Manifolds

Lecture Notes 2011–2016

Klaus Wirthmüller

Preface

These are notes of a course that I have taught at the Technische Universität Kaiserslautern thrice between 2011 and 2016. It is designed for students who have a solid knowlegde of analysis in one and several variables, and are acquainted with the standard notions of point set topology. The course introduces to the basic ideas and tools of differential topology, and includes a modern version of vector analysis.

The need to stay within the limits of a one semester course has not permitted to develop the vector analysis part beyond Stokes' theorem and Lie derivatives, and quite generally the course stays well away from Riemannian geometry. Likewise several common themes of basic differential topology like Sard's and transversality theorems are not treated. Having to select essentially one typical application of vector fields and flows I have opted for Ehresmann's theorem because it beautifully illustrates the technique of controlling the life span of flow lines by the maximality criterion.

In two versions of the course a bit of spare time had been left, which allowed to include some extra material according to the students' interests. This has resulted in the additional Sections 11 and 12, which have the character of two independent appendices.

In order to facilitate self-study of the text each section ends with a collection of exercises which the interested reader will have pleasure to work out.

Contents

Glossary of Symbols

$h\colon X \approx Y$	homeomorphism from X to Y, 1
(U, h)	chart h with domain U, 1
S^n	unit n-sphere, 2
$\widehat{T}$	omission of a term T, 2
U^n, D^n	open and closed unit balls in $\mathbb{R}^n$, 2, 75
$(X, \mathcal{A})$	manifold X with differentiable structure $\mathcal{A}$, 3
$\dim X$	dimension of a manifold X, 4
$\mathrm{rk}_o f$	rank of a differentiable map f at a point o, 5
Df	differential or Jacobian of a differentiable map f, 5
$h\colon X \simeq Y$	diffeomorphism from X to Y, 6
$\mathrm{Mat}(p\times n, \mathbb{R})$	space of real $p\times n$-matrices, 7
x^{t}	transpose of a matrix or vector x, 10
$\mathrm{Sym}(n, \mathbb{R})$	space of symmetric $n\times n$-matrices, 10
$SL(n, \mathbb{R})$	special linear group of real $n\times n$-matrices, 11
$\mathrm{adj}\, a$	adjugate of a matrix a, 11
$\mathbb{R}P^n$	n-dimensional real projective space, 12
(R, θ, φ), (r, φ, z)	spherical and cylinder coordinates on $\mathbb{R}^3$, 14, 29
$O(n)$	orthogonal group, 15
$SO(n)$	special orthogonal group, 15
T_oX	tangent space of X at o, 16
v_h	component of a tangent vector v with respect to a chart h, 16
T_of	differential of a map f at a point o of a manifold, 18
$L\cdot v$, Lv	value of a linear map L at a vector v, preferred over $L(v)$, 18
$\dot{f}(t)$	velocity vector of a curve f at time t, 20
$\mathfrak{sl}(n, \mathbb{R})$	Lie algebra of $SL(n, \mathbb{R})$, 21
TX, Tf	tangent bundle of a manifold X and differential of a map f, 27
$\mathcal{E}_{X,o}$, $\mathcal{E}_o$	algebra of germs of smooth functions at $o \in X$, 30
Ad	adjoint representation of a Lie group, 30
$SU(n)$	special unitary group of complex $n\times n$-matrices, 30
σ_x, σ_y, σ_z	Pauli matrices, 31
E_x	fibre of a bundle E over x, 32
$T \to \mathbb{R}P^n$	tautological bundle, 33
$GL(n, \mathbb{R})$	general linear group, 35
f^*F, f^*h	bundle induced from E and bundle homomorphism induced from h, 36
$F\vert X$	restriction of a bundle F to X, 37
$E \oplus F$, $g \oplus h$	Whitney sum of vector bundles E and F or of bundle homomorphisms, 39
$\mathrm{Hom}(E, F)$, $E \otimes F$	vector bundles constructed from E and F, 39
$E^{\check{}}$	dual of a vector bundle E, 39
S^kE, Λ^kE	symmetric and alternating powers of a bundle, 39
$\mathrm{Sym}\, E$, Alt^kE	bundles of bilinear and alternating forms, 39

ad	adjoint representation of the Lie algebra of a Lie group, 129
exp	exponential map of a Lie group, 131

1 Topological and Differential Manifolds

The real vector spaces $\mathbb{R}^n$ with its standard topology count among the most basic topological spaces. Curiously, topologists usually refer to them as *Euclidean spaces* even though their scalar product hardly ever plays a more than auxiliary role. While Euclidean spaces as such are not very interesting examples of topological spaces they do give rise to an extremely fruitful notion if they are considered not globally but as a local model for more general topological spaces. This is the idea underlying the notion of manifold.

1.1 Definition A topological space X is called an n-dimensional (*topological*) *manifold* or n-manifold for short if the following conditions are satisfied:

- For every point $o \in X$ there exist an open subset $U \subset X$ containing o, and a homeomorphism $h: U \to h(U)$ onto an open subset $h(U) \subset \mathbb{R}^n$;

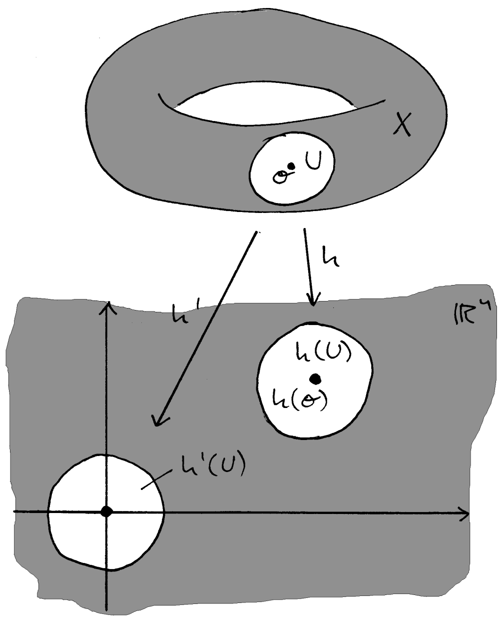

- X is a Hausdorff space, and
- the topology of X admits a countable base.

The homeomorphisms $h: U \approx h(U)$ are called *charts* and often quoted as (U, h) rather than just h in order to make the *chart domain* U explicit. To say that h is a chart *at* the point o simply means $o \in U$, while it is called *centred* at o if furthermore $h(o) = 0 \in \mathbb{R}^n$ — like h' of the figure. Of course any given chart h at o may be centred by composing it with the translation $z \mapsto z - h(o)$ of $\mathbb{R}^n$.

Even if the second and third condition have a mere technical feel they are quite important. While dropping the Hausdorff property would make for an immediate and profound change, at least an elementary part of the theory can be developed without the countability axiom, which will only come into play at a later stage.

1.2 Examples (0) A 0-dimensional manifold is the same as a discrete topological space with at most countably many points.

(1) Every open subspace $X \subset \mathbb{R}^n$ clearly is an n-manifold: the single chart (X, id) works for all $o \in X$.

(2) The sphere $S^n = \{x \in \mathbb{R}^{n+1} \mid |x| = 1\}$ is an n-manifold: the $2n+2$ subsets

$$U_j^- := \{x \in S^n \mid x_j < 0\} \text{ and } U_j^+ := \{x \in S^n \mid x_j > 0\} \quad (j = 0, \ldots, n)$$

form an open cover, and each is the domain of the chart[1]

$$U_j^\pm \ni x \xrightarrow{h_j^\pm} (x_0, \ldots, \widehat{x_j}, \ldots, x_n) \in U^n$$

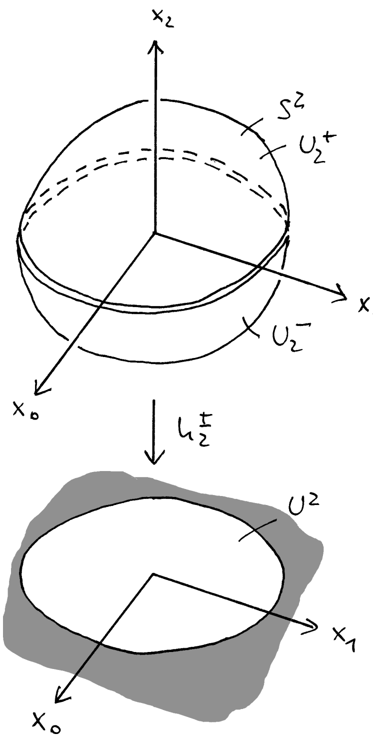

that projects onto the open unit ball $U^n = \{x \in \mathbb{R}^n \mid |x| < 1\}$: the formula

$$\Big(x_0, \ldots, x_{j-1}, \pm\sqrt{1 - \textstyle\sum_{i \neq j} x_i^2},\, x_{j+1}, \ldots, x_n\Big) \longleftarrow (x_0, \ldots, \widehat{x_j}, \ldots, x_n)$$

gives the inverse.

