MATH 8210, FALL 2011 LECTURE NOTES

MIKE USHER

1. Multivariable calculus without coordinates

The objects of study in this course are what are called "smooth manifolds." For the time being I won't give a precise definition of these (it will come later, or of course you can easily look it up), but for now suffice it to say that these are topological spaces which locally resemble Euclidean space and in which, in particular, it is possible to do something resembling calculus. The surface of the Earth is (to good approximation) an example of a two-dimensional smooth manifold. Of course, the Earth is not $\mathbb{R}^2$ but rather a closed surface (I was going to say a sphere, but then it occurred to me that if one looks closely enough there are some rock formations which cause the genus to be positive), yet locally it looks enough like $\mathbb{R}^2$ that it seems reasonable to speak for instance of the directional derivatives of a function (the temperature, say) defined on the Earth.

So how can we formulate calculus in such spaces? Part of the definition will be that a manifold M will have an open cover $\{U_\alpha | \alpha \in A\}$ by sets equipped with homeomorphisms ("charts") $\phi_\alpha\colon U_\alpha \to V_\alpha$ where $V_\alpha \subset \mathbb{R}^n$ is open. So we can try to do calculus on M by, roughly speaking, doing standard multivariable calculus in the open sets V_α and then transporting the constructions back to M by the maps ϕ_α (or their inverses). However, if $m \in M$, then m will typically belong to several of the sets U_α in the open cover of M, and one needs to make sure that one's constructions don't depend on which of the charts one is using. To compare between the αth chart and the βth chart, one needs to look at the "transition function"

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

This is a map between two open subsets of $\mathbb{R}^n$, and part of the definition of a smooth manifold will ensure that the map is smooth (*i.e.*, C^∞) and invertible (with a smooth inverse), but there won't be any restrictions on what $\phi_\beta \circ \phi_\alpha^{-1}$ other than that. So for example it doesn't make sense to "take the partial derivative of a function on M with respect to the first coordinate," since although we can differentiate a function on V_α with respect to the first coordinate, or we can do the same for a function on V_β, these operations won't be equivalent when we try to lift them up to M using the maps ϕ_α, ϕ_β.

So this makes it important to understand how notions of multivariable calculus behave under the action of *diffeomorphisms* (*i.e.*, smooth maps with smooth inverses) $\phi\colon U \to \tilde{U}$ where U and $\tilde{U}$ are open subsets of $\mathbb{R}^n$. You should think of the action of such a diffeomorphism as being the same as changing one's coordinate system, *e.g.* from Cartesian coordinates to polar coordinates. In particular I want to first discuss various notions of what a *tangent vector at a point* $p \in U$ is. (And we'll later generalize this to the notion of a tangent vector at a point in a smooth manifold.) Visually you're supposed to think of a tangent vector at p as being a little arrow whose base is at p, pointing in a possible direction of motion from p. The set of these tangent vectors will form a vector space called the *tangent space to U at p* and denoted T_pU. I'll give three characterizations, from most concrete to most abstract.

(1) The way to describe this notion that is used in undergraduate multivariable calculus courses is just to say that a tangent vector v at $p \in U$ is (or is represented by) an n-tuple of numbers $(v_1, \ldots, v_n) \in \mathbb{R}^n$. One can then draw the vector whose base is at p and whose first coordinate is v_1, second coordinate is v_2, and so on. (In somewhat more sophisticated language, the standard Cartesian coordinates on $\mathbb{R}^n$ determine a basis $\{e_1, \ldots, e_n\}$ of unit vectors, and one has $v = \sum v_i e_i$.)

This characterization is very good for computational purposes, but when one is interested in how tangent vectors behave under coordinate changes $\phi\colon U \to \tilde{U}$ it has some disadvantages. The tangent vector $v = (v_1, \ldots, v_n) \in T_pU$ should correspond under the coordinate change ϕ to a tangent vector $\phi_* v \in T_{\phi(p)}\tilde{U}$ at $\phi(p)$. Perhaps you've learned how this correspondence works: one constructs the Jacobian matrix at p of the map ϕ (with (i, j) entry given by $\frac{\partial \phi_i}{\partial x_j}$ where ϕ_i is the ith component of ϕ), and then the coordinates of $\phi_* v$ are obtained by multiplying the Jacobian matrix by the vector consisting of the components of v. This is a manageable computation, but it may not be very conceptually clear from this discussion what's going on here. In particular if we then want to say what a tangent vector to a point m on a smooth manifold is we'd have to say something like "an n-tuple of numbers for each chart containing m, such that the n-tuples for different charts are related by the Jacobians of the transition functions," which is much more opaque and less natural-sounding than it really should be.

(2) A more natural characterization of tangent vectors is the following. The idea is that the tangent space T_pU consists of all possible velocities of curves passing through p. If $p \in U$, consider all C^∞ paths $\gamma\colon (-\epsilon, \epsilon) \to U$ (for some $\epsilon > 0$) such that $\gamma(0) = p$. I would like to declare two of these to be equivalent if they have the same velocity, *i.e.*, $\gamma_1 \sim \gamma_2$ iff $\gamma_1'(0) = \gamma_2'(0)$ (or equivalently, and maybe less circularly, $\gamma_1 \sim \gamma_2$ if $\lim_{t\to 0} \frac{\gamma_1(t)-\gamma_2(t)}{t} = 0$). Then simply define a "tangent vector" at p to be an equivalence class $[\gamma]$ of C^∞ arcs through p (and so T_pU is just the set of equivalence classes). The way this behaves under coordinate changes is extremely simple, since I'm not using coordinates to define the notion: a tangent vector $v \in T_pU$ has the form $v = [\gamma]$ for some γ, and the corresponding tangent vector $\phi_* v \in T_{\phi(p)}\tilde{U}$ is just $[\phi \circ \gamma]$. We'll see later that this adapts to general smooth manifolds very simply and directly—a tangent vector at a point on a smooth manifold will just be a suitable equivalence class of curves passing through that point.

The one disadvantage of this characterization is that it's not so intuitively obvious how to do algebraic operations (like addition of tangent vectors) on equivalence classes of curves through a point (though you can make a suitable definition if you put your mind to it).

It shouldn't be hard to construct a natural correspondence between tangent vectors in this sense and tangent vectors in the sense of Definition (1) above, but again, the advantage of thinking about it this way is that it's less coordinate-dependent.

(3) Now for a characterization of tangent vectors that you almost certainly would not have thought of. To attempt to motivate it, note that a given tangent vector $v \in T_pU$ gives you the ability to differentiate smooth functions $f\colon U \to \mathbb{R}$ at p—namely you take the directional derivative at p:

$$(D_v f)(p) = \lim_{t\to 0} \frac{f(p + tv) - f(p)}{t}.$$

So we will *define* a tangent vector at p to be "a way of differentiating functions defined near p," *i.e.*, we will abstract some relevant properties of the operation of taking a directional derivative, and then define a tangent vector to be one of these operations.

To do this, first consider pairs (f, V) where V is an open neighborhood of p and $f\colon V \to \mathbb{R}$ is C^∞, and declare two such pairs (f, V) and (g, W) to be equivalent if there is a smaller neighborhood $Z \subset V \cap W$ of p such that $f|_Z = g|_Z$. Let O_p be the set of equivalence classes. Since we can set, for instance $[f, V] \cdot [g, W] = [fg, V \cap W]$, O_p is easily seen to be a commutative $\mathbb{R}$-algebra (*i.e.*, it is both a commutative ring and a vector space over $\mathbb{R}$, with appropriately compatible operations), called the "algebra of germs of functions at p." I'll tend to denote a germ by just f rather than $[f, V]$; it is to be understood that f is defined not necessarily throughout U but rather on some (varying) open neighborhood of p. Of course one always has a well-defined value $f(p)$ for $f \in O_p$.

A tangent vector at p will then be defined to be a *derivation* $v\colon O_p \to \mathbb{R}$, *i.e.* v is to satisfy

- ($\mathbb{R}$-linearity) $v(cf + g) = cv(f) + v(g)$ for $c \in \mathbb{R}$ and $f, g \in O_p$
- (Leibniz rule) $v(fg) = f(p)v(g) + g(p)v(f)$ for $f, g \in O_p$.

It's standard that the directional derivative operations D_v alluded to above satisfy these properties. It's not obvious that, conversely, any derivation on $\mathcal{O}_p$ is given by a directional derivative in some direction, but we'll prove this shortly.

Like the characterization of tangent vectors as equivalence classes curves, this formulation is completely coordinate free, making it easy to extend the definition to manifolds when the time comes. Unlike the situation with curve characterization, though, it's quite obvious that derivations form a vector space, which is another advantage.

To see how this notion behaves under diffeomorphisms (or indeed under more general smooth maps) $\phi\colon U \to \tilde{U}$, if $v \in T_pU$ (*i.e.*, if v is a derivation on $\mathcal{O}_p$), we need to construct a derivation $\phi_* v$ on $\mathcal{O}_{\phi(p)}$. Well, if $f \in \mathcal{O}_{\phi(p)}$ (really we should write $[f, V]$), so f is a smooth function defined near $\phi(p)$, then $f \circ \phi$ will be a smooth function defined near p (specifically, it will be defined on the open set $\phi^{-1}(V)$ around p), and so we can define

$$(\phi_* v)(f) = v(f \circ \phi)$$

So as with the curve formulation, it's quite simple to see how derivations transform under coordinate changes.

Among the three above characterizations of tangent vectors, it should be clear that (1) is equivalent to (2), under the correspondence which assigns to an equivalence class of curves $[\gamma]$ the vector $\gamma'(0)$ (expressed in coordinates using the standard basis for $\mathbb{R}^n$). We now set about proving that (1) and (3) are also equivalent. Let T_pU denote the space of tangent vectors as given by formulation (1) (*i.e.*, as elements of $\mathbb{R}^n$) and (for the moment) $\tilde{T}_pU$ that given by (3) (*i.e.*, as derivations). Write the coordinates of $p \in U \subset \mathbb{R}^n$ as $(p_1, \dots, p_n)$. Now we have a linear map $\alpha\colon T_pU \to \tilde{T}_pU$ given by

$$\alpha(v_1, \dots, v_n) = \sum_{i=1}^n v_i \frac{\partial}{\partial x_i},$$

i.e., α sends a vector (in the undergraduate multivariable calculus sense) to the operation given by directional differentiation in the direction of that vector. We claim that α is bijective, justifying our proposal to regard (3) as an equivalent definition of the tangent space at p. It should be clear that α is injective. Indeed, for each i we have an element $x_i - p_i \in \mathcal{O}_p$, and we see that, where $\beta\colon \tilde{T}_pU \to T_pU$ is given by

$$\beta(v) = (v(x_1 - p_1), \dots, v(x_n - p_n)),$$

we have $\beta \circ \alpha = 1$ (as $\frac{\partial}{\partial x_i}(x_j - p_j) = \delta_{ij}$). Thus α is injective, and β surjective. To see that α is surjective, we note the following, whenever $v \in \tilde{T}_pU$:

- $v(1) = v(1 \cdot 1) = 1v(1) + 1v(1) = v(1) + v(1)$. Hence $v(1) = 0$, and so by $\mathbb{R}$-linearity $v(c) = 0$ for every constant function c.
- For any i and j, if $f \in \mathcal{O}_p$ we have

$$v\Big((x_i - p_i)(x_j - p_j)f\Big) = (x_i - p_i)|_p v((x_j - p_j)f) + (x_j - p_j)|_p f(p) v((x_i - p_i)) = 0.$$

- By the multivariable Taylor formula, any (germ of a) function $g \in \mathcal{O}_p$ can be written (on some neighborhood of p)

$$g(x) = g(p) + \sum_{i=1}^n \frac{\partial g}{\partial x_i}(p)(x_i - p_i) + \sum_{i,j=1}^n (x_i - p_i)(x_j - p_j) f_{ij}(x)$$

for some $f_{ij} \in \mathcal{O}_p$. Hence by the first two items and the linearity of v, we get

$$v(g) = \sum_{i=1}^n \frac{\partial g}{\partial x_i}(p) v(x_i - p_i).$$

Thus
$$v = \sum v_i \frac{\partial}{\partial x_i} = \alpha(v_1, \dots, v_n),$$
where the numbers v_i are equal to $v(x_i - p_i)$.

In view of the above correspondence, we can drop the tilde in the notation $\tilde{T}_p U$, and always view tangent vectors as derivations on spaces of germs of functions. Even when we express a tangent vector in coordinates, we will often use notation consistent with the derivation interpretation and write the vector as
$$v_1 \frac{\partial}{\partial x_1} + \cdots + v_n \frac{\partial}{\partial x_n}$$
rather than $(v_1, \dots, v_n)$.

Of course, another familiar notion from multivariable calculus is that of a *vector field* on an open set U, which can be thought of as a smooth family of tangent vectors at all of the points of U, or as a smooth vector-valued function $X\colon U \to \mathbb{R}^n$, expressible in coordinates as $X(m) = (X_1(m), \dots, X_n(m))$. There is also a coordinate-free interpretation of what a vector field is: it is a map $X\colon C^\infty(U) \to C^\infty(U)$ which, as with tangent vectors, is a derivation, namely:

- $X(cf + g) = cX(f) + X(g)$ for all $c \in \mathbb{R}$, $f, g \in C^\infty(U)$, and
- $X(fg) = fX(g) + gX(f)$ for all $f, g \in C^\infty(M)$.

Note that while tangent vectors, when viewed as derivations, just take values in $\mathbb{R}$, vector fields take values in the space of smooth functions. Just as with tangent vectors, there's a natural one-to-one correspondence between the undergraduate versions of vector fields and the derivations on $C^\infty(U)$: simply assign to $(X_1(\cdot), \dots, X_n(\cdot))$ the derivation
$$f \mapsto \sum_{i=1}^n X_i \frac{\partial f}{\partial x_i}.$$

Again, the great advantage of the derivation interpretation is that it makes no direct reference to coordinates. So on a smooth manifold M, once have defined the space of smooth functions $C^\infty(M)$, we will effortlessly be able to define a vector field on M as a derivation $X\colon C^\infty(M) \to C^\infty(M)$.

Another nice feature of the derivation interpretation for vector fields (but not for tangent vectors) is that it points toward some additional structure on the space of vector fields that we wouldn't have noticed if we just worked in coordinates. Namely, given that a vector field is a certain kind of function $X\colon C^\infty(U) \to C^\infty(U)$, it becomes natural to think about composing such functions. Now a slight hitch with this is that the composition of two derivations will not typically be a derivation. For example, $\frac{\partial}{\partial x_1}$ is a derivation, but $\frac{\partial}{\partial x_1} \circ \frac{\partial}{\partial x_1}$ certainly is not: namely we have
$$\frac{\partial}{\partial x_1} \circ \frac{\partial}{\partial x_1}(x_1 x_1) = 2$$
but
$$x_1 \frac{\partial}{\partial x_1} \circ \frac{\partial}{\partial x_1}(x_1) + x_1 \frac{\partial}{\partial x_1} \circ \frac{\partial}{\partial x_1}(x_1) = 0.$$
So while we can "compose" two vector fields the result won't be a vector field. However:

Proposition 1.1. *Let $\mathcal{A}$ be a commutative $\mathbb{R}$-algebra and let $X, Y\colon \mathcal{A} \to \mathcal{A}$ be two derivations on $\mathcal{A}$. Then the* commutator $[X, Y] := X \circ Y - Y \circ X$ *is also a derivation on $\mathcal{A}$.*

Proof. The linearity of $[X, Y]$ is trivial, so we just need to check the Leibniz rule. We find, for $f, g \in \mathcal{A}$:
$$\begin{aligned}
[X, Y](fg) &= X(Y(fg)) - Y(X(fg)) = X(fYg + gYf) - Y(fXg + gXf) \\
&= (fXYg + (Xf)(Yg) + gXYf + (Xg)(Yf)) - (fYXg + (Yf)(Xg) + gYXf + (Yg)(Xf)) \\
&= f(XY - YX)g + g(XY - YX)f = f[X, Y](g) + g[Y, X](f),
\end{aligned}$$
which is precisely the Leibniz rule for $[X, Y]$. □

In local coordinates, if $X = \sum X_i \frac{\partial}{\partial x_i}$ and $Y = \sum Y_j \frac{\partial}{\partial x_j}$, then one finds

$$\begin{aligned}
[X,Y](f) &= \sum_{i=1}^{n} X_i \frac{\partial}{\partial x_i}\left(\sum_{j=1}^{n} Y_j \frac{\partial f}{\partial x_j}\right) - \sum_{i=1}^{n} Y_i \frac{\partial}{\partial x_i}\left(\sum_{j=1}^{n} X_j \frac{\partial f}{\partial x_j}\right) \\
&= \sum_{i,j=1}^{n}\left(X_i Y_j \frac{\partial^2 f}{\partial x_i \partial x_j} + X_i \frac{\partial Y_j}{\partial x_i}\frac{\partial f}{\partial x_j}\right) - \sum_{i,j=1}^{n}\left(Y_i X_j \frac{\partial^2 f}{\partial x_i \partial x_j} + Y_i \frac{\partial X_j}{\partial x_i}\frac{\partial f}{\partial x_j}\right) \\
&= \sum_{j=1}^{n}\left(\sum_{i=1}^{n} X_i \frac{\partial Y_j}{\partial x_i} - Y_i \frac{\partial X_j}{\partial x_i}\right)\frac{\partial f}{\partial x_j}.
\end{aligned}$$

Thus $[X,Y]$ is the vector field $\sum Z_j \frac{\partial}{\partial x_j}$ whose jth component is given by

$$Z_j = \sum_{i=1}^{n}\left(X_i \frac{\partial Y_j}{\partial x_i} - Y_i \frac{\partial X_j}{\partial x_i}\right) \tag{1}$$

This commutator operation on vector fields (also called the *Lie bracket*) turns out to be a fairly important one. Of course, if one wanted to work entirely in coordinates without taking a more abstract point of view, it would have been possible to just define the Lie bracket of two vector fields X and Y to be the vector field given by formula (1), but it's not clear why one would be motivated to do so.

In general, the commutator operation $[\cdot,\cdot]$ on the space of linear maps from a vector space to itself satisfies the **Jacobi identity**:

$$[X,[Y,Z]] + [Z,[X,Y]] + [Y,[Z,X]] = 0 \tag{2}$$

Indeed, the left hand side is equal to

$$X(YZ - ZY) - (YZ - ZY)X + Z(XY - YX) - (XY - YX)Z + Y(XZ - ZX) - (ZX - XZ)Y$$

and (using associativity of function composition) you can see that each of the six three-letter words made up of one each of the letters X,Y,Z appears above once positively and once negatively, so the sum is zero. Note that if $[\cdot,\cdot]$ were an associative operation we would instead have $[X,[Y,Z]] + [Z,[X,Y]] = [X,[Y,Z]] - [[X,Y],Z] = 0$; thus the Jacobi identity expresses a particular way for a binary operation to be non-associative. In general a vector space L equipped with a binary operation $[\cdot,\cdot]\colon A \times A \to A$ which is bilinear, which obeys $[X,Y] = -[Y,X]$, and which satisfies the Jacobi identity is called a *Lie algebra*; thus we have shown that, if $U \subset \mathbb{R}^n$ is open, then the space $\mathcal{X}(U)$ of vector fields on U is naturally a Lie algebra.

Exercise 1.2. a) Let $\phi\colon U \to V$ be a diffeomorphism between two open subsets of $\mathbb{R}^n$, and let X be a vector field on U. Prove that if $\phi_* X\colon C^\infty(V) \to C^\infty(V)$ is defined by $((\phi_* X)(f))(\phi(p)) = (X(f \circ \phi))(p)$, then $\phi_* X$ is a vector field on V. Why did we have to assume that ϕ was a diffeomorphism (or at least bijective) in order to do this (unlike the situation with tangent vectors, which can be pushed forward by any smooth map)?

b) Prove that if X, Y are two vector fields on U and if $\phi\colon U \to V$ is a diffeomorphism then

$$\phi_*[X,Y] = [\phi_* X, \phi_* Y].$$

Exercise 1.3. Define the following three vector fields[1] on $\mathbb{R}^3$:

$$I = z\frac{\partial}{\partial y} - y\frac{\partial}{\partial z}$$
$$J = x\frac{\partial}{\partial z} - z\frac{\partial}{\partial x}$$
$$K = y\frac{\partial}{\partial x} - x\frac{\partial}{\partial y}$$

a) Compute $[I, J]$, $[I, K]$, and $[J, K]$.
b) Deduce as a formal consequence of part (a) that the cross product on $\mathbb{R}^3$ satisfies the Jacobi identity.

2. Bump functions and partitions of unity in $\mathbb{R}^n$

In point-set topology one learns a result called Urysohn's Lemma, which states that given inclusions $A \subset U \subset X$ where X is a normal topological space, U is open, and A is closed, there is a continuous function $\chi\colon X \to [0, 1]$ identically equal to one on A and identically zero on $X \setminus U$. A version of this result is extremely important in differential topology (perhaps more important than in point-set topology); unfortunately, since we need our functions to be C^∞ and not just continuous, we can't just cite Urysohn's Lemma but rather need to prove a new, smooth, version of the result (of course, this smooth version will apply in a more limited context, if only because it doesn't make sense to speak of "smooth functions" on a general normal topological space). The good news is that the functions can be constructed in a more concrete fashion than one sees in the proof of Urysohn's Lemma.

We begin with a result in one-variable calculus.

Lemma 2.1. *Define the function* $f\colon \mathbb{R} \to \mathbb{R}$ *by*

$$f(t) = \begin{cases} e^{-1/t} & t > 0 \\ 0 & t \leq 0 \end{cases}$$

Then $f \in C^\infty(\mathbb{R})$. *Indeed, for all* $k \in \mathbb{N}$ *there is a polynomial* $P_k \in \mathbb{R}[t]$ *with the property that the kth derivative* $f^{(k)}$ *exists and is given by*

$$f^{(k)}(t) = \begin{cases} P_k(1/t)e^{-1/t} & t > 0 \\ 0 & t \leq 0 \end{cases} \tag{3}$$

Proof. First note that if (3) holds, then $f^{(k)}$ is continuous on all of $\mathbb{R}$: indeed continuity is obvious everywhere except zero, and at zero we have, by repeated applications of L'Hôpital's rule,

$$\lim_{t\to 0^+} P_k(1/t)e^{-1/t} = \lim_{s\to\infty} \frac{P_k(s)}{e^s} = \lim_{s\to\infty} \frac{c_k}{e^s} = 0$$

where c_k is some constant (which results from differentiating $degP_k$-many times the polynomial P_k), from which continuity at zero follows directly.

Thus we just need to prove (3), which we do by induction on k. So assume (3) holds for k; we prove it for $k + 1$. For $t < 0$ the formula is trivial. For $t = 0$ we see

$$\lim_{t\to 0^+} \frac{f^{(k)}(t) - f^{(k)}(0)}{t} = \lim_{t\to 0^+} \frac{1}{t}P_k(1/t)e^{-1/t} = \lim_{s\to\infty} \frac{sP_k(s)}{e^s} = 0$$

[1]Though it's not necessary in order to do the problem, you might convince yourself that if one interprets these vector fields in the standard multivariable calculus sense, I points in the direction of a rotation around the x-axis, J in the direction of a rotation around the y-axis, and K in the direction of a rotation around the z-axis.

by L'Hôpital's rule, and so (since the left-hand limit is trivially zero) we have $f^{(k+1)}(t) = 0$. Finally for $t > 0$ we have, by the product and chain rules,

$$f^{(k+1)}(t) = \frac{d}{dt}\left(P_k(1/t)e^{-1/t}\right) = -\frac{1}{t^2}P_k'\left(\frac{1}{t}\right)e^{-1/t} + \frac{1}{t^2}P_k\left(\frac{1}{t}\right)e^{-1/t},$$

and so the formula holds with

$$P_{k+1}(s) = s^2(P_k'(s) + P_k(s)).$$

□

Note that our function f is a surjection to the half-open interval $[0,1)$, with $f^{-1}(\{0\}) = (-\infty, 0]$. Out of this function we can build many other useful ones. For instance:

Corollary 2.2. *There is a C^∞ function $g\colon \mathbb{R} \to [0,1]$ with the property that $g^{-1}(\{1\}) = [1,\infty)$ and $g^{-1}(\{0\}) = (-\infty, 0]$.*

Proof. Note that the function $t \mapsto f(1-t)$ is smooth and nonnegative, and equals zero precisely on the interval $[1,\infty)$. In particular $f(t) + f(1-t)$ is positive everywhere. So we can let

$$g(t) = \frac{f(t)}{f(t) + f(1-t)}.$$

I leave it to you to check that this has the desired properties. □

Corollary 2.3. *For any real numbers $a < b$ there is a C^∞ function $g_{a,b}\colon \mathbb{R} \to [0,1]$ such that $g_{a,b}^{-1}(\{0\}) = (-\infty, a]$ and $g_{a,b}^{-1}(\{1\}) = [b,\infty)$.*

Proof. Let

$$g_{a,b}(t) = g\left(\frac{t-a}{b-a}\right).$$

□

Corollary 2.4. *For any real numbers $a < b < c < d$ there is a smooth "bump" function $h\colon \mathbb{R} \to [0,1]$ so that $h^{-1}(\{1\}) = [b,c]$ and $h^{-1}(\{0\}) = (-\infty, a] \cup [d,\infty)$.*

Proof. Let

$$h(t) = g_{a,b}(t)(1 - g_{c,d}(t)).$$

□

Corollary 2.5. *For $x \in \mathbb{R}^n$ and $r > 0$ let $B_r(x) = \{y \in \mathbb{R}^n | ||y - x|| < r\}$ denote the open ball of radius r around x. Then for any $0 < s < r$ there is a smooth function $\beta\colon \mathbb{R}^n \to [0,1]$ such that $\beta^{-1}(\{1\}) = \overline{B_s(x)}$ and $supp(\beta) = \overline{B_r(x)}$.*

(Here by $supp(\beta)$ we mean the *support* of β, *i.e.*, the closed set $\overline{\{y \in \mathbb{R}^n | \beta(y) \neq 0\}}$)

Proof. Let

$$\beta(y) = 1 - g_{s^2,r^2}(||y - x||^2).$$

□

Our goal now is the following theorem:

Theorem 2.6. *Let $U \subset \mathbb{R}^n$ be an open set, and let $\mathcal{V} = \{V_\alpha | \alpha \in A\}$ be an open cover of U. Then there are C^∞ functions $\chi_\alpha\colon U \to [0,1]$ obeying the following properties:*

(i) *$supp(\chi_\alpha) \subset V_\alpha$*
(ii) *Any $x \in U$ has a neighborhood W_x with the property that $\chi_\alpha|_{W_x} = 0$ for all but finitely many α.*
(iii) *For all $x \in U$ we have $\sum_\alpha \chi_\alpha(x) = 1$.*

Note that property (ii) ensures that $\sum_\alpha \chi_\alpha$ is well-defined and smooth (even if there are infinitely many—perhaps uncountably many—different α), since U is then covered by open sets on each of which the sum $\sum_\alpha \chi_\alpha$ is really a finite sum (all but finitely many terms are zero).

Definition 2.7. *A collection of functions $\{\chi_\alpha | \alpha \in A\}$ obeying properties (i)-(iii) of Theorem 2.6 is called a* partition of unity subordinate to the cover $\{V_\alpha\}$.

Theorem 2.6 has an analogue for general smooth manifolds (see Theorem 3.17); to make this more general version eventually easier to reach we present the proof for open sets in $\mathbb{R}^n$ in a fairly general way (a proof more specifically adapted to $\mathbb{R}^n$ can be found in Appendix A of Madsen-Tornehave). In particular we bring in the following definition from point-set topology:

Definition 2.8. *A topological space X is called* second-countable *if there is a countable basis for the topology of X.*

In other words, there should be a collection $\{O_n | n \in \mathbb{N}\}$ of open sets with the property that if U is open and $x \in U$ then $x \in O_n \subset U$ for some n. For example $\mathbb{R}^n$ has this property (take the base to consist of open balls centered at points with rational coordinates and having rational radius), as does any open subset of $\mathbb{R}^n$ (just use those rational balls that are contained in the open subset). Part of our eventual definition will require that any smooth manifold also has this property.

Lemma 2.9. *Let X be a second-countable locally compact Hausdorff space. Then there is a sequence of compact sets $\{K_i\}_{i=1}^\infty$ and a sequence of open sets $\{H_i\}_{i=1}^\infty$ such that*

- $K_i \subset H_i$
- $X = \cup_{i=1}^\infty K_i = \cup_{i=1}^\infty H_i$
- *If $j \geq i+3$ then $H_i \cap H_j = \emptyset$.*

Proof. First note that a second-countable, locally compact space has a countable base for its topology which consists of open sets with compact closure. Indeed, given a countable base $\mathcal{B}$, by local compactness any point $x \in X$ has a neighborhood O_x with compact closure, and there will be some $V \in \mathcal{B}$ such that $x \in V \subset O_x$; evidently $\overline{V}$ will be compact, and the set of all V that can be obtained in this fashion will still be a base for the topology (and will be contained in the original $\mathcal{B}$, so will be countable).

So let $\{U_i\}_{i=0}^\infty$ be a base for the topology which is countable and such that each $\overline{U_i}$ is compact. In particular the U_i cover X. We claim now that there is a sequence $\{G_i\}_{i=0}^\infty$ of open sets with each $\overline{G_i}$ compact, such that $\overline{G_i} \subset G_{i+1}$ and such that $\cup_{i=0}^\infty G_i = X$. Specifically, the G_i will have the form

$$G_i = U_0 \cup \cdots \cup U_{j_i}$$

for a certain increasing sequence of natural numbers $\{j_i\}$. To construct the sequence $\{j_i\}$, we let $j_0 = 0$ (so $G_0 = U_0$), and assuming that we have chosen j_k, so that $G_k = U_1 \cup \cdots \cup U_{j_k}$, we note that $\overline{G_k}$ is compact since the $\overline{U_i}$ are, and so since the U_i cover X there must be some $j_{k+1} > j_k$ so that $\overline{G_k} \subset \cup_{i=1}^{j_{k+1}} U_i$. Inductively choosing the j_k in this fashion results in a sequence G_i satisfying the required properties (the fact that the G_i cover X follows from the fact that the U_i do, and the fact that $j_i \to \infty$ since the j_i are a strictly increasing sequence of natural numbers).

To construct K_i and H_i, let $K_1 = \overline{G_1}$, $W_1 = G_2$, and, for $i \geq 2$, let $K_i = \overline{G_i} \setminus G_{i-1}$ and $H_i = G_{i+1} \setminus \overline{G_{i-2}}$. These are easily seen to satisfy the required properties.

