

SMOOTH MANIFOLDS

Notes to a course by
Eduard Looijenga

—

Fall 2010

Prerequisites

In this short chapter we review some basic notions and results from analysis in several variables and from elementary topology. It can be consulted as we go along, but occasionally it may be well worth returning to the sources where you learn(ed) these things.

1. What we need from calculus

We call a map f from an open subset U of \mathbb{R}^m to \mathbb{R}^n *differentiable at* $a \in U$ if it admits a first order approximation by a polynomial of degree one, more precisely, if there exists a linear map $\sigma \in \text{Hom}(\mathbb{R}^m, \mathbb{R}^n)$ with the property that

$$\lim_{h \rightarrow 0} \frac{\|f(a+h) - f(a) - \sigma(h)\|}{\|h\|} = 0$$

The map σ is unique, is called the *derivative* of f at a and usually denoted $D_a f$. The map $x \mapsto f(a) + \sigma(x-a)$ is the (best) first order approximation of f at a .

If f is differentiable in every point of U , then we have a matrix valued map $Df : U \rightarrow \text{Hom}(\mathbb{R}^m, \mathbb{R}^n)$ whose coefficients are the first order partial derivatives of f . We say that f is a C^1 -map if Df is continuous. And inductively, we say that f is a C^k -map ($k \geq 2$) if Df is a C^{k-1} -map. In order to make this also true for $k = 1$, we agree that C^0 stands for continuous. A map is C^k precisely when all its partial derivatives up to order k exist and are continuous. A C^∞ -map is a map which is C^k for all k .

THEOREM 1.1 (Chain rule). *Let f be a map from an open $U \subset \mathbb{R}^m$ to an open subset $V \subset \mathbb{R}^n$ and let g be a map from V to \mathbb{R}^p . Als f is differentiable at $a \in U$ and g is differentiable at $b := f(a)$, then gf is differentiable at a and we have $D_a(gf) = D_b g D_a f$.*

Notice that this theorem implies that if f and g are C^k , then so is gf .

A bijection f between two open subsets U and V of \mathbb{R}^m with the property that both f and f^{-1} are C^k -maps is called a C^k -diffeomorphism. In case $k \geq 1$, it follows from the chain rule that for every a , $D_a f$ is invertible with inverse $D_{f(a)}(f^{-1})$. The inverse function theorem yields a local converse:

THEOREM 1.2 (Inverse Function Theorem). *Let f be a $C^{k \geq 1}$ -map from an open neighborhood of $a \in \mathbb{R}^m$ to \mathbb{R}^m and suppose that $D_a f \in \text{Hom}(\mathbb{R}^m, \mathbb{R}^m)$ is nonsingular. Then f maps an open neighborhood of a C^k -diffeomorphically onto an open subset of \mathbb{R}^m .*

We formally derive the implicit function theorem (stated below) from the inverse function theorem as follows: write \mathbb{R}^{m+n} as a product $\mathbb{R}^m \times \mathbb{R}^n$ and denote its points accordingly: $x = (x', x'')$. Let f be a $C^{k \geq 1}$ -map from an open neighborhood U of $a = (a', a'') \in \mathbb{R}^{m+n}$ to \mathbb{R}^n and suppose that the restriction of

$D_a f \in \text{Hom}(\mathbb{R}^{m+n}, \mathbb{R}^n)$ to $0 \times \mathbb{R}^n$ (the map $v \in \mathbb{R}^n \mapsto D_a f(0, v) \in \mathbb{R}^n$) is non-singular. Then the map $\tilde{f} : U \rightarrow \mathbb{R}^{m+n}$, $\tilde{f}(x', x'') := (x', f(x', x''))$, satisfies the hypotheses of the inverse function theorem at a and hence maps a neighborhood $W \ni a$ C^k -diffeomorphically onto a neighborhood of $(a', f(b))$. Choose a product neighborhood $W' \times V$ of $(a', f(b))$ in $\tilde{f}(W)$ and replace W by $\tilde{f}^{-1}(W' \times V)$. Our gain is that $\tilde{f}(W)$ is now a product, namely $W' \times V$. The inverse of \tilde{f} , denoted $\tilde{g} : W' \times V \rightarrow W$, is a C^k -diffeomorphism which will have the form $\tilde{g}(x', y) = (x', g(x', y))$. So $(x', x'') \in W$ is mapped to $(x', y) \in W' \times V$ precisely when $y = f(x', x'')$, or equivalently, if $x'' = g(x', y)$. In particular:

- $\tilde{g}(x', y) = y$ and
- for every $y \in V$, $\tilde{f}^{-1}y \cap W$, is the graph of $x' \in W' \mapsto g(x', y) \in \mathbb{R}^n$.

The second property provides an *implicit* description of the map $g(-, y) : W' \rightarrow \mathbb{R}^n$ (whence the name implicit function theorem), whereas the first property tells us that if we compose f with a suitable diffeomorphism (namely \tilde{g}), we get a projection (namely the one of $W' \times V$ onto V). From this point of view the following formulation seems more natural:

COROLLARY 1.3. *Let f be a $C^{k \geq 1}$ -map from an open neighborhood of $a \in \mathbb{R}^{m+n}$ to \mathbb{R}^n and suppose that $D_a f \in \text{Hom}(\mathbb{R}^{m+n}, \mathbb{R}^n)$ is surjective. Then there exists a diffeomorphism h of an open neighborhood of a onto an open subset of \mathbb{R}^{m+n} such that fh^{-1} is on its domain the restriction of the projection $\mathbb{R}^{m+n} = \mathbb{R}^m \times \mathbb{R}^n \rightarrow \mathbb{R}^n$.*

PROOF. Since $D_a f$ is onto, there exist n basis vectors e_{i_1}, \dots, e_{i_n} of \mathbb{R}^{m+n} whose images under $D_a f$ make up a basis of \mathbb{R}^n . By renumbering coordinates we may just as well assume that these are the last n basis vectors of \mathbb{R}^{m+n} so that we find ourselves in the situation considered above. Then $h = \tilde{f}$ has the desired property, for $fh^{-1}(x', y) = y$. \square

There is a similar result for the case of an injective derivative:

PROPOSITION 1.4. *Let f be a $C^{k \geq 1}$ -map from an open neighborhood of $a \in \mathbb{R}^m$ to \mathbb{R}^{m+n} and suppose that $D_a f \in \text{Hom}(\mathbb{R}^m, \mathbb{R}^{m+n})$ is injective. Then there exists a diffeomorphism h from an open neighborhood of $f(a) \in \mathbb{R}^{m+n}$ onto an open neighborhood of $(a, 0)$ in $\mathbb{R}^m \times \mathbb{R}^n$ such that hf is a restriction of $\mathbb{R}^m \cong \mathbb{R}^m \times 0 \subset \mathbb{R}^m \times \mathbb{R}^n = \mathbb{R}^{m+n}$.*

PROOF. Since $D_a f$ is injective, the matrix of $D_a f$ has m linear independent rows, say rows i_1, \dots, i_m . By renumbering we may assume that these are the first m basis vectors. Then the composite of $D_a f$ with the projection $\mathbb{R}^m \times \mathbb{R}^n \rightarrow \mathbb{R}^m$ is nonsingular. Hence the map $F : U \times \mathbb{R}^n \rightarrow \mathbb{R}^{m+n}$, $F(x, y) := f(x) + (0, y)$, has the property that $D_{a,0}F(u, v) = D_a f(u) + (0, v)$, from which we see that $D_{a,0}F$ is nonsingular. By the inverse function theorem F maps a neighborhood of $(a, 0)$ in $U \times \mathbb{R}^n$ C^k -diffeomorphically onto a neighborhood of $f(a)$. If we take for h its inverse, so that $hF(x, y) = (x, y)$, then $hf(x) = (x, 0)$. Hence h is as desired. \square

Summing up: if a $C^{k \geq 1}$ -map f from a neighborhood of $a \in \mathbb{R}^p$ to \mathbb{R}^q has in a a derivative of maximal rank $\min\{p, q\}$, then f composed with some C^k -diffeomorphism at a (if $p \geq q$) resp. at $f(a)$ (if $p \leq q$) is the restriction of a linear projection resp. a linear inclusion. Phrased more suggestively (albeit less precise): the behavior of such a map at a is qualitatively no different from its derivative.

2. Recalling some notions from general topology

Topology addresses, among other things, the following question: if $f : X \rightarrow Y$ is a map between sets, what extra structure is needed on both X and Y in order to be able say that f is continuous? If we remember from calculus that a map f from a subset X of \mathbb{R}^m to \mathbb{R}^n is continuous precisely when the preimage of any open subset of \mathbb{R}^n is open in X (which means: the intersection of X with an open subset of \mathbb{R}^m), then it is clear that all that is required is a notion of ‘open set’.

Given a set X , then a *topology* on X is a collection \mathcal{O} of subsets of X that is closed under taking arbitrary unions and finite intersections and has both X and the empty set as members. A set equipped with a topology is called a *topological space* and the members of the topology are then referred to as *open subsets*. Although a topological space is strictly speaking a pair (like (X, \mathcal{O}) above), we often denote it by a single symbol X , for instance). A subset of a topological space is said to be *closed* if its complement is open (so the collection of such subsets is closed under taking arbitrary intersections and finite unions and contains both \emptyset and X).

The standard example is \mathbb{R}^m with its collection of usual open subsets. Recall what ‘usual’ means here: $U \subset \mathbb{R}^m$ is then open if for every $p \in U$ there is an open ball centered at that point and contained in U (the adjective open in *open ball* means that we take the ball defined by a strict inequality $\|x - p\| < r$, so the definition is not circular). In other words, a subset is open if and only if it is a union of open balls. If you check that the collection of subsets with this property is indeed a topology in the above sense, you readily discover that this is a formal consequence of two rather trivial properties of open balls: (i) every point is contained in an open ball and (ii) for every point in an intersection of two open balls there is an open ball centered at that point and contained in the intersection. This means that if you are given a set X and any collection \mathcal{B} of subsets of X

- (i) whose union is X and
- (ii) has the property that if $p \in B \cap B'$ with $B, B' \in \mathcal{B}$, then there is a $B'' \in \mathcal{B}$ with $p \in B'' \subset B \cap B'$,

then the subsets of X that can be written as a union of members of \mathcal{B} make up a topology on X . One calls this the *topology generated by \mathcal{B}* and \mathcal{B} is called a *basis* of this topology.

The notion of basis is quite helpful when we wish to give a product of topological spaces a topology: if X_1, \dots, X_m are topological spaces, then the collection of subsets of $U_1 \times \dots \times U_m$ with U_i open in X_i is usually not a topology on $X_1 \times \dots \times X_m$ because the union of two product subsets is in general not a product. What is however easily seen to be true is that this collection satisfies the properties (i) and (ii) above and hence generates a topology on $X_1 \times \dots \times X_m$. We call this the *product topology*. It is usually understood that $X_1 \times \dots \times X_m$ has been endowed with this topology; it is then called the *topological product* of the spaces X_i . For instance, \mathbb{R}^m is the topological product of m copies of \mathbb{R} .

Suppose X is a topological space.

A *neighborhood* of $p \in X$ is a subset $N \subset X$ for which there exists an open subset U with $p \in U \subset N$. More generally, a *neighborhood of a subset $A \subset X$* is subset $N \subset X$ for which there exists an open subset U with $A \subset U \subset N$.

Let $A \subset X$ be a subset.

The collection $\{A \cap U\}_{U \in \mathcal{O}}$ is a topology on A , called the *induced topology*. If a subset of X has been endowed with this topology we often refer to it as a *subspace*.

The *closure* of A , \bar{A} , the smallest closed subset of X containing A : $p \in \bar{A}$ if any neighborhood of p meets A . In case this closure equals X , then we say that A is *dense* in X . The *interior* of A , A° , is the biggest open subset of X contained in A : $p \in A^\circ$ if A is a neighborhood of p in X . It is clear that $A^\circ \subset A \subset \bar{A}$. The difference $\bar{A} - A^\circ$ is called the *boundary* of A in X , and sometimes denoted ∂A . (It is bit unfortunate that these notions have a slightly different meaning in the theory of manifolds, but there is no need to go into this now.)

We say that X is *connected* if X is nonempty and the only two open subsets that are also closed are X and \emptyset . It is proved in an analysis course that \mathbb{R} and in fact any interval in \mathbb{R} is connected.

We say that X is a *Hausdorff space* if X if any two of distinct points of X have disjoint neighborhoods. Notice that every subspace of Hausdorff space is also a Hausdorff space. For instance, \mathbb{R}^m and (hence) its subspaces are Hausdorff spaces.

We say that X is *compact* if for any collection $\{U_\alpha\}_\alpha$ of open subsets with union X (briefly, an *open covering* of X), there is a finite subcollection $U_{\alpha_1}, \dots, U_{\alpha_r}$ whose union is X (a *finite subcovering* of X). We know from analysis that the compact subspaces of \mathbb{R}^m are precisely those that are both closed and bounded and that such subsets have the property that every sequence in them has a convergent subsequence. This is an indication of this notion's nature and usefulness.

Let $f : X \rightarrow Y$ be map between topological spaces.

We make the expected definition and say that f is *continuous* if for every open $V \subset Y$, $f^{-1}V$ is open in X . It is easy to see that this property does not change if we replace open by closed. As noted above, if X is a subspace of \mathbb{R}^m and Y a subspace of \mathbb{R}^n , then this coincides with the notion familiar from calculus.

We say that f is *open* (resp. *closed*) if the image of every open (resp. closed) subset of X under f is open (resp. closed) in Y . For instance, the projection of a product onto a factor $X_1 \times X_2 \rightarrow X_1$ is open and continuous. It need not be closed: if we take $X_1 = X_2 = \mathbb{R}$, then the subset of $X_1 \times X_2$ defined by $x_1 x_2 = 1$ is closed, but its projection in X_1 , $\mathbb{R} - \{0\}$, isn't.

If f is bijective and both f and f^{-1} are continuous, then f is called a *homeomorphism*. Notice that f then also establishes a bijection between the open subsets of X and Y : $U \subset X$ is open in X if and only if $f(U)$ is open in Y and so this is the notion of isomorphism in topology: the topological spaces X and Y are then indistinguishable (for their topological properties) and we say that X and Y are *homeomorphic*. A deep theorem due to Brouwer says that a nonempty open subset of \mathbb{R}^m is not homeomorphic to an open subset of \mathbb{R}^n when $n \neq m$.

The proof of the items in the proposition below is no different from the special case (of subspaces of some \mathbb{R}^n) that you are familiar with.

- PROPOSITION 2.1. (i) *Every closed subset of a compact space is compact.*
 (ii) *Every compact subspace of a Hausdorff space is closed in that space.*
 (iii) *The image of compact (resp. connected) space under a continuous map is compact (resp. connected).*

PROOF. (i) Let X be a compact space and $A \subset X$ closed. An open covering of A can be given by a collection of open $U_\alpha \subset X$ with $\cup_\alpha U_\alpha \supset A$. Make it an open

covering of X by throwing in the open set $X - A$. Since X is compact, we can select a finite subcovering of X . This clearly yields a finite number of U_α 's covering A . So A is compact.

(ii) Let X be a Hausdorff space and $A \subset X$ compact. We show that $X - A$ is open in X by proving that every $p \in X - A$ has a neighborhood U disjoint with A . For every $q \in A$ choose disjoint X -open subsets $U_q \ni p, V_q \ni q$. The collection $\{V_q\}_{q \in A}$ is an open covering of A and so $A \subset V_{q_1} \cup \dots \cup V_{q_N}$ for certain q_1, \dots, q_N . Then $U := U_{q_1} \cap \dots \cap U_{q_N}$ is as desired.

(iii) Let $f : X \rightarrow Y$ be continuous. Suppose X compact. To prove that $f(X)$ is compact, let $\{V_\alpha\}_\alpha$ be a collection of Y -open subsets covering $f(X)$. Then $\{f^{-1}V_\alpha\}_\alpha$ is an open covering of X (for f is continuous) and hence has a finite subcovering: $X = f^{-1}V_{\alpha_1} \cup \dots \cup f^{-1}V_{\alpha_N}$ for certain $\alpha_1, \dots, \alpha_N$. Hence $f(X) \subset V_{\alpha_1} \cup \dots \cup V_{\alpha_N}$ and so $f(X)$ is compact.

Similarly it is proved that X connected implies $f(X)$ connected. \square

The third part gives among other things a generalization of the mean value theorem:

COROLLARY 2.2. *Let $f : X \rightarrow \mathbb{R}$ be a continuous function on a topological space X . If X is connected, then f is constant or $f(X)$ is an interval; if X is compact, then f has a minimum and a maximum; if X is both, then $f(X) = [a, b]$ for certain finite $a \leq b$.*

PROOF. If X is connected (resp. compact), then so is $f(X)$ and hence $f(X)$ is an interval or a singleton (resp. a bounded closed subset of \mathbb{R}). \square

Quite useful is:

COROLLARY 2.3. *Let $f : X \rightarrow Y$ be a continuous map from a compact space X to a Hausdorff space Y . Then f is closed. So if f is also bijective, then f is a homeomorphism.*

PROOF. If $A \subset X$ is closed, then A is compact (by 2.1-(i)), hence $f(A)$ is compact (by 2.1-(iii)) and therefore closed in Y (by 2.1-(ii)). If f is also bijective, then f being closed amounts to f^{-1} being continuous. \square

Let now $f : X \rightarrow Y$ be simply a map of sets.

If Y comes with a topology, then the collection $f^{-1}V$, with V running over the open subsets of Y defines a topology on X . We call it the *induced topology* on X (this is indeed a generalization of the case of an inclusion considered above); it is the smallest topology on X that makes f continuous.

EXERCISE 2.1. Let V be a real vector space of finite dimension m . If $f : V \rightarrow \mathbb{R}^m$ is an isomorphism of \mathbb{R} -vector spaces, then endow V with the f -induced topology. (a) Prove that this topology is independent of the choice of f so that V has a natural topology. (We will always assume that a real finite dimensional vector space is endowed with this topology.) (b) Prove that a linear map between finite dimensional real vector spaces is continuous. Prove also that the map is open resp. closed when it is onto resp. injective. (c) Prove that the topology on V as above is the smallest that makes all its linear functions $f : V \rightarrow \mathbb{R}$ continuous.

If, on the other hand, X (instead of Y) comes with a topology, then the collection of subsets $V \subset Y$ for which $f^{-1}V$ is open in X defines a topology on Y , sometimes called the *co-induced topology*. It is clearly the biggest topology on Y

that makes f continuous. When f happens to be surjective, we usually call it the *quotient topology*. Such a situation arises from any equivalence relation \sim on X : take $Y := X/\sim$ to be the set of equivalence classes and take for $f : X \rightarrow Y$ the map which assigns to x its equivalence class. Thus an equivalence relation \sim on a topological space X determines a topology on X/\sim .

Here is an example. Take $X = S^m$, the unit sphere in \mathbb{R}^{m+1} , and let \sim be the equivalence relation whose equivalence classes are the antipodal pairs $\{x, -x\}$. Then the quotient space is the *real projective m -space* P^m . This description of P^m is quite convenient if you are interested in its topology, but it does not very well exhibit its geometric origin as an extension of \mathbb{R}^m ‘with points at infinity’. The following exercise develops that point of view and at the same time helps you to familiarize yourself with the preceding.

EXERCISE 2.2. (a) Prove that P^m is connected and compact (you may use that S^m is connected for $m \geq 2$).

(b) Let $S_+^m \subset S^m$ be the closed northern hemisphere defined by $x_{m+1} \geq 0$ and identify its equator (defined by $x_{m+1} = 0$) with S^{m-1} . Prove that $S_+^m \subset S^m \rightarrow P^m$ is still surjective and the topology P^m picks up this way is the same topology as above.

(c) Prove that P^1 is homeomorphic to the circle S^1 .

(d) Prove that P^{m-1} is a subspace of P^m .

(e) Let $H \subset \mathbb{R}^{m+1}$ be the hyperplane obtained by lifting up \mathbb{R}^m in \mathbb{R}^{m+1} by the last unit vector: H is defined by $x_{m+1} = 1$. Prove that the map $x \in H \rightarrow x/\|x\|$ defines a homeomorphism of H onto $P^m - P^{m-1}$.

(f) Interpret the missing part $P^m - P^{m-1}$ in terms of H as the set of unoriented directions in H . More specifically, show that if $x \in H$ and $y \in S^{m-1}$, then the image of $x + ty \in H$ in P^m tends to the image of y in P^{m-1} as $t \rightarrow \pm\infty$. (In this way a line in H has a single point ‘at infinity’; that point lies in the ‘projective space at infinity’, P^{m-1} . Two lines in H have the same point at infinity if and only if they are parallel.)

Classical projective geometry is, as the name indicates, more concerned with geometry than topology and also tends to be more algebraic in character. In this setting a point of P^m is given by a nonzero vector in \mathbb{R}^{m+1} given up to scalar, rather than by an antipodal pair in S^m (there is no difference, of course, since any such vector spans a line which meets S^m in an antipodal pair). This leads to denoting a point of P^m by an $(m+1)$ -fold ratio, $[x^0 : \dots : x^m]$. A *projective linear subspace* of P^m of dimension k is simply the set of points coming from a $(k+1)$ -dimensional linear subspace of \mathbb{R}^{m+1} , but if $k = 1$ we rather speak of a (projective) line. Since any two distinct planes through the origin in \mathbb{R}^3 meet in a line, any two distinct lines in the projective plane P^2 will meet in a single point. This is in way simpler than what happens with distinct lines in an ordinary plane, as these can also be parallel. It is because of properties such as these that one often prefers to work with projective spaces rather than with affine spaces.

The language of manifolds

1. Topological manifolds

DEFINITION 1.1. An m -manifold (where m is an integer ≥ 0) is a Hausdorff space that is locally homeomorphic to \mathbb{R}^m (i.e., every point of this space has a neighborhood homeomorphic to an open subset of \mathbb{R}^m). We call m the *dimension* of M .

Before giving examples, we make a few general remarks. The space \mathbb{R}^m is Hausdorff and hence an m -manifold. The same must then also be true for any real finite dimensional vector space. Since a subspace of a Hausdorff space is also Hausdorff, any open subset of an m -manifold is an m -manifold as well; in particular, any open part of \mathbb{R}^m is one.

We remark here that there is a topological notion of dimension that assigns to a nonempty open subset of \mathbb{R}^m the number m (the definition is relatively simple, the proof anything but). This implies that a nonempty m -manifold cannot be at the same time an n -manifold if $n \neq m$.

Here are some more interesting examples.

EXAMPLES 1.2. (i) The m -sphere $S^m = \{x \in \mathbb{R}^{m+1} \mid \|x\| = 1\}$. Given $i = 0, \dots, m$ and $\epsilon \in \{+, -\}$ we have defined a half sphere $U_{i,\epsilon} := \{x \in S^m \mid \epsilon x^i > 0\}$. These half spheres form an open covering of S^m . The map $U_{i,\epsilon} \rightarrow \mathbb{R}^m$ which omits the i th coordinate is a homeomorphism onto the open unit ball $B^m \subset \mathbb{R}^m$. Hence S^m is a m -manifold.

(ii) The real projective space P^m , defined as the quotient of S^m by antipodal identification. Check that this is a Hausdorff space. Notice that the restriction of $S^m \rightarrow P^m$ to $U_{i,+}$ is injective and that its image U_i has preimage in S^m equal to $U_{i,+} \cup U_{i,-}$. This implies that U_i is open in P^m and that $U_{i,+} \rightarrow U_i$ is a homeomorphism. Hence P^m is a m -manifold.

(iii) Let f be a $C^{k \geq 1}$ -map from an open $U \subset \mathbb{R}^{m+n}$ to \mathbb{R}^n . Let $b \in \mathbb{R}^n$ and suppose that for every $a \in f^{-1}(b)$, Df_a is onto. Then it follows from Corollary 1.3 that $f^{-1}(b)$ is a m -manifold. (We may regard example (i) as a special case: take $f(x) := \|x\|^2$ and $b = 1$.)

(iv) This is a nonexample. Let M be obtained from two copies of \mathbb{R} , $\mathbb{R} \times \{1\}$ and $\mathbb{R} \times \{-1\}$ by identifying $(t, 1)$ and $(t, -1)$ for $t \neq 0$ (so M is a quotient of $\mathbb{R} \times \{1, -1\}$ by an equivalence relation). It is easy to verify that the image M_\pm of $\mathbb{R} \times \{\pm 1\}$ is open in M and that $\mathbb{R} \cong \mathbb{R} \times \{\pm 1\} \rightarrow M_\pm$ is a homeomorphism. Since $M = M_+ \cup M_-$, it follows that M is locally homeomorphic to \mathbb{R} . But M is not Hausdorff and hence not a manifold: if 0_\pm denotes the image of $(0, \pm)$ in M , and N_\pm is a neighborhood of 0_\pm , then N_\pm contains the image of $(-\epsilon_\pm, \epsilon_\pm) \times \{\pm 1\}$ for some $\epsilon_\pm > 0$ and it is then clear that $N_+ \cap N_-$ contains the image of $(0, \min\{\epsilon_+, \epsilon_-\}) \times \{1\}$.

Chart and atlas. Let M be a m -manifold. A *chart* for M is a couple (U, κ) , where U is an open part of M and $\kappa : U \rightarrow \mathbb{R}^m$ a homeomorphism of U onto an open subset of \mathbb{R}^m (the terminology is quite appropriate for this is like charting part of a territory).

Is given $p \in M$, then a *chart at* p simply a chart (U, κ) with $p \in U$ and $\kappa(p) = 0$.

Is given another chart (U', κ') , then we have two charts with domain $U \cap U'$ and the map

$$\kappa' \kappa^{-1} : \kappa(U \cap U') \rightarrow \kappa'(U \cap U')$$

is a homeomorphism between open subsets of \mathbb{R}^m . We call this homeomorphism a *coordinate change*.

An *atlas* for M is a collection charts $\mathcal{A} = (U_\alpha, \kappa_\alpha)_\alpha$ for M whose domains cover M : $\cup_\alpha U_\alpha = M$. If besides the atlas also all coordinate changes are given, then M may be reconstructed in much the same way as our globe can be recovered from satellite images: let \tilde{M} be the disjoint union of the open subsets $\kappa_\alpha(U_\alpha)$ of \mathbb{R}^m and obtain M from this as a quotient space: $a \in \kappa_\alpha(U_\alpha)$ is declared equivalent with $b \in \kappa_\beta(U_\beta)$ if $a \in \kappa_\alpha(U_\alpha \cap U_\beta)$ and $b \in \kappa_\beta(U_\alpha \cap U_\beta)$ and $\kappa_\beta \kappa_\alpha^{-1}(a) = b$. This is indeed an equivalence relation and the quotient space is M .

2. Smooth manifolds

Let M be an m -manifold. An atlas $\mathcal{A} = (U_\alpha, \kappa_\alpha)_\alpha$ for M is called a C^k -*atlas* if each of its coordinate changes is a C^k -map. Since the inverse of the coordinate change $\kappa_\beta \kappa_\alpha^{-1}$ is also one (namely $\kappa_\alpha \kappa_\beta^{-1}$) this implies that any coordinate change is then a C^k -diffeomorphism. It is clear that a C^k -atlas is also a C^l -atlas if $l < k$.

EXAMPLES 2.1. The atlases described in Examples 1.2-i, ii are C^∞ . The atlas in Example 1.2-iii that we get from an application of the implicit function theorem, is C^k .

Perhaps the main consequence of having a C^k -atlas \mathcal{A} is that it gives rise to a notion of C^k -differentiability for \mathbb{R} -valued functions defined on open subsets of M : we say that such a function $f : U \rightarrow \mathbb{R}$ is C^k *with respect to* \mathcal{A} if for every chart $(U_\alpha, \kappa_\alpha)$ of \mathcal{A} , $f \kappa_\alpha^{-1} : \kappa_\alpha(U_\alpha \cap U) \rightarrow \mathbb{R}$ is C^k .

PROPOSITION 2.2. For two C^k -atlases \mathcal{A} and \mathcal{B} of an m -manifold M are equivalent:

- (i) \mathcal{A} and \mathcal{B} determine the same notion of C^k -function: if $U \subset M$ is open, then $f : U \rightarrow \mathbb{R}$ is C^k with respect to \mathcal{A} if and only if it so with respect to \mathcal{B} ,
- (ii) $\mathcal{A} \cup \mathcal{B}$ is a C^k -atlas.

PROOF. Suppose (i) holds. We must show that for a chart (U, κ) from \mathcal{A} and a chart (V, λ) from \mathcal{B} the coordinate changes $\lambda \kappa^{-1}$ and $\kappa \lambda^{-1}$ are C^k . Since every component $\lambda_i : V \rightarrow \mathbb{R}$ of λ a C^k -function voor \mathcal{B} , it is, by our assumption, also one for \mathcal{A} . This means that $\lambda_i \kappa^{-1}$ is C^k ($i = 1, \dots, n$). In other words, $\lambda \kappa^{-1}$ is C^k . The same argument shows that $\kappa \lambda^{-1}$ is C^k .

The implication (ii) \Rightarrow (i) is left to you. □

If the equivalent properties of the preceding proposition are satisfied then we say that the two atlases are C^k -*equivalent*. Property (i) makes it plain that this is indeed an equivalence relation. Property (ii) shows that the union of all C^k -atlases belonging to the same C^k -equivalence class is itself also a C^k -atlas (and so there is a maximal one).

DEFINITION 2.3. A C^k -*structure* on a manifold M is the data of an equivalence class of C^k -atlases on M , or what boils down to the same: the data of a *maximal* C^k -atlas on M . A C^k -*manifold* is a manifold endowed with a C^k -structure; when $k = \infty$, we also call it a *smooth manifold* (but see Convention 2.8). A chart of a C^k -manifold is tacitly assumed to be taken from the C^k -atlas, unless otherwise stated.

Since S^m and P^m come in a natural manner with a C^∞ -structure, we will regard them as smooth manifolds. The M that appears in Example 1.2-iii comes with the structure of a C^k -manifold.

It is obvious that a C^k -manifold can be regarded as a C^l -manifold if $l < k$. A converse, due to H. Whitney, states that when $k \geq 1$, a C^k -atlas always contains a C^∞ -atlas (more than one usually), so that every C^k -manifold ‘comes from’ a C^∞ -manifold. On the other hand, R. Kirby and L. Siebenmann (1968) found manifolds that possess no C^1 -atlas at all!

It turns out that smooth manifolds offer the right setting for many results from calculus, although this sometimes forces us to state these results in a more geometric fashion (usually making them appear easier). But the chief reason for introducing this notion is its ubiquity: smooth manifolds are everywhere.

EXERCISE 2.1. Prove that the product of two C^k -manifolds is in a natural manner a C^k -manifold.