(3) Let (X_1, o_1) and (X_2, o_2) be two copies of the pointed space $(\mathbb{R}^n, 0)$, and let X be the quotient topological space of X_1+X_2 with respect to the gluing relation that identifies each point $x \in X_1 \smallsetminus \{o_1\}$ with the corresponding point $x \in X_2 \smallsetminus \{o_2\}$. Then X contains each of $X_1 \approx \mathbb{R}^n$ and $X_2 \approx \mathbb{R}^n$ as an open subset, thus in particular it is covered by chart domains. On the other hand it fails to be a manifold since the image points of o_1 and o_2 in X cannot be separated by neighbourhoods in X.

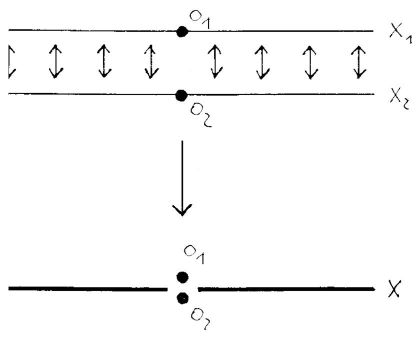

[1] We use the convention that terms covered by a hat are to be omitted.

A few further properties of manifolds follow directly from the definition: they are locally compact and locally path-wise connected, so that the path components of an arbitrary manifold are open subspaces, and the manifold can be recovered as the topological sum of its path components.

1.3 Definition (1) Let (U,h) and (V,k) be charts for a manifold X. The composition

$$h(U\cap V)\xrightarrow{k\circ h^{-1}}k(U\cap V)$$

is called their *transition map*. It clearly is a homeomorphism between open subsets of $\mathbb{R}^n$.

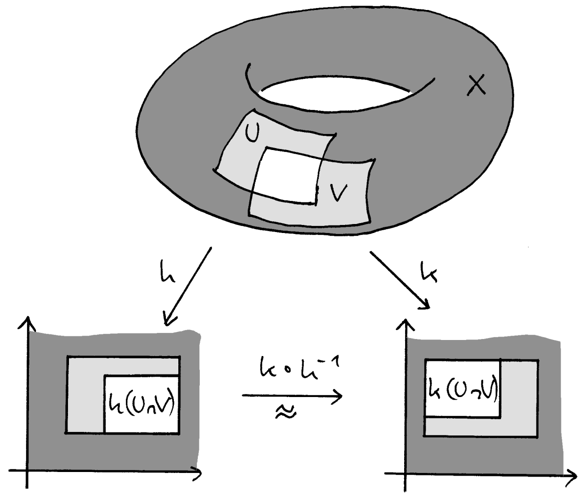

(2) An n-dimensional *atlas* for a manifold X is a set $\mathcal{A}$ of charts (U_h,h) with $h(U_h)\subset\mathbb{R}^n$ such that the chart domains U_h cover X:

$$X=\bigcup_{h\in\mathcal{A}}U_h$$

Note that by Definition 1.1 at least one atlas for X exists.

(3) An atlas $\mathcal{A}$ is called *differentiable* if for any two charts $h,k\in\mathcal{A}$ the transition map $k\circ h^{-1}$ is differentiable (and hence a diffeomorphism: simply swap the roles of h and k).

Note The atlases of our examples (0), (1), and (2) clearly are differentiable.

Differentiable atlases will allow us to perform differential calculus on a manifold. Some care has to be taken when adressing the question of how much structure should be included in the notion of differential manifold that we are now heading for. It is clearly not sufficient to require the mere existence of a differentiable atlas on a given topological manifold, since even the qualitative notions of differential calculus are likely to be depend on the choice of the atlas used. On the other hand fixing a particular differentiable atlas would include too much structure as our Example 1.2(2) with few charts shows: adding all restrictions of these charts to all open subdomains considerably enlarges this differentiable atlas, without really changing the structure.

A viable approach would pass from atlases to equivalence classes of atlases in such a way as to make the two atlases of the example equivalent. But in fact there is a much simpler solution, as follows.

1.4 Definition (1) Let X be a topological manifold. A *differentiable structure* on X is a differentiable atlas for X which is maximal with respect to inclusion (among differentiable atlases).

(2) A *differentiable* or *differential* n-manifold is a pair $(X,\mathcal{A})$ consisting of a topological manifold X and an n-dimensional differentiable structure $\mathcal{A}$ on it (usually dropped from the notation when implied by the context). The charts belonging to $\mathcal{A}$ are called the *differentiable charts* of $(X,\mathcal{A})$.

Note We have built into the definition that the *dimension* $n = \dim X$ of a non-empty differential n-manifold is well-defined. This is not truly restrictive since diffeomorphic transition maps between non-empty open subsets of Euclidean spaces in any case preserve the dimension. While the analogous statement for mere topological manifolds is true it is much harder to prove. You may already have noted that the dimension of a manifold is often included in its symbol as an upper index which does not necessarily indicate a Cartesian power: S^n, U^n.

1.5 Lemma Every differentiable atlas $\mathcal{A}$ is contained in a unique differentiable structure $\overline{\mathcal{A}}$.

Proof In spite of a superficial similarity with statements like Hahn Banach theorems this is not an application of Zorn's lemma but much simpler. Indeed the unique differential structure containing the n-dimensional differential atlas $\mathcal{A}$ is explicitly written down as

$$\overline{\mathcal{A}} := \left\{ (V,k) \;\middle|\; \begin{array}{l} (V,k) \text{ is an } n\text{-dimensional chart of } X\text{, and for} \\ \text{every chart } (U,h) \in \mathcal{A} \text{ the transition map} \\ k\circ h^{-1}\colon h(U\cap V) \to k(U\cap V) \text{ is a diffeomorphism} \end{array} \right\}.$$

Since $\mathcal{A}$ is a differentiable atlas it is clear that $\overline{\mathcal{A}}$ is a set of charts which contains $\mathcal{A}$, in particular that $\overline{\mathcal{A}}$ again is an atlas. To prove that it is differentiable consider two charts (V,k) and (W,l) of $\overline{\mathcal{A}}$, and any point $o \in V\cap W$. We pick a chart $(U,h) \in \mathcal{A}$ at o, and on the neighbourhood $k(U\cap V\cap W)$ of $k(o)$ may write

$$l\circ k^{-1} = (l\circ h^{-1})\circ(h\circ k^{-1}) = (l\circ h^{-1})\circ(k\circ h^{-1})^{-1}$$

as a composition of two differentiable mappings:

$$\begin{tikzcd}
 & U\cap V\cap W \arrow[dl, "k"'] \arrow[d, "h"] \arrow[dr, "l"] & \\
k(U\cap V\cap W) \arrow[r, "h\circ k^{-1}"'] & h(U\cap V\cap W) \arrow[r, "l\circ h^{-1}"'] & l(U\cap V\cap W)
\end{tikzcd}$$

In particular, by the chain rule $l\circ k^{-1}$ is differentiable at the point $k(o)$, and therefore throughout in $k(V\cap W)$ since $o\in V\cap W$ was arbitrary.

We now show that $\overline{\mathcal{A}}$ is maximal. Thus let us try and add to $\overline{\mathcal{A}}$ another chart (V,k) that makes $\overline{\mathcal{A}} \cup \{(V,k)\}$ a differentiable atlas. Then in view of $\mathcal{A}\subset\overline{\mathcal{A}}$ the transition map $k\circ h^{-1}$ must be a diffeomorphism for each $(U,h)\in\mathcal{A}$, and this means that we already had $(V,k)\in\overline{\mathcal{A}}$.