□

Proof of Theorem 2.6. Let K_i and H_i be subsets of U as in Lemma 2.9 (applied with $X = U$), and fix any i. For all $x \in K_i$ we may choose $\alpha_x \in A$ and $\epsilon_x > 0$ so that $B_{2\epsilon_x}(x) \subset V_{\alpha_x} \cap H_i$. Then the collection of open balls $\{B_{\epsilon_x}(x) | x \in K_i\}$ covers K_i, so it has a finite subcover.

Now letting i vary and taking the union of all of these finite subcovers, we have a countable collection of balls $\{B_k\}_{k=1}^\infty$ that covers X, and such that where $\tilde{B}_k$ denotes the ball with the same center as B_k but twice the radius, there are α_k and i_k such that $\tilde{B}_k \subset V_{\alpha_k} \cap H_{i_k}$. (While there may be more than one such α_k and i_k—there might even be uncountably many possible α_k—we specifically choose one α_k and i_k for every k. For convenience let us take i_k to be the i for which B_k was a member of the finite subcover of K_i, so that in particular for any i there are just finitely many k with $i_k = i$.)

I claim that the balls $\tilde{B}_k$ form a *locally finite* cover of U, *i.e.* that any point $x \in U$ has a neighborhood O_x which meets just finitely many of the $\tilde{B}_k$. Indeed we could use for O_x any neighborhood of x with compact closure. For then O_x is contained in the union of just finitely many of the sets H_i, say $O_x \subset H_1 \cup \cdots \cup H_r$. But the H_i have the property that $H_i \cap H_m = \emptyset$ whenever $m \geq i + 3$, and so $O_x \cap H_m = \emptyset$ for $m \geq r + 3$. Consequently $\tilde{B}_k \cap O_x = \emptyset$ unless k is one of the finitely many indices having $i_k \leq r + 2$.

We can now construct the desired functions. First, for each k, let $\psi_k \colon U \to [0, 1]$ be a smooth function identically equal to 1 on B_k and such that $supp(\psi_k) \subset \tilde{B}_k$; such ψ_k exist by Corollary 2.5. By the previous paragraph, any point in U has a neighborhood which is disjoint from the supports of all but finitely many of the ψ_k; consequently

$$\psi = \sum_{k=1}^{\infty} \psi_k$$

is a well-defined, smooth function. Moreover $\psi > 0$ everywhere, since the (smaller) balls B_k cover U. So for any k we have a well-defined, smooth function $\frac{\psi_k}{\psi}$, and obviously $\sum_k \frac{\psi_k}{\psi} = 1$.

Now define

$$\chi_\alpha = \sum_{k:\alpha_k=\alpha} \frac{\psi_k}{\psi}.$$

Since $\tilde{B}_k \subset V_\alpha$ whenever $\alpha = \alpha_k$, we have $supp(\chi_\alpha) \subset V_\alpha$ for all α. Since any point has a neighborhood intersecting the support of ψ_k for only finitely many k, there will be just finitely many χ_α whose supports intersect this neighborhood (namely, just those α which equal α_k for one of these k). Finally, we clearly have

$$\sum_\alpha \chi_\alpha = \sum_\alpha \sum_{k:\alpha_k=\alpha} \frac{\psi_k}{\psi} = \sum_k \frac{\psi_k}{\psi} = 1.$$

□

As essentially a special case we get a direct analogue of Urysohn's Lemma:

Corollary 2.10. *If $A \subset U \subset \mathbb{R}^n$ with A closed and U open, there is a C^∞ function $f \colon \mathbb{R}^n \to [0, 1]$ with $f|_A = 1$ and $supp(f) \subset U$.*

Proof. Let $\{\chi_1, \chi_2\}$ be a partition of unity subordinate to the cover $\{U, \mathbb{R}^n \setminus A\}$ of $\mathbb{R}^n$, and let $f = \chi_1$. I leave it to you to confirm the desired properties. □

Exercise 2.11. a) Let $U \subset \mathbb{R}^n$ be open, let $p \in U$, and let X be a vector field on U (use the interpretation of X as a derivation from $C^\infty(U)$ to itself). Prove that one can obtain a well-defined tangent vector (in the sense of a derivation $\mathcal{O}_p \to \mathbb{R}$) X_p by the following prescription: If $[f, V] \in \mathcal{O}_p$, let $\tilde{f} \in C^\infty(U)$ be a function such that $[\tilde{f}, U] = [f, V]$. Then $X\tilde{f} \in C^\infty(U)$, and we set

$$X_p([f, V]) = (X\tilde{f})(p)$$

(Part of the problem is showing that $\tilde{f}$ exists, and moreover that $X_p([f, V])$ is independent of the choice of such a $\tilde{f}$.)

b) If in coordinates we have $X = \sum_i f_i \frac{\partial}{\partial x_i}$, prove that $X_p = \sum_i f_i(p) \frac{\partial}{\partial x_i}$.

3. Smooth manifolds

Definition 3.1. *Let $n \in \mathbb{N}$. An n-*dimensional topological manifold *(or "topological n-manifold") is a second-countable Hausdorff space M with the property that, for all $m \in M$, there is a neighborhood $U \subset M$ of m and a homeomorphism $\phi\colon U \to V$ where $V \subset \mathbb{R}^n$ is an open subset.*

Remark 3.2. Of course, by replacing V with a small open ball $B \subset V$ around $\phi(p)$ and U with $\phi^{-1}(B)$, we could just as well require the image of ϕ is an open ball in $\mathbb{R}^n$ rather than an arbitrary open set. In turn, since any open ball in $\mathbb{R}^n$ is homeomorphic (and indeed diffeomorphic) to $\mathbb{R}^n$, we could equally well require the images of the maps ϕ in Defnition 3.1 to all be $\mathbb{R}^n$—*i.e.*, a topological n-manifold is a second-countable Hausdorff space in which every point has a neighborhood homeomorphic to $\mathbb{R}^n$.

Definition 3.3. *Let M be a topological n-manifold, and let k be either a positive integer or ∞. A C^k* atlas *on M is a collection $\mathcal{A} = \{(U_\alpha, \phi_\alpha) | \alpha \in A\}$ where*

- *The U_α are open subsets of M, and $\cup_{\alpha \in A} U_\alpha = M$.*
- *Each $\phi_\alpha\colon U_\alpha \to \mathbb{R}^n$ is a homeomorphism from U_α to the open subset $\phi_\alpha(U_\alpha) \subset \mathbb{R}^n$, and*
- *If $\alpha, \beta \in A$ are such that $U_\alpha \cap U_\beta \neq \varnothing$, then*
$$\phi_\beta \circ \phi_\alpha^{-1}\colon \phi_\alpha(U_\alpha \cap U_\beta) \to \phi_\beta(U_\alpha \cap U_\beta)$$
is of class C^k.

The maps $\phi_\alpha\colon U_\alpha \to \mathbb{R}^n$ are called coordinate charts *(or sometimes "coordinate patches") for the atlas $\mathcal{A}$.*

Exercise 3.4.
(a) If $\mathcal{A}$ and $\mathcal{B}$ are C^k atlases on a topological n-manifold, write $\mathcal{A} \sim \mathcal{B}$ if $\mathcal{A} \cup \mathcal{B}$ is also a C^k atlas. Prove that $\sim$ defines an equivalence relation on the set of all atlases.
(b) If $\mathcal{A} = \{(U_\alpha, \phi_\alpha)\}$ is a C^k atlas for M, let $\mathcal{A}_{max}$ denote the set of all pairs (U, ϕ) where $\phi\colon U \to \mathbb{R}^n$ is a homeomorphism from an open subset $U \subset M$ to an open subset $\phi(U) \subset \mathbb{R}^n$, and such that whenever $U \cap U_\alpha \neq \varnothing$ the map $\phi \circ \phi_\alpha^{-1}\colon \phi_\alpha(U \cap U_\alpha) \to \phi(U \cap U_\alpha)$ is C^k and has inverse which is C^k. Prove that $\mathcal{A}_{max}$ is an atlas containing $\mathcal{A}$, and is maximal in the sense that it contains every other atlas that contains $\mathcal{A}$. Deduce that if $\mathcal{A} \sim \mathcal{B}$ then $\mathcal{A}_{max} = \mathcal{B}_{max}$.

Definition 3.5. *A C^k-*differentiable structure *on a topological n-manifold is a maximal atlas $\mathcal{A}$ on M (*i.e., *an atlas such that, in the notation of Exercise 3.4(b), $\mathcal{A} = \mathcal{A}_{max}$). An n-*dimensional C^k manifold *is a topological n-manifold M equipped with a C^k-differentiable structure. A C^∞ manifold will also be called a* smooth manifold, *and a C^∞-differentiable structure will also be called a* smooth structure.

Remark 3.6. We will almost exclusively discuss *smooth* (*i.e.*, C^∞) manifolds in this course. This is partly justified by the fact that, for $1 \leq k < \infty$, any C^k manifold is C^k-diffeomorphic to a C^∞ manifold (there is a proof in Hirsch's book *Differential Topology*). On the other hand there is some real loss of generality in looking at C^∞ (or even just C^1) manifolds rather than just topological (C^0) manifolds, as there are topological manifolds which are not homeomorphic to any C^1 manifold. Examples of such are rather complicated—Kervaire constructed a 10-dimensional one in 1960, and the lowest dimension in which any occur is 4, where there are examples due to Freedman in the early 1980s.

Remark 3.7. The definition is that a smooth manifold is a certain kind of topological space equipped with a *maximal C^∞* atlas. A maximal atlas is a rather unwieldy object—except in trivial cases it will consist of uncountably many coordinate charts. But in view of Exercise 3.4 it is rarely if ever necessary to really work with a maximal atlas—you just have to specify *one* atlas (often with a small, finite number of charts), and then this canonically determines a maximal atlas by the construction in Exercise 3.4(b). One could equally well define a smooth manifold as a topological manifold equipped with an equivalence class of atlases, where the equivalence relation is the one from Exercise 3.4(a). One advantage of a maximal atlas is that "everything that could be a coordinate patch is," so that if you have to work in local coordinates you have a great variety of possible coordinate systems to work in and you can choose whichever works best for your purposes at the time.

Example 3.8. As the simplest possible example, we note that $\mathbb{R}^n$ is canonically a smooth manifold: take an atlas consisting of the single pair $(1_{\mathbb{R}^n}, \mathbb{R}^n)$ where $1_{\mathbb{R}^n}$ denotes the identity map. As noted in Remark 3.7 specifying this (very small!) atlas canonically determines a maximal atlas (*i.e.*, a differentiable structure).

Of course we could just as well have replaced $\mathbb{R}^n$ by any open subset U of $\mathbb{R}^n$, using the atlas $\{(1_U, U)\}$ to make U into a smooth manifold. More generally, if M is any smooth manifold with atlas $\{(\phi_\alpha, U_\alpha)\}$ and if $U \subset M$ is an open subset then we naturally get an atlas on U, namely $\{(\phi_\alpha|_{U\cap U_\alpha}, U \cap U_\alpha)\}$.

I promised at the outset that a smooth manifold would be the kind of space on which it is possible to do something resembling calculus. In particular if M is a smooth m-manifold it should be possible to speak of differentiable functions from M to $\mathbb{R}^n$, or vice versa, for any n (and, more generally, if M and N are two smooth manifolds we should be able to speak of differentiable functions from M to N). The principle is simple: one checks the differentiability of a function by using coordinate charts to turn the function into one whose domain and range are open subsets of Euclidean space, where we already have a notion of differentiability.

Definition 3.9. *Let M be an m-dimensional smooth manifold, with (maximal) atlas $\{(\phi_\alpha, U_\alpha)|\alpha \in A\}$.*

- *If $f\colon M \to \mathbb{R}^n$ is a continuous function, we say f is of class C^k, and write $f \in C^k(M, \mathbb{R}^n)$, if for every $\alpha \in A$ the function*
$$f \circ \phi_\alpha^{-1}\colon \phi_\alpha(U_\alpha) \to \mathbb{R}^n$$
is of class C^k (note that $f \circ \phi_\alpha^{-1}$ is a function from an open set in $\mathbb{R}^m$ to $\mathbb{R}^n$, so the notion of $f \circ \phi_\alpha^{-1}$ being of class C^k is well-defined from multivariable calculus).
- *If $V \subset \mathbb{R}^m$ is an open subset and $g\colon V \to M$ is a continuous function we say that g is of class C^k, and write $C^k(V, M)$, if for all $\alpha \in A$ the function*
$$\phi_\alpha \circ g\colon g^{-1}(U_\alpha) \to \mathbb{R}^m$$
is of class C^k.
- *Suppose that N is an n-dimensional smooth manifold, with (maximal) atlas $\{\psi_\beta, V_\beta)|\beta \in B\}$. If $f\colon M \to N$ is a continuous function, we say that f is of class C^k if, for all α,β such that $f(U_\alpha) \cap V_\beta \neq \emptyset$, the function*
$$\psi_\beta \circ f \circ \phi_\alpha^{-1}\colon \phi_\alpha(U_\alpha \cap f^{-1}(V_\beta)) \to \mathbb{R}^n$$
is of class C^k (as a function from an open subset of $\mathbb{R}^m$ to $\mathbb{R}^n$).

The appropriate notion of isomorphism of smooth manifolds is the following:

Definition 3.10. *Let M and N be C^k-manifolds. A C^k-*diffeomorphism *from M to N is a smooth, bijective map $f\colon M \to N$ such that f^{-1} is also smooth.*

As mentioned earlier, we will generally just consider the C^∞ case—as such a "diffeomorphism" will, unless otherwise indicated, mean a C^∞ diffeomorphism.

Of course, it would be a pain to actually check that Definition 3.9 is satisfied since maximal atlases are very large. But the following exercise shows that the C^k property can be checked more easily (and also implies that, viewing $\mathbb{R}^n$ as a smooth manifold, the third part of the above definition contains the first two as special cases). This exercise is intended in part to demonstrate the role of the assumption on the functions $\phi_\beta \circ \phi_\alpha^{-1}$ in the definition of an atlas.

Exercise 3.11. Let M and N be smooth manifolds, and let $f\colon M \to N$ be a continuous function. Prove that $f \in C^k(M, N)$ if and only if the following holds: For each $x \in M$, there exists a coordinate chart $\phi\colon U \to \mathbb{R}^m$ from the atlas for M and a coordinate chart $\psi\colon V \to \mathbb{R}^n$ from the atlas for N such that $x \in U$, $f(x) \in V$ and
$$\psi \circ f \circ \phi^{-1}\colon \phi(U \cap f^{-1}(V)) \to \mathbb{R}^n$$
is of class C^k.

Thus in practice to show that a map is C^k we just need to find collections of charts covering the manifolds in terms of which the map is a C^k map between Euclidean spaces, rather than checking the condition on the entire maximal atlas. Another way of saying this is that the two appearances of the word "(maximal)" in Definition 3.9 are unnecessary—we can just use any atlases (possibly quite small) to check the C^k condition.

Example 3.12. One can see that the n-dimensional sphere

$$S^n = \left\{ (x_0, x_1, \dots, x_n) \in \mathbb{R}^{n+1} \,\middle|\, \sum_{i=0}^{n} x_i^2 = 1 \right\}$$

is a smooth manifold by using stereographic projections. Of course the subspace topology on S^n induced by its inclusion into $\mathbb{R}^{n+1}$ makes S^n into a second-countable Hausdorff space. We construct a smooth atlas on S^n with two charts: define

$$\begin{aligned} U_- &= \{(x_0, \dots, x_n) \in S^n | x_0 \neq 1\} \\ U_+ &= \{(x_0, \dots, x_n) \in S^n | x_0 \neq -1\} \end{aligned}$$

In other words, U_- and U_+ are the complements of the north and south poles, respectively. Clearly $S^n = U_- \cup U_+$. Now define $\phi_- \colon U_- \to \mathbb{R}^n$ by

$$\phi_-(x_0, \dots, x_n) = \left(\frac{x_1}{1 - x_0}, \dots, \frac{x_n}{1 - x_0} \right)$$

and similarly define $\phi_+ \colon U_+ \to \mathbb{R}^n$ by

$$\phi_+(x_0, \dots, x_n) = \left(\frac{x_1}{1 + x_0}, \dots, \frac{x_n}{1 + x_0} \right)$$

So ϕ_- can be visualized as sending a point $p \in S^n \setminus \{\text{north pole}\}$ to the point of intersection between the hyperplane $\{x_0 = 0\}$ and the unique line through the north pole and p. It is clear from the formulas that ϕ_- and ϕ_+ are continuous. Both of them are in fact homeomorphisms to $\mathbb{R}^n$: one finds that the inverses $\phi_\pm^{-1}\mathbb{R}^n \to U_\pm$ are given by the formula

$$\phi_\pm^{-1}(y_1, \dots, y_n) = \left(\pm \frac{1 - \sum y_i^2}{1 + \sum y_i^2}, \frac{2y_1}{1 + \sum y_i^2}, \dots, \frac{2y_n}{1 + \sum y_i^2} \right).$$

Since the inverses are continuous the $\phi_\pm$ are indeed homeomorphisms to $\mathbb{R}^n$. What remains is to check that the "transition function" $\phi_+ \circ \phi_-^{-1} \colon \phi_-(U_+ \cap U_-) \to \phi_+(U_+ \cap U_-)$ is C^∞, and likewise that $\phi_- \circ \phi_+^{-1}$ is C^∞ (of course, the second of these is the inverse of the first). Now $U_+ \cap U_-$ is the complement of the two (north and south) poles of S^n, *i.e.* $U_+ \cap U_- = S^n \setminus \{(\pm 1, 0, \dots, 0)\}$. Now

$$\phi_+(1, 0, \dots, 0) = \phi_-(-1, 0, \dots, 0) = (0, \dots, 0),$$

so

$$\phi_-(U_+ \cap U_-) = \phi_+(U_+ \cap U_-) = \mathbb{R}^n \setminus \{(0, \dots, 0)\}.$$

For any $(y_1, \dots, y_n) \in \mathbb{R}^n \setminus \{(0, \dots, 0)\}$ we have

$$\begin{aligned} \phi_+ \circ \phi_-^{-1}(y_1, \dots, y_n) &= \phi_+ \left(\frac{\sum y_i^2 - 1}{\sum y_i^2 + 1}, \frac{2y_1}{\sum y_i^2 + 1}, \dots, \frac{2y_n}{\sum y_i^2 + 1} \right) \\ &= \left(\left(\frac{2 \sum y_i^2}{\sum y_i^2 + 1} \right)^{-1} \frac{2y_1}{\sum y_i^2 + 1}, \dots, \left(\frac{2 \sum y_i^2}{\sum y_i^2 + 1} \right)^{-1} \frac{2y_n}{\sum y_i^2 + 1} \right) \\ &= \left(\frac{y_1}{\sum y_i^2}, \dots, \frac{y_n}{\sum y_i^2} \right). \end{aligned}$$

Since this map is defined only on the complement of the origin, it is clearly C^∞ (the components are quotients of nonvanishing C^∞ functions), and its inverse (which as noted earlier is $\phi_- \circ \phi_+^{-1}$) is evidently C^∞ as well (actually if you look at the formula you see that it turns out that this map is equal to its own inverse). Thus we've shown that the transition functions for our atlas are C^∞, completing the proof that S^n is a smooth manifold.

Example 3.13. Recall that the n-dimensional *real projective space* $\mathbb{R}P^n$ is the space of lines through the origin in $\mathbb{R}^{n+1}$. This is given the structure of a (second-countable, Hausdorff) topological space by identifying it as

$$\mathbb{R}P^n = \frac{\mathbb{R}^{n+1} \setminus \{\vec{0}\}}{\vec{v} \sim \lambda\vec{v} \ \ \forall \vec{v} \in \mathbb{R}^{n+1} \setminus \{0\}, \lambda \in \mathbb{R} \setminus \{0\}}$$

and using the quotient topology. Thus a general element of $\mathbb{R}P^{n+1}$ can be written as an equivalence class $[x_0, \dots, x_n]$ for some $x_i \in \mathbb{R}$ with *not all* $x_i = 0$, and we have $[x_0 : \cdots : x_n] = [y_0 : \cdots : y_n]$ iff there is $\lambda \neq 0$ so that $y_i = \lambda x_i$ for all i. (The x_i are called "homogeneous coordinates.")

We now put a differentiable structure on $\mathbb{R}P^n$, making it a smooth n-manifold. For $i = 0, \dots, n$ let

$$U_i = \{[x_0, \dots, x_n] \in \mathbb{R}P^n | x_i \neq 0\}$$

(of course, the truth or falsehood of the statement that $x_i \neq 0$ is independent of which representative of the equivalence class we choose). The U_i are open sets (why?), and $\mathbb{R}P^n = \cup_{i=0}^n U_i$ since any element of $\mathbb{R}P^n$ has at least one of its homogeneous coordinates nonzero.

It shouldn't be too hard to convince yourself that each of the open sets U_i is homeomorphic to $\mathbb{R}^n$: for example for $i = n$, an element of $x \in U_n$ has form $[x_0 : \cdots : x_n]$ where $x_n \neq 0$, and since $x_n \neq 0$ we can simultaneously multiply all of the x_i by $\frac{1}{x_n}$—this doesn't change the equivalence class, but changes the last homogeneous coordinate to 1. Thus U_n can be identified with the set of tuples $(x_0, \dots, x_{n-1}, 1)$, which is equivalent to $\mathbb{R}^n$.

To make the discussion in the previous paragraph more precise, we introduce charts $\phi_i \colon U_i \to \mathbb{R}^n$. Namely, define

$$\phi_i \colon U_i \to \mathbb{R}^n$$
$$\phi_i([x_0 : \cdots : x_n]) = \left(\frac{x_0}{x_i}, \dots, \frac{x_{i-1}}{x_i}, \frac{x_{i+1}}{x_i}, \dots, \frac{x_n}{x_i}\right).$$

This map is certainly well-defined, since multiplying all entries of $(x_0, \dots, x_n)$ by the same scalar λ does not affect the ratios x_j/x_i. Moreover we see that ϕ_i is bijective, with inverse given by

$$\phi_i^{-1}(y_0, \dots, y_{i-1}, y_{i+1}, \dots, y_n) = [y_0 : \cdots, y_{i-1} : 1 : y_{i+1} : \cdots : y_n].$$

Both ϕ_i and ϕ_i^{-1} are continuous—of course to see this one has to think a little bit about the quotient topology, but it's not hard and is left to you.

So we have a covering $\mathbb{R}P^n = \cup_{i=0}^n U_i$ by open sets with homeomorphisms $\phi_i \colon U_i \to \mathbb{R}^n$. It remains to check that the transition functions $\phi_i \circ \phi_j^{-1} \colon \phi_j(U_i \cap U_j) \to \phi_i(U_i \cap U_j)$ are smooth. This follows quickly from the formulas that we've already written down: assuming that $i < j$

$$\begin{aligned}\phi_i \circ \phi_j^{-1}(y_0, \dots, y_{j-1}, y_{j+1}, \dots, y_n) &= \phi_i([y_0 : \cdots : y_{j-1} : 1 : y_{j+1} : \cdots : n]) \\ &= \left(\frac{y_0}{y_i}, \dots, \frac{y_{i-1}}{y_i}, \frac{y_{i+1}}{y_i}, \dots, \frac{y_{j-1}}{y_i}, \frac{1}{y_i}, \frac{y_{j+1}}{y_i}, \dots, \frac{y_n}{y_i}\right).\end{aligned}$$

Of course the case that $i > j$ differs from this only in the ordering of i and j in the above formula. Now on the open subset $\phi_j(U_i \cap U_j) \subset \mathbb{R}^n$ we will have $y_i \neq 0$, so $\phi_i \circ \phi_j^{-1}$ is indeed smooth on $\phi_j(U_i \cap U_j)$, as required. Thus $\{(\phi_i, U_i) : i = 0, \dots, n\}$ forms a C^∞ atlas for $\mathbb{R}P^n$, making $\mathbb{R}P^n$ into a smooth manifold.

Fairly easy modifications of this argument show that the complex projective space $\mathbb{C}P^n$ is a smooth $2n$-manifold, and that the quaternionic projective space $\mathbb{H}P^n$ is a smooth $4n$-manifold.

Exercise 3.14. Recall that another way of describing $\mathbb{R}P^n$ is as a quotient of S^n by the equivalence relation which identifies any $x \in S^n \subset \mathbb{R}^{n+1}$ with $-x$. Thus we have a quotient projection $\pi \colon S^n \to \mathbb{R}P^n$. Prove that $\pi \in C^\infty(S^n, \mathbb{R}P^n)$.

Exercise 3.15. (a) If M and N are smooth manifolds, construct a C^∞ atlas on the product $M\times N$ (thus $M\times N$ has the structure of a smooth manifold).

(b) Let M be a Hausdorff space, and suppose that we can write $M = U\cup V$ where U and V are open sets, and both U and V are smooth manifolds. Since $U\cap V$ is an open subset of U, it inherits a differentiable structure from U; likewise $U\cap V$ inherits a differentiable structure from V. Assume that these two differentiable structures on $U\cap V$ are the same. Prove that one can then construct a smooth structure on M such that the inclusions $U\to M$ and $V\to M$ are both smooth maps.

(c) Prove that for any g the compact surface of genus g (and no boundary) can be given the structure of a smooth manifold (Hint: The case $g = 0$ is covered by Example 3.12, and $g = 1$ follows from Example3.12 and part (a). Now repeatedly use (b) together with the fact that an open subset of a smooth manifold is naturally a smooth manifold.)

Remark 3.16. In our examples we've brushed over the question of whether the smooth structures on these spaces are unique. This is an important but difficult question; a fair amount is now known, but the proofs are generally beyond the scope of this course. It's known that in any dimension $n\leq 3$, every topological n-manifold has a unique smooth structure; in particular the smooth structures on surfaces from the exercise above are the only possible ones. Things become more complicated beginning in (and especially in) dimension 4: in fact there are uncountably many distinct smooth structures on $\mathbb{R}^4$, and there are many compact 4-manifolds with infinitely many smooth structures, and none that are currently known to have just one smooth structure (though as mentioned earlier there are some topological 4-manifolds with *no* smooth structures). For spheres, once $n\geq 7$ there is typically more than one smooth structure on S^n; the first "exotic" structure on S^7 was a big surprise when it was discovered by Milnor in 1956. It's still a major open question whether there are any smooth structures on S^4 other than the standard one.

We now record a result asserting the existence of partitions of unity subordinate to covers of smooth manifolds:

Theorem 3.17. *Let M be a smooth manifold and let $\{V_\alpha|\alpha\in A\}$ be a collection of open subsets of M with $\cup_{\alpha\in A}V_\alpha = M$. Then there is a smooth partition of unity on M subordinate to the cover $\{V_\alpha\}$,* i.e.*, a collection $\{\chi_\alpha|\alpha\in A\}$ where*

- *Each $\chi_\alpha\in C^\infty(M)$, with $0\leq\chi_\alpha(x)\leq 1$ for all $x\in M$*
- *For all α, $supp(\chi_\alpha)\subset V_\alpha$*
- *For any $x\in M$ there is a neighborhood O_x of x such that $O_x\cap supp(\chi_\alpha)=\varnothing$ for all but finitely many α*
- *$\sum_\alpha\chi_\alpha = 1$*

Proof. The special case in which M is an open subset of $\mathbb{R}^n$ was proven as Theorem 2.6. That proof carries over directly to the more general case now that we have the appropriate definitions. Indeed, a smooth manifold M is by definition second-countable and Hausdorff, and is certainly locally compact (any point has a neighborhood whose closure is homeomorphic to a closed ball in $\mathbb{R}^n$ and so is compact), so Lemma 2.9 applies to produce a sequence of compact sets K_i and open sets H_i. These sets can then be used just as they are used in the proof of Theorem 2.6. Basically all that needs to be changed is the first paragraph of that proof: if $x\in K_i$ we can find a neighborhood of x having the form $\phi^{-1}(B_{2r_x}(\phi(x)))$ which is contained in $V_{\alpha_x}\cap W_i$ for some α_x, where $\phi\colon U\to\mathbb{R}^n$ is some chart (depending on x) whose domain U contains x. The sets $\phi^{-1}(B_{r_x}(x))$ then cover K_i, and this cover has a finite subcover. Aggregating these finite subcovers gives a countable sequence $\{B_k\}$ of open sets covering M; the B_k are preimages of balls in $\mathbb{R}^n$ by local charts ϕ, and where $\tilde{B}_k$ is the preimage of the ball with the same center and twice the radius we will have $B_k\subset V_{\alpha_k}\cap W_{i_k}$ for appropriate α_k, i_k. Moreover there is a smooth function ψ_k supported in $\tilde{B}_k$ and identically equal to one on B_k—just precompose an appropriate smooth function on $\mathbb{R}^n$ given by Corollary 2.5 with ϕ^{-1}. The proof of Theorem 2.6 then applies verbatim. □

Partitions of unity are very useful in the study of smooth manifolds. For a brief indication of why, consider the case in which the cover $\{V_\alpha\}$ consists of the domains of coordinate charts $\phi_\alpha\colon V_\alpha\to\mathbb{R}^n$ (of course, by definition, any smooth manifold admits such a cover). If $f \in C^\infty(M)$, then we can write

$$f = \left(\sum_\alpha \chi_\alpha\right) f = \sum_\alpha (\chi_\alpha f).$$

Now for any α the function $\chi_\alpha f$ is supported in the set V_α, which is identified by ϕ_α with an open subset in $\mathbb{R}^n$. So we can hope to analyze f by decomposing it as a sum of smooth functions $\chi_\alpha f$, where each of these smooth functions can (at least individually) be treated as though it were just a compactly supported smooth function on $\mathbb{R}^n$. To get slightly ahead of myself, the same applies when f is, instead of a smooth function, a differential form.

3.1. **Tangent spaces.** If M is a smooth manifold and $m \in M$, we will define a vector space $T_m M$ called the *tangent space to M at m*. As suggested at the start of these notes, there are various ways of trying to do this, any of which can be considered to be inspired by the special case in which M is an open subset of $\mathbb{R}^n$. For instance we could define a tangent vector v at m to be an equivalence class $[\gamma]$ where $\gamma\colon (-\epsilon, \epsilon) \to M$ is a C^∞ map from an open interval around 0 to M with $\gamma(0) = m$, with two curves γ_1, γ_2 considered to be equivalent if $\frac{d}{dt}(\phi_\alpha \circ \gamma_1)(0) = \frac{d}{dt}(\phi_\alpha \circ \gamma_1)(0)$ (as vectors in $\mathbb{R}^n$) for one (and hence every—why?) chart $\phi_\alpha\colon U_\alpha \to \mathbb{R}^n$ whose domain contains m. However, for definiteness we will adopt the third interpretation from the start of the notes: a tangent vector at m will be, by definition, a derivation from the algebra of germs of smooth functions defined near m to $\mathbb{R}$.

