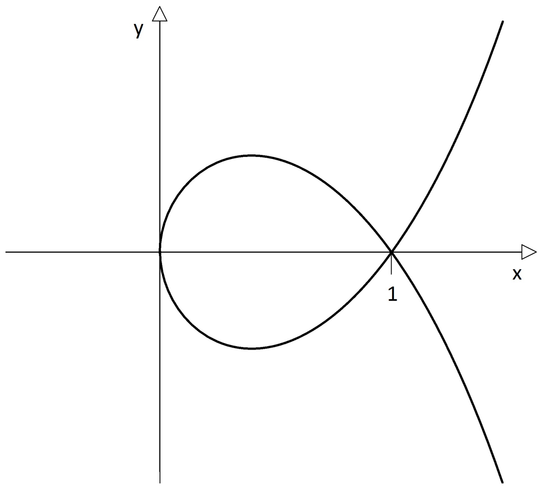

Figure 7: The image of an immersion need not be a submanifold.

One might hope that if F is an injective immersion then the image of F must be a submanifold, but this is not true either! For a counter-example, just restrict the function F from Example 4.19 to the open interval $(-\infty, 1) \subset \mathbb{R}$.

The following easy observation sometimes gives a convenient way of computing the rank of a smooth function.

Lemma 4.20. *Let $F : X \to Z$ be any smooth function and let $G : Z \to Y$ be an immersion. Pick any point $x \in X$. Then the rank of $G \circ F$ at x is the same as the rank of F at x.*

Proof. Exercise. □

In particular, suppose Z is a submanifold of Y and $F : X \to Z$ is a smooth function. Then it doesn't matter whether we view F as a function to Z, or as a function to Y that happens to land in Z; the rank is the same.

Example 4.21. In Example 4.5, and again in Example 4.9, we proved that the function

$$\begin{aligned} H : T^1 &\to S^1 \\ [t] &\mapsto (\cos 2\pi t, \sin 2\pi t) \end{aligned} \tag{4.22}$$

was smooth. Let's compute the rank of H at all points.

In Example 4.5 we saw that there are charts (U, f) on T^1 and (V, g) on S^1 in which H becomes:

$$\begin{aligned} \tilde{H} = g \circ H \circ f^{-1} \colon (0, \tfrac{1}{2}) &\to (-1, 1) \\ t &\mapsto \cos 2\pi t \end{aligned}$$

Then $D\tilde{H}|_t = 2\pi \sin 2\pi t$. This is never zero (in this domain), so H has rank 1 at all points in U. We can use other charts to check that in fact H has rank 1 at all points in T^1.

Alternatively, we can compose H with the inclusion $\iota : S^1 \hookrightarrow \mathbb{R}^2$ to get a function $\hat{H} =: T^1 \to \mathbb{R}^2$ (as in Example 4.9). Then Lemma 4.20 guarantees that the ranks of H and $\hat{H}$ are the same at all points, so let's compute the rank of $\hat{H}$ instead. This is slightly easier, since we don't need to pick charts on S^1.

If we pick any open set $\tilde{U} \subset \mathbb{R}$ on which the quotient map $q : \mathbb{R} \to T^1$ is an injection then we get a chart on T^1 with domain $U = q(\tilde{U})$ and co-ordinates q^{-1} (see problem sheets). If we write $\hat{H}$ in this chart it becomes a function

$$\hat{H} \circ q : \tilde{U} \to \mathbb{R}^2$$

defined by the same formula (4.22). This has derivative:

$$D(\hat{H} \circ q)|_t = 2\pi(-\sin 2\pi t, \cos 2\pi t) : \mathbb{R}^2 \to \mathbb{R}$$

This is never the zero matrix, so it has rank 1 at any t. Hence $\hat{H}$ has rank 1 at all points.

Our final class of smooth functions is perhaps the most important:

Definition 4.23. Let X and Y are two smooth manifolds. A function

$$F : X \to Y$$

is called a **diffeomorphism** if F is smooth, bijective, and the inverse function F^{-1} is also smooth. If there exists a diffeomorphism between X and Y then we say that X and Y are **diffeomorphic**.

If two manifolds are diffeomorphic then they are exactly the same, for all practical purposes (it may help to think of 'diffeomorphic' as another word for 'isomorphic'). Note that if X and Y are just open subsets of $\mathbb{R}^n$ then this reduces to our previous definition of 'diffeomorphism'.

Suppose $F : X \to Y$ is a diffeomorphism, and we want to look at in co-ordinates. So we pick a chart (U, f) on X, and a chart (V, g) on Y containing all of $F(U)$. But then $F(U)$ is an open subset of V (since F is a homeomorphism), so we can shrink V to this open set and get a smaller chart on on Y. In these co-ordinates, F becomes a diffeomorphism:

$$\tilde{F} = g \circ F \circ f^{-1} : \tilde{U} \xrightarrow{\sim} \tilde{V}$$

This implies that the derivative of $\tilde{F}$ at any point must be an isomorphism, so F is both a submersion and an immersion. In particular, diffeomorphic manifolds must have the same dimension, which is reassuring.

We can extend Corollary 3.10 to the following criterion for testing if a function is a diffeomorphism:

Lemma 4.24. *Let X and Y be n-dimensional manifolds, and let*

$$F : X \to Y$$

be a smooth bijection. If the rank of F is n at every point then F is a diffeomorphism.

Proof. We just need to show that the inverse function F^{-1} is smooth. Fix a point $y \in Y$, let $x = F^{-1}(y)$, and choose co-ordinate charts (U, f) containing x and (V, g) containing $F(U)$. Consider the function $\tilde{F} = g \circ F \circ f^{-1}$. Since the rank of F is n at the point $F(y)$, the derivative

$$D\tilde{F}|_{f(x)} : \mathbb{R}^n \to \mathbb{R}^n$$

is an isomorphism. By the Inverse Function Theorem, there is some open neighbourhood of $g(y)$ on which the function $\tilde{F}^{-1}$ is smooth. This proves that F^{-1} is smooth at y. □

Example 4.25. Recall from Example 4.21 that we have a smooth function

$$\begin{aligned} H : T^1 &\to S^1 \\ [t] &\mapsto (\cos 2\pi t, \sin 2\pi t) \end{aligned}$$

whose rank is 1 at all points. This H is obviously a bijection, it must be a diffeomorphism. Hence these two versions of the circle, T^1 and S^1, are diffeomorphic manifolds.

We leave our third version of the circle, $\mathbb{RP}^1$, as an exercise.

Example 4.26. Recall from Example 2.25 that we can find a 'non-standard' atlas $\mathcal{C}$ on the topological manifold $\mathbb{R}$ which is not compatible with the standard atlas $\mathcal{A}$. However, the two smooth manifolds $(\mathbb{R}, [\mathcal{A}])$ and $(\mathbb{R}, [\mathcal{C}])$ are diffeomorphic (exercise).

This leaves an interesting question: does there exist any smooth atlas $\mathcal{D}$ on $\mathbb{R}^n$ such that the resulting smooth manifold $(\mathbb{R}^n, [\mathcal{D}])$ is *not* diffeomorphic to the standard $\mathbb{R}^n$? This question was comprehensively answered in the 1980s, and the answer is one of the most astonishing results in all of mathematics.

5 Tangent spaces

If we have open sets $U \subset \mathbb{R}^n$ and $V \subset \mathbb{R}^m$ and a smooth function $F : U \to V$, we know that we can define the derivative of F at any point $x \in U$, and this is a linear map:

$$DF|_x : \mathbb{R}^n \to \mathbb{R}^m$$

If we replace U and V by arbitrary manifolds X and Y, then we don't yet know how to generalize this definition. We saw in the last section that it's possible to define the rank of the derivative of F at a point in X, but this is just a number. The aim of this section is to upgrade this number to a linear map.

We will show that to any point x in a manifold X there's an associated vector space, called the *tangent space* to X at x, and denoted T_xX. Then we'll show that if we have smooth map $F : X \to Y$ then we can define the derivative of F at x, and it's a linear map:

$$DF|_x : T_xX \to T_{F(x)}Y$$

In fact defining tangent spaces is the hard part, the fact that we can define $DF|_x$ will follow almost automatically.

5.1 Tangent vectors via curves

Roughly, a *tangent vector* to a point x in a manifold is a direction that you can go in from x. There are several ways to make this precise, the most intuitively appealing way is via equivalence classes of curves. We start by explaining how this works for the simplest kind of manifolds, namely open sets in $\mathbb{R}^n$.

Fix an open set $\tilde{U} \subset \mathbb{R}^n$, and pick a point $\tilde{x} \in \tilde{U}$. Let's declare that a *curve through* $\tilde{x}$ is a smooth function

$$\sigma = (\sigma_1, ..., \sigma_n) : (-\epsilon, \epsilon) \to \tilde{U}$$

with $\sigma(0) = \tilde{x}$, where here ϵ is some positive real number and $(-\epsilon, \epsilon) \subset \mathbb{R}$ is the corresponding open interval. This is indeed a smooth parametrized curve in $\tilde{U}$, passing through the point $\tilde{x}$. Note that we really mean the function σ and not just its image in $\tilde{U}$ (which would be an unparametrized curve).

The derivative of σ at the point $0 \in \mathbb{R}$ is a linear map:

$$D\sigma|_0 : \mathbb{R} \to \mathbb{R}^n$$

given by the n-by-1-matrix:

$$D\sigma|_0 = (\dot{\sigma}_1(0), ..., \dot{\sigma}_n(0))^\top$$

We're going to think of $D\sigma|_0$ as a column vector in $\mathbb{R}^n$ rather than as a matrix (or if you prefer, we're going to write $D\sigma|_0$ when we mean $D\sigma|_0(1)$). Of course, this is just the tangent vector to σ when it hits the point $\tilde{x}$. It's the 'direction' that σ is travelling when it passes through $\tilde{x}$, or more accurately it's the 'velocity' of σ, since we don't forget the length of $D\sigma|_0$.

Now suppose that we have two curves through $\tilde{x}$:

$$\sigma : (-\epsilon_1, \epsilon_1) \to \tilde{U}$$
$$\tau : (-\epsilon_2, \epsilon_2) \to \tilde{U}$$

Let's declare that σ and τ are *tangent at $\tilde{x}$* iff they have the same tangent vector at this point, so $D\sigma|_0 = D\tau|_0$. This is restrictive use of the word 'tangent', since we're requiring that the the tangent vectors are actually equal and not just proportional. For example if τ was a reparametrization of σ, then under our definition it probably wouldn't be tangent to σ at $\tilde{x}$.

Obviously being tangent at $\tilde{x}$ is an equivalence relation on curves through $\tilde{x}$, and by definition we have a well-defined function

$$\Delta : \{\text{curves through } \tilde{x}\} / (\text{tangency at } \tilde{x}) \longrightarrow \mathbb{R}^n$$

sending each equivalence class $[\sigma]$ to its tangent vector $D\sigma|_0$. By definition this function Δ is an injection. It's also a surjection, this is because for any vector $v \in \mathbb{R}^n$ we can consider a straight line

$$\begin{aligned} \sigma_v : \mathbb{R} &\to \mathbb{R}^n \\ t &\mapsto \tilde{x} + vt \end{aligned} \tag{5.1}$$

and if ϵ is small enough this defines a function:

$$\sigma_v : (-\epsilon, \epsilon) \to \tilde{U}$$

This is a curve through $\tilde{x}$, and obviously $\Delta(\sigma_v) = v$. So Δ is a bijection of sets.

Now we want to generalize this to other manifolds. Let X be an n-dimensional manifold, and pick a point $x \in X$.

Definition 5.2. A **curve through** x is a smooth function

$$\sigma : (-\epsilon, \epsilon) \to X$$

with $\sigma(0) = x$, where ϵ is any positive real number.

We want to find a definition of when two curves through x are tangent to each other. As usual, we need to look at our curves in co-ordinates.

Suppose we have a curve σ through a point $x \in X$, and we pick a chart (U, f) around the point x. Then in these co-ordinates σ becomes a curve

$$\tilde{\sigma} = f \circ \sigma$$

in $\tilde{U}$, through the point $f(x)$ (we might have to shrink ϵ to ensure that the image of σ lies within U). This curve has an associated tangent vector:

$$D\tilde{\sigma}|_0 \in \mathbb{R}^n$$

However, this vector in $\mathbb{R}^n$ is not independent of our co-ordinates, it will change when we change charts. If (U_1, f_1) and (U_2, f_2) are two charts around x, and we write σ in both sets of co-ordinates (possibly after shrinking ϵ to make sure that σ lands in $U_1 \cap U_2$), the answers are related by the equation

$$\tilde{\sigma}_2 = \phi_{21} \circ \tilde{\sigma}_1$$

where ϕ_{21} is the transition function between the two charts. If we use our first chart, the tangent vector to σ would be the vector $D\tilde{\sigma}_1|_0 \in \mathbb{R}^n$, and if we use our second chart, it would be $D\tilde{\sigma}_2|_0$. By the chain rule, we have that

$$D\tilde{\sigma}_2|_0 = D\phi_{21}|_{f_1(x)}\big(D\tilde{\sigma}_1|_0\big) \tag{5.3}$$

so these two vectors are related by the invertible linear map:

$$D\phi_{21}|_{f_1(x)} : \mathbb{R}^n \longrightarrow \mathbb{R}^n$$

So given a single curve σ, it makes no sense to ask what the tangent vector to σ is (as a vector in $\mathbb{R}^n$ at least) because the answer depends on your choice of chart. However, if you have two curves σ and τ, through the same point $x \in X$, it does make sense to ask if they have the *same* tangent vector.

Let's check this. Choose a chart (U_1, f_1) around x, and write both curves in this chart, so we get curves $\tilde{\sigma}_1$ and $\tilde{\tau}_1$ through the point $f_1(x) \in \tilde{U}_1$. Now pick a second chart (U_2, f_2) and again write both curves in co-ordinates, so we get curves $\tilde{\sigma}_2$ and $\tilde{\tau}_2$ through the point $f_2(x) \in \tilde{U}_2$. Then (5.3) implies that

$$D\tilde{\sigma}_1|_0 = D\tilde{\tau}_1|_0 \quad \iff \quad D\tilde{\sigma}_2|_0 = D\tilde{\tau}_2|_0$$

i.e. the curves $\tilde{\sigma}_1$ and $\tilde{\tau}_1$ are tangent at $f_1(x)$ if and only if the curves $\tilde{\sigma}_2$ and $\tilde{\tau}_2$ are tangent at $f_2(x)$. So we may make the following definition:

Definition 5.4. Fix a point x in a manifold X. We say that two curves σ, τ through x are **tangent at** x iff for any co-ordinate chart (U, f) containing x, we have:

$$D(f \circ \sigma)|_0 = D(f \circ \tau)|_0$$

As we have shown, if this holds in one chart then it holds in all charts. Obviously 'being tangent at x' is an equivalence relation on the set of all curves through x.

Definition 5.5 (*'Geometer's definition'*)**.** Fix a point x in a manifold X. A **tangent vector** to x is an equivalence class of curves through x. We denote the set of all tangent vectors to x by

$$T_xX = \{\text{curves through } x\} / (\text{tangency at } x)$$

and call it the **tangent space** to X at x.

If we fix a chart (U, f) around x then we can identify the tangent space T_xX with $\mathbb{R}^n$, because given a curve through x we can write in co-ordinates and then look at its tangent vector at $f(x)$. This defines a function:

$$\begin{aligned} \Delta_f : T_xX &\longrightarrow \mathbb{R}^n \\ [\sigma] &\mapsto D(f \circ \sigma)|_0 \end{aligned} \tag{5.6}$$

By the definition of T_xX, this function Δ_f is well-defined, and an injection. It's also a surjection, because given any $v \in \mathbb{R}^n$ we have a 'straight-line' curve

$$\begin{aligned} \tilde{\sigma}_v : (-\epsilon, \epsilon) &\to \tilde{U} \\ t &\mapsto f(x) + vt \end{aligned}$$

(for small-enough ϵ) as in (5.1), and then $\sigma_v = f^{-1} \circ \tilde{\sigma}_v$ is a curve through x such that $\Delta_f(\sigma_v) = v$. Hence Δ_f is a bijection of sets.

However, this identification of T_xX with $\mathbb{R}^n$ does depend on our choice of chart. If (U_1, f_1) and (U_2, f_2) are two charts containing x, then we get two different bjiections:

$$\Delta_{f_1} : T_xX \xrightarrow{\sim} \mathbb{R}^n \qquad \text{and} \qquad \Delta_{f_2} : T_xX \xrightarrow{\sim} \mathbb{R}^n$$

By (5.3), these two bijections are related by the equation

$$\Delta_{f_2} = D\phi_{21}|_{f_1(x)} \circ \Delta_{f_1} \tag{5.7}$$

where ϕ_{21} is the transition function between our two charts.

If we had some canonical bijection between T_xX to $\mathbb{R}^n$ (*i.e.* without making any choices) then obviously T_xX would have the structure of an n-dimensional vector space. But we don't have a canonical bijection, instead we have one bijection Δ_f for each choice of chart around x, and there is no way to single out any special one. Nevertheless, this is enough to define a vector space structure on T_xX.

Proposition 5.8. *If x is a point in an n-dimensional manifold X, then the tangent space T_xX is an n-dimensional vector space.*

Proof. Pick a co-ordinate chart (U, f) containing x, so we get a bijection Δ_f between T_xX and $\mathbb{R}^n$ as in (5.6). We can use this bijection to put a vector space structure on T_xX, *i.e.* we can define an addition operation

$$[\sigma] + [\tau] = \Delta_f^{-1}\big(\Delta_f(\sigma) + \Delta_f(\tau)\big)$$

and a scalar multiplication

$$\lambda[\sigma] = \Delta_f^{-1}\big(\lambda\Delta_f(\sigma)\big)$$

and these are guaranteed to satisfy the vector space axioms. We claim that this vector space structure on T_xX is independent of our choice of chart.

To see this, pick two charts (U_1, f_1) and (U_2, f_2) around x, and let ϕ_{21} be the transition function. We get two bijections Δ_{f_1} and Δ_{f_2}, related by the derivative of the transition function (5.7). Since $D\phi_{21}|_x$ is a linear isomorphism, it follows formally that the values of $[\sigma] + [\tau]$ and $\lambda[\sigma]$ are independent of whether we use Δ_{f_1} or Δ_{f_2}. ◻

So we have achieved our first aim for this section, namely to any point x in a manifold X we have attached a vector space T_xX. If we choose a chart (U, f) around x then we get a bijection

$$\Delta_f : T_xX \xrightarrow{\sim} \mathbb{R}^n$$

and this is a linear isomorphism (by definition). But if we have two different charts then we get two different isomorphisms of T_xX with $\mathbb{R}^n$, related by the equation (5.7).

Now we move on to our second aim: defining the derivative of a smooth function between two manifolds.

Firstly, suppose that $\tilde{U}$ and $\tilde{V}$ are open sets in $\mathbb{R}^n$ and $\mathbb{R}^m$ respectively, and that F is a smooth function:

$$F : \tilde{U} \to \tilde{V}$$

Pick a point $\tilde{x} \in \tilde{U}$ and let $\tilde{y} = F(\tilde{x})$. If we have a curve σ through $\tilde{x}$, then the composition $F \circ \sigma$ is a curve through $\tilde{y}$ (since the composition of two smooth functions is smooth). Furthermore, using the chain rule again tells us that:

$$D(F \circ \sigma)|_0 = DF|_{\tilde{x}}\big(D\sigma|_0\big) \tag{5.9}$$

In particular, the tangent vector to the curve $F \circ \sigma$ only depends on the tangent vector to σ. So we have a well-defined function:

$$\frac{\{\text{curves through } \tilde{x}\}}{(\text{tangency at } \tilde{x})} \longrightarrow \frac{\{\text{curves through } \tilde{y}\}}{(\text{tangency at } \tilde{y})}$$
$$[\sigma] \mapsto [F \circ \sigma]$$

We can identify the domain of this function with $\mathbb{R}^n$, and its codomain with $\mathbb{R}^m$, by sending curves to their tangent vectors. Then this function is just the derivative $DF|_{\tilde{x}}$.

Now we can generalize this to any smooth function between manifolds.

Proposition 5.10. *Let X and Y be manifolds of dimensions n and m, and let $F : X \to Y$ be a smooth function. Fix a point $x \in X$, and let $y = F(x)$. Then we have a well-defined function*

$$\begin{aligned} DF|_x : T_xX &\longrightarrow T_yY \\ [\sigma] &\mapsto [F \circ \sigma] \end{aligned}$$

and $DF|_x$ is linear.

We call this linear map the **derivative of F at x**.

Proof. Pick a chart (U, f) on X containing x, and a chart (V, g) on Y containing y. In these charts, F becomes the function:

$$\tilde{F} = g \circ F \circ f^{-1} : \tilde{U} \longrightarrow \tilde{V}$$

Now choose a curve σ through x. Using our chart, this becomes a curve $\tilde{\sigma} = f \circ \sigma$ through the point $\tilde{x} = f(x) \in \tilde{U}$, and its associated tangent vector is:

$$\Delta_f(\sigma) = D\tilde{\sigma}|_0 \ \in \mathbb{R}^n$$

Now form the composition $F \circ \sigma$, this is a curve through the point $y \in Y$. Using our chart on Y, it becomes a curve

$$g \circ (F \circ \sigma) = \tilde{F} \circ \tilde{\sigma}$$

through the point $\tilde{y} = g(y) \in \tilde{V}$. The tangent vector associated to this curve is

$$\Delta_g(F \circ \sigma) = D(\tilde{F} \circ \tilde{\sigma})|_0 \ \in \mathbb{R}^m$$

which by the chain rule (5.9) is equal to:

$$D\tilde{F}|_{\tilde{x}}\big(D\tilde{\sigma}|_0\big) = D\tilde{F}|_{\tilde{x}}\big(\Delta_f(\sigma)\big)$$

So the tangent vector $\Delta_g(F \circ \sigma)$ only depends on the tangent vector $\Delta_f(\sigma)$. This means that the equivalence class of the curve $F \circ \sigma$ in the space T_yY only depends on the equivalence class of the curve σ in the space T_xX, and thus we have a well-defined function $DF|_x : T_xX \to T_yY$ that sends $[\sigma] \to [F \circ \sigma]$. Furthermore the square

$$\begin{tikzcd} T_xX \arrow[r, "DF|_x"] \arrow[d, "\Delta_f"'] & T_yY \arrow[d, "\Delta_g"] \\ \mathbb{R}^n \arrow[r, "D\tilde{F}|_{\tilde{x}}"] & \mathbb{R}^m \end{tikzcd} \tag{5.11}$$

commutes, *i.e.*

$$DF|_x = \Delta_g^{-1} \circ D\tilde{F}|_{\tilde{x}} \circ \Delta_f$$

and therefore $DF|_x$ is a linear map, since it's the composition of three linear maps. $\square$

So now we have a way to talk about the derivative of a smooth map abstractly, without reference to any co-ordinate charts. If we do decide to pick co-ordinates, it reduces the ordinary notion of the derivative of a smooth map, via the square (5.11).

In particular, it's immediate that the rank of F at x (Definition 4.12) is exactly the rank of the linear map $DF|_x$. So we could restate our definitions of submersions, immersions, critical points, etc. just in terms of this abstract linear map, without ever picking charts.

5.2 Tangent spaces to submanifolds

If we have a submanifold Z of $\mathbb{R}^n$, and we choose a point $z \in Z$, then we have an intuitive idea of what it means for a vector $v \in \mathbb{R}^n$ to be 'tangent' to Z at the point z. This means that for submanifolds of $\mathbb{R}^n$ there should be a much more elementary definition of the tangent space T_zZ, it should be the subspace of $\mathbb{R}^n$ consisting of vectors that are tangent to Z at z. Let's show that our fancy definition agrees with this more elementary definition, in this special case.

The manifold $\mathbb{R}^n$ is very special in that for any point $z \in \mathbb{R}^n$, the tangent space $T_z\mathbb{R}^n$ can be *canonically* identified with $\mathbb{R}^n$, via the map $[\sigma] \mapsto D\sigma|_0$. Another way to say this is to observe that on any manifold X we can identify any tangent space T_xX with $\mathbb{R}^n$ once we've chosen a chart around x, but if $X = \mathbb{R}^n$ then there is a canonical choice of co-ordinates, namely the identity function. So for a point in $\mathbb{R}^n$ there is no harm in thinking of the tangent space as being literally the vector space $\mathbb{R}^n$.

Now let $Z \subset \mathbb{R}^n$ be a submanifold, and pick a point $z \in Z$. If we have a curve σ in Z through the point z, then we can think of σ as curve in $\mathbb{R}^n$ (through z). More formally, if $\iota : Z \hookrightarrow \mathbb{R}^n$ is the inclusion map, we can replace σ by $\iota \circ \sigma$. We saw above that this induces a well-defined map on tangency classes, this is how we defined the derivative:

$$\begin{aligned} D\iota|_z : T_zZ &\to T_z\mathbb{R}^n \\ [\sigma] &\mapsto [\iota \circ \sigma] \end{aligned}$$

If we identify $T_z\mathbb{R}^n$ with $\mathbb{R}^n$, then this map is simply sending the class of σ to the vector $D\sigma|_0 \in \mathbb{R}^n$. We checked in Lemma 4.18 that ι is an immersion, so the linear map $D\iota|_z$ is an injection. This means that we can view T_zZ as a subspace of $\mathbb{R}^n$; it is the subspace of vectors which are tangent to curves in Z. This is exactly our intuitive idea of a tangent space.

Example 5.12. Consider the submanifold $Z_1 = \{(x, \sin x)\} \subset \mathbb{R}^2$ from Example 3.1. For any point $y = (x, \sin x) \in Z_1$, we can define a curve in $\mathbb{R}^2$ through the point y by

$$\begin{aligned} \sigma : (-\epsilon, \epsilon) &\to \mathbb{R}^2 \\ t &\mapsto (t + x, \sin(t + x)) \end{aligned}$$

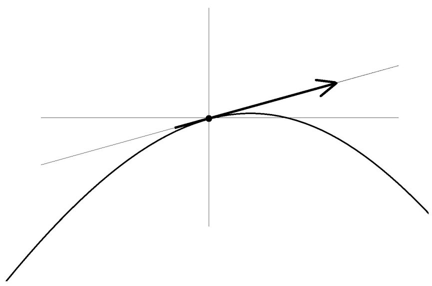

Figure 8: The tangent space to a point in Z_1.

(and ϵ can be any positive real number). The image of σ lies in Z_1, so σ is automatically a smooth function from $(-\epsilon, \epsilon)$ to Z_1 by Lemma 4.8. Hence $[\sigma]$ is a vector in $T_y Z_1$. If we want to view $[\sigma]$ as a vector in $T_y\mathbb{R}^2 \cong \mathbb{R}^2$ we compute:

$$D\sigma|_0 = \begin{pmatrix} 1 \\ \cos x \end{pmatrix} \in \mathbb{R}^2$$

Since Z_1 is only 1-dimensional, the tangent space $T_y Z_1$ is the line in $\mathbb{R}^2$ spanned by this vector (see Figure 8).

This is exactly what the tangent line ought to be, so our complicated definitions have reduced to a sensible answer.

We can generalize this picture from $\mathbb{R}^n$ to arbitrary manifolds. Suppose X is a manifold, and $Z \subset X$ is a submanifold. The inclusion map

$$\iota : Z \hookrightarrow X$$

is a smooth immersion (Lemma 4.18 again), so for any $z \in Z$ we have a linear injection:

$$D\iota|_z : T_z Z \hookrightarrow T_z X$$

So we can always view $T_z Z$ as a subspace of $T_z X$. We can see this very explicitly in co-ordinates: we know we can choose a chart (U, f) containing z such that

$$f(U \cap Z) = \tilde{U} \cap \mathbb{R}^m$$

for the standard subspace $\mathbb{R}^m \subset \mathbb{R}^n$. This fixes an isomorphism $\Delta_f : T_z X \xrightarrow{\sim} \mathbb{R}^n$, under which $T_z Z$ becomes the subspace $\mathbb{R}^m \subset \mathbb{R}^n$.

We saw in Proposition 4.14 that a good way to produce submanifolds is as the level sets of smooth functions.

Lemma 5.13. *Let $F : X \to Y$ be a smooth function, let $y \in Y$ be a regular value of F, and let $Z = \left\{F^{-1}(y)\right\}$ be the corresponding submanifold of X. For any $z \in Z$, the tangent space $T_z Z$ is the kernel of the linear map:*

$$DF|_z : T_z X \to T_y Y$$

Proof. Let the dimensions of X and Y be n and k. The Implicit Function Theorem implies (see problem sheets) that we can find a chart (U, f) on X containing z, and a chart (V, g) on Y containing y, such that the function

$$\tilde{F} = g \circ F \circ f^{-1} : \tilde{U} \to \tilde{V}$$

is simply the restriction of the standard linear projection:

$$\pi : \mathbb{R}^n \to \mathbb{R}^k$$

We can also assume (after translating) that $g(y) = 0$, and then $f(Z \cap U)$ is the intersection of $\tilde{U}$ with the subspace $\mathbb{R}^{n-k} \subset \mathbb{R}^n$. This is essentially how we proved Proposition 4.14.

Now since π is linear, if we look at $DF|_z$ in these charts we get

$$D\tilde{F}|_{f(z)} = \pi : \mathbb{R}^n \to \mathbb{R}^k$$

and the kernel of this is the subspace $\mathbb{R}^{n-k} \subset \mathbb{R}^n$, which is exactly the tangent space $T_z Z \subset T_z X$ written in this chart. $\square$

Example 5.14. Consider the submanifold $S^n \subset \mathbb{R}^{n+1}$. This is the level set of the function

$$h : (x_0, ..., x_n) \mapsto x_0^2 + ... x_n^2$$

at the regular value $h = 1$, as we saw in Example 3.19. At a point $x = (x_0, ..., x_n) \in S^n$, the derivative of h is the 1-by-n matrix:

$$(2x_0, ..., 2x_n) : \mathbb{R}^n \to \mathbb{R}$$

So the tangent space $T_x S^n$ is the subspace:

$$T_x S^n = \{v \; ; \; x.v = 0\} \subset \mathbb{R}^{n+1}$$

Example 5.15. If we specialize the previous Example to $S^1 \subset \mathbb{R}^2$, we see that the tangent space to a point $(x, y) \in S^1$ is the subspace:

$$\left\langle \begin{pmatrix} -y \\ x \end{pmatrix} \right\rangle \subset \mathbb{R}^2$$

Now let's derive this again using polar co-ordinates

$$\begin{aligned} f^{-1} : \; \tilde{U} = \mathbb{R}_{>0} \times (-\pi, \pi) \; &\xrightarrow{\sim} \; U = \mathbb{R}^2 \setminus \{(x, 0), \; x \leq 0\} \\ (r, \theta) &\mapsto (r\cos\theta, r\sin\theta) \end{aligned}$$

as in Example 3.3. In this chart, S^1 becomes the subspace $\{r = 1\} \subset \tilde{U}$, so if we pick a point $(1, \theta) \in f(S^1)$ then the tangent space is:

$$T_{(1,\theta)} f(S^1) = \left\langle \begin{pmatrix} 0 \\ 1 \end{pmatrix} \right\rangle \subset \mathbb{R}^2$$

The derivative of f^{-1} at this point is

$$D(f^{-1})|_{(1,\theta)} = \begin{pmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{pmatrix}$$

so the tangent space to the point $(\cos\theta, \sin\theta) \in S^1$ is the line spanned by the vector $(-\sin\theta, \cos\theta)^\top \in \mathbb{R}^2$.

5.3 A second definition of tangent vectors

We're now going to discuss a second way to define tangent vectors, and later on we'll introduce a third definition. These other definitions are precisely equivalent to the definition we've already introduced, but each one has its own advantages and disadvantages.

Fix a point x in a manifold X, and fix a tangent vector $[\sigma] \in T_x X$, the equivalence class of some curve σ through x. If we now choose a co-ordinate chart (U, f) containing x, we can turn $[\sigma]$ into an ordinary column vector $\Delta_f(\sigma) \in \mathbb{R}^n$. Furthermore, if we change co-ordinates between (U_1, f_1) and (U_2, f_2), we know that the 'transformation law'

$$\Delta_{f_2}(\sigma) = D\phi_{21}|_{f_1(x)}\big(\Delta_{f_1}(\sigma)\big) \tag{5.16}$$

holds (this was equation (5.7)). If we wish, we can take these properties to be the *definition* of a tangent vector.