Finally the uniqueness statement is clear: a differentiable atlas containing $\mathcal{A}$ can only be a subset of $\overline{\mathcal{A}}$, thus if it is maximal it must be $\overline{\mathcal{A}}$ itself. □

1.6 Definition Let X and Y be topological manifolds of dimensions n and p, let $f\colon X\to Y$ be a continuous map and $o\in X$ a point. Then for every chart (V,k) at $f(o)$ there exists a chart (U,h) at o with $f(U)\subset V$: indeed candidates for h and k can be chosen arbitrarily, and it then suffices to restrict h to $U\cap f^{-1}V$.

(1) The composition $k\circ f\circ h^{-1}$ is called the *representation* of f in the charts h and k:

$$\begin{tikzcd}
X \arrow[rrr, bend left, "f"] & U \arrow[l, hook'] \arrow[r] \arrow[d, "h"'] & V \arrow[r, hook] \arrow[d, "k"] & Y \\
\mathbb{R}^n & h(U) \arrow[l, hook'] \arrow[r, "k\circ f\circ h^{-1}"] & k(V) \arrow[r, hook] & \mathbb{R}^p
\end{tikzcd}$$

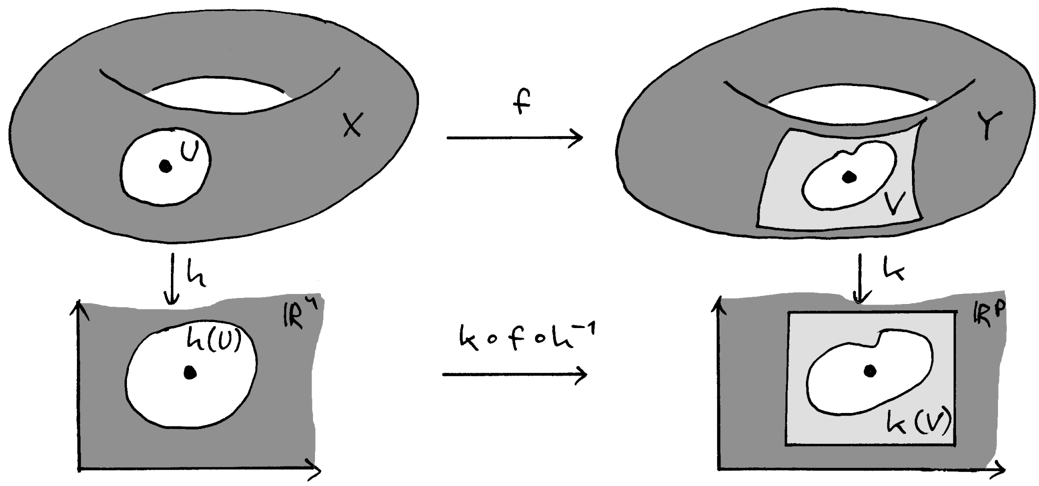

(2) Assume that X and Y are differentiable and that the charts h and k are chosen from the respective differentiable structures. Then f is called *differentiable* at o if $k \circ f \circ h^{-1}$ is differentiable at $h(o)$.

(3) Let $n = \dim X$ and $p = \dim Y$ denote the dimensions. If f is differentiable at o then its *rank* at o is the integer

$$\mathrm{rk}_o f := \mathrm{rk}\, D(k \circ f \circ h^{-1})(h(o)),$$

that is the rank of the linear mapping $D(k \circ f \circ h^{-1})(h(o))\colon \mathbb{R}^n \to \mathbb{R}^p$, or likewise the rank of this Jacobian matrix[2]. If the rank of f equals n at every point $o \in X$ then f is called an *immersion*; if it is equal to p then f is called a *submersion*.

Note Parts (2) and (3) make sense since for different choices of h and k the resulting transition maps are diffeomorphisms. Indeed, let h' and k' be if alternative choices of h and k. We may assume that the domains of h' and k' are U and V like those of h and k, and obtain a commutative diagram

$$\begin{tikzcd}
h(U) \arrow[rrr, "k\circ f\circ h^{-1}"] \arrow[dd, "h'\circ h^{-1}"', "\simeq"] & & & k(V) \arrow[dd, "\simeq"', "k'\circ k^{-1}"] \\
& U \arrow[ul, "h"'] \arrow[r, "f"] \arrow[dl, "h'"] & V \arrow[ur, "k"] \arrow[dr, "k'"'] & \\
h'(U) \arrow[rrr, "k'\circ f\circ (h')^{-1}"'] & & & k'(V)
\end{tikzcd}$$

where the verticals are diffeomorphisms. Assuming differentiability of $k \circ f \circ h^{-1}$ at $h(o)$ we conclude that of $k' \circ f \circ (h')^{-1}$ at $h'(o)$, and the chain rule

$$D(k' \circ f \circ (h')^{-1})(h'(o)) = \underbrace{D(k' \circ k^{-1})(k(o))}_{\in GL(p,\mathbb{R})} \cdot D(k \circ f \circ h^{-1})(h(o)) \cdot \underbrace{D(h \circ (h')^{-1})(h'(o))}_{\in GL(n,\mathbb{R})}$$

proves that the relevant ranks agree. — For a similar reason a qualitative version of the chain rule holds for mappings between manifolds: identity mappings are everywhere differentiable, and if two

[2] We use the same notation to denote a real $p{\times}n$ matrix and the linear mapping $\mathbb{R}^n \to \mathbb{R}^p$ defined by it. On the other hand the action of a differential $Df(a)$ on a vector v will be written as $Df(a) \cdot v$ rather than $Df(a)(v)$ even if we think of $Df(a)$ as a linear mapping.

composable maps f and g are differentiable then so is $g \circ f$, as is read off from the diagram

$$\begin{tikzcd}
U \arrow[r, "f"] \arrow[d, "h", "\simeq"'] & V \arrow[r, "g"] \arrow[d, "k", "\simeq"'] & W \arrow[d, "l", "\simeq"'] \\
h(U) \arrow[r, "k \circ f \circ h^{-1}"'] & k(V) \arrow[r, "l \circ g \circ k^{-1}"'] & l(W)
\end{tikzcd}$$

which represents f and g in charts (U, h), (V, k), and (W, l). To those familiar with the notion of category this means, of course, that we have a category whose objects are the differential manifolds and whose morphisms are the differentiable mappings. In this framework the first part of the following definition just gives a special name to the isomorphisms of that category.

1.7 Definition (1) Let X and Y be differential manifolds. A differentiable map $f: X \to Y$ is a *diffeomorphism* if there exists a differentiable map $g: Y \to X$ such that $g \circ f = \mathrm{id}_X$ and $f \circ g = \mathrm{id}_Y$. We then write $f: X \xrightarrow{\simeq} Y$ rather than $f: X \xrightarrow{\approx} Y$ as we would for a mere homeomorphism.

(2) If $o \in X$ is a point then a differentiable map $f: X \to Y$ is a *local diffeomorphism* at o if there exists an open neighbourhood $U \subset X$ of o such that f sends U diffeomorphically onto an open subset of Y. If this holds for all points $o \in X$ the reference to o will be dropped.

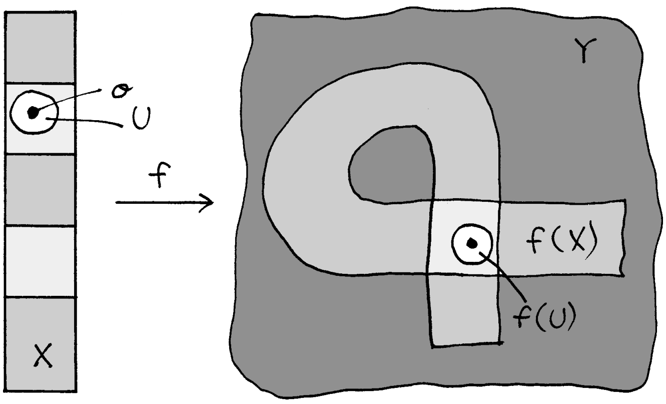

Notes Only now it makes sense to say — and is true — that the charts themselves are diffeomorphisms: the representation of a chart (U, h) in the charts h and $\mathrm{id}_{h(U)}$ is the identity mapping.