So just as earlier we consider pairs (f, V) where V is an open neighborhood of m in M and $f\colon V \to \mathbb{R}$ is C^∞ (this notion is well-defined since V, being an open set in a smooth manifold, is itself a smooth manifold, and we have defined the space of C^∞ functions on a smooth manifold). Say that $(f_1, V_1) \sim (f_2, V_2)$ if and only if there is an open set W with $m \in W \subset V_1 \cap V_2$ and $f_1|_W = f_2|_W$. Let $\mathcal{O}_m$ denote the set of equivalence classes; this inherits addition, multiplication, and scalar multiplication from $C^\infty(M)$ (for example, $[f, V][g, W] = [fg, V \cap W]$).

Definition 3.18. *$T_m M$ is defined as the space of derivations $v\colon \mathcal{O}_m \to \mathbb{R}$,* i.e.*, maps v such that*
- *$v(cf + g) = cv(f) + v(g)$ if $c \in \mathbb{R}$ and $f, g \in \mathcal{O}_p$*
- *$v(fg) = f(m)v(g) + g(m)v(f)$ if $f, g \in \mathcal{O}_m$*

As indicated in the above definition we will often abuse notation slightly by just writing f for $[f, V]$. Compatibly with this abuse of notation, if $\phi\colon M \to N$ is a smooth map where N is another smooth manifold and $m \in M$, if we write f for an element $[f, V] \in \mathcal{O}_{\phi(m)}$ (thus f is a function defined on a neighborhood of $f(m)$ in N), then we will write $f \circ \phi$ for the element $[f \circ \phi, \phi^{-1}(V)] \in \mathcal{O}_m$. These sorts of abuse of notation are justified by the fact that replacing the open set V by a different neighborhood of $\phi(m)$ will not change either the element $[f, V]$ (denoted f) or the element $[f \circ \phi, \phi^{-1}(V)]$ (denoted $f \circ \phi$).

We record here the fact that, if $U \subset M$ is an open subset and $m \in U$, there is a canonical identification of $T_m U$ with $T_m M$ (convince yourself of this if it's not obvious). Also, in case U is an open subset of $\mathbb{R}^n$, our definition coincides with the one from the start of these notes.

Definition 3.19. *If $\phi\colon M \to N$ is a smooth map between smooth manifolds and if $m \in M$, the* derivative of ϕ at m *(sometimes called the* linearization of ϕ at m *is the map*

$$\phi_*\colon T_m M \to T_{\phi(m)} N$$

defined by

$$(\phi_*(v))(f) = v(f \circ \phi)$$

whenever $f \in \mathcal{O}_{\phi(m)}$ and $v \in T_m M$.

Sometimes it's helpful to indicate m within the notation for ϕ_*, in which case we'll write $(\phi_*)_m$. One also sees the notation $d\phi$ or $d_m\phi$ used to denote what we have called ϕ_*.

Proposition 3.20. *Where 1_M is the identity map then for all $m \in M$, $(1_M)_*\colon T_m M \to T_m M$ is the identity map. Also, if $\phi\colon M \to N$ and $\psi\colon N \to P$ are smooth maps then*

$$(\psi \circ \phi)_* = \psi_* \circ \phi_*$$

Proof. The first statement (about the identity) is obvious from the definition. For the second, we have, if $f \in \mathcal{O}_{\psi\circ\phi(m)}$,

$$((\psi \circ \phi)_* v)(f) = v(f \circ (\psi \circ \phi)) = v((f \circ \psi) \circ \phi) = (\phi_* v)(f \circ \psi) = (\psi_* \phi_* v)(f).$$

□

Corollary 3.21. *If $m \in M$ where M is a smooth n-manifold, then* $\dim T_m M = n$.

Proof. We can choose a coordinate chart $\phi\colon U \to \phi(U)$ where U is an open neighborhood of m. As noted earlier we have $T_m M = T_m U$. By Proposition 3.20, $(\phi^{-1})_* \circ \phi_* = (\phi^{-1} \circ \phi)_*$ is the identity map from $T_m U = T_m M$ to itself, and $\phi_* \circ (\phi^{-1})_* = (\phi \circ \phi^{-1})_*$ is the identity map from $T_{\phi(m)}\phi(U)$ to itself. Thus ϕ_* is an isomorphism of vector spaces from $T_m M$ to $T_{\phi(m)}\phi(U)$, with inverse $(\phi^{-1})_*$. We showed in Section 1 that, since $\phi(U)$ is an open subset of $\mathbb{R}^n$, $\dim T_{\phi(m)}\phi(U) = n$, so the conclusion follows. □

Expanding a bit on the above proof, recall that we showed that $T_{\phi(m)}\phi(U)$ consists precisely of maps $\mathcal{O}_{\phi(m)} \to \mathbb{R}$ taking the form $g \mapsto \sum_{i=1}^n v_i \frac{\partial g}{\partial x_i}|_{\phi(m)}$. So since $(\phi^{-1})_*$ is an isomorphism, we conclude that, in the presence of a chosen coordinate chart $\phi\colon U \to \mathbb{R}^n$ around m, a general element $v \in T_m M$ will be given by the formula

$$v(f) = \sum_{i=1}^n v_i \frac{\partial}{\partial x_i}(f \circ \phi^{-1})|_{\phi(m)}.$$

When this is the case, we will say something along the lines of, "v is given in the coordinate chart ϕ by $v = \sum v_i \frac{\partial}{\partial x_i}$." Of course, the coefficients v_i will depend on the coordinate chart, not just on the tangent vector v.

Exercise 3.22. Let $\phi, \psi\colon U \to \mathbb{R}^n$ be two coordinate charts where U is an open subset of a smooth manifold M, and let $m \in U$. If v is given in the coordinate chart ϕ by $v = \sum v_i \frac{\partial}{\partial x_i}$, and is given in the coordinate chart ψ by $v = \sum w_i \frac{\partial}{\partial y_i}$, find, with proof, an expression for the w_i in terms of the v_i and the maps $\phi \circ \psi^{-1}$ and/or $\psi \circ \phi^{-1}$.

So if M is a smooth n-manifold, we have associated to every point $m \in M$ an n-dimensional vector space $T_m M$. A diffeomorphism $\phi\colon M \to M'$ induces an isomorphism of vector spaces $\phi_*\colon T_m M \to T_{\phi(m)} M'$. However there is (in general) no canonical way of identifying $T_{m_1} M$ with $T_{m_2} M$ for distinct point $m_1, m_2 \in M$ (of course, since the two vector spaces have the same dimension, they are isomorphic as vector spaces, just not canonically so).

Relatedly, while choosing the point $m \in M$ canonically determines the n-dimensional vector space $T_m M$, it does not canonically determine a basis for this vector space. One way of choosing a basis for $T_m M$ is suggested above: choose a local coordinate chart $\phi\colon U \to \mathbb{R}^n$ around U; then a basis is given by the derivations $f \mapsto \frac{\partial}{\partial x_i}(f \circ \phi^{-1})(p)$ for $i = 1, \ldots, n$ (the members of this basis are typically denoted by $\frac{\partial}{\partial x_i}$. Different choices of coordinate chart of course give rise to different bases; the relationship between the bases is determined by Exercise 3.22.

The *tangent bundle* of a smooth manifold is, as a set, defined to be the union

$$TM = \cup_{m \in M} \{m\} \times M.$$

For any subset $S \in M$ (typically S will be open or closed) we define the "restriction of the tangent bundle to S" as

$$TM|_S = \cup_{m \in S} \{m\} \times T_m M.$$

Given a coordinate chart $\phi\colon U \to \mathbb{R}^n$ where $U \subset M$ is open, we have a bijection $\Phi\colon TM|_U \to \phi(U) \times \mathbb{R}^n$ given by

$$\Phi\left(m, \sum v_i \frac{\partial}{\partial x_i}\right) = (\phi(m), v_1, \ldots, v_n).$$

We can then define a topology on TM by requiring that each of these bijections be homeomorphisms—more precisely, we take as a base for this topology the collection of subsets of the form $\Phi^{-1}(V)$ where $\Phi\colon TM|_U \to \phi(U) \times \mathbb{R}^n$ is a map as above constructed from a coordinate chart ϕ and $V \subset \phi(U) \times \mathbb{R}^n$ is open.

The various homeomorphisms $\Phi\colon TM|_U \to \phi(U)\times\mathbb{R}^n$ associated to coordinate charts $\phi\colon U\to\phi(U)$ in fact form a C^∞ atlas for TM. Indeed the domains $TM|_U$ certainly cover TM (since M is covered by coordinate charts) and so we just need to check that the transition functions are smooth. This latter fact follows from Exercise 3.22. Indeed, if $\phi_\alpha\colon U_\alpha\to\mathbb{R}^n$ and $\phi_\beta\colon U_\beta\colon U_\beta\to\mathbb{R}^n$ are two coordinate charts, then it should follow from your computation in Exercise 3.22 that the transition function

$$\Phi_\beta\circ\Phi_\alpha^{-1}\colon\ \phi_\alpha(U_\alpha\cap U_\beta)\times\mathbb{R}^n\to\phi_\beta(U_\alpha\cap U_\beta)\times\mathbb{R}^n$$

is given by

$$\Phi_\beta\circ\Phi_\alpha^{-1}(x,\vec{v}) = (\phi_\beta\circ\phi_\alpha^{-1}(x), g_{\alpha\beta}(x)\vec{v}) \tag{4}$$

where $g_{\alpha\beta}$ is a certain smooth function which takes values in the group of invertible $n\times n$ matrices. Thus the transition functions are smooth, and so determine a smooth manifold structure on TM.

Of course, we have a projection $\pi\colon TM\to M$ which sends (m,v) to m. In terms of the local coordinate charts Φ on TM and ϕ on M, π just acts by the projection of $\phi(U)\times\mathbb{R}^n$ onto its first factor; thus π is a smooth map.

Summing up, out of an n-dimensional smooth manifold M we have constructed a $2n$-dimensional smooth manifold TM, equipped with a projection $\pi\colon TM\to M$. The "fibers" $\pi^{-1}(\{m\})$ of π are canonically identified with the tangent spaces T_mM, and thus are n-dimensional vector spaces. Moreover there is an atlas on TM such that the transition functions respect the vector space structures on the fibers in the sense that they are given by a formula of the shape (4) where each $g_{\alpha\beta}(x)$ is a linear map. TM is thus an example of what is called a *vector bundle*; we will see more examples of vector bundles as the course proceeds.

3.2. **Vector fields.** Consistently with what was done in Section 1, we make the following definition:

Definition 3.23. *Let M be a smooth manifold and $U\subset M$ an open subset. A* vector field *on U is a derivation $X\colon C^\infty(U)\to C^\infty(U)$ (*i.e., *X obeys $X(cf+g)=cXf+Xg$ and $X(fg)=fXg+gXf$ if $f,g\in C^\infty(U)$, $c\in\mathbb{R}$). We denote the space of vector fields on U by $\mathcal{X}(U)$.*

Just as in Section 1, we can scalar multiply, add, and take the commutators of derivations from $C^\infty(U)$ to itself, so $\mathcal{X}(U)$ naturally has the structure of a Lie algebra.

A vector field on U should have another interpretation as a "smoothly-varying" choice of tangent vector at m for each $m\in M$. We now lay out how this works. For $U\subset M$ we have a (restricted) tangent bundle $\pi\colon TM|_U\to U$.

Definition 3.24. *A* smooth section *of TM over U is a smooth map $s\colon U\to TM|_U$ such that $\pi\circ s$ is the identity. We write $\Gamma(U,TM)$ for the space of smooth sections of TM over U.*

In other words, $s(m)\in T_mU$ for all $p\in U$; the notion that the tangent vectors should vary smoothly is encoded in the requirement that s should be a smooth map. Since T_mU is a vector space, we get vector space operations on $\Gamma(U,TM)$ defined by $(cs)(m)=c(s(m))$ and $(s_1+s_2)(m)=s_1(m)+s_2(m)$ (there's something to show here, namely that for instance the sum of two smooth sections is still smooth, but it's not hard to check this). One important example of a section of TM (or more generally of any vector bundle) is the *zero section*, defined by $s(m)=0\in T_mM$ for all p. (To see that this is smooth, just note that in the local coordinates $\phi(U)\times\mathbb{R}^n\subset\mathbb{R}^{2n}$ described earlier the map is given by $x\mapsto(x,0)$ which is obviously a smooth map from $\mathbb{R}^n$ to $\mathbb{R}^{2n}$).

Recall Exercise 2.11, to which the following gives a solution:

Proposition 3.25. *Let M be a smooth manifold, $U\subset M$ open, $m\in U$, and $X\in\mathcal{X}(U)$. Then the following prescription uniquely specifies an element $X_m\in T_mM$. For any $[f,V]\in\mathcal{O}_m$, choose a $\tilde{f}\in C^\infty(U)$ such that $[\tilde{f},U]=[f,V]$, and define $X_m([f,V])=(X\tilde{f})(m)$.*

Proof. First of all we need to show that for any $[f,V]\in\mathcal{O}_m$ (in other words, V is an open set around m and f is a smooth function on V) there is a smooth function $\tilde{f}$ defined throughout U and coinciding on with f on some neighborhood G of m. To see this, note that we can find a coordinate chart $\phi\colon W\to\mathbb{R}^n$ around m and $r>0$ so that $\overline{\phi^{-1}(B_{2r}(\phi(m)))}\subset V$. Take a partition of unity $\{\chi_1,\chi_2\}$ subordinate to the open cover

$\{\phi^{-1}(B_{2r}(\phi(m))), M \setminus \overline{\phi^{-1}(B_r(\phi(m)))}\}$ of M. Then let $\tilde{f} = \chi_1 f$; initially this function is only defined on V, but since it has support contained in a compact subset of V we may extend it by zero to obtain a smooth function on all of M. Since $\chi_1 + \chi_2 = 1$ and χ_2 vanishes on $\phi^{-1}(B_r(\phi(m)))$, $\tilde{f}$ coincides with f on $\phi^{-1}(B_r(\phi(m)))$, as desired.

We now show that the value $(X\tilde{f})(m)$ is independent of the choice of $\tilde{f}$ with $[\tilde{f}, U] = [f, V]$. If $\tilde{g}$ is another such choice, there is a neighborhood W of m such that $\tilde{f}|_W = \tilde{g}|_W$. Let O be a neighborhood of m such that $m \in \overline{O} \subset W$ (for instance take O to be the preimage of a small ball in a coordinate chart, as in the previous paragraph). Just as in the previous paragraph we can find a smooth function $\chi\colon M \to \mathbb{R}$ such that $\chi|_O = 1$ and $supp(\chi) \subset W$. Let $\beta = 1 - \chi$, so β vanishes identically on the neighborhood O of m and is equal to 1 outside W. Hence

$$(1 - \beta^2)\tilde{f} = (1 - \beta^2)\tilde{g}$$

(both sides are zero everywhere that $\tilde{f} \neq \tilde{g}$). On the other hand

$$\left(X(\beta^2 \tilde{f})\right)(m) = \beta(m)\left(X(\beta\tilde{f})\right)(m) + \beta(m)\tilde{f}(m)\,(X\beta)\,(m) = 0$$

and similarly

$$\left(X(\beta^2\tilde{g})\right)(m) = 0.$$

Hence

$$\begin{aligned}(X\tilde{f})(m) &= \left(X(\beta^2 \tilde{f})\right)(m) + \left(X((1-\beta^2)\tilde{f})\right)(m)\\ &= \left(X((1-\beta^2)\tilde{f})\right)(m) = \left(X((1-\beta^2)\tilde{g})\right)(m)\\ &= \left(X(\beta^2\tilde{g})\right)(m) + \left(X((1-\beta^2)\tilde{g})\right)(m) = (X\tilde{g})(m).\end{aligned}$$

This confirms that the prescription of the proposition gives a well-defined map $X_m\colon \mathcal{O}_m \to \mathbb{R}$. It remains to check that X_m is a derivation. But this follows easily from the derivation property for X. Given $[f, V], [g, W] \in \mathcal{O}_m$, if we use $\tilde{f} \in C^\infty(U)$ to compute $X_m[f, V] = (X\tilde{f})(m)$ and $\tilde{g} \in C^\infty(U)$ to compute $X_m[g, V] = (X\tilde{g})(m)$ then we can use $\widetilde{fg} = \tilde{f}\tilde{g}$ to compute $X_m([f, V][g, W])$ (of course we could make other choices for $\widetilde{fg}$, but the start of the proof ensures that this would result in the same value for $X_m([f, V][g, W])$). Then the derivation property for X shows

$$\begin{aligned}X_m([f, V][g, W]) = \left(X(\widetilde{fg})\right)(m) &= f(m)(X\tilde{g})(m) + g(m)(X\tilde{f})(m)\\ &= f(m)X_m[g, W] + g(m)X_m[f, V].\end{aligned}$$

$\mathbb{R}$-linearity is proved in essentially the same way, completing the proof that $X_m \in T_mM$.

□

We now show that giving a vector field (in the sense of a derivation on the space of smooth functions) is exactly the same as giving a smooth section of the tangent bundle.

Theorem 3.26. *Let U be an open subset of the smooth manifold M. A bijection $\mathcal{F}\colon \mathcal{X}(U) \to \Gamma(U, TM)$ may be defined as follows. For $X \in \mathcal{X}(U)$, set $\mathcal{F}(X)$ equal to the map $s_X\colon M \to TM$ defined by $s_X(m) = X_m$ (where X_m is given by Proposition 3.25).*

Proof. First we need to show that $\mathcal{F}$ is well-defined—we certainly have a well-defined function $s_X\colon M \to TM$ for any $X \in \mathcal{X}(U)$, and s_X is a section in the sense that $\pi \circ s_X = 1_M$, but we also need to check that s_X is smooth in order for $\mathcal{F}$ to take values in the space $\Gamma(U, TM)$ of smooth sections.

To see this, note first of all that a function f between two smooth manifolds is smooth if and only if the domain can be covered by open sets to each of which f restricts as a smooth function. If $m \in M$, let $\phi\colon V \to \mathbb{R}^n$ be a coordinate chart with $m \in V \subset U$, and for $r > 0$ small enough that $B_{2r}(\phi(m)) \subset \phi(V)$ let $W_m = \phi^{-1}(B_r(\phi(m)))$. We will show that $s_X|_{W_m}$ is smooth, which suffices since any point in M has a neighborhood of the form W_m.

In this direction, let $\chi\colon M \to \mathbb{R}$ be a smooth function with $\chi|_{\overline{W_m}} = 1$ and $supp(\chi) \subset V$. For any $q \in W_m$ and $f \in \mathcal{O}_q$ we have

$$(s_X(q))(f) = X_q(f) = X_q(\chi f)$$

since f and χf coincide on a neighborhood (namely W_m) of q.

Now for each $j = 1, \dots, n$ write $g_j = (x_j \circ \psi) \cdot \chi \in C^\infty(M)$. Then on W_m, g coincides with the jth coordinate of the chart $\psi|_{W_m} \colon W_m \to \mathbb{R}^n$. We know that, for each $q \in W_m$, since $X_q \in T_q M$ we can express X_q in the coordinate chart ψ as $X_q = \sum_i v_i(q) \frac{\partial}{\partial x_i}|_q$ for some $v_i(q) \in \mathbb{R}$. Evaluating on the functions g_j we see that, for each j,

$$v_j(q) = (X g_j)(q).$$

Thus the functions $v_j \colon W_m \to \mathbb{R}$ are each smooth. Now in terms of the local coordinates for the tangent bundle described at the end of the previous subsection, the map s_X is given within W_m by the formula (where $x \in \psi(W_m) \subset \mathbb{R}^n$)

$$x \mapsto \left(x, v_1(\psi^{-1}(x)), \dots, v_n(\psi^{-1}(x))\right).$$

This map is smooth since the v_j are smooth. Thus $s_X|_{W_m}$ is smooth, and so s_X is smooth since U can be covered by open sets of the form W_m.

Now that we have shown the map $\mathcal{F} \colon \mathcal{X}(U) \to \Gamma(U, TM)$ to be well-defined, we show that it is bijective. Suppose that $X, Y \in \mathcal{X}(U)$ are two distinct vector fields on U. Then there is $f \in C^\infty(U)$ and $m \in U$ such that $(Xf)(m) \neq (Yf)(m)$. But then $[f, U]$ is a well-defined element of $\mathcal{O}_m$ with $X_m([f, U]) \neq Y_m([f, U])$, and thus $X_m \neq Y_m$, *i.e.* $s_X(m) \neq s_Y(m)$. Thus $\mathcal{F}$ is injective.

Finally suppose that $s \in \Gamma(U, TM)$; we must find $X \in \mathcal{X}(U)$ so that $s_X = s$. If $f \in C^\infty(U)$ then for all m we have an element $[f, U] \in \mathcal{O}_m$ and so a real number $(s(m))([f, U])$. This determines a function $Xf \colon U \to \mathbb{R}$ by the formula $(Xf)(m) = (s(m))([f, U])$. The derivation properties $X(cf + g) = cXf + Xg$ and $X(fg) = fXg + gXf$ follow directly from the fact that each $s(m)$ is a derivation from $\mathcal{O}_m$ to $\mathbb{R}$; however we still need to check that $Xf \in C^\infty(U)$ for any $f \in C^\infty(U)$. In a local coordinate chart $\psi \colon V \to \mathbb{R}^n$, the tangent vectors $s(m)$ for $m \in V$ are represented as $s(m) = \sum v_i(m) \frac{\partial}{\partial x_i}$, where the functions v_i are C^∞ by the fact that s is a smooth map. But then $Xf|_V = \sum v_i \frac{\partial f}{\partial x_i}$, which is a smooth function. Thus Xf restricts to each coordinate chart as a smooth function, and so is smooth. It is clear from the definition that $s_X = s$. □

So we have two equivalent characterizations of vector fields on M: as derivations $C^\infty \to \mathbb{C}^\infty$, and as smooth sections $M \to TM$ (which in coordinate charts can be locally expressed in the form $\sum v_i \frac{\partial}{\partial x_i}$ for suitable smooth functions v_i). Both characterizations are often useful.

4. Differential forms

As the title of the course textbook suggests, a very important role will be played in the rest of the course by what are called the *differential forms* on a smooth manifold. If M is a smooth n-manifold, we will develop the notion of a "p-form" on M for $p = 0, 1, \dots, n$ (and also for $p > n$, but for algebraic reasons it turns out that the only p-forms with $p > n$ will be zero). These p-forms will form a vector space $\Omega^p(M)$, and we will have a very important map d, called the *exterior derivative*, which maps the space of all differential forms to itself and restricts for each p to a map $d \colon \Omega^p(M) \to \Omega^{p+1}(M)$.

To ease into this, let's start with $p = 0$ and $p = 1$.

Definition 4.1. *A 0-form on M is a smooth function $f \colon M \to \mathbb{R}$. In other words $\Omega^0(M) = C^\infty(M)$.*

The case of 1-forms is a bit more interesting. First we introduce the notion of the *cotangent space*:

Definition 4.2.
- *If M is a smooth manifold and $m \in M$, the* cotangent space at m, *denoted by $T_m^* M$, is the dual space to the tangent space $T_m M$.*
- *The* cotangent bundle *of M is*

$$T^* M = \cup_{m \in M} \{m\} \times T_m^* M.$$

In other words, $T_m^* M$ consists of linear functionals $\alpha \colon T_m M \to \mathbb{R}$. Since a vector space and its dual have the same dimension, if M is an n-manifold then $\dim T_p^* M = n$ for all $m \in M$.

Definition 4.2 identifies the cotangent bundle $T^* M$ as a set. One can equip it with a topology and then with a smooth manifold structure, in such a way that the projection $\pi \colon T^* M \to M$ (sending (m, α) to m if $\alpha \in T_m^* M$)

makes T^*M into a vector bundle, just like the situation with the tangent bundle. At least for now we won't really need to use this fact, but note that we have (at least at a set-theoretic level) the notion of a *section* $s\colon M \to T^*M$, *i.e.* a function $s\colon M \to T^*M$ such that $\pi \circ s = 1_M$. A section $s\colon M \to T^*M$ associates to each $m \in M$ an element $s_m \in T^*_pM$.

Definition 4.3. *A* differential 1-form *on a smooth manifold M is a section $\alpha\colon M \to T^*M$ which satisfies the following smoothness property: Whenever $X \in \mathcal{X}(M)$ is a vector field on M, the function*

$$\alpha(X)\colon m \mapsto \alpha_m(X_m)$$

is a C^∞ function on M. We denote by $\Omega^1(M)$ the vector space of differential 1-forms.

To unpack the above, note that the section α of the *cotangent* bundle determines *co*vectors $\alpha_m \in T^*_mM$ for all m, while the vector field X (which by Theorem 3.26) is equivalent to a section of the *tangent* bundle, determines for each m a tangent vector $X_m \in T_mM$. Hence we can evaluate $\alpha_m(X_m)$, and the smoothness requirement on α is that (as long as X is smooth) the result of this evaluation varies smoothly with m. If we had gone ahead and put a smooth manifold structure on T^*M it turns out that this would be equivalent to requiring $\alpha\colon M \to T^*M$ to be a smooth map.

As mentioned earlier, for all p we will define a map $d\colon \Omega^p(M) \to \Omega^{p+1}(M)$. I can now fulfill this promise for $p = 0$. Actually if one thinks of tangent vectors as derivations the definition may seem strangely simple:

To any $f \in \Omega^0(M)$, *i.e.*, any smooth function f, we are to associate a section $df\colon M \to T^*M$. In other words for each m we should obtain $(df)_m\colon T_mM \to \mathbb{R}$. Well, bearing in mind that an element of T_mM is a derivation from functions defined near m to $\mathbb{R}$, we use the formula

$$(df)_m(v) = v(f) \quad \text{if } v \in T_mM. \tag{5}$$

Suppose now that $\phi\colon U \to \mathbb{R}^n$ is a coordinate chart, where $U \subset M$ is open. Now U is a smooth manifold in its own right, so we can consider $\Omega^1(U)$. The coordinate chart ϕ distinguishes some special smooth functions on U, namely the *coordinate functions* $x_1, \ldots, x_n$ (perhaps we should really write $x_1 \circ \phi, \ldots, x_n \circ \phi$, or we could just agree that the decomposition of ϕ into coordinates is given by $\phi(m) = (x_1(m), \ldots, x_n(m))$). Since the x_i are smooth functions (*i.e.*, 0-forms) on U, we obtain 1-*forms* $dx_1, \ldots, dx_n \in \Omega^1(U)$. So for each $m \in U$ we have covectors $(dx_i)_m \in T^*_mU = T^*_mM$.

On the other hand, recall that the tangent space T_mM at m has basis given by $\frac{\partial}{\partial x_1}|_m, \ldots, \frac{\partial}{\partial x_n}|_m$. We have

$$(dx_i)_m\left(\frac{\partial}{\partial x_j}|_m\right) = \frac{\partial}{\partial x_j}(x_i) = \delta_{ij}.$$

Thus the $(dx_i)_m$ form a *dual basis* to the cotangent space T^*_mM with respect to the basis $\left\{\frac{\partial}{\partial x_i}|_m\right\}$ for T_pM.

Since the $(dx_i)_m$ form a basis for T^*_mM at all m, it follows that any 1-form $\alpha \in \Omega^1(U)$ can be written as

$$\alpha = \sum_{i=1}^{n} \alpha_i dx_i$$

for some functions $\alpha_i \in C^\infty(U)$ (which may be recovered by evaluating α on $\frac{\partial}{\partial x_i}$).

Exercise 4.4. Suppose that we have two different coordinate charts

$$\phi\colon m \mapsto (x_1(m), \ldots, x_n(m)) \quad \text{and} \quad \psi\colon m \mapsto (y_1(m), \ldots, y_n(m))$$

each with domain given by some open subset U of a smooth manifold. If $\alpha \in \Omega^1(U)$ can be written as

$$\alpha = \sum_{i=1}^{n} \alpha_i dx_i = \sum_{i=1}^{n} \beta_i dy_i$$

find a general formula (in terms of the derivatives of $\phi \circ \psi^{-1}$ and/or $\psi \circ \phi^{-1}$) for the relationship between the coefficients α_i and β_i.

The above exercise is designed to be compared to Exercise 3.22. A single coordinate chart around m produces distinguished bases $\left\{\frac{\partial}{\partial x_i}|_m\right\}$ for T_mM and $\{(dx_i)_m\}$ for T_m^*M, allowing one to parametrize T_mM or T_m^*M by $\mathbb{R}^n$. Changing the coordinate chart changes the appropriate parametrization for either T_mM or T_m^*M, and you should have found that the way in which the parametrization transforms under a coordinate change is different for T_mM than it is for T_m^*M. This reflects the fact that vector fields and 1-forms really are fundamentally different kinds of objects.

If $(x_1, \ldots, x_n)\colon U \to \mathbb{R}^n$ is a coordinate patch and $m \in U$, we see that

$$df_m\left(\frac{\partial}{\partial x_i}\right) = \frac{\partial f}{\partial x_i}(m) = \left(\sum_{j=1}^{n} \frac{\partial f}{\partial x_j}(dx_j)_m\right)\left(\frac{\partial}{\partial x_i}\right),$$

and thus, throughout the coordinate chart U, we have

$$df = \sum_{j=1}^{n} \frac{\partial f}{\partial x_j} dx_j. \tag{6}$$

In principle we could also have defined $d\colon \Omega^0(M) \to \Omega^1(M)$ by saying that if $f \in \Omega^0(M)$ has support in a coordinate chart then df is given by formula (6), and requiring that d be linear over $\mathbb{R}$—this would determine df for any f (not necessarily supported in a coordinate chart) since by using a partition of unity we can write an arbitrary function as a sum of functions each of which is supported in a coordinate chart. (Of course, with this approach one would need to make sure that df didn't depend on the way in which f is decomposed as such a sum—our more natural and coordinate-free definition of d evades this issue).

Having defined the map $d\colon \Omega^0(M) \to \Omega^1(M)$, one could ask whether it is surjective. A little thought should convince you that the answer must be no (if $\dim M \geq 2$)—indeed this may be familiar from multivariable calculus. Consider just a 1-form α which is supported in a coordinate chart U, so in coordinates $\alpha|_U = \sum_i \alpha_i dx_i$ for some smooth functions α_i supported in U, and α vanishes elsewhere. Evidently if $\alpha = df$ then, on U, we would have $\alpha_i = \frac{\partial f}{\partial x_i}$. Since f is assumed C^∞, its mixed partials are equal and so if we had $\alpha = df$ we would need $\frac{\partial \alpha_i}{\partial x_j} = \frac{\partial \alpha_j}{\partial x_i}$ for all i, j, and of course these equations have no reason to hold for a general collection of smooth functions α_i supported in U.