Definition 5.17 (*'Physicist's definition'*)**.** Fix a point x in an n-dimensional manifold X. Let $\mathcal{A}_x$ denote the set of all co-ordinate charts on X that contain the point x. A **tangent vector** to x is a function

$$\delta : \mathcal{A}_x \to \mathbb{R}^n$$

which we write

$$\delta : (U, f) \mapsto \delta_f$$

and which has the following property: for any two charts (U_1, f_1) and (U_2, f_2) in $\mathcal{A}_x$, the equation

$$\delta_{f_2} = D\phi_{21}|_{f_1(x)}\big(\delta_{f_1}\big) \tag{5.18}$$

holds.

Let's temporarily use the notation $\mathcal{T}_x X$ to denote the set of tangent vectors in the sense of Definition 5.17, although as we shall see in a moment the two definitions are equivalent and $\mathcal{T}_x X$ is the same thing as $T_x X$.

The vector space structure on $\mathcal{T}_x X$ is much more obvious than the one on $T_x X$. For two elements $\delta, \hat{\delta} \in \mathcal{T}_x X$, we can define $\delta + \hat{\delta}$ to be the function:

$$\delta + \hat{\delta} : (U, f) \mapsto \delta_f + \hat{\delta}_f \in \mathbb{R}^n$$

This still obeys the rule (5.18) for any two charts, because $D\phi_{21}|_{f_1(x)}$ is linear. Scalar multiplication is similar, and it's immediate that $\mathcal{T}_x X$ is a vector space. Now let's prove that it has dimension n.

Lemma 5.19. *For any chart $(U, f) \in \mathcal{A}_x$, the function 'evaluate in (U, f)'*

$$\begin{aligned} ev_f : \mathcal{T}_x X &\longrightarrow \mathbb{R}^n \\ \delta &\mapsto \delta_f \end{aligned}$$

is a linear isomorphism.

Proof. The fact that ev_f is linear follows instantly from the definition of the vector space structure on $\mathcal{T}_x X$. It's also clear that it's an injection, because if

δ and $\hat{\delta}$ give the same vector for (U, f) then they must give the same vector for all charts, because of the rule (5.18). So it only remains to prove surjectivity.

Pick any vector $v \in \mathbb{R}^n$. Let $(U_0, f_0) = (U, f)$, and define a function $\delta : \mathcal{A}_x \to \mathbb{R}^n$ by:

$$\delta : (U_i, f_i) \mapsto \delta_{f_i} = D\phi_{i0}|_{f_0(x)}(v)$$

If we pick any two charts $(U_1, f_1), (U_2, f_2) \in \mathcal{A}_x$, the transition functions obey the equation

$$\phi_{20} = \phi_{21} \circ \phi_{10}$$

(this makes sense in some neighbourhood of $f_0(x)$), so by the chain rule:

$$\delta_{f_2} = D\phi_{20}|_{f_0}(x)(v) = D\phi_{21}|_{f_1(x)}\big(D\phi_{10}|_{f_0(x)}(v)\big) = D\phi_{21}|_{f_1(x)}(\delta_{f_1})$$

So our function δ obeys the rule (5.18), hence it's an element of $\mathcal{T}_x X$, and by construction $ev_f(\delta) = v$. □

If we have two different charts in $\mathcal{A}_x$ then the isomorphisms ev_{f_1} and ev_{f_2} are related by

$$ev_{f_2} = D\phi_{21}|_{f_1(x)} \circ ev_{f_1}$$

which is the same rule as relates Δ_{f_1} and Δ_{f_2}.

Now we can prove that our two definitions of tangent vectors, Definition 5.5 and Definition 5.17, are equivalent.

Proposition 5.20. *There is a canonical linear isomorphism between $\mathcal{T}_x X$ and $T_x X$.*

Proof. If we have a tangent vector $[\sigma] \in T_x X$ in the 'geometer's sense', then we can get a tangent vector $\delta \in \mathcal{T}_x X$ in the 'physicist's sense' by considering the function:

$$\begin{aligned} \delta : \mathcal{A}_x &\to \mathbb{R}^n \\ (U, f) &\mapsto \Delta_f(\sigma) \end{aligned}$$

So there is a natural function:

$$T_x X \to \mathcal{T}_x X$$

If we fix any chart $(U, f) \in \mathcal{A}_x$, this function factors as the composition

$$T_x X \xrightarrow{\Delta_f} \mathbb{R}^n \xrightarrow{ev_f^{-1}} \mathcal{T}_x X$$

and since both factors are linear isomorphisms, so is their composition. □

6 Vector fields

6.1 Definition of a vector field

We have shown (twice over) that to any point x in a manifold X we can attach a vector space T_xX. A *vector field* on X is a function which assigns to any point $x \in X$ a vector in the corresponding tangent space T_xX.

We'll give the formal definition of a vector field shortly, but let's start by considering the case when $X = U$ is just an open set in $\mathbb{R}^n$. In this case, the tangent space T_xU to any point $x \in U$ is canonically isomorphic to $\mathbb{R}^n$, because we have a canonical set of co-ordinates on U. This means that a vector field on U is simply a function:

$$\tilde{\xi} : U \to \mathbb{R}^n$$

This is a definition you may have seen before, but it is potentially misleading. If x and y are two different points in U then we should really think of the vectors $\tilde{\xi}|_x$ and $\tilde{\xi}|_y$ as living in two different vector spaces; the first one lives in T_xU and the second one lives in T_yU. For example, it wouldn't normally be sensible to add these two vectors together. So we should visualize vector fields differently from other such functions, you should imagine that at every point $x \in U$ the function $\tilde{\xi}$ 'attaches' a vector $\tilde{\xi}|_x$ to the point x (see Figure 9).

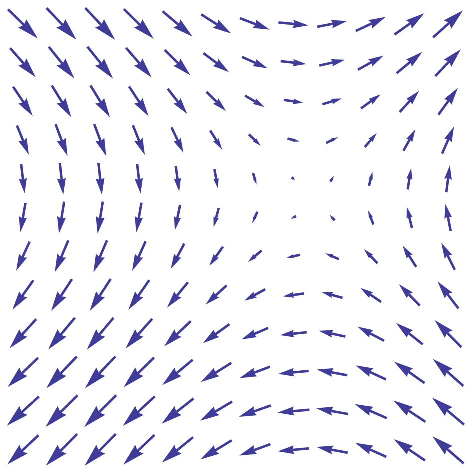

Figure 9: A vector field on $\mathbb{R}^2$.

To make this intuition more formal, we consider the set:

$$TU = \bigcup_{x \in U} T_xU \cong U \times \mathbb{R}^n$$

An element of TU is a pair (x, v) where x is a point in U, and v is a vector in T_xU. This set is exactly $U \times \mathbb{R}^n$, since each T_xU can be identified with $\mathbb{R}^n$. Projecting onto the first factor gives a function:

$$\pi : TU \to U$$

Now we can define a vector field on U to be a function

$$\xi : U \longrightarrow TU = U \times \mathbb{R}^n$$

such that $\pi \circ \xi = 1_U$. This just says that ξ is a function of the form

$$\xi : x \mapsto (x, \tilde{\xi}|_x)$$

for some function $\tilde{\xi} : U \to \mathbb{R}^n$. So the data of ξ and $\tilde{\xi}$ are exactly the same, but this second definition fits the intuitive picture better.

Now we generalize this second definition to an arbitrary manifold. For any manifold X, we can form a set

$$TX = \bigcup_{x \in X} T_x X$$

by taking the disjoint union of all the tangent spaces to all the points in X. So an element of TX is a pair (x, v) with x a point in X, and v a vector in $T_x X$. This set is called the *tangent bundle* to X.

In fact the tangent bundle TX is not just a set, it naturally has the structure of a manifold, whose dimension is twice the dimension of X. We're not going to use this, but it is explained in Appendix E.

If x and y are two different points in X then $T_x X$ and $T_y X$ are different vector spaces, and there is no canonical way to identify them. So in general, TX is *not* a cross-product of X with some other set (we could do this when X was an open set in $\mathbb{R}^n$, but this is a very special case). But we still have a projection function

$$\begin{aligned} \pi : TX &\to X \\ (x, v) &\mapsto x \end{aligned}$$

with $\pi^{-1}(x) = T_x X$ for all $x \in X$. Hence we can make the following definition:

Definition 6.1. A **vector field** on X is a function

$$\xi : X \to TX$$

such that $\pi \circ \xi = 1_X$.

So for each point $x \in X$, the function ξ selects a vector $\xi|_x \in T_x X$.

Example 6.2. If $X = S^1$, then we saw in Example 5.15 that the tangent space $T_{(x,y)} S^1$ to a point $(x, y) \in S^1$ can be identified with the subspace of $\mathbb{R}^2$ spanned by the vector $(-y, x)^\top$. So we can define a vector field on S^1 by:

$$\begin{aligned} \xi : S^1 &\to TS^1 \\ (x, y) &\mapsto \big((x, y), (-y, x)^\top\big) \end{aligned}$$

For a point in S^1 this vector field assigns the corresponding (anti-clockwise) unit angular vector (see Figure 10).

Of course, we are really interested in *smooth* vector fields, but we need to say what this means.

If our manifold is an open set $\tilde{U} \subset \mathbb{R}^n$ then the definition is obvious. We saw before that the tangent bundle $T\tilde{U}$ can be identified with $\tilde{U} \times \mathbb{R}^n$, and this is an open subset of $\mathbb{R}^{2n}$. So we can define a smooth vector field on $\tilde{U}$ to be a

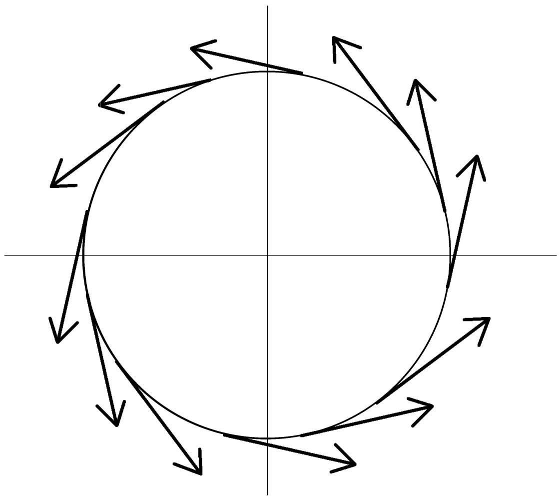

Figure 10: The vector field ξ on S^1.

smooth function $\xi : \tilde{U} \to T\tilde{U}$ satisfying $\pi \circ \xi = 1_{\tilde{U}}$. This is exactly the same data as a smooth function

$$\tilde{\xi} : \tilde{U} \to \mathbb{R}^n$$

since the first component of ξ has to be the identity, so ξ is smooth iff its second component $\tilde{\xi}$ is smooth.

Now we extend this definition to more general manifolds. Let X be a manifold, and let (U, f) be a chart on X. Now let TU be the subset:

$$TU = \pi^{-1}(U) = \bigcup_{x \in U} T_x X \quad \subset TX$$

For each point $x \in U$, our co-ordinates f give us an identification:

$$\Delta_f : T_x X \xrightarrow{\sim} \mathbb{R}^n$$

Putting these together gives us a bijection:

$$F : TU \xrightarrow{\sim} T\tilde{U} = \tilde{U} \times \mathbb{R}^n$$
$$(x, v) \mapsto \big(f(x), \Delta_f(v)\big)$$

So locally TX looks like an open subset of $\mathbb{R}^{2n}$, which is why we can give it the structure of a $2n$-dimensional manifold (see Appendix E).

Now let $\xi : X \to TX$ be a vector field. Since $\xi|_x \in T_x X$ for any point $x \in X$, restricting ξ to U defines a function:

$$\xi|_U : U \to TU$$

Then we can look at this in our co-ordinates, *i.e.* we can consider the function:

$$F \circ \xi|_U \circ f^{-1} : \tilde{U} \longrightarrow T\tilde{U} \cong \tilde{U} \times \mathbb{R}^n$$

This function is a vector field on $\tilde{U}$ (because the bijection F respects the projection maps), so we must have

$$F \circ \xi|_U \circ f^{-1} = \big(1_{\tilde{U}}, \tilde{\xi}\big)$$

for some function $\tilde{\xi} : \tilde{U} \to \mathbb{R}^n$. Now the definition of a smooth vector field on X should be clear.

Definition 6.3. A vector field $\xi : X \to TX$ on a manifold X is called **smooth** iff for any chart (U, f) on X, the associated function

$$F \circ \xi|_U \circ f^{-1} = \big(1_{\tilde{U}}, \tilde{\xi}\big) : \tilde{U} \longrightarrow \tilde{U} \times \mathbb{R}^n$$

is smooth, *i.e.* iff the function $\tilde{\xi} : \tilde{U} \to \mathbb{R}^n$ is smooth.

This definition is only useful if it is chart-independent. Let's check that this is the case.

Suppose we have two charts (U_1, f_1) and (U_2, f_2) on X, and suppose for simplicity that they have the same domain $U = U_1 = U_2$ (if not then shrink both of them to $U_1 \cap U_2$). For any point $x \in U$, we have two different linear isomorphisms

$$\Delta_{f_1} : T_x X \xrightarrow{\sim} \mathbb{R}^n \qquad \text{and} \qquad \Delta_{f_2} : T_x X \xrightarrow{\sim} \mathbb{R}^n$$

and they are related by the 'transformation law'

$$\Delta_{f_2} = D\phi_{21}|_{f_1(x)} \circ \Delta_{f_1}$$

where ϕ_{21} is the transition function. This means our two charts give bijections

$$F_1 : TU \xrightarrow{\sim} \tilde{U}_1 \times \mathbb{R}^n \qquad \text{and} \qquad F_2 : TU \xrightarrow{\sim} \tilde{U}_2 \times \mathbb{R}^n$$

and these are related by the bijection:

$$\begin{aligned} \Phi_{21} : \tilde{U}_1 \times \mathbb{R}^n &\xrightarrow{\sim} \tilde{U}_2 \times \mathbb{R}^n \\ (\tilde{x}, v) &\mapsto \big(\phi_{21}(\tilde{x}), D\phi_{21}|_{\tilde{x}}(v)\big) \end{aligned}$$

The important thing to observe is that this function Φ_{21} is smooth, because the partial derivatives of ϕ_{21} (*i.e.* the entries in the Jacobian matrix) are all smooth functions of $\tilde{x}$. The inverse function is just Φ_{12}, so this is also smooth, and hence Φ_{21} is a diffeomorphism.

Now suppose we have a vector field ξ on X. Looking at ξ in our two charts, we get vector fields

$$F_1 \circ \xi|_U \circ f_1^{-1} = \big(1_{\tilde{U}_1}, \tilde{\xi}_1\big) : \quad \tilde{U}_1 \longrightarrow \tilde{U}_1 \times \mathbb{R}^n$$

and:

$$F_2 \circ \xi|_U \circ f_2^{-1} = \big(1_{\tilde{U}_2}, \tilde{\xi}_2\big) : \quad \tilde{U}_2 \longrightarrow \tilde{U}_2 \times \mathbb{R}^n$$

Since $F_2 = \Phi_{21} \circ F_1$, and $f_2^{-1} = f_1^{-1} \circ \phi_{12}$, it's immediate that the vector field in the first chart is smooth iff the vector field in the second chart is smooth.

In particular if we only look at the second components then we have that $\tilde{\xi}_1$ is smooth iff $\tilde{\xi}_2$ is smooth. Indeed, these two functions are related by the 'transformation law'

$$\tilde{\xi}_2|_{\phi_{21}(\tilde{x})} = D\phi_{21}|_{\tilde{x}}\big(\tilde{\xi}_1|_{\tilde{x}}\big) \tag{6.4}$$

and $D\phi_{21}$ is a matrix whose entries are smooth functions of $\tilde{x}$.

So if we want to check if a vector field ξ is smooth then we don't have to check every chart (which is impossible!), we just pick an atlas for X, and compute the functions $\tilde{\xi}$ for every chart in the atlas.

Example 6.5. Let's check that the vector field ξ on S^1 from Example 6.2 is smooth. If we use polar co-ordinates on $\mathbb{R}^2$ then we get an induced chart on the submanifold $S^1 \subset \mathbb{R}^2$ with $U = S^1 \setminus (-1, 0)$ and $\tilde{U} = (-\pi, \pi) \subset \mathbb{R}$, and:

$$f^{-1} : \theta \mapsto (\cos\theta, \sin\theta)$$

We computed in Example 5.15 that this chart identifies the tangent space $T_{(\cos\theta,\, \sin\theta)}S^1$ with $T_\theta \tilde{U} \cong \mathbb{R}$ via the linear isomorphism:

$$\begin{aligned} \Delta_f^{-1} : \mathbb{R} &\xrightarrow{\sim} T_{(\cos\theta,\, \sin\theta)}S^1 \\ 1 &\mapsto (-\sin\theta, \cos\theta)^\top \end{aligned}$$

So in this chart, the vector field ξ is just a constant function:

$$\tilde{\xi} \equiv 1 : \tilde{U} \to \mathbb{R}$$

This is certainly smooth, which proves that ξ is smooth at every point in S^1 apart from $(-1, 0)$, and we can use another polar co-ordinate chart to check that ξ is smooth at that point too.

As usual we only care about smooth things, so from now on we're going to assume that all our vector fields are smooth, unless we need to specifically state otherwise.

6.2 Vector fields from their transformation law

Now we'll look at another definition of vector fields, by adopting the viewpoint of our 'physicist's definition' of a tangent vector.

If ξ is a vector field on X, then for any chart (U, f) we have a smooth function $\tilde{\xi} : \tilde{U} \to \mathbb{R}^n$. If we have two charts (U_1, f_1) and (U_2, f_2), then the two functions

$$\tilde{\xi}_1 : \tilde{U}_1 \to \mathbb{R}^n \qquad \text{and} \qquad \tilde{\xi}_2 : \tilde{U}_2 \to \mathbb{R}^n$$

are related by the transformation law

$$\tilde{\xi}_2|_{\phi_{21}(\tilde{x})} = D\phi_{21}|_{\tilde{x}}\big(\tilde{\xi}_1|_{\tilde{x}}\big)$$

(see (6.4)). In the style of Definition 5.17, we can take these properties to be the definition of a vector field.

Proposition 6.6. *Let ξ be a rule which assigns to any chart (U, f) on X a smooth function:*

$$\tilde{\xi} : \tilde{U} \to \mathbb{R}^n$$

Assume that ξ obeys the following property: for any two charts (U_1, f_1) and (U_2, f_2), and any point $\tilde{x} \in f_1(U_1 \cap U_2)$, the corresponding functions $\tilde{\xi}_1$ and $\tilde{\xi}_2$ satisfy

$$\tilde{\xi}_2|_{\phi_{21}(\tilde{x})} = D\phi_{21}|_{\tilde{x}}\big(\tilde{\xi}_1|_{\tilde{x}}\big)$$

where ϕ_{21} is the transition function between the two charts. Then ξ defines a (smooth) vector field on X.

Proof. Fix a point $x \in X$ and let $\mathcal{A}_x$ denote the set of charts containing x. Define a function

$$\xi|_x : \mathcal{A}_x \to \mathbb{R}^n$$

by:

$$\xi|_x : (U, f) \mapsto \tilde{\xi}|_{f_1(x)} \in \mathbb{R}^n$$

Then $\xi|_x$ is a tangent vector in the sense of Definition 5.17, so it defines an element of T_xX by Proposition 5.20. Hence we have a vector field:

$$\begin{aligned} \xi : X &\to TX \\ x &\mapsto \xi|_x \end{aligned}$$

In any single chart (U, f) this vector field becomes the corresponding function $\tilde{\xi} : \tilde{U} \to \mathbb{R}^n$, so ξ is smooth. $\square$

We can specify a vector field by choosing an atlas for X and then specifying the functions $\tilde{\xi}$ for every chart in the atlas. The values of the vector field in any other chart will then be determined by the transformation law.

Example 6.7. Let $X = S^1$, and recall our stereographic projection atlas from Example 2.4. So we have two charts whose codomains are

$$\tilde{U}_1 = \mathbb{R} \qquad \text{and} \qquad \tilde{U}_2 = \mathbb{R}$$

and the transition function is:

$$\begin{aligned} \phi_{21} : \mathbb{R} \setminus 0 &\xrightarrow{\sim} \mathbb{R} \setminus 0 \\ \tilde{x} &\mapsto \frac{1}{\tilde{x}} \end{aligned}$$

You can think of this data as alternative definition of the manifold S^1; it says we take two copies of $\mathbb{R}$ and 'glue them together' using ϕ_{21}.

By Proposition 6.6, a vector field on S^1 is equivalent to the data of two smooth functions

$$\tilde{\xi}_1 : \mathbb{R} \to \mathbb{R} \qquad \text{and} \qquad \tilde{\xi}_2 : \mathbb{R} \to \mathbb{R}$$

such that

$$\tilde{\xi}_2(\tfrac{1}{\tilde{x}}) = -\tfrac{1}{\tilde{x}^2}\tilde{\xi}_1(\tilde{x})$$

for all $\tilde{x} \neq 0$. It's very easy to construct vector fields on S^1 using this definition, just take any smooth function $\tilde{\xi}_1 : \mathbb{R} \to \mathbb{R}$ and define:

$$\begin{aligned} \tilde{\xi}_2 : \mathbb{R} \setminus 0 &\to \mathbb{R} \\ \tilde{x} &\mapsto -\tilde{x}^2 \tilde{\xi}_1(\tfrac{1}{\tilde{x}}) \end{aligned}$$

Provided that $\tilde{\xi}_1$ behaves sufficiently well as $|\tilde{x}| \to \infty$, we will be able to extend $\tilde{\xi}_2$ to a smooth function on $\mathbb{R}$. For example, the pairs

$$\begin{array}{lll} \tilde{\xi}_1 : \tilde{x} \mapsto 1 & \qquad \tilde{\xi}_1 : \tilde{x} \mapsto \tilde{x} & \qquad \tilde{\xi}_1 : \tilde{x} \mapsto \tilde{x}^2 \\ \tilde{\xi}_2 : \tilde{x} \mapsto -\tilde{x}^2 & \qquad \tilde{\xi}_2 : \tilde{x} \mapsto -\tilde{x} & \qquad \tilde{\xi}_2 : \tilde{x} \mapsto -1 \end{array}$$

all define vector fields on S^1.

6.3 Flows

Vector fields are closely related to a nice geometric idea called a *flow*, as we'll now explain.

If X and Y are two manifolds then we've defined a diffeomorphism between X and Y to be a smooth function $F : X \to Y$ with a smooth inverse (Definition 4.23). In particular, it makes sense to talk about diffeomorphisms

$$F : X \to X$$

from a manifold to itself. These are the symmetries of a manifold.

Example 6.8. Let $X = T^1$. For any constant $s \in \mathbb{R}$, we can define a bijection

$$F_s : T^1 \xrightarrow{\sim} T^1$$

by:

$$F_s : [t] \mapsto [t + s]$$

It's easy to check that F_s is smooth for any s, and since the inverse of F_s is F_{-s} this shows that each F_s is a diffeomorphism.

We checked in Example 4.25 that the function

$$\begin{aligned} H : T^1 &\to S^1 \\ [t] &\mapsto (\cos 2\pi t, \sin 2\pi t) \end{aligned}$$

is a diffeomorphism. This implies that for any s the function 'rotate by $2\pi s$'

$$\begin{aligned} G_s = H \circ F_s \circ H^{-1} :\ S^1 &\to S^1 \\ (\cos\theta, \sin\theta) &\mapsto \big(\cos(\theta + 2\pi s), \sin(\theta + 2\pi s)\big) \end{aligned}$$

is a diffeomorphism of S^1.

In the previous example, we didn't just write down one diffeomorphism, we wrote down a whole family of them, indexed by the parameter $s \in \mathbb{R}$, and in the middle we have the identity function $F_0 = 1_{T^1}$. Moreover, the diffeomorphism F_s 'depends smoothly on s', in the following sense. Put all of them together to form the function:

$$\begin{aligned} F : \mathbb{R} \times T^1 &\to T^1 \\ (s, [t]) &\mapsto F_s([t]) \end{aligned}$$

The set $\mathbb{R} \times T^1$ is fairly obviously a 2-dimensional manifold (an atlas is given in Example E.2), and it's easy to check that this total function F is smooth.

Let's abstract this example.

Definition 6.9. Let X be a manifold. A **1-parameter family of diffeomorphisms**, or **flow**, on X is a smooth map

$$F : (-\epsilon, \epsilon) \times X \to X$$

for some positive real number ϵ, such that for each $s \in (-\epsilon, \epsilon)$ the function

$$\begin{aligned} F_s : X &\to X \\ x &\mapsto F(s, x) \end{aligned}$$

is a diffeomorphism, and in particular the map F_0 is the identity on X.

It should be clear that that if take a smooth atlas for X then we can produce a smooth atlas for $(-\epsilon,\epsilon)\times X$, by simply crossing every chart with the interval $(-\epsilon,\epsilon)$. So $(-\epsilon,\epsilon)\times X$ is naturally a manifold (with dimension $1+\dim X$), and hence it does makes sense to ask for F to be smooth.

Instead of fixing the parameter s, we can instead choose to fix a point $x\in X$. Then restricting F to the subset $(-\epsilon,\epsilon)\times\{x\}$ gives us a smooth function:

$$\begin{aligned} F_x:(-\epsilon,\epsilon)&\to X\\ s&\mapsto F(s,x)\end{aligned}$$

Since $F_x(0)=F_0(x)=x$, this is a curve through x, so it determines a vector in the tangent space T_xX. If we do this at all points in x simultaneously then we produce a vector field on X, which we'll call ξ^F. This vector field is the 'infinitesimal version of the flow F', it tells us the direction that every point will move in if we start to apply the flow.

Let's look at this procedure in co-ordinates. If we pick a chart on X with codomain $\tilde{U}\subset\mathbb{R}^n$ then F will become a smooth function

$$\tilde{F}=(\tilde{F}_1,...,\tilde{F}_n):(-\epsilon,\epsilon)\times\tilde{U}\to\tilde{U}$$

such that $F_0=1_{\tilde{U}}$ (in fact $\tilde{F}$ isn't really defined on this domain because the flow might not preserve the subset $U\subset X$, but it does make sense on some open neighbourhood of the subset $\{0\}\times\tilde{U}$ and that is all we need). The associated vector $\widetilde{\xi^F}$ field on $\tilde{U}$ is:

$$\widetilde{\xi^F}=\left.\frac{\partial\tilde{F}}{\partial s}\right|_{s=0}=\left(\left.\frac{\partial\tilde{F}_1}{\partial s}\right|_{s=0},\dots,\left.\frac{\partial\tilde{F}_n}{\partial s}\right|_{s=0}\right):\ \tilde{U}\to\mathbb{R}^n$$

In particular, it's clear that the vector field ξ^F is smooth, since $\widetilde{\xi^F}$ is a smooth function.

Example 6.10. In Example 6.8 we defined the following flow on S^1:

$$\begin{aligned} G:\ &(-\epsilon,\epsilon)\times S^1\to S^1\\ &\big(s,(\cos\theta,\sin\theta)\big)\mapsto\big(\cos(\theta+2\pi s),\sin(\theta+2\pi s)\big)\end{aligned}$$

(here ϵ can be any positive real number). We claim that the vector field associated to this flow takes the value

$$\xi^G|_{(\cos\theta,\,\sin\theta)}=\begin{pmatrix}-2\pi\sin\theta\\ 2\pi\cos\theta\end{pmatrix}\in T_{(\cos\theta,\sin\theta)}S^1\subset\mathbb{R}^2$$

at the point $(\cos\theta,\sin\theta)\in S^1$. Up to the overall scale factor of 2π, this is the vector field that we saw in Example 6.2.

There are many ways to prove this claim, for example we can look at G in polar co-ordinates, where it becomes

$$\begin{aligned}\tilde{G}:(-\epsilon,\epsilon)\times(-\pi,\pi)&\longrightarrow(-\pi,\pi)\\ (s,\theta)&\mapsto\theta+2\pi s\end{aligned}$$

(which is really only defined in an open neighbourhood of $\{s = 0\}$). Then the associated vector field is constant:

$$\begin{aligned} \widetilde{\xi^G} : (-\pi, \pi) &\to \mathbb{R} \\ \theta &\mapsto 2\pi \end{aligned}$$

We saw in Example 6.5 that this the expression in polar co-ordinates of the angular vector field on S^1 (up the factor of 2π).

It's in interesting question to ask whether this process can be reversed: if we have a vector field ξ on X, can we construct a flow on X whose associated vector field is ξ?

This is a question about constructing solutions to partial differential equations, and dealing with it properly requires more analysis than we wish to introduce here. However, the answer is yes, provided that we assume that X is *compact*. If we fix a small neighbourhood $U \subset X$ then we can always construct a flow

$$F : (-\epsilon, \epsilon) \times U \to X$$

whose infinitesimal version is $\xi|_U$, for some value of ϵ. If X is non-compact then these values for ϵ might not be bounded above zero over the whole of X, meaning that we cannot find a global flow F for any positive value of ϵ. However if X is compact then there must some minimal $\epsilon_0 > 0$, and we have a global flow with $\epsilon = \epsilon_0$.

7 Cotangent spaces

We've seen that to any point x in a n-dimensional manifold X we can attach an n-dimensional vector space, the tangent space $T_x X$. In fact we've seen two equivalent ways to do this, the 'geometer's definition' and the 'physicist's definition'. In this section we're going to see another way to attach an n-dimensional vector space to x, called the *cotangent space*.

As we shall see, the cotangent space is not the same as the tangent space, it is its dual. This will give us yet another way to define $T_x X$.

7.1 Covectors

Let X be a manifold of dimension n. We let $C^\infty(X)$ denote the set of all smooth functions from X to $\mathbb{R}$, this forms an infinite-dimensional vector space under 'point-wise' addition and scalar multiplication of functions.

Now fix a point x in X. The rank of a function $h \in C^\infty(X)$ at the point x can either be one or zero, and we'll write

$$R_x(X) \ \subset C^\infty(X)$$

for the set of functions which have rank zero at x. We shall see in a moment that $R_x(X)$ is a subspace of $C^\infty(X)$.

Definition 7.1. The *cotangent space* to X at x is the quotient space:

$$T_x^\star X = C^\infty(X)\,/\,R_x(X)$$

We refer to the elements of $T_x^\star X$ as *covectors*.

This may look intimidating, we've taken an enormous vector space and quotiented by an enormous subspace. Let's start as usual by considering the case when X is an open subset in $\mathbb{R}^n$; we shall see that this quotient space is actually very mundane.

Lemma 7.2. *Let $X \subset \mathbb{R}^n$ be an open set, and let x be a point in X. Then $R_x(X) \subset C^\infty(X)$ is a subspace, and the quotient space $T_x^\star X$ has dimension n.*

Proof. Consider the map:

$$\begin{aligned} C^\infty(X) &\longrightarrow \mathbb{R}^n \\ h \mapsto Dh|_x &= \left.\left(\frac{\partial h}{\partial x_1},\ \dots\ ,\ \frac{\partial h}{\partial x_n}\right)\right|_x \end{aligned}$$

This map is linear, and by definition $R_x(X)$ is its kernel. Hence $R_x(X)$ is a subspace of $C^\infty(X)$, and by the first isomorphism theorem $T_x^\star X$ is isomorphic to the image of this map. So we need only show that the map is surjective. But for any vector $(a_1, ..., a_n) \in \mathbb{R}^n$ we can let h be the linear function

$$h = a_1 x_1 + ... + a_n x_n \ \in C^\infty(X)$$

and then the partial derivatives of h at x are the a_i's. □

Now let's move to a general manifold X. Suppose we have a function $h \in C^\infty(X)$ and we want to compute its rank at a point $x \in X$. What we must do is pick a chart (U, f) around x, and consider the function:

$$\tilde{h} = h \circ f^{-1} : \tilde{U} \to \mathbb{R}$$

Then h has rank zero at x iff all partial derivatives of $\tilde{h}$ vanish at the point $f(x) \in \tilde{U}$, and we know that if this is true in one chart then it must be true in all charts. So, once we've fixed our chart (U, f), we can define a linear map

$$\begin{aligned} \nabla_f : C^\infty(X) &\longrightarrow \mathbb{R}^n \\ h &\mapsto D\tilde{h}|_{f(x)} = \left(\frac{\partial \tilde{h}}{\partial x_1}, \ \dots \ , \ \frac{\partial \tilde{h}}{\partial x_n} \right)\Bigg|_{f(x)} \end{aligned}$$

and the kernel of this map is $R_x(X)$. This proves that $R_x(X)$ is a subspace of $C^\infty(X)$, and that the quotient space $T^\star_x(X)$ is isomorphic to some subspace of $\mathbb{R}^n$ - hence its dimension is at most n. The only tricky part is proving that this map ∇_f is surjective, this reason that this is not obvious is that we don't know a general procedure for producing any smooth functions in $C^\infty(X)$. To rectify this, we're going to introduce some very useful gadgets called *bump functions*.