In everything said so far the notion of differentiability may be replaced by continuous (C^1) or higher order C^k differentiability including C^∞; in topology the latter is often called *smoothness*[3]. Each of these choices gives rise to its own category of manifolds and differentiable mappings.[4]

More than on any other tool, the local theory of manifolds builds on a single result of differential calculus, the local inverse theorem. For manifolds it takes the following form:

1.8 Local Inverse Theorem Let X, Y be C^k manifolds with $k \geq 1$, let $f: X \to Y$ a C^k mapping, and $o \in X$ a point. Then f is a local C^k diffeomorphism at o if and only if $\dim X = \mathrm{rk}_o f = \dim Y$.

[3] This terminology is not compatible with that of algebraic geometry, where smoothness of a morphism f does not mean a regularity property of f itself but rather non-singularity of its fibres.

[4] It goes without saying that the mere existence of partial derivatives is not a viable substitute for differentiability, in spite of the foolish but widespread habit to call functions with this property partially differentiable.

Proof You must of course be familiar with the coordinate version (if not, do revise it now): there $X \subset \mathbb{R}^n$ and $Y \subset \mathbb{R}^p$ are open subsets, so that the number $\mathrm{rk}_o f$ directly becomes the rank of the Jacobian matrix

$$Df(o) = \begin{pmatrix} \frac{\partial f_1}{\partial x_1}(o) & \cdots & \frac{\partial f_1}{\partial x_n}(o) \\ \vdots & & \vdots \\ \frac{\partial f_p}{\partial x_1}(o) & \cdots & \frac{\partial f_p}{\partial x_n}(o) \end{pmatrix} \in \mathrm{Mat}(p{\times}n, \mathbb{R}).$$

Returning to the general case and assuming $\dim X = \mathrm{rk}_o f = \dim Y$ it suffices to represent f in differentiable charts (U, h) at o and (V, k) at $f(o)$ and apply the coordinate version to $k \circ f \circ h^{-1}$: it will yield a local inverse g of the latter map, and then $h^{-1} \circ g \circ k$ is a local inverse of f:

$$\begin{tikzcd} U \arrow[r, shift left, "f"] \arrow[d, "h"', "\simeq"] & V \arrow[l, dotted, shift left, "h^{-1}\circ g\circ k"] \arrow[d, "k", "\simeq"'] \\ h(U) \arrow[r, shift left, "k\circ f\circ h^{-1}"] & k(V) \arrow[l, dotted, shift left, "g"] \end{tikzcd}$$

(dotted arrows indicate maps that require shrinking V).

The converse statement is but a formality, and plain rather than C^k differentiability sufficient: Choosing centred charts at o and $f(o)$ we reduce to the case of open subsets $X \subset \mathbb{R}^n$ and $Y \subset \mathbb{R}^p$, with $o = 0 \in \mathbb{R}^n$ and $f(o) = 0 \in \mathbb{R}^p$. Then if g is a local inverse to f the chain rule yields

$$Dg(0) \circ Df(0) = D(g \circ f)(0) = D\,\mathrm{id}(0) = \mathrm{id} \quad \text{and} \quad Df(0) \circ Dg(0) = D(f \circ g)(0) = D\,\mathrm{id}(0) = \mathrm{id},$$

so that $Df(0)$ is an invertible linear map and in particular $n = p$. □

Note It follows from the easy direction of the local inverse theorem that diffeomorphic manifolds must have the same dimension.

1.9 Examples and Constructions Each of our examples (0), (1), and (2) comes along with a C^∞ atlas and thus is a smooth manifold.

(3) If X, Y are differential manifolds then the Cartesian products

$$h{\times}k\colon U \times V \longrightarrow h(U) \times k(V)$$

of all differentiable charts (U, h) of X, and (V, k) of Y form a differentiable atlas for the product space $X \times Y$ and thus give it a differential structure, with $\dim X{\times}Y = \dim X + \dim Y$ (if both X, Y are non-empty).

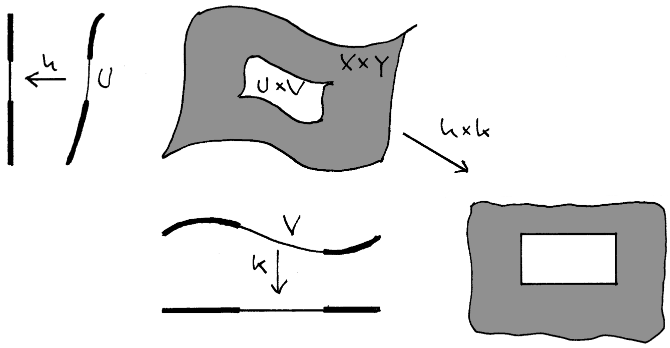

It has the expected universal property in the category of differential manifolds: a differentiable map into $X{\times}Y$ is, essentially, the same as one differentiable mapping into each of X and Y.

(4) Similarly, if X and Y have the same dimension then the topological sum $X + Y$ (disjoint union) again is a differential manifold of that dimension, with the dual universal property.

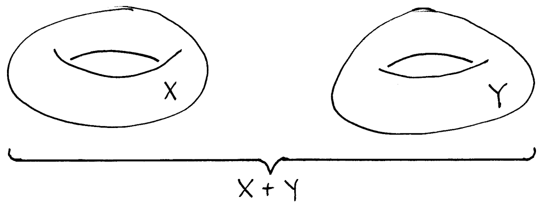

1.10 Definition (1) Let X be a differential $(n+p)$-manifold. A subset $S \subset X$ is called an n-dimensional (or p-codimensional) *submanifold* of X if for each $o \in S$ there exists a differentiable chart (U, h) for X at o such that

$$h(U \cap S) = h(U) \cap \big(\mathbb{R}^n \times \{0\}^p\big).$$

Such charts are often referred to as *submanifold charts* or *flattening charts* for S. Note that they make S a differentiable n-manifold in its own right: S inherits the subspace topology from X, and by definition the submanifold charts h restrict to homeomorphisms

$$S \supset U \cap S \xrightarrow{\approx} h(U) \cap (\mathbb{R}^n \times \{0\}^p) \subset \mathbb{R}^n \times \{0\}^p$$

which produce differentiable transition maps. Of course the inclusion map $S \subset X$ is differentiable, and submanifolds likewise have the expected universal property: differentiable maps into S are the same as differentiable maps into X with all values in S.

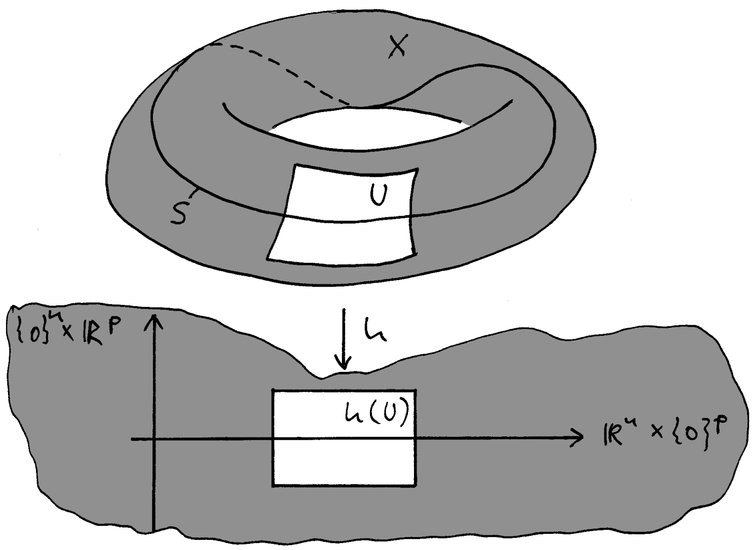

(2) An *embedding* of differential manifolds $e\colon W \to X$ is a map that sends W diffeomorphically onto a submanifold $e(W)$ of X.