Thus we obtain an *obstruction* to a 1-form α being in the image of d, which in local coordinates can be seen as coming from the partial derivatives of the various components of α. If α is in the image of d it is called *exact*. Once we define the space of 2-forms $\Omega^2(M)$ and the exterior derivative $d\colon \Omega^1(M) \to \Omega^2(M)$, we will see that the above obstruction vanishes in the sense that the relevant partial derivatives coincide if and only if $d\alpha = 0$. Indeed, $d \circ d\colon \Omega^0(M) \to \Omega^2(M)$ is zero (as, more generally, is $d \circ d\colon \Omega^p(M) \to \Omega^{p+2}(M)$). One can then ask whether every α for which the obstruction vanishes ($d\alpha = 0$) is indeed exact. We'll see that the answer to this question depends on the topology of M (as measured by the *de Rham cohomology groups*).)

4.1. **The alternating algebra.**

Definition 4.5. *Let V be a vector space over $\mathbb{R}$, and let p be a positive integer. An* alternating p-form *on V is a function $\eta\colon V^p \to \mathbb{R}$ with the following properties:*

- *η is p-linear: For any i, if $c \in \mathbb{R}$ and $v_1, \ldots, v_p \in V$ and $w_i \in V$ then*

$$\eta(v_1, \ldots, v_{i-1}, cv_i + w_i, \ldots, v_p) = c\eta(v_1, \ldots, v_{i-1}, v_i, \ldots, v_p) + \eta(v_1, \ldots, v_{i-1}, w_i, \ldots, v_p).$$

- *V is antisymmetric: if $v, w \in V$ then, for any $i < j$ and any $u_1, \ldots, u_{i-1}, u_{i+1}, \ldots, u_{j-1}, u_{j+1}, \ldots, u_p \in V$*

$$\eta(u_1, \ldots, u_{i-1}, v, u_{i+1}, \ldots, u_{j-1}, w, u_{j+1}, \ldots, u_p) = -\eta(u_1, \ldots, u_{i-1}, w, u_{i+1}, \ldots, u_{j-1}, v, u_{j+1}, \ldots, u_p).$$

We will denote the vector space of alternating p-forms on V by $\Lambda^p V^*$. We extend the notation $\Lambda^p V^*$ to $p = 0$ by setting $\Lambda^0 V^* = \mathbb{R}$.

Implicit in the above is that the alternating p-forms do indeed form a vector space, which should be clear. Our notation $\Lambda^p V^*$ reflects a number of algebraic facts, not all of which we will need or use: for any vector space V there is a certain standard vector space $\Lambda^p V$ ("the pth graded part of the exterior algebra"), and (at least assuming that V is finite-dimensional) what we denote by $\Lambda^p V^*$ can be canonically identified both with $(\Lambda^p V)^*$ and with $\Lambda^p(V^*)$ (so our lack of parentheses is in writing $\Lambda^p V^*$ is deliberate). There is an obvious identification of $\Lambda^1 V^*$ with V^*.

With this definition, there is for all $p, q \geq 0$ a map

$$\begin{aligned} \wedge\colon\ \Lambda^p V^* \times \Lambda^q V^* &\to \Lambda^{p+q} V^* \\ (\alpha, \beta) &\mapsto \alpha \wedge \beta \end{aligned}$$

called the *wedge product*, which satisfies various important properties. Let us give the definition gradually. The first interesting case is when $p = q = 1$: in this case we define the wedge product by, for $\alpha, \beta \in \Lambda^1 V^*$, and $v, w \in V$,

$$(\alpha \wedge \beta)(v, w) = \alpha(v)\beta(w) - \alpha(w)\beta(v).$$

It is not hard to see that, with this definition, $\alpha \wedge \beta$ does indeed belong to $\Lambda^2 V^*$ (the minus sign ensures that the antisymmetry condition holds). We then extend this to the case that $p = 1$ but q is arbitrary by, for $\alpha \in \Lambda^1 V^*, \beta \in \Lambda^q V^*$,

$$\begin{aligned} (\alpha \wedge \beta)(v_1, v_2, \ldots, v_{q+1}) &= \alpha(v_1)\beta(v_2, \ldots, v_{q+1}) - \alpha(v_2)\beta(v_1, v_3, \ldots, v_{q+1}) \\ &\quad + \alpha(v_3)\beta(v_1, v_2, v_4, \ldots, v_{q+1}) + \cdots + (-1)^l \alpha(v_{q+1})\beta(v_1, \ldots, v_q) \\ &= \sum_{j=1}^{q+1} (-1)^{j-1} \alpha(v_j)\beta(v_1, \ldots, v_{j-1}, v_{j+1}, \ldots, v_{q+1}) \end{aligned}$$

We introduce a notation for "omitting" inputs into k-forms as we often need to do: instead of writing $\beta(v_1, \ldots, v_{j-1}, v_{j+1}, \ldots, v_{q+1})$ we will write $\beta(v_1, \ldots, \hat{v}_j, \ldots, v_{q+1})$; thus the hat signifies that the jth term has been omitted.

We should check that $\alpha \wedge \beta$ as defined above is actually an element of $\Lambda^{q+1} V^*$. It's fairly obvious from this definition that $\alpha \wedge \beta$ is $(q+1)$-linear. As for antisymmetry, if we switch v_k and v_l with $k < l$ then the antisymmetry of β shows that all terms in the sum change sign except for those with $j = k, l$. Meanwhile the kth term changes from $(-1)^{k-1}\alpha(v_k)\beta(v_1, \ldots, \hat{v}_k, \ldots, v_l, \ldots, v_{q+1})$ to $(-1)^{k-1}\alpha(v_l)\beta(v_1, \ldots, \hat{v}_l, \ldots, v_k, \ldots, v_{q+1})$, and the lth term changes from $(-1)^{l-1}\alpha(v_l)\beta(v_1, \ldots, v_k, \ldots, \hat{v}_l, \ldots, v_{q+1})$ to $(-1)^{l-1}\alpha(v_k)\beta(v_1, \ldots, v_l, \ldots, \hat{v}_k, \ldots, v_{q+1})$. I claim that the new lth term is the negative of the old kth term, and vice versa. Indeed to convert the new lth term to something that looks like the old kth term we can "move the v_l past $v_{k+1}, \ldots, v_{l-1}$"—in other words we should switch v_l with v_{k+1}, then switch v_l with v_{k+2}, and so on, until we switch v_l with v_{l-1}. Since β is antisymmetric each of these switches produces a factor of -1, and so since there are a total of $l-k-1$ numbers from $k+1$ to $l-1$ the whole procedure produces a factor of $(-1)^{l-k-1}$. So the new lth term is equal to $(-1)^{l-1}(-1)^{l-k-1}\alpha(v_k)\beta(v_1, \ldots, \hat{v}_k, \ldots, v_{q+1})$, which is indeed equal to the negative of the old kth term. Similarly, the new kth term can be equated with the negative of the old lth term by "moving v_k $l-k-1$ slots to the left." Summing up, switching v_k with v_l causes all the terms with $j \notin \{k, l\}$ to change signs, and also causes the sum of the kth and lth terms to change sign. This proves that $\alpha \wedge \beta$ is alternating, so our map $\Lambda^1 V^* \times \Lambda^q V^* \to \Lambda^{q+1} V^*$ is well-defined.

Finally we extend the definition of the wedge product to general values of p and q. One way of characterizing this extension is that, given our definition for the case $p = 1$, there turns out to be a unique way of extending the definition to general p so that the operation $\wedge$ will be bilinear and associative (for instance, if $\alpha, \beta \in \Lambda^1 V^*$, so that $\alpha \wedge \beta \in \Lambda^2 V^*$, we take the wedge product with $\alpha \wedge \beta$ (on the left) by insisting that $(\alpha \wedge \beta) \wedge \gamma = \alpha \wedge (\beta \wedge \gamma)$ for $\gamma \in \Lambda^q V^*$—since we've already decided how to take wedge product with 1-forms the right-hand side is well-defined).

Instead of showing that this indirect argument gives a well-defined prescription, we give a formula. Given nonnegative integers p and q, let $\mathcal{S}_{p,q}$ denote the collection of p-element subsets of $\{1, \ldots, p+q\}$. Then for

$S \in \mathcal{S}_{p,q}$ let the positive integers $i_1^S < i_2^S < \cdots < i_p^S$ be the elements of S, and let the positive integers $j_1^S < \ldots < j_q^S$ be the elements of $\{1, \ldots, p+q\} \setminus S$. Define $\rho_S : \{1, \ldots, p+q\} \to \{1, \ldots, p+q\}$ by, for $1 \leq k \leq p$, $\rho_S(k) = i_k^S$, and for $p+1 \leq k \leq p+q$, $\rho_S(k) = j_{k-p}^S$. In other words ρ_S is the permutation of $\{1, \ldots, p+q\}$ gotten by writing all the elements of S in increasing order, and then all the elements of $\{1, \ldots, p+q\} \setminus S$ in increasing order. Let $(-)^S$ be 1 if the permutation ρ_S is even and -1 if ρ_S is odd. The general formula for the wedge product is then

$$(\alpha \wedge \beta)(v_1, \ldots, v_{p+q}) = \sum_{S \in \mathcal{S}_{p,q}} (-)^S \alpha(v_{i_1^S}, \ldots, v_{i_p^S}) \beta(v_{j_1^S}, \ldots, v_{j_q^S}) \tag{7}$$

In other words, $(\alpha \wedge \beta)(v_1, \ldots, v_{p+q})$ is gotten by looking at all the different products gotten by plugging in p of the v_i into α and q of them into β, and summing these up with a naturally associated sign. It's not hard to see that this coincides with our previous definition in case $p = 1$.

To help verify some other properties of the wedge product (in particular the fact that the wedge product of alternating forms is alternating) we rewrite (7) as a sum over all permutations on $p+q$ letters. Let $\mathfrak{S}_{p+q}$ denote the group of permutations on $p+q$ letters. Identify $\mathfrak{S}_p \times \mathfrak{S}_q$ with a subgroup of $\mathfrak{S}_{p+q}$ by associating to $(\sigma, \tau) \in \mathfrak{S}_p \times \mathfrak{S}_q$ with the permutation on $p+q$ letters (still denoted (σ, τ)) such that $(\sigma, \tau)(i) = \sigma(i)$ for $1 \leq i \leq p$ and $(\sigma, \tau)(p+j) = p + \tau(j)$ for $1 \leq j \leq q$ (in other words, σ acts on the first p letters and τ acts on the last q). Any permutation in $\eta \in \mathfrak{S}_{p+q}$ can be written uniquely in the form $\eta = \rho_S \circ (\sigma, \tau)$ where ρ_S is one of the permutations from the previous paragraph: namely, let $S = \{\eta(1), \ldots, \eta(p)\}$; let σ send j to r if $\eta(j)$ is the rth largest element of S; and let τ send j to s if $\eta(p+j)$ is the sth largest element of $S \setminus \{\eta(1), \ldots, \eta(p)\}$. If $\eta = \rho_S \circ (\sigma, \tau)$ we see that

$$\alpha(v_{\eta(1)}, \ldots, v_{\eta(p)}) = sgn(\sigma) \alpha(v_{\eta(\sigma^{-1}(1))}, \ldots, v_{\eta(\sigma^{-1}(p))}) = sgn(\sigma) \alpha(v_{i_1^S}, \ldots, v_{i_p^S})$$

where $sgn(\sigma)$ is one if σ is even and -1 if σ is odd, and similarly

$$\beta(v_{\eta(p+1)}, \ldots, v_{\eta(p+q)}) = sgn(\tau) \beta(v_{j_1^S}, \ldots, v_{j_q^S}).$$

Now evidently if $\eta = \rho_S \circ (\sigma, \tau)$ then $sgn(\eta) = (-)^S sgn(\sigma) sgn(\tau)$, and so we deduce

$$sgn(\eta) \alpha(v_{\eta(1)}, \ldots, v_{\eta(p)}) \beta(v_{\eta(p+1)}, \ldots, v_{\eta(p+q)}) = (-)^S \alpha(v_{i_1^S}, \ldots, v_{i_p^S}) \beta(v_{j_1^S}, \ldots, v_{j_q^S}) \quad \text{if } \eta = \rho_S \circ (\sigma, \tau).$$

Now as mentioned earlier any $\eta \in \mathfrak{S}_{p+q}$ can be expressed uniquely as $\rho_S \circ (\sigma, \tau)$ for some S, σ, τ, and so since the pair (σ, τ) varies through the group $\mathfrak{S}_p \times \mathfrak{S}_q$ which has order $p!q!$, we deduce the following (more symmetric and redundant) version of (7):

$$(\alpha \wedge \beta)(v_1, \ldots, v_{p+q}) = \frac{1}{p!q!} \sum_{\eta \in \mathfrak{S}_{p+q}} sgn(\eta) \alpha(v_{\eta(1)}, \ldots, v_{\eta(p)}) \beta(v_{\eta(p+1)}, \ldots, v_{\eta(p+q)}) \tag{8}$$

From (8) it is not difficult to see that $\alpha \wedge \beta$ (which is obviously $(p+q)$-linear) is antisymmetric and hence is an alternating $(p+q)$-form: indeed, let $\tau_{k,l}$ be the transposition which switches letters k and l; of course any permutation can be written uniquely in the form $\eta \circ \tau_{k,l}$, and so we have

$$\begin{aligned}
(\alpha \wedge \beta)(v_1, \ldots, v_{p+q}) &= \frac{1}{p!q!} \sum_{\eta \in \mathfrak{S}_{p+q}} sgn(\eta \circ \tau_{k,l}) \alpha(v_{\eta \circ \tau_{k,l}(1)}, \ldots, v_{\eta \circ \tau_{k,l}(p)}) \beta(v_{\eta \circ \tau_{k,l}(p+1)}, \ldots, v_{\eta \circ \tau_{k,l}(p+q)}) \\
&= \frac{1}{p!q!} \sum_{\eta \in \mathfrak{S}_{p+q}} (-1) sgn(\eta) \alpha(v_{\eta(1)}, \ldots, v_{\eta(p)}) \beta(v_{\eta(p+1)}, \ldots, v_{\eta(p+q)}) \text{ but with the places of } \eta(k) \text{ and } \eta(l) \text{ switched} \\
&= -(\alpha \wedge \beta)(v_1, \ldots, v_{k-1}, v_l, v_{k+1}, \ldots, v_{l-1}, v_k, v_{l+1}, \ldots, v_{p+q}).
\end{aligned}$$

This proves that the map $\wedge : \Lambda^p V^* \times \Lambda^q V^* \to \Lambda^{p+q} V^*$ defined by the equivalent formulas (7,8) is well-defined. The definition is still valid when p and/or q is zero (recalling that $\Lambda^0 V^* = \mathbb{R}$ by definition): wedge product with a 0-form is just multiplication by the corresponding number.

We define the *algebra of alternating forms on V* as the direct sum

$$\Lambda^* V^* = \oplus_{p=0}^{\infty} \Lambda^p V^*.$$

This is equipped with the obvious vector space structure, and also with a multiplication operation $\wedge$ induced by extending bilinearly from the above-defined operations $\wedge\colon\ \Lambda^pV^*\times\Lambda^qV^*\to\Lambda^{p+q}V^*$

Proposition 4.6. *The wedge product obeys:*

(a) *For $\alpha\in\Lambda^pV^*,\beta\in\Lambda^qV^*$,*

$$\beta\wedge\alpha=(-1)^{pq}\alpha\wedge\beta.$$

(b) *For all $\alpha,\beta,\gamma\in\Lambda^*V^*$,*

$$\alpha\wedge(\beta\wedge\gamma)=(\alpha\wedge\beta)\wedge\gamma.$$

Proof. (a) Let $\eta_{p,q}\in\mathfrak{S}_{p+q}$ be the permutation given by $\eta(i)=q+i$ for $1\le i\le p$ and $\eta(j)=j-p$ for $p+1\le j\le p+q$. Note that $sgn(\eta_{p,q})=(-1)^{pq}$ (why?). Any permutation in $\mathfrak{S}_{p+q}$ can be written uniquely in the form $\eta\circ\eta_{p,q}$, so we have

$$\begin{aligned}\alpha\wedge\beta(v_1,\dots,v_{p+q})&=\frac{1}{p!q!}\sum_{\eta\in\mathfrak{S}_{p+q}}sgn(\eta\circ\eta_{p,q})\alpha(v_{\eta\circ\eta_{p,q}(1)},\dots,v_{\eta\circ\eta_{p,q}(p)})\beta(v_{\eta\circ\eta_{p,q}(p+1)},\dots,v_{\eta\circ\eta_{p,q}(p+q)})\\&=\frac{1}{p!q!}\sum_{\eta\in\mathfrak{S}_{p+q}}(-1)^{pq}sgn(\eta)\beta(v_{\eta(1)},\dots,v_{\eta(q)})\alpha(v_{\eta(q+1)},\dots,v_{\eta(p+q)})\\&=(-1)^{pq}\beta\wedge\alpha,\end{aligned}$$

proving (a).

(b) Using the bilinearity of $\wedge$ we may assume that, for some p,q,r, we have $\alpha\in\Lambda^pV^*$, $\beta\in\Lambda^qV^*$, and $\gamma\in\Lambda^rV^*$. Consider ways of writing $\{1,\dots,p+q+r\}$ as a disjoint union $\{1,\dots,p+q+r\}=S_1\coprod S_2\coprod S_3$ where $\#S_1=p$, $\#S_2=q$, $\#S_3=r$. For any such decomposition, write the elements of S_1 in increasing order as $a_1<\cdots<a_p$, those of S_2 as $b_1<\cdots<b_q$, and those of S_3 as $c_1<\cdots<c_r$. Also let $(-)^{S_1S_2S_3}$ for the sign of the permutation obtained by sending i to a_i for $1\le i\le p$, to b_{i-p} for $p+1\le i\le p+q$, and to c_{i-p-q} for $p+q+1\le i\le p+q+r$. Then after repeatedly applying our original formula (7) and unraveling the notation it is easy to check that both

$$(\alpha\wedge(\beta\wedge\gamma))(v_1,\dots,v_{p+q+r})\quad\text{and}\quad((\alpha\wedge\beta)\wedge\gamma)(v_1,\dots,v_{p+q+r})$$

are equal to

$$\sum_{S_1,S_2,S_3}(-)^{S_1S_2S_3}\alpha(v_{a_1},\dots,v_{a_p})\beta(v_{b_1},\dots,v_{b_q})\gamma(v_{c_1},\dots,v_{c_r}).$$

□

Of course, one consequence of associativity is that if $\alpha_1,\dots,\alpha_m\in\Lambda^*V^*$ we can unambiguously write $\alpha_1\wedge\cdots\wedge\alpha_m$. The results of Proposition 4.6 can be summarized as sayingthat Λ^*V^* is an associative, graded commutative algebra.

We now observe that the exterior algebra behaves nicely under linear maps. Suppose that we have two real vector spaces V,W and a linear map $A\colon\ V\to W$. For any p, we obtain a linear map $A^*\colon\ \Lambda^pW^*\to\Lambda^pV^*$ (called the **pullback** of A) by setting

$$(A^*\alpha)(v_1,\dots,v_p)=\alpha(Av_1,\dots,Av_p).$$

Note that since we don't assume A to be invertible it is necessary for A^* to "go in the opposite direction" to get a well-defined map. Extending by linearity produces a linear map $A^*\colon\ \Lambda^*W^*\to\Lambda^*V^*$ defined on the whole alternating algebra.

Proposition 4.7. *If $A\colon\ V\to W$ is a linear map and $\alpha,\beta\in\Lambda^*W^*$ then*

$$A^*(\alpha\wedge\beta)=(A^*\alpha)\wedge(A^*\beta).$$

Proof. This is an immediate consequence of our formula (7) for the wedge product. □

In other words, a linear map $A\colon V \to W$ induces not just a linear map but in fact an algebra homomorphism $\Lambda^* W^* \to \Lambda^* V^*$. Looking at how compositions behave, one sees easily that the alternating algebra construction $V \mapsto \Lambda^* V^*$ defines a contravariant functor from the category of real vector spaces to the category of real associative graded commutative algebras. (Given what we've proven, one just needs to check that $1_V^* = 1_{\Lambda^* V^*}$ and that $(A \circ B)^* = B^* \circ A^*$.)

In the discussion of alternating forms so far, we have avoided choosing a basis for the vector space V (and we haven't even assumed that V is finite-dimensional). This has been deliberate, as we intend to apply this with V equal to the tangent space $T_m M$ at a point on a smooth manifold, and as mentioned before although we can impose a basis on $T_m M$ by choosing a coordinate chart around m, different coordinate charts yield different bases and so there is no canonical choice. However to actually do any computations on a specific vector space one typically does eventually have to choose a basis, and so we now turn to discussing how a basis for V allows one to do calculations in $\Lambda^* V^*$.

So let V be a real vector space with finite dimension n and basis $\{e_1, \ldots, e_n\}$. Let $\{e^1, \ldots, e^n\}$ denote the dual basis for V^* (so $e^i(e_j) = \delta_{ij}$), and recall that V^* is equal to $\Lambda^1 V^*$, so that the e^i can be viewed as elements of the alternating algebra $\Lambda^* V^*$.

Proposition 4.8. *Let $\eta \in \Lambda^p V^*$ and suppose that for all p-tuples of integers $(i_1, \ldots, i_p)$ with $1 \le i_1 < \cdots < i_p \le n$ we have*

$$\eta(e_{i_1}, \ldots, e_{i_p}) = 0.$$

Then $\eta = 0$.

Proof. Suppose to the contrary that $\eta \neq 0$. Then we can choose some $v_1, \ldots, v_p \in V$ with $\eta(v_1, \ldots, v_p) \neq 0$. Now the v_i can be written in the form $v = \sum_j v_{ji} e_j$ for some real numbers v_{ji}. Repeatedly using the p-linearity of η we then find that the nonzero number $\eta(v_1, \ldots, v_p)$ can be written as a linear combination of the real numbers $\eta(e_{j_1}, \ldots, e_{j_p})$ for various k-tuples $(j_1, \ldots, j_p)$. So the fact that $\eta(v_1, \ldots, v_p) \neq 0$ implies that some $\eta(e_{j_1}, \ldots, e_{j_p}) \neq 0$ where $j_1, \ldots, j_p \in \{1, \ldots, n\}$. Now if two of the numbers j_i are equal to each other then it follows directly from the antisymmetry property of η that $\eta(e_{j_1}, \ldots, e_{j_p})$ would be zero, so the numbers $j_1, \ldots, j_p$ making $\eta(e_{j_1}, \ldots, e_{j_p}) \neq 0$ must all be distinct. But again using the antisymmetry property, any reordering of the numbers $j_1, \ldots, j_p$ causes $\eta(e_{j_1}, \ldots, e_{j_p})$ to change only by multiplication by ± 1. So if we choose $i_1 < \cdots < i_p$ to be the result of writing $j_1, \ldots, j_p$ (which we know to be distinct) in strictly increasing order it will hold that $\eta(e_{i_1}, \ldots, e_{i_p}) \neq 0$. This proves (the contrapositive of) the proposition. $\square$

Proposition 4.9. *Suppose that $1 \le p \le n$ and that $1 \le i_1 < \cdots < i_p \le n$ and $1 \le j_1 < \cdots < j_p \le n$ are two strictly increasing sequences of integers from 1 to n. Then*

$$(e^{i_1} \wedge \cdots \wedge e^{i_p})(e_{j_1}, \ldots, e_{j_p}) = \begin{cases} 1 & \text{if } i_l = j_l \text{ for all } l \\ 0 & \text{otherwise} \end{cases}$$

Proof. We can use induction on p. For $p = 1$ this is just the definition of the dual basis, so assume the result holds for p and consider increasing sequences $i_1 < \cdots < i_{p+1}$ and $j_1 < \cdots < j_{p+1}$. If these sequences are not identical to each other, then there is some r such that $j_r \notin \{i_1, \ldots, i_{p+1}\}$. We have (using ˆ to signify omission)

$$(e^{i_1} \wedge \cdots \wedge e^{i_{p+1}})(e_{j_1}, \ldots, e_{j_{p+1}}) = \sum_{s=1}^{p+1} (-1)^{s-1} e^{i_1}(e_{j_s})(e^{i_2} \wedge \cdots \wedge e^{i_{p+1}})(e_{j_1}, \ldots, \widehat{e_{j_s}}, \ldots, e_{j_{p+1}}). \tag{9}$$

The rth term vanishes because $j_r \neq i_1$, and all of the other terms vanish by the inductive hypothesis because $j_r \notin \{i_2, \ldots, i_{k+1}\}$. This proves the "otherwise" part of the proposition.

On the other hand if each i_l coincides with j_l, then since the i_l form an increasing sequence it follows from the inductive hypothesis that, in (9), the first term (*i.e.* the one with $s = 1$) equals 1 and all others equal zero. $\square$

Corollary 4.10. *If $I = (i_1, \ldots, i_p)$ is a p-tuple of integers with $1 \le i_1 < \cdots < i_p \le n = \dim V$, and if we write*

$$e^I = e^{i_1} \wedge \cdots \wedge e^{i_p},$$

then the various e^I *form a basis for* $\Lambda^p V^*$*. In particular* $\dim \Lambda^p V^* = \binom{n}{p} = \frac{n!}{p!(n-p)!}$

Proof. The various e^I are linearly independent: if some linear combination $\sum_I c_I e^I = 0$ then, for any $J = (j_1, \ldots, j_p)$, evaluating both sides on the tuple $(e_{j_1}, \ldots, e_{j_p})$ shows that $c_J = 0$ by Proposition 4.9.

To see that the e^I span $\Lambda^p V^*$, if $\eta \in \Lambda^k V^*$ and $I = (i_1, \ldots, i_p)$ is an increasing sequence, let $\eta_I = \eta(e_{i_1}, \ldots, e_{i_p})$. Then by Proposition 4.9 we have

$$\left(\eta - \sum_I \eta_I e^I\right)(e_{j_1}, \ldots, e_{j_p}) = 0$$

for all increasing sequences $j_1 < \cdots < j_p$. So by Proposition 4.8 it follows that $\eta = \sum_I \eta_I e^I$.

The statement about $\dim \Lambda^p V^*$ just follows from counting the number of increasing sequences of p-tuples I drawn from the set $\{1, \ldots, n\}$, which is evidently the same as the number of p-element subsets of $\{1, \ldots, n\}$, which of course is $\binom{n}{p}$.

□

Of course, the formula $\dim \Lambda^p V^* = \binom{\dim V}{p}$ continues to hold for $p = 0$ for trivial reasons. We note in particular that, if $\dim V = n$, $\Lambda^p V^*$ is trivial for $p > n$, and *one-dimensional* for $p = n$. Evidently a generator for the one-dimensional vector space $\Lambda^n V^*$ is given by $e^1 \wedge \ldots \wedge e^n$ where the e^i form a dual basis to a basis $\{e_i\}$ for v. For some other basis $\{f_i\}$ the element $f^1 \wedge \cdots \wedge f^n$ will then be a multiple of $e^1 \wedge \cdots \wedge e^n$; this multiple is given by the determinant of a certain basis change matrix, as you may be able to see from the following exercise:

Exercise 4.11. Let $A\colon V \to V$ be a linear map, where V is an n-dimensional real vector space. We then have an induced map $A^*\colon \Lambda^n V^* \to \Lambda^n V^*$, which is a linear map from a one-dimensional vector space to itself and hence is given by the formula $A^* x = c_A x$ for all x where c_A is some number depending on A. Prove that $c_A = \det A$. (Hint: Choose a basis in terms of which A has Jordan normal form)

Exercise 4.12. Let V be a finite-dimensional real vector space and let $\alpha \in \Lambda^p V^*$, with $2 \le p \le \dim V$. Let us say that α is *decomposable* if there are $\alpha_1, \ldots, \alpha_p \in \Lambda^1 V^*$ so that $\alpha = \alpha_1 \wedge \cdots \wedge \alpha_p$.

(a) Prove that if α is decomposable then $\alpha \wedge \alpha = 0$.

(b) Prove that if $\dim V = 2$ or 3 then (for $2 \le p \le \dim V$) every $\alpha \in \Lambda^p V^*$ is decomposable.

(c) If $\dim V \ge 4$, construct (with proof, giving an explicit formula) some $\alpha \in \Lambda^2 V^*$ such that α is not decomposable. (Hint: By (a) it is enough to arrange that $\alpha \wedge \alpha \neq 0$.)

4.2. **Higher-degree differential forms.** If M is a smooth manifold and $m \in M$ we let $\Lambda^p T^*_m M$ denote the space of alternating p-forms on the tangent space $T_m M$ (strictly speaking in the notation of the previous subsection we should instead write $\Lambda^p T_m M^*$, but we do not), and let

$$\Lambda^p T^* M = \cup_{m \in M} \{m\} \times \Lambda^p T^*_m M.$$

Thus projection onto the first factor gives a function $\pi\colon \Lambda^p T^* M \to M$, and so we can consider the notion of a section $s\colon M \to \Lambda^p T^* M$, *i.e.* a map s obeying $\pi \circ s = 1_M$, and thus associating to each $m \in M$ an alternating p-form s_m on the tangent space $T_m M$.

Definition 4.13. *A* differential p-form *on M is a section $\eta\colon M \to \Lambda^p T^* M$ obeying the following smoothness property: If $X_1, \ldots, X_p$ are any smooth vector fields on M, then the function*

$$m \mapsto \eta_m\left((X_1)_m, \ldots, (X_p)_m\right)$$

is of class C^∞. We denote the vector space of differential p-forms on M by $\Omega^p(M)$.

Note that this coincides with the previous definition for $p = 1$, recalling the general fact that $\Lambda^1 V^* = V^*$. We also earlier defined $\Omega^0(M)$ to be the space of smooth functions from M to $\mathbb{R}$; since $\Lambda^0 V^* = \mathbb{R}$ this new definition is equivalent (albeit slightly notationally different, but this shouldn't cause a problem) to the previous one.