You may recall the remarkable function:

$$\begin{aligned} \phi : \mathbb{R} &\to \mathbb{R} \\ x &\mapsto \begin{cases} e^{-\frac{1}{x}}, \ \text{for } x > 0 \\ 0, \ \text{for } x \leq 0 \end{cases} \end{aligned}$$

This function is differentiable to arbitrary order, because all derivatives of $e^{-\frac{1}{x}}$ tend to zero as x tends to zero. It's a good example for showing that Taylor expansions are not necessarily to be trusted, because the Taylor expansion of ϕ at zero is identically zero, but ϕ is not equal to the zero function in any neighbourhood of zero.

By messing around with ϕ we can create some other nice functions, for example:

$$\psi(x) = \frac{\phi(x)}{\phi(x) + \phi(1-x)}$$

This function ψ is smooth, identically equal to zero for $x \leq 0$, and identically equal to 1 for $x \geq 1$ (see Figure 11).With some further modifications, it should be clear that we can create a smooth function of n variables

$$\widetilde{\psi} : \mathbb{R}^n \to \mathbb{R}$$

which is constantly equal to 1 inside the ball $B(0, r)$ and constantly equal to 0 outside the ball $B(0, r')$, for any $r' > r > 0$. This is called a *bump function*.

We can also create bump functions on arbitrary manifolds. Let X be a manifold, and let x be a point in X. Pick a chart (U, f) containing x, and for simplicity assume that $f(x) = 0 \in \mathbb{R}^n$. Let $\widetilde{\psi}$ be a bump function on $\mathbb{R}^n$ as above, chosen such that the closure of the larger ball $\overline{B(0, r')}$ is contained in $\tilde{U}$. Then we can define a function on the whole of X by

$$\psi(y) = \begin{cases} (\tilde{\psi} \circ f)(y), \ \text{for } y \in U \\ 0, \quad \text{for } y \notin U \end{cases}$$

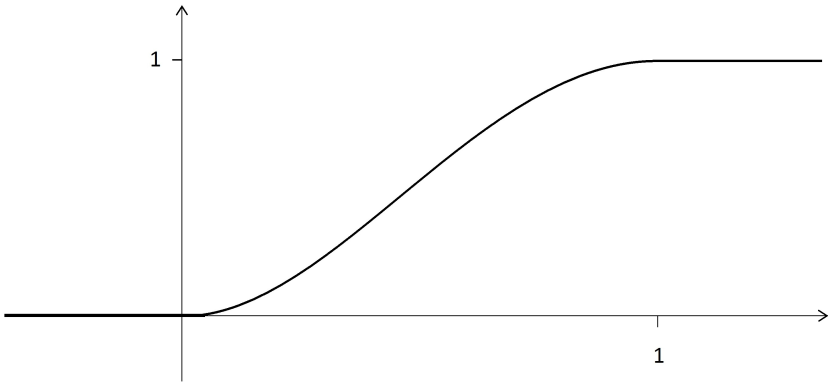

Figure 11: The function ψ.

So ψ is real-valued function which has constant value 1 in an open neighbourhood $W = f^{-1}(B(0,r))$ of x, and constant value 0 outside a larger open neighbourhood $W' = f^{-1}(B(0,r'))$. Moreover since ψ is smooth inside U and constant outside U it is obviously smooth.[1]

Bump functions are useful for many things, one of which is extending smooth functions from open sets to the whole of X. If we have a function $g \in C^\infty(U)$, then we can define a function $\widehat{g} \in C^\infty(X)$ by:

$$\widehat{g} = \begin{cases} g\psi & \text{inside } U \\ 0 & \text{outside } U \end{cases}$$

Then $\widehat{g}|_U$ and g agree on the open set W. Using this construction we can prove:

Proposition 7.3. *Let x be a point in an n-dimensional manifold X. Then the cotangent space $T_x^\star X$ has dimension n.*

Proof. As we saw above, if we fix a chart (U, f) containing x then we get a linear injection $\nabla_f : T_x^\star X \to \mathbb{R}^n$. This sends a function $h \in C^\infty(X)$ to the Jacobian matrix of h, computed in the chart (U, f). Fix a vector $a = (a_1, ..., a_n) \in \mathbb{R}^n$. On the codomain $\tilde{U}$ of our chart we can find a smooth function $\tilde{h} \in C^\infty(\tilde{U})$ such that $D\tilde{h}|_{f(x)} = a$ - for example we can set $\tilde{h}$ to be a linear function as we did in the proof of Lemma 7.2.

Now construct a bump function ψ on X which is identically equal to 1 on some open neighbourhood of x, and is identically zero outside U. Define a function $h \in C^\infty(X)$ by declaring h to be equal to $\psi(\tilde{h} \circ f)$ inside U, and zero outside of U. If we write h in the chart (U, f) we get a function which agrees with $\tilde{h}$ in some neighbourhood of $f(x)$, so $\nabla_f(h) = a$. This proves that ∇_f is an isomorphism. $\square$

[1] Actually it is *not* obvious that this function is smooth. It is only true because our manifolds are Hausdorff, without this condition ψ might not even be continuous (see Appendix C). This is the only point in this course where we will need to invoke the Hausdorff condition.

As an aside, this trick with bump functions shows that there are lots of smooth functions from X to $\mathbb{R}$, so the space $C^\infty(X)$ is very large indeed. You can also use this trick to prove that, for example, there are lots of smooth vector fields on X. Just pick a chart and write down any vector field on the codomain $\tilde{U}$ of that chart - this is just a smooth function from $\tilde{U}$ to $\mathbb{R}^n$. Then use a bump function to extend this to smooth vector field on X.

Given a function $h \in C^\infty(X)$, we'll use the notation

$$dh|_x \ \in T_x^\star X$$

for the associated covector at x, *i.e.* the equivalence class of h in the cotangent space. If X is just an open set in $\mathbb{R}^n$ then we can canonically identify $T_x^\star X$ with $\mathbb{R}^n$, and $dh|_x$ with the vector of partial derivatives of h at x. However on a more general manifold there is no canonical way to do this, it's only once we've picked a chart that we get an isomorphism $\nabla_f : T_x^\star X \xrightarrow{\sim} \mathbb{R}^n$.

This is exactly like the situation for the tangent space $T_x X$, but there is one important difference: the 'tranformation law' for covectors is different. Suppose we have two charts (U_1, f_1) and (U_2, f_2) both containing x, and we choose a smooth function $h \in C^\infty(X)$. If we write out the covector $dh|_x$ in these two charts, we get the vectors

$$\nabla_{f_1}(dh|_x) = D\tilde{h}_1|_{f_1(x)} \qquad \text{and} \qquad \nabla_{f_2}(dh|_x) = D\tilde{h}_2|_{f_2(x)}$$

where $\tilde{h}_1 = h \circ f_1^{-1}$ and $\tilde{h}_2 = h \circ f_2^{-1}$. Now $\tilde{h}_1$ and $\tilde{h}_2$ are related by the transition function ϕ_{21} between the two charts, in some neighbourhood of $f_2(x)$ we have that

$$\tilde{h}_2 = \tilde{h}_1 \circ \phi_{12}$$

and then the chain rule says that:

$$D\tilde{h}_2|_{f_2(x)} = D\tilde{h}_1|_{f_1(x)} D\phi_{12}|_{f_2(x)}$$

In this equation the matrix $D\tilde{h}_1|_{f_1(x)}$ is a row vector, and it is being transformed into another row vector by the n-by-n matrix $D\phi_{12}|_{f_2(x)}$ acting on the right. If you prefer your vectors to be column vectors, and your matrices to act on the left, then transpose the equation:

$$(D\tilde{h}_2|_{f_2(x)})^\top = (D\phi_{12}|_{f_2(x)})^\top (D\tilde{h}_1|_{f_1(x)})^\top$$

This says that the linear isomorphisms ∇_{f_1} and ∇_{f_2} are related by:

$$\nabla_{f_2} = \left(D\phi_{12}|_{f_2(x)}\right)^\top \circ \nabla_{f_1} \tag{7.4}$$

Compare this to the transformation law for tangent vectors (5.16), it is similar but not the same.

For tangent vectors, we showed that the transformation law could be used to give an alternative 'physicist's' definition. We can do the same thing for covectors.

Proposition 7.5. *Let x be a point in a manifold X, and let $\mathcal{A}_x$ denote the set of all charts containing x. A covector in $T_x^\star X$ is the same thing as a function*

$$\begin{aligned} \epsilon : \mathcal{A}_x &\to \mathbb{R}^n \\ (U, f) &\mapsto \epsilon_f \end{aligned}$$

such that, for any two charts $(U_1, f_1), (U_2, f_2) \in \mathcal{A}_x$, we have:

$$\epsilon_{f_2} = \left(D\phi_{12}|_{f_2(x)}\right)^\top (\epsilon_{f_1})$$

Proof. This is proved in exactly the same way as Proposition 5.20. Let $\mathcal{T}_x^\star X$ denote the set of all such functions ϵ, this has an obvious vector space structure. For any chart (U, f), the evaluation map

$$\begin{aligned} ev_f : \mathcal{T}_x^\star X &\to \mathbb{R}^n \\ \epsilon &\mapsto \epsilon_f \end{aligned}$$

is a linear isomorphism, by exactly the same argument that proved Lemma 5.19. Then the evident function from $T_x^\star X$ to $\mathcal{T}_x^\star X$ must be a linear isomorphism, since if we pick any chart then we can factor it as:

$$T_x^\star X \xrightarrow{\nabla_f} \mathbb{R}^n \xrightarrow{ev_f^{-1}} \mathcal{T}_x^\star X$$

$\square$

7.2 A third definition of tangent vectors

We've seen that we can define tangent vectors using curves, or as an operation turning charts into vectors. We're now going to introduce a third definition, which takes the point-of-view that a tangent vector is a direction in which we can take a partial derivative of a function. Using this definition, we shall see that the cotangent space $T_x^\star X$ is the dual to the tangent space $T_x X$.

As usual let's start with the easy case when our manifold X is an open subset in $\mathbb{R}^n$. Fix a point $x \in X$. In this easy case we can canonically identify the tangent space $T_x X$ with $\mathbb{R}^n$, so a tangent vector $v \in T_x X$ is just a vector in $\mathbb{R}^n$. Given v, we can consider the operation 'take the partial derivative at x in the direction v'. This is an operator

$$\partial_{x,v} : C^\infty(X) \longrightarrow \mathbb{R}$$

that sends a smooth function $h \in C^\infty(X)$ to the real number:

$$\partial_{x,v}(h) \;=\; Dh|_x(v) \;=\; v_1 \left.\frac{\partial h}{\partial x_1}\right|_x + ... + v_n \left.\frac{\partial h}{\partial x_n}\right|_x \tag{7.6}$$

This operator $\partial_{x,v}$ is a linear map, *i.e.* we have:

$$\partial_{x,v}(h_1 + h_2) = \partial_{x,v}(h_1) + \partial_{x,v}(h_2) \qquad \text{and} \qquad \partial_{x,v}(\lambda h_1) = \lambda \partial_{x,v}(h_1)$$

for any functions $h_1, h_2 \in C^\infty(X)$ and any scalar $\lambda \in \mathbb{R}$.

We can also express this operator using curves. If we want to think of $T_x X$ as the space of curves through x modulo tangency, then we should replace $v \in \mathbb{R}^n$ with (the class of) any curve σ such that $D\sigma|_0 = v$. Then σ is a function from some interval $(-\epsilon, \epsilon)$ into X, so $h \circ \sigma$ is a function from $(-\epsilon, \epsilon)$ to $\mathbb{R}$, and the chain rule tells us that:

$$\left.\frac{\mathrm{d(h \circ \sigma)}}{\mathrm{d}t}\right|_0 = \partial_{x,v}(h)$$

Now suppose h is a function in $C^\infty(X)$ that has rank zero at x. Then $\partial_{x,v}(h) = 0$ automatically. This means that $\partial_{x,v}$ induces a well-defined linear map

$$\begin{aligned} \partial_{x,v} : T_x^\star X &\longrightarrow \mathbb{R} \\ dh|_x &\mapsto \partial_{x,v}(h) \end{aligned}$$

since $T_x^\star X$ is the quotient space $C^\infty(X)/R_x(X)$, and adding an element of $R_x(X)$ onto h doesn't change the value of $\partial_{x,v}(h)$. So the operator $\partial_{x,v}$ is an element of the *dual* vector space $(T_x^\star X)^\star$ (a quick revision of dual vector spaces is given in Appendix B).

Since we're assuming X is just an open set in $\mathbb{R}^n$ we can canonically identify $T_x^\star X$ with $\mathbb{R}^n$, by sending $dh|_x$ to the vector of partial derivatives of h. If we make this identification then (7.6) says that $\partial_{x,v}$ is simply the linear map

$$\begin{aligned} \mathbb{R}^n &\to \mathbb{R} \\ u &\mapsto v.u \end{aligned}$$

given by taking the dot product with v.

Conversely, suppose we have some linear map $\delta : T_x^\star X \to \mathbb{R}$. We can identify $T_x^\star X$ with $\mathbb{R}^n$, and then this linear map δ must be given by taking the dot product with some vector $v \in \mathbb{R}^n$ (the dual space to $\mathbb{R}^n$ is $\mathbb{R}^n$). But then δ sends the class of a function h to the number $Dh|_x(v)$, so δ is the operator $\partial_{x,v}$.

So tangent vectors are the same things as linear maps from $T_x^\star X$ to $\mathbb{R}$, *i.e.* the tangent space is the dual space to the cotangent space.

Now let's repeat all the above on an arbitrary manifold X. Fix a point $x \in X$. The first step is to turn tangent vectors in $T_x X$ into operators on the space of functions $C^\infty(X)$.

Let's adopt our first definition of tangent vectors, so an element of $T_x X$ is the equivalence class of a curve $\sigma : (-\epsilon, \epsilon) \to X$ through x. We define an associated linear operator:

$$\begin{aligned} \partial_\sigma : C^\infty(X) &\to \mathbb{R} \\ h &\mapsto \left.\frac{\mathrm{d}(\mathrm{h} \circ \sigma)}{\mathrm{d}t}\right|_0 \end{aligned} \tag{7.7}$$

Let's look at this in co-ordinates, so pick a chart (U, f) around x. Then h becomes a function $\tilde{h} = h \circ f^{-1} \in C^\infty(\tilde{U})$, and σ becomes a curve

$$\tilde{\sigma} = f \circ \sigma : (-\epsilon, \epsilon) \to \tilde{U}$$

through the point $f(x)$ (possibly after shrinking ϵ so that the image of σ lies in U). But the composition $\tilde{h} \circ \tilde{\sigma}$ equals $h \circ \sigma$, so the the operator ∂_σ sends the function h to:

$$\left.\frac{\mathrm{d}(\tilde{h} \circ \tilde{\sigma})}{\mathrm{d}t}\right|_0 = D\tilde{h}|_{f(x)} D\tilde{\sigma}|_0 = \sum_{i=1}^n \left.\frac{\partial \tilde{h}}{\partial x_i}\right|_{f(x)} \left.\frac{d\tilde{\sigma}_i}{dt}\right|_0$$

This expression is a row vector applied to a column vector, or if you prefer it's a dot product of the two vectors:

$$\Delta_f([\sigma]) = D\tilde{\sigma}|_0 \qquad \text{and} \qquad \nabla_f(dh|_x) = D\tilde{h}|_{f(x)} \quad \in \mathbb{R}^n$$

Yet another way to express this is to say we have applied the operator $\partial_{f(x),v}$ to the function $\tilde{h}$, where $v = D\tilde{\sigma}|_0$.

So we have two ways to descibe what this operator ∂_σ does: we can either define it without using any charts as in (7.7), or we can say "pick a chart, write both h and σ in co-ordinates, then take the dot product of the vectors $D\tilde{h}|_{f(x)}$ and $D\sigma|_0$". These two prescriptions give the same answer, so in particular the second version is definitely chart-independent.

The version written in a chart makes two things explicitly clear. Firstly, the operator ∂_σ only depends on the tangency class of the curve σ, so for any tangent vector $[\sigma] \in T_xX$ we get a linear map from $C^\infty(X)$ to $\mathbb{R}$. Secondly, if $h \in C^\infty(X)$ has rank zero at x then $\partial_\sigma(h) = 0$, and consequently we have a well-defined linear map on the quotient space:

$$\partial_\sigma : T_x^\star X \to \mathbb{R}$$

This shows that for any any tangent vector in T_xX there is a corresponding element of the dual space of the cotangent space.

Proposition 7.8. *Let $x \in X$ be a point in a manifold. The map*

$$\begin{aligned} T_xX &\longrightarrow (T_x^\star X)^\star \\ [\sigma] &\to \partial_\sigma \end{aligned}$$

is a linear isomorphism.

Proof. Pick a chart (U, f) around x. Then we get isomorphisms $\Delta_f : T_xX \xrightarrow{\sim} \mathbb{R}^n$ and $\nabla_f : T_x^\star X \xrightarrow{\sim} \mathbb{R}^n$. Take a curve σ and a function $h \in C^\infty(X)$, so we have a tangent vector $[\sigma]$ and a covector $dh|_x$. The operator ∂_σ sends the covector $dh|_x$ to the dot product of $\Delta_f([\sigma])$ with $\nabla_f(dh|_x)$, so the map in the statement of the proposition becomes the map

$$\mathbb{R}^n \to (\mathbb{R}^n)^\star$$

sending a vector v to the operation 'dot with v'. This is a linear isomorphism. $\square$

So the tangent space T_xX is the dual to the cotangent space $T_x^\star X$. For any finite-dimensional vector space V we have a canonical isomorphism between V and $(V^\star)^\star$, so we can also interpret this an isomorphism

$$(T_xX)^\star \cong T_x^\star X \tag{7.9}$$

between the cotangent space and the dual of the tangent space. This explains the notation for the cotangent space.

It's worthwhile unpacking the isomorphism (7.9) explicitly, which means turning everything we did above around. Given a covector $dh|_x \in T_x^\star X$, we can define a function:

$$\begin{aligned} T_xX &\to \mathbb{R} \\ [\sigma] &\mapsto \partial_\sigma(h) \end{aligned}$$

Looking at this in co-ordinates shows that this is well defined and linear, and the resulting map

$$T_x^\star X \longrightarrow (T_xX)^\star$$

is a linear isomorphism (exercise). This isomorphism is the dual to the linear map in Proposition 7.8.

As another exercise, pick a chart around x, and consider the isomorphisms $\Delta_f : T_x X \xrightarrow{\sim} \mathbb{R}^n$ and $\nabla_f : T_x^\star X \xrightarrow{\sim} \mathbb{R}^n$. Then the proof of the proposition (if you unpack the definitions) is saying that ∇_f is the dual linear map to $(\Delta_f)^{-1}$, and vice-versa.

Given the previous proposition, we have a third way to define tangent vectors: we can declare that a tangent vector is a linear map from $T_x^\star X$ to $\mathbb{R}$, or equivalently it's a linear map from $C^\infty(X)$ to $\mathbb{R}$ which vanishes on the subspace $R_x(X)$. However, there is a slightly better version of this definition.

Definition 7.10. Let X be a manifold and let x be a point in X. A **derivation at** x is a linear map

$$\mathfrak{d} : C^\infty(X) \to \mathbb{R}$$

obeying the product rule

$$\mathfrak{d}(h_1 h_2) = h_1(x)\mathfrak{d}(h_2) + h_2(x)\mathfrak{d}(h_1) \tag{7.11}$$

for any two functions $h_1, h_2 \in C^\infty(X)$. We denote the set of all derivations at x by $\mathrm{Der}_x(X)$.

For any tangent vector $[\sigma] \in T_x X$ the operator ∂_σ is a derivation at x, this follows instantly from the ordinary product rule for functions of one variable. In fact these are the only derivations at x, because:

Proposition 7.12. *A linear map $\mathfrak{d} : C^\infty(X) \to \mathbb{R}$ is a derivation at x if and only if $\mathfrak{d}$ vanishes on the subspace $R_x(X)$ of functions having rank zero at x.*

The 'if' direction is a straight-forward exercise, but the other direction is tricky and we relegate it to Appendix D.

Definition 7.13 (*'Algebraist's definition'*)**.** Let x be a point in a manifold X. A **tangent vector** to x is a derivation at x.

We've already seen that tangent vectors are exactly the linear maps from $C^\infty(X)$ to $\mathbb{R}$ which vanish on $R_x(X)$, so Proposition 7.12 shows that we have an isomorphism

$$T_x X \xrightarrow{\sim} \mathrm{Der}_x(X)$$

given by sending $[\sigma]$ to ∂_σ, and hence this new definition is equivalent to our previous two. What's nice about this definition is that it only uses the fact that $C^\infty(X)$ is a ring (or more accurately an algebra over $\mathbb{R}$), so it can also be used in more algebraic contexts.

To summarize, a tangent vector in $T_x X$ can be thought of as either:

- A curve σ through x, up to tangency at x.
- An operation turning charts around x into vectors in $\mathbb{R}^n$, obeying the tangent vector transformation law.
- An operator from $C^\infty(X)$ to $\mathbb{R}$, obeying the product rule at x (7.11).
- An element of the dual of $T_x^\star X$.

If we fix a chart (U, f) around x then we can turn our tangent vector into an explicit vector $v \in \mathbb{R}^n$, or into an operator $\partial_{f(x),v}$ acting on $C^\infty(\tilde{U})$.

A covector in $T_x^\star X$ can be thought of as either:

- A function in $C^\infty(X)$, modulo the subspace $R_x(X)$.
- An operation turning charts around x into vectors in $\mathbb{R}^n$, obeying the covector transformation law.
- An element of the dual of $T_x X$.

Of course it can be confusing that so many different perspectives exist, but they are all useful.

7.3 Vector fields as derivations

Let's see what vector fields look like if we adopt our third definition of tangent vectors (Definition 7.13), as being operators on the space of functions.

We'll start with the easy case when X is an open set in $\mathbb{R}^n$. In this situation a vector field is just a smooth function:

$$\tilde{\xi} : X \to \mathbb{R}^n$$

At any point $x \in X$ we have a we have a partial derivative operator

$$\partial_{x,\tilde{\xi}|_x} : C^\infty(X) \to \mathbb{R}$$

which differentiates along the vector $\tilde{\xi}|_x$ at the point x. This means that if h is a smooth function in $C^\infty(X)$, then we can define a new function

$$\tilde{\xi}(h) \colon\ X \to \mathbb{R}$$

by:

$$\tilde{\xi}(h) \colon\ x \mapsto \partial_{x,\tilde{\xi}|_x}(h)$$

Explicitly, if the components of $\tilde{\xi}$ are $\tilde{\xi} = (\tilde{\xi}_1, ..., \tilde{\xi}_n)$ then:

$$\tilde{\xi}(h) = \tilde{\xi}_1 \frac{\partial h}{\partial x_1} \ + \ ... \ + \ \tilde{\xi}_n \frac{\partial h}{\partial x_n}$$

Each function $\tilde{\xi}_i$ is smooth, and the partial derivatives of h are themselves smooth functions, so this new function $\tilde{\xi}(h)$ is also smooth. This means we can view our vector field $\tilde{\xi}$ as an operator

$$\begin{aligned} \tilde{\xi} : C^\infty(X) &\to C^\infty(X) \\ h &\mapsto \tilde{\xi}(h) \end{aligned}$$

which turns functions into other functions. When we want to think in this way it's common to write vector fields in the form:

$$\tilde{\xi} = \sum_{i=1}^{n} \tilde{\xi}_i \frac{\partial}{\partial x_i}$$

For example, the notation $\frac{\partial}{\partial x_1}$ denotes the operator $h \mapsto \partial_{x_1} h$. It corresponds to a constant vector field

$$\begin{aligned} X &\to \mathbb{R}^n \\ x &\mapsto (1, 0, ..., 0) \end{aligned}$$

which sends any point to the first standard basis vector. Similarly $\frac{\partial}{\partial x_i}$ is a constant vector field that maps any point to the ith standard basis vector, and obvious any vector field is obtained by multplying these n constant vector fields by functions in $C^\infty(X)$ and then adding them together.

Recall that every operator $\partial_{x, \tilde{\xi}|_x}$ is a derivation at x, *i.e.* it satisfies the product rule (7.11). Letting x vary, this implies that the operator $\tilde{\xi}$ satisfies the following version of the product rule:

$$\tilde{\xi}(h_1 h_2) = h_1 \tilde{\xi}(h_2) + h_2 \tilde{\xi}(h_1) \tag{7.14}$$

for any two functions $h_1, h_2 \in C^\infty(X)$. An operator like this is called a *derivation.* Note the difference between a 'derivation' and a 'derivation at x'.

Now let's repeat this on an arbitrary manifold X. Suppose we have a smooth vector field:

$$\xi : X \to TX$$

At every point we have a tangent vector $\xi|_x \in T_x X$, and we have an associated operator

$$\partial_{\xi|_x} : C^\infty(X) \to \mathbb{R}$$

which is a derivation at x. If we take a function $h \in C^\infty(X)$ then we can define a new function

$$\xi(h) : X \to \mathbb{R}$$

by:

$$\xi(h) : x \mapsto \partial_{\xi|_x}(h)$$

We claim that (unsurprisingly) this new function $\xi(h)$ is smooth. To see this we have to look at $\xi(h)$ in co-ordinates, so pick a chart (U, f). For a given point $x \in U$, recall that we can evaluate $\partial_{\xi|_x}(h)$ by writing both $\xi|_x$ and h in co-ordinates, *i.e.* considering

$$v = \Delta_f(\xi_x) \qquad \text{and} \qquad \tilde{h} = h \circ f^{-1}$$

and then applying the operator $\partial_{f(x),v}$ to $\tilde{h}$ (or taking the dot product of v with $\nabla_f(dh|_x)$). We can also write ξ in this chart, it becomes the vector field:

$$\begin{aligned} \tilde{\xi} : \tilde{U} &\to \mathbb{R}^n \\ \tilde{x} &\mapsto \Delta_f(\xi|_{f^{-1}(\tilde{x})}) \end{aligned}$$

So if we write the function $\xi(h)$ in this chart we get the function:

$$\xi(h) \circ f^{-1} : \ \tilde{x} \mapsto \partial_{\tilde{x}, \tilde{\xi}|_{\tilde{x}}} \tilde{h}$$

This is exactly the function $\tilde{\xi}(\tilde{h}) \in C^\infty(\tilde{U})$, obtained by applying the vector field $\tilde{\xi}$ to the function $\tilde{h}$. Since $\tilde{\xi}(\tilde{h})$ is smooth, and we're working in an arbitrary chart, this shows that $\xi(h)$ is indeed smooth.

This means that our vector field ξ has given us an operator:

$$\xi : C^\infty(X) \to C^\infty(X)$$

Furthermore the product rule (7.14) holds, because for each point x the operator $\partial_{\xi|_x}$ is a derivation at x.

Definition 7.15. A **derivation** on a manifold X is a linear map

$$\mathfrak{D} : C^\infty(X) \to C^\infty(X)$$

such that the product rule

$$\mathfrak{D}(h_1 h_2) = h_1 \mathfrak{D}(h_2) + h_2 \mathfrak{D}(h_1)$$

holds for all $h_1, h_2 \in C^\infty(X)$. The set of all derivations on X is denoted by $\mathrm{Der}(X)$.

We've just seen that any smooth vector field ξ defines a derivation in $\mathrm{Der}(X)$. The converse is also true:

Proposition 7.16. *Any derivation $\mathfrak{D} \in \mathrm{Der}(X)$ defines a (smooth) vector field.*

Proof. Pick a $\mathfrak{D} \in \mathrm{Der}(X)$. For a fixed point $x \in X$, we can define a linear operator

$$\mathfrak{D}|_x : C^\infty(X) \to \mathbb{R}$$

by:

$$\mathfrak{D}|_x : h \mapsto \big(\mathfrak{D}(h)\big)|_x$$

The product rule (7.14) implies that $\mathfrak{D}|_x$ is a derivation at x, so by Proposition 7.12 it must be the partial derivative operator associated to some tangent vector in T_xX. Hence the function $\xi : x \mapsto \mathfrak{D}|_x$ is a vector field on X.

It remains to show that ξ is smooth. Pick any chart (U, f), then in this chart ξ becomes a vector field $\tilde{\xi} : \tilde{U} \to \mathbb{R}^n$, which we can write as

$$\sum_{i=1}^n \tilde{\xi}_i \frac{\partial}{\partial x_i}$$

for some functions $\tilde{\xi}_1, ..., \tilde{\xi}_n : \tilde{U} \to \mathbb{R}$. We need to show is that each of these functions $\tilde{\xi}_i$ is smooth. More specifically, let's fix a point $y \in U$, set $\tilde{y} = f(y) \in \tilde{U}$, and prove that each $\tilde{\xi}_i$ is smooth at $\tilde{y}$.

Let ϕ be a bump function on $\tilde{U}$ which is constantly equal to 1 inside some neighbourhood $\tilde{W}$ of $\tilde{y}$, and constantly equal to zero outside some larger neighbourhood. We can use ϕ to extend smooth functions in $C^\infty(\tilde{U})$ to smooth functions on X, just as we did in the proof of Proposition 7.3. In particular if we take one of the standard co-ordinates $x_i \in C^\infty(\tilde{U})$, we can get a function $\chi_i \in C^\infty(X)$, defined by:

$$\chi_i = \begin{cases} (x_i\phi) \circ f \text{ inside } U \\ 0 \text{ outside } U \end{cases}$$

Then if we write χ_i in the chart (U, f), we get a function $\tilde{\chi}_i \in C^\infty(\tilde{U})$ which agrees with x_i inside the open neighbourhood $\tilde{W}$ of $\tilde{y}$.

By definition, applying our derivation $\mathfrak{D}$ must send each χ_i to a smooth function $\mathfrak{D}(\chi_i) \in C^\infty(X)$. Let's compute this function $\mathfrak{D}(\chi_i)$ inside the open neighbourhood $W = f^{-1}(\tilde{W})$ of y. Its value at any point $x \in W$ is given by applying the operator $\mathfrak{D}|_x$ to χ_i, and we can compute this in the chart (U, f) and see that

$$\mathfrak{D}(\chi_i)|_x = \sum_{j=1}^{n} \tilde{\xi}_j|_{f(x)} \left.\frac{\partial \tilde{\chi}_i}{\partial x_j}\right|_{f(x)} = \tilde{\xi}_i|_{f(x)}$$

since $\tilde{\chi}_i \equiv x_i$ in the open set $\tilde{W}$. Since $\mathfrak{D}(\chi_i)$ is smooth, $\tilde{\xi}_i$ must be smooth inside $\tilde{W}$, and in particular smooth at $\tilde{y}$. □

So the set of all vector fields on X is exactly $\mathrm{Der}(X)$.

Example 7.17. Recall the angular vector field ξ on S^1 from Example 6.2, which takes values:

$$\xi|_{(x,y)} = (-y, x)^\top \in T_{(x,y)}S^1 \subset \mathbb{R}^2$$

Let's express ξ as a derivation on S^1.