Differentiable maps. We already agreed on what it means for an \mathbb{R} -valued function on a C^k -manifold to be a C^k -function. A minor extension yields the notion of a C^k -mapping between C^k -manifolds:

DEFINITION 2.4. Let $f : M \rightarrow N$ be a continuous map between C^k -manifolds. We say that f is a C^k -map if for every chart (U, κ) of M and every chart (V, λ) of N , the composite $\lambda f \kappa^{-1}$ (whose domain is the open $\kappa(U \cap f^{-1}V) \subset \mathbb{R}^{\dim M}$ and whose range is $\mathbb{R}^{\dim N}$) is a C^k -map. We say that f is a C^k -*diffeomorphism* if f is bijective and both f and f^{-1} are C^k ; if such a C^k -diffeomorphism exists, then M and N are said to be C^k -*diffeomorphic*.

Notice that this definition makes every chart (U, κ) of a C^k -manifold a C^k -diffeomorphism of U onto $\kappa(U)$.

EXERCISE 2.2. Prove that a continuous map $f : M \rightarrow N$ between C^k -manifolds is a C^k -map if and only if it takes C^k -functions to C^k -functions in the sense that if ψ is an \mathbb{R} -valued C^k -function on an open $V \subset N$, then ψf is an \mathbb{R} -valued C^k -function (with domain $f^{-1}V$).

EXERCISE 2.3. Prove that two $C^{k \geq 1}$ -manifolds that are C^k -diffeomorphic have the same dimension. (You are not allowed to use the here unproven fact that dimension is a topological notion.)

EXERCISE 2.4. Prove that a composite of C^k -maps is a C^k -map.

H. Whitney showed that if two smooth manifolds are C^1 -diffeomorphic, then they are C^∞ -diffeomorphic. This is no longer true if the smooth manifolds are only homeomorphic: J. Milnor found in 1956 examples of smooth manifolds homeomorphic, but not diffeomorphic to S^7 .

DEFINITION 2.5. Given a C^k -manifold N of dimension n , then we say that a subset $M \subset N$ is a C^k -submanifold of dimension m if M can be covered by charts (U, κ) of N with the property that $M \cap U = \kappa^{-1}(\mathbb{R}^m \times \{0\})$. (If $k = \infty$ we will also call this is a smooth submanifold.)

Here we must have $m \leq n$, of course, and \mathbb{R}^n is written as $\mathbb{R}^m \times \mathbb{R}^{n-m}$. The definition says that the diffeomorphism $\kappa : U \rightarrow \kappa(U)$ maps $M \cap U$ onto $\kappa(U) \cap (\mathbb{R}^m \times \{0\})$. In other words, locally M lies in N in the same way as $\mathbb{R}^m \times \{0\}$ lies in \mathbb{R}^n . The restriction of a chart as in the above definition to M maps an open subset of M homeomorphically onto an open subset of \mathbb{R}^m . The collection of these defines a C^k -atlas for M , so that M acquires the structure of a C^k -manifold of dimension m . The inclusion map $M \subset N$ is automatically C^k .

EXAMPLE 2.6. Let f be a $C^{k \geq 1}$ -map from an open $U \subset \mathbb{R}^{m+n}$ to \mathbb{R}^n and let $b \in \mathbb{R}^n$ be such that the derivative of f is surjective in any point of $f^{-1}(b)$. Then Corollary 1.3 implies that $f^{-1}(b)$ is a C^k -submanifold of U of dimension m .

Applying this to $U = \mathbb{R}^{m+1}$ and $f(x) := \|x\|^2$, shows that S^m is a submanifold of \mathbb{R}^{m+1} .

EXERCISE 2.5. Prove that for $m \leq n$, S^m is a smooth submanifold of S^n if we identify S^m with $S^n \cap (\mathbb{R}^{m+1} \times 0)$. Conclude that likewise P^m is a smooth submanifold of P^n .

EXERCISE 2.6. Let $f : \mathbb{R}^n \rightarrow \mathbb{R}$ be a homogeneous function of degree d . Prove that $f^{-1}(1)$ is a (possibly empty) submanifold of dimension $n - 1$.

EXERCISE 2.7. The set of $n \times n$ -matrices of determinant 1 makes up a group under composition, called the *special linear group* $SL(n, \mathbb{R})$. Prove that this group is a smooth submanifold and determine its dimension. Prove also that the map $SL(n, \mathbb{R}) \times SL(n, \mathbb{R}) \rightarrow SL(n, \mathbb{R})$, $(\sigma, \tau) \mapsto \sigma\tau^{-1}$ is C^∞ . Do the same for the orthogonal group $SO(n)$. (These are examples of so-called Lie groups: a *Lie group* is a group G which at the same time has the structure of a smooth manifold for which the map $(g, h) \in G \times G \mapsto gh^{-1} \in G$ is C^∞ .)

EXERCISE 2.8. Let M be a set endowed with a collection \mathcal{A} of pairs (U, κ) with $U \subset M$ and κ a bijection of U onto an open subset of \mathbb{R}^m with the property that (i) each element of M is in the domain of a member of \mathcal{A} and (ii) when $(U, \kappa), (U', \kappa') \in \mathcal{A}$, then $\kappa(U \cap U')$ is open in \mathbb{R}^m and $\kappa'\kappa^{-1} : \kappa(U \cap U') \rightarrow \mathbb{R}^m$ is a C^k -map.

Prove that the collection of subsets $\Omega \subset M$ with the property that $\kappa(\Omega \cap U)$ is open in \mathbb{R}^m for every $(U, \kappa) \in \mathcal{A}$, is a topology on M . Prove that then for any $(U, \kappa) \in \mathcal{A}$, κ is a homeomorphism of U (regarded as a subspace of M) onto $\kappa(U)$. Finally show that M satisfies the Hausdorff property (hence becomes a C^k -manifold) if and only if for each pair $(U, \kappa), (U', \kappa')$ in \mathcal{A} , the image of $U \cap U'$ in $\kappa(U) \times \kappa(U')$ under (κ, κ') is closed.

You may wonder whether every smooth manifold is diffeomorphic to a smooth submanifold of some \mathbb{R}^N . The answer is that this is the case, provided a rather mild countability property is satisfied (we shall return to this issue). So smooth manifolds are less abstract you might have feared and your subsequent question then might well be: why bother if it does not produce anything we did not know? The reason for introducing the notion of a smooth manifold is not because it might

be *more general*, but because it is *more convenient* than the notion of a smooth submanifold of some \mathbb{R}^N : quite often the context makes a possible ambient \mathbb{R}^N irrelevant and hence a drag, and indeed, it may even happen that no ambient \mathbb{R}^N is given. This means in practice that the theory becomes much more transparent and effective in the abstract setting.

AN ASIDE 2.7. A (*smooth*) *knot* is a smooth submanifold of \mathbb{R}^3 that is diffeomorphic to S^1 . Knot theory is among other things concerned with the question of when two knots are ‘the same’. This sameness can be formalized in several ways, one is saying that $K \subset \mathbb{R}^3$ and $K' \subset \mathbb{R}^3$ are equivalent if there exists a diffeomorphism of \mathbb{R}^3 onto \mathbb{R}^3 that takes K to K' . Incidentally, one usually prefers to regard a knot as a submanifold of the S^3 (obtained from \mathbb{R}^3 by adding a single ‘point at infinity’); the advantage of this is that S^3 is compact, while the passage to S^3 hardly affects the sameness issue. Knot theory is intimately tied with group theory and noncommutative algebra. Some remarkable developments in this area during the past decades were inspired by quantum field theory.

From now on we confine ourselves to C^∞ -notions:

CONVENTION 2.8. *In what follows we tacitly assume that all manifolds and all maps between them are C^∞ , unless it is explicitly stated otherwise.*

In particular, an \mathbb{R} -valued function on a manifold is automatically C^∞ .

3. Tangent space

We now have agreed on what is a (C^∞ -) map $f : M \rightarrow N$ between smooth manifolds, but still missing is the basic notion of derivative. We would like it to be at a point $p \in M$ a *linear* map $D_p f$ deserving of that name. Domain and range of a linear map are vector spaces, and so the question arises: which vector spaces? The answer is going to be: $D_p f$ will go from the tangent space M at p to the tangent space of N at $f(p)$. This makes it our first business to define these tangent spaces.

We begin our discussion in \mathbb{R}^n . Is γ a curve in \mathbb{R}^n , i.e., a (C^∞ -) map from an open interval I to \mathbb{R}^n , then for every $t \in I$, the derivative of γ at t is a linear map $D_t \gamma : \mathbb{R} \rightarrow \mathbb{R}^n$ whose value in $1 \in \mathbb{R}$ is the vector

$$\dot{\gamma}(t) := \begin{pmatrix} \dot{\gamma}^1(t) \\ \dot{\gamma}^2(t) \\ \dots \\ \dot{\gamma}^n(t) \end{pmatrix},$$

where $\dot{\gamma}^i(t)$ is of course the derivative of γ^i at t . This Newtonian notation purports to suggest that we regard $\dot{\gamma}(t)$ as a velocity vector based at $\gamma(t)$. If M is given a (smooth) submanifold $M \subset \mathbb{R}^n$ and a $p \in M$, then we may restrict our attention to curves γ with image in M , defined on an (unspecified) neighborhood of $0 \in \mathbb{R}$ with $\gamma(0) = p$. The velocity vector $\dot{\gamma}(0)$ of such a curve will be tangent to M and if we let γ run over all such curves, the vectors $\dot{\gamma}(0)$ run over all vectors that are tangent to M at p . The collection of these velocity vectors is a linear subspace of \mathbb{R}^n translated over p which we may well call the ‘tangent space of M at p ’.

This notion can be transferred without much difficulty to the abstract case of a manifold M and a point $p \in M$: let us first make an auxiliary definition and

agree that a *curve at p* is a (smooth) map γ from an open neighborhood of 0 in \mathbb{R} to M with $\gamma(0) = p$. On the collection of curves at p we define an equivalence relation: $\gamma_1 \equiv \gamma_2$ if for a chart κ at p , $\kappa\gamma_1$ and $\kappa\gamma_2$ have the same derivative in \mathbb{R}^m . The equivalence class of γ relative to this equivalence relation will be called the *derivative of γ at 0* and will be denoted accordingly by $\dot{\gamma}(0)$. This relation is in fact independent of the choice of κ .

EXERCISE 3.1. Prove this.

DEFINITION 3.1. A *tangent vector of M at p* is the derivative at 0 of a curve at p . The collection tangent vectors of M at p is the *tangent space of M at p* and will be denoted T_pM .

We now want to show that the term *tangent vector* is justified in the sense that T_pM is a vector space in a natural way. We will see this by means of the following general argument: if we are given a set T , then every bijection $K : T \rightarrow \mathbb{R}^m$ can be used (in the most banal way) to give T the structure of a real vector space for which K is a linear isomorphism of vector spaces: subtraction and scalar multiplication in T are just transferred through K : $v - w := K^{-1}(K(v) - K(w))$ and $\lambda.v := K^{-1}(\lambda K(v))$. Another bijection $K' : T \rightarrow \mathbb{R}^m$ yields the same vector space structure on T precisely when $K'K^{-1} : \mathbb{R}^m \rightarrow \mathbb{R}^m$ is a linear isomorphism (check this!).

Returning to the case at hand, we note that for every chart κ at p , the map $\gamma \mapsto D_0(\kappa\gamma)(0) \in \mathbb{R}^m$ defines a bijection $T_pM \rightarrow \mathbb{R}^m$: the injectivity follows from Exercise 3.1 and every $v \in \mathbb{R}^m$ arises as the image of the curve $t \mapsto \kappa^{-1}(\kappa(p) + tv)$ at p . We denote this bijection by $D_p\kappa$, so that $D_p\kappa(\dot{\gamma}(0)) = (\kappa\gamma)'(0)$. We use it to transfer the vector space structure of \mathbb{R}^m to T_pM so that $D_p\kappa$ becomes a linear isomorphism. This structure is *independent of κ* :

PROPOSITION 3.2. *The tangent space of a manifold M at a point $p \in M$, T_pM , has a natural (so chart independent) structure of a vector space of the same dimension as M . It is characterized by the property that every chart κ at p defines an isomorphism of vector spaces $D_p\kappa : T_pM \rightarrow \mathbb{R}^m$. If κ' is another chart at p , then $D_p\kappa' = D_{\kappa(p)}(\kappa'\kappa^{-1})D_p\kappa$ (where $D_{\kappa(p)}(\kappa'\kappa^{-1})$ is the usual derivative of $\kappa'\kappa^{-1}$ at $\kappa(p)$).*

PROOF. We first address the last statement. Since $\kappa'\gamma = (\kappa'\kappa^{-1})(\kappa\gamma)$, the chain rule tells us that

$$(\kappa'\gamma)'(0) = D_{\kappa(p)}(\kappa'\kappa^{-1})((\kappa\gamma)'(0)).$$

This is just saying that $D_p\kappa' = D_{\kappa(p)}(\kappa'\kappa^{-1})D_p\kappa$, so that the last assertion follows. The first assertion follows from the preceding applied to $K = D_p\kappa$ and $K' = D_p\kappa'$. \square

We can now present the derivative geometrically as a linear map between tangent spaces. Given a map $f : M \rightarrow N$ between manifolds, $p \in M$ and $q := f(p)$, we define $D_p f : T_pM \rightarrow T_qN$ as follows: if $v \in T_pM$, then choose a representative curve γ in M at p (in other words, $v = \dot{\gamma}(0)$), so that $f\gamma$ is a curve in N at $f(p)$ and stipulate that $f\gamma$ represents $D_p f(v)$. We must however check that this only depends on v (i.e., not on the way we represented v by γ) and that the map is linear. To see this, let κ resp. λ be a chart of M at p , resp. of N at q . Then $\lambda f\gamma = (\lambda f\kappa^{-1})(\kappa\gamma)$ and hence following the chain rule

$$(\lambda f\gamma)'(0) = D_{\kappa(p)}(\lambda f\kappa^{-1})((\kappa\gamma)'(0)).$$

By definition, the left hand side is $D_q\lambda((f\gamma)'(0))$. We also see that the right hand side equals $D_{\kappa(p)}(\lambda f\kappa^{-1})D_p\kappa(v)$ and hence $(f\gamma)'(0)$ only depends on v . If we denote $(f\gamma)'(0)$ by $D_p f(v) \in T_q N$ as we intended, then the left hand side is written $D_q\lambda D_p f(v)$. Thus $D_p f : T_p M \rightarrow T_q N$ is well-defined and obeys a formula that looks like the chain rule:

$$D_q\lambda \circ D_p f = D_{\kappa(p)}(\lambda f\kappa^{-1}) \circ D_p \kappa$$

(but bear in mind that here only $D_{\kappa(p)}(\lambda f\kappa^{-1})$ is a derivative of the old-fashioned kind). We claim that the *derivative of f at p* , $D_p f : T_p M \rightarrow T_q N$, is linear. In view of the way we put vector space structures on $T_p M$ and $T_q N$, this amounts to stating that $D_q\lambda \circ D_p f \circ (D_p \kappa)^{-1}$ is linear. This is clearly the case, because this composite equals the ordinary derivative $D_{\kappa(p)}(\lambda f\kappa^{-1}) : \mathbb{R}^m \rightarrow \mathbb{R}^n$.

We finally note that our somewhat *ad hoc* definition of $D_p \kappa : T_p M \rightarrow \mathbb{R}^m$ agrees with the final one, if we bear in mind that the tangent space of \mathbb{R}^m at $f(p)$ can be identified with \mathbb{R}^m .

PROPOSITION 3.3. *A map $f : M \rightarrow N$ between manifolds determines for every $p \in M$ a linear map between vector spaces, the derivative of f at p , $D_p f : T_p M \rightarrow T_{f(p)} N$. This derivative obeys the chain rule: if $g : N \rightarrow P$ is another map between manifolds, then $D_p(gf) = D_{f(p)}g \circ D_p f$.*

PROOF. The verification of the chain rule is easy: if γ is a curve in M at p (representing $v \in T_p M$), then $f\gamma$ is one in N at $f(p)$ (representing $D_p f(v) \in T_{f(p)} N$) and $gf\gamma$ one in P at $gf(p)$ (representing $D_{f(p)}g \circ D_p f(v)$, but also $D_p(gf)(v)$). \square

4. Submersions, immersions and embeddings

The notions just introduced enable us to state the inverse and implicit function theorems in a more canonical fashion.

PROPOSITION 4.1. *Let $f : M \rightarrow N$ be a map from an m -manifold to an n -manifold and let $p \in M$, $q := f(p)$.*

- (subm) *If $D_p f : T_p M \rightarrow T_q N$ is surjective (so $m \geq n$), then there exist charts κ of M at p and λ of N at $f(p)$ such that $\lambda f\kappa^{-1}$ is on its domain equal to the projection $\mathbb{R}^m = \mathbb{R}^n \times \mathbb{R}^{m-n} \rightarrow \mathbb{R}^n$. In particular, a neighborhood of p in $f^{-1}(q)$ is a submanifold of M of dimension $m - n$.*
- (imm) *If $D_p f : T_p M \rightarrow T_q N$ is injective (so $m \leq n$), then there exist charts κ of M at p and λ of N at $f(p)$ such that $\lambda f\kappa^{-1}$ is on its domain equal to the injection $\mathbb{R}^m \cong \mathbb{R}^m \times 0 \subset \mathbb{R}^m \times \mathbb{R}^{n-m} = \mathbb{R}^n$. In particular, f maps an open neighborhood of p diffeomorphically onto a submanifold of N .*
- (diff) *If $D_p f : T_p M \rightarrow T_q N$ is bijective (so $m = n$), then f maps an open neighborhood of p diffeomorphically onto an open neighborhood of q .*

PROOF. We only treat the first case since the proof of the other two is analogous or simpler. Let κ and λ be charts of M at p resp. of N at q . Then $D_p f$ is onto precisely when $D_{\kappa(p)}(\lambda f\kappa^{-1})$ is. If $D_{\kappa(p)}(\lambda f\kappa^{-1})$ is onto, then Corollary 1.3 applied to $\lambda f\kappa^{-1}$ tells us that there is a diffeomorphism h of an open neighborhood of $\kappa(p)$ in \mathbb{R}^m onto an open subset of \mathbb{R}^m such that $\lambda f\kappa^{-1}h^{-1}$ is on its domain the projection of \mathbb{R}^m on \mathbb{R}^n . So the assertion follows if we replace κ by $h\kappa$. \square

These properties deserve a name:

DEFINITION 4.2. A map $f : M \rightarrow N$ between manifolds is called a *submersion* resp. *immersion* at $p \in M$ if $D_p f : T_p M \rightarrow T_{f(p)} N$ is onto resp. injective. If this is the case for all $p \in M$, then f is said to be a submersion resp. immersion.

COROLLARY 4.3. Let $f : M \rightarrow N$ be a map between manifolds and $q \in N$. If f is in every point of $f^{-1}(q)$ a submersion, then $f^{-1}(q)$ is a submanifold of M of dimension $m - n$.

In order to determine whether a manifold can be realized as a submanifold of \mathbb{R}^n it is convenient to have the following notion.

DEFINITION 4.4. A map $f : M \rightarrow N$ between manifolds is an *embedding* if $f(M)$ is a (smooth) submanifold of N and f is a diffeomorphism of M onto $f(M)$.

This means that for every $p \in M$ there exists a chart $\lambda = (\lambda', \lambda'') : V \rightarrow \mathbb{R}^m \times \mathbb{R}^{n-m}$ of N at $f(p)$ such that $\lambda' \circ f : f^{-1}V \rightarrow \mathbb{R}^m$ is a chart of M .

Proposition 4.1 implies that an immersion $f : M \rightarrow N$ is locally an embedding in the sense, that M is covered by open subsets on which f is an embedding.

It is clear that an embedding is an injective immersion. The converse fails to hold, witness the following two examples.

EXAMPLE 4.5. Consider the map $f : \mathbb{R} \rightarrow \mathbb{R}^2$ defined by $f(t) = (t^2 - 1, t^3 - t)$. This is an immersion. We have $f(1) = f(-1)$, but apart from that, f is injective. The image of f is a curve with a simple self-intersection in the origin and is certainly no submanifold. That image does not change if we replace f by its restriction to $\mathbb{R} - \{-1\}$. But this restriction is an injective immersion.

EXAMPLE 4.6. Let $a \in \mathbb{R} - \mathbb{Q}$ and consider the map $f : \mathbb{R} \rightarrow S^1 \times S^1$ defined by $f(t) = (e^{\sqrt{-1}t}, e^{\sqrt{-1}at})$. This is an injective immersion. It is not hard to prove that the image of f is dense in the torus $S^1 \times S^1$ and so is certainly not a submanifold.

Nevertheless we have:

PROPOSITION 4.7. An injective immersion $f : M \rightarrow N$ between smooth manifolds is an embedding if and only if it maps M homeomorphically onto its image $f(M)$.

PROOF. We prove the nontrivial direction: we suppose that f maps M homeomorphically onto $f(M)$ and prove that $f(M)$ is a submanifold of N and f a diffeomorphism of M onto $f(M)$. Let $q \in f(M)$. If $p \in M$ is such that $f(p) = q$, then according to (imm) of Proposition 4.1, f maps an open neighborhood U of p diffeomorphically onto a submanifold of N . Since f is a homeomorphism onto its image, $f(U)$ will be open in $f(M)$, i.e., be the intersection of $f(M)$ with an open $V \subset N$. So $V \cap f(M) = f(U)$ is a submanifold of N and f maps U diffeomorphically onto $V \cap f(M)$. The proposition follows. \square

If we combine the preceding proposition with the fact that a continuous bijection from a compact space to a Hausdorff space is a homeomorphism (Corollary 2.3), we find:

COROLLARY 4.8. A closed injective immersion is an embedding. In particular, an injective immersion whose domain is compact is an embedding.

EXERCISE 4.1. Let $R > r > 0$. Prove that the map $f : S^1 \times S^1 \rightarrow \mathbb{R}^3$ defined by

$$f(e^{\sqrt{-1}\alpha}, e^{\sqrt{-1}\beta}) = ((R + r \cos \beta) \cos \alpha, (R + r \cos \beta) \sin \alpha, r \sin \beta)$$

is an embedding.

EXERCISE 4.2. Find an embedding of $S^n \times S^m$ in \mathbb{R}^{n+m+1} .

EXERCISE 4.3. Consider the map $f : \mathbb{R}^{m+1} - \{0\} \rightarrow \mathbb{R}^{\frac{1}{2}(m+1)(m+2)}$ which assigns to $(x^0, \dots, x^m) \in \mathbb{R}^{m+1}$ the vector $(x^i x^j)_{i \leq j}$ (in lexicographical order, say).

(a) Prove that f is an immersion.

(b) Prove that $f(a) = f(b)$ if and only if $b = \pm a$, so that $f|_{S^m}$ factors through an injective map $g : P^m \rightarrow \mathbb{R}^{\frac{1}{2}(m+1)(m+2)}$.

(c) Prove that g is an embedding.

REMARK 4.9. The preceding problem describes (for $m = 2$) an embedding of the projective plane P^2 in \mathbb{R}^6 . We can bring this down to an embedding in \mathbb{R}^4 by modifying the definition of $f|_{S^2}$ slightly: we take as its coordinates $(w^1, \dots, w^6) = (x^2, y^2, z^2, \sqrt{2}xy, \sqrt{2}yz, \sqrt{2}xz) \in \mathbb{R}^6$, and then note that the induced embedding $g : P^2 \rightarrow \mathbb{R}^6$ takes its values in the intersection N of the unit 5-sphere $S^5 \subset \mathbb{R}^6$ and the affine hyperplane defined by $w^1 + w^2 + w^3 = 1$. This N is in fact diffeomorphic to the 4-sphere (check this). The image will not be all of N and so misses a point of N , q say. Since the complement $N - \{q\}$ is (via stereographic projection) diffeomorphic to \mathbb{R}^4 , we thus have produced an embedding of P^2 in \mathbb{R}^4 . We shall see that we cannot do better: P^2 cannot be embedded in \mathbb{R}^3 .

THEOREM 4.10. Every compact manifold M can be embedded in some \mathbb{R}^n .

PROOF. We fix a function $\phi : [0, 1] \rightarrow \mathbb{R}$ with $\phi^{-1}(1) = [0, \frac{1}{2}]$ and ϕ zero on $[\frac{3}{4}, 1]$. For every $p \in M$, we choose a chart (U_p, κ_p) at p . By replacing κ_p by $\lambda\kappa_p$ if necessary, we may assume that $\kappa_p(U_p)$ contains the unit ball of \mathbb{R}^m . We extend $\phi(\|\kappa_p\|) : U_p \rightarrow \mathbb{R}$ and $\phi(\|\kappa_p\|) \cdot \kappa_p : U_p \rightarrow \mathbb{R}^m$ as differentiable maps defined on all of M by letting them be zero (resp. the origin) on $M - U_p$. Denote these extensions by $f_{p,0}$ resp. $(f_{p,1}, \dots, f_{p,m})$, so that we have a map

$$f_p := (f_{p,0}, \dots, f_{p,m}) : M \rightarrow \mathbb{R}^{m+1}.$$

Let $D_p \subset M$ be defined by $f_{p,0} = 1$. This is also the subset of U_p defined by $\|\kappa_p\| \leq \frac{1}{2}$ and so its interior D_p° is the subset of U_p defined by $\|\kappa_p\| < \frac{1}{2}$. Observe that on D_p , $(f_{p,1}, \dots, f_{p,m})$ coincides with κ_p . So $f_p|_{D_p}$ is injective and $f_p|_{D_p^\circ}$ is an immersion. Since $\{D_p^\circ\}_{p \in M}$ is an open covering of M and M is compact, there exists a finite subcovering $\{D_{p_i}^\circ\}_{i=1}^s$. We prove that

$$f = (f_{p_1}, \dots, f_{p_s}) : M \rightarrow (\mathbb{R}^{(m+1)})^s = \mathbb{R}^{(m+1)s}$$

is an embedding. In view of Corollary 4.8 it suffices to show that f is an injective immersion. Since the $D_{p_i}^\circ$ cover M and $f_{p_i}|_{D_{p_i}^\circ}$ is an immersion, f is an immersion on all of M . Now let $p, p' \in M$ be such that $f(p') = f(p)$. We have $p \in D_{p_i}$ for some i and so $f_{p_i,0}(p) = 1$. Since $f_{p_i,0}$ is a coordinate of f , we also have $f_{p_i,0}(p') = 1$. This implies that $p' \in D_{p_i}$. But f_{p_i} is injective on D_{p_i} and hence $p' = p$. \square

REMARK 4.11. The compactness assumption can be weakened considerably to: M has a countable basis for its topology. This condition is also necessary, because \mathbb{R}^N has a countable basis (for instance, the collection of open balls whose center has rational coordinates and whose radius is rational) and hence every subspace will have that property as well. H. Whitney proved that a manifold of dimension m possessing a countable basis for its topology can in fact be embedded in \mathbb{R}^{2m} .

Tangent bundles and vector fields

1. The tangent bundle

Let M be a manifold of dimension m . Our immediate aim is to endow the (disjoint) union

$$TM := \coprod_{p \in M} T_p M$$

with the structure of a manifold of dimension $2m$ so that the base point map $TM \rightarrow M, v \in T_p M \mapsto p$, is C^∞ . At the same time we want the vector space structure in the fibers (i.e., the subtraction and the scalar multiplication maps) to vary smoothly with the base point.

We begin with the simple observation that the derivatives of a map $f : M \rightarrow N$ between manifolds combine to define a map $Df : TM \rightarrow TN$. This is a bijection if f is a diffeomorphism (by the inverse function theorem the converse also holds).

For any $p \in \mathbb{R}^m$ we may identify $T_p \mathbb{R}^m$ with \mathbb{R}^m . So if $O \subset \mathbb{R}^m$ is open, then TO gets identified with $O \times \mathbb{R}^m$ and the map which assigns to a tangent vector its base point is simply the projection $O \times \mathbb{R}^m \rightarrow O$. If f is a map from O to an open $O' \subset \mathbb{R}^n$, then

$$Df : O \times \mathbb{R}^m \rightarrow O' \times \mathbb{R}^n, (a, v) \mapsto (f(a), D_a f(v))$$

is a C^∞ map (the matrix coefficients of $D_a f$ are the first order partial derivatives of f , hence C^∞). We observe that in case f is a diffeomorphism from O onto O' (so that $m = n$), the inverse of the C^∞ -map Df is the C^∞ -map $D(f^{-1})$ and hence Df is a diffeomorphism as well.

The derivatives of a chart (U, κ) of M define a bijection $D\kappa : TU \rightarrow \kappa(U) \times \mathbb{R}^m$ and if (U', κ') is another chart of M , then $D\kappa'$ 'differs' from $D\kappa$ by

$$(D\kappa')(D\kappa)^{-1} = D(\kappa'\kappa^{-1}) : \kappa(U \cap U') \times \mathbb{R}^m \rightarrow \kappa'(U \cap U') \times \mathbb{R}^m.$$

Since $\kappa'\kappa^{-1}$ is a diffeomorphism, so is $D(\kappa'\kappa^{-1})$. This suggests that we should regard the collection of $(TU, D\kappa)$ as an atlas for a manifold structure on TM . The only problem is that this presupposes a topology on TM that makes it a topological manifold. This is solved by:

PROPOSITION 1.1. *The set TM is in just one way a differentiable manifold of dimension $2m$ with the property that the collection $\{(TU, D\kappa)\}_{(U, \kappa)}$ is a C^∞ -atlas. It has the property that the base point map $TM \rightarrow M$ is C^∞ .*

PROOF. We first establish what the topology on TM should be. For any manifold N endowed with an atlas $\{(V_\alpha, \lambda_\alpha)\}_\alpha$, we have that a subset $\Omega \subset N$ is open for its topology if and only if $\lambda_\alpha(\Omega \cap V_\alpha)$ is open for all α . This tells us that there can be only one way to put a topology on TM as desired: it must consist of the collection

of subsets $\Omega \subset TM$ with the property that $D\kappa(\Omega \cap TU)$ is open in $T\mathbb{R}^m = \mathbb{R}^{2m}$ for all charts (U, κ) of M . It is easy to check that this is a topology indeed: it is closed under arbitrary unions and finite intersections and contains TM . Notice that then for any chart (U, κ) of M , TU is open for this topology, for if (U', κ') is any other chart, then $D\kappa'(TU \cap TU') = \kappa'(U \cap U') \times \mathbb{R}^m$ is open in $\mathbb{R}^m \times \mathbb{R}^m$. We now verify that $D\kappa : TU \rightarrow U \times \mathbb{R}^m$ is a chart for this topology, i.e., is a homeomorphism. It is clearly bijective and so we need to show that

$$\Omega \subset TU \text{ open in } TM \Leftrightarrow D\kappa(\Omega) \text{ open in } U \times \mathbb{R}^m.$$

The direction \Rightarrow is trivial, let us therefore assume that $D\kappa(\Omega)$ is open in $U \times \mathbb{R}^m$. Then for any chart (U', κ') of M , $D\kappa'(\Omega \cap TU') = D(\kappa'\kappa^{-1})(D\kappa)(\Omega \cap TU')$ is open in \mathbb{R}^{2m} (we use here that $D(\kappa'\kappa^{-1})$ is a homeomorphism). So Ω is open in TM .