Assume that $\dim M = n$. Choose a coordinate chart $(x_1, \dots, x_n)\colon U \to \mathbb{R}^n$ with $m \in U$. Recall that, for each $m \in M$, the covectors $(dx_1)_m, \dots, (dx_n)_m$ form a basis for $T^*_m M$, dual to the basis $\left\{\frac{\partial}{\partial x_i}|_m\right\}$ for $T_m M$. For $I = (i_1, \dots, i_p) \in \{1, \dots, n\}^p$ with $i_1 < \dots < i_p$, write

$$dx^I_m = (dx_1)_m \wedge \cdots \wedge (dx_n)_m.$$

According to Corollary 4.10, the various dx^I_m form a basis for $\Lambda^p T^*_m M$. Consequently, for any $\eta \in \Omega^p(M)$, for each q in the coordinate patch U we can write

$$\eta_q = \sum_I f_I(q) dx^I_q$$

for some functions $f_I\colon U \to \mathbb{R}$. Moreover, by evaluating η on tuples of vector fields whose restrictions to U coincide with some of the $\frac{\partial}{\partial x_i}$, we see that the functions f_I are smooth. Thus, a differential p-form restricts to a coordinate chart $(U, x_1, \dots, x_n)$ as an object of the form

$$\eta|_U = \sum_I f_I dx^I \text{ where } f_I \in C^\infty(U).$$

In less abbreviated notation, we could write

$$\eta|_U = \sum_{i_1 < \cdots < i_p} f_{i_1 \cdots i_p} dx_{i_1} \wedge \cdots \wedge dx_{i_p}.$$

Having defined the spaces of p-forms $\Omega^p(M)$, we can let $\Omega^*(M) = \oplus_{p=0}^\infty \Omega^p(M)$; a *differential form* on M is then simply an element of $\Omega^*(M)$.

For each $m \in M$ and $p, q \geq 0$ we have a wedge product operation $\wedge \Lambda^p T^*_m M \times \Lambda^q T^*_m M \to \Lambda^{p+q} T^*_m M$. This then induces a wedge product $\Omega^p(M) \times \Omega^q(M) \to \Omega^{p+q}(M)$ in an obvious way, setting $(\alpha \wedge \beta)_m = \alpha_m \wedge \beta_m$. So, extending bilinearly, we get a wedge product $\wedge\colon \Omega^*(M) \times \Omega^*(M) \to \Omega^*(M)$. In view of Proposition 4.6, the wedge product on differential forms is associative and graded commutative.

We now complete the definition of the *exterior derivative* $d\colon \Omega^*(M) \to \Omega^*(M)$.

Theorem 4.14. *There is a unique $\mathbb{R}$-linear map $d\colon \Omega^*(M) \to \Omega^*(M)$ obeying the following properties:*

(i) *For all p, the restriction $d|_{\Omega^p(M)}$ has image contained in $\Omega^{p+1}(M)$.*
(ii) *$d|_{\Omega^0(M)}$ coincides with the map $d\colon \Omega^0(M) \to \Omega^1(M)$ defined in (5).*
(iii) *If $\omega \in \Omega^p(M)$ and $\phi \in \Omega^q(M)$ we have*

$$d(\omega \wedge \phi) = (d\omega) \wedge \phi + (-1)^p \phi \wedge d\omega.$$

(iv) $d \circ d = 0$.

For any coordinate chart $(x_1, \dots, x_n)\colon U \to \mathbb{R}^n$, if $\omega|_U = \sum_I f_I dx^I$, then

$$d\omega|_U = \sum_{j=1}^n \sum_I \frac{\partial f_I}{\partial x_j} dx_j \wedge dx^I. \tag{10}$$

Proof. We start with the following lemma. Of course, the *support* $supp(\eta)$ of a p-form η is by definition the closure of the set of $m \in M$ for which $\eta_m \in \Lambda^p T^*_m M$ is nonzero.

Lemma 4.15. *Assume that the linear map $d\colon \Omega^*(M) \to \Omega^*(M)$ satisfies properties (i)-(iv) and suppose that $\omega \in \Omega^p(M)$ has $supp(\eta)$ equal to a closed subset of M which is contained in the domain U of a coordinate chart $(x_1, \dots, x_n)\colon U \to \mathbb{R}^n$. If $\omega|_U = \sum_I f_I dx^I$, then $d\omega$ has support contained in U and $d\omega|_U = \sum_{j=1}^n \sum_I \frac{\partial f_I}{\partial x_j} dx_j \wedge dx^I$. The same conclusion continues to hold if we only assume that conditions (i)-(iv) hold for d when d is restricted to forms whose supports are contained in U.*

Proof. Let $\beta\colon M\to\mathbb{R}$ be a smooth function such that $\beta|_{supp(\omega)}=1$ and $supp(\beta)\subset U$. Note then that for each i the smooth function $\beta x_i\colon U\to\mathbb{R}^n$ has closed support within U, and therefore extends to a smooth function on all of M by setting it equal to zero outside of U. Also the functions f_I each have support contained in the support of ω (on which $\beta=1$), so the f_I also extend by zero to smooth functions on all of M, and moreover if $I=(i_1,\dots,i_p)$ we have (at least on U, where both sides are defined)

$$f_I dx^I = f_I d(\beta x_{i_1})\wedge\cdots\wedge d(\beta x_{i_p}).$$

Thus

$$\omega=\sum_{I=(i_1,\dots,i_p)} f_I d(\beta x_{i_1})\wedge\cdots\wedge d(\beta x_{i_p})$$

(the two sides coincide on U, and are both zero outside of U).

Now by induction on the integer r it is easy to see from conditions (iii) and (iv) that, for any smooth functions $g_1,\dots,g_r$ we have

$$d\,(dg_1\wedge dg_2\wedge\cdots\wedge dg_r)=0.$$

Applying this fact together with (iii) again (and the linearity of d) shows that

$$d\omega=\sum_I df_I\wedge d(\beta x_{i_1})\wedge\cdots\wedge d(\beta x_{i_p}).$$

Since β is identically 1 on the union of the supports of the f_I (which is contained in U), and since $df_I=\sum_j\frac{\partial f_I}{\partial x_j}dx_j$ on U, the result follows. □

Motivated by this lemma, choose once and for all a cover $\{U_\alpha\}$ by domains of coordinate charts $(x_1^\alpha,\dots,x_n^\alpha)\colon U_\alpha\to\mathbb{R}$, and let $\{\chi_\alpha\}$ be a partition of unity subordinate to the cover $\{U_\alpha\}$. For $I=(i_1,\dots,i_p)$ let $dx_\alpha^I=dx_{i_1}^\alpha\wedge\cdots\wedge dx_{i_p}^\alpha$.

Lemma 4.16. *For any α let $\Omega_\alpha^*(M)$ denote the space of differential forms on M whose support is contained in α. Define $d_\alpha\colon\Omega_\alpha^*(M)\to\Omega_\alpha^*(M)$ by setting, if $\omega\in\Omega_\alpha^*(M)$ with $\omega|_{U_\alpha}=\sum_I f_I dx_\alpha^I$,*

$$d_\alpha\omega|_{U_\alpha}=\sum_I df_I\wedge dx_\alpha^I$$

(and $d_\alpha\omega=0$ outside U_α). Then $d_\alpha\colon\Omega_\alpha^*(M)\to\Omega_\alpha^*(M)$ satisfies (i)-(iv) of Theorem 4.14 when restricted to $\Omega_\alpha^*(M)$, and is the unique such map with these properties.

Proof. Uniqueness is already proven in (the last sentence of) Lemma 4.15, so we just need to check that (i)-(iv) are satisfied. (i) is obvious, and (ii) is given by Equation 6. The fact that (iii) holds outside of U_α is trivial (both sides are zero); inside of U_α let us write $\omega|_{U_\alpha}=\sum_I f_I dx_\alpha^I$ and $\phi|_{U\alpha}=\sum_J g_J dx_\alpha^J$ (where the multi-indices I have length p and the multi-indices J have length q). We then have, on U_α,

$$\begin{aligned}
d_\alpha(\omega\wedge\phi) &= d_\alpha\left(\sum_{I,J} f_I g_J dx_\alpha^I\wedge dx_\alpha^J\right)=\sum_{k,I,J}\frac{\partial(f_I g_J)}{\partial x_k^\alpha}dx_k^\alpha\wedge dx_\alpha^I\wedge dx_\alpha^J\\
&=\sum_{k,I,J}\left(\frac{\partial f_I}{\partial x_k^\alpha}g_J+f_I\frac{\partial g_J}{\partial x_k^\alpha}\right)dx_k^\alpha\wedge dx_\alpha^I\wedge dx_\alpha^J\\
&=\sum_{k,I,J}\left(\frac{\partial f_I}{\partial x_k^\alpha}dx_k^\alpha\wedge dx_\alpha^I\right)\wedge(g_J dx_\alpha^J)+\sum_{k,I,J}(-1)^p(f_I dx_\alpha^I)\wedge\left(\frac{\partial g_J}{\partial x_k^\alpha}\right)dx_k^\alpha\wedge dx_\alpha^J\\
&=(d_\alpha\omega)\wedge\phi+(-1)^p\omega\wedge d_\alpha\phi
\end{aligned}$$

where the $(-1)^p$ comes from applying Proposition 4.6 (a) to the wedge product $dx_k^\alpha \wedge dx_\alpha^I$. This proves that d_α satisfies (iii). As for (iv), if $\omega|_{U_\alpha} = \sum_I f_I dx_\alpha^I$, then clearly $d_\alpha(d_\alpha\omega)$ vanishes outside U_α, and on U_α we have

$$\begin{aligned} d_\alpha(d_\alpha\omega) &= d_\alpha\left(\sum_{k=1}^n \sum_I \frac{\partial f_I}{\partial x_k^\alpha} dx_k^\alpha \wedge dx_\alpha^I\right) \\ &= \sum_I \left(\sum_{l=1}^n \sum_{k=1}^n \frac{\partial^2 f_I}{\partial x_l^\alpha \partial x_k^\alpha} dx_l^\alpha \wedge dx_k^\alpha\right) \wedge dx_\alpha^I \\ &= \sum_I \left(\sum_{l=1}^n \sum_{k<l} \left(\frac{\partial^2 f_I}{\partial x_l^\alpha \partial x_k^\alpha} - \frac{\partial^2 f_I}{\partial x_k^\alpha \partial x_l^\alpha}\right) dx_l^\alpha \wedge dx_k^\alpha\right) \wedge dx_\alpha^I = 0 \end{aligned}$$

since the mixed partials of the smooth function f_I are equal (of course in the second-to-last equation we've switched the indices k and l in the terms that initially had $k > l$ and used the fact that $dx_k^\alpha \wedge dx_l^\alpha = -dx_l^\alpha \wedge dx_k^\alpha$). This proves (iv) and so completes the proof of the lemma. □

We now move from these local considerations to prove the global Theorem 4.14. We have fixed a (locally finite) partition of unity $\{\chi_\alpha\}$ subordinate to a cover U_α. Then if $\omega \in \Omega^*(M)$ we have

$$\omega = \sum_\alpha (\chi_\alpha \omega) \quad \text{where each} \quad \chi_\alpha\omega \in \Omega_\alpha^*(M).$$

So for each α we have a well-defined differential form $d_\alpha(\chi_\alpha\omega)$, whose support is contained in the support of χ_α (in particular any point in M has a neighborhood meeting the supports of only finitely many of the $d_\alpha(\chi_\alpha\omega)$, so the sum $\sum_\alpha d_\alpha(\chi_\alpha\omega)$ is a well-defined differential form). So define

$$d\omega = \sum_\alpha d_\alpha(\chi_\alpha\omega).$$

This is clearly $\mathbb{R}$-linear since each of the d_α are, and conditions (i), (ii), and (iv) are each also manifestly inherited from the corresponding facts for d_α (together, in the case of (ii), with the fact that the map $d\colon \Omega^0(M) \to \Omega^1(M)$ defined earlier in (5) is also $\mathbb{R}$-linear). Condition (iii) (the form version of the Leibniz rule) takes just a little more work. For each α let ψ_α be a smooth function which is equal to one on $supp(\chi_\alpha)$ but such that we still have $supp(\psi_\alpha) \subset U_\alpha$. If $\omega \in \Omega^p(M)$ and $\phi \in \Omega^q(M)$, we have by definition

$$d(\omega \wedge \phi) = \sum_\alpha d_\alpha(\chi_\alpha(\omega \wedge \phi)).$$

Note that $\chi_\alpha(\omega \wedge \phi) = (\chi_\alpha\omega) \wedge (\psi_\alpha\phi)$ (both factors of which have support in U_α), so

$$d_\alpha(\chi_\alpha(\omega \wedge \phi)) = d_\alpha(\chi_\alpha\omega) \wedge (\psi_\alpha\phi) + (-1)^p \chi_\alpha\omega \wedge d_\alpha(\psi_\alpha\phi)$$

and so (freely using associativity and distributivity of the wedge product, as well as the fact that $\psi_\alpha\phi = \phi$ wherever $d(\chi_\alpha\omega) \neq 0$)

$$\begin{aligned} d(\omega \wedge \phi) &= \sum_\alpha d_\alpha(\chi_\alpha\omega) \wedge \phi + (-1)^p \omega \wedge \left(\sum_\alpha \chi_\alpha d_\alpha(\psi_\alpha\phi)\right) \\ &= (d\omega) \wedge \phi + (-1)^p \omega \wedge \left(\sum_\alpha \chi_\alpha d_\alpha(\psi_\alpha\phi)\right) \end{aligned}$$

So evidently it remains only to show that

$$\sum_\alpha \chi_\alpha d_\alpha(\psi_\alpha\phi) \stackrel{?}{=} d\phi. \tag{11}$$

Note also that $\chi_\alpha\psi_\alpha = \chi_\alpha$ and $\psi_\alpha d\chi_\alpha = d\chi_\alpha$, so

$$\begin{aligned} d_\alpha(\chi_\alpha\phi) &= d_\alpha(\chi_\alpha\psi_\alpha\phi) \\ &= \chi_\alpha d_\alpha(\psi_\alpha\phi) + d\chi_\alpha \wedge (\psi_\alpha\phi) = \chi_\alpha d_\alpha(\psi_\alpha\phi) + d\chi_\alpha \wedge \phi, \end{aligned}$$

i.e.

$$\chi_\alpha d_\alpha(\psi_\alpha\phi) = d_\alpha(\chi_\alpha\phi) - d\chi_\alpha \wedge \phi.$$

Thus

$$\begin{aligned} \sum_\alpha \chi_\alpha d_\alpha(\psi_\alpha\phi) &= \sum_\alpha d_\alpha(\chi_\alpha\phi) - \sum_\alpha d\chi_\alpha \wedge \phi \\ &= d\phi - d\left(\sum_\alpha \chi_\alpha\right) \wedge \phi = d\phi \end{aligned}$$

since $\sum_\alpha d\chi_\alpha = 1$ and so $d\left(\sum_\alpha \chi_\alpha\right) = 0$.

This completes the proof that d, as we have defined it, satisfies the desired properties. The formula (10) given at the end of the theorem for the behavior of d on an arbitrary coordinate chart then follows from Lemma 4.15: If $m \in U$ choose a cutoff function $\beta\colon M \to \mathbb{R}$ equal to 1 on a neighborhood of m and with compact support contained in U; then $\omega = \beta\omega + (1-\beta)\omega$ and we have $(d((1-\beta)\omega))_m = 0$ while Lemma 4.15 ensures that $(d(\beta\omega))_m$ is given by evaluating the right-hand side of (10) at m. □

It is not initially obvious that the formula for d given in the proof, namely $d\omega = \sum_\alpha d_\alpha(\chi_\alpha\omega)$, would give an answer which is independent of the partition of unity $\{\chi_\alpha\}$ or of the open cover $\{U_\alpha\}$, but the uniqueness part of the theorem implies that this independence property holds.

In practice, one does not calculate $d\omega$ by choosing a partition of unity; rather one covers the manifold by coordinate charts U and uses the formula (10) to express $d\omega$ in each of these coordinate charts. Again, it is not initially obvious that, if V is another coordinate chart with $U \cap V = \emptyset$, the forms obtained by using (10) with reference to the two different coordinate charts would give both give the same answer when restricted to $U \cap V$. However, the theorem ensures that this is in fact the case (one can also verify this somewhat tediously by a direct computation).

Since $d \circ d = 0$, we can make the following definition:

Definition 4.17. *Let M be a smooth manifold, and p a nonnegative integer. The* pth de Rham cohomology *of M is the real vector space*

$$H^p_{dR}(M) = \frac{\ker(d\colon\ \Omega^p(M) \to \Omega^{p+1}(M))}{Im(d\colon\ \Omega^{p-1}(M) \to \Omega^p(M))}.$$

(For the case $p = 0$, we regard $\Omega^{-1}(M)$ as the trivial vector space, so that $H^0_{dR}(M) = \ker(d\colon\ \Omega^0(M) \to \Omega^1(M))$.)

Remark 4.18. A form ω such that $d\omega = 0$ is called *closed*, and a form ω such that $\omega = d\phi$ for some ϕ is called exact. Thus the fact that $d \circ d = 0$ expresses that every exact form is closed, and the pth de Rham cohomology group measures the extent to which it fails to be true that, conversely, every closed p-form is exact.

I would also like to record a fact which we will make use of shortly, and which basically was proven in the proof of Theorem 4.14:

Proposition 4.19. *Let $\omega \in \Omega^p(M)$. Then we can write ω as a locally finite sum $\omega = \sum_\gamma \omega_\gamma$ (*i.e.*, any point has an open set intersecting only finitely many of the $supp(\omega_\gamma)$) such that each ω_γ is given by*

$$\omega_\gamma = f_\gamma dg_{1,\gamma} \wedge \cdots \wedge dg_{p,\gamma}$$

for some functions $f_\gamma, g_{1,\gamma}, \ldots, g_{p,\gamma} \in C^\infty(M)$.

Proof. Let $\{U_\alpha\}$ be an open cover of M by domains of coordinate charts $(x_1^\alpha, \dots, x_n^\alpha)$ and $\{\chi_\alpha\}$ a (locally finite) partition of unity subordinate to $\{U_\alpha\}$. We can then write $\omega = \sum_\alpha (\chi_\alpha \omega)$ where each $\chi_\alpha \omega$ is supported in U_α. In turn, it was shown in the proof of Lemma 4.15 that each $\chi_\alpha\omega$ can be written as a finite sum of forms of the desired type $f_{\alpha,I} dg_{1,\alpha,I} \wedge \cdots \wedge dg_{p,\alpha,I}$ (as I varies over multi-indices $I = (i_1, \dots, i_p)$), namely one sets $g_{j,\alpha,I} = \beta x_{i_j}^\alpha$ where β is a smooth function supported in U_α and equal to 1 on $supp(\chi_\alpha)$. So by having the index γ vary over pairs (α, I) the result follows.

□

To get a sense of what the exterior derivative d is measuring, it is instructive to consider the special cases where the smooth manifold is an open subset U of $\mathbb{R}^2$ or $\mathbb{R}^3$. As mentioned earlier, for any open subset of $\mathbb{R}^n$ the degree-zero part of d acts by $df = \sum_{i=1}^n \frac{\partial f}{\partial x_i} dx_i$. So if we use the standard basis of $\mathbb{R}^n$ to identify vector fields with 1-forms[2], the exterior derivative of a function is essentially its *gradient* in the sense of multivariable calculus.

For open subsets $U \subset \mathbb{R}^2$, the only remaining interesting part of d is that acting on 1-forms. A general 1-form on U has the shape

$$\omega = P(x,y)dx + Q(x,y)dy$$

for functions $P, Q \in C^\infty(U)$, and we see that

$$\begin{aligned} d\omega &= \frac{\partial P}{\partial x} dx \wedge dx + \frac{\partial P}{\partial y} dy \wedge dx + \frac{\partial Q}{\partial x} dx \wedge dy + \frac{\partial Q}{\partial y} dy \wedge dy \\ &= \left(\frac{\partial Q}{\partial x} - \frac{\partial P}{\partial y} \right) dx \wedge dy. \end{aligned}$$

So if we consider ω as corresponding to the vector field with components P, Q, then $d\omega$ is obtained by multiplying the standard 2-form $dx \wedge dy$ by what is sometimes called the *scalar curl* of this vector field, $\frac{\partial Q}{\partial x} - \frac{\partial P}{\partial y}$, a function which is probably familiar from Green's theorem in multivariable calculus.

Moving up a dimension to open subsets $U \subset \mathbb{R}^3$, a general 1-form on U has the form

$$\omega = Pdx + Qdy + Rdz,$$

and we find that in this case

$$d\omega = \left(\frac{\partial R}{\partial y} - \frac{\partial Q}{\partial z} \right) dy \wedge dz + \left(\frac{\partial P}{\partial z} - \frac{\partial R}{\partial x} \right) dz \wedge dx + \left(\frac{\partial Q}{\partial x} - \frac{\partial P}{\partial y} \right) dx \wedge dy.$$

We see that the three coefficients above are the components of the *curl* of the vector field $\langle P, Q, R \rangle$.

Meanwhile, a general 2-form on U can be written $\eta = Pdy \wedge dz + Qdz \wedge dx + Rdx \wedge dy$ and so (because we are working in $\mathbb{R}^3$) also corresponds to a vector field $\langle P, Q, R \rangle$. We see that

$$d\eta = \left(\frac{\partial P}{\partial x} + \frac{\partial Q}{\partial y} + \frac{\partial R}{\partial z} \right) dx \wedge dy \wedge dz,$$

and recognize the coefficient from multivariable calculus as the *divergence* of the vector field $\langle P, Q, R \rangle$.

Thus in dimension 3 the maps $d \colon \Omega^0(U) \to \Omega^1(U)$, $d \colon \Omega^1(U) \to \Omega^2(U)$, and $d \colon \Omega^2(U) \to \Omega^3(U)$ correspond respectively to the gradient, curl, and divergence operators from multivariable calculus. The fact that $d \circ d = 0$ expresses the facts that the curl of a gradient is always zero, and that the divergence of a curl is always zero.

Again for open subsets $U \subset \mathbb{R}^3$, the first de Rham cohomology group $H_{dR}^1(U)$ will be zero if and only if, conversely, every vector field whose curl is equal to zero is in fact the gradient of a function. You probably learned

[2] As I've emphasized elsewhere, on a general smooth manifold vector fields and 1-forms are different kinds of objects and one shouldn't try to identify them since they transform differently under coordinate changes, but on $\mathbb{R}^n$ one can decide to only ever work in the standard coordinate chart and then there won't be any harm in making this identification

in multivariable calculus that if U is all of $\mathbb{R}^3$ then this statement holds. However if U is more topologically interesting it may not hold: for example there is the (misleadingly labeled) "$d\theta$" form, given by

$$d\theta = \frac{xdy - ydx}{x^2 + y^2}$$

defined on $U = \{(x, y, z) \in \mathbb{R}^3 | x^2 + y^2 \neq 0\}$, which you can verify to be closed, but which (despite the notation) is not exact since it has nonzero integral around closed curves which enclose the z-axis ($d\theta$ wants to be the exterior derivative of the polar coordinate θ, but θ is not a well-defined smooth function on U).

Similarly, the second de Rham cohomology group of an open subset $U \subset \mathbb{R}^3$ vanishes if and only if every vector field on U which has divergence equal to zero is in fact the curl of some other vector field. If $U = \mathbb{R}^3$ then this is true (we'll prove a much more general statement not too long from now), but this statement is false for $U = \mathbb{R}^3 \setminus \{(0, 0, 0)\}$. A standard example illustrating this is the form

$$\eta = \frac{xdy \wedge dz + ydz \wedge dx + zdx \wedge dy}{(x^2 + y^2 + z^2)^{3/2}}$$

Physically, η corresponds to the electric field on $\mathbb{R}^3 \setminus \{(0, 0, 0)\}$ generated by a point charge located at the origin. The statement that this vector field is not the curl of another vector field can be shown using Stokes' theorem, by taking the flux integral of the vector field over a sphere around the origin. Later we'll develop language for this that generalizes such arguments substantially and stays within the realm of differential forms rather than vector fields.

Exercise 4.20. (A coordinate-free formula for d): Let M be a smooth manifold, $\omega \in \Omega^p(M)$, and let $X^{(0)}, \ldots, X^{(p)}$ be vector fields on M. Prove that

$$(d\omega)(X^{(0)}, \ldots, X^{(p)}) = \sum_{i=0}^{p} (-1)^i X^{(i)} \left(\omega(X^{(0)}, \ldots, \widehat{X^{(i)}}, \ldots, X^{(p)})\right) + \sum_{i<j} (-1)^{i+j} \omega\left([X^{(i)}, X^{(j)}], X^{(0)}, \ldots, \widehat{X^{(i)}}, \ldots, \widehat{X^{(j)}}, \ldots, X^{(p)}\right).$$

(To clarify the notation, if we have a differential q-form α and vector fields $Y^{(1)}, \ldots, Y^{(q)}$, the function

$$m \mapsto \alpha_m(Y_m^{(1)}, \ldots, Y_m^{(q)})$$

is a smooth function, which we denote by $\alpha(Y^{(1)}, \ldots, Y^{(q)})$. In particular since vector fields are derivations on the space of smooth functions, if Z is another vector field we get another smooth function given by $Z\left(\alpha(Y^{(1)}, \ldots, Y^{(q)})\right)$. To do this problem, I would suggest first showing that the value of the function on the right-hand side at a point m is unchanged if some (or all) $X^{(i)}$ are replaced by another vector field $\bar{X}^{(i)}$ such that $X_m^{(i)} = \bar{X}_m^{(i)}$, and then proving the result when the $X^{(i)}$ are (at least on a neighborhood of a given point) equal to standard coordinate vector fields.)

4.3. **Pullbacks of differential forms and the naturality of d.** Let $\phi\colon M \to N$ be a smooth map between two smooth manifolds. Recall then that for each $m \in M$ we have a derivative map $\phi_*\colon T_mM \to T_{\phi(m)}N$, defined in terms of the derivation formalism by the simple formula

$$(\phi_* v)(f) = v(f \circ \phi)$$

whenever f is a germ of a C^∞ function defined near $\phi(m) \in N$. As described just before Proposition 4.7, this induces for all $m \in N$ a pullback operation

$$\phi^*\colon \Lambda^p T^*_{\phi(m)}N \to \Lambda^p T^*_m M$$

by setting, for $\alpha \in \Lambda^p T^*_{\phi(m)}N$ and $v_1, \ldots, v_p \in T_mM$,

$$(\phi^*\alpha)(v_1, \ldots, v_p) = \alpha(\phi_* v_1, \ldots, \phi_* v_p).$$

In particular, when $p = 1$, so that $\Lambda^p T_p^* M$ is just the cotangent space $T_p^* M$, ϕ^* coincides with the adjoint map to ϕ_* from linear algebra.

Theorem 4.21. *Let $\phi\colon M \to N$ be a smooth map and let $\omega \in \Omega^p(M)$ be a differential form. Define a section $\phi^*\omega$ of $\Lambda^p T^*M$ by*

$$(\phi^*\omega)_m = \phi^*(\omega_{\phi(m)}).$$

Then $\phi^*\omega$ is a differential form on M, and

$$d(\phi^*\omega) = \phi^*(d\omega). \tag{12}$$

The fact that $\phi^*\omega$ is a differential form requires proof, since there is a smoothness condition to check. In case $p = 0$ (so that $\omega \in C^\infty(M)$) the definition above should be read as saying that

$$\phi^*\omega := \omega \circ \phi \qquad (\text{if } \omega \in \Omega^0(M)).$$

Proof. Step 1: We prove the theorem when $p = 0$. Let $h \in \Omega^0(M) = C^\infty(N)$ be a 0-form. By definition $\phi^*h = h\circ\phi$, which is certainly a smooth function (*i.e.* a 0-form) on M since compositions of smooth functions are smooth. For all $v \in T_mM$ we have, by the definition of d on 0-forms:

$$(d(\phi^*h))_m(v) = v(\phi^*h) = v(h \circ \phi) = (\phi_*v)(h) = (dh)_{\phi(m)}(\phi_*v) = (\phi^*dh)_m(v).$$

This confirms that $d(\phi^*h) = \phi^*dh$ (It also confirms that ϕ^*dh satisfies the smoothness condition required of a 1-form, since $d(\phi^*h)$ certainly does so.)

Step 2: We prove the theorem in case $\omega = f dg_1 \wedge \cdots \wedge dg_p$ for some $f, g_1, \ldots, g_p \in C^\infty(N)$. In this case, if $m \in M$, we have (using Proposition 4.7 and Step 1)

$$\begin{aligned}(\phi^*\omega)_m &= f(\phi(m))\phi^*\left((dg_1)_{\phi(m)} \wedge \cdots \wedge (dg_p)_{\phi(m)}\right)\\ &= (f \circ \phi)(m)\left((\phi^*dg_1)_m \wedge \cdots \wedge (\phi^*dg_p)_m\right)\\ &= (f \circ \phi)(m)\left(d(g_1 \circ \phi)_m \wedge \cdots \wedge d(g_p \circ \phi)_m\right),\end{aligned}$$

i.e.

$$\phi^*\omega = (f \circ \phi)d(g_1 \circ \phi) \wedge \cdots \wedge d(g_p \circ \phi).$$

Now the space of differential forms is closed under wedge product (as the smoothness condition is easily seen to be preserved), and the zero-form $f \circ \phi$ and the 1-forms $d(g_i \circ \phi)$ are all differential forms by what we have already done, so this proves that $\phi^*\omega$ is a differential form. Using the Leibniz rule and the fact that $d^2 = 0$ we see that

$$\begin{aligned}d(\phi^*\omega) &= d\left((f \circ \phi)d(g_1 \circ \phi) \wedge \cdots \wedge d(g_p \circ \phi)\right)\\ &= d(f \circ \phi) \wedge d(g_1 \circ \phi) \wedge \cdots \wedge d(g_p \circ \phi)\\ &= (\phi^*df) \wedge (\phi^*dg_1) \wedge \cdots \wedge \phi^*(dg_p)\\ &= \phi^*\left(df \wedge dg_1 \wedge \cdots \wedge dg_p\right)\\ &= d\left(fdg_1 \wedge \cdots \wedge dg_p\right) = d\omega.\end{aligned}$$

Step 3: We prove the result in general. By Proposition 4.19, any differential form $\omega \in \Omega^p(N)$ can be written as a locally finite sum of forms of the type considered in Step 2. Now the smoothness condition required of a differential form is preserved under locally finite sums (since the smoothness of a function can be checked by looking at its restriction to each member of an open cover, we can reduce to the case of genuinely finite sums), so using the linearity of ϕ^* it follows that $\phi^*\omega$ is a differential form. Similarly the $\mathbb{R}$-linearity of d, together with Step 2, implies that $d\phi^*\omega = \phi^*d\omega$ □

Corollary 4.22. *A smooth map $\phi\colon M \to N$ between two smooth manifolds induces by the pullback operation a map $\phi^*\colon \Omega^*(N) \to \Omega^*(M)$. If $\omega \in \Omega^*(N)$ is closed, then $\phi^*\omega \in \Omega^*(M)$ is closed, and if $\omega \in \Omega^*(N)$ is exact, then $\phi^*\omega \in \Omega^*(M)$ is exact*

Proof. The first sentence has already been proven. If ω is closed, *i.e.* $d\omega = 0$, then $d(\phi^*\omega) = \phi^*d\omega = \phi^*0 = 0$. If ω is exact, *i.e.* $\omega = d\eta$ for some $\eta \in \Omega^*(N)$, then $\phi^*\omega = \phi^*d\eta = d(\phi^*\eta)$. □

Recall that we have defined the pth de Rham cohomology of a smooth manifold M as the quotient vector space

$$H^p_{dR}(M) = \frac{\{\text{closed } p\text{-forms}\}}{\{\text{exact } p\text{-forms}\}}.$$

If we write $H^*_{dR}(M) = \oplus_{p=0}^{\infty} H^p_{dR}(M)$, the wedge-product induces a *ring* structure on $H^*_{dR}(M)$: if $a \in H^p_{dR}(M)$ and $b \in H^q_{dR}(M)$, then we can find closed forms $\omega \in \Omega^p(M)$, $\eta \in \Omega^q(M)$, representing the classes a and b. Then $d(\omega \wedge \eta) = (d\omega) \wedge \eta + (-1)^p \omega \wedge (d\eta) = 0$, so $\omega \wedge \eta$ represents some cohomology class (denoted $a \cup b$) in $H^{p+q}_{dR}(M)$. Moreover this cohomology class is independent of our particular choice of representatives ω and η—for example if we replaced ω by some other form $\bar{\omega} = \omega + d\alpha$, then

$$\bar{\omega} \wedge \eta = (\omega + d\alpha) \wedge \eta = \omega \wedge \eta + (d\alpha) \wedge \eta = \omega \wedge \eta + d(\alpha \wedge \eta)$$

(we've used that $d\eta = 0$), *i.e.* the de Rham cohomology class of $\bar{\omega} \wedge \eta$ is the same as that of $\omega \wedge \eta$ (they differ by an exact form).