If we use polar co-ordinates (U, f) with domain $U = S^1 \setminus (-1, 0)$, then we saw in Example 6.5 that ξ becomes the constant vector field $\tilde{\xi} \equiv 1$ on $\tilde{U} = (-\pi, \pi)$. As a derivation on $\tilde{U}$, this is the operator:

$$\begin{aligned} \tilde{\xi} : C^\infty(\tilde{U}) &\to C^\infty(\tilde{U}) \\ \tilde{h} &\mapsto \frac{d\tilde{h}}{d\theta} \end{aligned}$$

Now let h be a function in $C^\infty(S^1)$. Thinking of ξ as an operator in $\mathrm{Der}(S^1)$, we can apply to h to get a new function $\xi(h) \in C^\infty(S^1)$. Within the open set U, this function is given by

$$\xi(h)|_U = \left(\frac{d(h \circ f^{-1})}{d\theta}\right) \circ f$$

(and a similar expression holds in other polar co-ordinate charts).

We can be more explicit if we assume that h is the restriction to S^1 of some smooth function $\hat{h} \in C^\infty(\mathbb{R}^2)$. At any point $(a, b) \in S^1$ we have a partial derivative operator:

$$\partial_{\xi|_{(a,b)}} : C^\infty(S^1) \to \mathbb{R}$$

However since $T_{(a,b)}S^1$ is naturally a subspace of $\mathbb{R}^2$, we can also view this as a partial derivative operator

$$\partial_{\xi|_{(a,b)}} : C^\infty(\mathbb{R}^2) \to \mathbb{R}$$

and this is given by:

$$\hat{h} \mapsto -b \left.\frac{\partial \hat{h}}{\partial x}\right|_{(a,b)} + a \left.\frac{\partial \hat{h}}{\partial y}\right|_{(a,b)}$$

If $\hat{h} \in C^\infty(\mathbb{R}^2)$ then we can either apply the operator $\partial_{\xi|_{(a,b)}}$, or we can first restrict to S^1 to get a function $h = \hat{h}|_{S^1} \in C^\infty(S^1)$ and then apply the operator;

it follows immediately from the definition that these two operations produce the same real number. Consequently, we have:

$$\xi(\hat{h}|_{S^1}) = \left(-y\frac{\partial \hat{h}}{\partial x} + x\frac{\partial \hat{h}}{\partial y}\right)\Bigg|_{S^1} \quad \in C^\infty(S^1)$$

8 Differential forms

8.1 One-forms

A vector field on X is a gadget that selects, for each point $x \in X$, an element of the tangent space $T_x X$. In an analogous way we can define a *covector field*, which is something that selects a covector in $T_x^\star X$ for each $x \in X$. Covector fields are also called *one-forms*, because they are the first case of a more general object called a p-form, where p can be any natural number. We will meet p-forms later on.

Let's define one-forms precisely, following the same procedure that we used to define vector fields in Section 6.1. Just as we did for the tangent bundle, we can take all the cotangent spaces $T_x^\star X$ for each $x \in X$, and assemble them together to get a set

$$T^\star X = \bigcup_{x \in X} T_x^\star X$$

called the *cotangent bundle*. This comes with a projection function

$$\pi : T^\star X \to X$$

which sends a covector $u \in T_x^\star X$ to the corresponding point $x \in X$.

Definition 8.1. Let X be a manifold. A **covector field**, or **one-form**, is a function

$$\alpha : X \to T^\star X$$

such that $\pi \circ \alpha = 1_X$.

So a one-form selects a covector $\alpha|_x \in T_x^\star X$ for every point $x \in X$.

In some ways one-forms are a lot like vector fields. For example, if our manifold is just an open set $U \subset \mathbb{R}^n$, then at each point $x \in U$ the cotangent space $T_x^\star U$ is canonically isomorphic to $\mathbb{R}^n$, so the cotangent bundle is just:

$$T^\star U \cong U \times \mathbb{R}^n$$

So a one-form on U is a function $\alpha : U \to U \times \mathbb{R}^n$, and it must be of the form $\alpha = (1_U, \tilde{\alpha})$ for some some function $\tilde{\alpha} : U \to \mathbb{R}^n$. This means that on U a one-form consists of exactly the same data as a vector field. However, on more complicated manifolds there is a difference, because the way that one-forms change when we change co-ordinates is different from the way that vector fields change.

Suppose X is any manifold, and α is a one-form on X. If we choose a chart (U, f) on X then for any point $x \in U$ our co-ordinates give us an isomorphism:

$$\nabla_f : T_x^\star X \xrightarrow{\sim} \mathbb{R}^n$$

If we restrict α to U then we can look at it in our co-ordinates, and it becomes a one-form on $\tilde{U}$ given by the function:

$$\begin{aligned} \tilde{\alpha} : \tilde{U} &\to \mathbb{R}^n \\ \tilde{x} &\mapsto \nabla_f(\alpha|_{f^{-1}(\tilde{x})}) \end{aligned}$$

Now suppose we have two charts, and we write α in each chart, so we get one-forms:

$$\tilde{\alpha}_1 : \tilde{U}_1 \to \mathbb{R}^n \qquad \text{and} \qquad \tilde{\alpha}_2 : \tilde{U}_2 \to \mathbb{R}^n$$

The transformation law for covectors (7.4) implies that these are related by:

$$\tilde{\alpha}_2|_{f_2(x)} = (D\phi_{12}|_{f_2(x)})^\top \tilde{\alpha}_1|_{f_1(x)} \tag{8.2}$$

In particular $\tilde{\alpha}_1$ is smooth if and only if $\tilde{\alpha}_2$ is smooth. So just as we did for vector fields, we may define a one-form to be **smooth** iff its expression in any chart is smooth. From now on we will assume that all our one-forms are smooth.

As we did for vector fields (Proposition 6.6), we could take the transformation law (8.2) as the definition of a one-form. This means that we can specify a one-form on X by choosing an atlas $\mathcal{A} = \{(U_i, f_i)\}$ for X and choosing functions

$$\tilde{\alpha}_i : \tilde{U}_i \to \mathbb{R}^n$$

such that the correct transformation law holds.

Despite their many similiarities, one-forms are probably more important than vector fields. One reason for this is the following: for any smooth function $h \in C^\infty(X)$, there is an associated one-form on X, denoted by dh.

Let's understand this first in the easy case when X is an open subset of $\mathbb{R}^n$. Fix a function $h \in C^\infty(X)$. Then for any fixed point $x \in X$, our function h determines a covector $dh|_x \in T^\star_x X$, which is is just the equivalence class of h modulo the subspace $R_x(X)$. Since X is an open subset of $\mathbb{R}^n$ we can identify $T^\star_x X$ with $\mathbb{R}^n$, and then $dh|_x$ is the vector of partial derivatives of h at the point x. So we have a one-form:

$$\begin{aligned} dh : X &\to \mathbb{R}^n \\ x &\mapsto dh|_x = \left(\left.\frac{\partial h}{\partial x_1}\right|_x, \ldots, \left.\frac{\partial h}{\partial x_n}\right|_x \right) \end{aligned}$$

This is a smooth one-form, since the partial derivatives of h are smooth functions of x. Notice that if we set h to be one of the co-ordinate functions $x_i \in C^\infty(X)$ then dx_i is the constant one-form sending every point in U to the standard basis vector $e_i \in \mathbb{R}^n$, for example dx_1 is the one-form:

$$dx_1 : x \mapsto (1, 0, ..., 0)$$

This means that if we have any one-form $\alpha : X \to \mathbb{R}^n$, we can write it as

$$\alpha \;=\; \alpha_1 dx_1 + ... + \alpha_n dx_n \tag{8.3}$$

where $\alpha_1, ..., \alpha_n \in C^\infty(X)$ are the components of α. In particular, we have

$$dh = \frac{\partial h}{\partial x_1} dx_1 + ... + \frac{\partial h}{\partial x_n} dx_n \tag{8.4}$$

which is a very attractive equation.

Now suppose X is an arbitrary manifold. We want to show that for a function $h \in C^\infty(X)$, there is an associated (smooth) one-form dh. The definition is clear; for any $x \in X$ we have a covector $dh|_x \in T^\star_x X$, so we define:

$$\begin{aligned} dh : X &\to T^\star X \\ x &\mapsto dh|_x \end{aligned}$$

We just need to check that this one-form is smooth. Pick any chart (U, f), then h becomes a function $\tilde{h} \in C^\infty(\tilde{U})$. The one-form dh becomes a one-form on $\tilde{U}$, given by the function

$$\begin{aligned}\tilde{U} &\to \mathbb{R}^n \\ \tilde{x} &\mapsto \nabla_f(dh|_{f^{-1}(\tilde{x})}) = D\tilde{h}|_{\tilde{x}}\end{aligned}$$

i.e. if we write dh in co-ordinates, we get the one-form $d\tilde{h}$ on $\tilde{U}$. This is always smooth, so dh is indeed a smooth one-form on X.

Example 8.5. Let $X = S^2$, and let $h \in C^\infty(S^2)$ be the function:

$$h : (x, y, z) \mapsto z^2$$

This is the restriction to S^2 of a smooth function on $\mathbb{R}^3$, so it is smooth. Therefore there is an associated one-form dh on S^2.

Let's examine dh in co-ordinates. If we use the chart with domain $U_1 = S^2 \cap \{x > 0\}$ and co-ordinates $f_1 : (x, y, z) \mapsto (y, z)$, then h becomes the function $\tilde{h}_1 = z^2 \in C^\infty(\tilde{U}_1)$. Then $d\tilde{h}_1$ is the one-form

$$\begin{aligned}d\tilde{h}_1 : \tilde{U}_1 &\to \mathbb{R}^2 \\ (y, z) &\mapsto (0, 2z)\end{aligned}$$

which we could also write as

$$d\tilde{h}_1 = 2z\, dz$$

(since dz is the constant function taking the value $(0, 1)$ at all points). Alternatively we could use the chart with domain $U_2 = S^2 \cap \{z > 0\}$ and co-ordinates $f_2 : (x, y, z) \mapsto (x, y)$. In this chart h becomes $\tilde{h}_2 : (x, y) \to 1 - x^2 - y^2$, and $d\tilde{h}_2$ becomes the one-form:

$$d\tilde{h}_2 = -2x\, dx - 2y\, dy$$

If we wanted to we could also write down the transition function between these two charts, and verify the transformation law (8.2).

However, it's important to realize that not every one-form on X arises as dh for some $h \in C^\infty(X)$.

Example 8.6. Let $X = \mathbb{R}^2$, then any one-form on X is of the form

$$\alpha = \alpha_1 dx + \alpha_2 dy$$

for two smooth functions $\alpha_1, \alpha_2 \in C^\infty(\mathbb{R}^2)$. Now suppose that there is some $h \in C^\infty(\mathbb{R}^2)$ such that $\alpha = dh$. From (8.4) we know that

$$dh = \frac{\partial h}{\partial x} dx + \frac{\partial h}{\partial y} dy$$

which implies:

$$\frac{\partial \alpha_1}{\partial y} = \frac{\partial^2 h}{\partial x \partial y} = \frac{\partial \alpha_2}{\partial x}$$

So we cannot hope to find such an h unless α_1 and α_2 obey this equation (in fact on $\mathbb{R}^2$ this equation is a sufficient condition, but that is harder to show).

The previous example gives an easy 'local' condition showing why one-forms need not be d of any function; the next example shows a more subtle 'global' condition.

Example 8.7. Let $X = T^1 = \mathbb{R}/\mathbb{Z}$, and recall the smooth atlas from Example 2.11. This has two charts (U_1, f_1) and $(U_2, f_2$, having codomains $\tilde{U}_1 = (0,1)$ and $\tilde{U}_2 = (-\frac{1}{2}, \frac{1}{2})$, and the transition function is:

$$\begin{aligned} \phi_{21} : (0, \tfrac{1}{2}) \sqcup (\tfrac{1}{2}, 1) &\longrightarrow (-\tfrac{1}{2}, 0) \sqcup (0, \tfrac{1}{2}) \\ x &\mapsto \begin{cases} x, & \text{for } x < \frac{1}{2} \\ x - 1, & \text{for } x > \frac{1}{2} \end{cases} \end{aligned}$$

A function $h : X \to \mathbb{R}$ is, by definition, a function $\hat{h} : \mathbb{R} \to \mathbb{R}$ which is periodic, *i.e.*:

$$\hat{h}(x+n) = \hat{h}(x), \quad \forall x \in \mathbb{R},\ n \in \mathbb{Z}$$

If we look at h in our two charts we just get the function $\hat{h}$ (restricted to $\tilde{U}_1$ or $\tilde{U}_2$), so h is smooth precisely if $\hat{h}$ is smooth. So $C^\infty(T^1)$ is the space of periodic smooth functions on $\mathbb{R}$.

Now suppose α is a (smooth) one-form on T^1. If we write it in our two charts we get smooth functions:

$$\tilde{\alpha}_1 : \tilde{U}_1 \to \mathbb{R} \qquad \text{and} \qquad \tilde{\alpha}_2 : \tilde{U}_2 \to \mathbb{R}$$

We could also write these as $\tilde{\alpha}_1 dx$ and $\tilde{\alpha}_2 dx$, since in either chart dx is the constant one-form with value $1 \in \mathbb{R}$. Now the transition law for one-forms (8.2) says that these two expressions are related by

$$\tilde{\alpha}_2|_{\phi_{21}(x)} = \tilde{\alpha}_1|_x$$

since the derivative of ϕ_{21} is 1 at all points. So $\tilde{\alpha}_2 = \tilde{\alpha}_1$ on the interval $(0, \frac{1}{2})$, and on the interval $(-\frac{1}{2}, \frac{1}{2})$ the function $\tilde{\alpha}_2$ is obtained by translating the function $\tilde{\alpha}_1$. Specifying this data is exactly the same as specifying a (smooth) periodic function $\hat{\alpha} \in C^\infty(\mathbb{R})$, with:

$$\tilde{\alpha}_1 = \hat{\alpha}|_{\tilde{U}_1} \qquad \text{and} \qquad \tilde{\alpha}_2 = \hat{\alpha}|_{\tilde{U}_2}$$

So on T^1, a one-form is also the same thing a smooth periodic function on $\mathbb{R}$.

Now pick $h \in C^\infty(T^1)$, and let $\hat{h} \in C^\infty(\mathbb{R})$ be the associated periodic function. Then

$$d\hat{h} = \frac{d\hat{h}}{dx} dx$$

is a one-form on $\mathbb{R}$, and the coefficient $\frac{d\hat{h}}{dx}$ is a periodic function. If we look at h in either of our charts we get $\hat{h}$, so if we look at the one-form dh in either chart we must get $d\hat{h}$. So $\frac{d\hat{h}}{dx}$ is the periodic function associated to the one-form dh.

Now let $\hat{\alpha} \in C^\infty(\mathbb{R})$ be any periodic function. Is there a periodic function $\hat{h}$ such that $\hat{\alpha} = \frac{d\hat{h}}{dx}$? In other words, if α is the associated one-form on T^1, is there a function $h \in C^\infty(T^1)$ such that $dh = \alpha$? If there is, then by the fundamental theorem of calculus we have:

$$\int_0^1 \hat{\alpha}\, dx = \int_0^1 \frac{d\hat{h}}{dx}\, dx = h(1) - h(0) = 0$$

But there are plenty of periodic functions that do not satisfy this, for example the constant function $\hat{\alpha} \equiv 1$. So most one-forms on T^1 do not come from functions in $C^\infty(T^1)$.

The second reason that one-forms are important is because they behave nicely with respect to smooth functions between manifolds. Suppose

$$F : X \to Y$$

is a smooth function between two manifolds, and we have a one-form α on the manifold Y. We claim that there is an associated one-form $F^\star\alpha$ on the manifold X, called the *pull-back* of α along F.

To understand this, we first have to understand the dual of the derivative. Fix a point $x \in X$, and let $y = F(x) \in Y$. We have a linear map

$$DF|_x : T_xX \to T_yY$$

so we have a dual linear map:

$$DF|_x^\star : T_y^\star Y \to T_x^\star X$$

If we unpack the definitions, this turns out to be something very simple.

Lemma 8.8. *Let $F : X \to Y$ be smooth, fix $x \in X$ and let $y = F(x)$. Take a function $h \in C^\infty(Y)$, so we have a smooth function $h \circ F \in C^\infty(X)$. Then $DF|_x^\star$ sends the covector $dh|_y \in T_y^\star Y$ to the covector $d(h \circ F)|_x \in T_x^\star X$.*

Proof. Exercise. □

If we choose charts around x and y then we can identify both T_xX and $T_x^\star X$ with $\mathbb{R}^n$, and both T_yY and $T_y^\star Y$ with $\mathbb{R}^k$ (here n is the dimension of X and k is the dimension of Y). We know that if we do this the derivative $DF|_x$ becomes identified with the Jacobian matrix $D\tilde{F}|_{f(x)}$ where $\tilde{F}$ is F written in these charts; consequently $DF|_x^\star$ must become identified with transposed matrix $D\tilde{F}|_{f(x)}^\top$. One can also deduce this easily from the above lemma: in these charts $d(h \circ F)|_x$ becomes the vector

$$D(\tilde{h} \circ \tilde{F})|_{\tilde{x}} = D\tilde{h}|_{\tilde{y}} D\tilde{F}|_{\tilde{x}}$$

where $\tilde{x}, \tilde{y}$ and $\tilde{h}$ are the points x and y and the function h written in our chosen charts. This equation expresses a row-vector being transformed into a new row vector by a matrix acting on the right, so we must transpose it to express it in terms of column vectors.

Now we define the pull-back of a one-form.

Definition 8.9. Let $F : X \to Y$ be smooth, and let α be a one-form on Y. The **pull-back of α along** F is the one-form on X defined by:

$$F^\star\alpha : \ x \mapsto DF|_x^\star(\alpha|_{F(x)})$$

Of course we need to check that $F^\star\alpha$ is a smooth one-form. Let's start with the case where X is an open set $U \subset \mathbb{R}^n$ and Y is an open set $V \subset \mathbb{R}^k$, and we have a smooth function:

$$F = (F_1, ..., F_k) : U \to V$$

Then a one-form α on V is just a smooth function $\alpha : V \to \mathbb{R}^k$. If we pick a point $z \in U$ then the linear map $DF|_z^\star$ is just the transpose of the Jacobian matrix of F at z, so the pull-back of α along F is given by:

$$\begin{aligned} F^\star\alpha :\ & U \to \mathbb{R}^n \\ & z \mapsto (DF|_z)^\top \alpha|_{F(z)} \end{aligned}$$

This is smooth, since we are multiplying a vector of smooth functions by a matrix of smooth functions. Now we can let X and Y be any two manifolds; to check that $F^\star\alpha$ is smooth we have to pick charts, but then the question reduces to the situation that we just considered. So $F^\star\alpha$ is indeed smooth.

It's helpful to examine what this pull-back operation looks like when we write our one-forms in the style (8.3). Assume again that $F : U \to V$ is a smooth function between open subsets of $\mathbb{R}^n$ and $\mathbb{R}^k$, and let $(F_1, ..., F_k)$ be the components of F. Write $x_1, ..., x_n$ for the standard co-ordinates on U, and $y_1, ..., y_k$ for the standard co-ordinates on V.

We saw before that dy_j is the constant one-form on V which sends every point to the jth standard basis vector in $\mathbb{R}^k$. If we pull it back along F, it becomes a one-form on U given by

$$F^\star dy_j :\ z \mapsto \left(\frac{\partial F_j}{\partial x_1}, \ ... \ , \frac{\partial F_j}{\partial x_n}\right)\Bigg|_z^\top \in \mathbb{R}^n$$

since this expression is the jth column of $(DF|_z)^\top$. We may also write this as

$$F^\star dy_j = \frac{\partial F_j}{\partial x_1}\, dx_1 + ... + \frac{\partial F_j}{\partial x_n}\, dx_n \tag{8.10}$$

which is another nice equation. A general one-form α on V can be written as

$$\alpha \ = \ \alpha_1\, dy_1 + ... + \alpha_k\, dy_k$$

for some $\alpha_1, ..., \alpha_n \in C^\infty(V)$, and then

$$F^\star\alpha = (\alpha_1 \circ F)(F^\star dy_1) + ... + (\alpha_k \circ F)(F^\star dy_k) \tag{8.11}$$

since the value of $F^\star\alpha$ at a point $z \in U$ depends linearly on $\alpha|_{F(z)}$.

If your one-form happens to be of the form dh for some function h then there is another way to say what its pull-back is.

Lemma 8.12. *If $F : X \to Y$ is a smooth function and $h \in C^\infty(Y)$ then:*

$$F^\star dh = d(h \circ F)$$

Proof. This follows instantly from Lemma 8.8. □

Notice that the equation (8.10) is actually a special case of this lemma. If we take a smooth function $F : U \to V$ between open subsets of $\mathbb{R}^n$ and $\mathbb{R}^k$, then the component $F^j \in C^\infty(U)$ is the composition of F with the jth co-ordinate function $y_j \in C^\infty(V)$. So:

$$F^* dy_j = dF_j = \sum_{i=1}^n \frac{\partial F_j}{\partial x_i}\, dx_i$$

Example 8.13. Consider the inclusion map $\iota : S^2 \to \mathbb{R}^3$. Let x, y, z be the standard co-ordinates on $\mathbb{R}^3$, and let α be the one-form:

$$\alpha = 2z\, dz$$

Then $\iota^\star \alpha$ is a one-form on S^2.

Let's look at $\iota^\star \alpha$ in co-ordinates. Let's use the chart (U_2, f_2) on S^2 from Example 8.5, and the trivial chart on $\mathbb{R}^3$, and then the map ι becomes:

$$\tilde{\iota} : (x, y) \mapsto \big(x, y, \sqrt{1 - x^2 - y^2}\big)$$

By (8.10) we have

$$\tilde{\iota}^\star dz = \tfrac{-1}{\sqrt{1-x^2-y^2}} \big(x\, dx + y\, dy\big)$$

so by (8.11) we have:

$$\tilde{\iota}^\star \alpha = 2(z \circ \tilde{\iota}) \tilde{\iota}^\star dz = -2\big(x\, dx + y\, dy\big)$$

Alternatively, we can observe that $\alpha = d\hat{h}$, where $\hat{h} \in C^\infty(\mathbb{R}^3)$ is the function $\hat{h}(x, y, z) = z^2$. If we let $h = \hat{h} \circ \iota = \hat{h}|_{S^3}$, then Lemma 8.12 says that $\iota^\star dz = dh$. So to find the expression for $\iota^\star \alpha$ in the chart (U_2, f_2) we can write down the expression $\tilde{h}_2$ for h in these co-ordinates, and then compute $d\tilde{h}_2$. This is what we did in Example 8.5, and we got the same answer:

$$d\tilde{h}_2 = -2x\, dx - 2y\, dy$$

It's worth noticing that the transformation law for one-forms (8.2) is actually a special case of this pull-back operation. Suppose we have two charts (U_1, f_1) and (U_2, f_2) on X, so we have a transition function:

$$\phi_{21} : f_1(U_1 \cap U_2) \xrightarrow{\sim} f_2(U_1 \cap U_2)$$

Now pick a one-form α on X, which in our charts becomes:

$$\tilde{\alpha}_1 : \tilde{U}_1 \to \mathbb{R}^n \qquad \text{and} \qquad \tilde{\alpha}_2 : \tilde{U}_2 \to \mathbb{R}^n$$

Then on the overlap, the one-forms $\tilde{\alpha}_1$ and $\tilde{\alpha}_2$ are related by pull-back along the inverse transition function ϕ_{12}, *i.e.* we have

$$\tilde{\alpha}_2 = \phi_{12}^\star \tilde{\alpha}_1$$

on the open set $f_2(U_1 \cap U_2)$.

8.2 Antisymmetric multi-linear maps and the wedge product

Our next goal is generalize one-forms to p-forms, where p can be any natural number. We have to start by discussing quite a lot of multi-linear algebra for vector spaces, but once we've got that out the way the definition of p-forms on manifolds will be quite easy.

Let V be a vector space, of dimension n. We've previously consider the dual space of all linear maps $u : V \to \mathbb{R}$, now we're going to consider *bilinear* maps

$$b : V \times V \to \mathbb{R}$$

i.e. maps which are linear in each argument. If we choose a basis $e_1, ..., e_n$ for V then b is determined by its values on each pair of basis elements, and we can view this data as an n-by-n matrix B with entries:

$$B_{ij} = b(e_i, e_j)$$

Conversely any such matrix B specifies a bilinear map b, by extending linearly in each argument. The space of all such bilinear maps forms a vector space under point-wise addition and scalar multiplication, and once we choose a basis for V this vector space becomes identified with the vector space $\mathrm{Mat}_{n\times n}(\mathbb{R})$. So its dimension is n^2.

In fact we're only going to be interested in *antisymmetric* bilinear maps, which means that

$$b(v, \hat{v}) = -b(\hat{v}, v)$$

for all $v, \hat{v} \in V$ (which means in particular that $b(v, v) = 0$ for all v). If we choose a basis for V, these correspond to antisymmetric matrices. It's easy to check that the antisymmetric maps form a subspace of the space of all bilinear maps, so the set of all antisymmetric bilinear maps is a vector space. We denote it by:

$$\bigwedge^2 V^\star$$

Now suppose we have two elements $u, \hat{u}$ of the space $V^\star$. We can combine them to form an element of $\bigwedge^2 V^*$, by setting:

$$\begin{aligned} u \wedge \hat{u} : V \times V &\mapsto \mathbb{R} \\ (v, \hat{v}) &\mapsto u(v)\hat{u}(\hat{v}) - u(\hat{v})\hat{u}(v) \end{aligned}$$

Clearly $u \wedge \hat{u}$ is an antisymmetric bilinear map. We call it the *wedge product* of u and $\hat{u}$. This wedge product is an important structure, we can think of it as a kind of 'multiplication':

$$\begin{aligned} \wedge : V^\star \times V^\star &\to \bigwedge^2 V^\star \\ (u, \hat{u}) &\mapsto u \wedge \hat{u} \end{aligned}$$

It's a straight-forward exercise to check that this product is itself bilinear. It's also antisymmetric, since it's clear from the definition that:

$$u \wedge \hat{u} = -\hat{u} \wedge u$$

Now pick a basis $e_1, ..., e_n$ for V, and let $\epsilon_1, ..., \epsilon_n$ be the dual basis for $V^\star$. Choose a pair $i, j \in [1, n]$ with $i < j$, and form the bilinear map $\epsilon_i \wedge \epsilon_j \in \bigwedge^2 V^\star$. Applying this to pairs of basis vectors in V we get:

$$\epsilon_i \wedge \epsilon_j : (e_s, e_t) \mapsto \begin{cases} 1, & s = i \text{ and } t = j \\ -1, & s = j \text{ and } t = i \\ 0, & \text{otherwise} \end{cases}$$

So $\epsilon_i \wedge \epsilon_j$ corresponds to the matrix with a 1 in the (i, j) position (which is above the diagonal), a -1 in the (j, i) position (which is below the diagonal), and zeroes everywhere else. Clearly this set of matrices forms a basis for the space of all antisymmetric n-by-n matrices, so the set

$$\{\epsilon_i \wedge \epsilon_j \ ; \ i < j\} \subset \bigwedge^2 V^\star$$

is a basis. In particular, we have:

$$\dim \textstyle\bigwedge^2 V^\star = \binom{n}{2}$$

We can use these bases to describe the wedge product explicitly. If we take two elements

$$u = \lambda_1 \epsilon_1 + ... + \lambda_n \epsilon_n \qquad \text{and} \qquad \hat{u} = \mu_1 \epsilon_1 + ... + \mu_n \epsilon_n \quad \in V^\star$$

(here $\lambda_1, ..., \lambda_n$ and $\mu_1, ..., \mu_n$ are just real numbers) then their wedge product is:

$$\begin{aligned} u \wedge \hat{u} \; &= \big(\lambda_1\mu_2 - \lambda_2\mu_1\big)\epsilon_1 \wedge \epsilon_2 \; + \; \big(\lambda_1\mu_3 - \lambda_3\mu_1\big)\epsilon_1 \wedge \epsilon_3 \; + ... \\ &\quad ... + \; \big(\lambda_{n-1}\mu_n - \lambda_n\mu_{n-1}\big)\epsilon_{n-1} \wedge \epsilon_n \end{aligned}$$

Example 8.14. Let $V = \mathbb{R}^3$, and e_1, e_2, e_3 be the standard basis. Then $V^\star$ is also $\mathbb{R}^3$, and $\epsilon_1, \epsilon_2, \epsilon_3$ is again the standard basis. The dimension of $\bigwedge^2 V^\star$ is $\binom{3}{2} = 3$, and it has a basis:

$$\{\epsilon_2 \wedge \epsilon_3, \; \epsilon_1 \wedge \epsilon_3, \; \epsilon_1 \wedge \epsilon_2\}$$

Using these bases, the wedge product of two vectors $(\lambda_1, \lambda_2, \lambda_3)$ and (μ_1, μ_2, μ_3) is:

$$(\lambda_2\mu_3 - \lambda_3\mu_2, \; \lambda_1\mu_3 - \lambda_3\mu_1, \; \lambda_1\mu_2 - \lambda_2\mu_1)$$

If we flip the sign of the basis vector $\epsilon_1 \wedge \epsilon_3$, then the formula above becomes the usual cross-product of vectors in $\mathbb{R}^3$. This explains why there is no direct analogue of the cross-product in other dimensions, since if $n \neq 3$ then $\binom{n}{2} \neq n$. In other dimensions, the 'cross-product' of two vectors is really the wedge product, and it lands in $\mathbb{R}^{\binom{n}{2}}$.

We know that a linear map $F : V \to W$ induces a dual linear map $F^\star : W^\star \to V^\star$. It also induces a linear map

$$\textstyle\bigwedge^2 F^\star : \bigwedge^2 W^\star \to \bigwedge^2 V^\star$$

defined by

$$\textstyle\bigwedge^2 F^\star(b) : (v, \hat{v}) \mapsto b\big(F(v), F(\hat{v})\big) \;\; \in \mathbb{R}$$

for $b \in \bigwedge^2 W^\star$ and $v, \hat{v} \in V$. It's easy to check that $\bigwedge^2 F^\star(b)$ really is an element of $\bigwedge^2 V^\star$, and that $\bigwedge^2 F^\star$ really is linear in b. In terms of matrices, moving from F to $\bigwedge^2 F^\star$ is a rather complicated operation that turns a k-by-n matrix into an $\binom{n}{2}$-by-$\binom{k}{2}$ matrix (you have to take the determinants of each 2×2 minor), it's generally easier to work with the abstract definition.

It follows immediately from the definition that if $F : V \to W$ and $G : W \to U$ are two linear maps then we have:

$$\textstyle\bigwedge^2 (G \circ F)^\star = \bigwedge^2 F^\star \circ \bigwedge^2 G^\star$$

In particular, if F is an isomorphism, then so is $\bigwedge^2 F^\star$. If you know what a *functor* is, this says that the operation which sends V to $\bigwedge^2 V^\star$ is a contravariant functor (as is the operation which sends V to $V^\star$).

Now fix a natural number p. We're going to generalize from $\wedge^2 V^\star$ to $\wedge^p V^\star$. For any p we can consider p-linear maps from V to $\mathbb{R}$:

$$c : V^{\times p} \to \mathbb{R}$$

where $V^{\times p}$ means $V \times ...(p \text{ times})... \times V$. If we choose a basis $e_1, ..., e_n$ for V then c gives us a 'p-dimensional array' of numbers

$$C_{i_1,...,i_p} = c(e_{i_1}, ..., e_{i_p})$$

by evaluating c on each p-tuple of basis vectors, and conversely any such array of numbers determines a p-linear map by extending linearly in each argument. Hopefully it's clear that the set of all p-linear maps from V to $\mathbb{R}$ is a vector space, of dimension n^p.

We say that c is *antisymmetric* if c flips sign when we swap any two of its arguments, *i.e.*

$$c(v_1, ..., v_p) = -c(v_{\sigma(1)}, ..., v_{\sigma(p)})$$

for any transposition σ acting on the set $\{1, ..., p\}$. In particular this means that if we set any two of its arguments to be the same vector, then c must give the answer zero. If we apply a more general permutation $\sigma \in S_p$, we must have:

$$c(v_1, ..., v_p) = (-1)^\sigma c(v_{\sigma(1)}, ..., v_{\sigma(p)})$$

where $(-1)^\sigma$ is our notation for the sign of the permutation σ.

The antisymmetric maps form a subspace of the space of all p-linear maps from V to $\mathbb{R}$, and we denote this vector space by:

$$\wedge^p V^\star$$

Note that in the case $p = 1$ the antisymmetry condition is vacuous, so $\wedge^1 V^\star$ is just $V^\star$.