In order to prove that TM is a topological manifold, it remains to show that it is a Hausdorff space. Suppose $v \in T_p M$ and $v' \in T_{p'} M$ are distinct. If $p \neq p'$, then we can choose disjoint chart domains $U \ni p$, $U' \ni p'$ and then $TU \ni v$ and $TU' \ni v'$ are disjoint open subsets. If $p = p'$, we choose a chart (U, κ) at p . Then $D\kappa(v)$ and $D\kappa(v')$ are distinct in $\mathbb{R}^m \times \mathbb{R}^m$, so if we choose disjoint neighborhoods of these points, then their preimages in TU are disjoint neighborhoods of v and v' . So TM is a Hausdorff space.

We have already seen that a coordinate change $D(\kappa'\kappa^{-1})$ is a diffeomorphism. So TM becomes a smooth manifold. The proof that the projection $TM \rightarrow M$ is C^∞ is straightforward. \square

EXERCISE 1.1. Prove that if M is a submanifold of a manifold N , then TM is a submanifold of TN . (So if $M \subset \mathbb{R}^n$, then TM may be regarded as a submanifold of $M \times \mathbb{R}^n$.)

EXAMPLE 1.2. The tangent bundle TS^m of S^m is the collection pairs $(p, v) \in \mathbb{R}^{m+1} \times \mathbb{R}^{m+1}$ with $\|p\| = 1$ and $p \perp v$. This is a $2m$ -dimensional submanifold of \mathbb{R}^{2m+2} .

EXERCISE 1.2. Prove that the tangent bundle of P^m can be obtained from the preceding example by identifying in TS^m the antipodal pairs (p, v) and $(-p, -v)$.

Vector fields. In this subsection M is an m -manifold and so TM is a $2m$ -manifold.

DEFINITION 1.3. A *vector field* on M is a (by our Convention 2.8, automatically C^∞) map $V : M \rightarrow TM$, $p \mapsto V_p$ with $V_p \in T_p M$ for all $p \in M$ (in other words, V is a C^∞ section of $TM \rightarrow M$).

The vector fields on M make up a \mathbb{R} -vector space (of infinite dimension unless M is a finite set): if V, W are vector fields and $\lambda, \mu \in \mathbb{R}$, then $\lambda V + \mu W : p \mapsto \lambda V_p + \mu W_p$ is also one. We remark that scalar multiplication is not only defined for real constants, but even makes sense for functions: first notice that the \mathbb{R} -valued functions on M form a ring (under pointwise addition and multiplication). Is V a vector field and $f : M \rightarrow \mathbb{R}$ a function (so C^∞ by convention), then $f \cdot V : p \mapsto f(p)V_p$ is a vector field and this product satisfies the same properties as scalar

multiplication:

$$\begin{aligned}(f.g).V &= f.(g.V), \\ 1.V &= V, \\ (f + g).V &= f.V + g.V, \\ f.(V + W) &= f.V + f.W.\end{aligned}$$

(In the language of algebra: the vector fields on M make up a *module* over the ring of functions on M .) With this in mind, we say that an m -tuple of vector fields (V_1, \dots, V_m) on M is a *basis* of the tangent bundle if for every $p \in M$, $(V_1)_p, \dots, (V_m)_p$ is a basis for T_pM . This means that every vector field V on M can be written as $V = f^1.V_1 + \dots + f^m.V_m$ for (unique) functions f^1, \dots, f^m on M . (An algebraicist would then say that the vector fields on M form a *free module of rank m* over the ring of functions on M .) Such a basis always exists locally: if for $a \in \mathbb{R}^m$, we let $\frac{\partial}{\partial x^i}|_a \in T_a\mathbb{R}^m$ be the tangent vector represented by the curve $\gamma(t) := a + te_i$, then the vector fields

$$\frac{\partial}{\partial x^i} : a \in \mathbb{R}^m \mapsto \frac{\partial}{\partial x^i} \Big|_a, \quad i = 1, \dots, m,$$

are a basis for the vector fields on \mathbb{R}^m (every vector field V on \mathbb{R}^m can be written as $V = \sum_{i=1}^m f^i \frac{\partial}{\partial x^i}$). If (U, κ) is a chart of M , then let $\frac{\partial}{\partial \kappa^i}$ be the vector field on U that $D\kappa$ takes to $\frac{\partial}{\partial x^i}$: $\frac{\partial}{\partial \kappa^i}|_p$ is defined by the curve $\gamma(t) := \kappa^{-1}(\kappa(p) + te_i)$. The identity

$$\frac{\partial}{\partial x^i} \Big|_{\kappa(p)} = D\kappa \left(\frac{\partial}{\partial \kappa^i} \Big|_p \right)$$

shows that this is a vector field on U and that $(\frac{\partial}{\partial \kappa^1}, \dots, \frac{\partial}{\partial \kappa^m})$ is a basis of vector fields on U .

REMARK 1.4. In general such a basis need not exist on all of M . Take for instance S^2 : if a basis (V_1, V_2) of vector fields were to exist for S^2 , then the map $f : S^2 \rightarrow S^2$ which assigns to $p \in S^2$ the unit vector $(V_1)_p / \|(V_1)_p\| \in S^2$ is such that $f(p) \perp p$ for all p . In particular, f leaves no point fixed. It is also homotopic to the identity map of S^2 by means of the family $f_t(p) := \cos(\frac{1}{2}\pi t)p + \sin(\frac{1}{2}\pi t)f(p)$, $0 \leq t \leq 1$. But algebraic topology tells us that a continuous self-map of S^2 homotopic to the identity must have a fixed point.

For S^1 there is no problem: just think of S^1 as the unit circle in \mathbb{C} and take $V_z := \sqrt{-1}z$.

EXERCISE 1.3. Prove that S^3 has a basis of vector fields. (Hint: use that S^3 is the unit sphere in the skew field of quaternions.)

With the help of the so-called *octonians* (also known as *Cayley numbers*) one can prove that S^7 too has a basis of vector fields. A deep theorem due to Frank Adams (1958) says that there are no other spheres with that property.

EXERCISE 1.4. Let G be a Lie group (see Exercise 2.7) with unit e . For every $g \in G$, let $L_g : G \rightarrow G$, $h \mapsto gh$, be *left multiplication*.

(a) Prove that the map $G \times T_eG \rightarrow TG$, $(g, v) \mapsto D_eL_g(v)$ is a diffeomorphism. Conclude that G has a basis of vector fields.

(b) Prove that for every $v \in T_eG$ there exists a unique vector field V on G that is left invariant (in the sense that $L_{g*}V = V$ for all $g \in G$) with $V_e = v$.

Observe that the same holds if we replace *left* by *right* (the *right multiplication* $R_g : G \rightarrow G$ is defined by $h \mapsto hg$).

2. Derivations and the Lie bracket

We first give an algebraic characterization of a vector field. A tangent vector $v \in T_p M$ assigns to every function f on a neighborhood of p a 'directional derivative' $D_p f(v) \in \mathbb{R}$. Since we here fix v and let f vary, we prefer to denote this as $v(f)$. Let us then investigate the properties of the map $f \mapsto v(f) = D_p f(v)$. It is clearly linear in f : $v(\lambda_1 f_1 + \lambda_2 f_2) = \lambda_1 v(f_1) + \lambda_2 v(f_2)$. It also satisfies the so-called *Leibniz rule*

$$v(f_1 \cdot f_2) = f_1(p)v(f_2) + v(f_1)f_2(p),$$

which is an incarnation of the product rule: if we represent v by a curve γ at p and apply that rule to $(f_1 \cdot f_2) \circ \gamma$ we find $v(f_1 \cdot f_2) = ((f_1 \cdot f_2) \circ \gamma)'(0) = (f_1 \circ \gamma)'(0) \cdot f_2(\gamma(0)) + f_1(\gamma(0)) \cdot (f_2 \circ \gamma)'(0) = f_1(p)v(f_2) + v(f_1)f_2(p)$.

So if V is a vector field on M , then for every function $f : U \rightarrow \mathbb{R}$ on an open $U \subset M$ we get another function $V(f) : p \in U \mapsto V_p(f) = D_p f(V_p)$. For an open part of \mathbb{R}^m it is given as follows: if $V = \sum_{i=1}^m V^i \frac{\partial}{\partial x^i}$, then

$$V(f) : p \in U \mapsto \sum_{i=1}^m V^i(p) \frac{\partial f}{\partial x^i}(p).$$

This makes us to view a vector field as a simple kind of (one might say, first order) differential operator on functions. We observe the following properties:

local character: If f is a function on an open $U \subset M$ and $U' \subset U$ is open, then $V(f|_{U'}) = V(f)|_{U'}$,

\mathbb{R} -linearity: if f_1, f_2 are functions on the open $U \subset M$ and $\lambda_1, \lambda_2 \in \mathbb{R}$, then $V(\lambda_1 f_1 + \lambda_2 f_2) = \lambda_1 V(f_1) + \lambda_2 V(f_2)$ and

Leibniz rule: $V(f_1 \cdot f_2) = f_1 \cdot V(f_2) + V(f_1) \cdot f_2$.

These turn out to characterize vector fields as operators on functions:

PROPOSITION 2.1. *A map D which assigns to a real valued function f with domain an open subset of M a real valued function $D(f)$ with the same domain and obeys the three properties above is defined by a vector field on M . That vector field is unique.*

PROOF. We prove this for an open $U \subset \mathbb{R}^m$; the general case is then dealt with by means of an atlas. Consider the i th coordinate x^i of \mathbb{R}^m as a function on U and let $V^i := D(x^i|_U)$. We prove that D is defined by the vector field $V := \sum_{i=1}^m V^i \frac{\partial}{\partial x^i}$.

We first show that D is zero on the constant functions: we have $D(1) = D(1 \cdot 1) = 1 \cdot D(1) + 1 \cdot D(1) = 2D(1)$ so that $D(1) = 0$. The linearity of D then implies that $D(c) = 0$ for every constant c .

Is f a function defined on a neighborhood of $a \in U$, then we must show that $V(f)(a) = D(f)(a)$. Notice that we have

$$f(x) = f(a) + \int_0^1 \left(\frac{d}{dt} f(a + t(x - a)) \right) dt = \sum_i g_i(x) \cdot (x^i - a^i),$$

where $g_i(x) := \int_0^1 \frac{\partial f}{\partial x^i}(a + t(x - a)) dt$. Notice that $g_i(a) = \frac{\partial f}{\partial x^i}(a)$. It follows that

$$\begin{aligned} D(f) &= D\left(f(a) + \sum_i g_i \cdot (x^i - a^i)\right) = D(f(a)) + \sum_i \left(D(g_i) \cdot (x^i - a^i) + D(x^i - a^i) \cdot g_i\right) \\ &= \sum_i \left(D(g_i) \cdot (x^i - a^i) + D(x^i) \cdot g_i\right) = \sum_i \left(D(g_i) \cdot (x^i - a^i) + V^i \cdot g_i\right). \end{aligned}$$

Taking the value in a gives

$$D(f)(a) = \sum_i V^i(a) \cdot g_i(a) = \sum_i V^i(a) \frac{\partial f}{\partial x^i}(a) = V(f)(a). \quad \square$$

This characterization turns out to be extremely useful. We also remark that in commutative algebra the \mathbb{R} -linearity and the Leibniz rule make up the notion of an \mathbb{R} -derivation (so that a vector field is an \mathbb{R} -derivation of the ring of functions to itself). The following proposition is in fact a general property of derivations.

PROPOSITION 2.2. *Let V and W be vectorfields on the manifold M . Then we have a vector field $[V, W]$ on M characterized by*

$$[V, W](f) := VW(f) - WV(f).$$

The binary operation $[\ , \]$ thus defined, the Lie product, obeys:

anti-symmetry: $[V, W] = -[W, V]$ and

Jacobi identity: $[[V, W], X] = [V, [W, X]] - [W, [V, X]]$.

PROOF. We first show by means of Proposition 2.1 that $[V, W]$ is a vector field. Local character and \mathbb{R} -linearity are clear. We verify the Leibniz rule:

$$\begin{aligned} VW(f \cdot g) &= V(W(f) \cdot g + f \cdot W(g)) = \\ &= VW(f) \cdot g + W(f) \cdot V(g) + V(f) \cdot W(g) + f \cdot VW(g) \end{aligned}$$

and subtracting from this the identity gotten by interchanging V and W yields

$$(VW - WV)(f \cdot g) = (VW - WV)(f) \cdot g + f \cdot (VW - WV)(g).$$

The antisymmetry of the Lie bracket is obvious and the Jacobi identity is a formal consequence of the definition:

$$\begin{aligned} [V, [W, X]] - [W, [V, X]] &= (V(WX - XW) - (WX - XW)V) \\ &\quad - (W(VX - XV) - (VX - XV)W) = \\ &= (VW - WV)X - X(VW - WV) = [[V, W], X]. \quad \square \end{aligned}$$

EXERCISE 2.1. Prove that for vector fields V, W and a function f on a manifold:

$$[V, f \cdot W] = f \cdot [V, W] + V(f) \cdot W.$$

EXERCISE 2.2. Prove that for vector fields $V = \sum_i V^i \frac{\partial}{\partial x^i}$ and $W = \sum_i W^i \frac{\partial}{\partial x^i}$ on an open $U \subset \mathbb{R}^m$

$$[V, W] = \sum_{i,j} \left(V^j \frac{\partial W^i}{\partial x^j} \frac{\partial}{\partial x^i} - W^j \frac{\partial V^i}{\partial x^j} \frac{\partial}{\partial x^i} \right)$$

EXERCISE 2.3. Let G be a Lie group with unit e . Prove that the Lie product of two left invariant vector fields on G is left invariant. Conclude that the Lie bracket defines a product $T_e G \times T_e G \rightarrow T_e G$ (that we also denote by $[\cdot, \cdot]$) (see Exercises I-2.7 and 1.4). Compute this bracket for $G = GL(n, \mathbb{R})$ (where we identify $T_e G$ with the space of $n \times n$ -matrices).

REMARK 2.3 (Push-forward and pull-back). Given a diffeomorphism $h : M \rightarrow M'$, then objects on M (such as a function $f : M \rightarrow \mathbb{R}$ or a vector field $V : M \rightarrow TM$) can be transferred to M' to give a corresponding object on M' . This is called *push-forward* and the transferring map is often denoted by h_* . For instance, $h_* f := fh^{-1}$ and $h_* V := Dh \circ V \circ h^{-1} : M' \rightarrow TM'$.

Similarly, the *pull-back* refers to transferring objects on M' via h to objects on M and is denoted by h^* . This is in fact the push-forward of h^{-1} . The difference is however that a pull-back often makes sense when h only is a C^∞ -map (without insisting that it be a diffeomorphism). For instance, the pull-back of a function g on M' is simply $h^* g := gh$. We will later encounter more interesting examples.

Is $g : M' \rightarrow \mathbb{R}$ a function, then

$$\begin{aligned} (h_* V)(g) &= Df'(h_* V) \text{ (by definition)} \\ &= Dg \circ Dh \circ V \circ h^{-1} = D(gh) \circ V \circ h^{-1} = V(gh) \circ h^{-1} = h_*(V(h^* g)). \end{aligned}$$

The last expression can be understood as the push-forward of V , regarded as a derivation, applied to f' . So whether we push forward V as a vector field or as a derivation, the result is the same (as it should be). Likewise, if W is a vector field on M' , then its pull-back $h^* W$ is the vector field on M , which when viewed as derivation, assigns to a function f on M the function $h^*(W(h_* f))$.

3. Flows

You may think of a flow in the definition below as a fluid in motion which does not change in time (the map in the definition below then describes how a given fluid particle moves in time).

DEFINITION 3.1. A *flow* on a manifold M is a (C^∞ -) map $H : \mathbb{R} \times M \rightarrow M$ with the property that if $H_t : M \rightarrow M$ is defined by $H_t(p) = H(t, p)$, then H_0 is the identity and $H_s H_t = H_{s+t}$ for all $s, t \in \mathbb{R}$.

The fact that $H_t H_{-t} = H_0$ and $H_{-t} H_t = H_0$ are the identity, implies that H_t is a diffeomorphism of M onto M with inverse H_{-t} . (A flow is therefore simply a differentiable action of the additive group \mathbb{R} on M .) The *orbit* of a point $p \in M$ is the image of the map

$$\gamma_p : \mathbb{R} \rightarrow M, \quad \gamma_p(t) := H(t, p).$$

If this map is constant (so that the orbit of p is reduced to $\{p\}$), then we call p a *stationary point* of the flow. The orbit map γ_p defines a tangent vector $\dot{\gamma}_p(0) \in T_p M$, which in the physical situation of a fluid in motion, can be interpreted as the velocity of the particle at p . By letting p vary, we get a vector field on M . We call this vector field the *infinitesimal generator* of the flow, and denote it, appropriately, by $\frac{\partial H}{\partial t}|_{t=0}$. It is clearly zero at stationary points. We expect to know the flow of a fluid, once we know the velocity of each of its particles and therefore hope that this is also true in the abstract setting, i.e., that a flow is determined by its infinitesimal generator. We shall see that this is the case.

EXAMPLE 3.2. For $a \in \mathbb{R} - \{0\}$, a flow on \mathbb{R} is given by $H(t, x) := e^{at}x$. It has the origin as stationary point. The infinitesimal generator is the vector field $ax \frac{\partial}{\partial x}$. For $a < 0$ every point of $\mathbb{R} - \{0\}$ flows to the origin (without it ever being reached) and for $a > 0$ every point in $(0, \infty)$ (resp. $(-\infty, 0)$) flows to $+\infty$ (resp. $-\infty$).

EXAMPLE 3.3 (Example Ch. 1, 4.6 revisited). We have for every $a \in \mathbb{R}$ a flow on the torus $S^1 \times S^1$ defined by $H_t(z_1, z_2) := (e^{\sqrt{-1}t}z_1, e^{\sqrt{-1}at}z_2)$. If $a \in \mathbb{Q}$, then every orbit closes up to become an embedded circle, but if $a \notin \mathbb{Q}$, then every orbit is dense in $S^1 \times S^1$.

EXERCISE 3.1. For $a \in \mathbb{C}$, a flow on \mathbb{C} is defined by $H(t, z) := e^{at}z$. Describe the qualitative behavior of this flow for the different values of $a \in \mathbb{C}$ (distinguish in particular the cases when the real part of a is positive, zero or negative).

EXERCISE 3.2. Let W be real vector space of finite dimension and denote by $\text{End}(W)$ the vector space of all linear self-maps of W . For $A \in \text{End}(W)$ define

$$e^A := \sum_{n=0}^{\infty} \frac{A^n}{n!}.$$

We suppose as known that this series converges absolutely on every bounded part of $\text{End}(W)$ (relative to some norm on W) and that $e^{sA}e^{tA} = e^{(s+t)A}$ for $s, t \in \mathbb{R}$. Then for every $A, B \in \text{End}(W)$, $(t, X) \mapsto e^{tA}X$, $(t, X) \mapsto Xe^{tB}$ and $(t, X) \mapsto e^{tA}Xe^{tB}$ are flows on $\text{End}(W)$. Determine their infinitesimal generators.

EXERCISE 3.3. Let H be a flow on M with infinitesimal generator V . Prove that $H_{t*}(V) = V$ for all $t \in \mathbb{R}$.

Many processes in nature are mathematically modeled by a vector field on a manifold. The predictions of the model—properties that we can observe—usually concern a flow of which the vector field is an infinitesimal generator. This explains the interest of finding the flow when the vector field is given. Unfortunately, such a flow need not exist. To see this, suppose we are given a flow H on M and an open $U \subset M$ that is *not* preserved by the flow ($H_{t_0}(U) \neq U$ for some t_0). Then we clearly cannot restrict the flow to U , whereas there is no problem in restricting the infinitesimal generator to U . The latter is then a vector field on U that does not come from a flow; all we can say is that we have a map whose domain is the set of $(t, p) \in \mathbb{R} \times U$ with $H(t, p) \in U$; this is an open subset of $\mathbb{R} \times U$ containing $\{0\} \times U$. This motivates the following

DEFINITION 3.4. A *local flow* is a map $H : \Omega \rightarrow M$, where Ω is a neighborhood of $\{0\} \times M$ in $\mathbb{R} \times M$ with the property that Ω meets every line $\mathbb{R} \times \{p\}$ in an interval $\Omega_p \times \{p\}$ and H is such that $H(0, p) = p$, and $H(s, H(t, p)) = H(s + t, p)$ whenever that makes sense (i.e., when (t, p) , $(s, H(t, p))$ and $(s + t, p)$ all lie in Ω).

We should perhaps emphasize that the domain of a local flow need not contain a product neighborhood $(-\varepsilon, \varepsilon) \times M$ of $\{0\} \times M$ in $\mathbb{R} \times M$ and so H_t need not be defined for any $t \neq 0$ (for the consequences, see Exercise 3.6 below).

A local flow still defines a vector field, which we shall continue to call its infinitesimal generator. The following converse turns out to be a geometric interpretation of the existence and uniqueness of the solution of an ordinary differential equation with initial conditions.

PROPOSITION 3.5. *Every vector field V on M is the infinitesimal generator of a local flow. This flow is unique in the sense that if (Ω, H) and (Ω', H') both have V as infinitesimal generator, then H and H' coincide on $\Omega \cap \Omega'$ and yield a flow on $\Omega \cup \Omega'$ (in particular, there is local flow whose domain is maximal; it is called the evolution of V).*

PROOF. If $(U_\alpha, \kappa_\alpha)_\alpha$ is an atlas for M , then it suffices to prove the proposition for the restrictions $V|_{U_\alpha}$, for if the domain of the local flow on U_α is Ω_α , then the uniqueness part of the proposition guarantees that the two local flows defined on $\Omega_\alpha \cap \Omega_\beta$ will coincide so that we end up with a local flow whose domain is $\cup_\alpha \Omega_\alpha$. We use the charts to transfer the issue to open subsets of \mathbb{R}^m , so that now M has been replaced by an open $U \subset \mathbb{R}^m$. We then write the vector field V on U as

$$V_x = \sum_{i=1}^m V^i(x) \frac{\partial}{\partial x^i}$$

with $V^i : U \rightarrow \mathbb{R}$. The existence and uniqueness theorem for ordinary differential equations says that for every $p \in U$ there is a neighborhood U_p of p in U and an $\varepsilon > 0$ such that for every $a \in U_p$ there exists a unique solution $\gamma_a : (-\varepsilon, \varepsilon) \rightarrow U$ of the differential equation

$$(*) \quad \dot{\gamma}(t) = V_{\gamma(t)} \quad (\text{in other words, } \dot{\gamma}^i(t) = V^i(\gamma(t)) \text{ for } i = 1, \dots, n),$$

with initial condition $\gamma(0) = a$ whose dependence on a is C^∞ , in the sense that

$$\Gamma : (-\varepsilon, \varepsilon) \times U_p \rightarrow U, \quad (t, a) \mapsto \gamma_a(t),$$

is a C^∞ -map. We prove that Γ is part of a local flow that has $V|_{U_p}$ as infinitesimal generator. The initial condition requirement says that $\Gamma(0, a) = a$. We next observe that for $|s| < \varepsilon$, the curve $t \mapsto \gamma_a(s+t)$ also satisfies the differential equation (*). This solution has a different initial condition as it starts for $t = 0$ in $\gamma_a(s)$. The uniqueness property then implies that

$$\gamma_a(s+t) = \gamma_{\gamma_a(s)}(t)$$

as long as $|s+t| < \varepsilon$. This boils down to $\Gamma(s+t, a) = \Gamma(t, \Gamma(s, a))$ for $a \in U_p$ and $|s|, |t|$ and $|s+t|$ all $< \varepsilon$.

We finally verify uniqueness: suppose $\Gamma' : (-\varepsilon, \varepsilon) \times U_p \rightarrow U$ has $V|_{U_p}$ as infinitesimal generator. Let $|t| < \varepsilon$ and $a \in U_p$. It follows from $\Gamma'(s+t, a) = \Gamma'(s, \Gamma'(t, a))$ (by taking the derivative at s in $s = 0$) that $\frac{\partial \Gamma'}{\partial t}(t, a) = V_{\Gamma'(t, a)}$. If we write this out, we see that this amounts to $\gamma'_a(t) := \Gamma'(t, a)$ satisfying the differential equation (*) with the same initial condition. So $\Gamma'(t, a) = \gamma'_a(t) = \gamma_a(t) = \Gamma(t, a)$. \square

REMARK 3.6. A *time dependent vector field* on a manifold M is of course a vector field on M that depends on a real variable, hence is a map $V : \mathbb{R} \times M \rightarrow TM$ of the form $(s, p) \mapsto V_p(s) \in T_p M$. Such a field determines an ordinary vector field \tilde{V} on $\mathbb{R} \times M$, whose \mathbb{R} -component is the constant unit vector field:

$$\tilde{V}_{(s,p)} := \left(\frac{d}{ds} \Big|_s, V_p(s) \right).$$

Its evolution is of the form $(t, (s, p)) \mapsto (s+t, H(t, s, p))$.

EXERCISE 3.4. Let V be the vector field $x \frac{d}{dx}$ on the interval $(0, 1)$. Determine its evolution.

EXERCISE 3.5. Let V be a vector field on a manifold M , $p \in M$ and N a submanifold of M through p of codimension 1 ($\dim N = \dim M - 1$) with the property that $V_p \notin T_p N$. Prove that there exists a diffeomorphism of a neighborhood of $(0, p)$ in $\mathbb{R} \times N$ onto a neighborhood of p in M which takes the vector field $(\frac{\partial}{\partial t}, 0)$ to V . Hint: Prove that a local flow H for V maps a product neighborhood of $(0, p)$ in $\mathbb{R} \times N$ diffeomorphically onto a neighborhood of p in M .

EXERCISE 3.6. Let V be a vector field on a manifold M with a local flow H defined on $(-\varepsilon, \varepsilon) \times M$ for some $\varepsilon > 0$. Prove that V is the infinitesimal generator of a flow. (Hint: use that H_t is defined for $|t| < \varepsilon$ and that $H_s H_t = H_{s+t}$ wherever that makes sense.)

EXERCISE 3.7. Let V be the vector field on a manifold M with *compact support* (this means that V is zero outside a compact subset). Prove that V generates a flow. Hint: Use Exercise 3.6.

Local flows help us understand the Lie bracket of two vector fields in geometric terms. Let V be a vector field on M and let be given an open $U \subset M$ and some $\varepsilon > 0$ such that $(-\varepsilon, \varepsilon) \times U$ is in the domain of a local flow H of V . So $H_t : U \rightarrow M$ is defined for all $t \in (-\varepsilon, \varepsilon)$. If we fix our position in some $p \in U$, then after a time interval t , $0 < t < \varepsilon$, the particle that passes p was at time zero at $H_{-t}(p)$. So if $f : M \rightarrow \mathbb{R}$ is a function, then after time t its push-forward relative to H_t is the function that takes in p value $fH_{-t}(p)$ (in a physical setting this may become relevant if we have assigned to any given fluid particle a value and we wish to record how the corresponding function on M evolves with time: an observer at p will see at time t pass the particle that was at time 0 in $H_{-t}(p)$). We thus found a family of functions $H_{t*}f = fH_{-t}|_U : U \rightarrow \mathbb{R}$, $|t| < \varepsilon$. If we do the same for a vector field W on M , we find for $|t| < \varepsilon$ the push-forward field $H_{t*}W = DH_t \circ W \circ H_{-t}|_U$.

Instead of fixing our position at p , we could decide to keep track of the particle's vicissitudes as it moves along. Then it makes more sense to pull back the attributes of M at $H_t(p)$ to p and to consider H_t^*f and H_t^*W . This is in fact the stand-point that is embodied in the Lie-derivative that we will discuss later. All these expressions can be derived with respect to t .

PROPOSITION 3.7 (Infinitesimal transport). *We have on U :*

$$\begin{aligned} \frac{\partial}{\partial t} \Big|_{t=0} H_t^* f &= V(f), & \frac{\partial}{\partial t} \Big|_{t=0} H_{t*} f &= -V(f), \\ \frac{\partial}{\partial t} \Big|_{t=0} H_t^* W &= [V, W], & \frac{\partial}{\partial t} \Big|_{t=0} H_{t*} W &= -[V, W]. \end{aligned}$$

PROOF. We only verify the assertions on the left. Choose $p \in U$ and let $\gamma_p(t) := H_t(p)$. So $\dot{\gamma}_p(0) = V_p$. Now is $H_t^*f(p) = f\gamma_p(t)$; the derivative of $f\gamma_p(t)$ with respect to t in $t = 0$ is $D_p(f)(\dot{\gamma}_p(0)) = D_p(f)(V_p) = V(f)(p)$ by definition.

We check the second identity for the associated derivations. So given open subset $U \subset M$, we consistently think of both W and H_t^* as (time dependent) operators in the linear space of functions on U (which therefore can be differentiated with respect to t). Given $f : U \rightarrow \mathbb{R}$, then according to 2.3 we have

$(H_t^*W)(f) = H_t^*(W(H_{t^*}f)) = H_t^*(W(H_{-t}^*f))$ and we find

$$\begin{aligned} \frac{\partial}{\partial t} \Big|_{t=0} (H_t^*W)(f) &= \frac{\partial}{\partial t} \Big|_{t=0} H_t^*(W(H_{-t}^*f)) \\ &= \frac{\partial}{\partial t} \Big|_{t=0} W(H_{-t}^*f) + \frac{\partial}{\partial t} \Big|_{t=0} H_t^*W(f) \text{ (since } H_0 \text{ is the identity)} \\ &= W\left(\frac{\partial}{\partial t} \Big|_{t=0} H_{-t}^*f\right) + \frac{\partial}{\partial t} \Big|_{t=0} H_t^*W(f) \\ &= W(-V(f)) + V(W(f)) = [V, W](f), \end{aligned}$$

where we used the chain rule for time dependent operators (and the already proven assertion for the penultimate identity). \square

REMARK 3.8. The proof of the preceding proposition shows that properties of vector fields are sometimes best proved using their interpretation as a derivation. This is not the case for a proof that a vector field generates a local flow, but this interpretation can still offer a different perspective. For instance, if X is a vector field on the manifold M , then view it as an operator on functions f defined on M . The exponential of this operator is formally given by the series

$$\exp(tX)(f) := \sum_{n=0}^{\infty} \frac{t^n}{n!} X^n(f).$$

(So $\exp(tX)$ assigns to a function on M a series whose coefficients are functions on M .) Assuming that the righthand side converges (in a sense we do not wish to specify, but for which at least termwise differentiation can be justified), then we find that

$$\begin{aligned} \frac{d}{dt} \Big|_{t=0} \exp(tX)(f) &= \frac{d}{dt} \Big|_{t=0} \sum_{n=0}^{\infty} \frac{t^n}{n!} X^n(f) = \\ &= \sum_{n=1}^{\infty} \frac{t^{n-1}}{(n-1)!} X^n(f) = \sum_{n=0}^{\infty} \frac{t^n}{n!} X^{n+1}(f) = X(\exp(tX)(f)). \end{aligned}$$

If we now compare this with the first formula of Proposition 3.7, then we conclude that we must have $H_t^* = \exp(tX)$.