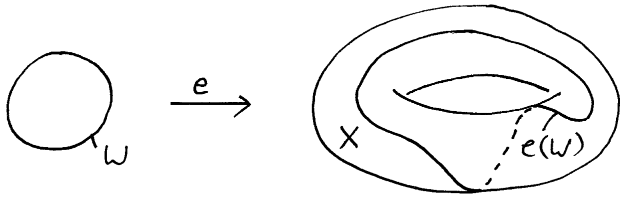

1.11 Examples (0) A 0-codimensional submanifold is the same as an open subset.

(1) The sphere S^n is a smooth submanifold of $\mathbb{R}^{n+1}$: each of the $2n{+}2$ charts discussed before extends to a submanifold chart for $S^n \subset \mathbb{R}^{n+1}$, for instance h_0^+ to

$$(0,\infty) \times U^n \ni (x_0,\ldots,x_n) \longmapsto \Big(x_0 - \sqrt{1-\textstyle\sum_{j>0} x_j^2},\, x_1,\ldots,x_n\Big) \in \mathbb{R} \times U^n.$$

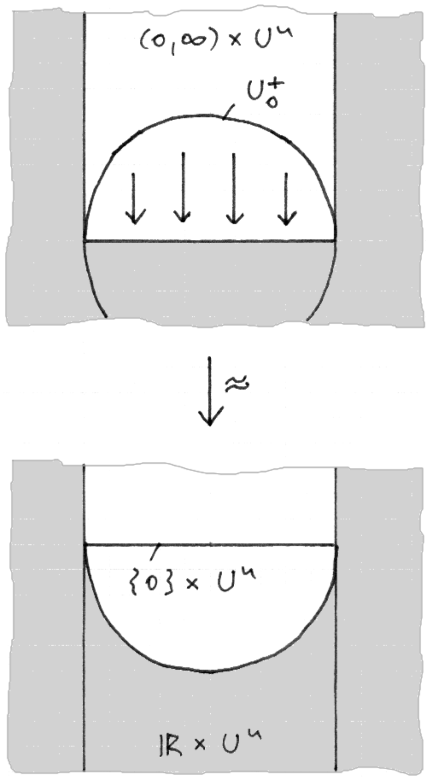

(2) The mapping $\mathbb{R} \ni x \mapsto (x^2, x^3) \in \mathbb{R}^2$ is a topological embedding — inverted by $\sqrt[3]{z} \leftarrow\!\!\shortmid (y,z)$ — as well as a differentiable map, but it fails to be an embedding of differential manifolds. It is safer to avoid the mildly ambiguous term of differentiable embedding in such situations.

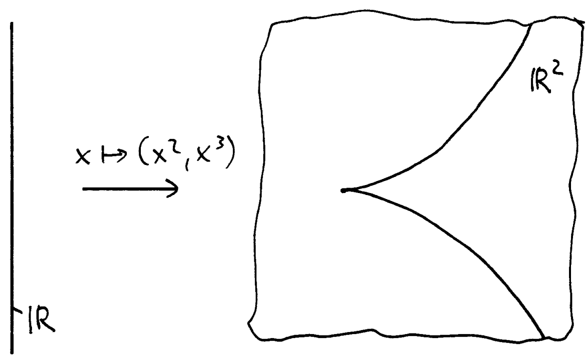

A classical consequence of the local inverse theorem is the implicit function theorem. In the theory of manifolds it is at the heart of the following result, which provides many further examples of differential manifolds.

1.12 Definition Let $f\colon X \to Y$ be differentiable. A point $b \in Y$ is called a *regular value* of f if

$$\operatorname{rk}_a f = \dim Y$$

holds for every $a \in f^{-1}\{b\}$; otherwise b is called a *critical value* of f.

1.13 Regular Value Theorem Let X, Y be C^k manifolds with $k \geq 1$ and $\dim Y = p$, say. Let further $f\colon X \to Y$ be a C^k mapping, and $b \in Y$ a regular value of f. Then the fibre

$$S := f^{-1}\{b\} \subset X$$

is a p-codimensional C^k submanifold of X.

Proof The statement is local in S, and choosing differentiable charts for X at a point $a \in S$, and for Y at $f(a) = b$ we reduce it to the special case where $X \subset \mathbb{R}^{n+p}$ and $Y \subset \mathbb{R}^p$ are open subsets, and $b = 0 \in \mathbb{R}^p$ is the origin: thus $S \subset X$ is the set of solutions of the equation $f(x) = 0$.

By assumption the Jacobian matrix $Df(a) \in \mathrm{Mat}\big(p\times(n+p), \mathbb{R}\big)$ has rank p, and after a suitable permutation of coordinates in $\mathbb{R}^{n+p}$ we may assume that the right hand $p\times p$ submatrix of $Df(a)$ is invertible. Now the classical theorem on implicit functions applies: it provides open balls $U \subset \mathbb{R}^n$ and $V \subset \mathbb{R}^p$ such that $a \in U\times V \subset X$ and such that $S \cap (U\times V)$ is the graph of a C^k function $\varphi\colon U \to V$. Then the map

$$\begin{array}{ccc} U\times V & \xrightarrow{h} & U \times \mathbb{R}^p \\ (x, y) & \longmapsto & \big(x\ ,\ y-\varphi(x)\big) \end{array}$$

is a C^k diffeomorphism onto an open subset of $\mathbb{R}^n \times \mathbb{R}^p$, and in fact a flattening chart for S. $\square$

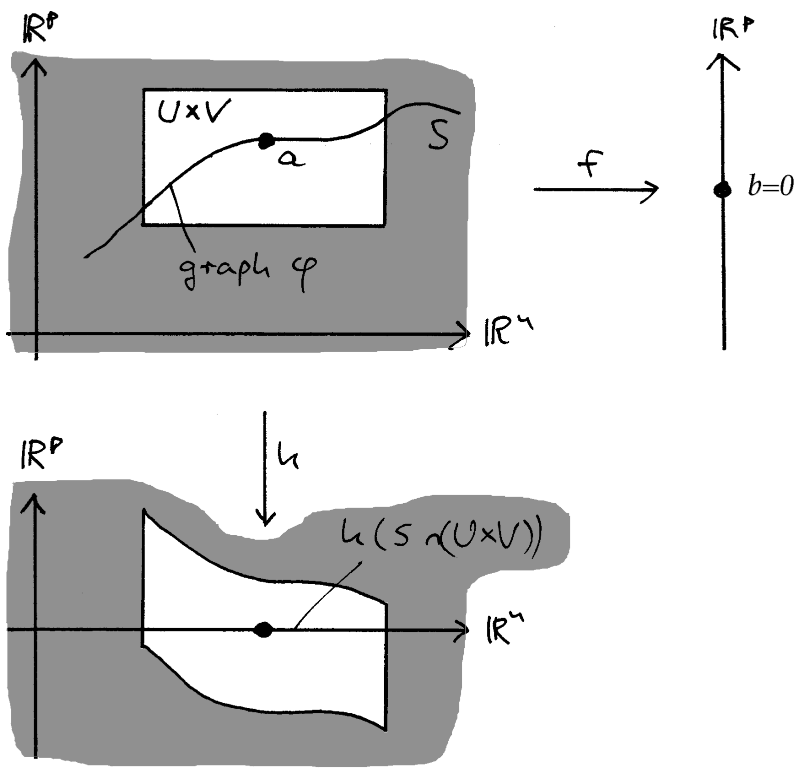

Note If $b \in Y$ is not a value of f at all it is a regular value in the sense of the definition. While this sounds weird it is useful, and quite the intended meaning. In this case the regular value theorem correctly predicts that the empty fibre is an n-manifold for all $n \in \mathbb{N}$.

1.14 Examples (1) Let $f\colon \mathbb{R}^{n+1} \to \mathbb{R}$ be a quadratic form. The number 0 is the only critical value of f, for if f is written out as a product

$$f(x) = x^{\mathrm{t}} s x$$

with a symmetric matrix $s \in \mathrm{Sym}(n+1, \mathbb{R})$ then according to the product rule the differential of f at a sends the vector $x \in \mathbb{R}^{n+1}$ to $x^{\mathrm{t}} s a + a^{\mathrm{t}} s x = 2a^{\mathrm{t}} s x$, so that $Df(a) = 2a^{\mathrm{t}} s$.

In particular the affine quadric $f^{-1}\{1\} \subset \mathbb{R}^{n+1}$ always is a smooth submanifold. This includes the case of the sphere S^n but also other quadrics like hyperboloids.