We've seen that two elements of $V^\star$ can be 'wedged' together to get an element of $\wedge^2 V^\star$. Similarly, if we have p elements $u_1, ..., u_p$ of $V^\star$, then we can combine them to get an element of $\wedge^p V^\star$, which we denote by $u_1 \wedge ... \wedge u_p$. We define it to be the map

$$u_1 \wedge ... \wedge u_p : V^{\times p} \to \mathbb{R}$$

which sends a p-tuple $(v_1, ..., v_p)$ to the real number:

$$\sum_{\sigma \in S_p} (-1)^\sigma u_1(v_{\sigma(1)}) \dots u_p(v_{\sigma(p)}) \quad \in \mathbb{R} \tag{8.15}$$

This map is clearly linear in each argument, and by construction it's antisymmetric, so it is indeed an element of $\wedge^p V^\star$.

Hence we've defined a 'p-fold wedge product':

$$\begin{aligned} (V^\star)^{\times p} &\to \wedge^p V^\star \\ (u_1, ..., u_p) &\mapsto u_1 \wedge ... \wedge u_p \end{aligned}$$

This product is p-linear and antisymmetric, since the expression (8.15) is linear in each u_i, and changes sign if we switch any u_i and u_j. In particular if any two u_i and u_j are equal then we get the zero element of $\wedge^p V^\star$.

Now suppose we choose a basis $e_1, ..., e_n$ for V, so we get a dual basis $\epsilon_1, ..., \epsilon_n$ for $V^\star$. We can produce elements in $\wedge^p V^\star$ by picking an p-tuple $(i_1, ..., i_p)$ of integers in $[1, n]$ and then forming the wedge product $\epsilon_{i_1} \wedge ... \wedge \epsilon_{i_p}$. If any entries in our p-tuple are repeated then this product must be zero. If our p-tuple contains no repeated entries, then it must be of the form $(\sigma(j_1), ..., \sigma(j_p))$ for some 'correctly-ordered' p-tuple $j_1 < ... < j_p$ and some permutation $\sigma \in S_p$. Then antisymmetry implies that:

$$\epsilon_{\sigma(j_1)} \wedge ... \wedge \epsilon_{\sigma(j_p)} = (-1)^\sigma \epsilon_{j_1} \wedge ... \wedge \epsilon_{j_p}$$

So up to sign, this procedure creates one element of $\wedge^p V^\star$ for each *subset* of $[1, n]$ of size p.

Proposition 8.16. *Let $e_1, ..., e_n$ be a basis for V, and let $\epsilon_1, ..., \epsilon_n$ be the dual basis for $V^\star$. Then the set of elements*

$$\left\{ \epsilon_{i_1} \wedge \epsilon_{i_2} \wedge .... \wedge \epsilon_{i_p} \;\middle|\; 1 \leq i_1 < i_2 < ... < i_p \leq n \right\} \subset \wedge^p V^\star$$

is a basis. In particular:

$$\dim \wedge^p V^\star = \binom{n}{p}$$

If $V = \mathbb{R}^n$ and $e_1, ..., e_n$ is the standard basis, then this proposition provides us with a basis for the space $\wedge^p (\mathbb{R}^n)^\star$. This means that we can identify $\wedge^p (\mathbb{R}^n)^\star$ with $\mathbb{R}^{\binom{n}{p}}$ if we wish, but this is not quite canonical, since there's no preferred way to order the basis vectors in $\wedge^p (\mathbb{R}^n)^\star$.

Proof. Choose a 'correctly-ordered' p-tuple $i_1 < ... < i_p$ with each entry in $[1, n]$, and form the p-linear map $\epsilon_{i_1} \wedge ... \wedge \epsilon_{i_p} \in \wedge^p V^\star$. Now take an arbitrary p-tuple $(j_1, ..., j_p)$ of numbers from the set $[1, n]$, and consider evaluating the map $\epsilon_{i_1} \wedge ... \wedge \epsilon_{i_p}$ on the p-tuple of basis vectors

$$(e_{j_1}, ..., e_{j_p}) \in V^{\times p}$$

using the defining formula (8.15). We can only get a non-zero result if the p-tuple $(j_1, ..., j_p)$ is a permutation of the p-tuple $(i_1, ..., i_p)$, in particular there must be no repetitions in the first p-tuple. So we have

$$\epsilon_{i_1} \wedge ... \wedge \epsilon_{i_p} : \; (e_{i_{\sigma(1)}}, ..., e_{i_{\sigma(p)}}) \mapsto (-1)^\sigma$$

for each $\sigma \in S_p$, and it vanishes on every other p-tuple of basis vectors.

A general antisymmetric p-linear map c is determined by its values on each p-tuple of basis vectors for V. It must vanish on p-tuples containing any repetition, and if $j_1 < ... < j_p$ is a correctly-ordered p-tuple then we must have

$$c : (e_{j_{\sigma(1)}}, ..., e_{j_{\sigma(p)}}) \mapsto (-1)^\sigma c(e_{j_1}, ..., e_{j_p})$$

for each $\sigma \in S^p$. This means that c can be written as a linear combination

$$c = \sum_{i_1 < ... < i_p} c(e_{i_1}, ..., e_{i_p}) \epsilon_{i_1} \wedge ... \wedge \epsilon_{i_p}$$

since both sides agree on any p-tuple of basis vectors for V.

This shows that this set of elements span $\wedge^p V^\star$. Futhermore if some linear combination of them gives the zero map then each coefficent must be zero, so they're linearly independent. $\square$

So the spaces $\bigwedge^p V^\star$ initially get larger as we increase p, but once $p > n/2$ then they start to get smaller again, and indeed we have a symmetry

$$\dim \bigwedge^p V^\star = \dim \bigwedge^{n-p} V^\star$$

for $1 \leq p \leq n-1$. The space $\bigwedge^n V^\star$ is only 1-dimensional, so there is a unique antisymmetric n-linear map from V to $\mathbb{R}$, up to scale. For $p > n$ the space $\bigwedge^p V^*$ is zero-dimensional, so there are no antisymmetric p-linear maps at all, except for the zero map.

We can extend the definition down to $p = 0$ by *declaring* that $\bigwedge^0 V^\star = \mathbb{R}$. This is analogous to the rule that $x^0 = 1$ for real numbers, and it satisfies $\dim \bigwedge^0 V^\star = 1 = \binom{n}{0}$.

An element of $\bigwedge^p V^\star$ is called *decomposable* if it lies in the image of the 'p-fold wedge product' $(V^\star)^{\times p} \to \bigwedge^p V^\star$, *i.e.* if it can be written in the form $u_1 \wedge ... \wedge u_p$. Proposition 8.16 implies that any element of $\bigwedge^p V^\star$ can be written as a linear combination of decomposable elements. However, it is not true that every element is decomposable (see Problem Sheets).

If $F : V \to W$ is a linear map, then we get an induced linear map

$$\bigwedge^p F^\star : \bigwedge^p W^\star \to \bigwedge^p V^\star$$

just as we did in the case $p = 2$, by defining:

$$\bigwedge^p F^\star(c) : (v_1, ..., v_p) \mapsto c\big(F(v_1), ..., F(v_p)\big)$$

Again it's immediate that $\bigwedge^p (G \circ F)^\star = \bigwedge^p F^\star \circ \bigwedge^p G^\star$. For a decomposable element $c = u_1 \wedge ... \wedge u_p$, it's easy to check that

$$\bigwedge^p F^\star(u_1 \wedge ... \wedge u_p) = \big(F^\star(u_1)\big) \wedge ... \wedge \big(F^\star(u_p)\big) \tag{8.17}$$

(apply both sides to any p-tuple $(v_1, ..., v_p) \in V^{\times p}$ and check that they give the same number). We extend down to the case $p = 0$ by declaring that for any F, the map $\bigwedge^0 F^\star$ is just the identity map from $\bigwedge^0 W^\star = \mathbb{R}$ to $\bigwedge^0 V^\star = \mathbb{R}$.

If we try to understand $\bigwedge^p F^\star$ in terms of matrices then it gets rather complicated, but there is one simple special case which we'll now explain. Assume that $\dim V = \dim W = n$, and consider $\bigwedge^n F^\star$. Pick bases $e_1, ..., e_n$ for V and $f_1, ..., f_n$ for W, so we can express F as a matrix $M \in \mathrm{Mat}_{n\times n}(\mathbb{R})$ where:

$$F : e_i \mapsto M_{1,i} f_1 + ... + M_{n,i} f_n$$

Now let $\epsilon_1, ..., \epsilon_n \in V^\star$ and $\phi_1, ..., \phi_n \in W^\star$ be the dual bases. We know that $\bigwedge^n V^\star$ and $\bigwedge^n W^\star$ are both one-dimensional, with basis vectors $\epsilon_1 \wedge ... \wedge \epsilon_n$ and $\phi_1 \wedge ... \wedge \phi_n$, so the 'matrix' describing $\bigwedge^n F^\star$ is a single real number. To compute this number, we observe that

$$\begin{aligned}
\bigwedge^n F^\star(\phi_1 \wedge ... \wedge \phi_n) :\ (e_1, ..., e_n) &\mapsto \phi_1 \wedge ... \wedge \phi_n\big(F(e_1), ..., F(e_n)\big) \\
&= \sum_{\sigma \in S_p} (-1)^\sigma \phi_1\big(F(e_{\sigma(1)})\big) ... \phi_n\big(F(e_{\sigma(n)})\big) \\
&= \sum_{\sigma \in S_p} (-1)^\sigma M_{1,\sigma(1)} ... M_{n,\sigma(n)} \\
&= \det(M)
\end{aligned}$$

so:

$$\bigwedge^n F^\star(\phi_1 \wedge ... \wedge \phi_n) = \det(M)\, \epsilon_1 \wedge ... \wedge \epsilon_n \tag{8.18}$$

This is an important observation that we will use later on.

The next thing we want to do is extend the wedge product to give a bilinear product

$$\bigwedge^p V^\star \times \bigwedge^q V^\star \to \bigwedge^{p+q} V^* \tag{8.19}$$

for any p and q. It is possible to write down an explicit formula for this product, but it's not very enlightening, so instead we're going to approach it indirectly.

Suppose we have two decomposable elements:

$$u_1 \wedge ... \wedge u_p \in \bigwedge^p V^\star \qquad \text{and} \qquad \hat{u}_1 \wedge ... \wedge \hat{u}_q \in \bigwedge^q V^\star$$

Clearly we would like their wedge product to be given by:

$$\big(u_1 \wedge ... \wedge u_p\big) \wedge \big(\hat{u}_1 \wedge ... \wedge \hat{u}_q\big) = u_1 \wedge ... \wedge u_p \wedge \hat{u}_1 \wedge ... \wedge \hat{u}_q$$

Since everything is a linear combination of decomposable elements we should be able to extend this rule bilinearly, and define the wedge product of any two elements. However it's not immediately obvious that this is well-defined, because the expression of an element in $\bigwedge^p V^\star$ as a linear combination of decomposable elements is not at all unique.

Alternatively, we could fix a basis for V, and use the induced bases for all the $\bigwedge^p V^\star$'s as above. Then if we take a basis vector $\epsilon_{i_1} \wedge ... \wedge \epsilon_{i_p}$ in $\bigwedge^p V^\star$, and a basis vector $\epsilon_{j_1} \wedge ... \wedge \epsilon_{j_q}$ in $\bigwedge^q V^\star$, their wedge product should be:

$$\epsilon_{i_1} \wedge ... \wedge \epsilon_{i_p} \wedge \epsilon_{j_1} \wedge ... \wedge \epsilon_{j_q} \in \bigwedge^{p+q} V^\star \tag{8.20}$$

This product is zero if the two subsets $\{i_1, ..., i_p\}$ and $\{j_1, ..., j_q\} \subset [1, n]$ are not disjoint. If they are disjoint, there's some 'shuffle' permutation $\sigma \in S_{p+q}$ that returns the $(p+q)$-tuple $(i_1, ..., i_q, j_1, ..., j_q)$ to its correct order, and the product (8.20) is equal to $(-1)^\sigma$ times the corresponding basis vector in $\bigwedge^{p+q} V^\star$.

This rule defines the wedge product (8.19) on each pair of basis vectors, then we can extend bilinearly. This approach is perfectly well-defined, but it's not obvious that it doesn't depend on our choice of basis. The next lemma says that both of these approaches work and give the same thing; so the first is well-defined, and the second is basis-independent.

Lemma 8.21. *For each p, q there is a unique bilinear map from $\bigwedge^p V^\star \times \bigwedge^q V^\star$ to $\bigwedge^{p+q} V^\star$ which makes the following triangle commute:*

$$\begin{tikzcd} (V^\star)^{\times p} \times (V^\star)^{\times q} \arrow[r] \arrow[dr] & \bigwedge^p V^\star \times \bigwedge^q V^\star \arrow[d] \\ & \bigwedge^{p+q} V^\star \end{tikzcd}$$

Proof of Lemma 8.21. Suppose we've found such a bilinear map (the vertical arrow in the triangle). Choose a basis $\epsilon_1, ..., \epsilon_n$ of $V^\star$. If we take any two basis vectors $\epsilon_{i_1} \wedge ... \wedge \epsilon_{i_p} \in \bigwedge^p V^\star$ and $\epsilon_{j_1} \wedge ... \wedge \epsilon_{j_q} \in \bigwedge^q V^\star$ then commutativity of the triangle forces the product of these two basis vectors to be given by the expression (8.20), and then bilinearity determines all other products. This proves uniqueness.

Now we prove existence. Pick a basis again, and define a map $\wedge^p V^\star \times \wedge^q V^\star$ to $\wedge^{p+q} V^\star$ by applying the rule (8.20) for each pair of basis vectors, and then extending bilinearly. We have to check that this product really does make the triangle commute, and by multi-linearity it's sufficient to check this on any $(p+q)$-tuple $(\epsilon_{i_1}, ..., \epsilon_{i_p}, \epsilon_{j_1}, .., \epsilon_{j_q})$ of basis vectors for $V^\star$. If any entries are repeated in this $(p+q)$-tuple then going either way around the triangle gives the answer zero. If no entries are repeated then going either way around the triangle gives $\pm$ the corresponding basis vector in $\wedge^{p+q} V^\star$, so we just need to check that the signs match. If go diagonally across the triangle then we get the sign of the permutation $\sigma \in S_{p+q}$ that restores this $(p+q)$-tuple to its correct order. We can factor σ as 'first correctly order $(i_1, ..., i_p)$, then correctly order $(j_1, ..., j_q)$, then shuffle them together', and this corresponds exactly to the sign that we pick up by going the other way around the triangle. □

This extended version of the wedge product behaves very nicely, as the next proposition shows.

Proposition 8.22. *(i) For any $c \in \wedge^p V^\star$, $\hat{c} \in \wedge^q V^\star$ and $\overline{c} \in \wedge^r V^\star$, we have:*

$$c \wedge (\hat{c} \wedge \overline{c}) = (c \wedge \hat{c}) \wedge \overline{c} \ \in \wedge^{p+q+r} V^\star$$

(ii) For any $c \in \wedge^p V^\star$ and $\hat{c} \in \wedge^q V^\star$ we have:

$$c \wedge \hat{c} = (-1)^{pq} \hat{c} \wedge c \ \in \wedge^{p+q} V^\star$$

(iii) If we have a linear map $F : U \to V$ then for any $c \in \wedge^p V^\star$ and $\hat{c} \in \wedge^q V^\star$ we have:

$$\wedge^{p+q} F^\star (c \wedge \hat{c}) = \big(\wedge^p F^\star(c)\big) \wedge \big(\wedge^q F^\star(\hat{c})\big)$$

Proof. By bilinearity it's sufficient to check all properties on decomposable elements:

$$c = u_1 \wedge ... \wedge u_p, \qquad \hat{c} = \hat{u}_1 \wedge ... \wedge \hat{u}_q, \qquad \overline{c} = \overline{u}_1 \wedge ... \wedge \overline{u}_r$$

Property (i) is obvious. Property (ii) just says that the sign of the permutation $\big((p+1)...(p+q)1....p\big) \in S_{p+q}$ is $(-1)^{pq}$. Property (iii) follows immediately from the observation (8.17). □

We can also extend the wedge product down to the case when $p = 0$ (or $q = 0$), by declaring that if $\lambda \in \wedge^0 V^\star = \mathbb{R}$ and $c \in \wedge^q V^\star$ then $\lambda \wedge c$ is just λc, the scalar multiple. It's trivial to check that the properties in Proposition 8.22 continue to hold in this case.

If we take the direct sum of all our $\wedge^p V^\star$'s we get a single vector space

$$\wedge^\bullet V^\star = \bigoplus_{p=0}^{n} \wedge^p V^\star$$

called the *exterior algebra* of $V^\star$. We can give $\wedge^\bullet V^\star$ a (bilinear) multiplication by using our wedge product for each component. Property (i) in Proposition 8.22 says that this structure is an *associative algebra*, and it has a unit $1 \in \wedge^0 \mathbb{R}$. Property (ii) in the proposition says that this algebra is *supercommutative*.

Property (iii) says that $\bigwedge^\bullet$ is a contravariant functor from vector spaces to algebras.

If these words are unfamiliar, don't worry, you may safely ignore the last paragraph.

8.3 p-forms

Finally we can return to manifolds!

If X is a manifold of dimension n, and x is a point in X, then we have a vector space T_xX. Hence for any $p \geq 0$ we may form the vector space of all p-linear antisymmetric maps from T_xX to $\mathbb{R}$, and this is conventionally denoted by:

$$\bigwedge^p T_x^\star X$$

(rather than $\bigwedge^p (T_xX)^\star$). In the case $p = 1$ we know that $\bigwedge^1 (T_xX)^*$ is just $(T_xX)^*$, but this is precisely the cotangent space $T_x^\star X$ (Proposition 7.8). So at each point $x \in X$ we have a sequence of vector spaces, starting with the cotangent space to x, and ending with the 1-dimensional space $\bigwedge^n T_x^\star X$. We can also go down to $p = 0$, since our convention says that $\bigwedge^0 T_x^\star X = \mathbb{R}$.

If we let x vary, we can form the set

$$\bigwedge^p T^\star X = \bigcup_{x \in X} \bigwedge^p T_x^\star X$$

which is called the *p-th wedge power of the cotangent bundle*. This comes with a projection map $\pi : \bigwedge^p T^\star X \to X$ which maps an element of $\bigwedge^p T_x^\star X$ down to the corresponding point $x \in X$.

Definition 8.23. A **p-form** on X is a function $\alpha : X \to \bigwedge^p T^\star X$ such that $\pi \circ \alpha = 1_X$.

So a p-form selects an element of $\bigwedge^p T_x^\star X$ for each $x \in X$. If we don't wish to specify p we can use the phrase **differential form**, which means a p-form, for some p. In the case $p = 1$ we recover one-forms. Also notice what happens in the case $p = 0$: since $\bigwedge^0 T_x^\star X = \mathbb{R}$ for all points $x \in X$, the set $\bigwedge^0 T^\star X$ is simply $X \times \mathbb{R}$, and a 0-form is just a function from X to $\mathbb{R}$.

We still need to give a definition of a *smooth* p-form, but this is very similar to the story that we saw for vector fields and one-forms. If our manifold is just an open set $U \subset \mathbb{R}^n$ then each tangent space T_xU is just $\mathbb{R}^n$, so $\bigwedge^p T_x^\star U$ is canonically isomorphic to $\bigwedge^p (\mathbb{R}^n)^\star$. Recall that this is a vector space of dimension $\binom{n}{p}$, and it has a 'standard basis' provided by Proposition 8.16. So

$$\bigwedge^p T^\star U \;\cong\; U \times \bigwedge^p (\mathbb{R}^n)^\star$$

and a p-form on U is the same thing as a function from U to $\bigwedge^p (\mathbb{R}^n)^\star$. If we want to we can identify $\bigwedge^p (\mathbb{R}^n)^\star$ with $\mathbb{R}^{\binom{n}{p}}$ by choosing an ordering of the standard basis, then a p-form on U is just a function from U to $\mathbb{R}^{\binom{n}{p}}$. So we know what it means for a p-form on U to be smooth.

Now suppose X is an arbitrary manifold, and we choose a chart (U, f) on X. Then for each $x \in U$ we have an isomorphism $\Delta_f : T_xX \xrightarrow{\sim} \mathbb{R}^n$, so we get an induced isomorphism

$$\bigwedge^p \left(\Delta_f^{-1}\right)^\star : \bigwedge^p T_x^\star X \xrightarrow{\sim} \bigwedge^p (\mathbb{R}^n)^\star$$

(in the case $p = 1$ this the usual isomorphism $(\Delta_f^{-1})^\star = \nabla_f$). If α is a p-form on X then we can look at it in these co-ordinates, and it becomes a p-form

$$\begin{aligned}\tilde{\alpha} : \tilde{U} &\longrightarrow \wedge^p(\mathbb{R}^n)^\star \\ x &\mapsto \wedge^p \left(\Delta_f^{-1}\right)^\star \left(\alpha|_{f^{-1}(x)}\right)\end{aligned}$$

on $\tilde{U}$. If we change co-ordinates by some transition function ϕ_{21}, then the two expressions for α will be related by applying the linear map

$$\wedge^p \left(D\phi_{12}|_{f_2(x)}\right)^\star : \wedge^p(\mathbb{R}^n)^\star \to \wedge^p(\mathbb{R}^n)^\star \tag{8.24}$$

at each point. This is a matrix of smooth functions (of size $\binom{n}{p} \times \binom{n}{p}$), so if α looks smooth in the first chart then it will also look smooth in the second chart. So just as we did for vector fields and one-forms, we may define a smooth p-form on X to be a p-form that becomes smooth when we write it in any chart.

If we have a p-form α, and a q-form β, then we can 'wedge them together' by forming their wedge-product at every point. This gives us a $(p+q)$-form:

$$\begin{aligned}\alpha \wedge \beta :\ X &\to \wedge^{p+q} T^\star X \\ x &\mapsto \alpha|_x \wedge \beta|_x \quad \in \wedge^{p+q} T^\star_x X\end{aligned}$$

We claim that if α and β are smooth then $\alpha \wedge \beta$ is also smooth. As usual we only need to check this claim in the special case that our manifold is an open set in $\mathbb{R}^n$, because on any other manifold we just pick a chart and then it reduces to that case. So let's examine this 'wedging' process in that special case.

Let $X = U \subset \mathbb{R}^n$ be an open subset, and let $x_1, ..., x_n$ be the standard co-ordinates on U. Recall that each co-ordinate function $x_j \in C^\infty(U)$ gives us a constant one-form dx_j, and at any point $z \in U$ the covector $dx_j|_z$ is the jth standard basis vector in $T^\star_z U \cong \mathbb{R}^n$. This means that if we take a correctly-ordered p-tuple $i_1 < ... < i_p$, and form the function

$$dx_{i_1} \wedge ... \wedge dx_{i_p} :\ U \to \wedge^p(\mathbb{R}^n)^\star$$

then at any point $z \in U$ this function just gives us one of the standard basis vectors in $\wedge^p(\mathbb{R}^n)^\star$. This is a constant function, so it's certainly smooth, and so this is a smooth p-form.

A general p-form on U is given by some smooth function $\alpha : U \to \wedge^p(\mathbb{R}^n)^\star$, so it has one component $\alpha_{\mathbf{i}} \in C^\infty(U)$ for each correctly-ordered p-tuple $\mathbf{i} = \{i_1 < ... < i_p\}$. Then we may write α as

$$\alpha = \sum_{\mathbf{i}} \alpha_{\mathbf{i}}\, dx_{i_1} \wedge ... \wedge dx_{i_p}$$

where $\mathbf{i}$ runs over all correctly-ordered p-tuples.

Now we can look at the wedge product of two differential forms on U. For example, if we have two 1-forms

$$\alpha = \alpha_1 dx_1 + ... + \alpha_n dx_n \qquad \text{and} \qquad \beta = \beta_1 dx_1 + ... + \beta_n dx_n$$

then their wedge-product is the two-form:

$$\alpha \wedge \beta = \sum_{i<j} (\alpha_i \beta_j - \alpha_j \beta_i)\, dx_i \wedge dx_j$$

All the coefficient functions are obviously smooth. In general, if we wedge together a p-form α and a q-form β then the coefficient functions for $\alpha \wedge \beta$ will be linear combinations of products of the coefficient functions for α and β, so this is inded a smooth $(p+q)$-form.

So on a general manifold, we have a 'wedge product' on smooth differential forms. Moreover it follows instantly from parts (i) and (ii) of Proposition 8.22 that we have

$$\alpha \wedge \beta = (-1)^{pq} \beta \wedge \alpha$$

and

$$(\alpha \wedge \beta) \wedge \gamma = \alpha \wedge (\beta \wedge \gamma)$$

for any p-form α, any q-form β, and any r-form γ. As a special case we can 'wedge' a p-form α with a 0-form, and this just means we multiply α by a function $h \in C^\infty(X)$.

If have a smooth function $F : X \to Y$ between two manifolds, then we've already seen that we can pull-back 1-forms along F. This is also true for p-forms, since if α is a p-form on Y then we can define a p-form on X by:

$$\begin{aligned} F^\star \alpha : X &\to \wedge^p T^\star X \\ x &\mapsto \wedge^p (DF|_x)^\star \left(\alpha|_{F(x)} \right) \end{aligned}$$

We need to check that $F^\star \alpha$ is smooth, but we should first note that it follows immediately from Proposition 8.22(iii) that

$$F^\star(\alpha \wedge \beta) = F^\star(\alpha) \wedge F^\star(\beta) \tag{8.25}$$

for any two differential forms α and β on Y. Now it is easy to check that $F^\star \alpha$ is smooth, because in co-ordinates we can write α as a linear combination of wedge products of one-forms, but we know that the pull-back of a one-form is smooth, and that the wedge-product of any differential forms is always smooth.

Notice that if α is just a zero-form, *i.e.* an element of $C^\infty(Y)$, then $F^\star \alpha$ is just $\alpha \circ F \in C^\infty(X)$. This is because by definition $\wedge^0 (DF|_x)^\star$ is always the identity map on $\mathbb{R}$.

Example 8.26. Let $U = \mathbb{R}_{>0} \times (-\pi, \pi) \subset \mathbb{R}^2$, and consider the smooth function

$$\begin{aligned} F :\ U &\to \mathbb{R}^2 \\ (r, \theta) &\mapsto (r \cos\theta, r \sin\theta) \end{aligned}$$

(the inverse to polar co-ordinates). The one-forms dx and dy on $\mathbb{R}^2$ pull-back via F to give one-forms

$$F^\star dx = \cos\theta \, dr - r \sin\theta \, d\theta \qquad \text{and} \qquad F^\star dy = \sin\theta \, dr + r \cos\theta \, d\theta$$

on U, so the 2-form $dx \wedge dy$ must pull-back to give the 2-form:

$$\begin{aligned} F^*(dx \wedge dy) &= \big(\cos\theta \, dr - r \sin\theta \big) \wedge \big(\sin\theta \, dr + r \cos\theta \, d\theta \big) \\ &= r \cos^2\theta \, dr \wedge d\theta - r \sin^2\theta \, d\theta \wedge dr \\ &= r \, dr \wedge d\theta \end{aligned}$$

We noted earlier that the transformation law for one-forms is actually just 'pull-back along the transition function', and the same thing is true for p-forms. Suppose we have a p-form α on X, and we write it in two different charts, so we get p-forms $\tilde{\alpha}_1$ on $\tilde{U}_1$ and $\tilde{\alpha}_2$ on $\tilde{U}_2$. Then the transformation law (8.24) says exactly that

$$\tilde{\alpha}_2 = \phi_{12}^\star \tilde{\alpha}_1$$

(on the open set $f_2(U_1 \cap U_2) \subset \tilde{U}_2$). As for vector fields and one-forms, we could take this transformation law as the *definition* of a p-form if we wished.

If we try and write this transformation law explicitly in terms of the Jacobian matrix $D\phi_{12}|_{f_2(x)}$ then it generally gets rather complicated. However, it is very easy in the special case of n-forms (where $n = \dim X$). If α is an n-form, then in each chart α just becomes smooth functions:

$$\tilde{\alpha}_1 : \tilde{U}_1 \to \mathbb{R} \qquad \text{and} \qquad \tilde{\alpha}_2 : \tilde{U}_2 \to \mathbb{R}$$

By our observation (8.18), on the overlap these functions are related by:

$$\tilde{\alpha}_2|_{f_2(x)} = \det\big(D\phi_{12}|_{f_2(x)}\big)\, \tilde{\alpha}_1|_{f_1(x)} \tag{8.27}$$

So n-forms transform by the determinant of the Jacobian matrix of the transition function.

Example 8.28. Look again at Example 8.26. We have $\det DF|_{(r,\theta)} = r$, so

$$F^\star(dx \wedge dy) = \det(DF)\, dr \wedge d\theta = r\, dr \wedge d\theta$$

as we saw before.

8.4 The exterior derivative

There is an extremely important operation that can be performed on differential forms, called the *exterior derivative*.

Let's denote the set of all (smooth) p-forms on X by:

$$\Omega^p(X)$$

This is a vector space, and it's clearly infinite-dimensional, because we can produce a huge number of p-forms on X by picking a p-form in some chart and then extending it to X with a bump function. In the case $p = 0$, we have that $\Omega^0(X) = C^\infty(X)$, since 0-forms are just smooth functions from X to $\mathbb{R}$. Now recall that for any smooth function $h \in C^\infty(X)$ we produced a one-form dh, and in co-ordinates this is given by:

$$dh = \frac{\partial h}{\partial x_1} dx_1 + \dots + \frac{\partial h}{\partial x_n} dx_n$$

So we have an operator:

$$\begin{aligned} d :\ \Omega^0(X) &\to \Omega^1(X) \\ h &\mapsto dh \end{aligned}$$

What we want to do is extend this to an operator

$$d :\ \Omega^p(X) \to \Omega^{p+1}(X)$$

for all p.

We'll begin by assuming X is an open set $U \subset \mathbb{R}^n$, and we'll let $x_1, ..., x_n$ be the co-ordinate functions on U. Suppose we have a p-form α which has only one non-zero component, so

$$\alpha = \alpha_{\mathbf{i}}\, dx_{i_1} \wedge ... \wedge dx_{i_p}$$

where $\mathbf{i} = (i_1 < ... < i_p)$ is a single correctly-ordered p-tuple, and $\alpha_{\mathbf{i}} \in C^\infty(U)$. Then we define:

$$\begin{aligned} d\alpha &= d\alpha_{\mathbf{i}} \wedge dx_{i_1} \wedge ... \wedge dx_{i_p} \\ &= \sum_{j=1}^{n} \frac{\partial \alpha_{\mathbf{i}}}{\partial x_j}\, dx_j \wedge dx_{i_1} \wedge ... \wedge dx_{i_p} \end{aligned} \tag{8.29}$$

Some of the terms in this sum will be zero, since when j is equal to one of the i_t then that wedge product of one-forms is zero. So there will be one (potentially) non-zero component of $d\alpha$ for each j which does not appear in the p-tuple $\mathbf{i}$, and if we want to write it in terms of our standard basis then we pick up a ± 1 when we apply the permutation which returns the $(p+1)$-tuple $(j, i_1, ..., i_p)$ to its correct order.

The expression (8.29) is linear in $\alpha_{\mathbf{i}}$, so we can extend it to a linear operator:

$$d : \Omega^p(U) \to \Omega^{p+1}(U)$$

Together, these operators are called the *exterior derivative* (or the *de Rham differential*).

Example 8.30. Let $X = \mathbb{R}^3$, with co-ordinates x, y and z. If we have a one-form $\alpha = \alpha_2\, dy$ for some $\alpha_2 \in C^\infty(\mathbb{R}^2)$, then:

$$d\alpha = \frac{\partial \alpha_2}{\partial x}\, dx \wedge dy - \frac{\partial \alpha_2}{\partial z}\, dy \wedge dz$$

More generally, if

$$\alpha = \alpha_1\, dx + \alpha_2\, dy + \alpha_3\, dz$$

then:

$$d\alpha = \left(\frac{\partial \alpha_2}{\partial x} - \frac{\partial \alpha_1}{\partial y}\right) dx \wedge dy \;+\; \left(\frac{\partial \alpha_3}{\partial x} - \frac{\partial \alpha_1}{\partial z}\right) dx \wedge dz \;+\; \left(\frac{\partial \alpha_3}{\partial y} - \frac{\partial \alpha_2}{\partial z}\right) dy \wedge dz$$

You might recognise this formula - if we flip the sign of middle term, which we can do by deciding to write things in terms of $dz \wedge dx$ instead of $dx \wedge dz$, then this is the formula for the curl operator $\nabla\times$ which turns a vector field on $\mathbb{R}^3$ into another vector field on $\mathbb{R}^3$. However it's more natural to interpret it as an operator that turns one-forms into 2-forms.