EXERCISE 3.8. Let V (resp. W) be the infinitesimal generator of the flow H (resp. I).

(a) Prove that

$$\frac{\partial^2}{\partial s \partial t} \Big|_{s=0, t=0} H_s I_t H_s^{-1} I_t^{-1} = -[V, W].$$

(Hint: use the derivation interpretation for vector fields.)

(b) Prove that if $H_s I_t = I_t H_s$ for all $s, t \in \mathbb{R}$, then $[V, W] = 0$.

(c) Now assume that $[V, W] = 0$. Prove that $H_s^* W = W$ for all s . Prove subsequently that $H_s I_t = I_t H_s$ for all s, t . (Hint: use Exercise 3.3.)

EXERCISE 3.9. The unit vector fields $\frac{\partial}{\partial x^i}$ on \mathbb{R}^m mutually commute. Prove a converse: let V_1, \dots, V_m be mutually commuting vector fields on an m -dimensional manifold M : $[V_i, V_j] = 0$ for all i, j . Let $p \in M$ be such that $(V_1)_p, \dots, (V_m)_p$ make up a basis of $T_p M$. Prove that there exists a chart (U, κ) at p such that $V_i|_U = \frac{\partial}{\partial \kappa^i}$, $i = 1, \dots, m$. Hint: Let H_i be a local flow for V_i and consider the map

$$h(x^1, \dots, x^m) = H_1(x^1, H_2(x^2, H_3(\dots H_m(x^m, p) \dots)))$$

in a neighborhood of $0 \in \mathbb{R}^m$ to a neighborhood of p in M . Prove that this map is a diffeomorphism onto a neighborhood B of 0 and show subsequently that if we take B to be a block, then $h_*\left(\frac{\partial}{\partial x^i}\Big|_B\right) = V_i|_{h(B)}$, $i = 1, \dots, m$.

EXERCISE 3.10. Let W be a finite dimensional vector space. Every $A \in \text{End}(W)$ determines vector fields \tilde{A}^l and \tilde{A}^r on $\text{End}(W)$ by $\tilde{A}_X^l := XA$ and $\tilde{A}_X^r := AX$ (the superscript r resp. l stands for 'right invariant' resp. 'left invariant'). If $A, B \in \text{End}(W)$, then determine $[\tilde{A}^r, \tilde{B}^r]$, $[\tilde{A}^l, \tilde{B}^l]$ and $[\tilde{A}^r, \tilde{B}^l]$.

Vector bundles and differential forms

1. Vector bundles

In the appendix we discuss a number of constructions (such as ‘dualizing’ and ‘tensor product’) which produce from one or more vector spaces another one. Our goal is to imitate this with the tangent bundle of a manifold in such a manner that for every tangent space this is the given construction. We thus obtain manifolds E that are endowed with a map $E \rightarrow M$ such that every fiber E_p , $p \in M$, has the structure of a vector space (e.g., $\wedge^k T_p M^*$). Such objects are conveniently described in terms of the following definition.

DEFINITION 1.1. A *vector bundle of rank r* over the manifold M is a map of manifolds $\pi : E \rightarrow M$ with the property that every fiber E_p has the structure of a real vector space of dimension r . This structure must depend C^∞ on p in the following sense: M can be covered by open subsets $U \subset M$ for which exists a map $\rho : E_U \rightarrow \mathbb{R}^r$ with the property that

- (i) the restriction of ρ to any fiber, $\rho|_{E_p} : E_p \rightarrow \mathbb{R}^r$ ($p \in U$), is an isomorphism of vector spaces and
- (ii) the map $\hat{\rho} = (\rho, \pi|_{E_U}) : E_U \rightarrow \mathbb{R}^r \times U$, (which is bijective by (i)) is a diffeomorphism.

The manifold E is called the *total space* of the vector bundle, M its *basis* and π the *projection map*.

As you will have guessed, E_p and E_U simply denote the preimages of p and U in E . We often denote a vector bundle by single (usually Greek) letter.

The definition simply says that the restriction $E_U \rightarrow U$ looks like the projection $\mathbb{R}^r \times U \rightarrow U$ in a way that is compatible with the vector space structure on the fibers. A pair (U, ρ) as in the definition is called a *local trivialization* of the vector bundle. If one exists with $U = M$, then the vector bundle is said to be *trivial* and ρ is called a *trivialization*.

A collection of local trivializations $(U_\alpha, \rho_\alpha)_\alpha$ which covers M (in the sense that $M = \cup_\alpha U_\alpha$) is sometimes called an *atlas* for the vector bundle.

EXAMPLES 1.2. (i) The trivial example is: the projection $\mathbb{R}^r \times M \rightarrow M$.

(ii) The tangent bundle τ_M of the manifold M is a vector bundle of rank $m = \dim M$. For a chart (U, κ) of M defines a local trivialization (U, ρ) of the tangent bundle: take for ρ the last m components of $D\kappa : TU \rightarrow \kappa(U) \times \mathbb{R}^m$. We have seen that such a bundle need not be trivial (the tangent bundle of S^2 is an example).

(iii) The *tautological line bundle* γ_m^1 over the real-projective space P^m : every point of P^m determines a line in \mathbb{R}^{m+1} through 0 and this line will be the fiber above this point. Precisely, the total space E is the set of $(x, \{\pm y\}) \in \mathbb{R}^{m+1} \times P^m$ with $x \in \mathbb{R}y$ and the projection map is projection on the second factor. We can

also obtain this bundle from the trivial line bundle $\mathbb{R} \times S^m \rightarrow S^m$ by taking the quotient of total space and base by the antipodal map (it sends (t, y) to $(-t, -y)$).

Other examples are obtained by means of constructions that we discuss shortly.

EXERCISE 1.1. We define the Möbius band M as the quotient space of $[0, 2\pi] \times \mathbb{R}$ obtained by identifying two vertical boundary lines via $(2\pi, t) \leftrightarrow (0, -t)$. Prove that the map $(\theta, t) \mapsto e^{i\theta} \in S^1$ induces a continuous projection $M \rightarrow S^1$. Show also that M is homeomorphic to the total space of γ_1^1 .

Given two local trivializations (U, ρ) and (U', ρ') of a rank r vector bundle $E \rightarrow M$, then we have a 'coordinate change' over $U \cap U'$:

$$\tilde{\rho}'\tilde{\rho}^{-1} : \mathbb{R}^r \times (U \cap U') \xleftarrow{\cong} E_{U \cap U'} \xrightarrow{\cong} \mathbb{R}^r \times (U \cap U').$$

It has the form $(v, p) \mapsto (h(p)(v), p)$, where $h(p)$ is a $r \times r$ matrix. This matrix is nonsingular and depends C^∞ on p . The resulting map $h : U \cap U' \rightarrow GL(r, \mathbb{R})$ is called the *transition function* of the pair (ρ, ρ') . Thus an atlas $(U_\alpha, \rho_\alpha)_\alpha$ of the vector bundle yields a collection transition functions $\{h_{\alpha, \beta} : U_\alpha \cap U_\beta \rightarrow GL(r, \mathbb{R})\}_{\alpha, \beta}$.

The following lemma, which gives a kind of converse construction, will prove quite useful. Among other things, it enables us to show that natural constructions involving vector spaces generalize to vector bundles.

LEMMA 1.3. Let $\{E_p\}_{p \in M}$ be a collection of real vector spaces of dimension r indexed by M . Suppose given a formal atlas in the sense that we have an open covering $(U_\alpha)_\alpha$ of M and for every α and every $p \in U_\alpha$ an isomorphism of vector spaces $\rho_{\alpha, p} : E_p \cong \mathbb{R}^r$. If the associated transition functions

$$h_{\alpha, \beta} : U_\alpha \cap U_\beta \rightarrow GL(r, \mathbb{R}), \quad h_{\alpha, \beta}(p) = \rho_{\beta, p} \rho_{\alpha, p}^{-1},$$

are C^∞ , then $E := \coprod_{p \in M} E_p \rightarrow M$ has precisely one structure of a rank r vector bundle for which the pairs $(U_\alpha, E_{U_\alpha} := \coprod_{p \in U_\alpha} E_p \xrightarrow{\rho_{\alpha, \cdot}} \mathbb{R}^r)$ make up an atlas.

PROOF. We proceed in much the same way as we defined a manifold structure on a tangent bundle. For every α , the map $\tilde{\rho}_\alpha := (\rho_{\alpha, p})_{p \in U_\alpha} : E_{U_\alpha} \rightarrow \mathbb{R}^r \times U_\alpha$ is a bijection. Give E_{U_α} the structure of a manifold by stipulating that this bijection be a diffeomorphism. It is clear that then for every β , $E_{U_\alpha \cap U_\beta}$ is open in E_{U_α} . The maps $\tilde{\rho}_\alpha$ and $\tilde{\rho}_\beta$ differ on $E_{U_\alpha \cap U_\beta}$ by the self-map $\tilde{h}_{\alpha, \beta}$ on $\mathbb{R}^r \times U_\alpha \cap U_\beta$ defined by $(v, p) \mapsto (h_{\alpha, \beta}(p)(v), p)$. This map is C^∞ as is its inverse $\tilde{h}_{\beta, \alpha}$. Hence it is a diffeomorphism. This shows that the two manifold structures on $E_{U_\alpha \cap U_\beta}$ coincide. A topology on E is gotten by declaring $\Omega \subset E$ open if its intersection with every E_{U_α} is open in E_{U_α} . This topology is easily checked to be Hausdorff. The lemma follows. \square

Change of basis. Given a rank r vector bundle $\xi = (E \rightarrow N)$ over a manifold N , then it is clear from the definition that if $U \subset N$ is open, then the restriction $E_U \rightarrow U$ has still the structure of a rank r vector bundle. This is true more generally if we restrict instead to a submanifold $M \subset N$: $E_M \rightarrow M$ is in a natural manner a rank r vector bundle (that we shall denote by $\xi|_M$). There is even a further generalization for which the inclusion $M \subset N$ is replaced by an arbitrary map (but of course C^∞) $f : M \rightarrow N$ of manifolds: the vector bundle $f^*\xi$ that we want to

define on M , the so-called *f-pull-back* of ξ , puts over $p \in M$ the vector space $E_{f(p)}$. In precise terms, the total space of $f^*\xi$ will be the subspace of $E \times M$ defined by

$$f^*E := \{(z, p) \in E \times M \mid f(p) = \pi(z)\}$$

and the projection to M is the projection on the second factor. We may apply Lemma 1.3 to verify that this is a vector bundle: is $(V_\alpha, \rho_\alpha)_\alpha$ an atlas for ξ with transition functions $h_{\alpha, \beta}$, then let $f^*\rho_\alpha : f^*E_{V_\alpha} \rightarrow \mathbb{R}^r$ be the composite $f^*E_{V_\alpha} \rightarrow E_{V_\alpha}$ followed by ρ_α . Then $(f^{-1}V_\alpha, f^*\rho_\alpha)_\alpha$ is an atlas for $f^*\xi$ with transition functions $f^*h_{\alpha, \beta} = h_{\alpha, \beta}f$.

Change of fiber. Here is another corollary to Lemma 1.3.

COROLLARY 1.4. *Let $\xi : E \rightarrow M$ and $\xi' : E' \rightarrow M$ be vector bundles of rank r and r' . Then the disjoint union over $p \in M$ of the vector spaces E_p^* , resp. $\wedge^k E_p$, $E_p \oplus E'_p$, $E_p \otimes E'_p$, $\text{Hom}(E_p, E'_p), \dots$ have in a natural manner the structure of a vector bundle ξ^* , resp. $\wedge^k \xi$, $\xi \oplus \xi'$, $\xi \otimes \xi'$, $\text{Hom}(\xi, \xi'), \dots$*

PROOF. We do this first for $\xi \otimes \xi'$. Choose atlases for ξ and ξ' . By intersecting their domains, we arrange that they are subordinate to the same open covering $\{U_\alpha\}_\alpha$ of M : $(U_\alpha, \rho_\alpha)_\alpha$ and $(U_\alpha, \rho'_\alpha)_\alpha$. These have transition functions $h_{\alpha, \beta} : U_\alpha \cap U_\beta \rightarrow \text{GL}(r, \mathbb{R})$ and $h'_{\alpha, \beta} : U_\alpha \cap U_\beta \rightarrow \text{GL}(r', \mathbb{R})$. We want to apply Lemma 1.3 to the collection $E_p \otimes E'_p$, $p \in M$, and the isomorphisms $\rho_{\alpha, p} \otimes \rho'_{\alpha, p} : E_p \otimes E'_p \rightarrow \mathbb{R}^r \otimes \mathbb{R}^{r'} (\cong \mathbb{R}^{rr'})$, $p \in U_\alpha$. The hypothesis of the lemma is satisfied, because the transition function for (α, β) ,

$$p \in U_\alpha \cap U_\beta \mapsto h_{\alpha, \beta}(p) \otimes h'_{\alpha, \beta}(p) \in \text{GL}(\mathbb{R}^r \otimes \mathbb{R}^{r'}) \cong \text{GL}(rr', \mathbb{R}),$$

is C^∞ . So we have a vector bundle $\xi \otimes \xi'$ over M .

The other cases are handled likewise, for all we need is that the matrix valued functions on $U_\alpha \cap U_\beta$ that assign to $p \in U_\alpha \cap U_\beta$

$$(h_{\alpha, \beta}(p))^t \in \text{GL}(r, \mathbb{R}),$$

$$\wedge^k h_{\alpha, \beta}(p) \in \text{GL}(\wedge^k \mathbb{R}^r),$$

$$h_{\alpha, \beta}(p) \oplus h'_{\alpha, \beta}(p) \in \text{GL}(\mathbb{R}^r \oplus \mathbb{R}^{r'}),$$

$$(s \in \text{Hom}(\mathbb{R}^r, \mathbb{R}^{r'})) \mapsto h'_{\alpha, \beta}(p) s h_{\alpha, \beta}^{-1}(p) \in \text{Hom}(\mathbb{R}^r, \mathbb{R}^{r'}) \in \text{GL}(\text{Hom}(\mathbb{R}^r, \mathbb{R}^{r'})),$$

...

are C^∞ . □

We generalize the notion of a vector subspace:

DEFINITION 1.5. Let $\xi = (E \rightarrow M)$ be a vectorbundle on the manifold M of rank r . A *vector subbundle* of ξ of rank $r' \leq r$ is (given by) a subset $E' \subset E$ with the property that there is an atlas (U_α, ρ_α) for ξ such that for all α , $E' \cap E_{U_\alpha} = \rho_\alpha^{-1}(\mathbb{R}^{r'} \times 0)$.

It is clear that then $E' \rightarrow M$ indeed defines a vector bundle ξ' of rank r' . Using Lemma 1.3 we easily see that the union of the quotients $E/E' := \coprod_{p \in M} E_p/E'_p$ with its projection $E/E' \rightarrow M$ has the structure of a vector bundle of rank $r - r'$ that has $(U_\alpha, \bar{\rho}_\alpha)$ as atlas. Here $\bar{\rho}_\alpha$ is the union over $p \in U_\alpha$ of the isomorphisms defined by $\rho_\alpha, E_p/E'_p \cong \mathbb{R}^r/\mathbb{R}^{r'} \times 0 \cong \mathbb{R}^{r-r'}$. We call this the *quotient bundle* of ξ by ξ' .

That ξ' is a vector subbundle of ξ is simply denoted $\xi' \subset \xi$ and the associated quotient bundle is then denoted ξ/ξ' .

EXAMPLE 1.6. Let M be a submanifold of a manifold N . Then the tangent bundle τ_M is a subbundle of the restriction $\tau_N|_M$ of τ_N to M . The quotient bundle $\nu_{N,M} := \tau_N|_M/\tau_M$ is called the *normal bundle* of M in N . Although the name is appropriate, beware that we have no obvious way of saying that a vector in $T_p N$ with $p \in M$ is 'normal' to $T_p M$ (for that would require an inner product in $T_p N$). All we can say here is that a vector in $T_p N - T_p M$ defines a nonzero vector in $T_p N/T_p M$ and this should then be thought of as a normal vector in an abstract sense.

Sections and homomorphisms. If $\xi : E \rightarrow M$ is a rank r vector bundle, then a *section* of ξ is a map $s : M \rightarrow E$ with $s(p) \in E_p$ for all $p \in M$. For instance, a section of the tangent bundle of M is simply a vector field. The sections of ξ form an \mathbb{R} -vector space and admit—just as the vector fields—scalar multiplication by functions on M .

It will be convenient to have some notation here. A *local section* of ξ is a section of ξ_U for nonempty open $U \subset M$ (so a map $s : U \rightarrow E$ with $s(p) \in E_p$ for all $p \in U$). For $U \subset M$ open we denote the collection of sections of ξ_U by $\mathcal{E}(U, \xi)$. The collection of all local sections of ξ (i.e., the union of the $\mathcal{E}(U, \xi)$ over the nonempty open $U \subset M$) will be denoted $\mathcal{E}_M(\xi)$. Analogously, we write $\mathcal{E}(U)$ for the algebra of (as ever, C^∞ -) functions $U \rightarrow \mathbb{R}$ and \mathcal{E}_M for their union.

Given vector bundles $\xi = (E \rightarrow M)$ and $\xi' = (E' \rightarrow M)$ over the same manifold M , then a (*vector bundle*) *homomorphism* $\phi : \xi \rightarrow \xi'$ is a map $\phi : E \rightarrow E'$ that maps every fiber E_p linearly to the fiber E'_p . If we have (by means of local trivializations over a common open subset $U \subset M$) $\xi|_U$ and $\xi'|_U$ identified with $\mathbb{R}^r \times U \rightarrow U$ and $\mathbb{R}^{r'} \times U \rightarrow U$, then ϕ is over U given by a map of the form $(v, p) \mapsto (s_p(v), p)$, where $s_p \in \text{Hom}(\mathbb{R}^r, \mathbb{R}^{r'})$ depends C^∞ on p . We thus see that ϕ is simply a section of the bundle $\text{Hom}(\xi, \xi')$. Clearly, a composite of homomorphisms of vector bundles over M is a homomorphism of vector bundles. We remark that for every nonempty open U , ϕ determines a map $\mathcal{E}(U, \xi) \rightarrow \mathcal{E}(U, \xi')$ (a homomorphism of $\mathcal{E}(U)$ -modules).

We have already encountered some simple examples in the case of vector subbundle: both $\xi' \subset \xi$ and the projection are $\xi \rightarrow \xi/\xi'$ homomorphisms. Here are two more examples.

EXAMPLE 1.7. Let V be a vector field on M . Inner contraction with V defines a homomorphism of vector bundles

$$\iota_V : \wedge^p \tau_M^* \rightarrow \wedge^{p-1} \tau_M^*.$$

EXAMPLE 1.8. Let $f : M \rightarrow N$ be a map between manifolds. The pull-back of τ_N , $f^* \tau_N$, is a vector bundle over M whose fiber over $p \in M$ is $T_{f(p)} N$. So the fiber of $\text{Hom}(\tau_M, f^* \tau_N)$ over p is $\text{Hom}(T_p M, T_{f(p)} N)$, which is just the receiving space of the derivative $D_p f$. Thus Df may be viewed as a homomorphism $\tau_M \rightarrow f^* \tau_N$.

A homomorphism of vector bundles is called an *isomorphism* if it is bijective on the total spaces. (Normally we would also require that its inverse be a homomorphism of vector bundles, but that is here automatic—why?) We then say that the vector bundles are *isomorphic*. For instance, a vector bundle over M of rank r is trivial precisely when it is isomorphic to $\mathbb{R}^r \times M \rightarrow M$.

EXERCISE 1.2. Prove that γ_m^1 is naturally a subbundle of the trivial bundle over \mathbb{P}^m of rank $m + 1$.

EXERCISE 1.3. Prove that a linear form on \mathbb{R}^{m+1} defines a section of the dual bundle of γ_m^1 . Show also that, conversely, any section of this dual determines a function $f : \mathbb{R}^{m+1} - \{0\} \rightarrow \mathbb{R}$ with the property that $f(\lambda x) = \lambda f(x)$ for all $x \in \mathbb{R}^{m+1} - \{0\}$ and $\lambda \in \mathbb{R} - \{0\}$.

2. Differentials and differential forms

Differentials. You will have encountered in analysis or elsewhere the notion of a differential. It is usually sort of clear how to calculate with these objects, but what they are is often left unsaid. We can now remedy this:

DEFINITION 2.1. Given a manifold M , then the *cotangent space* of M at $p \in M$ is the dual T_p^*M of the tangent space T_pM . The *cotangent bundle* of M is the dual of the tangent bundle, and denoted here by $\tau_M^* : T^*M \rightarrow M$. A *differential* on M is a section of its cotangent bundle.

Is $f : M \rightarrow \mathbb{R}$ a function, then for every $p \in M$, the derivative $D_p f : T_pM \rightarrow T_{f(p)}\mathbb{R} = \mathbb{R}$ may be regarded as an element of T_pM^* . If denote that element by df_p , then we thus get a differential $df : p \mapsto df_p$ on M .

Consider the basic case when M is an open $U \subset \mathbb{R}^m$. Then for a function $f : U \rightarrow \mathbb{R}$, the value of df_p on $\frac{\partial}{\partial x^i}|_p$ is $\frac{\partial f}{\partial x^i}(p)$. If we take $f = x^j$, then we see that dx^j takes on $\frac{\partial}{\partial x^i}$ the value δ_i^j . So (dx^1, \dots, dx^m) is (fiberwise) the basis dual to $(\frac{\partial}{\partial x^1}, \dots, \frac{\partial}{\partial x^m})$. Hence every differential ω on U is uniquely written as

$$\omega = \sum_{i=1}^m \omega_i dx^i$$

with $\omega_i : U \rightarrow \mathbb{R}$ being defined by $\omega(\frac{\partial}{\partial x^i})$. For $f : U \rightarrow \mathbb{R}$ as above we thus get $df = \sum_{i=1}^m \frac{\partial f}{\partial x^i} dx^i$. In general, if (U, κ) is a chart of the manifold M , then $d\kappa^1, \dots, d\kappa^m$ is a basis of differentials on U in the sense that every differential ω on U is uniquely written as a $\mathcal{E}(U)$ -linear combination of $d\kappa^1, \dots, d\kappa^m$:

$$\omega = \sum_{i=1}^m \omega_i d\kappa^i \text{ with } \omega_i = \omega\left(\frac{\partial}{\partial \kappa^i}\right).$$

In particular, we have for $f \in \mathcal{E}(U)$

$$df = \sum_{i=1}^m \frac{\partial f}{\partial \kappa^i} d\kappa^i,$$

where $\partial f / \partial \kappa^i$ is the value of the vector field $\partial / \partial \kappa^i$ applied to f .

Formation of the differential defines a map $d : \mathcal{E}_M \rightarrow \mathcal{E}_M(\tau_M^*)$ which obeys the following three properties:

local character: if $U' \subset U \subset M$ are open, then for any function $f : U \rightarrow \mathbb{R}$ we have $df|_{U'} = d(f|_{U'})$,

linearity: the map $f \mapsto df$ is \mathbb{R} -linear,

Leibniz rule: if $f, g : U \rightarrow \mathbb{R}$, then $d(f.g) = fdg + gdf$.

The last property is essentially the product rule. That these properties agree with those we observed earlier for a vector field is not a coincidence: if V is a vector field on M , then we have for an \mathbb{R} -valued function f on an open $U \subset M$ that $V(f)(p) = D_p f(V_p) = df_p(V_p) = \iota_{V_p}(df_p)$ for all $p \in U$. This shows that $V(f)$ only depends on df . More precisely, that V , viewed as a map from \mathcal{E}_M to itself, is the composite of d and inner contraction with ι_V :

$$V : \mathcal{E}_M \xrightarrow{d} \mathcal{E}_M(\tau_M^*) \xrightarrow{\iota_V} \mathcal{E}_M.$$

Inner contraction with V is linear over \mathcal{E}_M (not just over \mathbb{R}) and this causes the operator d to share with V the properties listed above. (The operator d is what is called in algebra the *universal derivation* of \mathcal{E}_M .)

EXERCISE 2.1. Let $U \subset \mathbb{R}^m$ be an open block and let $\omega := \sum_{i=1}^m \omega_i dx^i$ be a differential on U . Prove that ω is the differential of a function on U precisely when $\partial \omega_j / \partial x^i = \partial \omega_i / \partial x^j$ for all i, j .

EXERCISE 2.2. A differential on $\mathbb{R}^2 - \{0\}$ is defined by $\omega := (x^2 + y^2)^{-1}(x dy - y dx)$. Prove that ω can be locally written as the differential of a function, but not of a function defined on all of $\mathbb{R}^2 - \{0\}$. (Hint: compute the integral of ω over the unit circle.)

Let $F : M \rightarrow N$ be a map of manifolds. Then functions on open subsets of N can be pulled back along F to give functions on open subsets of M : if g is a \mathbb{R} -valued function on an open $V \subset N$, then $F^*g = g \circ F$ is one on the open subset $F^{-1}V \subset M$. We thus have defined $F^* : \mathcal{E}(V) \rightarrow \mathcal{E}(F^{-1}V)$ (a homomorphism of commutative \mathbb{R} -algebra's). We can do likewise for differentials: if ω is a differential on the open subset $V \subset N$, then for every $p \in F^{-1}V$, the composite of $D_p F : T_p M \rightarrow T_{F(p)} N$ with $\omega_{F(p)} : T_{F(p)} N \rightarrow \mathbb{R}$ is a linear function on $T_p M$, in other words, an element of $T_p M^*$. This is C^∞ in p and we thus find a differential on $F^{-1}V$, which we shall denote by $F^*\omega$. These two pull-backs are compatible:

LEMMA 2.2. *The operator d commutes with pull-back: $dF^* = F^*d$.*

PROOF. Let $g : V \rightarrow \mathbb{R}$ be a function on an open subset of N . If $p \in F^{-1}V$ and $v \in T_p M$, then the chain rule gives

$$(F^*dg)_p(v) = dg_{F(p)}((D_p F)(v)) = D_{F(p)}g D_p F(v) = D_p(g \circ F)(v) = d(F^*g)_p(v). \quad \square$$

Differential forms. We begin with the definition.

DEFINITION 2.3. A *differential form of degree k* , or briefly, a *k -form*, on the manifold M , $k = 0, 1, 2, \dots$, is a section of $\wedge^k \tau_M^*$. We denote the space of k -forms on an open $U \subset M$ by $\mathcal{E}^k(U)$ (so this is also $\mathcal{E}(U, \wedge^k \tau_M^*)$) and the union (over U) of the $\mathcal{E}^k(U)$'s by \mathcal{E}_M^k .

Hence a 0-form is a function on M and a 1-form is a differential on M . A k -form α on M can also be understood as a C^∞ -map which assigns to every k -tuple of tangent vectors (v_1, \dots, v_k) taken from the same tangent space $T_p M$ a real number, linear in every v_i (if restricted to a tangent space) and zero if two adjacent arguments coincide. We note that $\mathcal{E}^\bullet(U) := \bigoplus_{k=0}^m \mathcal{E}^k(U)$ is an algebra: if $\alpha, \beta \in \mathcal{E}^\bullet(U)$, then $\alpha \wedge \beta$ is defined pointwise.

If (U, κ) is a chart of M , then we observed that $d\kappa^1, \dots, d\kappa^m$ is a basis of differentials on U . In the same way the collection $d\kappa^I = d\kappa^{i_1} \wedge \dots \wedge d\kappa^{i_k}$, where

$I = (1 \leq i_1 < i_2 < \cdots < i_k \leq m)$ runs over the k -element subsets of $\{1, \dots, m\}$, is a *basis* of k -forms on U : every k -form α on U is uniquely written as

$$\alpha = \sum_{|I|=k} \alpha_I d\kappa^I \text{ with } \alpha_I \in \mathcal{E}(U).$$

Forms can be pulled back along a map $F : M \rightarrow N$ of manifolds like we did this for differentials: if $V \subset N$ is open and $\alpha \in \mathcal{E}^k(V)$, then $F^*\alpha \in \mathcal{E}^k(F^{-1}V)$ is characterized by

$$F^*\alpha(v_1, \dots, v_k) = \alpha(D_p F(v_1), \dots, D_p F(v_k)),$$

where $p \in F^{-1}V$ and $v_1, \dots, v_k \in T_p M$. A special case is that when M is a submanifold of N : a form α on N is mapped to one on M ; this is *restriction of α as a form*.

REMARK 2.4. In physical applications one may encounter the terms *covariant tensor* and *contravariant tensor*. These are sections of $(\tau_M^*)^{\otimes k}$ and $\tau_M^{\otimes k}$ respectively (for some manifold M and some k). More generally, a notion is there said to be covariant if a diffeomorphism pulls it back and contravariant if a diffeomorphism pushes it down. So a vector field would be a contravariant and a k -form a covariant notion. This clashes with what is customary in mathematics, as the conventions are there just the other way around.

3. The exterior derivative

We shall extend the d -operator as we defined on functions to one defined on forms. This extension comes up if we want to state certain laws of physics (such as the Maxwell equations) as naturally as possible. We also need this extension in Stokes' theorem, an elegant and powerful generalization of the theorem package from calculus that goes by the same name.

THEOREM 3.1. *For a manifold M there is precisely one \mathbb{R} -linear map $d : \mathcal{E}_M^k \rightarrow \mathcal{E}_M^{k+1}$, $k = 0, 1, \dots$ which for $k = 0$ is the formation of the differential and obeys:*

local character: *d preserves the domain of forms and commutes with restriction:*

if $U' \subset U \subset M$ are open, then for every $\alpha \in \mathcal{E}^\bullet(U)$ we have $d\alpha|_{U'} = d(\alpha|_{U'})$,

Leibniz rule: *if $\alpha \in \mathcal{E}^k(U)$ and $\beta \in \mathcal{E}^\bullet(U)$, then*

$$d(\alpha \wedge \beta) = (d\alpha) \wedge \beta + (-1)^k \alpha \wedge d\beta,$$

flatness: $dd = 0$.