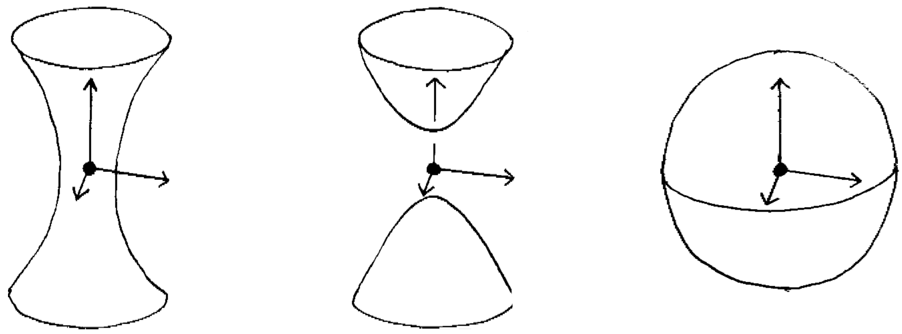

$f(x,y,z) = x^2+y^2-z^2$ $\qquad$ $f(x,y,z) = -x^2-y^2+z^2$ $\qquad$ $f(x,y,z) = x^2+y^2+z^2$

(2) The special linear group $SL(n,\mathbb{R})$ may be written as

$$SL(n,\mathbb{R}) = \bigl\{x \in \mathrm{Mat}(n\times n,\mathbb{R}) \bigm| \det x = 1\bigr\}$$

and is a 1-codimensional smooth submanifold of the space of matrices $\mathrm{Mat}(n\times n,\mathbb{R}) = \mathbb{R}^{n^2}$. To see this we prove the formula

$$D\det(a)\cdot x = \mathrm{tr}\,(\mathrm{adj}\, a\cdot x)$$

which expresses the differential at a of the determinant function $\mathrm{Mat}(n\times n,\mathbb{R}) \xrightarrow{\det} \mathbb{R}$ in terms of the adjugate matrix of a. Indeed the formula is true at $a=1$ since in that case the n^2 partial derivatives are

$$\frac{\partial\det(x)}{\partial x_{ij}} = \delta_{ij}$$

so that $D\det(1)\cdot x = \mathrm{tr}\, x$. Given an arbitrary invertible[5] a we factor the determinant function according to the commutative diagram

$$\begin{tikzcd}
\mathrm{Mat}(n\times n,\mathbb{R}) \arrow[r, "\det"] \arrow[d, "a^{-1}\cdot"'] & \mathbb{R} \arrow[d, "\frac{1}{\det a}\cdot"] \\
\mathrm{Mat}(n\times n,\mathbb{R}) \arrow[r, "\det"] & \mathbb{R}
\end{tikzcd}$$

where the left hand arrow — left multiplication by the constant matrix a^{-1} — sends a to 1. Since both vertical maps are linear the induced diagram of differentials reads

$$\begin{tikzcd}
\mathrm{Mat}(n\times n,\mathbb{R}) \arrow[r, "D\det(a)"] \arrow[d, "a^{-1}\cdot"'] & \mathbb{R} \arrow[d, "\frac{1}{\det a}\cdot"] \\
\mathrm{Mat}(n\times n,\mathbb{R}) \arrow[r, "\mathrm{tr}"] & \mathbb{R},
\end{tikzcd}$$

and in view of $a^{-1}\cdot\det a = \mathrm{adj}\, a$ this confirms the formula in general. For invertible a the adjugate $\mathrm{adj}\, a$, and hence the differential $D\det(a)$ is non-zero, and in particular we conclude that 1 is a regular value of the determinant function. By Theorem 1.13 it follows that $SL(n,\mathbb{R}) \subset \mathrm{Mat}(n\times n,\mathbb{R})$ is a submanifold as stated.

Of course the group operations multiplication and inversion are differentiable: this makes $SL(n,\mathbb{R})$ a so-called *Lie group* — see Section 11.

[5] The formula would extend further to non-invertible $a \in \mathrm{Mat}(n\times n,\mathbb{R})$ by continuity.

(3) Identifying each pair $\{\pm x\} \subset S^n$ of opposite points of the sphere S^n to a single point $[x]$ yields the real projective space $\mathbb{R}P^n$.

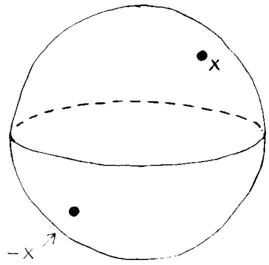

As to its topological properties $\mathbb{R}P^n$ is at once seen to be a compact Hausdorff space with a countable base, and furthermore the $n+1$ charts of the sphere

$$h_j^+ \colon \{x \in S^n \mid x_j > 0\} \longrightarrow U^n \quad (j = 0, \ldots, n)$$

drop to form a smooth atlas. Thus $\mathbb{R}P^n$ is a new example of a smooth manifold, and the two-to-one quotient map $S^n \to \mathbb{R}P^n$ is, essentially by definition, a local diffeomorphism.

(4) The quotient group $\mathbb{R}^n/\mathbb{Z}^n$, in which n-tuples are identified whenever the difference is integral, likewise is a compact Hausdorff space called the n-dimensional *torus*. The quotient homomorphism $q\colon \mathbb{R}^n \to \mathbb{R}^n/\mathbb{Z}^n$ sends sufficiently small open cubes in $\mathbb{R}^n$ homeomorphically to open subsets of the torus, so that inverting these homeomorphisms we obtain an atlas for the torus. If (U, h) and (V, k) are charts from the atlas the transition map restricts to a translation on each connected component of $h(U \cap V)$.

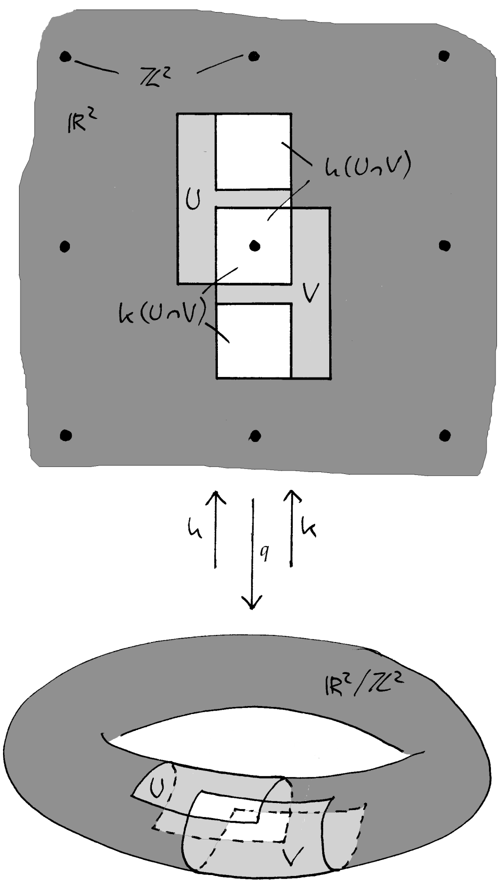

In particular the atlas is smooth: thus the torus is a smooth n-manifold, and once more, the quotient mapping $q\colon \mathbb{R}^n \to \mathbb{R}^n/\mathbb{Z}^n$ a local diffeomorphism.

The torus may be alternatively described as the Cartesian product

$$\mathbb{R}^n/\mathbb{Z}^n = (\mathbb{R}/\mathbb{Z})^n \simeq S^1 \times \cdots \times S^1;$$

this point of view would reduce the previous reasoning to the 1-dimensional case.

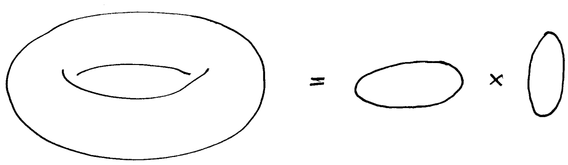

The Story of Differential Structures There are obvious questions about differential structures that we neither have the time nor the means to discuss in detail. They include:

(1) Given a topological manifold, does there exist a differentiable structure on it?

(2) If so, is it unique?