The exterior derivative has many nice properties.

Proposition 8.31. *(i) For any $\alpha \in \Omega^p(U)$ we have:*

$$d(d\alpha) = 0 \quad \in \Omega^{p+2}(U)$$

(ii) For any $\alpha \in \Omega^p(U)$ and $\beta \in \Omega^q(U)$ we have:

$$d(\alpha \wedge \beta) = d\alpha \wedge \beta + (-1)^p \alpha \wedge d\beta \quad \in \Omega^{p+1}(U) \tag{8.32}$$

(iii) If V is an open set in $\mathbb{R}^k$ and $F : U \to V$ is a smooth function, then for any $\alpha \in \Omega^p(V)$ we have:

$$d(F^\star \alpha) = F^\star d\alpha \quad \in \Omega^p(U)$$

A good way to remember the sign in part (ii) is to pretend that the symbol d behaves a bit like a one-form, so if we want to permute it past the p-form α then we pick up a sign $(-1)^p$. Notice that (8.32) is formally similar to the product rule for derivations, and in fact d is indeed a 'derivation' in a more general sense.

Proof. (i). By linearity it's enough to prove the result for an α which has a single component $\alpha_{\mathbf{i}}$ for some $\mathbf{i}$. Applying the formula (8.29) twice, we get the $(p+2)$-form:

$$d(d\alpha) = \sum_{m=1}^{n} \sum_{j=1}^{n} \frac{\partial^2 \alpha_{\mathbf{i}}}{\partial x_m \partial x_j} \, dx_m \wedge dx_j \wedge dx_{i_1} \wedge \ldots \wedge dx_{i_p}$$

This sum has a (potentially) non-zero term for every pair m, j such that $m \neq j$ and neither m nor j appear in $\mathbf{i}$. However, the double partial derivative $\frac{\partial^2 \alpha_{\mathbf{i}}}{\partial x_m \partial x_j}$ is symmetric in m and j, and the wedge product $dx_m \wedge dx_j$ is antisymmetric in m and j, so these terms cancel in pairs.

(ii). Firstly suppose that α and β are just zero-forms, *i.e.* elements of $C^\infty(\tilde{U})$. Then (8.32) says that

$$d(\alpha\beta) = \beta \, d\alpha + \alpha \, d\beta$$

(since for zero-forms the wedge-product is ordinary point-wise multiplication), and this is true by the product rule for partial differentiation.

Now let $\alpha \in C^\infty(U)$ be a zero-form, and let β be the constant q-form $dx_{i_1} \wedge \ldots \wedge dx_{i_q}$. Then the formula (8.29) says that $d\beta = 0$, and it also says that

$$d(\alpha dx_{i_1} \wedge \ldots \wedge dx_{i_q}) = d\alpha \wedge dx_{i_1} \wedge \ldots \wedge dx_{i_q}$$

so this special case of (8.32) is also true.

Now let α be a p-form with a single component, and β be a q-form with a single component, so

$$\alpha = \alpha_{\mathbf{i}} \, dx_{i_1} \wedge \ldots \wedge dx_{i_p} \qquad \text{and} \qquad \beta = \beta_{\mathbf{j}} \, dx_{j_1} \wedge \ldots \wedge dx_{j_q}$$

with $\alpha_{\mathbf{i}}, \beta_{\mathbf{j}} \in C^\infty(U)$. Then:

$$\begin{aligned}
d(\alpha \wedge \beta) &= d\left(\alpha_{\mathbf{i}} \beta_{\mathbf{j}} \, dx_{i_1} \wedge \ldots \wedge dx_{i_p} \wedge dx_{j_1} \wedge \ldots \wedge dx_{j_p}\right) \\
&= (\beta_{\mathbf{j}} \, d\alpha_{\mathbf{i}} + \alpha_{\mathbf{i}} \, d\beta_{\mathbf{j}}) \wedge dx_{i_1} \wedge \ldots \wedge dx_{i_p} \wedge dx_{j_1} \wedge \ldots \wedge dx_{j_p} \\
&= (d\alpha_{\mathbf{i}} \wedge dx_{i_1} \wedge \ldots \wedge dx_{i_p}) \wedge (\beta_{\mathbf{j}} \, dx_{j_1} \wedge \ldots \wedge dx_{j_q}) \\
&\qquad + (-1)^p (\alpha_{\mathbf{i}} \, dx_{i_1} \wedge \ldots \wedge dx_{i_p}) \wedge (d\beta_{\mathbf{j}} \wedge dx_{j_1} \wedge \ldots \wedge dx_{j_q}) \\
&= d\alpha \wedge \beta + (-1)^p \alpha \wedge d\beta
\end{aligned}$$

So by linearity (8.32) holds for any p-form and any q-form.

(iii). We've already proved the case $p = 0$ in Lemma 8.12, since if α is a zero-form then $F^\star \alpha$ is just $\alpha \circ F$. Now suppose $\alpha = \alpha_{\mathbf{i}}\, dx_{i_1} \wedge ... \wedge dx_{i_p}$ is a p-form with a single component. Then

$$F^\star \alpha = (\alpha_{\mathbf{i}} \circ F)\, (F^\star dx_{i_1}) \wedge ... \wedge (F^\star dx_{i_p})$$

by (8.25). Now for any k, we have $F^\star dx_k = d(x_k \circ F)$ by Lemma 8.12, so $d(F^\star dx_k) = 0$ by part (i) of this proposition. Then repeatedly applying part (ii) of this proposition shows that:

$$d(F^\star \alpha) = d(\alpha_{\mathbf{i}} \circ F) \wedge (F^\star dx_{i_1}) \wedge ... \wedge (F^\star dx_{i_p})$$

Hence

$$\begin{aligned} d(F^\star \alpha) &= (F^\star d\alpha_{\mathbf{i}}) \wedge (F^\star dx_{i_1}) \wedge ... \wedge (F^\star dx_{i_p}) \\ &= F^\star \left(d\alpha_{\mathbf{i}} \wedge dx_{i_1} \wedge ... \wedge dx_{i_p} \right) \\ &= F^\star d\alpha \end{aligned}$$

as required.

□

Now we want to define the exterior derivative on an arbitrary manifold X. Obviously, when we work in co-ordinates it should reduce to the operations that we've just defined. In fact, this requirement is a valid way to *define* the operation d on X.

Lemma 8.33. *Let $\alpha \in \Omega^p(X)$ be a p-form. Then there exists a unique $(p+1)$-form $d\alpha$ on X such that, for any chart (U, f) on X, when we write $d\alpha$ in this chart we get*

$$d\tilde{\alpha} \;\in \Omega^{p+1}(\tilde{U})$$

where $\tilde{U}$ is the codomain of the chart and $\tilde{\alpha}$ is the expression for α in this chart.

Proof. Consider the rule that assigns to any chart (U, f), the $(p+1)$-form $d\tilde{\alpha}$ on $\tilde{U}$. We claim that this rule is actually a $(p+1)$-form on X in the 'physicist's' sense, *i.e.* it obeys the correct transformation law when we change co-ordinates.

So we just need to check that $d\tilde{\alpha}$ transforms correctly. If we have two charts (U_1, f_1) and (U_2, f_2), then in each chart the p-form α turns into two smooth functions $\tilde{\alpha}_1$ and $\tilde{\alpha}_2$ from $\tilde{U}$ to $\wedge^p(\mathbb{R}^n)^\star$. These functions are related by the transformation law

$$\tilde{\alpha}_2 = \phi_{12}^\star \tilde{\alpha}_1$$

where the two charts overlap. But by Proposition 8.31(iii), we have

$$d\tilde{\alpha}_2 = \phi_{12}^\star d\tilde{\alpha}_1$$

and this is the correct transformation law for a $(p+1)$-form on X. □

We have shown that on any manifold X, and any p, we have a exterior derivative:

$$d: \; \Omega^p(X) \to \Omega^{p+1}(X)$$

Futhermore it's immediate that the three properties listed in Proposition 8.31 all hold, since each of them can be checked in co-ordinates.

You might (in fact you should) find the previous proof a bit unsatisfactory, it would be better to find a definition of d that didn't rely on picking charts. It is possible to give such a definition, but it's not very straight-forward.

The main use for the exterior derivative is to define *de Rham cohomology*.

Definition 8.34. Let $\alpha \in \Omega^p(X)$ be a differential form. We say that α is **closed** if $d\alpha = 0$. We say that α is **exact** if there is some $\beta \in \Omega^{p-1}(X)$ such that $d\beta = \alpha$.

Since $d \circ d = 0$ (Proposition 8.31 (i)), any exact form is automatically closed, so we have inclusions:

$$\{\text{exact } p\text{-forms}\} \subset \{\text{closed } p\text{-forms}\} \subset \Omega^p(X)$$

Also both subsets are actually subspaces, since they are the image or kernel of a linear map.

Definition 8.35. For each $p \in [0, n]$, the p-th **de Rham cohomology group** of X is defined to be the quotient vector space:

$$H^p_{dR}(X) = \frac{\{\text{closed } p\text{-forms}\}}{\{\text{exact } p\text{-forms}\}}$$

The word 'group' here is just historical; this is object indeed an abelian group, but it's also a real vector space. De Rham cohomology is a fundamental topological invariant for manifolds.

9 Integration

If we have an open set $U \subset \mathbb{R}^n$, and a smooth function $h \in C^\infty(U)$, we can try to compute the multiple integral:

$$\int_U h\, dx_1 ... dx_n$$

As long as the function h doesn't grow too big as we approach the edges of U, this integral will converge. For example, if we know that h is identically zero outside some closed ball $\overline{B(0,r)} \subset U$ (*i.e.* it is a bump function) then the integral certainly converges.

Now suppose we have a diffeomorphism $F : V \xrightarrow{\sim} U$ where V is some other open set in $\mathbb{R}^n$. This is just a change of co-ordinates, so we can compute the integral in the new co-ordinates if we wish. You should recall the 'change-of-variables' formula:

$$\int_U h\, dx_1 ... dx_n = \int_V (h \circ F)\, |\det DF|\ dy_1 ... dy_n \tag{9.1}$$

The factor $|\det DF|$, the absolute value of the determinant of the Jacobian matrix of F, keeps track of how volume gets distorted when we change variables.

This is strikingly similar to the the way that n-forms behave. If we decide that the symbols after the integral sign $\int_U$ are really the n-form

$$\alpha = h\, dx_1 \wedge ... \wedge dx_n \ \in \Omega^n(U)$$

then we have

$$\tilde{F}^\star \alpha = (h \circ F) \det(DF)\, dy_1 \wedge ... \wedge dy_n$$

which is almost the same as the symbols after $\int_V$. This is very strong evidence that the correct thing to integrate over U is not functions, but n-forms. Indeed, all n-forms on U are of the form $\alpha = h\, dx_1 \wedge ... \wedge dx_n$ for some $h \in C^\infty(U)$, so we can simply define the integral of α to be:

$$\int_U \alpha = \int_U h\, dx_1 ... dx_n$$

The formula (9.1) almost says that $\int_V F^\star \alpha = \int_U \alpha$, *i.e.* the value of the integral doesn't change when you apply a diffeomorphism. The only problem is the occurence of $|\det(DF)|$ instead of $\det(DF)$. However, since F is a diffeomorphism, $\det(DF)$ can never be zero. So if we assume that U is connected, then $\det(DF)$ must either be always positive or always negative, and we have that

$$\int_V F^\star \alpha = \big(\text{sign of } \det(DF)\big) \int_U \alpha \tag{9.2}$$

(if U is not connected then split it into its connected components, and this equation holds for each component).

Now suppose we have an arbitrary n-dimensional manifold X, and an n-form $\alpha \in \Omega^n(X)$. The goal of this section is to define the integral:

$$\int_X \alpha \in \mathbb{R}$$

If we pick a chart (U, f) on X, then we can at least try to define the integral of α over the region U by writing α in co-ordinates and then integrating it over $\tilde{U}$. There are two problems with this definition: (1) the integral might not converge, and (2) the value of the integral might not be co-ordinate independent. We will solve (2) in Section 9.1 by introducing *orientations*.

Then in Section 9.2 we will explain how we define the integral of α over the whole of X. The strategy is that we cover X by an atlas, then use a gadget called a *partition-of-unity* to 'cut up' α into pieces, one for each chart in the atlas, such that the integral of each piece is guaranteed to converge. Finally we get the total integral by adding up the contribution from each piece.

9.1 Orientations

Let X be an n-dimensional manifold, and let $\alpha \in \Omega^n(X)$ be an n-form on X. Let (U_1, f_1) and (U_2, f_2) be two charts on X, and for simplicity let's assume they have the same domain $U = U_1 = U_2$ and that U is connected. Writing α in these two sets of co-ordinates gives n-forms $\tilde{\alpha}_1 \in \Omega^n(\tilde{U}_1)$ and $\tilde{\alpha}_2 \in \Omega^n(\tilde{U}_2)$, and the transformation law for n-forms (8.27) says that

$$\tilde{\alpha}_2 = \phi_{12}^\star \tilde{\alpha}_1$$

where $\phi_{12} : \tilde{U}_2 \xrightarrow{\sim} \tilde{U}_1$ is the transition function. We would like to define the integral of α over U by writing α in either chart, and then integrating it. Unfortunately, even if this integral converges, it is not quite chart-independent. Instead, we have that

$$\int_{\tilde{U}_2} \tilde{\alpha}_2 = \big(\text{sign of } \det(D\phi_{12})\big) \int_{\tilde{U}_1} \tilde{\alpha}_1$$

by (9.2). So we need to find a way to guarantee that $\det(D\phi_{12})$ is positive.

Definition 9.3. A **volume form** on an n-dimensional manifold X is an n-form $\omega \in \Omega^n(X)$ such that ω is not zero at any point. If there exists a volume form on X then we say that X is **orientable**.

Example 9.4. If $U \subset \mathbb{R}^n$ is an open set, then

$$\omega_0 = dx_1 \wedge \ldots \wedge dx_n \in \Omega^n(U)$$

is a volume form, so $\mathbb{R}^n$ is orientable. We'll call ω_0 the *standard volume form* on U. Any n-form on U is of the form $h\omega_0$ for some $h \in C^\infty(U)$, and this is a volume form iff h is never zero.

One feature of the previous example generalizes to any manifold: if we have a volume form ω on X, then any n-form $\alpha \in \Omega^n(X)$ can be written as $\alpha = h\omega$ for some $h \in C^\infty(X)$, because the vector space $\wedge^n T_x^\star X$ is 1-dimensional at each point $x \in X$. If h is never zero then α will be another volume form, and vice-versa.

Example 9.5. Let $X = T^1$. Recall from Example 8.7 that a one-form on T^1 is the same thing as a periodic one-form on $\mathbb{R}$, *i.e.* a one-form $\hat{\alpha}\, dx$ where $\hat{\alpha} \in C^\infty(\mathbb{R})$ is a function such that $\hat{\alpha}(x+1) = \hat{\alpha}(x)$ for all x. If we set $\hat{\alpha}$ to be the constant function $\hat{\alpha} \equiv 1$ then we get a volume form $\omega \in \Omega^1(T^1)$, corresponding to the periodic 1-form $dx \in \Omega^1(\mathbb{R})$. This shows that T^1 is orientable.

The previous example can easily be generalized to show that the torus T^n is orientable for any n. Constructing volume forms on the n-sphere S^n is a little harder, and we'll use the following general result:

Proposition 9.6. *Let X be an orientable manfold, and let*

$$Z = h^{-1}(y) \subset X$$

be a level set of some function $h \in C^\infty(X)$ at a regular value $y \in \mathbb{R}$. Then Z is orientable.

In particular if we set $X = \mathbb{R}^{n+1}$ and $Z = S^n$ then we see that S^n is orientable, for any n.

Proof. Let $\dim X = n$, and fix a point $z \in Z$, so the volume form ω at this point gives us a non-zero element $\omega|_z \in \bigwedge^n T_z^\star X$. We saw in Lemma 5.13 that the tangent space $T_z Z$ is the subspace of $T_z X$ given by the kernel of the linear map $Dh|_z : T_z X \to T_z \mathbb{R} \cong \mathbb{R}$. Let $\mathbf{n} \in T_z X$ be any vector such that $Dh|_z(\mathbf{n}) = 1$. Then we can define an element $\omega'|_z \in \bigwedge^{n-1} T_z^\star Z$ by declaring that

$$\omega'|_z : (v_1, ..., v_{n-1}) \mapsto \omega|_z(v_1, ..., v_{n-1}, \mathbf{n}) \quad \in \mathbb{R}$$

for any vectors $v_1, ..., v_{n-1} \in T_z Z$. This map $\omega'|_z$ is automatically $(n-1)$-linear and antisymmetric, so it is indeed an element of $\bigwedge^{n-1} T_z^\star Z$. Futhermore we claim that it's independent of our choice of $\mathbf{n}$. To see this, recall that $\bigwedge^{n-1} T_z^\star Z$ is only 1-dimensional, so if we pick a basis $e_1, ..., e_{n-1}$ for $T_z Z$ then $\omega'|_z$ is determined by the single real number :

$$\omega'|_z(e_1, ..., e_{n-1}) = \omega|_z(e_1, ..., e_{n-1}, \mathbf{n}) \tag{9.7}$$

This number will not change if we change $\mathbf{n}$ by adding on any linear combination of the e_i's, because $\omega|_z$ is anti-symmetric and linear in each argument. However, any vector in $Dh|_z^{-1}(1)$ must differ from $\mathbf{n}$ by some linear combination of the e_i's, so $\omega'|_z$ is indeed independent of our choice of $\mathbf{n}$.

Also, the number (9.7) cannot be zero, because the vectors $e_1, ..., e_{n-1}, \mathbf{n}$ form a basis of $T_z X$, and we know that $\omega|_z$ is not zero. Therefore $\omega'|_z$ is not the zero element of $\bigwedge^{n-1} T_z^\star Z$.

So for every point $z \in Z$, we have constructed a non-zero element $\omega'|_z \in \bigwedge^{n-1} T_z^\star Z$. If we can show that these elements vary smoothly, then we have found a volume form ω' on Z. So we need to look at this construction in co-ordinates.

We can assume that Z is the level set of $y = 0$, since we can always replace h by $h - y$. Then for any point $z \in Z$, we know that we can find a chart (U, f) containing z such that when we write h in this chart it is just the last co-ordinate $\tilde{h} = x_n$ on $\tilde{U}$. In such a chart, the submanifold Z becomes $f(Z \cap U) = \mathbb{R}^{n-1} \cap \tilde{U}$, and we may choose our vector $\mathbf{n}$ to be the tangent vector $\frac{\partial}{\partial x_n}$ at any point in $f(Z \cap U)$. Now write the volume form ω in these co-ordinates, so it becomes

$$\tilde{\omega} = \tilde{g}\, dx_1 \wedge ... \wedge dx_n$$

for some $\tilde{g} \in C^\infty(\tilde{U})$. It follows that ω', in these co-ordinates, is given by

$$\tilde{\omega}' = \tilde{g}|_{\{x_n = 0\}}\, dx_1 \wedge ... \wedge dx_{n-1}$$

which is indeed a smooth $(n-1)$-form on $f(Z \cap U)$. $\square$

With a little more work, this proposition can be generalised to level sets (at regular values) of smooth functions $h : X \to Y$, where Y is any other orientable manifold. However, it is not true that *any* submanifold of an orientable manifold is orientable.

In practice most manifolds that one cares about are orientable, but it is not too hard to find non-orientable manifolds. We leave the next example as an exercise:

Example 9.8. The Klein bottle K (see problem sheets) is a 2-dimensional manifold which is not orientable.

The manifold $\mathbb{RP}^n$ is orientable iff n is odd.

Now let X be an orientable manifold, and let $\omega \in \Omega^n(X)$ be a volume form. Pick a chart (U, f) on X. In this chart ω becomes a volume form

$$\tilde{\omega} = h\omega_0 \in \Omega^n(\tilde{U})$$

for some $h \in C^\infty(\tilde{U})$, where ω_0 is the standard volume form on $\tilde{U}$. This function h must be non-zero at all points.

Definition 9.9. Let $\omega \in \Omega^n(X)$ be a volume form on a manifold X. Let (U, f) be a chart on X, and let ω_0 be the standard volume form on $\tilde{U}$. We say that (U, f) is **oriented** (with respect to ω) if when we write ω in this chart it becomes

$$\tilde{\omega} = h\omega_0$$

where the function h is always positive.

It's not hard to find oriented charts. If U is connected, then h must be either always positive or always negative. If it's negative, just compose the co-ordinates f with the reflection:

$$F : \mathbb{R}^n \to \mathbb{R}^n$$
$$(x_1, x_2, ..., x_n) \mapsto (-x_1, x_2, ..., x_n)$$

Since $F^\star\omega_0 = -\omega_0$, the chart $(U, F \circ f)$ will be oriented. We could also replace F here with any other diffeomorphism F satisfying $\det DF < 0$ at all points. If U is not connected, then we can split U into its connected components, and perform this trick on each component where h is negative. So for any chart (U, f), we can find an oriented chart which has the same domain U.

Now suppose that $\omega \in \Omega^n(X)$ is a volume form, and that $g \in C^\infty(X)$ is a real-valued function which is positive at all points. Then $g\omega$ is another volume form on X, and asking for a chart to be oriented with respect to $g\omega$ is exactly the same condition as asking for it to be oriented with respect to ω. This leads us to the following definition:

Definition 9.10. An **orientation** on a manifold X is an equivalence class of volume forms on X, where we declare that two volume forms ω_1 and ω_2 are equivalent iff $\omega_2 = g\omega_1$ for some $g \in C^\infty(X)$ which is positive everywhere. If we've fixed an orientation on X we say that X is **oriented**.

Obviously, we can find an orientation for X iff X is orientable. If we fix an orientation $[\omega]$ on X then that determines which charts are oriented, it is not necessary to choose a specific volume form in the equivalence class $[\omega]$.

For our purposes, the reason for introducing oriented manifolds, and oriented charts, is the following easy observation:

Lemma 9.11. *Let X be an oriented manifold. Let (U_1, f_1) and (U_2, f_2) be two oriented charts, and let ϕ_{12} be the transition function between them. Then for any point $\tilde{x} \in f_2(U_1 \cap U_2)$ we have:*

$$\det D\phi_{12}|_{\tilde{x}} > 0$$

Proof. Pick any volume form ω representing the given equivalence class. In our two charts, ω becomes

$$\tilde{\omega}_1 = h_1\omega_0 \in \Omega^n(\tilde{U}_1) \qquad \text{and} \qquad \tilde{\omega}_2 = h_2\omega_0 \in \Omega^n(\tilde{U}_2)$$

where both h_1 and h_2 are positive everywhere since both charts are oriented. For a point $x \in U_1 \cap U_2$ we have

$$h_2|_{f_2(x)} = \det\big(D\phi_{12}|_{f_2(x)}\big)\, h_1|_{f_1(x)}$$

by the transformation law for n-forms (8.27). Hence $\det\big(D\phi_{12}|_{f_2(x)}\big) > 0$. □

This solves part of our problem of defining integration on manifolds. If we stick to oriented manifolds, and oriented charts, then the determinant of the derivative of the transition function will always be positive. Then the integral of an n-form over a chart U (if it converges) will be independent of the choice of co-ordinates.

9.2 Partitions-of-unity and integration

For any (oriented) manifold X, and any n-form α on X, we would like to be able to define the integral

$$\int_X \alpha$$

as some real number. However this is not going to work in general, because integrals do not always converge. For example if we take $X = \mathbb{R}$, and $\alpha \in \Omega^1(\mathbb{R})$ to be the constant one-form dx, then we are trying to evaluate

$$\int_{-\infty}^{\infty} 1\, dx$$

which doesn't converge to a finite answer. So we have to put some restrictions on either X or α.

To start with, we'll let X be any oriented manifold, but we'll only consider a restricted class of n-forms.

Definition 9.12. Let $\alpha \in \Omega^p(X)$ for some p. We'll call α a **bump form** if there exists some chart (U, f) on X, and some compact subset $W \subset U$, such that α is identically zero outside of W.

Warning: this is not standard terminology, but it will be a convenient definition for us.

If $\alpha \in \Omega^n(X)$ is a bump form, then we can define the integral $\int_X \alpha \in \mathbb{R}$ in the following way. By definition, there is some chart (U, f) such that α vanishes outside of a compact subset $W \subset U$. We can also assume that (U, f) is an oriented chart (if not then compose with it with a reflection). Now look at the

n-form $\tilde{\alpha} \in \Omega^n(\tilde{U})$ that we get by writing α in this chart. The subset $f(W) \subset \tilde{U}$ is compact, so it's contained in some closed ball $\overline{B(0,R)} \subset \tilde{U}$, so $\tilde{\alpha}$ is identically zero outside this ball. Therefore we can define

$$\int_X \alpha = \int_{\tilde{U}} \tilde{\alpha} \quad \in \mathbb{R}$$

since this integral converges. Now let's prove that this definition is independent of our choice of chart.

Proposition 9.13. *Let (U_1, f_1) be an oriented chart such that α vanishes outside of some compact subset $W_1 \subset U_1$. Let (U_2, f_2) be another oriented chart such that α vanishes outside of some compact subset $W_2 \subset U_2$. Let $\tilde{\alpha}_1 \in \Omega^n(\tilde{U}_1)$ and $\tilde{\alpha}_2 \in \Omega^n(\tilde{U}_2)$ be the n-forms that we get by writing α in the two charts. Then we have:*

$$\int_{\tilde{U}_1} \tilde{\alpha}_1 = \int_{\tilde{U}_2} \tilde{\alpha}_2$$

Proof. Let $U = U_1 \cap U_2$, and $W = W_1 \cap W_2$. Then W is a compact subset of U, and α vanishes outside of W. Consequently

$$\int_{\tilde{U}_1} \tilde{\alpha}_1 = \int_{f_1(U)} \tilde{\alpha}_1 \qquad \text{and} \qquad \int_{\tilde{U}_2} \tilde{\alpha}_2 = \int_{f_2(U)} \tilde{\alpha}_2$$

since $\tilde{\alpha}_1$ vanishes outside $f_1(U)$ and $\tilde{\alpha}_2$ vanishes outside $f_2(U)$, and all these integrals converge. The transition function is a diffeomorphism

$$\phi_{12} : f_2(U) \xrightarrow{\sim} f_1(U)$$

and $\tilde{\alpha}_2 = \phi_{12}^* \tilde{\alpha}_1$, so the formula (9.2) shows that

$$\int_{f_1(U)} \tilde{\alpha}_1 = \int_{f_2(U)} \tilde{\alpha}_2$$

since both charts are oriented. □

So for any bump form $\alpha \in \Omega^n(X)$ we have a well-defined integral $\int_X \alpha \in \mathbb{R}$. We can calculate this integral using any chart that contains the locus where α is non-zero.

We now want to consider integrating arbitrary n-forms. This means that we have to put some restriction on X, and the correct restriction is to insist that X itself is compact. On a compact manifold, there is a way to 'chop-up' an arbitrary n-form into a finite number of bump forms. Then we can then integrate each piece, and add the answers together. The 'chopping-up' step is done with the following gadget:

Definition 9.14. Let X be a manifold. A **partition-of-unity** on X is a set of functions $\varphi_\bullet = \{\varphi_i,\ i \in I\} \subset C^\infty(X)$ (indexed by some set I) with the following properties:

(i) For each $i \in I$ the function $\varphi_i \in \Omega^0(X)$ is a bump form, so there exists a chart (U_i, f_i) and a compact subset $W_i \subset U_i$ such that φ_i vanishes outside of W_i.

(ii) At any point $x \in X$, only finitely-many of the φ_i have $\varphi_i(x) \neq 0$.

(iii) The sum

$$\sum_{i \in I} \varphi_i$$

is the constant function with value $1 \in \mathbb{R}$.

Note that property (ii) implies that the sum in (iii) is finite at every point, so there are no convergence issues.

Suppose that we choose a specific chart (U_i, f_i) for each φ_i, satisfying the condition (i). Then this set of charts gives an atlas $\mathcal{A} = \{(U_i, f_i),\ i \in I\}$ for X, because at any point $x \in X$ at least one of the φ_i's must be non-zero, so the corresponding U_i contains x. We say that $\varphi_\bullet$ is *subordinate* to the atlas $\mathcal{A}$. Sometimes we want to specify the atlas $\mathcal{A}$ in advance, and then construct a partition-of-unity subordinate to $\mathcal{A}$.

It's possible to prove that a partition-of-unity exists on any manifold X, using the technical assumption that the toplogical space underlying X is second-countable. In fact, given any atlas $\mathcal{A}$, there exists a partition-of-unity subordinate to $\mathcal{A}$. We are not going to prove these statements, but we will prove the following easy special case:

Proposition 9.15. *If X is compact then there exists a partition-of-unity*

$$\varphi_\bullet = \{\varphi_i,\ i \in I\} \subset C^\infty(X)$$

on X, where the set I is finite.

It's fairly easy to show that a finite partition-of-unity can only exist if X is compact, so the proposition is really 'if-and-only-if'.

Proof. For any point $x \in X$, we can find a bump function ψ_x which is constantly equal to 1 on some open neighbourhood V_x of x, never negative, and vanishes outside of some compact set which is contained within a chart. Choose this data of ψ_x and V_x and for each point x. The open sets $\{V_x,\ x \in X\}$ form an open cover of X, so since X is compact there is some finite subcover $\{V_{x_1}, ..., V_{x_r}\}$. Let $\psi_{x_r}, ..., \psi_{x_r}$ be the corresponding set of bump functions. Then the sum $\psi_{x_1} + ... + \psi_{x_r}$ is strictly positive at all points of X, since no term is negative and at all points at least one term is equal to 1. Hence we can define

$$\varphi_i = \frac{\psi_{x_i}}{\psi_{x_1} + ... + \psi_{x_r}} \ \in C^\infty(X)$$

and then $\varphi_1, ..., \varphi_r$ is a partition-of-unity. □

Partitions-of-unity can be used for many things, one of which is defining integration. Suppose X is oriented and compact, and we have found a finite partition-of-unity $\varphi_\bullet = \{\varphi_1, ..., \varphi_r\}$. Then if $\alpha \in \Omega^n(X)$ is any n-form, we have that:

$$\alpha = \varphi_1\alpha + ... + \varphi_r\alpha$$

Each $\varphi_i\alpha$ is a bump form, so we have a well-defined integral $\int_X \varphi_i\alpha \in \mathbb{R}$, as we saw before. So we can define the 'integral of α over X, using $\varphi_\bullet$', as:

$$\int_X^{\varphi_\bullet} \alpha \ = \ \sum_{i=1}^r \left(\int_X \varphi_i \alpha \right) \quad \in \mathbb{R}$$

We just need to check that this definition is independent of which partition-of-unity we chose.