Moreover d commutes with pull-back: if $F : M \rightarrow N$ is a map between manifolds and $\alpha \in \mathcal{E}^\bullet(N)$, then $F^(d\alpha) = d(F^*\alpha)$.*

PROOF. We first prove uniqueness. Let $\alpha \in \mathcal{E}^k(U)$, where $U \subset M$ is open. To verify that the above properties determine $d\alpha$, we may by locality assume that U is the domain of a chart κ . Then we can write $\alpha = \sum_I \alpha_I d\kappa^I$. Flatness gives $dd\kappa^i = 0$ for $i = 1, \dots, m$ and iterated use of the Leibniz rule shows that $d(d\kappa^I) = 0$ for every $I \subset \{1, \dots, m\}$ (there is something to check here because the notation $d\kappa^I$ is a bit misleading: it stands for an expression of the form $d\kappa^{i_1} \wedge \cdots \wedge d\kappa^{i_k}$ and not for the d of a form named κ^I). The Leibniz rule now implies that $d\alpha = \sum_I d(\alpha_I) \wedge d\kappa^I$.

As to existence, an obvious way to proceed is to define for a form $\alpha = \sum_I \alpha_I d\kappa^I$ on the chart domain U , $d\alpha$ by the preceding formula. This at least yields for functions on U ($k = 0$) the differential introduced before. It is also clear that

$d : \mathcal{E}^\bullet(\mathcal{U}) \rightarrow \mathcal{E}^\bullet(\mathcal{U})$ thus defined is \mathbb{R} -linear. It remains to verify the flatness and the Leibniz rule and it suffices to do this for 'monomial' forms of the type $f d\kappa^I$. Flatness is straightforward:

$$dd(f d\kappa^I) = d(df \wedge d\kappa^I) = d\left(\sum_i \frac{\partial f}{\partial \kappa^i} d\kappa^i \wedge d\kappa^I\right) = \sum_{j,i} \frac{\partial^2 f}{\partial \kappa^j \partial \kappa^i} d\kappa^j \wedge d\kappa^i \wedge d\kappa^I.$$

Since $\partial^2 f / \partial \kappa^j \partial \kappa^i$ is symmetric in i and j and $d\kappa^j \wedge d\kappa^i = -d\kappa^i \wedge d\kappa^j$ it follows that the last expression is zero.

We next verify the Leibniz rule for $f d\kappa^I$ and $g d\kappa^J$. If $I \cap J \neq \emptyset$, then the rule holds as all its terms are zero. We therefore assume that $I \cap J = \emptyset$ and let $\epsilon \in \{\pm 1\}$ be the sign determined by $d\kappa^I \wedge d\kappa^J = \epsilon d\kappa^{I \cup J}$. Then

$$\begin{aligned} d(f d\kappa^I \wedge g d\kappa^J) &= d(\epsilon f g d\kappa^{I \cup J}) = \\ &= d(\epsilon f g) \wedge d\kappa^{I \cup J} = \epsilon (g df + f dg) \wedge d\kappa^{I \cup J} = \\ &= (g df + f dg) \wedge d\kappa^I \wedge d\kappa^J = \\ &= (df \wedge d\kappa^I) \wedge g d\kappa^J + (-1)^{|I|} f d\kappa^I \wedge dg \wedge d\kappa^J = \\ &= d(f d\kappa^I) \wedge g d\kappa^J + (-1)^{|I|} f d\kappa^I \wedge d(g d\kappa^J), \end{aligned}$$

showing that the Leibniz rule holds. This combined with uniqueness guarantees that on an intersection of two chart domains the definition of d is unambiguous. This also implies the local nature of d .

Finally the commutation of d and F^* . We already know that this is so on functions. A k -form on N is locally a linear combination of forms α of the type $g dg_1 \wedge \cdots \wedge dg_k$. Since $F^*(\alpha) = F^*g F^*dg_1 \wedge \cdots \wedge F^*dg_k = F^*g dF^*g_1 \wedge \cdots \wedge dF^*g_k$, we have

$$\begin{aligned} dF^*(\alpha) &= dF^*g \wedge dF^*g_1 \wedge \cdots \wedge dF^*g_k = F^*dg \wedge F^*dg_1 \wedge \cdots \wedge F^*dg_k = \\ &= F^*(dg \wedge dg_1 \wedge \cdots \wedge dg_k) = F^*d(g dg_1 \wedge \cdots \wedge dg_k) = F^*d\alpha. \quad \square \end{aligned}$$

The operator d is called the *exterior derivative*. Let us write it out once more on an open $\mathcal{U} \subset \mathbb{R}^m$: a k -form α on \mathcal{U} has the form $\alpha = \sum_{|I|=k} \alpha_I dx^I$ and then

$$d\alpha = \sum_I \sum_{i=1}^m \frac{\partial \alpha_I}{\partial x^i} dx^i \wedge dx^I = \sum_I \sum_{i \notin I} \epsilon(i, I) \frac{\partial \alpha_I}{\partial x^i} dx^{I \cup \{i\}},$$

where $\epsilon(i, I) = (-1)^l$ if $i \notin I$ and I contains exactly l indices $< i$, and is zero if $i \in I$. We consider a few special cases. For $k = 1$ we have:

$$d\left(\sum_{j=1}^m \omega_j dx^j\right) = \sum_{1 \leq i < j \leq m} \left(\frac{\partial \omega_j}{\partial x^i} - \frac{\partial \omega_i}{\partial x^j}\right) dx^i \wedge dx^j.$$

For $m = 3$ this resembles the curl of a vector field (we might come back to this).

If $V = \sum_{i=1}^m V^i \partial / \partial x^i$ is a vector field on \mathcal{U} , then:

$$\begin{aligned} d\iota_V(dx^1 \wedge \cdots \wedge dx^m) &= d\left(\sum_{i=1}^m (-1)^{i-1} V^i dx^1 \wedge \cdots \widehat{dx^i} \cdots \wedge dx^m\right) = \\ &= \sum_{i=1}^m (-1)^{i-1} d(V^i) \wedge dx^1 \wedge \cdots \widehat{dx^i} \cdots \wedge dx^m = \left(\sum_{i=1}^m \frac{\partial V^i}{\partial x^i}\right) dx^1 \wedge \cdots \wedge dx^m. \end{aligned}$$

So the divergence of V appears here as the coefficient of the ‘volume-element’ $dx^1 \wedge \cdots \wedge dx^m$. This has an interpretation as the infinitesimal change of the volume under the local flow generated by V (see Exercise 3.2) and explains why the infinitesimal generator of the flow of an incompressible fluid (in an open subset of \mathbb{R}^m) must have zero divergence. The divergence of a vector field V on a manifold M of dimension m can only be defined relative to an m -form μ on M which is nowhere zero (and which we therefore regard as an infinitesimal volume element): the preceding suggests the definition $\operatorname{div}(V) \cdot \mu = d\iota_V(\mu)$.

EXERCISE 3.1. With the help of Remark 3.3 of the appendix we regard a k -form α on a manifold M as a function on $\wedge^k TM$. The goal is to show that $d\alpha$ has the property (which in fact characterizes it) that for every $(k+1)$ -tuple of vector fields V_0, \dots, V_k on M ,

$$\begin{aligned} d\alpha(V_0 \wedge \cdots \wedge V_k) &= \sum_{i=0}^k (-1)^i V_i(\alpha(V_0 \wedge \cdots \widehat{V}_i \cdots \wedge V_k)) \\ &\quad - \sum_{0 \leq i < j \leq k} (-1)^{i+j-1} \alpha([V_i, V_j] \wedge V_0 \wedge \cdots \wedge \widehat{V}_i \wedge \cdots \wedge \widehat{V}_j \wedge \cdots \wedge V_k). \end{aligned}$$

(We here see at work the *Koszul rule*, a heuristic principle that says that two variables of odd degree can only be interchanged at the expense of a minus sign; the V_i 's have odd degree.) We do this in stages:

- Check this on an open $U \subset \mathbb{R}^m$ for $\alpha = f dx^I$, $|I| = k < m$, and $V_i = \partial/\partial x^{i+1}$, $i = 0, \dots, k$.
- Same for α arbitrary on U and V_i of the form $\partial/\partial x^{j_i}$ for some j_i .
- Prove that the right hand side is multilinear over the functions on M : if $f : U \rightarrow \mathbb{R}$ and we replace V_i by fV_i for some i , then the whole expression gets multiplied by f .
- Prove the identity in general.
- Give another proof of this formula by exploiting the properties in Theorem 3.1 that characterize the exterior derivative.

If V a vector field, then we define the *Lie-derivative* $\mathcal{L}_V : \mathcal{E}_M^k \rightarrow \mathcal{E}_M^k$, $k = 0, 1, \dots$, by

$$\mathcal{L}_V := \iota_V d + d\iota_V.$$

Since $dd = 0$ and $\iota_V \iota_V = 0$, we have $\mathcal{L}_V = (\iota_V + d)^2$. Observe that for a function f , $\mathcal{L}_V(f) = (\iota_V d + d\iota_V)(f) = \iota_V df = V(f)$ and for an m -form α , $\mathcal{L}_V(\alpha) = d\iota_V \alpha$. If M is an open subset of \mathbb{R}^m and $\alpha = dx^1 \wedge \cdots \wedge dx^m$, then this can also be written as $\operatorname{div}(V)\alpha$.

PROPOSITION 3.2. *The Lie derivative obeys the following three properties.*

commutation with d: $d\mathcal{L}_V = \mathcal{L}_V d$,

Leibniz rule: if $\alpha \in \mathcal{E}^k(U)$ and $\beta \in \mathcal{E}^l(U)$, then

$$\mathcal{L}_V(\alpha \wedge \beta) = \mathcal{L}_V(\alpha) \wedge \beta + \alpha \wedge \mathcal{L}_V(\beta),$$

infinitesimal pull-back: if V the infinitesimal generator of a flow H on a manifold M , then for every form α on M ,

$$\left. \frac{\partial}{\partial t} \right|_{t=0} H_t^* \alpha = \mathcal{L}_V \alpha.$$

PROOF. It is clear that $d\mathcal{L}_V = dt_V d = \mathcal{L}_V d$. The Leibniz rule is a direct consequence of the Leibniz rule for the exterior derivative and a similar property of the interior product 3.5. We therefore concentrate on the last property.

For a function f we have $\mathcal{L}_V(f) = V(f)$, and so the pull-back formula holds for 0-forms because of Ch. 2., 2.3. Then the pull-back formula also holds for df :

$$\frac{\partial}{\partial t} \Big|_{t=0} H_t^* df = \frac{\partial}{\partial t} \Big|_{t=0} dH_t^* f = d \left(\frac{\partial}{\partial t} \Big|_{t=0} H_t^* f \right) = dV(f) = dt_V df = \mathcal{L}_V(df).$$

The general case now follows from the product formula: A k -form with $k \geq 1$ is locally the sum of forms of the type $\alpha = f df_1 \wedge \cdots \wedge df_k$. For such a form we have $H_t^* \alpha = H_t^*(f df_1 \wedge \cdots \wedge df_k)$. Differentiation with respect to t and then taking $t = 0$ yields:

$$\begin{aligned} & \frac{\partial}{\partial t} \Big|_{t=0} H_t^* \alpha = \\ & = \left(\frac{\partial}{\partial t} \Big|_{t=0} H_t^* f \right) df_1 \wedge \cdots \wedge df_k + \sum_i f df_1 \wedge \cdots \wedge \left(\frac{\partial}{\partial t} \Big|_{t=0} H_t^* df_i \right) \wedge \cdots \wedge df_k = \\ & = \mathcal{L}_V(f) df_1 \wedge \cdots \wedge df_k + \sum_i f df_1 \wedge \cdots \wedge \mathcal{L}_V(df_i) \wedge \cdots \wedge df_k = \\ & = \mathcal{L}_V(f df_1 \wedge \cdots \wedge df_k) = \mathcal{L}_V(\alpha), \end{aligned}$$

where the last transition used the Leibniz rule. \square

EXERCISE 3.2. Let H be a local flow on an open $U \subset \mathbb{R}^m$ with infinitesimal generator V . Prove that

$$\frac{\partial}{\partial t} \Big|_{t=0} H_t^*(dx^1 \wedge \cdots \wedge dx^m) = \operatorname{div}(V) dx^1 \wedge \cdots \wedge dx^m.$$

4. The Frobenius theorem

We will be interested in subbundles of the tangent bundle of a manifold of a particular type.

If $U \subset \mathbb{R}^m$ is open and $\pi : U \rightarrow \mathbb{R}^m \rightarrow \mathbb{R}^{m-r}$ is the projection onto the last $m-r$ coordinates (where $0 \leq r \leq m$, of course), then the kernel of $D\pi$ is a rather trivial kind of subbundle of τ_U : under the obvious identification of τU with $\mathbb{R}^m \times U$ this will correspond to $\mathbb{R}^r \times \{0\} \times U$; a section of this subbundle is a vector field of the form $\sum_{i=1}^r f^i(x) \partial / \partial x^i$. It is clear that the Lie bracket of two such vector fields is of the same form: the sections of our subbundle are closed under Lie bracket. The Frobenius theorem states a converse.

PROPOSITION 4.1 (Frobenius). *Let M be an m -manifold and $\xi \subset \tau_M$ a subbundle of rank r . Denote by $\xi^\perp \subset \tau_M^*$ the subbundle (of rank $(m-r)$) of linear forms on tangent spaces that vanish on this subbundle. Then the following are equivalent at $p \in M$:*

- (i) *there is a neighborhood $U \ni p$ and a submersion $\pi : U \rightarrow \mathbb{R}^{m-r}$ such that $\xi|_U$ equals the kernel of $D\pi$,*
- (ii) *there is a neighborhood $U \ni p$ such that $[\mathcal{E}(U, \xi), \mathcal{E}(U, \xi)] \subset \mathcal{E}(U, \xi)$,*
- (iii) *there is a neighborhood $U \ni p$ such that $d\mathcal{E}(U, \xi^\perp) \subset \mathcal{E}^1(U) \wedge \mathcal{E}(U, \xi^\perp)$,*
- (iv) *there is a neighborhood $U \ni p$ and a generating section α of $\wedge^{m-r} \xi^\perp|_U$ such that $d\alpha \in \mathcal{E}^1(U) \wedge \alpha$.*

PROOF. (i) \Rightarrow (ii). If ξ is defined at p by the submersion π in p , then choose a chart (U, κ) at p such that $\pi = (\kappa^{r+1}, \dots, \kappa^m)$. Then ξ has as a basis of sections $\partial/\partial\kappa^1, \dots, \partial/\partial\kappa^r$. These have zero Lie brackets. It is easily checked that then the Lie bracket of two $\mathcal{E}(U)$ -linear combinations of these is also of that form.

(ii) \Rightarrow (iii). Let V_1, \dots, V_r be a basis of sections of ξ on a neighborhood U of p . By assumption, $[V_i, V_j]$ is a $\mathcal{E}(U)$ -linear combination of V_1, \dots, V_r . Let $\theta \in \mathcal{E}(U, \xi^\perp)$. This means that $\theta(V_i) = 0$ for all i . According to Exercise 3.1 we have

$$d\theta(V_i, V_j) = \theta([V_i, V_j]) + V_i(\theta(V_j)) + V_j(\theta(V_i))$$

It is clear that all three terms on the right vanish, so that $d\theta(V_i, V_j) = 0$. After perhaps shrinking U we can extend V_1, \dots, V_r to a basis of sections of τ_U : V_1, \dots, V_m . If $\theta_1, \dots, \theta_m$ is the basis of τ_U^* dual to this: $\theta_i(V_j) = \delta_{ij}$ and if we write $d\theta = \sum_{i < j} f_{ij} \theta_i \wedge \theta_j$, then $f_{ij} = 0$ for $i < j \leq r$. So $d\theta$ is an $\mathcal{E}^1(U)$ -linear combination of $\theta_{r+1}, \dots, \theta_m$. Since $\theta_{r+1}, \dots, \theta_m$ is a basis of $\xi^\perp|_U$, the assertion follows.

(iii) \Rightarrow (iv). Let $\theta_1, \dots, \theta_m$ be as above so that $\theta_{r+1}, \dots, \theta_m$ is a basis of $\xi^\perp|_U$. Our assumption says that if $i > r$, then $d\theta_i = \sum_{j > r} \theta_{ij} \wedge \theta_j$ for certain $\theta_{ij} \in \mathcal{E}^1(U)$. Now $\alpha := \theta_{r+1} \wedge \dots \wedge \theta_m$ is a generating section of $\wedge^{m-r} \tau_U^*$ and it follows that

$$d\alpha = \sum_{i > r} (-1)^{i-r-1} \theta_{r+1} \wedge \dots \wedge d\theta_i \wedge \dots \wedge \theta_m = \left(\sum_{i > r} \theta_{ii} \right) \wedge \alpha.$$

(iv) \Rightarrow (i). We prove with induction on r that after perhaps shrinking U , there is a $f \in \mathcal{E}(U)$, an open $W \subset \mathbb{R}^{m-r}$, a submersion $\pi : U \rightarrow W$ and a $\beta \in \mathcal{E}^{m-r}(W)$ such that $e^f \alpha = \pi^* \beta$. This will suffice, for then π is as desired. For $r = 0$ there is nothing to show, so let us assume $r > 0$ and the implication verified for smaller values of r . Let $V := V_1$ be as above. According to Exercise Ch.2 3.5, we may shrink U to the domain of a chart κ at p such that $V = \partial/\partial\kappa^1$ and $\kappa(U) = (-\varepsilon, \varepsilon) \times U'$ for some open $U' \subset \mathbb{R}^{m-1}$. The fact that $\iota_V \alpha = 0$ implies that α does not involve $d\kappa^1$. Since $d\alpha = \theta \wedge \alpha$ for some $\theta \in \mathcal{E}^1(U)$, we also have

$$\mathcal{L}_V \alpha = \iota_V d\alpha = \iota_V(\theta \wedge \alpha) = \theta(V) \cdot \alpha.$$

Let $g \in \mathcal{E}(U)$ be a solution to the ordinary differential equation $V(g) = -\theta(V)$. Then

$$\mathcal{L}_V(e^g \alpha) = e^g V(g) \alpha + e^g \theta(V) \alpha = 0.$$

This simply means that $e^g \alpha$ is invariant under the local flow generated by V . So $e^g \alpha$ is independent of κ^1 . Since $e^g \alpha$ does not involve $d\kappa^1$ either, it follows that $e^g \alpha$ is the pull-back under $\rho := (\kappa_2, \dots, \kappa_m) : U \rightarrow U'$ of a $(m-r)$ -form α' defined on U' . This form α' satisfies our induction hypothesis, for $\rho^* d\alpha' = d(e^g \alpha) = e^g (dg + \theta) \wedge \alpha = (dg + \theta) \wedge \rho^* d\alpha'$ and since ρ faithfully preserves divisibility, this implies that $d\alpha'$ is divisible by α' (this is straightforward to check using the coordinates $\kappa_1, \dots, \kappa_m$). So by induction we may shrink U' and find $f' \in \mathcal{E}(U')$, an open $W \subset \mathbb{R}^{m-r}$, a submersion $\pi' : U' \rightarrow W$ and a $\beta^{m-r}(W)$ such that $e^{f'} \alpha' = \pi'^* \beta$. The assertion follows with $\pi := \pi' \rho$ and $f := g + \rho^* f'$. \square

Let M be a manifold. A rank r vector subbundle of the tangent bundle of M is also called a *rank r distribution*. Such a distribution is said to be *integrable* (or called a *rank r foliation*) if satisfies the equivalent properties of Proposition 4.1 at every point of M ; property (ii) is sometimes referred to as saying that $\mathcal{E}(\xi)$ is *involutive*. In that case M has as a special subatlas, consisting of the charts (U, κ) for which the

foliation on U is defined by the submersion $(\kappa^{r+1}, \dots, \kappa^m) : U \rightarrow \mathbb{R}^{m-r}$. Its coordinate changes are of the form $h = (h', h'')$ with h'' a (local) diffeomorphism of an open subset of \mathbb{R}^{m-r} onto another such subset. A rank r foliation decomposes M into subsets (called *leaves*) that look like connected r -dimensional submanifolds, but aren't in general: rather they are the images of connected r -manifolds under an injective immersion. For instance, a flow without stationary points defines a rank one foliation whose leaves are the orbits and the orbits of the irrational flow defined in Example Ch. 2, 3.3 are not submanifolds.

EXERCISE 4.1. Prove that any rank 1 distribution is a foliation.

EXERCISE 4.2. Let $V = f\partial/\partial x + g\partial/\partial y + h\partial/\partial z$ be a vector field on an open $U \subset \mathbb{R}^3$. Prove that the rank 2 distribution on U that is perpendicular to V (relative to the standard inner product on \mathbb{R}^3) is integrable if and only if V has zero curl.

5. De Rham cohomology

We say that a form α on a manifold M is *closed* resp. *exact* if $d\alpha = 0$ resp. $\alpha = d\beta$ for a form β . So for connected M , the closed 0-forms on M are the constant functions. The flatness property of d implies that an exact form is closed. 'Being closed' is a local property in the sense that it only needs to be verified on the members of an open covering of M , but not so with 'being exact'. Still, for forms of positive degree, being closed is equivalent to being locally exact:

PROPOSITION 5.1 (Poincaré Lemma). *Every closed form of positive degree on \mathbb{R}^m is exact.*

We derive this from a more general result. In order to state it, it is convenient to have at our disposal some terminology.

DEFINITION 5.2. The k -th *De Rham cohomology group* $H_{DR}^k(M)$ of a manifold M is the vector space of closed k -forms on M modulo the subspace of exact k -forms:

$$H_{DR}^k(M) = \ker(d : \mathcal{E}^k(M) \rightarrow \mathcal{E}^{k+1}(M)) / d\mathcal{E}^{k-1}(M)$$

(despite the fact that it is a vector space rather than just a group).

So the Poincaré lemma says that $H_{DR}^k(\mathbb{R}^m) = 0$ when $k \geq 1$.

It follows from 2.2 that pulling back along a map $f : M \rightarrow N$ of manifolds takes a closed resp. exact form on N to a ditto form on M . For instance, if ω is a closed k -form on N , then $df^*\omega = f^*d\omega = 0$ and so $f^*\omega$ is closed. This implies that f^* induces a linear map between De Rham cohomology groups:

$$H^k(f^*) : H_{DR}^k(N) \rightarrow H_{DR}^k(M).$$

The Poincaré lemma turns out to be a consequence of:

PROPOSITION 5.3. *Let M be a manifold. Then the projection $\pi : \mathbb{R} \times M \rightarrow M$ induces an isomorphism of De Rham cohomology groups in every degree. The inverse is induced by any embedding $i_t : p \in M \mapsto (t, p) \in \mathbb{R} \times M, t \in \mathbb{R}$.*

Before we begin the proof proper, let us make a few remarks about differential forms on $\mathbb{R} \times M$. The tangent space of $\mathbb{R} \times M$ at (t, p) decomposes as $T_t\mathbb{R} \oplus T_pM$. Hence the cotangent space of $\mathbb{R} \times M$ at (t, p) decomposes as $T_t^*\mathbb{R} \oplus T_p^*M$. Notice

that the first summand has the generator dt . The differential forms on $\mathbb{R} \times M$ decompose accordingly: a k -form α on $\mathbb{R} \times M$ is uniquely written as

$$\alpha = \alpha' + dt \wedge \alpha'',$$

where α' and α'' are forms (of degree k resp. $k-1$) in which dt does not occur. This means that α' (and likewise α'') can be considered as a form on M that depends on t , that is, as a map $t \in \mathbb{R} \mapsto \alpha'(t) \in \mathcal{E}^k(M)$. Let us write $d_M \alpha'$ for the $(k+1)$ -form that is given by $t \in \mathbb{R} \mapsto d\alpha'(t) \in \mathcal{E}^{k+1}(M)$ ('exterior derivation in the M -direction'). Observe that α is the pull-back of a k -form on M precisely when $\alpha'' = 0$ and $\alpha'(t)$ is constant in t (informally: neither dt nor t occurs in α) so that α' can be thought of as a k -form on M (whose pull-back to $\mathbb{R} \times M$ is then α).

It is worthwhile to understand the exterior derivative on $\mathcal{E}^k(\mathbb{R} \times M)$ with respect to this decomposition. A local verification in terms of a chart of M shows that

$$d\alpha = d_M \alpha' + dt \wedge \left(\frac{\partial \alpha'}{\partial t} - d_M \alpha'' \right).$$

Let now $\bar{\alpha}$ be the $(k-1)$ -form that we get by integrating α'' with respect to t :

$$\bar{\alpha}_p(t) := \int_0^t \alpha_p''(\tau) d\tau.$$

Here the right hand side is an ordinary Riemann integral of a function from \mathbb{R} to the vector space $\wedge^{\bullet} T_p^* M$. (In terms of a chart (U, κ) at p this integration takes a familiar form: if $\alpha_p''(t) = \sum_I \alpha_I''(t, p) d\kappa^I$, then we integrate every coefficient $\alpha_I'' : \mathbb{R} \times U \rightarrow \mathbb{R}$ with respect to t .) Notice, that since $\frac{\partial}{\partial t} \bar{\alpha} = \alpha''$, we have

$$d\bar{\alpha} = d_M \bar{\alpha} + dt \wedge \frac{\partial \bar{\alpha}}{\partial t} = d_M \bar{\alpha} + dt \wedge \alpha'' = d_M \bar{\alpha} + \alpha - \alpha'.$$

PROOF OF PROPOSITION 5.3. First observe, that since $\pi_{i_t} : M \rightarrow M$ is the identity, so is $H^k(i_t^* \pi^*) = H^k(i_t^*) H^k(\pi^*)$ (as a transformation in $H_{DR}^k(M)$). In particular, $H^k(\pi^*)$ is injective.

So it remains to prove that $H^k(\pi^*)$ is surjective, in other words, that every closed form $\alpha \in \mathcal{E}^k(\mathbb{R} \times M)$ is the sum of an exact form on $\mathbb{R} \times M$ and a closed form that comes from M . The preceding discussion makes this easy: we have that $\tilde{\alpha} := \alpha - d\bar{\alpha}$ is equal to $\alpha' - d_M \bar{\alpha}$ and hence is a form in which dt does not occur. We may therefore think of it as a k -form on M that depends on t . We have $d\tilde{\alpha} = d\alpha - dd\bar{\alpha} = 0$ and from the decomposition

$$0 = d\tilde{\alpha} = d_M \tilde{\alpha} + dt \wedge \frac{\partial \tilde{\alpha}}{\partial t}$$

we see that $\tilde{\alpha}$ is constant in t and closed, when regarded as a form on M . So $\tilde{\alpha} := \alpha - d\bar{\alpha}$ is the pull-back of a closed k -form on M . \square

PROOF OF THE POINCARÉ LEMMA. With induction we see that the constant map $\mathbb{R}^m \rightarrow \{0\}$ induces an isomorphism on De Rham cohomology. The assertion follows from this. \square

The De Rham cohomology groups can in fact be identified with the *singular cohomology groups* of M with real coefficients. These are vector spaces that are defined for every topological space (see any course on Algebraic Topology). We make this plausible for spaces of the homotopy type of a manifold in Exercise 5.5.

EXERCISE 5.1. Prove that for a connected manifold M , $H_{\text{DR}}^0(M) = \mathbb{R}$.

EXERCISE 5.2. Prove that $\mathbb{R}^{m+1} - \{0\}$ and S^m have the same De Rham cohomology.

EXERCISE 5.3. Prove that the closed forms on a manifold M form a subalgebra (for the exterior product). Prove also that the exact forms make up a two-sided ideal in this subalgebra. Conclude that $H_{\text{DR}}^*(M) := \bigoplus_{k=0}^{\dim M} H_{\text{DR}}^k(M)$ has the structure of a graded-commutative algebra in the sense that if $a \in H_{\text{DR}}^k(M)$ and $b \in H_{\text{DR}}^l(M)$, then $ba = (-1)^{kl}ab$. (So $aa = 0$ if a has odd degree.) Prove also that a map of manifolds $f : M \rightarrow N$ induces an algebra homomorphism $H_{\text{DR}}^*(f) : H_{\text{DR}}^*(N) \rightarrow H_{\text{DR}}^*(M)$.

EXERCISE 5.4. Prove that for every differential form α on $\mathbb{R} \times M$ we have

$$\alpha - \pi^* i_0^* \alpha = d\bar{\alpha} + \overline{d\alpha}.$$

Show that Proposition 5.3 is a formal consequence of this identity.

EXERCISE 5.5. In part (b) you may use the following approximation theorem: Let $f : M \rightarrow N$ be a continuous map between manifolds that is C^∞ on an open neighborhood of a closed subset $G \subset M$. Then there is a continuous map $F : [0, 1] \times M \rightarrow N$ such $F_0 = f$, F_1 is C^∞ and $F_t|_G = f|_G$ for all $t \in [0, 1]$ (and we call F_1 a C^∞ -representative of the homotopy class modulo G of f).

(a) We say that two maps of manifolds $f, f' : M \rightarrow N$ are C^∞ -homotopic if there exists a C^∞ -map $F : \mathbb{R} \times M \rightarrow N$ such that $F_t = f$ voor $t \leq 0$ and $F_t = f'$ voor $t \geq 1$. Prove that this is an equivalence relation and prove that C^∞ -homotopic maps induce the same map on De Rham cohomology.

(b) Let $f : M \rightarrow N$ be a continuous map between manifolds. Prove that any two C^∞ -representatives of the homotopy class of f are C^∞ -homotopic. Conclude that a homotopy class of maps $M \rightarrow N$ determines a well-defined graded-algebra homomorphism $H_{\text{DR}}^*(N) \rightarrow H_{\text{DR}}^*(M)$ and in particular, that a homotopy-equivalence of manifolds induces an isomorphism on De Rham cohomology.

6. A generalization of Stokes' theorem

Integration of volume forms. We shall interpret an m -form on an m -manifold as an integrand, that is, as an object that can be integrated. For this a bit of extra structure on the manifold in question is needed, namely an orientation. We do not want to deal with improper integrals and that is why we confine ourselves to forms with compact support. (The *support* of a form on M is the closure of the set of $p \in M$ where the form is nonzero.)