(3) The same questions for C^1, C^k, and C^∞ differential structures, if a differential manifold is given.

Question (2) requires interpretation in order to be interesting. If taken at face value the answer (for positive dimension) will always be no. For instance the homeomorphism $\mathbb{R} \ni x \overset{h}{\longmapsto} x^3 \in \mathbb{R}$ fails to be a diffeomorphism. Choosing $\{h\}$ as an atlas will put a differential structure on the real line that differs from the ordinary one; nevertheless the resulting manifold X still is diffeomorphic to $\mathbb{R}$, just via $h\colon X \to \mathbb{R}$. Both statements are confirmed at once when the set-theoretic identity map and h are represented in the relevant charts (X, h) and $(\mathbb{R}, \mathrm{id})$:

$$\begin{tikzcd} X \arrow[r, "x\mapsto x", "\not\simeq"'] \arrow[d, "h"', "\simeq"] & \mathbb{R} \arrow[d, equal] \\ \mathbb{R} \arrow[r, "h^{-1}"] & \mathbb{R} \end{tikzcd} \qquad \begin{tikzcd} X \arrow[r, "h", "\simeq"'] \arrow[d, "h"', "\simeq"] & \mathbb{R} \arrow[d, equal] \\ \mathbb{R} \arrow[r, "\mathrm{id}"] & \mathbb{R} \end{tikzcd}$$

In order to be interesting, question (2) rather should ask whether on a given topological manifold there may exist several *non-diffeomorphic* differential structures.

Let us turn to (3): I am not aware of any work concerning the difference between C^1 and plain differentiable structures — which does not surprise since the question is hardly exciting. On the other hand it is known that every C^1 manifold admits an essentially unique C^∞ structure. By and large this result justifies the habit of considering C^∞ manifolds the true objects of differentiable manifold theory.

Questions (1) and (2) have turned out to be much more fruitful. By 1950 affirmative answers to both had been found for manifolds of small dimension, not exceeding three. By contrast, in 1960 the Swiss mathematician Michel Kervaire found an example of a topological 7-manifold that does not admit any C^1 differentiable structure. Already in 1956 the American John Milnor had found a C^∞ structure on the sphere S^7 which is not diffeomorphic to the ordinary one, and in 1963 in a joint work both mathematicians succeeded to give a complete and beautiful classification of the differentiable structures on the topological manifold S^7: besides the standard sphere there are exactly 14 such *exotic spheres*, and in a somewhat finer classification including orientations — introduced in 6.4 below — the set of diffeomorphism classes of 7-spheres turns out to carry the structure of a cyclic group of order 28. In any case it was established that the answer to both questions (1) and (2) is negative in principle, and for the next two decades not much attention was payed to the particular question of whether seven was the smallest dimension where these phenomena occur. One reason is the curious fact that among all manifolds those of the small dimensions three and four seem to be the most difficult to investigate.

New information came in the form of by-results from important progress in the investigation of topological and differential 4-manifolds that was made in the 1980s. For instance it was found that most simply-connected compact topological 4-manifold do not admit a C^1 structure — but do so as soon as a single point is removed. On the other side by a spectacular and completely unexpected result the Euclidean space $\mathbb{R}^4$ admits, besides its standard differential structure, infinitely many exotic ones. Nevertheless no systematic theory is available, and by contrast to the situation in higher dimensions it is unknown whether an exotic 4-sphere exists.

Convention For the sake of simplicity and convenience all manifolds and mappings will forthwith be understood to be C^∞ differentiable, and the terms *differentiable* and *smooth* be used as synonymous with C^∞ differentiable. Indeed we will often go as far as to drop these attributes altogether, so that mappings between manifolds are implicitly meant to be differentiable. In particular the term chart will always refer to a differentiable chart.

Exercises

1.1 Describe diffeomorphisms between $\mathbb{R}^n$, the open cube $(-1,1)^n$, and the open n-ball

$$U^n = \left\{x \in \mathbb{R}^n \,\middle|\, |x| < 1\right\}.$$

1.2 Let X be a topological space that obeys the first axiom of Definition 1.1 — note that the notion of atlas introduced in 1.3 already makes sense for such X. Prove that the following are equivalent:

(1) The topology of X admits a countable base.

(2) There exists a countable atlas for X.

(3) Every atlas of X contains an atlas which is countable.

1.3 The assignment

$$(R,\theta,\varphi) \overset{\Phi}{\longmapsto} (R\sin\theta\cos\varphi,\ R\sin\theta\sin\varphi,\ R\cos\theta)$$

defines the spherical coordinate mapping. How can it be used to define differentiable charts of the manifolds $\mathbb{R}^3$ and S^2?

1.4 Prove that the n-sphere admits a differentiable atlas with just two charts. Does there exist an atlas consisting of a single chart?

1.5 Let $S^n = \left\{x = (x_0,\dots,x_n) \in \mathbb{R}^{n+1} \,\middle|\, |x| = 1\right\}$ be the sphere as usual, and denote by $\pi_0\colon S^n \to \mathbb{R}^n$ the projection which suppresses the 0-th coordinate. Determine the points of S^n where π_0 is a local diffeomorphism.

More generally let $\pi_k\colon S^n \to \mathbb{R}^{n-k}$ be the projection which suppresses the coordinates $x_0,\dots,x_k$ and compute the rank $\operatorname{rk}_a \pi_k$ at all points $a \in S^n$.

1.6 Here X is an $(n{+}p)$-manifold. Let $S \subset X$ be an n-dimensional submanifold. Explain why there is an open subset $U \subset X$ which contains S as a *closed* submanifold.

Let $S \subset X$ and $T \subset X$ be two disjoint closed submanifolds of the same dimension n. Explain why $V := S \cup T$ is a submanifold. Show by an example that the assumption of closedness cannot be removed.

1.7 Prove that neither of the following subsets of $\mathbb{R}^3$ is a differentiable submanifold of $\mathbb{R}^3$.

- $R = \big\{(x, y, z) \in \mathbb{R}^3 \,\big|\, x^2 + y^2 = z^2\big\}$
- $S = \big\{(x, y, z) \in \mathbb{R}^3 \,\big|\, \sqrt{x^2 + y^2} = z\big\}$

1.8 Consider the subset $S := \big([0, \infty)\times\{0\}\big) \cup \big(\{0\}\times[0, \infty)\big) \subset \mathbb{R}^2$.

- Prove that S is not a differentiable submanifold of $\mathbb{R}^2$.
- Construct a C^∞-differentiable map $f\colon \mathbb{R} \to \mathbb{R}^2$ which is a topological embedding with $f(\mathbb{R}) = S$.

1.9 Determine all $\lambda \in \mathbb{R}$ for which the set

$$S := \big\{(x, y) \in \mathbb{R}^2 \,\big|\, y^2 = x\,(x-1)\,(x-\lambda)\big\}$$

is a submanifold of $\mathbb{R}^2$.

1.10 Prove that the orthogonal group

$$O(n) = \big\{x \in \mathrm{Mat}(n\times n, \mathbb{R}) \,\big|\, x^{\mathrm{t}}x = 1\big\}$$

is a $\frac{1}{2}n(n-1)$-dimensional submanifold of $\mathrm{Mat}(n\times n, \mathbb{R})$ — it thus is another example of a Lie group.

Explain why the special orthogonal group $SO(n) = \big\{x \in O(n) \,\big|\, \det x = 1\big\}$ is a Lie group of the same dimension. Also explain why both $O(n)$ and $SO(n)$ are compact.

2 Tangent Spaces and Tangent Bundle

What we have said so far about differentiable mappings clearly is unsatisfactory: we do understand the qualitative notion of differentiability but are unable to actually differentiate a map $f\colon X \to Y$ between manifolds at a point $o \in X$, say. Any attempt at differential calculus would require to represent f in charts h at o and k at $f(o)$, but then the differential $Df(o)$ certainly will depend on the largely arbitrary choices of h and k. The purpose of the following construction is to set up a framework for chart-independent differential calculus.