Proposition 9.16. *Let X be a compact oriented manifold, and let $\varphi_\bullet = \{\varphi_1, ..., \varphi_r\}$ and $\widehat{\varphi}_\bullet = \{\widehat{\varphi}_1, ..., \widehat{\varphi}_s\}$ be two finite partitions-of-unity on X. Then for any n-form $\alpha \in \Omega^n(X)$, we have that:*

$$\int_X^{\varphi_\bullet} \alpha \;=\; \int_X^{\widehat{\varphi}_\bullet} \alpha$$

Proof. First we claim that if β is a bump-form, and $\varphi_\bullet$ is a finite partition-of-unity, then:

$$\int_X^{\varphi_\bullet} \beta = \int_X \beta$$

The proposition follows immediately from this claim, because if α is any n-form, and $\varphi_\bullet, \widehat{\varphi}_\bullet$ are two finite partitions-of-unity, then:

$$\int_X^{\varphi_\bullet} \alpha = \sum_i \left(\int_X \varphi_i \alpha \right) = \sum_i \left(\int_X^{\widehat{\varphi}_\bullet} \varphi_i \alpha \right) = \sum_{i,j} \left(\int_X \widehat{\varphi}_j \varphi_i \alpha \right) = \int_X^{\widehat{\varphi}_\bullet} \alpha$$

So, suppose β is a bump-form, and let (U, f) be an oriented chart such that β vanishes outside of a compact subset $W \subset U$. Now let $\varphi_\bullet$ be any a partition-of-unity. Each bump-form $\varphi_i \beta$ also vanishes outside of W, so we may evaluate each $\int_X \varphi_i \beta$ using the chart (U, f). Writing everything in these co-ordinates, we have that

$$\tilde{\beta}\, dx_1 \wedge ... \wedge dx_n = \sum_i \tilde{\varphi}_i \tilde{\beta}\, dx_1 \wedge ... \wedge dx_n$$

and since integrating over $\tilde{U}$ is a linear operation, we see that $\int_X \beta = \sum_i \int_X \varphi_i \beta$. □

So given a compact, oriented manifold X, and an n-form $\alpha \in \Omega^n(X)$, we have a well-defined definition of the integral $\int_X \alpha$, namely pick a finite partition-of-unity $\varphi_\bullet$, and define:

$$\int_X \alpha \;=\; \int_X^{\varphi_\bullet} \alpha$$

This doesn't depend on our choice of partition-of-unity. It defines an operation

$$\int_X :\; \Omega^n(X) \to \mathbb{R}$$
$$\alpha \mapsto \int_X \alpha$$

and it is easy to check that this is linear in α.

Example 9.17. Let $X = T^1$, and let's use our usual atlas $\{(U_1, f_1), (U_2, f_2)\}$ (as in Example 2.11). We observed in Example 9.5 that there is a volume form $\omega \in \Omega^1(T^1)$ which corresponds to the periodic one-form $dx \in \Omega^1(\mathbb{R})$, so if we write ω in either chart we just get dx. Any other one-form α on T^1 can be written as $\alpha = h\omega$ for some $h \in C^\infty(T^1)$.

Let's fix the orientation $[\omega]$ on T^1, this means that both of our charts are oriented. Now pick a one-form $h\omega \in \Omega^1(T^1)$, and let's evaluate the integral:

$$\int_{T^1} h\omega$$

First we need a partition-of-unity. Let $\widehat{\varphi}_1 \in C^\infty\big((0,1)\big)$ be some function which is never negative, constantly equal to 1 in some interval containing $\frac{1}{2}$, and vanishes outside some larger interval. Now extend $\widehat{\varphi}_1$ to a periodic function on the whole of $\mathbb{R}$ by setting $\widehat{\varphi}_1(0) = 0$ and insisting that $\widehat{\varphi}_1(x) = \widehat{\varphi}_1(x+1)$ for all $x \in \mathbb{R}$. This defines a function $\varphi_1 \in C^\infty(T^1)$, which vanishes outside a compact subset of U_1. If we let $\widehat{\varphi}_2 = 1 - \widehat{\varphi}_1 \in C^\infty(\mathbb{R})$, then this defines a function $\varphi_2 \in C^\infty(T^1)$ which vanishes outside a compact subset of U_2, and the pair (φ_1, φ_2) is a partition-of-unity on T^1. Then our integral is the sum of two terms:

$$\int_{T^1} h\omega = \int_{T_1} \varphi_1 h\omega + \int_{T^1} \varphi_2 h\omega$$

Lift the function $h \in C^\infty(T^1)$ to a periodic function $\widehat{h} \in C^\infty(\mathbb{R})$, then the expression for h in either chart is just given by restricting $\widehat{h}$ to the corresponding interval. Now let's evaluate the first term in our integral, which we can do in the chart (U_1, f_1). It gives the answer:

$$\int_{T^1} \varphi_1 h\omega = \int_0^1 \widehat{\varphi}_1(x)\widehat{h}(x)\, dx$$

The second term can be evaluated in (U_2, f_2), and gives

$$\int_{T^1} \varphi_2 h\omega = \int_{-\frac{1}{2}}^{\frac{1}{2}} \widehat{\varphi}_2(x)\widehat{h}(x)\, dx = \int_0^1 \widehat{\varphi}_2(x)\widehat{h}(x)\, dx$$

since $\widehat{h}$ and $\widehat{\varphi}_2$ are periodic. So adding the two terms together gives:

$$\int_{T^1} h\omega = \int_0^1 \big(\widehat{\varphi}_1(x) + \widehat{\varphi}_2(x)\big)\widehat{h}(x)\, dx = \int_0^1 \widehat{h}(x)\, dx$$

(this is probably what you would have guessed initially, but this shows that our complicated definition has reduced to a sensible answer in this case). In particular, we have that:

$$\int_{T^1} \omega = 1$$

There's an easy general observation we can make here: if X is any compact manifold and $\omega \in \Omega^n(X)$ is a volume form, then if we use the orientation $[\omega]$ on X we must have

$$\int_X \omega > 0$$

since the integral will be a sum of strictly positive terms. This means that integration over X defines a surjective linear map

$$\int_X : \Omega^n(X) \to \mathbb{R}$$

since if multiply ω by a scalar $\lambda \in \mathbb{R}$ then we can arrange $\int_X \lambda\omega$ to take any value we wish.

9.3 Stokes' Theorem

Suppose we have a function $h \in C^\infty(\mathbb{R})$ that vanishes outside of some closed interval $[-r, r] \subset \mathbb{R}$. Then the one-form

$$dh = \frac{\partial h}{\partial x}\, dx$$

also vanishes outside of this interval, so the integral $\int_\mathbb{R} dh$ certainly converges, and in fact $\int_\mathbb{R} dh = \int_{-r}^{r} dh$. So by the fundamental theorem of calculus, we have

$$\int_\mathbb{R} dh = \int_{-r}^{r} \frac{\partial h}{\partial x}\, dx = h(r) - h(-r) = 0$$

since $h(r) = h(-r) = 0$.

Let's generalise this observation to higher dimensions. If we work on $\mathbb{R}^n$, then the objects that we can integrate are n-forms, so we must replace h by an $(n-1)$-form. For example, let's work on $\mathbb{R}^3$, and consider a 2-form

$$\alpha = \alpha_1\, dy \wedge dz$$

which only has one non-zero component. Then:

$$d\alpha = \frac{\partial \alpha_1}{\partial x}\, dx \wedge dy \wedge dz$$

Suppose that α vanishes outside some cube, *i.e.* the function $\alpha_1 \in C^\infty(\mathbb{R}^3)$ is zero unless $x, y, z \in [-r, r]$. Then we have that

$$\begin{aligned}
\int_{\mathbb{R}^3} d\alpha &= \int_{-r}^{r} \left(\int_{-r}^{r} \left(\int_{-r}^{r} \frac{\partial \alpha_1}{\partial x}\, dx \right) dy \right) dz \\
&= \int_{-r}^{r} \left(\int_{-r}^{r} \Big(\alpha_1(r, y, z) - \alpha_1(-r, y, z) \Big) dy \right) dz \\
&= 0
\end{aligned}$$

since $\alpha_1(r, y, z) = \alpha_1(-r, y, z) = 0$ for any values of y and z.

If α has more than one component then it's still true that $\int_{\mathbb{R}^3} d\alpha = 0$, since we can evaluate each component individually (both d and $\int$ are linear) and each piece will be zero. We can also perform this argument in exactly the same way in higher dimensions:

Lemma 9.18. *Let $U \subset \mathbb{R}^n$ be an open set, and let $\alpha \in \Omega^{n-1}(U)$ be an $(n-1)$-form that vanishes outside some compact subset $W \subset U$. Then:*

$$\int_U d\alpha = 0$$

Proof. By linearity it's enough to consider the case when α has only one non-vanishing component, and without loss of generality we can assume that

$$\alpha = \alpha_1\, dx_2 \wedge \ldots \wedge dx_n$$

for some $\alpha_1 \in C^\infty(U)$ (which vanishes outside W). We can extend α to a smooth $(n-1)$-form on the whole of $\mathbb{R}^n$ by setting it to be zero on $\mathbb{R}^n \setminus U$, and if we pick a large enough r then W will be contained in the n-dimensional cube $[-r,r]^{\times n}$. Then we have:

$$\begin{aligned}\int_U d\alpha &= \int_{[-r,r]^{\times n}} \frac{\partial \alpha_1}{\partial x_1}\, dx_1 dx_2 ... dx_n \\ &= \int_{[-r,r]^{\times (n-1)}} \Big(\alpha_1|_{\{x_1=r\}} - \alpha_1|_{\{x_1=-r\}} \Big)\, dx_2 .... dx_n \\ &= 0\end{aligned}$$

since $\alpha \equiv 0$ on the faces $\{x_1 = r\}$ and $\{x_1 = -r\}$. $\square$

Futhermore, this statement generalizes very easily to more interesting manifolds.

Theorem 9.19 (Stokes' Theorem)**.** *Let X be a compact oriented manifold of dimension n. Then for any $\alpha \in \Omega^{n-1}(X)$ we have:*

$$\int_X d\alpha = 0$$

In fact this is not the full version of Stokes' Theorem, there is a more interesting version that uses the full power of the fundamental theorem of calculus. The stronger version requires the concept of a *manifold-with-boundary*, and is explained in Appendix F.

Proof. Choose a finite partition of unity $\varphi_\bullet$ on X, so $\alpha = (\varphi_1 + ... + \varphi_r)\alpha$ and $d\alpha = d(\varphi_1\alpha) + ... + d(\varphi_r\alpha)$. For each i we know that $\varphi_i\alpha$ is a bump form vanishing outside some compact set W_i, hence $d(\varphi_i\alpha)$ is also a bump form, since it also vanishes outside W_i. Also, we can see that

$$\int_X d(\varphi_i\alpha) = 0$$

by passing to some oriented chart containing W_i and applying Lemma 9.18. Since integrating over X is linear, it follows that $\int_X d\alpha = 0$. $\square$

Recall that the n-th de Rham cohomology group of X is defined to be the quotient vector space

$$H^n_{dR}(X) = \Omega^n(X) \;/\; \{\text{exact } (n-1)\text{-forms}\}$$

(since all n-forms are automatically closed). Stokes' Theorem says precisely that integration gives a well-defined map:

$$\int_X : H^n_{dR}(X) \to \mathbb{R}$$

Example 9.20. Let $X = T^1$, and let $\alpha \in \Omega^1(T^1)$ be the one-form corresponding to a periodic one-form $\hat{\alpha}\, dx \in \Omega^1(\mathbb{R})$. We saw in Example 9.17 that (after fixing the orientation $[\omega]$ on T^1) the integral of α is given by:

$$\int_{T^1} \alpha = \int_0^1 \widehat{\alpha}(x)\, dx$$

If $\alpha = dh$ for some $h \in C^\infty(T^1)$, *i.e.* if α is an exact one-form, then Stokes' Theorem says that $\int_{T^1} \alpha = 0$. This is precisely the observation we made in Example 8.7, it says that if $\hat{\alpha} = \frac{d\hat{h}}{dx}$ for a periodic function $\hat{h} \in C^\infty(\mathbb{R})$, then:

$$\int_0^1 \hat{\alpha}\, dx = \hat{h}(1) - \hat{h}(0) = 0$$

For this example, it's easy to show that the converse to Stokes' Theorem is also true. If $\int_0^1 \hat{\alpha}(x)\, dx = 0$, then the function

$$\begin{aligned} \hat{h} : \mathbb{R} &\to \mathbb{R} \\ x &\mapsto \int_0^x \hat{\alpha}(y)\, dy \end{aligned}$$

will satisfy $\hat{h}(x) = \hat{h}(x+1)$ for all $x \in \mathbb{R}$. Hence it defines a corresponding function $h \in C^\infty(T^1)$, and $dh = \alpha$. So on the circle T^1, we have shown that the integral of a one-form is zero if and only if the one-form is exact. In other words, integration gives an isomorphism:

$$\int_{T^1} : \ H^1_{dR}(T^1) \xrightarrow{\sim} \mathbb{R}$$

The observation from the previous example is actually a general phenomenon. If X is a *connected* compact oriented n-dimensional manifold, and $\alpha \in \Omega^n(X)$ is an n-form, then

$$\int_X \alpha = 0$$

if-and-only-if there exists some $\beta \in \Omega^{n-1}(X)$ such that $\alpha = d\beta$. Hence integration gives an isomorphism:

$$\int_X : \ H^n_{dR}(X) \xrightarrow{\sim} \mathbb{R}$$

This result is quite hard to prove, it's a first glimse of an extremely deep and important fact about manifolds called *Poincaré duality*.

A Topological spaces

In this section we provide a very brief revision of the basics of topological spaces.

If X is a set we let $\mathcal{P}X$ denote the power set of X, i.e. the set of all subsets of X.

Definition A.1. Let X be a set. A **topology** on X is a collection

$$\mathcal{T} \subset \mathcal{P}X$$

of subsets of X, satisfying the list of axioms below. We refer to elements of $\mathcal{T}$ as *open sets*. The axioms are:

(i) The empty subset ϕ is open, and the whole of X is open.

(ii) The intersection of two open sets is open.

(iii) Given any collection of open sets, their union is also open.

Axiom (ii) implies that the intersection of any finite collection of open sets is open. Axiom (iii) applies to *any* collection of open sets, including infinite ones. If we have chosen a topology on X, then we call X a **topological space**.

A subset V of a topological space X is called *closed* iff its complement $V^c = X \setminus V$ is open. Note that most subsets of X are neither open nor closed.

Example A.2. Let $X = \mathbb{R}^n$, equipped with the usual (Euclidean) norm. For a point $x \in \mathbb{R}^n$, and a real number $r \in \mathbb{R}_{\geq 0}$, the open ball around x of radius r is the set:

$$B(x, r) = \{y \in \mathbb{R}^n;\ |y - x| < r\}$$

We declare that a subset $U \subset \mathbb{R}^n$ is open iff for any point $x \in U$ there exists some $\epsilon > 0$ such that:

$$B(x, \epsilon) \subset U$$

Equivalently, we can say that a subset $U \subset X$ is open iff U can be written as a union of some collection of open balls. It is easy to prove that this defines a topology on $\mathbb{R}^n$.

Definition A.3. Let X and Y be topological spaces, and let f be a function:

$$f : X \to Y$$

We say that f is **continuous** iff whenever $U \subset Y$ is an open set then its pre-image

$$f^{-1}(U) \subset X$$

is also open. Equivalently, we can require that the pre-image of every closed set is closed.

Its easy to show that the composition of two continuous functions is continuous.

Definition A.4. If $f : X \to Y$ is a continuous function between two topological spaces then we say that f is a **homeomorphism** iff f is a bijection and the inverse function

$$f^{-1} : Y \to X$$

is also continuous. If there exists a homeomorphism between X and Y then we say that X and Y are **homeomorphic**.

Definition A.5. Let X be a topological space, and let $Z \subset X$ be any subset. We define the **subspace topology** on Z by declaring that a subset $U \subset Z$ is open iff there exists some open set $\tilde{U} \subset X$ such that:

$$U = Z \cap \tilde{U}$$

It's easy to prove that this really does define a topology on Z, and that the inclusion map $Z \hookrightarrow X$ is continuous. It follows that if $f : X \to Y$ is continuous then the restriction $f|_Z : Z \to Y$ is also continuous. It's also easy to prove that a function $g : Y \to Z$ is continuous iff g is continuous when viewed as a function $g : Y \to X$.

Definition A.6. Let X be a topological space, let Y be a set, and let

$$q : X \to Y$$

be a surjective function. We define the **quotient topology** on Y by declaring that $U \subset Y$ is open iff $q^{-1}(U)$ is open in X.

It's easy to check that this really is a topology on Y, and that it makes q continuous.

Let X and Y be two topological spaces. We can put a topology on their cross-product

$$X \times Y = \{(x, y);\ x \in X, y \in Y\}$$

by declaring that if U_1 is an open set in X and U_2 is an open set in Y then

$$U_1 \times U_2 \subset X \times Y$$

is an open set, and further declaring that any union of sets of this form is also an open set. It's easy to check that this defines a topology, and that the projection map from $X \times Y$ to either X or Y is continuous.

We can also put a topology on the disjoint union

$$X \sqcup Y$$

by declaring that a subset $U \subset X \sqcup Y$ is open iff $U \cap X$ is open in X and $U \cap Y$ is open in Y (again it's easy to check that this is a topology). This means that both X and Y are subspaces of $X \sqcup Y$.

Definition A.7. A topological space X is called **compact** iff, whenever we have a collection of open sets $\{U_i,\ i \in I\}$ (indexed by some set I) such that

$$\bigcup_{i \in I} U_i = X$$

then it is possible to find a finite subset $J \subset I$ such that we still have:

$$\bigcup_{j \in J} U_j = X$$

A collection $\{U_i,\ i \in I\}$ like this is called an *open cover* of X, and the sub-collection $\{U_j,\ j \in J\}$ is called a *finite sub-cover*. A subset $Z \subset X$ is called *compact* iff Z is compact in the subspace topology.

Example A.8. If $X = \mathbb{R}^n$ (with the usual topology), then a subset $Z \subset \mathbb{R}^n$ is compact iff Z is both closed and *bounded*, *i.e.*

$$Z \subset B(0, R)$$

for some large-enough R. We won't give a proof of this fact, to find one consult any first course on topological spaces.

Definition A.9. A topological space X is called **Hausdorff** if for any two distinct points $x, y \in X$ we can find open sets U and V with $x \in U$ and $y \in V$ and $U \cap V = \phi$.

So x and y can be 'housed-off' from each other by these open neighbourhoods.

Example A.10. The space $\mathbb{R}^n$ is Hausdorff. Take x and y distinct points in $\mathbb{R}^n$, and let $d = |x - y|$. Then the open balls $B(x, d/2)$ and $B(y, d/2)$ don't intersect.

It's easy to show that any subspace of a Hausdorff space is also Hausdorff.

Example A.11. Here is a rather strange example of a space that would be 1-dimensional topological manifold, except that it is not Hausdorff.

Take the disjoint union $\mathbb{R} \sqcup \mathbb{R}$ of two copies of $\mathbb{R}$, and for any $x \in \mathbb{R}$ let's write x_1 or x_2 for the corresponding points in either component. Let X be the quotient

$$X = (\mathbb{R} \sqcup \mathbb{R}) \;/\; (x_1 \sim x_2 \text{ for } x \neq 0)$$

(with the quotient topology). Then X looks a lot like $\mathbb{R}$, but the origin has been replaced with two points 0_1 and 0_2. If we let U_1 and U_2 be the images in X of the two copies of $\mathbb{R}$, then it's easy to show that they are the domains of two co-ordinate charts, so this is an atlas (it's even a smooth atlas).

However, X is not Hausdorff. If U is any open set containing the first 'origin' 0_1, and V is any open set containing the second 'origin' 0_2, then U and V must have a non-empty intersection.

Definition A.12. A topological space X is called **second-countable** if there is a countable base for the topology $\mathcal{T}$ on X, *i.e.* there is some countable collection of open sets $\mathcal{B} \subset \mathcal{T}$ such that any set in $\mathcal{T}$ can be written as a union of sets from $\mathcal{B}$.

Example A.13. The space $\mathbb{R}^n$ is second-countable. The most obvious base for the topology on $\mathbb{R}^n$ is the set of all open balls (this is how we defined the topology!), and this is an uncountable set. However, there is a countable base, given by the set

$$\mathcal{B} = \big\{ B\big((x_1, ..., x_n), r\big),\ x_1, ..., x_n, r \in \mathbb{Q} \big\}$$

of 'rational balls'. To see that $\mathcal{B}$ is a base, just observe that any open ball $B(x, r)$ can be written as the union of all rational balls that are contained in $B(x, r)$.

It's easy to show that any subspace of a second-countable space is also second-countable.

Example A.14. Take a copy of $\mathbb{R}^n$ for every real number $r \in \mathbb{R}$, and let X be the disjoint union:

$$X = \bigcup_{r \in \mathbb{R}} \mathbb{R}^n$$

Then X would be a n-dimensional topological manifold (or smooth manifold), except that it is not second-countable.

B Dual vector spaces

In this section we provide a brief revision of dual vector spaces and dual linear maps.

Let V be any vector space, of dimension n. The *dual vector space* to V is the space

$$V^\star = \mathrm{Hom}(V, \mathbb{R})$$

of all linear maps from V to $\mathbb{R}$. If we pick a basis $\{e_1, ..., e_n\}$ for V, then $V^\star$ has a corresponding *dual basis* $\{\epsilon_1, ..., \epsilon_n\}$, where $\epsilon_i \in V^\star$ is the linear map defined by:

$$\begin{aligned} \epsilon_i : V &\to \mathbb{R} \\ e_j &\mapsto \begin{cases} 1, & j = i \\ 0, & i \neq j \end{cases} \end{aligned}$$

In particular V^* also has dimension n.

If $V = \mathbb{R}^n$ then we can identify $V^\star$ with $\mathbb{R}^n$; if we think of V as column vectors then elements of $V^\star$ are row vectors. Under this identification the standard basis becomes its own dual basis, and the operation of evaluating a map in $(\mathbb{R}^n)^\star$ on a vector in $\mathbb{R}^n$ becomes the dot product.

If V is not $\mathbb{R}^n$ then there is no canonical isomorphism between V and $V^\star$ - they are isomorphic, but to get an isomorphism we have to choose a basis for V and then identify both V and $V^\star$ with $\mathbb{R}^n$. However, there is a canonical isomorphism between V and $(V^\star)^\star$. It takes a vector $v \in V$ to the linear map:

$$\begin{aligned} ev_v : V^\star &\to \mathbb{R} \\ u &\mapsto u(v) \end{aligned}$$

So V is the dual space to $V^\star$.

Now let W be a second vector space (of dimension m) and let

$$F : V \to W$$

be a linear map. There is a corresponding *dual linear map*

$$F^\star : W^\star \to V^\star$$

which sends a vector $u \in W^\star$ to a vector $F^\star(u) \in V^\star$ defined by:

$$\begin{aligned} F^\star(u) : V &\to \mathbb{R} \\ v &\mapsto u(F(v)) \end{aligned}$$

If we compose two linear maps F and G then the dual of the composed map is:

$$(G \circ F)^\star = F^\star \circ G^\star$$

It follows easily from this that F is an isomorphism if and only if $F^\star$ is an isomorphism. If we pick a basis for both V and W then the linear map F can be expressed as an m-by-n matrix:

$$\tilde{F} : \mathbb{R}^n \to \mathbb{R}^m$$

The dual map $F^\star$ can be an expressed as an n-by-m matrix, using the corresponding dual bases of $V^\star$ and $W^\star$, and it's easy to calculate that it becomes the transpose matrix:

$$\tilde{F}^\top : \mathbb{R}^m \to R^n$$

This is consistent with the fact that $(MN)^\top = N^\top M^\top$.

C Bump functions and the Hausdorff condition

In this short section we show why the Hausdorff condition is necessary for constructing bump functions. We start with a basic fact about Hausdorff spaces:

Lemma C.1. *If X is Hausdorff, and $Z \subset X$ is compact, then Z is closed in X.*

Proof. Fix a point $y \in X \setminus Z$. For any point $z \in Z$ we can find an open set U_z containing z and an open set V_z containing y such that $U_z \cap V_z = \phi$. The union of all the U_z's is an open cover of Z, so it contains a finite subcover $\{U_{z_1}, ..., U_{z_t}\}$. The intersection $V_{z_1} \cap ... \cap V_{z_t}$ of the corresponding open neighbourhoods of y is an open neighbourhood of y, and it does not intersect any U_{z_i}, so it is contained in $X \setminus Z$. Therefore $X \setminus Z$ is open. □

Now we recall the idea of bump functions, discussed in Section 7.1. Fix r and r' with $0 < r < r'$.We observed that it is possible to find a smooth function $\psi \in C^\infty(\mathbb{R}^n)$ such that ψ is constantly equal to 1 inside the ball $B(0, r)$ and constantly equal to 0 outside the larger ball $B(0, r')$.

Now let X be a manifold, and (U, f) a chart on X, such that the codomain $\tilde{U}$ contains the closed ball $\overline{B(0, r')}$. We extend ψ to a function on the whole of X by defining:

$$\widehat{\psi}(y) = \begin{cases} (\psi \circ f)(y), & \text{for } y \in U \\ 0, & \text{for } y \notin U \end{cases}$$

Lemma C.2. *This function $\widehat{\psi}$ is smooth.*

Proof. The closed ball $\overline{B(0, r')}$ is a compact subset of $\tilde{U}$, so $W = f^{-1}(\overline{B(0, r')})$ is a compact subset of X, contained in U. Since X is Hausdorff, Lemma C.1 says that W is closed in X. Then $\widehat{\psi}$ is smooth inside the open set U, and it is certainly smooth inside the open set $X \setminus W$ since it's constant in this locus. Therefore $\widehat{\psi}$ is smooth. □

To understand why the Hausdorff condition was necessary here, consider the following:

Example C.3. Let X be the 'line with two origins' from Example A.11. Let ψ be a bump function on $\mathbb{R}$ which is constantly equal to 1 in some open inteval around 0, and constantly equal to 0 outside some larger open interval. View this is a bump function in the chart U_1, and extend it to a function $\widehat{\psi}$ on the whole of X as we did above. Notice that $X \setminus U_1$ is a single point, the other 'origin' 0_2. If we restrict $\widehat{\psi}$ to the chart $U_2 \cong \mathbb{R}$ we get a function which is equal to 0 at the origin, but constantly equal to 1 inside the set $(-r, 0) \cup (0, r)$ for some $r > 0$. So $\widehat{\psi}$ is not even continuous.

D Derivations at a point

In this section we prove Proposition 7.12, which says that, given a point $x \in X$ in a manifold, a linear map

$$\mathfrak{d} : C^\infty(X) \to \mathbb{R}$$

is a derivation at x if and only if $\mathfrak{d}$ vanishes on the subspace $R_x(X)$ of functions which have rank zero at x.

As usual we start with the easy case when X is an open subset in $\mathbb{R}^n$. Fix a point $x \in X$.

Lemma D.1. *If $\mathfrak{d}$ is a derivation at x and $h \in C^\infty(X)$ has rank zero at x then $\mathfrak{d}(h) = 0$.*

Proof. Recall that $\mathfrak{d} \in \mathrm{Der}_x(X)$ obeys the product rule:

$$\mathfrak{d}(h_1 h_2) = h_1(x)\mathfrak{d}(h_2) + h_2(x)\mathfrak{d}(h_1) \tag{D.2}$$

Firstly we show that $\mathfrak{d}$ vanishes on constant functions. Let $1 \in C^\infty(X)$ denote the constant function with the value $1 \in \mathbb{R}$. Let $h \in C^\infty(X)$ be any function such that $h(x) \neq 0$, then the product rule implies that

$$\mathfrak{d}(h) = \mathfrak{d}(h.1) = h(x)\mathfrak{d}(1) + \mathfrak{d}(h)$$

and hence $\mathfrak{d}(1) = 0$. Since $\mathfrak{d}$ is linear it must send any constant function to zero.

Now let $h \in C^\infty(X)$ be any function. Let $x_1, .., x_n$ be the standard co-ordinate functions on X, and let $(a_1, ..., a_n) \in \mathbb{R}^n$ be the co-ordinates of our fixed point x. Let the Jacobian of h at x be $Dh|_x = (v_1, ..., v_n)$. Then Taylor's theorem says that we can write

$$h = h(x) + \sum_{i=1}^{n} (x_i - a_i)\big(v_i + H_i\big)$$

where $H_1, ..., H_n \in C^\infty(X)$ are functions such that $H_i(x) = 0$ for each i (this is the 'second-order' version of Taylor's theorem, the full theorem says that we can do something similar with higher-order expansions). In particular if h has rank zero at x then each v_i is zero, and:

$$h = h(x) + \sum_{i=1}^{n} (x_i - a_i) H_i$$

In this expression $h(x)$ is a constant, which we could view as a constant function in $C^\infty(X)$, and for each i the expression $(x_i - a_i)$ is also a function in $C^\infty(X)$ and it obeys $(x_i - a_i)|_x = 0$. Then the product rule, and the fact that $\mathfrak{d}$ vanishes on constant functions, implies that $\mathfrak{d}(h) = 0$. □

Now we want to prove a similar result on a general manifold X. Fix a point $x \in X$. Obviously we'd like to take a chart (U, f) around x, and reduce the case of open subsets in $\mathbb{R}^n$. However, there is a subtlety here.

Consider the linear map:

$$C^\infty(X) \to C^\infty(U)$$
$$h \mapsto h|_U$$

This respects multiplication of functions *i.e.* it's a map of rings (or $\mathbb{R}$-algebras), and it also respects the map 'evaluate at the point x'. So if we have a linear map $\mathfrak{d}' : C^\infty(U) \to \mathbb{R}$ which is a derivation at x then we can compose it with the restriction map, and we'll get an operator in $\mathrm{Der}_x(X)$. This provides us with a map:

$$\mathrm{Der}_x(U) \to \mathrm{Der}_x(X)$$

We want to show this map is an isomorphism, but it's not immediately obvious how to invert it, because not every function in $C^\infty(U)$ is the restriction of a function in $C^\infty(X)$. Here is the first step:

Lemma D.3. *Let $\mathfrak{d}$ be a derivation at x. If $h \in C^\infty(X)$ is identically zero on some open neighbourhood of x then $\mathfrak{d}(h) = 0$.*

Proof. Suppose that h is identically zero on a neighbourhood U of x. Let $\psi \in C^\infty(X)$ be a bump function such that we have neighbourhoods

$$x \in V \subset W \subset U$$

with $\psi|_V \equiv 1$ and $\psi|_{X\setminus W} \equiv 0$. Then $(1-\psi)h = h$, so $\mathfrak{d}(h) = \mathfrak{d}((1-\psi)h) = 0$ by the product rule. □

So if $\mathfrak{d} \in \mathrm{Der}_x(X)$ then the value of $\mathfrak{d}(h)$ doesn't really depend on the whole function h, it only depends on the behaviour of h near the point x. This suggests we should consider the following subset of $C^\infty(X)$:

$$\{h \in C^\infty(X)\ ;\ \exists \text{ an open neighbourhood } U \text{ of } x \text{ with } h|_U \equiv 0\} \tag{D.4}$$

It's easy to see that this is a subspace of $C^\infty(X)$, so we may form the quotient space, which we'll denote by:

$$\widehat{C^\infty_x}(X)$$

Elements of $\widehat{C^\infty_x}(X)$ are called *germs* of smooth functions at x. A germ is an equivalence class of functions, where we declare that two functions are equivalent if they agree in some open neighbourhood of x. The lemma we just proved says that any $\mathfrak{d} \in \mathrm{Der}_x(X)$ induces a linear map:

$$\widehat{\mathfrak{d}} : \widehat{C^\infty_x}(X) \to \mathbb{R}$$

Now notice that if $[h] \in \widehat{C^\infty_x}(X)$ is a germ, we can evaluate $[h]$ at the point x to get a real number $h(x) \in \mathbb{R}$. This is well-defined, it doesn't depend on which representative of the germ we choose (but we cannot evaluate $[h]$ at any point other than x). Also notice that we have a well-defined multiplication in $\widehat{C^\infty_x}(X)$

$$([h_1], [h_2]) \mapsto [h_1 h_2]$$

because if h_1 vanishes in a neighbourhood of x, and h_2 is any function, then $h_1 h_2$ vanishes in a neighbourhood of x (*i.e.* the subspace (D.4) is an *ideal* in $C^\infty(X)$).

This means that it if we have a linear map from $\widehat{C^\infty_x}(X)$ to $\mathbb{R}$ then it makes sense to ask if it obeys the product rule (D.2). If we consider such a map $\widehat{\mathfrak{d}}$ induced from a $\mathfrak{d} \in \mathrm{Der}_x(X)$ then it will obey this rule; conversely if we take any linear map from $\widehat{C^\infty_x}(X)$ to $\mathbb{R}$ which obeys the product rule then by composing it with the quotient map

$$C^\infty(X) \to \widehat{C^\infty_x}(X)$$

we will get a derivation at x. So derivations at x are precisely linear maps from $\widehat{C^\infty_x}(X)$ to $\mathbb{R}$ which obey the product rule.

Now we show that the space of germs doesn't change if we restrict from X to an open neighbourhood of x.