We first consider the case of an open $U \subset \mathbb{R}^m$. Let α be an m -form on U with compact support (so α zero outside a compact subset of U). If $\alpha = f dx^1 \wedge \dots \wedge dx^m$, then we *define*

$$\int_U \alpha := \int_U f(x) dx^1 dx^2 \dots dx^m,$$

where the righthand side has the usual meaning. Is h a diffeomorphism of an open $V \subset \mathbb{R}^m$ onto U , then the transformation rule for integrals says that

$$\int_U f(x) dx^1 dx^2 \dots dx^m = \int_V \left| \det \left(\left(\frac{\partial h^i}{\partial y^j} \right)_j^i \right) \right| f(h(y)) dy^1 dy^2 \dots dy^m.$$

The Jacobian of h , $\det((\frac{\partial h^i}{\partial y^j})^i_j)$, is of course nowhere zero. Let us say that h is *orientation preserving* if it is everywhere positive. This means that for every $q \in V$, the derivative $D_q h : \mathbb{R}^m \rightarrow \mathbb{R}^m$ has positive determinant, in other words, takes every basis of \mathbb{R}^m to one that defines the same orientation. If this is the case, then we can omit the absolute value sign in the transformation formula and as $h^* \alpha = \det((\frac{\partial h^i}{\partial y^j})^i_j) dy^1 \wedge \cdots \wedge dy^m$, that formula takes the elegant form

$$\int_V h^* \alpha = \int_U \alpha.$$

In other words, integration of m -forms is not sensitive to orientation preserving diffeomorphisms. This is the point of departure for transferring this notion to manifolds.

DEFINITION 6.1. An *orientation* of a manifold M is a choice of an orientation of every tangent space $T_p M$ that (when $m > 0$) is locally constant in the sense that we can cover M by charts (U, κ) with the property that the orientation at every $p \in U$ is defined by $(\partial/\partial \kappa^1|_p, \dots, \partial/\partial \kappa^m|_p)$, or equivalently, by $(d\kappa_1 \wedge \cdots \wedge d\kappa_m)_p$ (such charts are then called *oriented*). If M admits an orientation (this need not be the case), then we say that M is *orientable*.

Notice that a coordinate change $\kappa' \kappa^{-1} : \kappa(U \cap U') \rightarrow \kappa'(U \cap U')$ between two oriented charts of an oriented manifold has a Jacobian that is positive everywhere.

For $m = 0$, M is discrete, and so each of its tangent spaces is the zero vector space. An orientation then amounts to giving each point of M a sign. Here are some examples of more substance.

EXAMPLE 6.2. The m -sphere S^m is orientable: for $p \in S^m$ we orient $T_p S^m$ as follows: if (v_1, \dots, v_m) is a basis of $T_p S^m$, then (p, v_1, \dots, v_m) is a basis of \mathbb{R}^{m+1} and we declare the former to be oriented when the latter is.

EXAMPLE 6.3. The projective plane P^2 is not orientable: a point $q \in P^2$ corresponds to an antipodal pair $\{\pm p\}$ on S^2 . The projection $\pi : S^2 \rightarrow P^2$ maps $T_p S^2$ and $T_{-p} S^2$ isomorphically onto $T_q P^2$. The given orientations of $T_p S^2$ and $T_{-p} S^2$ define different orientations of $T_q P^2$ (for the antipodal map does not preserve the orientation of S^2). So if P^2 is oriented, then we would find for any $q \in P^2$ a preferred point p in the fiber $\pi^{-1} q$ that depends continuously on q (namely the point for which the tangent map to $T_q P^2$ is orientation preserving). We thus obtain a continuous section σ of π . But then we obtain a splitting of S^2 in $\sigma(P^2)$ and $-\sigma(P^2)$, which contradicts the connectedness of S^2 .

EXERCISE 6.1. Show that the Möbius band (which we can obtain from P^2 by omitting a point) is not orientable.

EXERCISE 6.2. For which m is P^m orientable?

EXERCISE 6.3. Prove that every submanifold of \mathbb{R}^{m+k} of dimension m obtained via the implicit function theorem is orientable: if $U \subset \mathbb{R}^{m+k}$ is open and $f : U \rightarrow \mathbb{R}^k$ is a submersion, then $f^{-1}(0)$ is orientable.

EXERCISE 6.4. Prove that a connected non-orientable manifold M has a degree 2 covering $\pi : \tilde{M} \rightarrow M$ with \tilde{M} connected and orientable.

In order to define integrals over manifolds we need:

LEMMA 6.4 (Partition of unity). *Let K be a compact subset of a manifold M and let $\{U_\alpha\}_\alpha$ be a collection of open subsets of M covering K . Then there exist finitely many functions $(\phi_i : M \rightarrow [0, 1])_{i=1}^n$, each of which has its support compact and contained in some U_{α_i} , such that $\sum_{i=1}^n \phi_i$ is constant 1 on a neighborhood of K . (We then say that this collection of functions is a partition of unity for K subordinate to $\{U_\alpha\}_\alpha$.)*

PROOF. For every $p \in K$ we choose an index α_p with $p \in U_{\alpha_p}$. With the help of a chart at p it is easy to construct a function $f_p : M \rightarrow [0, 1]$ whose support is compact and contained in U_{α_p} and such that f_p is constant 1 on an open neighborhood V_p of p . The collection $\{V_p\}_{p \in K}$ covers the compact K and so there exists a finite subcovering $\{(V_{p_i})_{i=1}^n\}$ of K . Then $f := \prod_{i=1}^n (1 - f_{p_i})$ is everywhere ≥ 0 with zero set $\cup_i f_{p_i}^{-1}(1)$. Notice that this last set contains $V := \cup V_{p_i}$, which is a neighborhood of K . So $g := f + \sum_i f_{p_i}$ is positive on all of M and is on V equal to $\sum_i f_{p_i}$. It follows that the collection $(\phi_i := f_{p_i}/g)_{i=1}^n$ is as desired, for $\sum_i \phi_i$ is constant 1 on V . \square

Let now M be an oriented m -manifold and α an m -form on M whose support is a compact subspace $K \subset M$. Applying the lemma to the collection domains of oriented charts yields a finite collection of functions $(\phi_i)_{i=1}^n$ with sum 1 on K and such that the support of each ϕ_i is a compact part of the domain U_i of an oriented chart κ_i . Then for every i , $(\kappa_i^{-1})^*(\phi_i \cdot \alpha)$ is an m -form on $\kappa_i(U_i)$ with compact support, and hence integrable.

PROPOSITION-DEFINITION 6.5. *The sum $\sum_{i=1}^n \int_{\kappa_i(U_i)} (\kappa_i^{-1})^*(\phi_i \cdot \alpha)$ only depends on α and the orientation of M . We call this the integral of α over the (oriented) manifold M and denote it $\int_M \alpha$. This integral depends linearly on α .*

PROOF. We must show that if $(\psi_j, V_j, \lambda_j)_j$ is another system with which we can form such a sum, then

$$\sum_i \int_{\kappa_i(U_i)} (\kappa_i^{-1})^*(\phi_i \cdot \alpha) = \sum_j \int_{\lambda_j(V_j)} (\lambda_j^{-1})^*(\psi_j \cdot \alpha).$$

The left hand side equals

$$\begin{aligned} & \sum_i \int_{\kappa_i(U_i)} (\kappa_i^{-1})^* \left(\sum_j \psi_j \cdot \phi_i \cdot \alpha \right) \text{ (for } \sum_j \psi_j = 1 \text{ where } \alpha \neq 0) \\ &= \sum_{i,j} \int_{\kappa_i(U_i)} (\kappa_i^{-1})^*(\psi_j \cdot \phi_i \cdot \alpha) \\ &= \sum_{i,j} \int_{\kappa_i(U_i \cap V_j)} (\kappa_i^{-1})^*(\psi_j \cdot \phi_i \cdot \alpha) \text{ (for } \text{supp}(\psi_j \phi_i) \text{ is a compact subset of } U_i \cap V_j). \end{aligned}$$

Since $\kappa_i \lambda_j^{-1}$ is a diffeomorphism of $\lambda_j(U_i \cap V_j)$ onto $\kappa_i(U_i \cap V_j)$ which preserves the orientation, the transformation formula implies that

$$\int_{\kappa_i(U_i \cap V_j)} (\kappa_i^{-1})^*(\psi_j \cdot \phi_i \cdot \alpha) = \int_{\lambda_j(U_i \cap V_j)} (\lambda_j^{-1})^*(\psi_j \cdot \phi_i \cdot \alpha).$$

The first part of the proposition then follows after summation over (i, j) . Linearity is obvious. \square

Stokes' theorem. This concerns a *manifold with boundary*. We will not formally introduce that notion and get around it using the following set-up: let \tilde{M} be an m -manifold with $m \geq 1$ and M an open subset of \tilde{M} with the property that its boundary ∂M is a submanifold of \tilde{M} of dimension $m - 1$. We suppose that the boundary is two-sided in the sense that \tilde{M} can be covered by charts (U, κ) such that $\partial M \cap U$ is given by $\kappa^1 = 0$ and κ^1 has constant sign on $M \cap U$. So $(\kappa^2, \dots, \kappa^m)|_{\partial M \cap U}$ is then a chart for ∂M . Suppose now \tilde{M} oriented. We orient the boundary ∂M as follows: if $p \in \partial M$, then $T_p(\partial M)$ is a hyperplane in $T_p\tilde{M}$ that is in fact the boundary of the half space of the tangent vectors that point towards M . According to the appendix (A.3) this determines an orientation of $T_p(\partial M)$. If we take the preceding κ oriented with respect to the given orientation of M , then $(\kappa^2, \dots, \kappa^m)|_{\partial M \cap U}$ is an oriented chart for ∂M precisely when $M \cap U$ is given by $\kappa^1 < 0$. (This means that if $\partial M \cap U$ is given by $\kappa^1 > 0$, then we get the opposite orientation.) Such a choice for κ is always possible unless $m = 1$: if $M \cap U$ is given by $\kappa^1 > 0$ and $d\kappa^1$ defines the orientation of M , then $\partial M \cap U$ is a singleton whose orientation is given by a minus sign.

EXAMPLES 6.6. (i) Take $\tilde{M} = \mathbb{R}$ and take for M a finite open interval (a, b) so that $\partial M = \{a, b\}$. If we orient \mathbb{R} in the obvious way (so by dx), then b gets oriented with a plus sign and a with a minus sign.

(ii) Take $M = \mathbb{R}^m$, $m \geq 1$, with the orientation defined by (dx^1, \dots, dx^m) and take for M the open unit ball B^m defined by $\|x\| \leq 1$ so that ∂M is the unit $(m-1)$ -sphere S^{m-1} . The orientation for S^{m-1} that results from the above construction is called the standard orientation. For $m = 2$ this yields the orientation of the unit circle by the counterclockwise direction (given by an angular coordinate). For $m = 3$, the orientation of S^2 (when regarded as the surface on the earth) at the tangent plane of our location is given by the direction straight ahead, followed by the direction left).

(iii) Any 3-manifold M inhabited by a human being is naturally oriented: he orients the tangent space at his point of location by the basis given by (1) the direction straight ahead, (2) the direction to his left and (3) the direction above his head. Unless of course this person can make a hike through his universe M and friends welcoming him back find him mirrored (in the sense that left and right have interchanged) upon return—left and right then no longer make sense.

THEOREM 6.7 (Stokes' theorem). *In this situation we have for every $(m-1)$ -form β on \tilde{M} with compact support*

$$\int_{\tilde{M}} d\beta = \int_{\partial M} \iota^* \beta,$$

where $\iota : \partial M \subset \tilde{M}$ is the inclusion.

The following simple special case is also the crucial case:

LEMMA 6.8. *Let β be a $(m-1)$ -form on \mathbb{R}^m with compact support. Then $\int_{\mathbb{R}^m} d\beta = 0$. If \mathbb{R}_{\pm}^m is the part of \mathbb{R}^m defined by $\pm x^1 \geq 0$ and $\iota_{\pm} : \mathbb{R}^{m-1} \rightarrow \mathbb{R}^m$ is the map $(y^1, \dots, y^{m-1}) \mapsto (0, y^1, \dots, y^{m-1})$, then $\int_{\mathbb{R}_{\pm}^m} d\beta = \mp \int_{\mathbb{R}^{m-1}} \iota_{\pm}^* \beta$.*

PROOF. We write β as $\sum_i \beta_i dx^1 \wedge \cdots \wedge \widehat{dx^i} \wedge \cdots \wedge dx^m$. Then

$$d\beta := \sum_i (-1)^{i-1} \frac{\partial \beta_i}{\partial x^i} dx^1 \wedge \cdots \wedge dx^m$$

and so for every block $B \subset \mathbb{R}^m$ we have

$$\int_B d\beta = \sum_i (-1)^{i-1} \int_B \frac{\partial \beta_i}{\partial x^i} dx^1 \cdots dx^m.$$

We compute the integral of $\partial \beta_i / \partial x^i$ over \mathbb{R}^m resp. \mathbb{R}_\pm^m by first integrating over the i th variable. The latter always yields zero (since $\beta_i(x) = 0$ when $|x^i|$ is sufficiently large), unless we integrate over \mathbb{R}_\pm^m and $i = 1$: then we get $\mp \beta_1(0, x^2, \dots, x^m)$. Integration of the latter with respect to the variables x^2, \dots, x^m gives $\mp \int_{\mathbb{R}^{m-1}} \iota_1^* \beta$ by definition. \square

PROOF OF STOKES' THEOREM. Let K be the support of β . Choose a partition of unity for K , $(\phi_i)_{i=1}^n$, such that support of ϕ_i is a compact subset of the domain U_i of an oriented chart $\kappa_i = (\kappa_i^1, \dots, \kappa_i^m)$, such that $M \cap U_i$ is defined by $\kappa_i^1 \leq 0$ or $\kappa_i^1 \geq 0$. Notice that the restriction of $(\kappa_i^2, \dots, \kappa_i^m)$ to $\partial M \cap U_i$ is in the first case orientation preserving and in the second case orientation reversing.

Since $\sum_i \phi_i$ equals 1 where $\beta \neq 0$, we have $\beta = \sum_i \phi_i \beta$. Integration is \mathbb{R} -linear, so we only need to check the theorem for $\phi_i \beta$. In other words, we may assume that the support of β is a compact subset of the domain of an oriented chart (U, κ) with $U \cap \overline{M}$ defined by $\mp \kappa^1 \geq 0$. Then

$$\begin{aligned} \int_{\overline{M}} d\beta &= \int_{\kappa(U \cap \overline{M})} (\kappa^{-1})^*(d\beta) = \int_{\kappa(U) \cap (\mp x^1 \geq 0)} d(\kappa^{-1})^*(\beta) \\ &= \pm \int_{\kappa(U) \cap (x^1 = 0)} \iota_1^*(\kappa^{-1})^*(\beta) \text{ (by Lemma 6.8)} \\ &= \int_{U \cap \partial M} \iota^* \beta \text{ (because of the way } \partial M \text{ has been oriented)} \\ &= \int_{\partial M} \iota^* \beta. \end{aligned}$$

\square

REMARK 6.9. Let M be a compact, nonempty oriented manifold. Then \int_M is a linear functional on the m -forms on M . This functional is evidently nonzero. Stokes' theorem tells us that it is zero on $d\mathcal{E}^{m-1}(M)$ and so \int_M factors over a nonzero linear functional

$$\int_M : H_{\mathbb{D}\mathbb{R}}^m(M) \rightarrow \mathbb{R}$$

REMARK 6.10. Let M be an arbitrary m -manifold and let $n \in \{0, \dots, m\}$. If $f : N \rightarrow M$ is a map from a compact oriented n -manifold N , then we have a linear functional

$$H_{\mathbb{D}\mathbb{R}}^n(M, \mathbb{R}) \xrightarrow{H^n(f^*)} H_{\mathbb{D}\mathbb{R}}^n(N, \mathbb{R}) \xrightarrow{\int_N} \mathbb{R}.$$

A theorem due to R. Thom (1954) implies that these functionals span the dual of $H_{\mathbb{D}\mathbb{R}}^n(M, \mathbb{R})$. In other words, if a closed n -form α on M fails to be exact, then there exist N and f as above such that $\int_N f^* \alpha \neq 0$.

REMARK 6.11. In the preceding an 'integrand' on a m -manifold M required us to give for every tangent space $T_p M$ an orientation and an anti-symmetric multilinear form $\alpha_p : (T_p M)^m \rightarrow \mathbb{R}$. We get the same result if we change the orientation and at the same time replace α_p by $-\alpha_p$. We may merge these two ingredients in what is perhaps a simpler notion: a *Haar measure* on a real m -dimensional vector space V is a function $\mu : V^m \rightarrow \mathbb{R}$ with the property that for every linear map $s : V \rightarrow V$ we have:

$$\mu(s(v_1), \dots, s(v_m)) = |\det(s)|\mu(v_1, \dots, v_m).$$

It is clear that μ is determined by its value on a single basis (e_1, \dots, e_m) , for if $(v_1, \dots, v_m) \in V^m$, then $\mu(v_1, \dots, v_m) = |\det(s)|\mu(e_1, \dots, e_m)$, where $s : V \rightarrow V$ is defined by $s(e_i) = v_i$. In particular is $\mu(v_1, \dots, v_m) = 0$ when (v_1, \dots, v_m) is linearly dependent. So if $\alpha : V^m \rightarrow \mathbb{R}$ is the anti-symmetric multilinear form that takes the same value as μ on (e_1, \dots, e_m) , then α and μ agree or are opposite on every basis according to whether or not that basis defines the same orientation as (e_1, \dots, e_m) . This is just the ambiguity that disappears in our interpretation as an integrand. We say that μ is *positive* if μ takes positive values on all bases of V (there is no such notion for m -forms on V unless V has been endowed with an orientation).

The Haar measures on V form a one dimensional vector space. Likewise, the Haar measures of the tangent spaces of M define a vector bundle of rank one over M (which is in fact trivial if M has a countable basis). An *integrand* (or *C^∞ -measure*) on M is a section μ of this line bundle. If it has compact support, then it is easy to give $\int_M \mu$ a sense, but while this notion of integral is perhaps more direct, there is no reasonable formulation of Stokes' theorem in these terms.

Topics in linear algebra

We shall usually need the results proved in this appendix only for real vector spaces, although it is sometimes convenient to have them also available in the complex case. We shall find however that much of the discussion does not even require our base field to be \mathbb{R} or \mathbb{C} and hence we fix a field K and assume all vector spaces to be K -vector spaces and all linear maps between them to be K -linear. Only at certain loci we make restrictive assumptions regarding K .

1. Dualization

Let V be a vector space. The collection of linear functions $l : V \rightarrow K$ is a vector space as well (for subtraction of two functions and multiplication of a function by a scalar), called the *dual* of V and denoted V^* . A linear map $s : V \rightarrow W$ of vector spaces determines a map $s^* : W^* \rightarrow V^*$ by composition: $s^*(l) := ls$. This map is linear (clear) and is called the *dual* of s .

EXAMPLES 1.1. We give a couple of infinite dimensional examples borrowed from functional analysis. These examples arise in the mathematical literature (and so are not farfetched) and illustrate a point we wish to make, namely that a vector space and its dual are different things.

(i) Take for V the \mathbb{R} -vector space of all (so not necessarily continuous) functions $v : \mathbb{R} \rightarrow \mathbb{R}$. Then for every $a \in \mathbb{R}$, evaluation in a , $\delta_a : v \in V \mapsto v(a)$ is a linear function on V .

(ii) Let V be the \mathbb{R} -vector space of continuous functions $v : [a, b] \rightarrow \mathbb{R}$. Given a continuous function $u : [a, b] \rightarrow \mathbb{R}$, then $v \in V \mapsto \int_a^b v(t)u(t) dt$ is a linear function on V .

(iii) Let $I \subset \mathbb{R}$ be open and nonempty and V be the \mathbb{R} -vector space of C^k -functions $v : I \rightarrow \mathbb{R}$. Then for $a_0, \dots, a_k \in \mathbb{R}$ and $p \in I$, the map $v \in V \mapsto \sum_{i=0}^k a_i v^{(i)}(p)$ is a linear function on V .

LEMMA 1.2. *Suppose that V is of finite dimension m and let (e_1, \dots, e_m) be a basis of V . Let for $i = 1, \dots, m$, $x^i : V \rightarrow \mathbb{R}$ be the linear function for which $x^i(e_j) = \delta_j^i$. Then (x^1, \dots, x^m) is a basis of V^* (called the basis dual to the given one on V).*

Here δ_j^i is the *Kronecker delta*; it takes the value 1 if the two indices coincide and is 0 otherwise.

PROOF OF LEMMA 1.2. They are linearly independent: suppose that $\sum_i \lambda_i x^i = 0$ for certain scalars $\lambda_1, \dots, \lambda_m$. The value of $\sum_i \lambda_i x^i$ on e_j is λ_j , hence $\lambda_j = 0$ for all j .

They generate V^* : for every linear $l : V \rightarrow \mathbb{R}$, $\sum_i l(e_i)x^i$ takes the same value as l on the basis e_1, \dots, e_m and so these two must be equal: $l = \sum_i l(e_i)x^i$. \square

The *transpose* of a matrix is best (and intrinsically) understood as the matrix of the dual of a linear transformation: Let $s : V \rightarrow W$ be a linear map between finite dimensional vector spaces. Suppose given bases (e_1, \dots, e_m) of V resp. (f_1, \dots, f_n) of W . Then the matrix $\text{mat}(s) := (s_i^j)$ of s relative to these bases is defined by $s(e_i) = \sum_{j=1}^n s_i^j f_j$. (The column index is always one of a basisvector of the domain and the row index one of the range, so s_i^j is in the i th column and the j th row.) Denote by (x^1, \dots, x^m) resp. (y^1, \dots, y^n) the bases of V^* resp. W^* dual to given bases. Then for $i = 1, \dots, n$, $s^*(y^i) = y^i s$ is the linear function on V which takes in e_j the value $y^i s(e_j) = s_j^i$. In other words, $s^*(y^i) = \sum_{j=1}^m s_j^i x^j$. So $\text{mat}(s^*) = \text{mat}(s)^t$.

If we pick a vector $v \in V$, then *evaluation in v* , given by $l \in V^* \mapsto l(v)$, is a linear function on V^* . We shall denote it by ι_v :

$$\iota_v(l) := l(v).$$

It is clear that ι_v is linear in v : $\iota_{v+v'}(l) = l(v+v') = l(v) + l(v') = \iota_v(l) + \iota_{v'}(l)$ and $\iota_{\lambda v}(l) = l(\lambda v) = \lambda l(v) = \lambda \iota_v(l)$. So we have a linear map

$$\iota : V \rightarrow V^{**}, \quad v \mapsto \iota_v$$

LEMMA 1.3. *If V is finite dimensional, then the map $\iota : V \rightarrow V^{**}$ is an isomorphism of vector spaces. In fact, it maps any basis of V onto its double-dual (which is indeed a basis of V^{**}).*

PROOF. Let (e_1, \dots, e_m) be a basis of V . We have $\iota_{e_j}(x^i) = x^i(e_j) = \delta_j^i$, which shows that $(\iota_{e_1}, \dots, \iota_{e_m})$ is the basis of V^{**} dual to x^1, \dots, x^m . \square

This lemma is reflected by the property that transposing a matrix twice leaves it unchanged.

EXERCISE 1.1. Let $s : V \rightarrow W$ and $u : W \rightarrow X$ be linear maps of vector spaces. Prove that $(us)^* = s^*u^*$. Connect this with the fact that if A and B are matrices whose sizes are such that the matrix product $A \cdot B$ is defined, then $(A \cdot B)^t = B^t \cdot A^t$.

EXERCISE 1.2. Let V and W be finite dimensional vector spaces and $B : V \times W \rightarrow K$ a bilinear map. This means that for every $v \in V$, the function $w \in W \mapsto B(v, w)$ is linear and that the map $b' : V \rightarrow W^*$ thus obtained is linear as well. Similarly we find a linear map $b'' : W \rightarrow V^*$.

(a) Prove that b' and b'' are each other's dual.

(b) Consider the special case that $W = V^*$ and $B(v, l) := l(v)$ (this is called the *natural pairing* between V and V^*). What are b' and b'' in this case?

(c) Prove that b' is an isomorphism precisely when b'' is an isomorphism. (We then call B a *perfect pairing* between V and W .) Prove also that for every linear transformation $s : V \rightarrow V$ there is precisely one linear transformation $s^t : W \rightarrow W$ such that $B(s(v), w) = B(v, s^t(w))$ for all $v \in V$ and $w \in W$. (We call s^t the *adjoint* of s relative to B .)

2. The tensor product

Let V, W, X be vector spaces. We recall that a map $B : V \times W \rightarrow X$ is said to be *bilinear* if it is linear in every argument, so

$$\begin{aligned} B(\lambda v, w) &= \lambda B(v, w), & B(v, \lambda w) &= \lambda B(v, w) \\ B(v + v', w) &= B(v, w) + B(v', w), & B(v, w + w') &= B(v, w) + B(v, w'). \end{aligned}$$

for all $v, v' \in V, w, w' \in W$ and $\lambda \in K$. This motivates the following construction.

Consider the K -vector space $K^{(V \times W)}$ spanned by the elements of the set $V \times W$. This vector space has by definition a basis indexed by the set $V \times W$: $(e_{v,w})_{v \in V, w \in W}$ so that an arbitrary element of $K^{(V \times W)}$ is a *finite* sum $\sum_{i=1}^k \lambda_i e_{v_i, w_i}$ with $\lambda_i \in K, v_i \in V$ and $w_i \in W$. A bilinear map B as above determines a linear map $\tilde{b} : K^{(V \times W)} \rightarrow X$ by $\sum_i \lambda_i e_{v_i, w_i} \mapsto \sum_i \lambda_i B(v_i, w_i)$. This map has always in its kernel the elements

$$\begin{aligned} e_{\lambda v, w} - \lambda e_{v, w}, & \quad e_{v, \lambda w} - \lambda e_{v, w}, \\ e_{v+v', w} - e_{v, w} - e_{v', w}, & \quad e_{v, w+w'} - e_{v, w} - e_{v, w'}, \end{aligned}$$

where $\lambda \in K, v, v' \in V$ and $w, w' \in W$. So if $H \subset K^{(V \times W)}$ is the subspace spanned by them, then $\tilde{b} : K^{(V \times W)} \rightarrow X$ will factor through a linear map $b : K^{(V \times W)}/H \rightarrow X$. We denote the quotient $K^{(V \times W)}/H$ by $V \otimes W$ and call it the *tensor product* of V and W ; the image of $e_{v,w}$ in $V \otimes W$ is written $v \otimes w$. So we have

$$\begin{aligned} (\lambda v) \otimes w &= \lambda(v \otimes w) = v \otimes (\lambda w), \\ (v + v') \otimes w &= v \otimes w + v' \otimes w, & v \otimes (w + w') &= v \otimes w + v \otimes w' \end{aligned}$$

where $\lambda \in K, v, v' \in V$ and $w, w' \in W$. In other words, the map

$$\otimes : V \times W \rightarrow V \otimes W, \quad (v, w) \mapsto v \otimes w$$

is bilinear. This bilinear map is *universally* so in the sense that every bilinear map $B : V \times W \rightarrow X$ is the composite of the bilinear $\otimes : V \times W \rightarrow V \otimes W$ and a (unique) linear map $V \otimes W \rightarrow X$.

LEMMA 2.1. *Suppose V and W are finite dimensional vector spaces that come with bases (e_1, \dots, e_m) resp. (f_1, \dots, f_n) . Then $(e_i \otimes f_j)_{i,j}$ is a basis of $V \otimes W$, in particular, $V \otimes W$ has (finite) dimension mn .*

PROOF. The elements $e_i \otimes f_j$ generate $V \otimes W$: given $v \in V$ and $w \in W$, then write them out in the bases: $v = \sum_i v^i e_i, w = \sum_j w^j f_j$ and observe that $v \otimes w = (\sum_i v^i e_i) \otimes (\sum_j w^j f_j) = \sum_{i,j} v^i w^j (e_i \otimes f_j)$.

The elements $e_i \otimes f_j$ are linearly independent: if (x^1, \dots, x^m) resp. (y^1, \dots, y^n) are their dual bases of V^* resp. W^* , then for a given pair (i, j) , the function $B^{ij} : V \times W \rightarrow K, (v, w) \mapsto x^i(v)y^j(w)$ is bilinear. For the associated $b^{ij} : V \otimes W \rightarrow K$ we have $b^{ij}(e_{i'} \otimes f_{j'}) = \delta_{i'}^i \delta_{j'}^j$, i.e., 1 if $(i', j') = (i, j)$ and 0 otherwise. This proves that $(e_i \otimes f_j)_{i,j}$ is linearly independent and hence a basis. (This argument also shows that $(b^{ij})_{i,j}$ is the basis dual to it.) \square

We denote the set of linear maps $V \rightarrow W$ by $\text{Hom}(V, W)$. This is a vector space. A pair $(l, w) \in V^* \times W$ defines a simple kind of linear map $V \rightarrow W, v \mapsto l(v)w$.

We thus obtain a map $V^* \times W \rightarrow \text{Hom}(V, W)$, which is easily verified to be bilinear (check!) and hence gives rise to a linear map

$$\varepsilon : V^* \otimes W \rightarrow \text{Hom}(V, W).$$

LEMMA 2.2. *If V and W are finite dimensional, then ε is an isomorphism of vector spaces.*

PROOF. Pick bases as before. Then $\varepsilon(x^i \otimes f_j)$ is the linear map that takes e_i to f_j and every other basis vector $e_{i'}$ to 0. So its matrix has a single nonzero entry (in the i th column and the j th row) and that entry has value 1. So the set of $\varepsilon(x^i \otimes f_j)$ makes up a basis of $\text{Hom}(V, W)$. \square

EXERCISE 2.1. (a) Prove that if V is finite dimensional with basis e_1, \dots, e_m , then any element of $V \otimes W$ is uniquely written $\sum_{i=1}^m e_i \otimes w_i$ with $w_i \in W$ (so that $V \otimes W$ gets identified with W^m).

(b) Prove that $\varepsilon : V^* \otimes W \rightarrow \text{Hom}(V, W)$ is still an isomorphism if only one of V and W is finite dimensional.

(c) Prove that in general $\varepsilon : V^* \otimes W \rightarrow \text{Hom}(V, W)$ is injective and that its image consists of the $\phi \in \text{Hom}(V, W)$ with $\dim(\phi(W)) < \infty$.

EXERCISE 2.2. (a) Let $s : V \rightarrow W$ and $s' : V' \rightarrow W'$ be linear maps. Prove that $v \otimes v' \mapsto s(v) \otimes s'(v')$ defines a linear map $V \otimes V' \rightarrow W \otimes W'$ (which we usually denote by $s \otimes s'$).

(b) Prove that this construction defines a linear map

$$\text{Hom}(V, W) \otimes \text{Hom}(V', W') \rightarrow \text{Hom}(V \otimes V', W \otimes W')$$

and show that this is an isomorphism if all four vector spaces are finite dimensional.

(c) Suppose all four vector spaces come with a (finite) basis and endow both $V \otimes V'$ and $W \otimes W'$ with the associated bases. If (s_i^j) is the matrix of s and (u_i^k) the one of u , then determine the matrix coefficients of $s \otimes s'$.