Using Proposition 4.6, one easily checks that this multiplication on $H^*_{dR}(M)$ (called the *cup product*) gives $H^*_{dR}(M)$ the structure of an associative, graded commutative $\mathbb{R}$-algebra.

Corollary 4.23. *If M and N are smooth manifolds and $\phi\colon M \to N$ is a smooth map, we obtain a homomorphism of graded $\mathbb{R}$-algebras (in particular a ring homomorphism) $\phi^*\colon H^*_{dR}(N) \to H^*_{dR}(M)$ by setting $\phi^*[\omega] = [\phi^*\omega]$ for any closed form ω on N. If ϕ is a diffeomorphism then ϕ^* is an isomorphism.*

Proof. The first sentence follows directly from various things that we have already done (check this for yourself if it's not clear). For the second, note that ϕ^* (acting either on forms or on cohomology) satisfies the functoriality conditions $(Id)^* = (Id)$ and $(\phi \circ \psi)^* = \psi^* \circ \phi^*$ (note the order on the right hand side, reflecting that ϕ^* "goes in the opposite direction" to ϕ). From this it follows immediately that if ϕ is a diffeomorphism then ϕ^* is an isomorphism with inverse $(\phi^{-1})^*$. □

Exercise 4.24. If M is a smooth manifold, give an explicit formula, in terms of the point-set topology of M, for the degree-zero de Rham cohomology $H^0_{dR}(M)$. (As a point of convention, since there is no such thing as a (-1)-form, we regard the exact 0-forms on M to consist only of 0.)

MATH 8210 LECTURE NOTES, PART 2

1. Submanifolds

Throughout this section fix a smooth m-dimensional manifold M.

Definition 1.1. Let N be a smooth manifold and let $\phi\colon N \to M$ be a smooth map. Then

- ϕ is called a *submersion* if, for all $x \in N$, the linearization $\phi_*\colon T_xN \to T_{\phi(x)}M$ is surjective.
- ϕ is called an *immersion* if, for all $x \in N$, the linearization $\phi_*\colon T_xN \to T_{\phi(x)}M$ is injective.
- ϕ is called an *embedding* if it is an immersion and, moreover, the map ϕ is a homeomorphism from N to $\phi(N)$, where $\phi(N)$ is equipped with the subspace topology.

Here are some examples; if these notions are unfamiliar to you then you should check for yourself that they satisfy the respective definitions.

Example 1.2. (i) *The projection $\pi\colon \mathbb{R}^n \to \mathbb{R}^m$ onto the first m coordinates (assuming $m \le n$) is a submersion; in fact this provides a local model for all submersions, as will follow from the proof of Theorem 1.8. For a more interesting global example, the projection $\pi\colon \mathbb{R}^{n+1} \setminus \{\vec{0}\} \to \mathbb{R}P^n$ is a submersion (as is the projection $S^n \to \mathbb{R}P^n$).*

(ii) *Dually, if $n \le m$, then the inclusion $i\colon \mathbb{R}^n \to \mathbb{R}^m$ (defined by $i(\vec{x}) = (\vec{x}, \vec{0})$ where $\mathbb{R}^m$ is split as $\mathbb{R}^n \times \mathbb{R}^{m-n}$) is an example of an embedding.*

(iii) *A simple example of an immersion which is not an embedding is the map $\phi\colon \mathbb{R} \to \mathbb{C}$ given by $\phi(x) = e^{ix}$.*

(iv) *Of course the problem with (iii) was that it wasn't injective, but one can also construct examples of injective immersions which are not embeddings. For instance, take two smooth functions $f, g\colon \mathbb{R} \to \mathbb{R}$ such that for all $t < 0$ one has $f(t) = t$ and $g(t) = 0$, and such that there is no $t \in \mathbb{R}$ such that $f'(t) = g'(t) = 0$. Then the map $\phi\colon \mathbb{R} \to \mathbb{R}^2$ defined by $\phi(t) = (f(t), g(t))$ will be an immersion. If one chooses f and g so that $\lim_{t\to\infty} f(t) = -1$ and $\lim_{t\to\infty} g(t) = 0$ and so that ϕ is injective (as can easily be done—you might draw a picture if this isn't obvious to you), then ϕ won't be an embedding, since by looking at neighborhoods of $(-1, 0)$ in $\psi(\mathbb{R})$ one sees that the image isn't a topological manifold when equipped with the subspace topology.*

Declare two embeddings $\phi_1\colon N_1 \to M$ and $\phi_2\colon N_2 \to M$ to be equivalent if there is a diffeomorphism $\psi\colon N_1 \to N_2$ such that $\phi_1 = \phi_2 \circ \psi$. An overly formal definition of a *submanifold* is that a submanifold is an equivalence class of embeddings under this equivalence relation. Of course, part of the point of the above equivalence relation is that if $\phi_1 \sim \phi_2$ then $\phi_1(N_1) = \phi_2(N_2)$; when one thinks of a submanifold one should think of the subset of M formed as the image of any representative embedding. If one has a subset $N \subset M$, it inherits a subspace topology, and one can ask whether or not this subspace topology makes N a topological manifold. One can then ask whether the topological space N admits smooth structures (this is now an intrinsic question about N), and how these are related to the ambient space M. Accordingly I prefer the following definition:

Definition 1.3. A *submanifold* of M is a subset N which is a topological manifold with respect to its subspace topology, equipped moreover with a smooth structure such that the inclusion $i: N \to M$ is an embedding.

This is equivalent to the definition using equivalence classes of embeddings: if $\phi_1: N_1 \to M$ and $\phi_2: N_2 \to M$ are equivalent embeddings (with common image $N \subset M$) then one can get a smooth atlas on N by constructing charts by precomposition with either $\phi_1^{-1}: N \to N_1$ or $\phi_2^{-1}: N \to N_2$ (since $\phi_1 \sim \phi_2$ the atlases so obtained will be equivalent), and with this atlas the inclusion $i: N \to M$ will be a distinguished member of the equivalence class of ϕ_1 and ϕ_2. So we can (and do) identify the equivalence class with this distinguished member.

Accordingly let $N \subset M$ be a submanifold, with $i: N \to M$ the inclusion. In particular i is an immersion, so for each $x \in N$ we have an induced injective linear map $i_*: T_xN \to T_xM$. We can then identify T_xN with its image under this map—in other words, for every $x \in N$ we have a natural identification of T_xN with a subspace of T_xM.

Theorem 1.4. *If $N \subset M$ is a submanifold where $\dim N = n$ and $\dim M = m$ and $x_0 \in N$, there exists a coordinate chart $\phi: U \to \mathbb{R}^m$ for M such that $x_0 \in U$ and $\phi^{-1}(\mathbb{R}^n \times \{\vec{0}\}) = N \cap U$.*

Proof. As can easily be seen from the definition of the subspace topology, there is a neighborhood $U_0 \subset M$ of x_0 which is the domain of a coordinate chart $\phi_0: U_0 \to \mathbb{R}^m$ for M, such that where $V_0 = U_0 \cap N$, V_0 is the domain of some coordinate chart $\psi_0: V_0 \to \mathbb{R}^n$ for N. By replacing ϕ_0 and ψ_0 by their compositions with translations we may as well assume that $\phi_0(x_0)$ and $\psi_0(x_0)$ are the origins of $\mathbb{R}^m$ and $\mathbb{R}^n$ respectively. Also, by composing ϕ_0 with an appropriate linear map, we may as well assume that the composition $\phi_0 \circ \psi_0^{-1}: \psi_0(V_0) \to \phi_0(U_0)$ (which is a smooth map with injective linearization from a neighborhood of the origin in $\mathbb{R}^n$ to a neighborhood of the origin in $\mathbb{R}^m$) has the property that its linearization at 0 is given by $\vec{v} \mapsto (\vec{v}, 0)$ where we split $\mathbb{R}^m$ as $\mathbb{R}^n \times \mathbb{R}^{m-n}$.

Now define a map $\alpha: \psi_0(V_0) \times \mathbb{R}^{m-n} \to \mathbb{R}^n$ by $\alpha(x, y) = (\phi_0 \circ \psi_0^{-1})(x) + (\vec{0}, y)$. This map α is C^∞, and its linearization at $(\vec{0}, \vec{0})$ is the identity. The inverse function theorem from multivariable calculus then asserts that α is a local diffeomorphism near $\vec{0}$, *i.e.* that there is a neighborhood W of the origin in $\mathbb{R}^m$ and a smooth map $\beta: \alpha(W) \to W$ so that $\beta \circ \alpha: W \to W$ and $\alpha \circ \beta: \alpha(W) \to \alpha(W)$ are the respective identities.

Now set $U = \phi_0^{-1}(\alpha(W)) \subset M$ and $\phi = \beta \circ \phi_0$. ϕ is a composition of two maps which are diffeomorphisms to their images in $\mathbb{R}^m$, so ϕ is a coordinate chart in M (in the maximal atlas for M). Moreover since by construction we have $\alpha(\psi_0(V_0) \times \{0\}) = \phi_0(N \cap U_0)$, we see that $\beta(\phi_0(N \cap U)) = W \cap (\psi_0(V_0) \times \{0\})$. In other words, ϕ maps the points of its domain which lie in N precisely to the points of its range (namely W) which lie in $\mathbb{R}^n \times \{0\}$, as desired. □

Remark 1.5. Conversely, suppose that $N \subset M$ is a subset such that every point $x_0 \in N$ is contained in a coordinate chart for M as in Theorem 1.4, so there is an M-neighborhood U for x_0 and a coordinate chart $\phi: U \to \mathbb{R}^m$ so that $\phi^{-1}(\mathbb{R}^n \times \{0\}) = N \cap U$. For any such coordinate chart ϕ, the restriction $\phi|_{N \cap U}$ is a homeomorphism to an open subset of $\mathbb{R}^n \times \{0\} \cong \mathbb{R}^n$; this shows that N is a topological manifold. Moreover if $\phi_\alpha: U_\alpha \to \mathbb{R}^m$ and $\phi_\beta: U_\alpha: U_\beta \to \mathbb{R}^n$ are two such coordinate charts, so that by restricting ϕ_α, ϕ_β to $U_\alpha \cap N$ and $U_\beta \cap N$ we obtain homeomorphisms ψ_α, ψ_β from open subsets of N to open sets in $\mathbb{R}^n$, then the transition map $\psi_\beta \circ \psi_\alpha^{-1}$ is just the restriction to $\phi_\alpha(U_\alpha \cap (\mathbb{R}^n \times \{0\}))$ of $\phi_\beta \circ \phi_\alpha^{-1}$, which is smooth. This proves that such charts ψ_α give N the structure of a smooth manifold. Since in terms of the charts ψ and ϕ the inclusion $i: N \to M$ is just given my the inclusion of $\mathbb{R}^n$ into $\mathbb{R}^m$ as the first n coordinates, i is an immersion. Thus, as a converse to Theorem 1.4, a subset of M which can be covered by charts of the type described there is a submanifold of M.

If N is a submanifold of M and $\phi\colon U \to \mathbb{R}^m$ is a chart as in Theorem 1.4, note that if we define a map $f\colon U \to \mathbb{R}^{m-n}$ by taking the last $m-n$ coordinates of ϕ (*i.e.*, $f = (x_{m+1}, \dots, x_n)$), then $f\colon U \to \mathbb{R}^{m-n}$ is a submersion and $f^{-1}(\{\vec{0}\}) = N \cap U$. Moreover, within U, the tangent space to N is given by the kernel of the linearization of f. This is a sort of converse to an important method of constructing submanifolds.

To prepare for this, we make the following definitions:

Definition 1.6. Let $f\colon M \to P$ be a smooth map between two smooth manifolds.

- A *critical point* of f is a point $x \in M$ such that the linearization $f_*\colon T_xM \to T_{f(x)}P$ is *not* surjective.
- A *critical value* of f is a point $y \in P$ such that $y = f(x)$ for some critical point x of f.
- A *regular value* of f is any point $y \in P$ which is not a critical value.

Note in particular that a point $y \in P$ which is not in the image of f is still a regular value.

An important fact, which we will not prove, is the following:

Theorem 1.7 (Sard's Theorem). *If $f\colon M \to P$ is a smooth map between two smooth manifolds then the set of critical values of f has measure zero in P.*

(To make sense of this statement one has to know what "measure zero" means for a subset of a smooth manifold—to interpret this, note that a diffeomorphism between two open sets in Euclidean space preserves the class of sets of measure zero (even though it generally isn't measure preserving), so we can define a set of measure zero in a smooth manifold to be one whose intersection with the domain of every coordinate chart is mapped by that coordinate chart to a set of measure zero. If you prefer a statement that does not appeal to measure theory, it is also true that the set of regular values is residual in the sense of Baire—*i.e.*, it contains a countable intersection of open dense sets.)

Note that if $\dim M < \dim P$ and $f\colon M \to P$ is smooth, then since the linearization of f is never surjective every point of $f(M) \subset P$ is a critical value. So in this case Sard's theorem amounts to the statement that $f(M)$ has measure zero in P, *i.e.* that the image of f misses almost every point of P.

Whether we find a regular value of f by appealing to Sard's theorem or by directly examining the map, the following gives a useful way of producing submanifolds:

Theorem 1.8. *Let $f\colon M \to P$ be a smooth map between two smooth manifolds and let $y_0 \in P$ be a regular value. Then $N = f^{-1}(\{y_0\})$ is a submanifold of M. Moreover if $n_0 \in N$ then $T_{n_0}N = \ker(f_*\colon T_{n_0}M \to T_{f(n_0)}P)$. (In particular, $\dim N = \dim M - \dim P$.)*

Proof. Let $m = \dim M$, and $p = \dim P$, and $n_0 \in N$. Let $(y_1, \dots, y_p)\colon V \to \mathbb{R}^p$ be a coordinate chart for P around $f(n_0)$ which sends $f(n_0)$ to the origin For $i = 1, \dots, p$ define $z_i\colon f^{-1}(V) \to \mathbb{R}$ by $z_i = y_i \circ f$. By the surjectivity of f_* at n_0, we may choose tangent vectors $v_1, \dots, v_p \in T_{x_0}M$ so that $dz_i(v_j) = dy_i(f_*v_j) = \delta_{ij}$. Let $S \leq T_{n_0}M$ be the span of $v_1, \dots, v_p$. We may then choose linearly independent cotangent vectors $\alpha_{p+1}, \dots, \alpha_m \in T^*_{n_0}M$ so that each $\alpha_i|_S = 0$. Let $\phi = (x_1, \dots, x_m)\colon U \to \mathbb{R}^m$ be a coordinate chart around n_0. In terms of these coordinates, the cotangent vectors α_i at n_0 can be written as $\sum_{k=1}^m \alpha_{ki} dx_k$ for some real numbers α_{ki}. Define functions $z_{p+1}, \dots, z_m\colon U \to \mathbb{R}$ by $z_i = \sum_{k=1}^m \alpha_{ki} x_k$.

We now claim that the functions $z_1, \dots, z_m$ together provide a coordinate chart on a neighborhood of n_0. First note that the covectors $(dz_1)|_{n_0}, \dots, (dz_m)|_{n_0}$ are linearly independent elements of $T^*_{n_0}M$. For if $\sum c_i dz_i$ were to vanish at n_0, then by evaluating both sides on v_j for $j = 1, \dots, p$ we obtain that $c_1 = \cdots = c_p = 0$, from which it also follows that $c_{p+1} = \cdots = c_m = 0$ since we chose the $\alpha_i = (dz_i)|_{n_0}$ for $i \geq p+1$ to be linearly independent. Since the dz_i are linearly independent at n_0, they form a basis for $T^*_{n_0}M$ by a dimension count, and in particular there is a unique, bijective linear map of $T^*_{n_0}M$ which sends $(dz_i)_{n_0}$ to $(dx_i)_{n_0}$ for $i = 1, \ldots, m$. But the this linear map is the transpose of the Jacobian at $\phi(n_0)$ of the map which sends $(x_1, \ldots, x_m) \in \phi(U)$ to $(z_1, \ldots, z_m)$, and so the Jacobian of this map $F\colon (x_1, \ldots, x_n) \mapsto (z_1, \ldots, z_n)$ is invertible at $\phi(n_0)$. So by the inverse function theorem there is an open set W around $\phi(n_0)$ so that $F|_W$ is a diffeomorphism to its image. Recalling that $\phi = (x_1, \ldots, x_m)$ was a coordinate chart for M, it follows from this that $(z_1, \ldots, z_m)\colon \phi^{-1}(W) \to \mathbb{R}^n$ is a diffeomorphism to its image, and so it contained in the maximal atlas defining the smooth structure on M.

So given a point $n_0 \in N \subset M$, we have constructed coordinate charts $\tilde{\phi}\colon \phi^{-1}(W) \to \mathbb{R}^m$ around n_0 and $(y_1, \ldots, y_p)\colon V \to \mathbb{R}^p$ around $f(n_0)$ in terms of which the map f is given by $(z_1, \ldots, z_m) \mapsto (z_1, \ldots, z_p)$. In particular for any such coordinate chart $N \cap W$ is, in local coordinates, given by the preimage under the coordinate chart of $\{0\} \times \mathbb{R}^{m-p}$. By Remark 1.5, this suffices to establish that N is a submanifold of M. $\square$

Theorem 1.9. *Let K be a compact subset of a smooth manifold M. Then there exists an open set $V \subset M$ with $K \subset V$, a positive number q, and an embedding $\psi\colon V \to \mathbb{R}^q$. In particular if M is a compact manifold then there is an embedding $\psi\colon M \to \mathbb{R}^q$ for some q.*

Remark 1.10. In fact the compactness assumption is not necessary—any smooth manifold M embeds into Euclidean space of some dimension q, and indeed a result called the Whitney Embedding Theorem implies that one can take $q = 2\dim M$ (Whitney also showed that $\mathbb{R}P^{2^k}$ does not embed in any Euclidean space of dimension less than $2\cdot 2^k$, so this is generally the best one can do).

Proof. Write $m = \dim M$. Any point $x \in K$ is contained in the image of a surjective coordinate chart $\phi^{(x)}\colon U^{(x)} \to B^m(2)$ with $\phi^{(x)}(x) = \vec{0}$ (where $B^m(2)$ denotes the open ball of radius 2 around the origin in $\mathbb{R}^m$). If we write $V^{(x)} = (\phi^{(x)})^{-1}(B^m(1))$, then the $V^{(x)}$ still cover K, and so by compactness they have a finite subcover $\{V^{(x_1)}, \ldots, V^{(x_n)}\}$. Rename the $V^{(x_i)}, U^{(x_i)}$, and $\phi^{(x_i)}$ as V_i, U_i, ϕ_i. For each i let $\chi_i\colon M \to [0,1]$ be a smooth function such that $\chi_i^{-1}(1) = \bar{V}_i$ and which is supported in U_i. Also, define $\psi_i\colon M \to \mathbb{R}^n$ by $\psi_i(x) = \chi_i(x)\phi_i(x)$ if $x \in U_i$ and $\psi_i(x) = 0$ otherwise; of course this is smooth since χ_i is supported in U_i. Now define

$$\psi\colon M \to \mathbb{R}^{n(m+1)}$$

by

$$\psi(x) = \big(\chi_1(x), \psi_1(x), \chi_2(x), \psi_2(x), \ldots, \chi_n(x), \psi_n(x)\big)$$

I claim that the restriction of ψ to the open subset $V = \cup_{i=1}^n V_i$ is an embedding.

First, $\psi|_V$ is an immersion. For if $x \in V$ then $V \in V_i$ for some i, and since m of the coordinates of $\psi(x)$ are given by $\psi_i(x)$ and $\psi_i\colon V_i \to B^m(1)$ is (the restriction of) a coordinate chart, if we had $\psi_* v = 0$ for some $v \in T_xM$ then it would hold that $(\psi_i)_* v = 0$ and so $v = 0$. We must now show that $\psi|_V$ is a homeomorphism to its image. Of course ψ is continuous since all of its coordinates are. To see that $\psi|_V$ is injective, suppose that $\psi(x) = \psi(y)$. For some i we have $x \in V_i$, so $\chi_i(x) = 1$, and so $\chi_i(y) = 1$. We chose χ_i to be 1 precisely on $\bar{V}_i$, so this forces $y \in \bar{V}_i$. But then since $\psi(x) = \psi(y)$ and $x, y \in \bar{V}_i \subset U_i$ we have $\phi_i(x) = \psi_i(x) = \psi_i(y) = \phi_i(y)$, forcing $x = y$ since ϕ_i is a coordinate chart on U_i.

Finally we must show that the inverse of $\psi|_V\colon V \to \psi(V)$ is continuous. Let $x \in V$; we should show that for any neighborhood W of x there is an open set in $\mathbb{R}^{n(m+1)}$, containing $\psi(x)$, whose preimage under ψ is contained in W. To do this, let i be such that $x \in V_i$, and let $\epsilon > 0$ be small enough that the preimage under ϕ_i of the ball of radius 2ϵ around $\phi_i(x)$ is contained in $W \cap V_i$.

If $\delta > 0$, there is an open set $W' \subset \mathbb{R}^{n(m+1)}$ so that

$$(\psi|_V)^{-1}(W') = \{y \in V | \chi_i(y) > 1-\delta, |\psi_i(y) - \psi_i(x)| < \epsilon\}.$$

If we take $\delta = \frac{2}{2+\epsilon}$, any $y \in (\psi|_V)^{-1}(W')$ belonging to this latter set will obey

$$\begin{aligned}|\phi_i(y) - \phi_i(x)| &= \left|\frac{1}{\chi_i(y)}\phi_i(y) - \phi_i(x)\right| \\ &\leq \left|\frac{1}{\chi_i(y)} - 1\right| |\phi_i(y)| + |\phi_i(y) - \phi_i(x)| < \frac{\epsilon}{2}2 + \epsilon = 2\epsilon\end{aligned}$$

and so y will belong to W, as desired.

□

2. Vector bundles and tubular neighborhoods

A vector bundle E of rank k over a smooth manifold M is, to be brief (and to leave out some important details), a family of vector spaces E_x parametrized by the points $x \in M$.

More precisely:

Definition 2.1. Let M be a smooth manifold, and k a positive integer. A (smooth, real) *vector bundle of rank k over M* is a smooth map $\pi\colon E \to M$ where E is a smooth manifold, with the following additional structure

- For all $x \in M$, the preimage $\pi^{-1}(\{x\})$ (also denoted E_x) has the structure of a real vector space of dimension k.
- There is an open cover $\cup_\alpha U_\alpha$ of M and, for each α, a diffeomorphism $\Phi_\alpha\colon \pi^{-1}(U_\alpha) \to U_\alpha \times \mathbb{R}^k$ such that, for each $x \in U_\alpha$, Φ_α restricts to E_x as a linear isomorphism to the vector space $\{x\} \times \mathbb{R}^k$.

For a definition more closely analogous to our definition of a smooth manifold, and in order to resolve concerns about uniqueness, one could insist that the collection of transition functions is maximal; just as in the smooth manifold case any collection of local trivializations as in the definition can be enlarged in a unique way to a maximal such collection.

The vector space E_x is called the *fiber of E at x*, and the Φ_α are called *local trivializations* of E over U_α. Note that the map $\pi\colon E \to M$ is automatically a surjective submersion.

The smooth manifold E carries a distinguished copy of M embedded inside it as the *zero section* 0_M, whose intersection with each E_x consists of just the zero element of the vector space E_x (in terms of the local trivializations, $0_M = \cup_\alpha \Phi_\alpha^{-1}(U_\alpha \times \{0\})$).

Example 2.2. *If M is a smooth m-dimensional manifold then the* tangent bundle $\pi\colon TM \to M$ *is a vector bundle of rank m. For all intents and purposes we showed this in the first part of the course: M is covered by coordinate charts* $(x_1^\alpha, \dots, x_n^\alpha)\colon U_\alpha \to \mathbb{R}^m$, *and for a local trivialization of* TM *over* U_α *we can take the inverse of the map* $U_\alpha \times \mathbb{R}^k \to \pi^{-1}(U_\alpha)$ *which sends* $(x, v_1, \dots, v_m)$ *to the tangent vector* $\sum_{i=1}^m v_i \frac{\partial}{\partial x_i}$ *at x.*

Example 2.3. *Let* $f: N \to M$ *be a smooth map between two smooth manifolds, and let* $\pi\colon E \to M$ *be a vector bundle. We can then form the* pullback bundle $\Pi\colon f^*E \to N$ *as follows. Set theoretically, define*

$$f^*E = \left\{(n, e) \in N \times E | e \in E_{f(n)}\right\}.$$

We have local trivializations $\Phi_\alpha\colon \pi^{-1}(U_\alpha) \to U_\alpha \times \mathbb{R}^k$; *for* $x \in U_\alpha$ *define* $\phi_{\alpha x}\colon E_x \to \mathbb{R}^k$ *to be the linear isomorphism such that for* $e \in E_x$ *we have* $\Phi_\alpha(e) = (x, \phi_{\alpha x}(e))$. *Now for each* α *define* $\Psi_\alpha\colon \Pi^{-1}(f^{-1}(U_\alpha)) \to f^{-1}(U_\alpha)$ *by, for* $e \in \Pi^{-1}(n)$ *where* $n \in f^{-1}(U_\alpha)$, *setting* $\Psi_\alpha(e) =$

$(n, \phi_{\alpha f(n)}(e))$. *Now the various $f^{-1}(U_\alpha)$ cover N, and it's not hard to check that there is a unique smooth structure imposed on f^*E by requiring that the maps Ψ_α be diffeomorphisms. This gives the map $\Pi\colon f^*E \to N$ the structure of a vector bundle, and we have a commutative diagram*

$$\begin{tikzcd} f^*E \arrow[r] \arrow[d, "\Pi"] & E \arrow[d, "\pi"] \\ N \arrow[r, "f"] & M \end{tikzcd}$$

where the upper map just sends $(n, e) \in f^*E \subset N \times E$ to e, and maps fibers of f^*E isomorphically to fibers of E.

As an important special case, we can let f be the inclusion of a submanifold $N \subset M$. *In this case f^*E is more often just denoted by $E|_N$. In particular, we have $TM|_N$, the restriction of the tangent bundle of the ambient manifold M to the submanifold N; its fiber over $n \in N$ consists of the whole tangent space T_nM.*

Example 2.4. *If $N \subset M$ is a submanifold we have the vector bundles TN and $TM|_N$; these give rise to a third vector bundle over N, the* normal bundle $\nu_{N,M} \to N$, *whose fiber over a point $n \in N$ is naturally identified with $\frac{T_nM}{T_nN}$. Perhaps the easiest way of constructing this bundle is to make use of the adapted coordinate charts from Theorem 1.4. We cover a neighborhood of N in M by charts $\phi_\alpha\colon U_\alpha \to \mathbb{R}^m$, such that for each α we have $U_\alpha \cap N = \phi_\alpha^{-1}(\{0\} \times \mathbb{R}^n)$. So the $\psi_\alpha := \phi_\alpha|_{U_\alpha \cap N}$ form an atlas for N. Let $\pi_1\colon \mathbb{R}^m \to \mathbb{R}^{m-n}$ be the projection onto the first $m-n$ coordinates. So if $v \in T_nN$ where $n \in V_\alpha$, then since ϕ_α sends N to $\{0\} \times \mathbb{R}^n$, we will have (borrowing the notation of the previous example) $\pi_1 \circ \phi_{\alpha n} v = 0$. Thus $\pi_1 \circ \phi_{\alpha n}$ descends to a linear isomorphism from $(\nu_{N,M})_n$ to $\mathbb{R}^{m-n}$. Consequently the $\pi_1 \circ \phi_{\alpha n}$ give rise to local trivializations over V_α for $\nu_{N,M}$, confirming that $\nu_{N,M}$ is a vector bundle (again, to get the smooth manifold structure on $\nu_{N,M}$ one can just require that these local trivializations are diffeomorphisms).*

Note that, if $\dim M = m$ *and* $\dim N = n$, *then the rank of $\nu_{N,M}$ is $m-n$, and so the dimension of $\nu_{N,M}$ as a smooth manifold is $n+(m-n) = m$, the same as the dimension of the ambient manifold. The tubular neighborhood theorem (Theorem 2.11) will show that, in fact, $\nu_{N,M}$ is diffeomorphic to an open neighborhood of N in M.*

Remark 2.5. It is possible to formulate the notion of a subbundle of a vector bundle, and then show quite generally that if $F \leq E$ is a subbundle then one can form the quotient bundle E/F (with fiber over x canonically identified with E_x/F_x). In the case of a submanifold $N \subset M$, one can show that TN is a subbundle of $TM|_N$, and so the normal bundle can be identified with the quotient bundle of the latter by the former.