Lemma D.5. *If $U \subset X$ is an open neighbourhood of x then we have a linear isomorphism*

$$\widehat{C^\infty_x}(X) \xrightarrow{\sim} \widehat{C^\infty_x}(U)$$

given by sending $[h]$ to $[h|_U]$.

Proof. Firstly note that if $h \in C^\infty(X)$ vanishes in an open neighbourhood of x then so does the function $h|_U \in C^\infty(U)$, so this map is well defined. Now let's find an inverse map. Choose a bump function $\psi \in C^\infty(X)$ which is constant with value 1 on some open neighourhood V of x, and vanishes outside some larger closed neighbourhood $W \subset U$. If $g \in C^\infty(U)$, then we can extend g to a function on the whole of X by defining:

$$\widehat{g} = \begin{cases} g\psi \text{ inside } U \\ 0 \text{ outside } U \end{cases}$$

Then $\widehat{g}$ is smooth, since it's smooth in the open sets U and $X \setminus W$. Then we have a linear map:

$$\begin{aligned} C^\infty(U) &\to C^\infty(X) \\ g &\mapsto \widehat{g} \end{aligned}$$

Furthermore if g vanishes on some neighbourhood of x then so does $\widehat{g}$, so we get an induced map:

$$\widehat{C^\infty_x}(U) \to \widehat{C^\infty_x}(X)$$

If h is a function on X then $\widehat{(h|_U)}$ agrees with h on the neighbourhood W. Also if g is a function on U then $(\widehat{g})|_U$ agrees with g on the neighbourhood W. This proves that the function $[g] \mapsto [\widehat{g}]$ is the inverse to the function $[h] \mapsto [h|_U]$. $\square$

Corollary D.6. *We have an isomorphism*

$$\mathrm{Der}_x(U) \xrightarrow{\sim} \mathrm{Der}_x(X)$$

which sends $\mathfrak{d}' \in \mathrm{Der}_x(U)$ to the operator $\mathfrak{d} : h \mapsto \mathfrak{d}'(h|_U)$.

Now we can prove the 'only if' direction in Proposition 7.12.

Proposition D.7. *Let $x \in X$ be a point in a manifold, and let $\mathfrak{d}$ be a derivation at x. If $h \in C^\infty(X)$ has rank zero at x then $\mathfrak{d}(h) = 0$.*

Proof. Pick any chart (U, f) around x. Then by Corollary D.6 $\mathfrak{d}$ determines an operator $\mathfrak{d}' \in \mathrm{Der}_x(U)$, with $\mathfrak{d}(h) = \mathfrak{d}'(h|_U)$. The co-ordinate function f gives an isomorphism between $C^\infty(U)$ and $C^\infty(\tilde{U})$, which induces an isomorphism between $\mathrm{Der}_x(U)$ and $\mathrm{Der}_{f(x)}(\tilde{U})$. The function h has rank zero at x, so the function $\tilde{h} = h \circ f^{-1} \in C^\infty(\tilde{U})$ has rank zero at $f(x)$. Now use Lemma D.1. $\square$

E Vector bundles

In Section 6 we introduced the tangent bundle TX to a manifold X, which is defined to be the set:

$$TX = \bigcup_{x \in X} T_x X$$

It comes with a surjection $\pi : TX \to X$ which sends a tangent vector $v \in T_x X$ to the point $x \in X$.

Proposition E.1. *The tangent bundle TX naturally has the structure of a manifold, with dimension $2(\dim X)$.*

Proof. We'll define our manifold structure by writing down a pseudo-atlas on TX, and invoking Proposition 2.26 and Corollary 2.28. In fact we've done most of the work already, in Section 6.1. Let (U, f) be any chart on X. We saw that this induces a bijection:

$$\begin{aligned} F : TU = \pi^{-1}(U) &\xrightarrow{\sim} \tilde{U} \times \mathbb{R}^n \\ (x, v) &\mapsto \big(f(x), \Delta_f(v)\big) \end{aligned}$$

The codomain here is an open subset in $\mathbb{R}^{2n}$, so this is a pseudo-chart on TX. If we do this for all charts on X then the corresponding set of pseudo-charts obviously cover TX, so they form a pseudo-atlas.

Now we must check conditions 1 and 2 from Proposition 2.26. Let (U_1, f_1) and (U_2, f_2) be two charts on X, and let $U = U_1 \cap U_2$ denote their intersection. Then the intersection of the pseudo-charts TU_1 and TU_2 is exactly TU, and

$$F_1(TU) = f_1(U) \times \mathbb{R}^n \ \subset \tilde{U}_1 \times \mathbb{R}^n$$

which is an open subset. Furthermore the transition function between these two pseudo-charts is

$$\begin{aligned} \Phi_{21} : f_1(U) \times \mathbb{R}^n &\xrightarrow{\sim} f_2(U) \times \mathbb{R}^n \\ (\tilde{x}, v) &\mapsto \big(\phi_{21}(\tilde{x}), D\phi_{21}|_{\tilde{x}}(v)\big) \end{aligned}$$

which is smooth. $\square$

Example E.2. Let X be the manifold T^1 from Example 2.11. We claim that the tangent bundle to T^1 is the infinite cylinder

$$T(T^1) \cong T^1 \times \mathbb{R}$$

(this is clearly a 2-dimensional manifold). We'll prove this claim carefully a bit later, but it's also quite easy to see using an atlas.

Recall that we have an atlas for T^1 with two charts,

$$f_1 : \ U_1 = T^1 \setminus [0] \ \xrightarrow{\sim} \ \tilde{U}_1 = (0, 1)$$

and

$$f_2 : \ U_2 : T^1 \setminus [\tfrac{1}{2}] \ \xrightarrow{\sim} \ \tilde{U}_2 = (-\tfrac{1}{2}, \tfrac{1}{2})$$

both of which simply lift an equivalence class to its representative in the given interval. The transition function between these two charts is:

$$\phi_{21} : (0, \tfrac{1}{2}) \sqcup (\tfrac{1}{2}, 1) \longrightarrow (-\tfrac{1}{2}, 0) \sqcup (0, \tfrac{1}{2})$$
$$x \mapsto \begin{cases} x, & \text{for } x < \frac{1}{2} \\ x - 1, & \text{for } x > \frac{1}{2} \end{cases}$$

So we have an atlas for the tangent bundle $T(T^1)$ which has two charts

$$T\tilde{U}_1 = (0, 1) \times \mathbb{R}$$

and:

$$T\tilde{U}_2 = (-\tfrac{1}{2}, \tfrac{1}{2}) \times \mathbb{R}$$

The derivative of the transition function ϕ_{21} at any point is just the identity map from $\mathbb{R}$ to $\mathbb{R}$, so the transition function between our two charts on $T(T^1)$ is just:

$$\Phi_{21} = (\phi_{21}, 1) : \Big((0, \tfrac{1}{2}) \sqcup (\tfrac{1}{2}, 1) \Big) \times \mathbb{R} \;\longrightarrow\; \Big((-\tfrac{1}{2}, 0) \sqcup (0, \tfrac{1}{2}) \Big) \times \mathbb{R}$$

A manifold which has this atlas must be $T^1 \times \mathbb{R}$.

This example is rather unusual, it's not normally true that the tangent bundle to X is simply $X \times \mathbb{R}^n$. Manifolds like this are called *parallelizable*, we'll come back to this later on. Also, this is unfortunately the only non-trivial example for which it's easy to visualise the tangent bundle, because if X has dimension ≥ 2 then TX has dimension ≥ 4.

For any manifold, the projection function $\pi : TX \to X$ is smooth, because if we look at it in one of the charts (TU, F) then it becomes the projection $\tilde{U} \times \mathbb{R}^n \to \tilde{U}$. The level sets of π are the individual tangent spaces $T_x X$, and it's immediately clear that these are n-dimensional submanifolds of TX; moreover it's clear that π is a submersion.

A vector field on X is, by definition, a function

$$\xi : X \to TX$$

such that $\pi \circ \xi = 1_X$. Given Proposition E.1, we don't need to do any extra work to define a *smooth* vector field, since we already have a definition of when a function from X to TX is smooth. If you look back at Section 6.1, this is precisely the definition of a smooth vector field that we wrote down there.

So for any manifold X we have an associated tangent bundle, which is a smooth manifold TX coming with a smooth surjection

$$\pi : TX \to X$$

such that every level set $\pi^{-1}(x) = T_x X$ is a vector space. This is a very rich mathematical structure, it's an example of something called a *vector bundle*. We're now going to describe this more general structure, but the definition is fairly complicated.

Let X be a manifold of dimension n. Informally, a vector bundle over a manifold X is a collection of vector spaces $\{E_x\}$, indexed by the points of x.

These vector spaces have to fit together to give a smooth manifold E, equipped with a smooth map $\pi : E \to X$ whose level set over $x \in X$ is the associated vector space $\pi^{-1}(x) = E_x$. For example, for any manifold X, and any integer r, there is a vector bundle

$$\pi : \; E = X \times \mathbb{R}^r \to X$$

where π is the projection map $\pi : (v, x) \mapsto x$. It's easy to show that there's a smooth structure on E making it into a manifold (of dimension $n+r$), and that π is smooth. Obviously the level set of π at any point $x \in X$ is the vector space $\mathbb{R}^r$. This is called the *trivial vector bundle* of rank r.

The tangent bundle is not usually of this form, in general we can't canonically identify T_xX with $\mathbb{R}^n$ so it's not usually true that $TX = X \times \mathbb{R}^n$. However if we pick a chart (U, f), then within the open set $U \subset X$ it *is* true that $TU = U \times \mathbb{R}^n$, since our co-ordinates give us this bijection. So within small neighbourhoods in X, the tangent bundle looks like the trivial bundle of rank n. This condition, of being 'locally trivial', is one of the key properties of a vector bundle.

Definition E.3. Let X be a manifold of dimension n. A **vector bundle** over X is the following data:

- A manifold E, of dimension $n + r$.
- A smooth surjection $\pi : E \to X$, whose level sets we denote $E_x = \pi^{-1}(x)$.
- For each $x \in X$, the structure of an r-dimensional vector space on the level set E_x.

We require that it is possible to find an atlas $\{(U_i, f_i), i \in I\}$ for X, and an atlas $\{(V_i, g_i), i \in I\}$ for E (indexed by the same set I), with the following properties:

(i) $V_i = \pi^{-1}(U_i)$, for each $i \in I$.

(ii) $\tilde{V}_i = \tilde{U}_i \times \mathbb{R}^r \subset \mathbb{R}^{n+r}$, for each $i \in I$.

(iii) For each $i \in I$, the square

$$\begin{array}{ccc} V_i & \xrightarrow{g_i} & \tilde{U}_i \times \mathbb{R}^r \\ \Big\downarrow{\scriptstyle \pi} & & \Big\downarrow \\ U_i & \xrightarrow{f_i} & \tilde{U}_i \end{array}$$

commutes, where the right-hand vertical arrow is the obvious projection map.

(iv) For any $x \in X$ and any $i \in I$, the map

$$g_i|_{E_x} : E_x \to \mathbb{R}^n$$

is an isomorphism of vector spaces.

The integer r is called the **rank** of the vector bundle, and the vector spaces E_x are called the **fibres** of the vector bundle.

A vector bundle is a lot of data: we need to specify E, π, and the vector space structure on each fibre E_x. However it's common to just write it as the map $\pi : E \to X$, or sometimes just as E.

Example E.4. For any manifold X, we have met the following examples of vector bundles on X:

(i) The tangent bundle $TX \to X$. When we defined the smooth structure on TX we used an atlas of exactly this form, so the tangent bundle is a vector bundle of rank n.

(ii) If we take E to be the trivial vector bundle of rank r, so $E = X \times \mathbb{R}^r$, then any atlas $\{(U_i, f_i), i \in I\}$ for X will produce an atlas for E of the required form, just by setting $V_i = U_i \times \mathbb{R}^r$ and $\tilde{V}_i = \tilde{U}_i \times \mathbb{R}^r$. So fortunately the trivial vector bundle of rank r is indeed an example of a vector bundle, of rank r.

(iii) In Section 8.1 we introduced the cotangent bundle:

$$T^\star X = \bigcup_{x \in X} T^\star_x X$$

Using an argument which is essentially identical to the proof of Proposition E.1, the cotangent bundle is a manifold of dimension $2n$. The same argument shows that it is a vector bundle over X, of rank n.

(iv) In Section 8.3 we generalized this to the set:

$$\wedge^p T^\star X = \bigcup_{x \in X} \wedge^p T^\star_x X$$

This is a vector bundle over X, of rank $\binom{n}{p}$.

But these are not the only examples of vector bundles.

Example E.5. Let E be the quotient of $\mathbb{R}^2$ by the equivalence relation

$$(x, y) \sim (x + n, (-1)^n y),\ n \in \mathbb{Z}$$

(these are the orbits of a group action generated by a horizontal glide reflection). This is an infinite Möbius strip. We have a well-defined map:

$$\begin{aligned} \pi : E &\to T^1 \\ [(x, y)] &\mapsto [x] \end{aligned}$$

Notice that the usual vector space structure on $\mathbb{R}^2$ does give a well-defined vector space structure on each fibre $E_{[x]}$. Futhermore if we take the atlas $\{(U_1, f_1), (U_2, f_2)\}$ on T^1 from Example 2.11 then it's easy to find a corresponding atlas $\{(V_1, g_1), (V_2, g_2)\}$ on E with all the properties required by Definition E.3. For example, $\tilde{V}_1 = (0, 1) \times \mathbb{R} \subset \mathbb{R}^2$, and V_1 is the image of $\tilde{V}_1$ in E.

The concept of a vector field can be easily generalized to other vector bundles:

Definition E.6. Let $\pi : E \to X$ be a vector bundle. A **section** of E is a smooth map

$$\sigma : X \to E$$

such that $\pi \circ \sigma = 1_X$.

So a section σ sends any point $x \in X$ to a vector $\sigma|_x \in E_x$ lying in the fibre over x, and this vector varies smoothly with x. A section of the tangent bundle TX is precisely a vector field. Similarly, a section of the cotangent bundle $T^\star X$ is a one-form, and a section of the bundle $\wedge^p T^\star X$ is a p-form.

Any vector bundle has one obvious section, the *zero section*, which maps any point $x \in X$ to the zero vector $0 \in E_x$. By looking at this in a chart it's clear that it is smooth, and furthermore it gives us an injective immersion

$$X \hookrightarrow E$$

whose image is a submanifold which is diffeomorphic to X.

We saw earlier (Example 6.2) that the manifold T^1 has the rather special property that the its tangent bundle looks like the trivial bundle $T^1 \times \mathbb{R}$. We now want to say this precisely, but we first we need to say what it means for two vector bundles over X to be *isomorphic*.

Definition E.7. Let $\pi_1 : E_1 \to X$ and $\pi_2 : E_2 \to X$ be two vector bundles over X. An **isomorphism** between E_1 and E_2 is a diffeomorphism

$$F : E_1 \xrightarrow{\sim} E_2$$

such that $\pi_2 \circ F = \pi_1$, and such that the induced function

$$F_x : (E_1)_x \to (E_2)_x$$

is a linear isomorphism, for each $x \in X$.

So an isomorphism of vector bundles is a bijection that preserves all the structure of a vector bundle. In particular if two vector bundles over X are isomorphic they must obviously have the same rank.

Definition E.8. A rank r vector bundle $\pi : E \to X$ is called **trivial** if it is isomorphic to the trivial vector bundle $X \times \mathbb{R}^r$.

Here is one way to tell if a vector bundle is trivial:

Proposition E.9. *Let $\pi : E \to X$ be a vector bundle of rank r. Then E is trivial iff there exist r sections $\sigma^1, ..., \sigma^r$ of E such that, for every point $x \in E$, the vectors*

$$\sigma^1|_x, ..., \sigma^r|_x \in E_x$$

form a basis of E_x.

Proof. If E is the trivial vector bundle $X \times \mathbb{R}^r$ then we can just pick any basis $e_1, ..., e_r$ for $\mathbb{R}^r$ and consider the constant sections $\tilde{\sigma}^i : x \mapsto e_i$ for each i, which are evidently smooth. More generally if $F : X \times \mathbb{R}^r \to E$ is an isomorphism of vector bundles then we can define r sections of E by:

$$\sigma^i = F \circ \tilde{\sigma}^i : \; x \mapsto F_x(e_i) \in E_x$$

These are smooth since both F and $\tilde{\sigma}^i$ are smooth, and give a basis of E_x since F_x is an isomorphism of vector spaces.

Conversely, suppose that we have such a set of sections $\sigma^1, ..., \sigma^r$. Define a function

$$F : X \times \mathbb{R}^r \to E$$

by:

$$F : \big(x, (v_1, ..., v_r)\big) \mapsto \left(x, \sum_{i=1}^{r} v_i \sigma^i|_x\right)$$

Obviously F commutes with the projection maps, and for each $x \in X$ the map F_x is linear and sends the standard basis of $\mathbb{R}^r$ to the basis $\sigma^1|_x, ..., \sigma^r|_x$ of E_x. Hence each F_x is an isomorphism of vector spaces, and it follows that F is a bijection.

Now pick a chart (U, f) on X and a chart (V, g) on E of the form specified in Definition E.3, and choose the corresponding chart $U \times \mathbb{R}^r$ on $X \times \mathbb{R}^r$. In these charts, each section σ^i is a smooth function

$$\tilde{\sigma}^i = (\tilde{\sigma}^i_1, ..., \tilde{\sigma}^i_r) : \; \tilde{U} \to \mathbb{R}^r$$

and F is the function

$$\tilde{F} : \tilde{U} \times \mathbb{R}^r \to \tilde{U} \times \mathbb{R}^r$$

given by the smooth family of invertible r-by-r matrices $M_{\tilde{x}}$ whose entries are $\tilde{\sigma}^i_j|_{\tilde{x}}$. The inverse function $\tilde{F}^{-1}$ is given by the family of matrices $M^{-1}_{\tilde{x}}$, whose entries will also vary smoothly with $\tilde{x}$ since they are rational functions of the entries in $M_{\tilde{x}}$. Hence both F and F^{-1} are smooth, so we have shown that F is an isomorphism of vector bundles. □

Definition E.10. A manifold X is said to be **parallelizable** iff its tangent bundle TX is trivial.

Example E.11. Let $X = S^1$. In Example 6.2 we found a smooth vector field ξ on S^1 which was not equal to zero at any point. Since S^1 is 1-dimensional, this means that ξ gives a basis of the tangent space at every point. So by Proposition E.9 the bundle TS^1 is trivial, and S^1 is parallelizable.

The manfold S^2 is not parallelizable, because of the following fact:

Theorem E.12 ('Hairy ball theorem')**.** *Any vector field on S^2 must be equal to zero at some point.*

Consequently it is impossible to find a pair of vector fields ξ_1, ξ_2 on S^2 that form a basis of the tangent space at every point.

Theorem E.12 is is a very nice result. It implies for example that at any moment in time there must be a point on the Earth where the wind speed is zero, and also that you cannot groom a spherical dog without leaving a protruding tuft of hair at one point. The proof is not very difficult, but unfortunately it requires some algebraic topology that doesn't form a part of this course.

Theorem E.12 is true for any even-dimensional sphere S^{2n}, so no even-dimensional sphere is parallelizable. In fact the only parallelizable spheres are S^1, S^3 and S^7, but this is rather harder to prove.

F Manifolds-with-boundary

In Section 9.3 we proved a version of Stokes' Theorem, which is based fundamentally on the fact that if a function $h \in C^\infty(\mathbb{R})$ vanishes outside some interval then we must have $\int dh = 0$. However, the fundamental theorem of calculus is much more precise than this, it says that for any function h, and any interval $[a, b]$, we have:

$$\int_a^b dh = h(b) - h(a)$$

We want to generalize this statement to n-forms on an arbitrary manifold. However, the interval $[a, b]$ is *not* a manifold, because of the end points a and b. The statement we are after requires us to generalize the notion of a manifold, to allow 'boundary points'.

Definition F.1. A (second-countable, Hausdorff) topological space X is called an **n-dimensional topological manifold-with-boundary** if for all points $x \in X$ we can find an open neighbourhood U of x, an open set

$$\tilde{U} \subset \mathbb{R}_{\leq 0} \times \mathbb{R}^{n-1}$$

and a homeomorphism $f : U \xrightarrow{\sim} \tilde{U}$.

So a manifold-with-boundary is a space that 'locally looks like' the half-space $\{x_1 \leq 0\} \subset \mathbb{R}^n$. We continue to use the name 'co-ordinate chart' for these homeomorphisms $f : U \xrightarrow{\sim} \tilde{U}$.

There are two kinds of point in a manifold-with-boundary:

- For some $x \in X$, we can find a chart $f : U \to \tilde{U}$ around x with $\tilde{U}$ entirely contained in $\mathbb{R}_{<0} \times \mathbb{R}^n$, so $\tilde{U}$ is an open subset of $\mathbb{R}^n$. The set of such x is called the *interior* of X.
- If $x \in X$ is not in the interior, then one can prove (using algebraic topology) that any chart around x must send x to a point on the hyperplane $\{x_1 = 0\} \cong \mathbb{R}^{n-1} \subset \mathbb{R}^n$. The set of these points is called the *boundary* of X, and denoted ∂X.

The interior of X is obviously a (non-compact) n-dimensional topological manifold. The boundary ∂X is an $(n-1)$-dimensional topological manifold, since it can be covered by the charts:

$$f : U \cap \partial X \xrightarrow{\sim} \tilde{U} \cap \{x_1 = 0\} \ \subset \mathbb{R}^{n-1}$$

However, note that X itself is not a topological manifold, of any dimension.

The entire theory of manifolds - smooth structures, tangent spaces, differential forms, etc - can be generalized to manifolds-with-boundary fairly easily. The only extra input needed is the definition of a smooth function between two open sets in a half-space, and for this we adopt the following convention: a function $F : U \to V$ between open sets

$$U \subset \mathbb{R}_{\leq 0} \times \mathbb{R}^{n-1} \qquad \text{and} \qquad V \subset \mathbb{R}_{\leq 0} \times \mathbb{R}^{k-1}$$

is called *smooth* if there exist open sets $\widehat{U} \subset \mathbb{R}^n$ and $\widehat{V} \subset \mathbb{R}^k$, with $U \subset \widehat{U}$ and $V \subset \widehat{V}$, and a smooth function $\widehat{F} : \widehat{U} \to \widehat{V}$ such that $\widehat{F}|_U = F$. With this definition we can define the derivative $DF|_x$ at any point in U, even the boundary points $\{x_1 = 0\}$, by taking the derivative of $\widehat{F}$. This doesn't depend on our choice of extension $\hat{F}$, because any partial derivative can be computed from the $x_1 \geq 0$ side.

Example F.2.
- The closed unit ball $X = \overline{B(0,1)} \subset \mathbb{R}^n$ is an n-dimensional manifold-with-boundary, with a smooth stucture inherited from $\mathbb{R}^n$. The interior of X is the open unit ball, and its boundary ∂X is S^{n-1}.
- Generalizing the previous example, if $h : \mathbb{R}^{n+1} \to \mathbb{R}$ is a smooth function and $\alpha \in \mathbb{R}$ is a regular value, then the proof of Proposition 3.16 adapts to prove that the set
$$X = \left\{x \in \mathbb{R}^{n+1},\ h(x) \leq \alpha\right\}$$
is an n-dimensional (smooth) manifold-with-boundary. Its boundary ∂X is the level set $h^{-1}(\alpha)$.
- If Y is any $(n-1)$-dimensional manifold, and I is the closed interval $I = [0, 1]$, then $X = Y \times I$ is an n-dimensional manifold-with-boundary. The boundary of X is the disjoint union $\partial X = Y \sqcup Y$ of two copies of Y.

The full version of Stokes' Theorem is about comparing the integral of differential forms over X and over the boundary ∂X. As usual, we start by seeing what happens in co-ordinates.

Let $U \subset \mathbb{R}_{\leq 0} \times \mathbb{R}^{n-1}$ be an open set, so U is a manifold-with-boundary, and its boundary is $\partial U = U \cap \{x_1 = 0\}$. Let $\iota : \partial U \hookrightarrow U$ denote the inclusion of the boundary. If we are given an $(n-1)$-form $\alpha \in \Omega^{n-1}(U)$, then we can pull-back α along ι to a get an $(n-1)$-form on ∂U. Let's write α explicitly as

$$\alpha = \alpha_1\, dx_2 \wedge ... \wedge dx_n \ + \ \alpha_2\, dx_1 \wedge dx_3 \wedge ... \wedge dx_n \ + \ ... \\ + \alpha_n dx_1 \wedge ... \wedge dx_{n-1}$$

for $\alpha_1, ..., \alpha_n \in C^\infty(U)$, and compute the pull-back $\iota^\star \alpha$.

For a point $z \in \partial U$, the tangent space to ∂U is the subspace:

$$T_z(\partial U) = \{x_1 = 0\} \ \subset \mathbb{R}^n = T_z U$$

Pulling back α to ∂U just means restricting it to this subspace at all points, *i.e.* evaluating it on $(n-1)$-tuples of vectors from the subspace $\left\langle \frac{\partial}{\partial x_2}, ..., \frac{\partial}{\partial x_n} \right\rangle$. This only picks up the first component of α, so:

$$\iota^\star \alpha = \alpha_1|_{\{x_1=0\}}\, dx_2 \wedge ... \wedge dx_n \quad \in \Omega^{n-1}(\partial U)$$

The local version of Stokes's Theorem on a manifold-with-boundary is the following statement:

Lemma F.3. *Let $U \subset \mathbb{R}_{\geq 0} \times \mathbb{R}^{n-1}$ be an open set, and let $\iota : \partial U \hookrightarrow U$ denote the inclusion of the boundary. Let $\alpha \in \Omega^{n-1}(U)$ be an $(n-1)$ form that vanishes outside some compact subset $W \subset U$. Then:*

$$\int_U d\alpha = \int_{\partial U} \iota^\star \alpha$$

Proof. This is very similar to the proof of Lemma 9.18. We can extend α to the whole of the half-space $\{x_1 \geq 0\}$, and then perform the integral over a 'half-cube' $[-r, 0] \times [-r, r]^{\times(n-1)}$ if r is big enough. Now suppose that α has only one non-vanishing component, and that component is

$$\alpha = \alpha_j \, dx_1 \wedge ... \wedge dx_{j-1} \wedge dx_{j+1} \wedge ... \wedge dx_n$$

for some $j > 1$. Then $\iota^\star \alpha$ vanishes, but the integral $\int_U d\alpha$ also vanishes, by the argument from Lemma 9.18. The remaining case to consider is when

$$\alpha = \alpha_1 \, dx_2 \wedge ... \wedge dx_n$$

and then we have

$$\begin{aligned}
\int_U d\alpha &= \int_{[-r,r]^{\times(n-1)}} \left(\int_{[-r,0]} \frac{\partial \alpha_1}{\partial x_1} \, dx_1 \right) dx_2 ... dx_n \\
&= \int_{[-r,r]^{\times(n-1)}} \alpha_1|_{\{x_1=0\}} \, dx_2 .... dx_n \\
&= \int_{\partial U} \iota^\star \alpha
\end{aligned}$$

since $\alpha_1 \equiv 0$ on the set $\{x_1 = -r\}$. □

To generalize this statement to more interesting manifolds, we just need one further fact:

Proposition F.4. *If X is an oriented manifold-with-boundary, then there is a canonical orientation on the boundary ∂X.*

Sketch proof: Let $\omega \in \Omega^n(X)$ be a volume form. For a point $x \in \partial X$, pick a chart around X, so locally we have an open set $\tilde{U} \subset \mathbb{R}_{\leq 0} \times \mathbb{R}^{n-1}$ and the boundary is $\partial \tilde{U} = \tilde{U} \cap \{x_1 = 0\}$. In this chart, ω becomes

$$\tilde{\omega} = h \, dx_1 \wedge ... \wedge dx_n$$

for some $h \in C^\infty(\tilde{U})$. Now choose a vector field along the boundary $\partial \tilde{U}$ which points 'out of' $\tilde{U}$, *i.e.* choose a function

$$\xi = (\xi_1, ..., \xi_n) : \ \partial \tilde{U} \longrightarrow \mathbb{R}^n$$

such that $\xi_1 > 0$ at all points. Contracting $\tilde{\omega}$ with ξ (see Problem Sheets) produces an $(n-1)$-form

$$i_\xi \tilde{\omega} = h\xi_1 \, dx_2 \wedge ... \wedge dx_n \ - \ h\xi_2 \, dx_1 \wedge dx_3 \wedge ... \wedge dx_n \ + \ ... \quad \in \Omega^{n-1}(\tilde{U})$$

and pulling this back to $\partial \tilde{U}$ gives an $(n-1)$ form

$$\iota^\star(i_\xi \tilde{\omega}) = h\xi_1 \, dx_2 \wedge ... \wedge dx_n$$

on $\partial \tilde{U}$. This is a volume form, since neither h nor ξ_1 vanish at any point. Furthermore, choosing a different ξ will have the effect of multiplying this volume form by some always-positive function in $C^\infty(\partial \tilde{U})$, so the orientation that we get on $\partial \tilde{U}$ is independent of ξ.

To do this over the whole of ∂X, we use a partition-of-unity to build a vector field ξ along ∂X which points outwards at all points. Then contracting the volume form on X with ξ gives a volume form on ∂X, and the orientation class is independent of ξ.

$\square$

Note that if (U, f) is an oriented chart, so $\tilde{\omega}$ is a positive multiple of the standard volume form, then this induced volume form $\iota^\star(i_\xi \tilde{\omega})$ is a positive multiple of the standard volume form on $\partial\tilde{U} \subset \mathbb{R}^{n-1}$. So the associated chart $(\partial U, f|_{\partial U})$ on ∂X is also oriented.

Example F.5.

- The closed ball $X = \overline{B(0,1)} \subset \mathbb{R}^n$ carries a volume form given by just restricting the standard volume form on $\mathbb{R}^n$. Hence the boundary $\partial X = S^{n-1}$ gets an induced orientation, and this is the same orientation that is produced by Proposition 9.6.
- If X is an oriented manifold, and I is the closed interval $[0,1]$, then $X \times I$ can be given an orientation by wedging the volume form on X with the standard volume form on I. Then the boundary $X \sqcup X$ gets an induced orientation, but on the first component this is the opposite orientation to the one that we started with.

Combining Lemma F.3 with the proof of Theorem 9.19 immediately proves:

Theorem F.6 (Stokes' Theorem, Version 2). *Let X be a compact oriented n-dimensional manifold-with-boundary, and let $\iota : \partial X \hookrightarrow X$ denote the inclusion of the boundary. For any $\alpha \in \Omega^{n-1}(X)$, we have:*

$$\int_X d\alpha = \int_{\partial X} \iota^\star \alpha$$

Note that ∂X is compact (it's a closed subset of a compact space) and oriented (by Proposition F.4) so integrating over ∂X does make sense.

Example F.7.

- Let D be the closed disc $D = \overline{B(0,1)} \subset \mathbb{R}^2$, and let $F : D \hookrightarrow \mathbb{R}^3$ be an injective immersion. For any one-form $\alpha \in \Omega^1(\mathbb{R}^3)$, we have that

  $$\int_D F^\star d\alpha = \int_{S^1} F^\star \alpha$$

  (since $F^\star d\alpha = d(F^\star \alpha)$. This is the classical version of Stokes' Theorem, although normally one thinks of α as a vector field, and $d\alpha$ as the curl of the vector field (see Example 8.30). As we've seen, it actually works when F is any smooth function.

- Let $F_0, F_1 : X \to Y$ be two smooth functions, and let

  $$H : X \times [0,1] \longrightarrow Y$$

  be a smooth function such that $H|_{X\times\{0\}} = F_0$ and $H|_{X\times\{1\}} = F_1$. This is called a (smooth) *homotopy*. Now suppose that X is compact and orientable, and that we have a closed n-form $\alpha \in \Omega^n(X)$, where n is the dimension of X. Then since $d\alpha = 0$, we have

$$\int_{X\times[0,1]} H^\star d\alpha = \int_X F_1^\star \alpha \; - \; \int_X F_0^\star \alpha = 0$$

and so $\int_X F_1^\star \alpha = \int_X F_0^\star \alpha$ (the minus sign appears in the above equation because the first boundary component carries the opposite orientation).