EXERCISE 2.3. Let V be a finite dimensional vector space. The map $(l, v) \in V^* \times V \mapsto l(v) \in K$ is bilinear (check this) and hence determines a linear function $V^* \otimes V \rightarrow K$. To which familiar function on $\text{Hom}(V, V) = \text{End}(V)$ does this correspond?

The tensor product is associative in the sense that for three vector spaces V, W, X , the iterated tensor products $(V \otimes W) \otimes X$ and $V \otimes (W \otimes X)$ can be identified via $(v \otimes w) \otimes x \leftrightarrow v \otimes (w \otimes x)$. So we may suppress parentheses: $v \otimes w \otimes x \in V \otimes W \otimes X$. In fact, for a finite number of vector spaces V_1, \dots, V_k we have a direct definition (which we shall not spell out) of the k -fold tensor product $V_1 \otimes \dots \otimes V_k$ in terms of (the K vector space spanned by) the set of *multilinear* maps $M : V_1 \times \dots \times V_k \rightarrow W$. Such a multilinear map is then the composite of a universal multilinear map $V_1 \times \dots \times V_k \rightarrow V_1 \otimes \dots \otimes V_k$ followed by a linear map $m : V_1 \otimes \dots \otimes V_k \rightarrow W$.

3. Tensor algebra and exterior algebra

The tensor algebra. If V is a vector space, then we may abbreviate the k -fold tensor product $V \otimes \dots \otimes V$ by $V^{\otimes k}$, stipulating that for $k = 0$ this will stand for the

base field K . Notice that the multilinear map $V^k \times V^l = V^{k+l} \rightarrow V^{\otimes(k+l)}$ factors through a bilinear map $V^{\otimes k} \times V^{\otimes l} \rightarrow V^{\otimes(k+l)}$. This makes that the direct sum

$$T(V) := \bigoplus_{k \geq 0} V^{\otimes k}$$

is a ring if we take \otimes for product. Since this ring contains the base field K as a subring, it is in fact a K -algebra. We call it the *tensor algebra* of V .

If V is finite dimensional and (e_1, \dots, e_m) is a basis of V , then for $k > 0$ a basis of $V^{\otimes k}$ is the collection $e_{i_1} \otimes e_{i_2} \otimes \dots \otimes e_{i_k}$, where (i_1, \dots, i_k) runs over the sequences of length k in $\{1, \dots, m\}$. We may regard such expressions as ‘non-commutative monomials’ in e_1, \dots, e_m . As these are linearly independent, we can think of $T(V)$ as the ‘free (non-commutative) K -algebra’ generated by e_1, \dots, e_m .

The exterior algebra. The algebra $\wedge(V)$ that we are about to define is the ‘biggest’ quotient of the tensor algebra $T(V)$ with the property that for every $v \in V$, the image of $v \otimes v$ in $\wedge(V)$ is zero. Here is how to construct it. Let $I(V) \subset T(V)$ be the two-sided ideal generated by the elements $v \otimes v$, $v \in V$. So an element of this ideal is a linear combination of tensors of the form $v_1 \otimes v_2 \otimes \dots \otimes v_k$, for which the sequence (v_1, \dots, v_k) ‘stammers’ at least once: $v_\kappa = v_{\kappa+1}$ for some κ . It is clear that $I(V)$ is *homogeneous* in the sense that it is the direct sum of its homogeneous parts:

$$I(V) = \bigoplus_{k \geq 2} I(V)^k \text{ with } I(V)^k \subset V^{\otimes k}.$$

The quotient algebra $T(V)/I(V)$ is called the *exterior algebra* of V and denoted $\wedge(V)$. The degree k part $V^{\otimes k}/I(V)^k$ is written $\wedge^k V$, so that

$$\wedge(V) = \bigoplus_{k \geq 0} \wedge^k V.$$

In particular, $\wedge^0 V = K$ and $\wedge^1 V = V$. The product in $\wedge(V)$ is called the *exterior product* and denoted \wedge . So $v_1 \wedge \dots \wedge v_k$ is the image of $v_1 \otimes \dots \otimes v_k$. From the definition we see that $v \wedge v = 0$ for all $v \in V$. This implies that for all $v, v' \in V$, $v' \wedge v = -v \wedge v'$, because $v' \wedge v + v \wedge v' = (v + v') \wedge (v + v') - v \wedge v - v' \wedge v' = 0$. With induction we find that for every permutation $\sigma \in S_k$,

$$v_{\sigma(1)} \wedge v_{\sigma(2)} \wedge \dots \wedge v_{\sigma(k)} = \text{sign}(\sigma) v_1 \wedge v_2 \wedge \dots \wedge v_k,$$

where $\text{sign}(\sigma) \in \{\pm 1\}$ is of course the sign of σ . In particular, $v_1 \wedge v_2 \wedge \dots \wedge v_k = 0$ when $v_i = v_j$ for some $i \neq j$ (for we can arrange by means of a permutation that these vectors become neighbors). In particular, if $\alpha \in \wedge^k V$ and $\beta \in \wedge^l V$, then $\beta \wedge \alpha = (-1)^{kl} \alpha \wedge \beta$ (for this is true if α resp. β is of the form $v_1 \wedge \dots \wedge v_k$ resp. $v'_1 \wedge \dots \wedge v'_l$).

EXERCISE 3.1. Let V and W be vector spaces and $F : V^k \rightarrow W$ a multilinear map. Prove that F factors through a linear map $\wedge^k V \rightarrow W$ precisely when F is zero on any sequence in V of length k that ‘stammers’ (has two adjacent terms equal).

PROPOSITION 3.1. Suppose that V is finite dimensional and comes with a basis (e_1, \dots, e_m) . For any subset $I \subset \{1, \dots, m\}$, let

$$e_I := \begin{cases} e_{i_1} \wedge e_{i_2} \wedge \dots \wedge e_{i_k} & \text{when } \emptyset \neq I = \{1 \leq i_1 < i_2 < \dots < i_k \leq m\}; \\ 1 & \text{when } I = \emptyset. \end{cases}$$

Then:

- (i) The collection $(e_I)_I$, where I runs over the k -element subsets of $\{1, \dots, m\}$ is a K -basis of $\wedge^k V$, in particular, $\dim(\wedge^k V) = \binom{m}{k}$.

(ii) The K -algebra $\wedge(V)$ is generated as such by e_1, \dots, e_m and subject to the relations $e_i \wedge e_i = 0$ and $e_i \wedge e_j = -e_j \wedge e_i$ ($i < j$). This is a presentation of $\wedge(V)$ (so that any other relation follows from these).

(iii) Finally,

$$e_I \wedge e_J = \begin{cases} 0 & \text{in case } I \cap J \neq \emptyset \text{ and} \\ \pm e_{I \cup J} & \text{otherwise,} \end{cases}$$

where \pm is the sign of the permutation that turns the sequence IJ (obtained by juxtaposition) into an increasing sequence.

PROOF. We only prove (i), because this implies (ii), whereas (iii) is clear. Let us first show that the e_I 's generate $\wedge(V)$ as a K -vector space. The tensors $e_{i_1} \otimes \dots \otimes e_{i_k}$ generate $T(V)$ and so the $e_{i_1} \wedge \dots \wedge e_{i_k}$ generate $\wedge(V)$. But $e_{i_1} \wedge \dots \wedge e_{i_k}$ is zero unless all indices are distinct in which case it equals $\pm e_I$, where I is the index set.

We next verify that the e_I 's are linearly independent: given $I = (1 \leq i_1 < \dots < i_k \leq m)$, consider the function $V^k \rightarrow K$ defined by

$$(v_1, \dots, v_k) \in V^k \mapsto \sum_{\sigma \in \mathcal{S}_k} \text{sign}(\sigma) x^{i_1}(v_{\sigma(1)}) x^{i_2}(v_{\sigma(2)}) \dots x^{i_k}(v_{\sigma(k)}) \in K$$

Since this function is multilinear, it determines a linear function $V^{\otimes k} \rightarrow K$. The latter vanishes on the tensors $v_1 \otimes \dots \otimes v_k$ that 'stammer' (i.e., for which $v_\kappa = v_{\kappa+1}$ for some κ) and as such tensors span $I(V)^k$, this function factors through a linear function $\wedge^k V \rightarrow K$. It is easily checked that this function (which we denote by x^I) has the property that $x^I(e_J) = \delta_J^I$ (i.e., equals 1 if $J = I$ and is 0 otherwise). So the e_I 's are linearly independent. \square

EXAMPLES 3.2. For small values of m , the preceding may look familiar:

(i) For $m = 1$ we get $\wedge(K) = K[e]/(e^2)$.

(ii) A generator of $\wedge^2(K^2)$ is e_{12} . If we use this generator to identify $\wedge^2(K^2)$ with K , then \wedge is given by:

$$\begin{pmatrix} a_1 \\ a_2 \end{pmatrix} \wedge \begin{pmatrix} b_1 \\ b_2 \end{pmatrix} = a_1 b_2 - a_2 b_1 = \det \begin{pmatrix} a_1 & b_1 \\ a_2 & b_2 \end{pmatrix}.$$

(iii) A basis of $\wedge^2(K^3)$ is $f_1 := e_{23}, f_2 := e_{31}, f_3 := e_{12}$. If we use that basis to identify $\wedge^2(K^3)$ with K^3 , then

$$\begin{pmatrix} a_1 \\ a_2 \\ a_3 \end{pmatrix} \wedge \begin{pmatrix} b_1 \\ b_2 \\ b_3 \end{pmatrix} = \begin{pmatrix} a_2 b_3 - a_3 b_2 \\ a_3 b_1 - a_1 b_3 \\ a_1 b_2 - a_2 b_1 \end{pmatrix},$$

which reminds us of the formula for the classical exterior product.

REMARK 3.3. The proof of Proposition 3.1 shows that there is a natural pairing between $\wedge^k V$ and $\wedge^k(V^*)$, which is perfect in case V is finite dimensional. To be precise, the map

$$(V^*)^k \times V^k \rightarrow K, \quad (l_1, \dots, l_k, v_1, \dots, v_k) \mapsto \sum_{\sigma \in \mathcal{S}_k} \text{sign}(\sigma) l_{\sigma(1)}(v_1) \dots l_{\sigma(k)}(v_k),$$

is linear in each of its arguments and takes the value zero whenever two v_i 's or two l_i 's are equal, hence factors through a bilinear map

$$\wedge^k(V^*) \times \wedge^k V \rightarrow K.$$

Thus an element of $\wedge^k(V^*)$ can be regarded as a linear function on $\wedge^k V$. In case V is finite dimensional, the map thus obtained $\wedge^k(V^*) \rightarrow (\wedge^k V)^*$ is a linear isomorphism, because in the notation of Proposition 3.1 the image of the pair (x^I, e_I) equals δ_I^I : $(x^I)_I$ may be understood as the basis dual to $(e_I)_I$.

EXERCISE 3.2. Show that for a vector space V of (finite) dimension m and $k = 0, \dots, m$, the map

$$P : \wedge^k V \rightarrow \text{Hom}(\wedge^{m-k} V, \wedge^m V), \quad P(a) : b \mapsto a \wedge b,$$

is a linear isomorphism.

EXERCISE 3.3. Prove that for finite dimensional vector spaces V and W , the map $V \oplus W \rightarrow \wedge(V) \otimes \wedge(W)$, $(v, w) \mapsto v \otimes 1 + 1 \otimes w$ is naturally extended to an isomorphism of (graded) vector spaces $\wedge(V \oplus W) \cong \wedge(V) \otimes \wedge(W)$ that takes the wedge product on the domain to a product on the range with the property that for homogeneous elements, $(a \otimes b) \cdot (a' \otimes b') = (-1)^{\deg(b) \deg(a')} (a \wedge a') \otimes (b \wedge b')$.

A linear map $s : V \rightarrow W$ of vector spaces determines for $k = 0, 1, 2, \dots$ a linear map $s^{\otimes k} : V^{\otimes k} \rightarrow W^{\otimes k}$. These maps are the homogeneous parts of a map $T(s) : T(V) \rightarrow T(W)$ that is clearly a homomorphism of algebras. Since this takes $I(V)$ to $I(W)$, we also find for every k a linear map

$$\wedge^k s : \wedge^k V \rightarrow \wedge^k W, \quad v_1 \wedge \dots \wedge v_k \mapsto s(v_1) \wedge \dots \wedge s(v_k)$$

that is the degree k -part of an algebra homomorphism $\wedge(s) : \wedge(V) \rightarrow \wedge(W)$. Of particular interest is the case $W = V$:

PROPOSITION 3.4. If V is of finite dimension m , then $\wedge^m s : \wedge^m V \rightarrow \wedge^m V$ is multiplication by $\det(s)$.

PROOF. Choose a basis e_1, \dots, e_m for V . If $s(e_i) = \sum_j s_{ij} e_j$, then

$$\begin{aligned} \wedge^m(s)(e_{12\dots m}) &= s(e_1) \wedge \dots \wedge s(e_m) \\ &= \left(\sum_j s_{1j} e_j \right) \wedge \dots \wedge \left(\sum_j s_{mj} e_j \right) \\ &= \sum_{j_1, j_2, \dots, j_m} s_{1j_1} e_{j_1} \wedge s_{2j_2} e_{j_2} \wedge \dots \wedge s_{mj_m} e_{j_m}. \end{aligned}$$

In the latter sum only the terms for which (j_1, \dots, j_m) all differ can contribute. Such a sequence (j_1, \dots, j_m) is the image of $(1, 2, \dots, m)$ under a permutation $\sigma \in \mathcal{S}_m$, and hence this sum equals

$$\sum_{\sigma \in \mathcal{S}_m} \text{sign}(\sigma) s_1^{\sigma(1)} s_2^{\sigma(2)} \dots s_m^{\sigma(m)} e_1 \wedge \dots \wedge e_m = \det(s) e_{12\dots m}.$$

□

This suggests we should *define* the determinant of s this way: namely as the scalar by which $\wedge^m s$ acts on the one dimensional vector space $\wedge^m V$, for it is then a priori clear that this is independent of the choice of a basis.

EXERCISE 3.4. Let V be a vector space of finite dimension m and let $s : V \rightarrow V$ be an endomorphism of V .

(a) Prove that the trace of $\wedge(s) : \wedge(V) \rightarrow \wedge(V)$ equals $\det(s + 1_V)$. (Upon replacing s by λs , $\lambda \in K$, we then see that the trace of $\wedge^k s : \wedge^k V \rightarrow \wedge^k V$ is equal to the

coefficient of λ^k in the modified characteristic polynomial $\det(\lambda s + \mathbf{1}_V)$.)

(b) Choose a generator μ of $\wedge^m V$ and let $B_k : \wedge^k V \times \wedge^{m-k} V \rightarrow K$ be the bilinear map defined by $\alpha \wedge \beta = B_k(\alpha, \beta)\mu$. Prove that $B_k(\wedge^k s(\alpha), \wedge^{m-k} s(\beta)) = \det(s)B_k(\alpha, \beta)$. Let $(\wedge^{m-k} s)^t : \wedge^k V \rightarrow \wedge^k V$ be the adjoint of $\wedge^{m-k} s$ relative to B_k (see Exercises 1.2 and 3.2). Prove that the composite $(\wedge^{m-k} s)^t \circ (\wedge^k s)$ is scalar multiplication in $\wedge^k V$ by $\det(s)$. (For $k = 1$ this is an abstract form of Cramer's rule.)

Internal contraction. Recall that every $v \in V$ defines a linear function $\iota_v : V^* \rightarrow K, l \mapsto l(v)$.

PROPOSITION 3.5. *For every $v \in V$ there is precisely one linear map*

$$\iota_v : \wedge(V^*) \rightarrow \wedge(V^*)$$

which decreases the degree by one and enjoys the following two properties:

- (i) *if $l \in V^*$, then $\iota_v(l) = l(v) \in K = \wedge^0 V$ (in other words, ι_v agrees on V^* with the earlier definition) and*
- (ii) *if $\alpha \in \wedge^k(V^*)$ and $\beta \in \wedge(V^*)$, then*

$$\iota_v(\alpha \wedge \beta) = \iota_v(\alpha) \wedge \beta + (-1)^k \alpha \wedge \iota_v(\beta).$$

We then have $\iota_v^2 = 0$ and the map $V \times \wedge(V^) \rightarrow \wedge(V^*), (v, \alpha) \mapsto \iota_v(\alpha)$ is bilinear.*

PROOF. The value of ι_v on $l_1 \wedge \cdots \wedge l_k, l_i \in V^*$ is determined by the properties above:

$$\iota_v(l_1 \wedge \cdots \wedge l_k) = \iota_v(l_1)l_2 \wedge \cdots \wedge l_k - l_1 \wedge \iota_v(l_2 \wedge \cdots \wedge l_k)$$

If we substitute $\iota_v(l_1) = l_1(v)$ on the right and proceed inductively with the factor $\iota_v(l_2 \wedge \cdots \wedge l_k)$, we find

$$\iota_v(l_1 \wedge \cdots \wedge l_k) = \sum_{\kappa=1}^k (-1)^{\kappa-1} l_\kappa(v) l_1 \wedge \cdots \wedge \widehat{l_\kappa} \wedge \cdots \wedge l_k,$$

where as usual the hatted vector must be omitted. This proves the uniqueness assertion.

On the other hand, we may define ι_v this way. To be precise, consider for $k = 1, \dots, m$ the map

$$(V^*)^k \rightarrow \wedge^{k-1}(V^*), \quad (l_1, \dots, l_k) \mapsto \sum_{\kappa=1}^k (-1)^{\kappa-1} l_\kappa(v) l_1 \wedge \cdots \wedge \widehat{l_\kappa} \wedge \cdots \wedge l_k.$$

This map is multilinear (clear) and it is easy to see that it vanishes on (l_1, \dots, l_k) whenever the latter stammers. So the map factors through a linear map $\iota_v : \wedge^k(V^*) \rightarrow \wedge^{k-1}(V^*)$ (see Exercise 3.1). It is easy to verify that it has all the claimed properties. \square

REMARK 3.6. By interchanging the role of V and V^* we find for every $l \in V^*$ a map:

$$\iota_l : \wedge(V) \rightarrow \wedge(V)$$

enjoying similar properties.

EXERCISE 3.5. Let V be a vector space.

(a) Given $(v_1, \dots, v_k) \in V^k$, consider the map

$$\iota_{v_1} \iota_{v_2} \cdots \iota_{v_k} : \wedge(V^*) \rightarrow \wedge(V^*)$$

Prove that this map only depends on $v_1 \wedge v_2 \wedge \cdots \wedge v_k$ and thus defines a linear map

$$\wedge^k(V) \times \wedge(V^*) \rightarrow \wedge(V^*), \quad (\alpha, \alpha) \mapsto \iota_\alpha(\alpha).$$

(b) Prove that for $(v_1, \dots, v_k) \in V^k$ and $(l_1, \dots, l_k) \in (V^*)^k$ we have

$$\iota_{v_k} \wedge \iota_{v_{k-1}} \wedge \cdots \wedge \iota_{v_1} (l_1 \wedge l_2 \wedge \cdots \wedge l_k) = \sum_{\sigma \in \mathcal{S}_k} \text{sign}(\sigma) l_1(v_{\sigma(1)}) \cdots l_k(v_{\sigma(k)})$$

and hence equals the image of $(l_1 \wedge \cdots \wedge l_k, v_1 \wedge \cdots \wedge v_k)$ under the pairing described in Remark 3.3.

Orientation. We here assume $K = \mathbb{R}$. If V is a real vector space of finite dimension m , then $\wedge^m V$ is a 1-dimensional real vector space. This is also true for $m = 0$, because we simply agreed that in that case $\wedge^0 V = \mathbb{R}$. So $\wedge^m V - \{0\}$ has two components (two elements lie in the same component if their ratio is positive). For $m = 0$, there is a privileged component (namely the interval of positive reals). But for $m > 0$ this is no longer the case and so if we single one out, then we are in fact introducing (literally!) one bit of extra structure on V . The choice of such a component is what we call an *orientation* of V . To rephrase the preceding: V is not naturally oriented unless it is zero dimensional.

If $m > 0$ and V is oriented and (e_1, \dots, e_m) is an (ordered) basis of V , then either $e_1 \wedge \cdots \wedge e_m$ or $-e_1 \wedge \cdots \wedge e_m$ is in the chosen component of $\wedge^m V$. In the first case we say that the basis is *oriented* relative to the given orientation. Conversely, the choice of an ordered basis (e_1, \dots, e_m) of V determines an orientation of V as being the component of $\wedge^m V - \{0\}$ that contains $e_1 \wedge \cdots \wedge e_m$. For $m = 0$ (so $V = \{0\}$) an orientation is simply given by a sign (\pm) that refers to the corresponding component of $\wedge^0 V - \{0\} = \mathbb{R} - \{0\}$.

EXERCISE 3.6. Prove that two bases (e_1, \dots, e_m) and (e'_1, \dots, e'_m) of a real vector space define the same orientation precisely when the matrix $(a_i^j)_{i,j}$ which expresses one in the other: $e'_i = \sum_j a_i^j e_j$, has positive determinant.

Suppose that V is oriented and $H \subset V$ is a hyperplane (we assume $m > 0$ here). If H is given as the boundary of a half space $V_- \subset V$ (i.e., there is some $l \in V^* - \{0\}$ such that H resp. V_- is defined by $l = 0$ resp. $l \leq 0$), then H may be oriented as follows. If $m > 1$, choose a vector $v \in V - V_-$ and stipulate that a basis (e_1, \dots, e_{m-1}) of H is oriented if (v, e_1, \dots, e_{m-1}) is an oriented basis of V (this is indeed independent of the choice of v). If $m = 1$ (so that the orientation is a component of $V - \{0\}$ and $H = \{0\}$), then we give H the sign \pm that we need to make $\pm v$ lie in the orientation component of $V - \{0\}$. In either case we call this the *induced orientation* of H .

EXERCISE 3.7. Prove that in the situation above, $\iota_l : \wedge^m V \rightarrow \wedge^{m-1} H$ is an isomorphism of one-dimensional vector spaces that takes the orientation of V to the induced orientation of H .

EXERCISE 3.8. Let V and W be oriented real finite dimensional vector spaces. Show that $V \oplus W$ and $V \otimes W$ can be oriented in a natural manner.

If V is oriented, then there are (at least) two conventions for orienting the dual V^* . We choose the following one: is (e_1, \dots, e_m) an oriented basis of V , then the dual basis (x^1, \dots, x^m) is by definition an oriented basis of V^* . (Another convention takes (x^m, \dots, x^1) instead.)

4. Quadratic forms

In this section all vector spaces are finite dimensional and real (so $K = \mathbb{R}$).

Let V be a vector space. A bilinear function $g : V \times V \rightarrow \mathbb{R}$ that is symmetric in the sense that $g(v, w) = g(w, v)$ determines a function $q = q_g : V \rightarrow \mathbb{R}, v \mapsto g(v, v)$ satisfying

- q is homogeneous of degree 2: $q(\lambda v) = \lambda^2 q(v)$,
- the map $(v, w) \in V^2 \mapsto q(v+w) - q(v) - q(w) \in \mathbb{R}$ is bilinear (and equal to $2g$).

DEFINITION 4.1. We call a *quadratic form* on V a function $q : V \rightarrow \mathbb{R}$ that satisfies these two properties.

A quadratic form q defines in turn a symmetric bilinear form g on V by the formula $g_q(v, w) := \frac{1}{2}(q(v+w) - q(v) - q(w))$. The assignments $g \mapsto q_g$ and $q \mapsto g_q$ are each others inverse, so that giving a symmetric bilinear form on V is equivalent to giving a quadratic form on V .

REMARK 4.2. The notion of a quadratic form makes of course sense for any field K and not just \mathbb{R} . But the correspondence between such forms and symmetric bilinear functions requires that we can divide by 2 and hence is only valid if the characteristic of K is $\neq 2$. It is then usual to use a slightly different convention that makes the quadratic form the stronger notion in the sense that the associated bilinear form is given by $q(v+w) - q(v) - q(w)$ (rather than half of this). So in that case, the quadratic form attached to $g : V \times V \rightarrow K$ is only defined if the characteristic of K is $\neq 2$ and then given by $q(v) = \frac{1}{2}b(v, v)$. It turns out that this convention (we shall not use) would also be the preferred one in some physical contexts (as the formula for the kinetic energy shows).

Let g be a symmetric bilinear form on the vector space V . Given a basis (e_1, \dots, e_m) of V , then we may form the matrix $(g_{ij} := g(e_i, e_j))_{i,j}$. This matrix is clearly symmetric. If (x^1, \dots, x^m) is the basis of V^* dual to (e_1, \dots, e_m) , then $g(v, w) = \sum_{i,j} g_{ij} x^i(v) x^j(w)$ and hence the associated quadratic form is given by

$$q = \sum_{i,j} g_{ij} x^i x^j = \sum_i g_{i,i} (x^i)^2 + 2 \sum_{i<j} g_{ij} x^i x^j.$$

LEMMA 4.3 (Diagonalization). *There is a basis for V such that the matrix of g is a diagonal matrix. So if (x^1, \dots, x^m) is the corresponding dual basis, then $q = \sum_i \lambda_i (x^i)^2$.*

PROOF. With induction on $m := \dim V$. If g is identically zero, there is nothing to show. If not, then q is not identically zero either and so there is a $e_1 \in V$ with $q(e_1) \neq 0$. Let $V' \subset V$ be the zero set of the linear function $v \in V \mapsto g(v, e_1)$. Since that function is nonzero (it is nonzero in e_1), $\dim(V') = m - 1$. The induction hypothesis applied to $g|_{V' \times V'}$ yields a basis (e_2, \dots, e_m) of V' for which $(g(e_i, e_j))_{2 \leq i, j \leq m}$ has the diagonal form. Then (e_1, \dots, e_m) is a basis of V as desired, for $g(e_i, e_0) = g(e_0, e_i) = 0$ voor $i = 2, \dots, m$. \square

The real numbers λ_i that appear in this lemma have no intrinsic meaning for we may adapt the basis further as follows: if $\lambda_i \neq 0$ then replace e_i by $|\lambda_i|^{-1/2}e_i$ such that now all nonzero diagonal coefficients become ± 1 and then re-order the basis elements so that the diagonal begins with 1's and ends with 0's. On such a basis q takes the so-called *standard form*:

$$q = \sum_{i=1}^p (x^i)^2 - \sum_{i=p+1}^{p+n} (x^i)^2.$$

Lemma 4.5 below shows that the pair (p, n) does have an intrinsic (basis independent) meaning.

DEFINITION 4.4. We call a quadratic form $q : V \rightarrow \mathbb{R}$ (and the associated symmetric bilinear form) *positive definite* (resp. *negative definite*) if $q(v) > 0$ (resp. $q(v) < 0$) for all $v \in V - \{0\}$.

LEMMA 4.5. *The number p resp. n is the maximal dimension of a subspace of V on which q is positive (resp. negative) definite.*

PROOF. Let $W \subset V$ be a subspace on which q is positive definite. We prove that $\dim(W) \leq p$ and that the value p is attained. Let e_1, \dots, e_m be a basis as above, V_+ the span of the basis vectors e_i with $i \leq p$ (so of dimension p) and $\pi : V \rightarrow V_+$ the projection on V_+ with kernel the span of the remaining basis vectors. It is clear that q is positive definite on V_+ and ≤ 0 on e_{p+1}, \dots, e_m . So $W \cap \text{Ker}(\pi) = 0$. In other words, π maps W injectively to V_+ . In particular, $\dim W \leq p$. On the other hand, the value p is attained for $W = V_+$.

The analogous statement for n follows by applying the preceding to $-q$. \square

It is clear that q is positive definite precisely when $p = m$.

Let $g : V \times V \rightarrow \mathbb{R}$ be symmetric bilinear. Every $v \in V$ defines a linear function $g_v : V \rightarrow \mathbb{R}$ by $g_v(v') = g(v, v')$. There results a map $\tilde{g} : V \rightarrow V^*$, $v \mapsto g_v$, which satisfies $\tilde{g}(v)(v') = g(v, v')$ and which is obviously linear. We say that g is a *pseudo-inner product* if $g_v = 0$ implies $v = 0$, in other words, if \tilde{g} is injective. Since V and V^* have the same dimension this amounts to: \tilde{g} is an isomorphism. If we substitute $v' = \tilde{g}^{-1}(l)$ in the formula $\tilde{g}(v')(v) = g(v', v)$, then we see that

$$l(v) = g(\tilde{g}^{-1}(l), v).$$

We may also use \tilde{g} to transfer g to V^* :

$$\check{g} : V^* \times V^* \rightarrow \mathbb{R}, \quad \check{g}(l, l') := g(\tilde{g}^{-1}(l), \tilde{g}^{-1}(l')) = l(\tilde{g}^{-1}(l')).$$

This is perhaps easier understood in terms of a basis (e_1, \dots, e_m) : \tilde{g} is then given by $\tilde{g}(e_i) = \sum_j g_{ij}x^j$ and so g is a pseudo-inner product precisely when the matrix $(g_{ij})_{ij}$ is nonsingular. The matrix $\check{G} := (g^{ij} := \check{g}(x^i, x^j))$ is then simply the inverse of the matrix $G := (g_{ij})$, because it follows from

$$\begin{aligned} g_{ij} &= g(e_i, e_j) = \check{g}(\tilde{g}(e_i), \tilde{g}(e_j)) \\ &= \check{g}\left(\sum_k g_{ik}x^k, \sum_l g_{jl}x^l\right) = \sum_{k,l} g_{ik}g^{kl}g_{lj}, \end{aligned}$$

that $G = G\check{G}G$ so that $\check{G} = G^{-1}$. For this reason the quadratic form \check{g} is often called (by a slight abuse of language) the *inverse* of g .

Things simplify considerably if g is diagonalized by the basis (e_1, \dots, e_m) : $g(e_i, e_j) = \lambda_i \delta_{ij}$, for then $\tilde{g}(e_i) = \lambda_i x^i$, and so g is a pseudo-inner product precisely when $\lambda_i \neq 0$ for all i . In that case is $\tilde{g}(x^i, x^j) = \lambda_i^{-1} \delta^{ij}$. This makes it plain that \tilde{g} has the same signature as g .

Inner products. A symmetric bilinear form g on V is called an *inner product* when it is positive definite. We then refer to the vector space V endowed with such a g as an *inner product space*. According to the preceding there exists a basis (e_1, \dots, e_m) of V such that $g(e_i, e_j) = \delta_{i,j}$ (and hence $\tilde{g}(e_i) = x^i$). Such a basis is called *orthonormal*.