Definition 2.6. An *orthogonal structure* on a vector bundle $\pi\colon E \to M$ is a map

$$\langle \cdot, \cdot \rangle\colon \cup_{x \in M} (E_x \times E_x) \to \mathbb{R}$$

whose restriction to each $E_x \times E_x$ defines an inner product on E_x, and such that whenever $s_1, s_2\colon M \to E$ are two smooth sections (i.e. smooth maps so that $\pi \circ s_i = 1_M$), the map $x \mapsto \langle s_1(x), s_2(x) \rangle$ is smooth.

In other words, an orthogonal structure is a smoothly varying family of inner products on the fibers of E. In the case that E is the tangent bundle TM of M an orthogonal structure on TM is called a Riemannian metric on M. (Indeed, sometimes one uses the term "Riemannian metric" to refer to an orthogonal structure on any vector bundle.)

Proposition 2.7. *If $\pi\colon E \to M$ is a vector bundle there exists a orthogonal structure on E.*

Proof. Recall that the vector bundle structure on E gives us an open cover $\{U_\alpha\}$ of M and diffeomorphisms $\Phi_\alpha\colon \pi^{-1}(U_\alpha) \to U_\alpha \times \mathbb{R}^k$ which commute with the projections to U_α and restrict to the fibers E_x as a linear isomorphism $\phi_{\alpha x}$ to $\mathbb{R}^k$. So if we denote the standard inner product on $\mathbb{R}^k$ by $(\cdot,\cdot)_0$ we can define $\langle,\cdot,\cdot\rangle_\alpha\colon \cup_{x\in U_\alpha} E_x \times E_x \to \mathbb{R}$ by, for $e_1, e_2 \in E_x$, setting $\langle e_1, e_2\rangle_\alpha = \left(\phi_{\alpha x}e_1, \phi_{\alpha x}e_2\right)_0$.

Now let $\{\chi_\alpha\}$ be a partition of unity subordinate to the cover $\{U_\alpha\}$ and define $\langle\cdot,\cdot\rangle\colon \cup_{x\in E} E_x \times E_x \to \mathbb{R}$ to be equal to $\sum_\alpha \chi_\alpha(x)\langle\cdot,\cdot\rangle_\alpha$, where we have extended $\chi_\alpha(x)\langle\cdot,\cdot\rangle_\alpha$ by zero outside of U_α. Since convex combinations of inner products on vector spaces are still inner products, it's easy to see that this satisfies the requirements. $\square$

Note that an orthogonal structure $\langle\cdot,\cdot\rangle$ on E gives rise to a smooth function $\|\cdot\|^2\colon E \to [0,\infty)$ defined by $\|e\|^2 = \langle e, e\rangle$. The square root of this function, $\|\cdot\|\colon E \to \mathbb{R}$, is smooth on the complement of the zero section.

Using orthogonal structures one can show:

Proposition 2.8. *Let $\pi\colon E \to M$ and let U be any neighborhood of the zero section 0_M. Then there is an open set V with $0_M \subset V \subset U$ and a diffeomorphism $\psi\colon E \to V$ which restricts to the identity on 0_M.*

Thus the entire total space of a vector bundle can be shrunk by a diffeomorphism to an arbitrarily small neighborhood of the zero-section. This is basically a parametrized version of the statement that $\mathbb{R}^k$ is diffeomorphic to an arbitrarily small ball around the origin.

Proof of Proposition 2.8. This is easier if M is compact, since then there is $r > 0$ such that the open set $V = E_r = \{e \in E | \|e\|^2 < r^2\}$ is contained in U. (Proof: The subset $\bar{E}_1 := \{\|e\|^2 \leq 1\}$ is in this case also compact, as one can see by writing it as a union of compact sets obtained from a finite cover by local trivializations, so $\bar{E}_1 \setminus U$ is also compact. If the statement were false then one could find a sequence $e_i \in \bar{E}_1 \setminus U$ with $\|e_i\|^2 \to 0$. But since $\bar{E}_1 \setminus U$ is compact a subsequence of the e_i would converge to some e, which would have the contradictory properties that $\|e\| = 0$ and $e \notin U$.) In this case one can choose a diffeomorphism $f\colon [0,\infty) \to [0,1)$ such that f is equal the identity on a neighborhood of the origin (for instance, take $f(t) = \int_0^t g(s)ds$ where $g(s) = 1$ for small s, $g(s) > 0$ for all s and $\int_0^\infty g(s) = 1$). Then define $\psi\colon E \to V$ by

$$\psi(e) = rf(\|e\|)\frac{e}{\|e\|}$$

(so ψ rescales each fiber in such a way that a point with norm n now has norm $rf(n)$). This is easily seen to satisfy the required properties (the only place where smoothness either of ψ or ψ^{-1} might seem to be an issue is at the zero section, but in fact we have arranged for ψ to be just scalar multiplication by r on a neighborhood of the zero section).

If M is noncompact then there might not be a single number $r > 0$ as above. However, we shall construct below a smooth function $r\colon M \to (0,\infty)$ so that $V := \{e \in E | e \in E_x \Rightarrow \|e\|^2 < r(x)^2\}$ is contained in the given open set U. If we can do this, then a simple modification of the ψ constructed above works: just define $\psi(e) = r(x)f(\|e\|)\frac{e}{\|e\|}$ for $e \in E_x$; since r is smooth and positive (as, therefore, is $\frac{1}{r}$) this ψ will be a diffeomorphism from E to V just as before.

To construct the desired $r\colon M \to (0,\infty)$ we can just proceed as follows. Cover M by open sets O_β such that $\bar{O}_\beta$ is compact. Then for each β there will, as earlier, be a number $r_\beta > 0$ so that if $x \in \bar{O}_\beta$ and $e \in E_x$ has $\|e\| < r_\beta$ then $e \in U$. Now let $\{\chi_\beta\}$ be a partition of unity subordinate to $\{O_\beta\}$ and let $r(x) = \sum_\beta r_\beta \chi_\beta(x)$. For each x, $r(x)$ will then be a convex combination of those r_β with $x \in O_\beta$, and hence will be less than or equal to one of them, so any $e \in E_x$ with $\|e\| < r(x)$ will lie in U. $\square$

Definition 2.9. Let $N \subset M$ be a submanifold. A *tubular neighborhood of N in M* consists of an open subset $U \subset M$ with $N \subset U$, and a diffemorphism $\Phi\colon\ \nu_{N,M} \to U$, where $\nu_{N,M}$ is the normal bundle of N in M, such that the restriction of Φ to the zero section $N \cong 0_N \subset \nu_{N,M}$ is the identity map to N.

Remark 2.10. In view of Proposition 2.8, to construct a tubular neighborhood it is enough to construct a diffeomorphism $\Phi'\colon\ U' \to U$ restricting as the identity on N, where $U' \subset \nu_{N,M}$ is some neighborhood of the zero section. For then we can find a subneighborhood $V \subset U'$ and a diffeomorphism $\psi\colon\ \nu_{N,M} \to V$ as in Proposition 2.8, and then $\Phi' \circ \psi$ will give a tubular neighborhood (with image $\Phi'(U')$, which will still be an open neighborhood of N in M).

The rest of this section will be concerned with proving the following:

Theorem 2.11. *If $N \subset M$ is any compact submanifold then there exists a tubular neighborhood of N in M.*

(In fact, the compactness assumption is not strictly necessary–its main role in the proof given here will be to allow us to embed a neighborhood of N in M into $\mathbb{R}^q$ for some q, and as mentioned after Theorem 1.9 this can be done without the compactness assumption. Near the end of the proof we will also use the compactness of N to find the limit of a sequence, but one can get around this as long as one arranges for the embedding of M into $\mathbb{R}^q$ to be proper, as can be arranged in the Whitney embedding theorem.) To construct the tubular neighborhood, one needs some systematic way of "moving in directions normal to N in M." There are two common ways of doing this—either by choosing a Riemannian metric on M and using the theory of geodesics, or by embedding a neighborhood of N in M into Euclidean space and using the special structure of Euclidean space. To avoid a digression into Riemannian geometry, we'll take the latter approach.

Throughout the following discussion, for $x \in \mathbb{R}^q$ we will make the standard identification of $T_x\mathbb{R}^q$ with $\mathbb{R}^q$ (and so if $X \subset \mathbb{R}^q$ is a submanifold and $x \in X$ then T_xX is identified with a subspace of $\mathbb{R}^q$).

To begin the proof of the theorem note that we may as well replace M by a small neighborhood of the compact submanifold N, and then by Theorem 1.9 (applied with $K = N$) we can assume that M is embedded in $\mathbb{R}^q$. We can then define

$$\tilde{\nu}_{M,\mathbb{R}^q} = \{(x, v) \in M \times \mathbb{R}^q | v \in T_xM^\perp\}$$

and

$$\tilde{\nu}_{N,M} = \{(x, v) \in N \times \mathbb{R}^q | v \in (T_xM) \cap (T_xN)^\perp\}.$$

Note that there is a bijection $\alpha_{M,\mathbb{R}^q}\colon\ \tilde{\nu}_{M,\mathbb{R}^q} \to \nu_{M,\mathbb{R}^q}$, taking (x, v) to the equivalence class of v in $T_x\mathbb{R}^q/T_xM$. Similarly there is a bijection $\alpha_{N,M}\colon\ \tilde{\nu}_{N,M} \to \nu_{N,M}$ sending (x, v) to its equivalence class in T_xM/T_xN. These bijections commute with the projections to M (in the case of $\alpha_{M,\mathbb{R}^q}$) or N (in the case of $\alpha_{N,M}$). It should at least appear that $\tilde{\nu}_{M,\mathbb{R}^q}$ is a vector bundle over M, and likewise that $\tilde{\nu}_{N,M}$ is a vector bundle over N, and that these are isomorphic to the respective normal bundles. This is indeed true, but so far we have not even shown that $\tilde{\nu}_{M,\mathbb{R}^q}$ and $\tilde{\nu}_{N,M}$ are smooth manifolds. We now remedy this:

Lemma 2.12. *Let $M \subset \mathbb{R}^q$ be a submanifold, and $N \subset M$ a submanifold. Then $\tilde{\nu}_{M,\mathbb{R}^q}$ and $\tilde{\nu}_{N,M}$ are smooth manifolds, and the bijections $\alpha_{M,\mathbb{R}^q}\colon\ \tilde{\nu}_{M,\mathbb{R}^q} \to \nu_{M,\mathbb{R}^q}$ and $\alpha_{N,M}\colon\ \tilde{\nu}_{N,M} \to \nu_{N,M}$ are diffeomorphisms.*

Proof. Actually the statements about $\tilde{\nu}_{M,\mathbb{R}^q}$ are, after renaming, just special cases of those about $\tilde{\nu}_{N,M}$, but for clarity's sake we prove the results about $\tilde{\nu}_{M,\mathbb{R}^q}$ first. To show that $\tilde{\nu}_{M,\mathbb{R}^q}$ is a smooth manifold it suffices to show that for any point $m_0 \in M$ there is a neighborhood $U \subset M$ of m_0 so that $\tilde{\nu}_{M,\mathbb{R}^q} \cap (U \times \mathbb{R}^q)$ is a submanifold of $U \times \mathbb{R}^q$.

To do this, note that on a sufficiently small neighborhood $U \subset M$ we there will be smooth functions $a_{ij} \colon U \to \mathbb{R}$ $(1 \le i \le m, q \le j \le q)$ so that, for all $X \in U$,

$$T_x M = span\left\{\sum_{j=1}^{q} a_{ij}(m)\frac{\partial}{\partial x_j} : 1 \le i \le m\right\}$$

(for instance, one could take an adapted coordinate chart as in Theorem 1.4 and use for $\sum a_{ij}\frac{\partial}{\partial x_j}$ the vector fields that are mapped by the coordinate chart to the standard coordinate vector fields on $\mathbb{R}^m \times \{0\}$) So at each $x \in U$ the matrix $A(x) = \{a_{ij}(x)\}$ has full rank m. Define

$$\begin{aligned} F \colon U \times \mathbb{R}^q &\to \mathbb{R}^m \\ (x, v_1, \ldots, v_q) &\mapsto \left(\sum_{j=1}^{k} a_{1j}(x)v_j, \ldots, \sum_{j=1}^{k} a_{mj}(x)v_j\right). \end{aligned}$$

Identifying $T_{(x,v)}(M \times \mathbb{R}^q)$ with $T_x M \oplus \mathbb{R}^q$, we see that the linearization $F_* \colon T_{(x,v)}(M \times \mathbb{R}^q) \to \mathbb{R}^m$ has $F_*(0, w) = A(x)w$. In particular since $A(x)$ has full rank, F_* is surjective. By Theorem 1.8 this proves that $\tilde{\nu}_{M,\mathbb{R}^q} \cap (U \times \mathbb{R}^q)$ is a submanifold of $U \times \mathbb{R}^q$ for each member U of an open cover of a neighborhood of M in $\mathbb{R}^q$, and hence that $\tilde{\nu}_{M,\mathbb{R}^q}$ is a smooth manifold.

Moreover, inspection of the coordinate charts constructed in the proof of Theorem 1.8 shows that the smooth structure on $\tilde{\nu}_{M,\mathbb{R}^q}$ is consistent with that of $\nu_{M,\mathbb{R}^q}$ under the obvious bijection between them. Indeed, in the intersection of $\tilde{\nu}_{M,\mathbb{R}^q}$ with $U \times \mathbb{R}^q$ where U is a sufficiently small open set as in the previous paragraph, we can define a coordinate system whose first m coordinates (parametrizing M) are the same as those of an adapted coordinate chart for $M \subset \mathbb{R}^q$, and whose last $q - m$ coordinates depend only on the $\mathbb{R}^q$ factor. It's not hard to see that such a coordinate chart is diffeomorphic via $\alpha_{M,\mathbb{R}^q}$ to a corresponding local trivialization for $\nu_{M,\mathbb{R}^q}$ as described in Example 2.4. So since the bijection $\alpha_{M,\mathbb{R}^q}$ restricts to each member of an open cover as a diffeomorphism it is a diffeomorphism.

Now we turn to the slightly more complicated case of $\tilde{\nu}_{N,M}$. In this case, for any $n_0 \in N$ we can find a neighborhood of n_0 in $\mathbb{R}^q$ and smooth functions a_{ij} $(1 \le n + q - m,\ 1 \le j \le q)$ on U so that, for each $x \in U$, $T_x N$ is spanned by the $\sum_{j=1}^{q} a_{ij}(x)\frac{\partial}{\partial x_j}$ for $1 \le i \le n$, and $T_x M^{\perp}$ is spanned by the $\sum_{j=1}^{q} a_{ij}(x)\frac{\partial}{\partial x_j}$ for $n+1 \le i \le n+q-m$. Namely, as before we can use an adapted coordinate chart for the vector fields spanning $TN|_U$, while for the vector fields spanning $TM^{\perp}|_U$ we can start with a similar such basis of vector fields spanning $T_x M$ at every $x \in U$, extend this to a basis for $\mathbb{R}^q$ (say using vector fields with constant coefficients, and perhaps shrinking U in the process), and then modify this basis using the Gram-Schmidt procedure to get a basis for all of $T_x\mathbb{R}^q$ at every point of U consisting of smooth vector fields, the last $q - m$ of which span $T_x M^{\perp}$ at every $x \in U$.

Now since $T_x N \cap T_x M^{\perp} = \{0\}$ for all $x \in M$ (as $TN \subset TM$), our entire set of vector fields $\left\{\sum_{j=1}^{q} a_{ij}(x)\frac{\partial}{\partial x_j} : 1 \le i \le n + q - m\right\}$ is linearly independent at each $x \in U$. So just as before we can define

$$G\colon U\times\mathbb{R}^q\to\mathbb{R}^{n+m-q}$$
$$(x,v_1,\ldots,v_q)\mapsto\left(\sum_{j=1}^k a_{1j}(x)v_j,\ldots,\sum_{j=1}^k a_{mj}(x)v_j\right)$$

The preimage of 0 under this map consists of those pairs (x,v) where v is perpendicular both to the subspace T_xN and to the subspace $T_xM^\perp$, *i.e.* where $v\in T_xM\cap T_xN^\perp$, so $G^{-1}(\{0\})=\{(x,v)\in\tilde{\nu}_{N,M}|x\in U\}$. As in the case of $\tilde{\nu}_{M,\mathbb{R}^q}$, the linear independence of the vector fields that we have chosen implies that G_* is surjective, so $G^{-1}(\{0\})$ is a submanifold, and indeed following the proof of Theorem 1.8 we can take a coordinate system on $G^{-1}(\{0\})$ so that the first n coordinates depend only on the N factor and the last $m-n$ depend only on the $\mathbb{R}^q$ factor, so that in this coordinate system the projection $\tilde{\nu}_{N,M}\to N$ appears as the projection onto the first n coordinates. Allowing U to vary through sufficiently small open neighborhoods in $\mathbb{R}^q$ of points of N produces an atlas for $\tilde{\nu}_{N,M}$ each member of which can be seen as a local trivialization for the bundle $\tilde{\nu}_{N,M}\to N$. Once again, these trivializations are compatible under the bijection $\alpha_{N,M}$ with the standard normal bundle trivializations as given in Example 2.4, completing the proof. □

The following (when combined with Proposition 2.8) proves the tubular neighborhood theorem for submanifolds of $\mathbb{R}^q$, and will also be used in the more general case. To prepare for the statement, note that the space

$$\tilde{\nu}_{M,\mathbb{R}^q}=\{(x,v)|x\in M,v\in T_xM^\perp\}$$

of the previous lemma contains a distinguished "zero section" consisting of points of form $(x,0)$.

Lemma 2.13. *If $M\subset\mathbb{R}^q$ is a submanifold, define*

$$\epsilon_{M,\mathbb{R}^q}\colon\tilde{\nu}_{M,\mathbb{R}^q}\to\mathbb{R}^q$$

by

$$\epsilon_{M,\mathbb{R}^q}(x,v)=x+v.$$

Then there is a neighborhood V of the zero section of $\tilde{\nu}_{M,\mathbb{R}^q}$ such that $\epsilon_{M,\mathbb{R}^q}$ restricts to V as a diffeomorphism to its image, which is an open neighborhood of M in $\mathbb{R}^q$.

Proof. Writing $0_M=\{(m,0)\}$ for the zero section of $\tilde{\nu}_{M,\mathbb{R}^q}$, for any $(x,0)\in 0_M$ the tangent space $T_{(x,0)}\tilde{\nu}_{M,\mathbb{R}^q}$ splits naturally as $T_xM\oplus(T_xM)^\perp$, where the first factor is tangent to 0_M and the second is tangent to the fibers of the bundle projection $\tilde{\nu}_{M,\mathbb{R}^q}$. Of course, since T_xM is identified via the embedding as a subspace of $\mathbb{R}^q$, $T_xM\oplus(T_xM)^\perp$ in turn may be identified with all of $\mathbb{R}^q$. As should be clear from the definition of $\epsilon_{M,\mathbb{R}^q}$, with respect to these identifications the linearization of $\epsilon_{M,\mathbb{R}^q}$ at $(x,0)$ is just the identity from $\mathbb{R}^q$ to itself and in particular is invertible. So by the inverse function theorem every point 0_M has a neighborhood to which $\epsilon_{M,\mathbb{R}^q}$ restricts as a diffeomorphism to its image.

If $x\in M$ and $\delta>0$, define

$$V_{x,\delta}=\{(y,v)\in\tilde{\nu}_{M,\mathbb{R}^q}||y-x|+|v|<\delta\},$$

where $|\cdot|$ refers to the standard distance in Euclidean space $\mathbb{R}^q$. By the previous paragraph, for any $x\in M$ there is $\delta>0$ so that $\epsilon_{M,\mathbb{R}^q}|_{V_{x,\delta}}$ is a diffeomorphism to its image. So define a function $\delta\colon M\to\mathbb{R}$ by setting $\delta(x)$ equal to the supremum of all numbers δ such that $\epsilon_{M,\mathbb{R}^q}|_{V_{x,\delta}}$ is a diffeomorphism to its image. So evidently $\delta(x)>0$ for all x.

I now claim that $\delta\colon M \to \mathbb{R}$ is continuous. Indeed, one has a relationship

$$\delta(y) \geq \delta(x) - |x-y|,$$

resulting from the fact that $V_{y,\delta-|x-y|} \subset V_{x,\delta}$ for all $\delta > 0$. Combining this relationship with the same one where x and y are reversed shows that

$$|\delta(x) - \delta(y)| \leq |x-y|,$$

so δ is indeed continuous. Now define

$$V = \left\{ (x,v) \in \tilde{\nu}_{M,\mathbb{R}^q} \,\middle|\, |v| < \frac{1}{3}\delta(x) \right\}$$

V is then an open subset (since $(x,v) \mapsto |v| - \frac{1}{3}\delta(x)$ is continuous and V is the preimage of an open set under this map), and we will show that it has the property stated in the lemma.

The main issue is to show that $\epsilon_{M,\mathbb{R}^q}|_V$ is injective. So we must show that if $(x,v),(y,w) \in \tilde{\nu}_{M,\mathbb{R}^q}$ with $|v| < \frac{\delta(x)}{3}$, $|w| < \frac{\delta(y)}{3}$, and $x+v = y+w$, then $(x,v) = (y,w)$ Without loss of generality assume that $\delta(x) \leq \delta(y)$. Now the assumed relation $x+v = y+w$ is equivalent to $x - y = w - v$. But $|w-v| \leq |w| + |v| < 2\delta(x)/3$, and so

$$|x-y| + |w| = |w-v| + |w| < \delta(x).$$

So for some $\delta < \delta(x)$ we have $(x,v),(y,w) \in V_{x,\delta}$. But by the definition of $\delta(x)$, $\epsilon_{M,\mathbb{R}^q}$ restricts injectively to $V_{x,\delta}$ for all $\delta < \delta(x)$. So indeed $(x,v) = (y,w)$.

So we have shown that $\epsilon_{M,\mathbb{R}^q}$ restricts injectively to V. By the construction of V and by what was done at the start of the proof, V is covered by open sets on which $\epsilon_{M,\mathbb{R}^n}$ is a local diffeomorphism, and so $\epsilon_{M,\mathbb{R}^n}|_V$ is also continuous and open. Thus $\epsilon_{M,\mathbb{R}^q}|_V$ is a diffeomorphism to its image, which is open in $\mathbb{R}^q$. □

Corollary 2.14. *Let $M \subset \mathbb{R}^q$ be a submanifold. Then there is an open neighborhood W of M and a smooth map $r\colon W \to M$ so that $r|_M$ is the identity.*

(Indeed, r can be taken to be a deformation retraction, as you can check.)

Proof. Let V be a neighborhood of $0_M \subset \tilde{\nu}_{M,\mathbb{R}^q}$ as in Lemma 2.13, so that $\epsilon_{M,\mathbb{R}^q}\colon V \to \mathbb{R}^q$ is a diffeomorphism to its image. Denote this image by $W \subset M$. Then where $\pi\colon \tilde{\nu}_{M,\mathbb{R}^q} \to M$ is the bundle projection and where we identify M with 0_M, define $r\colon W \to M$ by $r = \pi \circ (\epsilon_{M,\mathbb{R}^q}|_V)^{-1}$. Since $\epsilon_{M,\mathbb{R}^q}$ restricts to $0_M \cong M$ as the identity, r is easily seen to satisfy the desired property. □

End of the proof of Theorem 2.11. We let $N \subset M$ be any compact submanifold, and by replacing M by a sufficiently small open set containing N and applying Theorem 1.9 we assume M to be embedded as a submanifold of $\mathbb{R}^q$. Where again

$$\tilde{\nu}_{N,M} = \{(x,v) \in N \times \mathbb{R}^q | v \in T_xM \cap (T_xN)^\perp\}$$

define $f_0\colon \tilde{\nu}_{N,M} \to \mathbb{R}^q$ by $f(x,v) = x + v$. Where r and W is as in Corollary 2.14, let $U_0 = f^{-1}(W)$, and define

$$f\colon U_0 \to M \qquad \text{by } f = r \circ f_0.$$

We have the zero section $0_N = \{(x,0)\} \subset \tilde{\nu}_{N,M}$; clearly for $(x,0) \in 0_N$, $f_0(x,0) = x \in N \subset M$ and so $f(x,0) = x$ also. Moreover, for $(x,0) \in 0_N$, $T_{(x,0)}U_0$ splits (compatibly with the splitting of $N \times \mathbb{R}^q$) as a direct sum $T_xN \oplus (T_xM \cap (T_xN)^\perp)$ (which is the same as T_xM), and with respect to this splitting the linearization $(f_0)_*\colon T_{(x,0)}U_0 \to T_x\mathbb{R}^q \cong \mathbb{R}^q$ acts as the inclusion. Now since the map r acts as the identity on M, and since $(f_0)_*$ sends $T_{(x,0)}U_0$ isomorphically to $T_xM \leq \mathbb{R}^q$, it follows by the chain rule that $f_* = r_* \circ (f_0)_*$ also sends $T_{(x,0)}U_0$ isomorphically to T_xM for all $(x,0) \in 0_N$. Thus around any point $(x,0) \in 0_N \subset U_0$ there is a neighborhood V_x to which $f\colon U_0 \to M$ restricts as a diffeomorphism to its image, which is open in M.

So much like the proof of Lemma 2.13 we should now show that, perhaps after shrinking the neighborhood U_0 of 0_N to some smaller neighborhood U_1, f restricts injectively to U_1. For this we exploit the compactness of N. If there were no neighborhood of 0_N to which f restricted injectively, we could find $(x_i, v_i), (y_i, w_i) \in U_0$ such that $(x_i, v_i) \neq (y_i, w_i)$ and $v_i, w_i \to 0$ but $f(x_i, v_i) = f(y_i, w_i)$ for all i. After passing to subsequences, the sequences x_i, y_i would converge in N by compactness, say to x and y, and we would have $f(x,0) = f(y,0)$ and hence $x = y$. But then (x_i, v_i) and (y_i, w_i) would eventually both lie in the neighborhood V_x from the previous paragraph, contradicting the fact that f is injective on that neighborhood. This contradiction shows that there is some neighborhood U_1 of 0_N, which we may as well take to be contained in $\cup_{x\in N} V_x$, such that $f|_{U_1}$ is injective.

Since $f\colon U_1 \to M$ is injective and U_1 is covered by sets to which f restricts as a diffeomorphism to its image, it follows that $f\colon U_1 \to f(U_1)$ is a global diffeomorphism to its image (since smoothness of f and of f^{-1} can be checked on these open sets).

This shows that a neighborhood U_1 of 0_N in $\tilde{\nu}_{N,M}$ is diffeomorphic to a neighborhood of N in M by a diffeomorphism restricting to the identity on M. By Remark 2.10 and the fact that $\tilde{\nu}_{N,M}$ is diffeomorphic to $\nu_{N,M}$ by a diffeomorphism acting as the identity on the zero section, this suffices to yield a tubular neighborhood $\Phi\colon \nu_{N,M} \to M$. □

3. Vector fields and flows

The following is a basic result from the theory of ordinary differential equations:

Theorem 3.1. *Let $F\colon \mathbb{R}^n \to \mathbb{R}^n$ be a compactly supported smooth function. Then for any $x_0 \in \mathbb{R}^n$ there is a unique solution $\gamma_{x_0}\colon \mathbb{R}\vec{\ }\mathbb{R}^n$ to the initial value problem*

$$\gamma'(t) = F(\gamma(t))$$
$$\gamma(0) = x_0$$

Moreover if $I \subset \mathbb{R}$ is any open interval, if $t_0 \in \mathbb{R}$, and if $\gamma\colon I \to \mathbb{R}$ obeys $\gamma'(t) = F(\gamma(t))$ and $\gamma(t_0) = \gamma_{x_0}(t_0)$, then $\gamma = \gamma_{x_0}|_I$. Furthermore, the map

$$\Phi\colon \mathbb{R} \times \mathbb{R}^n \to \mathbb{R}^n$$
$$(t,x) \mapsto \gamma_x(t)$$

is a smooth map.

Sketch of proof. (See [Lee, Chapter 17] for details.) First of all, suppose that we can show that there is $\epsilon > 0$ such that for every $x_0 \in \mathbb{R}^n$ and every $t_0 \in \mathbb{R}$ there is a solution $\gamma_{x_0}\colon (t_0 - \epsilon, t_0 + \epsilon) \to \mathbb{R}^n$ to $\gamma'(t) = F(\gamma(t))$ with $\gamma'(t_0) = x_0$, and such that any other solution γ on some subinterval of $(t_0 - \epsilon, t_0 + \epsilon)$ such that $\gamma(t_1) = \gamma_{x_0}(t_1)$ for some t_1 coincides is equal to γ_{x_0} everywhere. From this the existence of the *all-time* solution $\gamma_{x_0}\colon \mathbb{R} \to \mathbb{R}^n$ would follow. Indeed, we could initially apply the result to get a solution γ_0 on $(-\epsilon, \epsilon)$ with $\gamma_0(0) = x_0$. But we could then also get a solution γ_1 on $(-ep/2, 3\epsilon/2)$ with $\gamma_1(\epsilon/2) = \gamma_0(\epsilon/2)$. The uniqueness statement would then force γ_1 and γ_0 to be equal everywhere that they are both defined; hence they would combine to give a solution (still denoted γ_0) on $(-\epsilon, 3\epsilon/2)$. But there would also be a solution γ_2 on $(0, 2\epsilon)$ with $\gamma_2(\epsilon) = \gamma_0(\epsilon)$, and by uniqueness this then coincides with γ_0 everywhere, allowing the domain of γ_0 to be extended to $(-\epsilon, 2\epsilon)$. This can be repeated indefinitely, and the union of all of the solutions so obtained gives a map $\gamma_{x_0} \to \mathbb{R} \to \mathbb{R}^n$.