If g is merely a pseudo-inner product, then we call the pair (V, g) a *pseudo-inner product space*. The linear transformations $s : V \rightarrow V$ which preserve such a pseudo-inner product, i.e., which satisfy $g(sv, sv') = g(v, v')$, are called *orthogonal* and form a group, the *orthogonal group* of (V, g) , under composition. We denote that group by $O(V, g)$. Suppose (e_1, \dots, e_m) is a basis of V on which g takes the standard form: if $(p, m-p)$ is the signature of g , then $(g(e_i, e_j))_{i,j}$ is the diagonal matrix $I_{p, m-p}$ which has on the diagonal $(1, \dots, 1, -1, \dots, -1)$. Assigning to $s \in O(V, g)$ its matrix on this basis identifies $O(V, g)$ with a matrix group that is usually denoted by $O(p, m-p)$. The latter consists of the matrices X which satisfy $X^t I_{p, m-p} X = I_{p, m-p}$. In the case of an inner product: $p = m$ this is the standard orthogonal group $O(m)$, and in case $(p, m-p) = (1, 3)$, this is known as the *Lorentz group*.

LEMMA 4.6. *A symmetric bilinear form g on V extends to one $\wedge(V)$ such that*

$$g(v_1 \wedge \dots \wedge v_k, v'_1 \wedge \dots \wedge v'_k) = \begin{cases} \sum_{\sigma \in S_k} \text{sign}(\sigma) g(v_1, v'_{\sigma(1)}) \cdots g(v_k, v'_{\sigma(k)}) & \text{when } k' = k; \\ 0 & \text{otherwise.} \end{cases}$$

This extension is a (pseudo-)inner product when g is. Precisely, is (e_1, \dots, e_m) a basis of V such that $g(e_i, e_j) = \lambda_i \delta_{i,j}$, then $g(e_I, e_J) = (\prod_{i \in I} \lambda_i) \delta_{I,J}$.

PROOF. The map $V^k \times V^k \rightarrow \mathbb{R}$ defined by the formula

$$(v_1, \dots, v_k; v'_1, \dots, v'_k) \mapsto \sum_{\sigma \in S_k} \text{sign}(\sigma) g(v_1, v'_{\sigma(1)}) \cdots g(v_k, v'_{\sigma(k)})$$

is multilinear and zero whenever one of the sequences stammers. An application of Exercise 3.1 then yields the desired extension. The symmetry follows from the definition.

The second part is easily checked. \square

We emphasized that a quadratic form g on a real vector space V of dimension m has no other intrinsic invariants than its signature. But if the vector space already comes with a fixed inner product g_o , then as the following proposition will show, things are different and we can speak of the eigen values of g . The proposition in question is equivalent to the fact that for every real symmetric $(m \times m)$ -matrix there is an orthonormal basis consisting of eigen values.

PROPOSITION 4.7. *Let (V, g_o) be inner product space. Then for every symmetric bilinear form g on V there is a g_o -orthonormal basis (e_1, \dots, e_m) of V on which g takes the diagonal form: $g(e_i, e_j) = \lambda_i \delta_{ij}$ for certain λ_i . The basis (e_1, \dots, e_m) consists of*

eigen vectors of the transformation $(\tilde{g}_o)^{-1}\tilde{g} : V \rightarrow V$ with eigen values $(\lambda_1, \dots, \lambda_m)$. In particular, these numbers are intrinsic to the triple (V, g_o, g) (and are called the eigen values of g on (V, g_o)).

PROOF. We show the existence of the basis with induction on m . For $m = 0$ there is nothing to prove. So suppose that $m > 0$ and that the proposition has been proved in dimension $< m$. Let $S \subset V$ be the unit sphere defined by $q_o(v) = 1$. This is a compact subset of V , and so $q|_S$ has a minimum, that is taken in $e_m \in S$, say. With the method of Lagrange multipliers it follows that the derivatives of the functions q and q_o are proportional in e_m . This means that there is a $\lambda_m \in \mathbb{R}$ such that $g(e_m, v) = \lambda_m g_o(e_m, v)$ for all v . Applying the induction assumption to the g_o -orthogonal complement of e_m , gives a basis (e_1, \dots, e_m) orthonormal for g_o on which g takes the a diagonal form.

The last assertion follows from the fact that $\tilde{g}_o(e_i) = x^i$ and $\tilde{g}(e_i) = \lambda_i x^i$. \square

EXERCISE 4.1. Show that every linear map $V \rightarrow V$ which preserves a pseudo-inner product has determinant ± 1 .

EXERCISE 4.2. Prove that two quadratic forms g, g' on an inner product space (V, g_o) can be transformed into each other by an orthogonal transformation (i.e., there is some $s \in O(V, g_o)$ such that $g'(v, w) = g(sv, sw)$ for all $v, w \in V$) precisely when they have the same eigen values up to order relative to (V, g_o) .

EXERCISE 4.3. An analogue of Proposition 4.7 no longer holds if g_o is merely a pseudo-inner product. Illustrate this by proving that the quadratic forms $x^2 - y^2$ and xy on \mathbb{R}^2 cannot be simultaneously diagonalized.

EXERCISE 4.4. Let (V, g) be a pseudo-inner product space of signature (p, n) . Show that the line $\wedge^m V$ is positive definite if n is even and negative definite if n is odd.

If g is an inner product, then we write $\|v\| := \sqrt{g(v, v)}$. Note that then for all $v, v' \in V$ and $\lambda \in \mathbb{R}$:

- (i) $\|v\| \geq 0$ with $\|v\| = 0$ precisely when $v = 0$,
- (ii) $\|\lambda v\| = |\lambda| \cdot \|v\|$,
- (iii) $\|v + v'\| \leq \|v\| + \|v'\|$.

These properties amount to saying that $\|\cdot\|$ is a *norm* for V . This makes V a metric space if we let $\|v - v'\|$ be the distance between v and v' .

In a pseudo-inner product space (V, g) the notion of a gradient makes sense: For $U \subset V$ open, we have $TU = U \times V$ and since g defines an isomorphism $\tilde{g} : V \rightarrow V^*$ we may identify $TU = U \times V$ with $U \times V^* = T^*U$. The *gradient* of a function $f : U \rightarrow \mathbb{R}$ is then the vector field $\text{grad}(f)$ on U that under this isomorphism corresponds to df . So if (e_1, \dots, e_m) is a basis on which g takes the form $\sum_i \lambda_i (x^i)^2$, then \tilde{g} makes the basis vector field $\partial/\partial x^i$ correspond to the differential $\lambda_i dx^i$ and hence

$$\text{grad}(f) = \sum_{i=1}^m \lambda_i^{-1} \frac{\partial f}{\partial x^i} \frac{\partial}{\partial x^i}.$$

If g is an inner product and (e_1, \dots, e_m) is orthonormal (so then $\lambda_i = 1$ for all i), then we get the familiar form.

5. The star operator

Let (V, g) be a pseudo-inner product space of finite dimension m and of signature $(m - n, n)$ and suppose V endowed with an orientation. In the previous section we extended the pseudo-inner product to all of $\wedge(V)$. By Exercise 4.4 this extension is on the line $\wedge^m V$ positive or negative according to whether n is even or odd. In either case we have on this line two generators on which the quadratic form is 1 in absolute value. These generators are each other antipode and so the given orientation of V singles out one of them: we let μ be the unique generator which is in the privileged component and for which $|g(\mu, \mu)| = 1$. We call this generator the *oriented volume element* of V . Concretely: is (e_1, \dots, e_m) an oriented basis for V for which $g(e_i, e_j) = \lambda_i \delta_{i,j}$, then

$$\mu = |\lambda_1 \lambda_2 \cdots \lambda_m|^{-\frac{1}{2}} e_1 \wedge \cdots \wedge e_m.$$

Let be given $\beta \in \wedge^k V$ and consider the linear map

$$\alpha \in \wedge^k V \mapsto g(\alpha, \beta)\mu \in \wedge^m V.$$

This map is according to Exercise 3.2 obtained by wedging with a unique element of $\wedge^{m-k} V$. We denote that element by $\star\beta$. So the map

$$\star : \wedge^k V \rightarrow \wedge^{m-k} V$$

thus defined is characterized by the property that

$$\alpha \wedge \star\beta = g(\alpha, \beta)\mu \text{ for all } \alpha, \beta \in \wedge^k V.$$

This formula shows that \star is linear and that $\star(1) = \mu$. Let us compute $\star(e_I)$, where $I \subset \{1, \dots, m\}$ is a k -element subset. We must have

$$e_J \wedge \star(e_I) = g(e_J, e_I)\mu$$

for all k -element subsets $J \subset \{1, \dots, m\}$. The right hand side is nonzero only if $J = I$. If we write $\star(e_I)$ out on the basis $(e_K)_{|K|=m-k}$ of $\wedge^{m-k} V$, then we see that $\star(e_I)$ is proportional to $e_{I'}$, where $I' = \{1, \dots, m\} - I$: $\star(e_I) = c_I e_{I'}$ for some $c_I \in \mathbb{R}$. Substitution then yields $c_I e_I \wedge e_{I'} = g(e_I, e_I)\mu$ or

$$c_I = \text{sign}(I, I') \lambda_I |\lambda_{12\dots m}|^{-\frac{1}{2}}.$$

(Here (I, I') is the permutation which puts the juxtaposed sequence II' in increasing order $(1, 2, \dots, m)$.) In particular, \star is an isomorphism of $\wedge^k V$ onto $\wedge^{m-k} V$. If (e_1, \dots, e_m) happens to be a standard basis (so one for which λ_i is 1 or -1 according to whether $i \leq p$ or $i > p$), then we see from this formula that $c_I = \text{sign}(I, I')(-1)^{n(I)} \in \{1, -1\}$, where $n(I) := |I \cap \{p+1, \dots, m\}|$.

The formula also shows that $c_{I'} = \text{sign}(I', I) \lambda_{I'} |\lambda_{12\dots m}|^{-\frac{1}{2}}$ so that

$$c_{I'} c_I = \text{sign}(I', I) \text{sign}(I, I') \frac{\lambda_{I'} \lambda_I}{|\lambda_{12\dots m}|} = (-1)^{k(m-k)} \frac{\lambda_{12\dots m}}{|\lambda_{12\dots m}|} = (-1)^{k(m-k)+n}.$$

This implies that $\star\star$ acts on $\wedge^k V$ as multiplication by $(-1)^{k(m-k)+n}$.

EXAMPLE 5.1. Is $m = 3$ and is g an inner product, then the map $\star : V \wedge V \rightarrow V$ yields a product on V : $v \times v' := \star(v \wedge v')$. This is in fact a basis independent characterization of the classical exterior product, for if (e_1, e_2, e_3) is an oriented orthonormal basis, then $e_1 \times e_2 = e_3$ etc.

EXAMPLES 5.2. Many familiar differential operators can be expressed in terms of the exterior derivative and the star operator. Suppose (V, g) is an oriented pseudo-inner product space. If $U \subset V$ is open, then we let \star act pointwise on the tangent bundle $U \times V$ and the cotangent bundle $U \times V^*$ and thus on the space of differential forms $\mathcal{E}^k(U)$. Many familiar differential operators can be expressed in terms of \star and the exterior derivative. Let (e_1, \dots, e_m) be an oriented standard basis so that $g(e_i, e_j) = \lambda_i \delta_{i,j}$ with λ_i equaling ± 1 depending on whether $i \leq p$ or $i > p$. We denote by (x^1, \dots, x^m) the dual basis of V^* .

(i) The divergence of a vector field W on U : We have a volume element $\mu := \star(1) = dx^1 \wedge \dots \wedge dx^m$ and hence a divergence $\operatorname{div}(W) : U \rightarrow \mathbb{R}$ defined by

$$\operatorname{div}(W) \star(1) := dt_W \star(1).$$

The right hand side is $\sum_i \partial W^i / \partial x^i$. So this recovers the usual definition.

(ii) The composite $\operatorname{div} \operatorname{grad}$ yields the Laplace operator acting on functions up to sign. We let it be $\star d \star d$:

$$\begin{aligned} \star d \star df &= \star d \left(\sum_i (-1)^{i-1} \lambda_i^{-1} \frac{\partial f}{\partial x^i} dx^1 \wedge \dots \wedge \widehat{dx^i} \wedge \dots \wedge dx^m \right) \\ &= \star \left(\sum_i \lambda_i^{-1} \frac{\partial^2 f}{\partial x^i \partial x^i} dx^1 \wedge \dots \wedge dx^m \right) = \sum_i \lambda_i^{-1} \frac{\partial^2 f}{\partial x^i \partial x^i}, \end{aligned}$$

For the case $(p, m-p) = (1, 3)$, this operator plays an important role in the theory of special relativity and is known as the *Alembertian*.

In both cases (i) and case (ii), the \star -operator appears twice in the definition and so neither the divergence nor the Laplace operator depends on the orientation. This is different in the next example.

(iii) For the case $m = 3$ with g an inner product, we have the operator rot that acts on vector fields. The inner product turns vector fields into differentials and vice versa and so if we use this to regard rot as an operator on differentials, then

$$\operatorname{rot}(\alpha) = \star d\alpha,$$

because

$$\star d(fdx + gdy + hdz) = \star \left(\left(\frac{\partial h}{\partial y} - \frac{\partial g}{\partial z} \right) dy \wedge dz + \dots \right) = \left(\frac{\partial h}{\partial y} - \frac{\partial g}{\partial z} \right) dx + \dots.$$

Hermitian forms. We mentioned in Remark 4.2 that the notion of a quadratic form makes sense over other fields than \mathbb{R} , so in particular over \mathbb{C} . But in the complex setting there is another notion that is more like an inner product.

DEFINITION 5.3. Let V be a complex vector space. A map $h : V \times V \rightarrow \mathbb{C}$ is called a *Hermitian form* if it is \mathbb{C} -linear in the first variable and $\overline{H(v, v')} = H(v', v)$ for all $v, v' \in V$.

Notice that then H is \mathbb{C} -antilinear in the second variable, because $H(v, \lambda v') = \overline{H(\lambda v', v)} = \overline{\lambda H(v', v)} = \overline{\lambda} H(v, v')$ and $H(v, v' + v'') = H(v', v) + H(v'', v)$. It is clear that $H(v, v)$ is always real.

For every $v' \in V$, the function $h(\cdot, v')$ defines a linear form on V . We thus get a map $\tilde{h} : V \rightarrow V^*$. This map is \mathbb{C} -antilinear and we therefore like to write this a map $\tilde{h} : \overline{V} \rightarrow V^*$. Here \overline{V} has the same underlying real vector space as V , but the multiplication by complex scalars has changed: multiplication by $\lambda \in \mathbb{C}$ in \overline{V} is by

definition multiplication by $\bar{\lambda}$ in V . This makes $\tilde{h} : \bar{V} \rightarrow V^*$ a complex-linear map. We say h is *nondegenerate* or that h is a *pseudo-inner product* if \tilde{h} is injective.

We say that h is *positive (resp. negative) definite* if $h(v, v) > 0$ resp. $h(v, v) < 0$ for all $v \in V - \{0\}$. A positive Hermitian form is also called an *inner product*; $\|v\| := \sqrt{h(v, v)}$ is then a norm on V (it satisfies the same properties of the norm in the real case; the difference is now that in the property $\|\lambda v\| = |\lambda| \cdot \|v\|$, we may take $\lambda \in \mathbb{C}$). If V is finite dimensional, then V is *complete* for this norm (which means that every Cauchy sequence converges). In general this need not be the case, but if so, then we say that V is a *Hilbert space*.

If e_1, \dots, e_m is a complex basis of V , then $h(e_i, e_j)$ is a Hermitian matrix: its transpose equals its complex conjugate. We can always find a basis (e_1, \dots, e_m) for which this matrix takes the (diagonal) form: $h(e_i, e_j) = \lambda_i \delta_{i,j}$. We can even arrange that the sequence $(\lambda_1, \dots, \lambda_m)$ begins with 1's, is followed by -1 's and ends with 0's. The number of 1's (resp. -1 's) is the dimension of a maximal positive (resp. negative) definite subspace and the pair (p, n) is called the *signature* of h .

There is an analogue of Lemma 4.6 which we will not bother to state.

6. Affine spaces, Euclidian spaces and Minkowski space

An affine space is essentially a vector space without an origin: its points are not vectors, but the two points of such space differ by one. In other words we can translate the space over a vector.

DEFINITION 6.1. Let V be a vector space. An *affine space over V* is a set \mathbb{V} on which the additive group of V acts freely and transitively: in other words for any pair $O, P \in \mathbb{V}$ there is precisely one $v \in V$ that takes O into P . The *dimension* of \mathbb{V} will be that of V .

We denote the action by a $+$: $P = v + O$; we also write $P - O \in V$ for v so that $P = (P - O) + P$, but a sum like $O + P$ is not defined. For a given $v \in V$, the map $P \in \mathbb{V} \mapsto v + P \in \mathbb{V}$ is called the *translation over v* . We therefore call \mathbb{V} the translation space of \mathbb{V} .

For a given $O \in \mathbb{V}$, the map $v \in V \mapsto v + O \in \mathbb{V}$ is a bijection and via this bijection, \mathbb{V} becomes a vector space with origin O .

Clearly, a vector space is an affine space over itself.

EXAMPLE 6.2. Is $f : V \rightarrow W$ a linear map of vector spaces, then for every $w \in W$, $V_w := f^{-1}(w)$ is not a linear subspace (unless $w = 0$), but is an affine space over $V_0 = \text{Ker}(f)$ (because the difference of two vectors in V_w lies in V_0).

We have obvious notions of affine subspace and affine-linear map: is \mathbb{V} an affine space over V , then a subset $\mathbb{W} \subset \mathbb{V}$ is called an *affine subspace* if it is the orbit of a linear subspace $W \subset V$ (so $\mathbb{W} = O + W$ for any $O \in \mathbb{W}$). A map $F : \mathbb{V} \rightarrow \mathbb{W}$ of affine spaces is called *affine-linear* if there is a linear map of translation spaces $f : V \rightarrow W$ such that $F(v+O) = f(v)+F(O)$ for all $v \in V$ and $O \in \mathbb{V}$ (or equivalently, for which $F(Q) - F(P) = f(Q - P)$ for all $P, Q \in \mathbb{A}$). This linear map is of course unique and is called the *linear part* of F (we might also call it the derivative of F). Notice that if we identify V with \mathbb{V} via $v \mapsto v + O$ and W with \mathbb{W} via $w \mapsto w + F(O)$, then F becomes the linear map $f : V \rightarrow W$. If F is a bijection, then its inverse is also affine-linear: f is bijective, hence its inverse f^{-1} is linear and so if $O' = F(O)$

and $w = f(v)$, then applying F^{-1} to $F(v + O) = f(v) + F(O) = w + O'$ yields $f^{-1}w + F^{-1}(O') = v + O = F^{-1}(w + O')$.

An *automorphism* of \mathbb{V} is of course a bijective affine-linear map $F : \mathbb{V} \rightarrow \mathbb{V}$. These form a group $\text{Aut}(\mathbb{V})$ under composition. This group contains V as the subgroup of translations. If we fix some $O \in \mathbb{A}$, then any automorphism of \mathbb{V} that fixes O is of the form $P \mapsto s(P - O) + O$ for some $s \in \text{GL}(V)$ and thus we find an injective group homomorphism $j_O : \text{GL}(V) \rightarrow \text{Aut}(\mathbb{V})$ whose image is the $\text{Aut}(\mathbb{V})$ -stabilizer of O . If $F \in \text{Aut}(\mathbb{V})$, then $-F(O) + F$ stabilizes O and hence is of the form $j_O(s)$ for some $s \in \text{GL}(V)$. We thus find a bijection $\text{GL}(V) \times V \rightarrow \text{Aut}(\mathbb{V})$, $(v, s) \mapsto v + j_O(s)$ (the latter sends P to $v + s(P - O) + O$). In these terms, composition is:

$$\begin{aligned} (v' + j_O(s')) \circ (v + j_O(s))(P) &= v' + s'(v + s(P - O)) + O = \\ &= v' + s'(v) + s's(P - O) + O = (v + s'v) + j_O(s's)(P). \end{aligned}$$

and hence composition is given by $(v', s') \circ (v, s) = (v' + s'(v), s's)$. Since this deviates from a componentwise product, we refer to this as a *semi-direct product*. The notation $V \rtimes \text{GL}(V)$ is in use to indicate this.

If $K = \mathbb{R}$ and \mathbb{V} is finite dimensional, then we may regard \mathbb{V} as a manifold. Its tangent bundle is just the projection $\mathbb{V} \times V \rightarrow \mathbb{V}$.

EXERCISE 6.1. Let \mathbb{V} be an affine space with translation space V . Prove that for any linear subspace $W \subset V$, the set of orbits \mathbb{V}/W of W in \mathbb{V} has the structure of an affine-linear space over V/W .

EXERCISE 6.2. Let \mathbb{V} be an affine space of finite dimension n with translation space V . Prove that the set $\text{Aff}(\mathbb{V}, K)$ of affine-linear functions $\phi : \mathbb{V} \rightarrow K$ is a vector space of dimension $n + 1$. This space contains as a special element the constant $\mathbf{1} : \mathbb{V} \rightarrow K$. We now have a map $\text{Aff}(\mathbb{V}, K) \times \mathbb{V} \rightarrow K$, $(\phi, P) \mapsto \phi(P)$. Prove that the map from \mathbb{V} to the dual of $\text{Aff}(\mathbb{V}, K)$ that is thus defined is an affine-linear injection and show that the image is the hyperplane in $\text{Aff}(\mathbb{V}, K)^*$ of linear functions $\text{Aff}(\mathbb{V}, K) \rightarrow K$ that take on $\mathbf{1}$ the value 1 .

EXERCISE 6.3. Here \mathbb{V} is an affine space with translation space V .

(a) For $P, Q \in \mathbb{V}$ and $\lambda, \mu \in K$ such that $\lambda + \mu = 1$, we define $\lambda P + \mu Q := P + \mu(Q - P)$. Show that this also equals $Q + \lambda(P - Q)$ and prove that the collection of all such points is an affine subspace with translation space spanned by $Q - P$.

(b) Generalize the preceding to an interpretation of a finite linear combination $\sum_{i=1}^n \lambda_i P_i$ with $P_i \in \mathbb{A}$, $\lambda_i \in K$ and $\sum_i \lambda_i = 1$ as an element of \mathbb{V} . (If $K \supset \mathbb{Q}$, dan heet then $\sum_i \frac{1}{n} P_i$ is called the *barycenter* of $(P_i)_i$.)

(c) Suppose now $K = \mathbb{R}$. If in (b) we let run every λ_i only over $[0, 1]$, then we denote the corresponding set by $[P_1, \dots, P_n]$. We call this the *convex hull* of $(P_i)_i$. Try to show that for every $Q \in \mathbb{V} - [P_1, \dots, P_n]$ there exists an affine-linear function $\phi : \mathbb{V} \rightarrow \mathbb{R}$ with $\phi(P_i) \geq 0$ and $\phi(Q) < 0$ (this is harder than it seems).

Euclidean spaces. A *Euclidean space* is an affine space \mathbb{E} of which the translation space E is a real finite dimensional vector space endowed with an inner product g . Classical Euclidean geometry concerns the case of dimension two, (we then speak of the *Euclidean plane*) and the space we seem to live in is one of dimension three.

Such a space \mathbb{E} is a metric space with distance $d(P, Q) := \sqrt{g(P - Q, P - Q)}$. There is also a notion of angle: if $O, P, Q \in \mathbb{E}$ are such that $P \neq O \neq Q$, then the angle $\angle(OP, OQ) \in [0, \pi]$ is characterized by

$$\cos(\angle(OP, OQ)) = \frac{g(P - O, Q - O)}{d(P, O)d(Q, O)}.$$

An affine-linear transformation in \mathbb{E} whose linear part is orthogonal preserves the distance and vice versa. These make up a subgroup of $\text{Aut}(\mathbb{E})$. The choice of an origin gives this subgroup the structure of a semi-direct product of \mathbb{E} and $O(\mathbb{E})$. An automorphism of \mathbb{E} whose linear part is orthogonal and of determinant one is called a *Euclidean motion*.

Minkowski space. This is a four dimensional real affine space \mathbb{M} whose translation space M (which we prefer to call the *Lorentz space*) is endowed with a pseudo-inner product g of signature $(1, 3)$. We usually also insist that this Lorentz space be oriented and that we have singled out a connected component (the *future cone*) of the set of $v \in M$ with $g(v, v) > 0$. The hypersurface defined by $g(v, v) = 1$ has also two connected components, one of which lies in the future cone. This is the hypersurface of velocity vectors.

The *Lorentz group* $L(M)$ is the group of linear transformations of M that preserve the pseudo-inner product, the orientation (so is of determinant one) and the future cone. We can always choose an oriented basis (e_0, e_1, e_2, e_3) for M such that $g(e_i, e_j)$ takes the diagonal form with $(1, -1, -1, -1)$ on the diagonal. If (t, x^0, x^1, x^3) is the dual basis, then $g(v, v) > 0$ is defined by $t^2 > (x^0)^2 + (x^1)^2 + (x^3)^2$. Then on the future cone either $t > 0$ or $t < 0$. Upon replacing our basis by $(-e_0, -e_1, e_2, e_3)$, we can also arrange that $t > 0$ on the future cone. The *Poincaré group* $P(\mathbb{M})$ is the group of affine-linear transformations of \mathbb{M} whose linear part lies in the Lorentz group. If we are in addition given an origin $O \in \mathbb{M}$, then the chart $\mathbb{M} \rightarrow \mathbb{R}^4, P \mapsto (t(P - O), x^1(P - O), x^2(P - O), x^3(P - O))$ is called in the special theory of relativity an *inertia system*. The Poincaré group acts on the collection of inertia systems and does so transitively. In fact, for any other inertia system $(O'; t', x'^1, \dots, x'^3)$, there is precisely one element of the Poincaré group that takes it to $(O; t, x^1, x^2, x^3)$. A tenet of the special theory of relativity is that physical laws must be invariant under the Poincaré group. This we can now understand in more geometric terms as saying that these laws must pertain to Minkowski space (so that their formulation does not require the choice of an inertia system!).

A particle with positive mass m_o defines curve $\gamma : I \rightarrow \mathbb{M}$ for which $\dot{\gamma}(\tau)$ is a velocity vector for all $\tau \in I$ (this is a way of saying that nothing goes faster than light) and the forces f governing its behaviour (via Newton's law: $f = m_o \ddot{\gamma}$) can be described intrinsically in terms of \mathbb{M} (so without reference to a specific inertia system; this amounts to saying that these laws are independent of the observer). Perhaps more important than the velocity vector is the *energy-momentum vector* $p := m_o \dot{\gamma}$, as it carries more information (it allows us to recover $\dot{\gamma}$ as the unit vector along p). It is clear that Newton's law becomes $\dot{p} = f$.) We illustrate the preceding with the Maxwell equations without bothering too much about their physical interpretation.

The Maxwell equations. These are in vacuum (and in a suitable system of units) as follows:

$$(1) \quad \operatorname{rot}(\mathbf{E}) + \frac{\partial}{\partial t} \mathbf{B} = 0, \quad \operatorname{div}(\mathbf{B}) = 0,$$

$$(2) \quad \operatorname{rot}(\mathbf{B}) - \frac{\partial}{\partial t} \mathbf{E} = \mathbf{I}, \quad \operatorname{div}(\mathbf{E}) = \rho.$$

Here \mathbf{E} , \mathbf{B} and \mathbf{I} are time dependent vector fields on \mathbb{R}^3 equal to the *electrical field*, the *magnetic field* and the *electrical current* respectively, and ρ is a time dependent function on \mathbb{R}^3 , called the *charge density*. We rewrite these formulas in terms of the exterior derivative and the star operator, not in \mathbb{R}^3 , but in Minkowski space \mathbb{M} . In terms of an inertia system $(O; t, x^1, x^2, x^3)$, the star operator acts on typical 1-forms and 2-forms as follows

$$\begin{aligned} \star dt &= dx^1 \wedge dx^2 \wedge dx^3, & \star dt \wedge dx^1 &= -dx^2 \wedge dx^3, \\ \star dx^1 &= dt \wedge dx^2 \wedge dx^3, & \star dx^1 \wedge dx^2 &= dt \wedge dx^3 \end{aligned}$$

We combine electrical and magnetic field in a single 2-form on \mathbb{M} , the *electromagnetic field* (a 2-form on \mathbb{M}):

$$F := B_1 dx^2 \wedge dx^3 + B_2 dx^3 \wedge dx^1 + B_3 dx^1 \wedge dx^2 + \sum_{i=1}^3 E_i (-dt) \wedge dx^i$$

and we combine charge density and current in the *relativistic charge density* (a 3-form on \mathbb{M}):

$$j := \rho dx^1 \wedge dx^2 \wedge dx^3 + I_1 (-dt) \wedge dx^2 \wedge dx^3 + I_2 (-dt) \wedge dx^3 \wedge dx^1 + I_3 (-dt) \wedge dx^1 \wedge dx^2.$$

(We combined the minus sign with dt only in order to make these formulae easier to remember.) Now the equations (1) and (2) come down to

$$(1') \quad dF = 0,$$

$$(2') \quad d \star F = j.$$

This formulation is clearly basis independent; this reflects a requirement of special relativity that the model be observer independent. The force that this system exerts on a charged particle (the *Lorentz force*) can also be expressed that way: a particle with velocity vector v at $p \in \mathbb{M}$ and charge q_0 feels a force $f_L(v)$ (perceived by an acceleration defined by Newton's law: $f_L(v)$ equals its contribution to the particle's acceleration times the particle's rest mass), that is characterized by $f_L(v) = q_0 \tilde{g}^{-1}(t_v F_p)$.

Is $\Omega \subset \mathbb{M}$ a compact 3-dimensional submanifold of \mathbb{M} with boundary $\partial\Omega$, then according to Stokes

$$\int_{\partial\Omega} F = 0, \quad \int_{\partial\Omega} \star F = \int_{\Omega} j.$$

(If Ω is 'space-like' is in the sense that g is negative definite on the tangent spaces $T_p\Omega$, then the right hand side of the second equation can be interpreted as the charge carried by Ω .)

REMARK 6.3. According to the Poincaré lemma we have $F = dA$ for some 1-form A . Such an A is not unique, because for any function f on A , $A + df$ also has that property. A choice for A is called a *potential* for F . In terms of a potential the Maxwell equations reduce to the single equation: $d \star dA = j$. The *Yang-Mills* theory regards the potential A (rather than F) as the primary object.

Maxwell's equations are *not* invariant under Galileo transformations: in the Galileo model two inertial observers who think they are at rest, but have nevertheless constant nonzero velocity with respect to each other, will find different equations. So in hindsight one might say that these equations were discovered prematurely and were waiting for the theory of special relativity to acquire their proper formulation. Indeed, it is their incompatibility with the Galileo model that lead to the development of special relativity.

EXERCISE 6.4. What happens to the Maxwell equations under the transformation $(t, x^1, x^2, x^3) \mapsto (-t, -x^1, -x^2, -x^3)$ which preserves both orientation and g , but interchanges past and future?