In other words, existence and uniqueness for all time (*i.e.*, all of the theorem except the last sentence) will follow if we can prove existence and uniqueness on all time intervals I of length at most 2ϵ for some fixed ϵ. We will do this by converting the differential equation to a fixed point problem and applying the contractive mapping principle. Namely, observe that the fundamental theorem of calculus implies that the following two statements about a map $\gamma\colon I \to \mathbb{R}^n$, where I is an interval containing a point t_0 are equivalent:

$$\gamma \text{ is a differentiable map such that } \gamma'(t) = X(\gamma(t)) \text{ and } \gamma(t_0) = x_0$$

and

$$\gamma \text{ is a continuous map such that } \gamma(t) = x_0 + \int_{t_0}^{t} F(\gamma(s))ds \text{ for all } t \in I$$

Let $C(I, \mathbb{R}^n)$ denote the space of continuous functions from I to $\mathbb{R}^n$, endowed with the uniform ("sup") norm. A standard fact in analysis is that $C(I, \mathbb{R}^n)$ is a Banach space (basically this is because a uniform limit of continuous functions is continuous). Define

$$\mathscr{A}\colon C(I, \mathbb{R}^n) \to C(I, \mathbb{R}^n)$$

by

$$\big(\mathscr{A}\gamma\big)(t) = x_0 + \int_{t_0}^{t} F(\gamma(s))ds$$

Now F was assumed compactly supported and smooth—in particular F is Lipschitz (actually F being Lipschitz is all that is needed for the conclusion of the theorem), *i.e.*, there is C such that $|F(x) - F(y)| \le C|x - y|$ for all $x, y \in \mathbb{R}^n$ (C can be taken to be the maximum norm of the gradient of F). An easy computation shows that, *provided the length of I is less than $\frac{1}{C}$*, the above map $\mathscr{A}$ is *contractive*, *i.e.*, there is $r < 1$ such that $\|\mathscr{A}\gamma - \mathscr{A}\eta\| \le r\|\gamma - \eta\|$. But the contractive mapping mapping principle asserts that a contractive map from a Banach space to itself always has a unique fixed point (if γ_0 is chosen arbitrarily and we define $\gamma_i = \mathscr{A}\gamma_{i-1}$, the fixed point is $\lim_{i=1}^{\infty} \mathscr{A}\gamma_i$). This precisely gives the desired existence and uniqueness of solutions on sufficiently short time intervals, and hence by the first paragraph proves all of the theorem except the last sentence.

The smoothness of Φ relies on some somewhat subtle estimates which can be found in [Lee]; I'll just prove the fact that Φ is continuous, which is a first step in the smoothness proof. First of all observe that for any smooth $u\colon \mathbb{R} \to \mathbb{R}^n$ such that $u(t)$ is nonzero for all t, one has, using the Cauchy–Schwarz inequality and the chain and product rules,

$$\frac{d}{dt}|u(t)| = \frac{d}{dt}\sqrt{u(t)\cdot u(t)} = \frac{2u(t)\cdot u'(t)}{2\sqrt{u(t)\cdot u(t)}} \le \frac{|u(t)||u'(t)|}{|u(t)|} = \left|\frac{du}{dt}\right|. \tag{1}$$

Now if $x, y \in \mathbb{R}^n$ are distinct points, by uniqueness of solutions we have $\gamma_x(t) \ne \gamma_y(t)$ for all t, so we can apply (1) with $u(t) = \gamma_x(t) - \gamma_y(t)$ to get

$$\begin{aligned}\frac{d}{dt}|\gamma_x(t) - \gamma_y(t)| &\le \left|\frac{d}{dt}\big(\gamma_x(t) - \gamma_y(t)\big)\right| = \big|F(\gamma_x(t)) - F(\gamma_y(t))\big| \\ &\le C|\gamma_x(t) - \gamma_y(t)|\end{aligned}$$

where as before C is the Lipschitz constant of F. Dividing by $|\gamma_x(t) - \gamma_y(t)|$ (which as noted earlier is nowhere zero) and recalling the identity $\frac{d}{dt}(\ln f) = \frac{f'}{f}$ then gives

$$\frac{d}{dt}\ln|\gamma_x(t) - \gamma_y(t)| \le C.$$

The Fundamental Theorem of Calculus (and then exponentiation of both sides) then shows that

$$|\gamma_x(t) - \gamma_y(t)| \le e^{Ct}|\gamma_x(0) - \gamma_y(0)|,$$

i.e.

$$|\gamma_x(t) - \gamma_y(t)| \le e^{Ct}|x - y|.$$

This shows that Φ is continuous as a function of x for fixed t. To take into account the varying of t we can just note that, if $D > 0$ is such that $|F(x)| \le D$ for all x (such D exists since we assumed F was compactly supported), then $|\gamma_x(s) - \gamma_x(t)| \le D|s - t|$. So we get

$$|\Phi(s,x) - \Phi(t,y)| \le |\Phi(s,x) - \Phi(t,x)| + |\Phi(t,x) - \Phi(t,y)| \le D|s - t| + e^{Ct}|x - y|,$$

and this proves that Φ is continuous at the (arbitrary) point (t, y).

As mentioned earlier, smoothness as opposed to continuity takes more work; I'll just mention that part of the idea is to differentiate the equation

$$\frac{\partial}{\partial t}(\Phi(t,x)) = F(\Phi(t,x))$$

which is satisfied by Φ with respect to t and/or x, in order to get a differential equation satisfied by a partial derivative of Φ; one can work inductively on the order of the derivative.

□

In other words, for compactly supported vector fields F on $\mathbb{R}^n$, there is always a unique integral curve of a vector field passing through any given point, and this curve varies smoothly with the point. This can easily be exported to smooth manifolds to yield the following corollary:[1]

Corollary 3.2. *Let M be a smooth manifold and let X be a compactly supported vector field on M. Then there is a unique family of diffeomorphisms, parametrized by $t \in \mathbb{R}$,*

$$\phi_X^t \colon M \to M$$

such that for all $m \in M$ we have

$$\phi_X^0(m) = m \text{ and } \frac{d}{dt}\phi_X^t(m) = X(\phi_X^t(m)).$$

These diffeomorphisms obey

$$\phi_X^t \circ \phi_X^s = \phi_X^{t+s} \tag{2}$$

and the map $(t, m) \mapsto \phi_X^t(m)$ is smooth.

The equation 2 should be easy to see: both sides represent the effect of starting at a point and flowing along the flow of the vector field for a time $t + s$. This equation is part of what leads to the conclusion that the ϕ_X^t are diffeomorphisms rather than just being smooth maps, since evidently ϕ_X^{-t} is an inverse to ϕ_X^t. The family $\{\phi_X^t\}$ is called the *flow* of the vector field X.

Remark 3.3. It is sometimes useful to allow the vector field X to itself depend on t, *i.e.* one can have a family of vector fields X_t varying with the parameter t. As long as this dependence is smooth and X_t has a uniform Lipschitz constant for all t then the proofs of Theorem 3.1 and Corollary 3.2 go through essentially without change in order to show that one still gets diffeomorphisms ϕ_X^t so that $\phi_X^0 = id_M$ and $\frac{d}{dt}\phi_X^t(m) = X_t(\phi_X^t(m))$. In fact, any smooth path of diffeomorphisms starting at the identity can be described as such a "time-dependent" flow—given such a path ϕ_t one can define $X_t(m) = \frac{d\phi_t}{dt}(\phi_t^{-1}(m))$ and, more or less tautologically, the flow of X_t will recover ϕ_t. Of course, for one of these time-dependent flows the homomorphism property (2) typically will not hold.

[1]One can approach the derivation of the corollary from Theorem 3.1 in either of a couple of different ways, either by directly working in local coordinate charts or by embedding M in $\mathbb{R}^q$ for some large q and using the tubular neighborhood theorem to construct a vector field on $\mathbb{R}^q$ which restricts to M as the given vector field X; details are left to you

Remark 3.4. If one drops the hypothesis that X is compactly supported (or, more generally, Lipschitz in a suitable sense) then Corollary 3.1 will no longer be true as stated. However a "local" statement can be made: for any $m \in M$ there will still be $\epsilon > 0$ and a neighborhood U of m in M on which there exists a unique "partial flow"

$$(-\epsilon,\epsilon) \times U \to M$$
$$(t,y) \mapsto \phi_X^t(y)$$

so that $\frac{d}{dt}\phi_X^t(y) = X(\phi_X^t(y))$ and $\phi_X^0(y) = y$. In other words, while a long-time existence result along the lines of Theorem 3.1 will typically fail, one still has uniqueness and short-time existence, for a time ϵ which depends on the point of interest in M.

The classic example of failure of long-time existence comes in the case $M = \mathbb{R}$, where the vector field X is given by $X(x) = x^2$. Thus the relevant differential equation is

$$\frac{dx}{dt} = x^2.$$

This equation can be solved by separation of variables to yield, where $x_0 = x(0)$,

$$x(t) = \frac{x_0}{1 - x_0 t}.$$

So for any x_0 we have a unique integral curve $x(t)$ through x_0, but this solution "blows up in finite time"—it ceases to be well-defined at time $t = \frac{1}{x_0}$ (but is a perfectly good solution until then).

3.1. **The Lie Derivative.** Given a (say compactly supported for convenience, but this is not really necessary for this section) vector field X on a smooth manifold M, the flow of X as described above provides a path of diffeomorphisms $\phi_X^t \colon M \to M$. The *Lie derivative* of a vector field or of a differential form along X is meant to be a measurement of how that vector field or differential form changes as one moves along the flow of X.

We'll start with the definition for vector fields:

Definition 3.5. Let X and Y be vector fields on M The Lie derivative of Y along X is the vector field $\mathscr{L}_X Y$ whose value at a point $m \in M$ is the element of $T_m M$ defined by

$$(\mathscr{L}_X Y)_m = \lim_{t \to 0} \frac{(\phi_X^{-t})_*(Y_{\phi_X^t(m)}) - Y_m}{t}$$

Note that this definition makes sense: recall that ϕ_X^{-t} is inverse to ϕ_X^t, and therefore we have a map $(\phi_X^{-t})_* \colon T_{\phi_X^t(m)} M \to T_m M$. Thus the two tangent vectors in the numerator belong to the same vector space, namely $T_m M$.

There is a similar definition for differential forms, but actually it can be rewritten in a somewhat simpler way because one moves from $T^*_{\phi_X^t(m)} M$ to $T^* M$ by pullback by the map ϕ_X^t. So if $\omega \in \Omega^p(M)$ and we wish to compare $\omega_{\phi_X^t(m)}$ to ω_m we can hit the first of these with the transpose of the linearization of ϕ_X^t. But recall that pullback of differential forms was defined precisely to so that the differential form $((\phi_X^t)^*\omega)_m$ would be equal to the result of applying the transpose of the linearization of ϕ_X^t to $\omega_{\phi_X^t(m)}$. So we define:

Definition 3.6. Let $\omega \in \Omega^p(M)$ be a differential form and let X be a vector field on M. The Lie derivative of ω along X is the differential p-form defined by

$$\mathscr{L}_X \omega = \lim_{t \to 0} \frac{(\phi_X^t)^*\omega - \omega}{t} = \left.\frac{d}{dt}\right|_{t=0} (\phi_X^t)^*\omega$$

Remark 3.7. It can be shown (either directly from the definition or from the formulas that we are about to prove) that our definitions of the Lie derivative of a vector field and of a differential form are compatible in the following sense. Suppose that $X, Y_1, \ldots, Y_p$ are vector fields and ω is a p-form. Then $\omega(Y_1, \ldots, Y_p)$ is a smooth function, *i.e.* a 0-form, so we can take its Lie derivative along X. On the other hand we can take the Lie derivatives along X of ω and of the Y_i. These obey the Leibniz rule:

$$\mathscr{L}_X\left(\omega(Y_1, \ldots, Y_p)\right) = (\mathscr{L}_X\omega)(Y_1, \ldots, Y_p) + \sum_{j=1}^{p} \omega\left(Y_1, \ldots, Y_{j-1}, \mathscr{L}_X Y_j, Y_{j+1}, \ldots, Y_p\right).$$

The definition of the Lie derivative along X makes it look somewhat impossible to compute; however we will presently give formulas which allow it to be quite easily computed from local coordinate expressions of X and of the object being differentiated. We start with 0-forms, and remind the reader that a vector field can be viewed as a derivation on the space of C^∞ functions; in particular if X is a vector field and $f \in C^\infty(M)$ we have a well-defined function Xf.

Proposition 3.8. *If $f \in \Omega^0(M) = C^\infty(M)$ then $\mathscr{L}_X f = Xf$.*

Proof. For any point $m \in M$ we have, using the chain rule

$$\begin{aligned}(\mathscr{L}_X f)(m) &= \left.\frac{d}{dt}\right|_{t=0} (\phi_X^{t*} f)(m) = \left.\frac{d}{dt}\right|_{t=0} (f \circ \phi_X^t)(m) \\ &= \left.\frac{d}{dt}\right|_{t=0} f\left(\phi_X^t(m)\right) = df\left(\left.\frac{d}{dt}\right|_{t=0} \phi_X^t(m)\right) = df(X_m) = (Xf)_m\end{aligned}$$

□

Recall that, since vector fields are derivations on $C^\infty(M)$, they have well-defined commutators ($[X,Y] = X \circ Y - Y \circ X$), which are also vector fields. Interestingly, commutators fit into the story of Lie derivatives:

Theorem 3.9. *If X and Y are vector fields on M then $\mathscr{L}_X Y = [X,Y]$*

Proof. Let $f \in C^\infty(M)$ and $m \in M$; we are to show that $\big((\mathscr{L}_X Y)(f)\big)(m) = \big(X(Yf)\big)(m) - \big(Y(Xf)\big)(m)$.

We see (recalling that an element of, *e.g.* T_mM is a derivation from the algebra of germs of C^∞ functions around m to $\mathbb{R}$, so if $v \in T_mM$ and $f \in C^\infty(M)$ we have a number $v(f)$):

$$\begin{aligned}\big((\mathscr{L}_X Y)(f)\big)(m) &= \lim_{t\to 0} \frac{((\phi_X^{-t})_* Y_{\phi_X^t(m)})f - Y_m f}{t} = \lim_{t\to 0} \frac{Y_{\phi_X^t(m)}(f \circ \phi_X^{-t}) - Y_m f}{t} \\ &= \left.\frac{d}{dt}\right|_{t=0} Y_{\phi_X^t(m)}(f \circ \phi_X^{-t})\end{aligned}$$

where in the first inequality we have used the definition of the pushforward in terms of derivations: $(\phi_* v)(f) = v(f \circ \phi)$. Now define a function of two variable H by

$$H(s,t) = (f \circ \phi_X^{-t})(\phi_Y^s(\phi_X^t(m))).$$

We observe

$$\frac{\partial H}{\partial s}(0,t) = d(f \circ \phi_X^{-t})(Y_{\phi_X^t(m)}) = Y_{\phi_X^t(m)}(f \circ \phi_X^{-t}).$$

Combining this we the previous displayed equation we see that

$$\big((\mathscr{L}_X Y)(f)\big)(m) = \frac{\partial^2 H}{\partial t \partial s}(0,0).$$

Now (setting $u = -t$ and using the chain rule)[2]

$$\begin{aligned}\frac{\partial^2 H}{\partial t \partial s}(0,0) &= \left.\frac{\partial^2 H}{\partial s \partial t}\right|_{(0,0)} f(\phi_Y^s(\phi_X^t(m))) - \left.\frac{\partial^2 H}{\partial s \partial u}\right|_{(0,0)} f(\phi_X^u(\phi_Y^s(m))) \\ &= \left.\frac{\partial}{\partial t}\right|_{t=0} \left(\left.\frac{\partial}{\partial s}\right|_{s=0} f \circ \phi_Y^s \right)(\phi_X^t(m)) - \left.\frac{\partial}{\partial s}\right|_{s=0} \left(\left.\frac{\partial}{\partial u}\right|_{u=0} f \circ \phi_X^u \right)(\phi_Y^s(m)) \\ &= \mathscr{L}_X(Yf)_m - \mathscr{L}_Y(Xf)_m = \big((XY - YX)f\big)(m),\end{aligned}$$

where in the last inequality we use Proposition 3.8. Since f and m are arbitrary this proves that $\mathscr{L}_X Y = XY - YX$. □

We now turn to the Lie derivative on differential forms; just as with vector fields there turns out to be a rather simple formula, which is quite useful for geometric applications. First we observe:

Lemma 3.10. *The Lie derivative $\mathscr{L}_X$ (defined by $\mathscr{L}_X \omega = \left.\frac{d}{dt}\right|_{t=0} \phi_X^{t*}\omega$) is the unique linear map $\mathscr{L}\colon \Omega^*(M) \to \Omega^*(M)$ obeying the following properties:*

1. *For all $f \in \Omega^0(M) = C^\infty(M)$, $\mathscr{L} f = Xf$.*
2. *$d\mathscr{L}\omega = \mathscr{L}(d\omega)$ for all $\omega \in \Omega^*(M)$.*
3. *For all $\omega, \theta \in \Omega^*(M)$, $\mathscr{L}(\omega \wedge \theta) = (\mathscr{L}\omega) \wedge \theta + \omega \wedge (\mathscr{L}\theta)$, and*
4. *If U is an open set and $\omega, \omega' \in \Omega^*(M)$ are such that $\omega|_U = \omega'|_U$, then $(\mathscr{L}\omega)|_U = (\mathscr{L}\omega')|_U$.*

Proof. First we should check that $\mathscr{L}_X$ obeys properties (1)-(4).

Property (1) is Proposition 3.8.

For property (2), simply note that, since d commutes with pullback,

$$d\left(\phi_X^{t*}\omega - \omega\right) = \phi_X^{t*} d\omega - d\omega,$$

and then (2) follows by dividing by t and taking the limit as $t \to 0$

For property (3) we have

$$\begin{aligned}\left.\frac{d}{dt}\right|_{t=0} \phi_X^{t*}(\omega \wedge \theta) &= \left.\frac{d}{dt}\right|_{t=0} (\phi_X^{t*}\omega) \wedge (\phi_X^{t*}\theta) \\ &= \left.\frac{d}{dt}\right|_{t=0} (\phi_X^{t*}\omega) \wedge \theta + \omega \wedge (\phi_X^{t*}\theta) = (\mathscr{L}_X\omega) \wedge \theta + \omega \wedge (\mathscr{L}_X\theta).\end{aligned}$$

Property (4) is easily verified: for any point $m \in U$ we will have $\phi_X^t(m) \in U$ for sufficiently small t, and so $(\phi_X^{t*}\omega)_m = (\phi_X^{t*}\omega')_m$ for all sufficiently small t, from which the conclusion immediately follows.

It remains to show that properties (1)-(4) uniquely specify a linear map. If $\mathscr{L}$ is any map obeying (1)-(3), and if $f, g_1, \ldots, g_p \in C^\infty(M)$, then we will have $\mathscr{L} f = \mathscr{L}_X f$ and $\mathscr{L} g_i = \mathscr{L}_X g_i$ by (1), and then $\mathscr{L}(dg_i) = \mathscr{L}_X(dg_i)$ by (2), and then

$$\mathscr{L}\left(f\, dg_1 \wedge \cdots \wedge dg_p\right) = \mathscr{L}_X\left(f\, dg_1 \wedge \cdots \wedge dg_p\right)$$

by (3). So by linearity $\mathscr{L}$ and $\mathscr{L}_X$ coincide on any forms which are finite linear combinations of forms of the shape $f\, dg_1 \wedge \cdots \wedge dg_p$. Now we proved earlier (Proposition 4.19 of Part 1) that any differential form ω can be written as a locally finite sum of forms of the shape $f\, dg_1 \wedge \cdots \wedge dg_p$,

[2]Here and below I will make use of the following point (and similar ones) without comment: if $f(x, y)$ is some function and if we set $g(z) = f(-z, z)$, then the chain rule gives that $g'(0) = (\nabla f) \cdot \langle -1, 1\rangle = \frac{\partial f}{\partial y}(0,0) - \frac{\partial f}{\partial x}(0,0)$

i.e., M is covered by open sets on each of which ω is a finite linear combinations of forms of the shape $f\,dg_1\wedge\cdots\wedge dg_p$. Now if $\mathscr{L}$ (like $\mathscr{L}_X$) obeys condition (4) then the restriction of $\mathscr{L}\omega$ to any open set is determined by the restriction of ω to that set, so by considering the restriction of ω to the open sets in the cover in the previous paragraph we see that $\mathscr{L}\omega=\mathscr{L}_X\omega$. □

Accordingly if we find a simple formula for an operation obeying (1)-(4) above then we can deduce that $\mathscr{L}_X$ is given by that formula. To prepare for this, recall the operation of "interior multiplication" of a form by a vector field: For any vector field X we get a map $\iota_X\colon \Omega^p(M)\to\Omega^{p-1}(M)$ defined by

$$(\iota_X\omega)(v_1,\dots,v_{p-1})=\omega(X,v_1,\dots,v_{p-1}).$$

Here is the promised formula:

Theorem 3.11 (Cartan's Magic Formula). *For any $\omega\in\Omega^p(M)$ we have*

$$\mathscr{L}_X\omega=d\iota_X\omega+\iota_X d\omega.$$

Proof. We just have to show that $\mathscr{L}'_X:=d\iota_X+\iota_X d$ obeys condition (1)-(4) above.

If $f\in\Omega^0(M)$ we see that

$$\mathscr{L}'_X f=0+\iota_X df=df(X)=Xf,$$

confirming (1).

(2) holds, since both $d\mathscr{L}'_X$ and $\mathscr{L}'_X d$ are equal to $d\iota_X d$ (as $d^2=0$).

(4) is immediate from the definition of $\mathscr{L}'_X$.

So the only nontrivial part is (3). And this isn't too hard: the key point is (as I will leave you to verify, using formula (7) on p. 23 of part 1) the identity

$$\iota_X(\omega\wedge\theta)=(\iota_X\omega)\wedge\theta+(-1)^p\omega\wedge(\iota_X\theta)\quad\text{if }\omega\in\Omega^p(M).$$

Of course this is the same "anti-derivation" property as is satisfied by d. Combining these we get, if $\omega\in\Omega^p(M)$,

$$\begin{aligned}\big(d\iota_X+\iota_X d\big)(\omega\wedge\theta)&=d\left((\iota_X\omega)\wedge\theta+(-1)^p\omega\wedge(\iota_X\theta)\right)+\iota_X\left((d\omega)\wedge\theta+(-1)^p\omega\wedge(d\theta)\right)\\&=(d\iota_X\omega)\wedge\theta+(-1)^{p-1}\iota_X\omega\wedge d\theta+(-1)^p d\omega\wedge(\iota_X\theta)+(-1)^{2p}\omega\wedge d\iota_X\theta\\&\quad+(\iota_X d\omega)\wedge\theta+(-1)^{p+1}d\omega\wedge\iota_X\theta+(-1)^p\iota_X\omega\wedge d\theta+(-1)^{2p}\omega\wedge\iota_X d\theta,\end{aligned}$$

and after cancellation one ends up with precisely $\big(d\iota_X\omega+\iota_X d\omega\big)\wedge\theta+\omega\wedge\big(d\iota_X\theta+\iota_X d\theta\big)$, as desired. □

We have, by definition,

$$\mathscr{L}_X\omega=\left.\frac{d}{dt}\right|_{t=0}\phi_X^{t*}\omega;$$

to find the derivative at a time other than zero, we compute, using (2),

$$\left.\frac{d}{dt}\right|_{t=s}\phi_X^{t*}\omega=\lim_{h\to0}\frac{\phi_X^{(s+h)*}\omega-\phi_s^*\omega}{h}=\lim_{h\to0}\frac{\phi_X^{s*}(\phi_X^{h*}\omega-\omega)}{h}=\phi_X^{s*}\mathscr{L}_X\omega \tag{3}$$

Corollary 3.12. *Suppose that $\omega\in\Omega^*(M)$ is closed: $d\omega=0$. Then a vector field X has the property that $\phi_X^{t*}\omega=\omega$ for all t if and only if $d(\iota_X\omega)=0$.*

Proof. We have $\phi_X^{t*}\omega = \omega$ for all t if and only if $\frac{d}{dt}\Big|_{t=s}\phi_X^{t*}\omega = 0$ for all s, which by (3) is equivalent to $\phi_X^{s*}\mathscr{L}_X\omega = 0$ for all s, which of course is equivalent to $\mathscr{L}_X\omega = 0$. Cartan's Magic Formula reveals that this in turn is equivalent to

$$d\iota_X\omega + \iota_X d\omega = 0,$$

and of course the second term on the left is zero since we assume ω is closed. □

3.2. **Volume forms and the Moser argument.** From now on let M be a compact oriented n-manifold (without boundary). A *volume form* on M is by definition a differential form $\omega \in \Omega^n(M)$ which is nowhere zero. In particular since volume forms have top degree they are obviously automatically closed. If ω is a volume form, then for any open subset $U \subset M$, by restricting ω to U and then integrating we can define the volume of U:

$$vol_\omega(U) = \int_U \omega$$

(this is a finite number by virtue of the ambient manifold M being compact). A diffeomorphism ϕ is called *volume-preserving* (with respect to the volume form ω) if $\phi^*\omega = \omega$; this terminology is justified by recalling the behavior of integrals under pullbacks by diffeomorphisms: we have

$$vol_\omega(\phi(U)) = \int_{\phi(U)} \omega = \int_U \phi^*\omega = \int_U \omega = vol_\omega(U)$$

if ϕ is volume-preserving.

Corollary 3.12 shows how to construct many volume-preserving diffeomorphisms. Namely, the time-t flow of a vector field X will be volume-preserving provided that $d(\iota_X\omega) = 0$. To get a feel for this condition, note that we can write ω in local coordinates $(x_1, \dots, x_n)$ as

$$\omega = g\,dx_1 \wedge \cdots \wedge dx_n$$

for some smooth nowhere-zero function g. Then if a vector field X is given locally by $X = \sum_i f_i \frac{\partial}{\partial x_i}$, we will have

$$\iota_X\omega = \sum_i g f_i \iota_{\frac{\partial}{\partial x_i}} dx_1 \wedge \cdots \wedge dx_n = \sum_i (-1)^{i-1} g f_i dx_1 \wedge \cdots \wedge \widehat{dx_i} \wedge \cdots \wedge dx_n$$

and so

$$d(\iota_X\omega) = \left(\sum_i \frac{\partial}{\partial x_i}(g f_i)\right) dx_1 \wedge \cdots \wedge dx_n.$$

Thus the condition for the flow of X to preserve ω is just that

$$\sum_{i=1}^n \frac{\partial}{\partial x_i}(g f_i) = 0.$$

In case the function g is 1 (*i.e.*, ω restricts to the coordinate chart as the standard volume form $dx_1 \cdots \wedge \cdots dx_n$; it is actually always possible to find coordinates around any given point such that this holds—see Remark 3.14), this condition just reads that the divergence (in the standard multivariable calculus sense) of $X = \sum f_i \frac{\partial}{\partial x_i}$ should be zero. Thus *the flow of a divergence-free vector field preserves volume*.

It should be clear from the local coordinate formulas above that, given a volume form ω and *any* $\alpha \in \Omega^{n-1}(M)$, a unique vector field X can be chosen so that $\iota_X\omega = \alpha$ (this can be done locally in coordinate charts, and then the local solutions can be pieced together with a partition of unity). Of course, there are many closed $(n-1)$-forms (for instance, the derivative of any $(n-2)$-form will do), and so there are many vector fields X with $d\iota_X\omega = 0$. As such, from Cartan's Magic Formula we have seen that for any volume form on a compact oriented manifold there are many diffeomorphisms which preserve the volume form.

Now let us call two volume forms ω_0 and ω_1 on M *equivalent* if there is a diffeomorphism $f: M \to M$ so that $f^*\omega_1 = \omega_0$. In view of the behavior of the integral under pullbacks, if ω_1 is equivalent to ω_0 it is obviously necessary to have $\int_M \omega_1 = \int_M \omega_0$. An argument of Moser shows that this condition is also sufficient:

Theorem 3.13 ([Mos]). *Let M be a compact connected oriented manifold without boundary and $\omega_0, \omega_1 \in \Omega^n(M)$ two volume forms such that $\int_M \omega_0 = \int_M \omega_1$. Then there is a diffeomorphism $f: M \to M$ so that $f^*\omega_1 = \omega_0$.*

Sketch of proof. First of all note that since the only closed 0-forms on a connected manifold are the constants, the assumption says that the ω_i have the same integral when wedged with any closed 0-form. By Poincaré duality, this then implies that ω_0 and ω_1 represent the same cohomology class in $H^n(M)$. So there is some $\alpha \in \Omega^{n-1}(M)$ such that $\omega_1 = \omega_0 + d\alpha$. Now for $0 \leq t \leq 1$ let

$$\omega_t = \omega_0 + t\,d\alpha = (1-t)\omega_0 + t\omega_1.$$

The fact that ω_0 and ω_1 have equal integrals (or even just integrals of the same sign) means that they induce the same orientation on M. So if $m \in M$ and $\{e_1, \dots, e_n\}$ is a basis for T_mM such that $\omega_0(e_1, \dots, e_n) > 0$, then it will also hold that $\omega_1(e_1, \dots, e_n) > 0$. But then for all $t \in [0,1]$

$$\omega_t(e_1, \dots, e_n) = ((1-t)\omega_0 + t\omega_1)(e_1, \dots, e_n) > 0.$$

Since $m \in M$ was an arbitrary point this shows that *the $\omega_t = \omega_0 + t\,d\alpha$ are all volume forms*.

The plan now is to find a *time-dependent* vector field X_t on M so that where $\{\phi_t\}$ is the flow of X_t (*i.e.* $\phi_0 = id_M$ and $\frac{d}{dt}\phi_t(m) = X_t(\phi_t(m))$) we have $\phi_t^*\omega_t = \omega_0$ for all t. If we can do this then $f = \phi_1$ would be our desired diffeomorphism.

In this direction, Cartan's Magic Formula together with the chain rule can be seen to imply that, if X_t has flow ϕ_t:

$$\begin{aligned}\frac{d}{dt}(\phi_t^*\omega_t) &= \phi_t^*\frac{d\omega_t}{dt} + \phi_t^*\mathscr{L}_{X_t}\omega_t \\ &= \phi_t^*\left(d\alpha + d\iota_{X_t} + \iota_{X_t}d\omega_t\right) = \phi_t^* d\left(\alpha + \iota_{X_t}\omega_t\right)\end{aligned}$$

So one need only solve the equation $\iota_{X_t}\omega_t = -\alpha$, which by local coordinate considerations as above can be done in a unique way, producing a vector field X_t which depends smoothly on t. So indeed we can just set f equal to the time-one map of the flow of X_t. □

Remark 3.14. With sufficient care, one can localize this argument to show that for any volume form $\omega = g(x_1, \dots, x_n)dx_1 \wedge \cdots \wedge dx_n$ on a neighborhood U of the origin in $\mathbb{R}^n$, there are coordinates $(y_1, \dots, y_n)$ on a smaller neighborhood U' of the origin so that $\omega|_{U'} = dy_1 \wedge \cdots \wedge dy_n$. This justifies a statement made earlier that for any volume form, the manifold is covered by coordinate charts in which the volume form is given by the standard formula $dx_1 \wedge \cdots \wedge dx_n$.

References

[Lee] John M. Lee. *Introduction to Smooth Manifolds* GTM 218, Springer, 2003.
[Mos] Jürgen Moser. *On the volume elements of a manifold*. Trans. Amer. Math. Soc. **120** (1965), 286–294.