2.1 Definition Let o be a point of the n-manifold X. We consider the (huge) real vector space

$$(\mathbb{R}^n)^{\{\text{all charts at } o\}} = \prod_h \mathbb{R}^n,$$

where we prefer the latter notation, it being understood that the Cartesian product is taken over all charts h at o. The subspace

$$T_oX := \Big\{ v \in \prod_h \mathbb{R}^n \;\Big|\; D(k \circ h^{-1})\big(h(o)\big) \cdot v_h = v_k \text{ for all } h \text{ and } k \Big\}$$

is the *tangent space* of X at o.

Explanation The idea behind this quite abstract definition comes from physics. Asked what a (tangent) vector in three-dimensional space is, a physicist would, quite characteristically, avoid a direct answer but rather say that a vector is described by three independent real numbers which transform linearly with the differential of any coordinate change. To a mathematician this is not directly acceptable as a definition since it does not tell you what a tangent vector is as an object; nevertheless the idea can be translated.

By a coordinate change physicists of course mean what we call the transition map $k \circ h^{-1}$ from one chart h at o to another, k. Assume for the moment that one chart h has been fixed. Then our physicist's definition plainly states that a vector is a triple $v_h \in \mathbb{R}^3$ in the 3-dimensional standard Euclidean space. Now if we want to free ourselves from the particular choice of h we must read a vector as a function v that assigns to *every* chart h at o a triple $v_h \in \mathbb{R}^3$. The statement that upon passing from h to k this triple transforms in a particular way means that the assignment $h \mapsto v_h$ cannot be arbitrary but that the values v_h and v_k mutually determine each other; specifically they must obey the rule

$$D(k \circ h^{-1})\big(h(o)\big) \cdot v_h = v_k.$$

Writing this out we arrive at the definition of T_oX, since an element of the Cartesian product $\prod_h \mathbb{R}^n$ after all is nothing but a function assigning to each chart h an element of $\mathbb{R}^n$.

Of course since (for positive n) there are many charts at o — given one we already obtain a vast collection of them by mere restriction to open subsets — the vector space $\prod_h \mathbb{R}^n$ is frighteningly large. But the formula defining the tangent space only allows the value v_h for one particular chart to be specified freely, and then determines v_k for all further charts k:

2.2 Lemma Let X be an n-manifold and $o \in X$. For every chart h at o the projection

$$T_oX \ni v \longmapsto v_h \in \mathbb{R}^n$$

is a linear isomorphism. In particular $\dim T_oX = n = \dim X$.

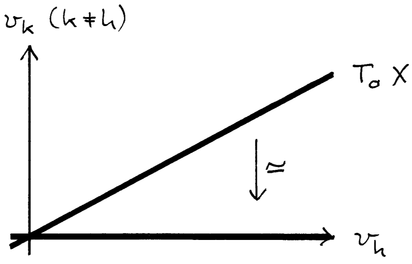

Proof If $v \in T_o X$ projects to zero, that is if $v_h = 0$ then

$$v_k = D(k \circ h^{-1})\big(h(o)\big) \cdot v_h = 0 \quad \text{for all } k,$$

which means that $v = 0$: therefore the projection is injective. On the other hand given an arbitrary $v_h \in \mathbb{R}^n$ we define $v \in \prod_k \mathbb{R}^n$ putting

$$v_k = D(k \circ h^{-1})\big(h(o)\big) \cdot v_h$$

for every chart k at o. This does not change the value v_h, and if g and k are any two charts at o then the identity[1]

$$D(k \circ g^{-1}) \cdot v_g = D(k \circ h^{-1}) \cdot D(g \circ h^{-1})^{-1} \cdot v_g = D(k \circ h^{-1}) \cdot v_h = v_k$$

shows that $v \in T_o X$. Therefore the projection is surjective too. □

Thus while the huge Cartesian product $\prod_h \mathbb{R}^n$ is essential in order to give a precise meaning to the concept of tangent vector, working and calculating with them may, and usually is, done in terms of an — abstractly or concretely specified — particular chart, and thus fits in the framework of finite dimensional linear algebra.

If after all the definition of $T_o X$ is conceptually simple it is not very intuitive in general, and I feel that the best justification eventually lies in the fact that the $T_o X$ thus defined has all the formal properties that one would expect the tangent space at o to have: many of them we will work out below.

It is worth pointing out that the special case of an open subset $X \subset \mathbb{R}^n$ is at once understood. The identity mapping of X is a preferred chart for X, and for each point $a \in X$ the corresponding projection $v \mapsto v_{\mathrm{id}}$ identifies $T_a X$ with $\mathbb{R}^n$. This coincides with the intuitive notion of a tangent vector at a: an observer placed at a may look within X into any direction of $\mathbb{R}^n$. Here you may prefer to think of the origin of $T_a X$ as shifted to the point a so that tangent vectors will be "based" at a rather than the origin. While this latter point of view does help intuition it would be technically awkward.

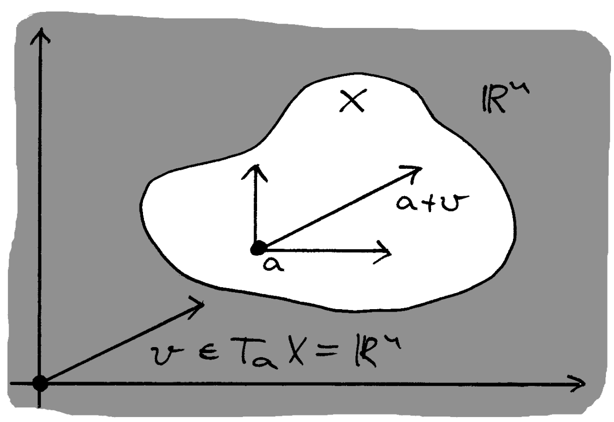

[1] We omit the points where the differentials are taken when they are clear form the context. Neither is it necessary to specify the exact domains of the functions to be differentiated as long as it is clear that they are defined in some neighbourhood of the relevant point.

2.3 Lemma and Definition Let $f: X \to Y$ be a differentiable map between manifolds, and $a \in X$ a point with image $b := f(a) \in Y$. Then f induces a linear mapping

$$T_a f: T_a X \longrightarrow T_b Y$$

which is called the *differential* of f at a, and obeys the chain rule:

$$T_a \operatorname{id}_X = \operatorname{id}_{T_a X} \quad \text{and} \quad T_a(g \circ f) = T_b g \circ T_a f$$

The differential is defined as follows. Using that f is continuous represent f in suitable charts h at a and k at b.

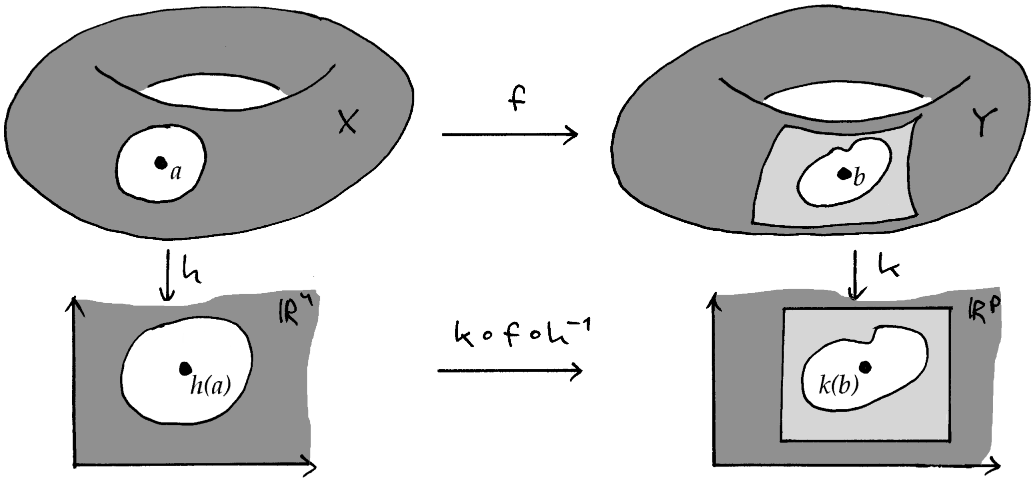

Now if $v \in T_a X$ is a tangent vector then the component of $T_a f(v)$ labelled by the chart k is

$$T_a f(v)_k := D(k \circ f \circ h^{-1})(h(a)) \cdot v_h \in \mathbb{R}^p.$$

Proof By Lemma 2.2 the defining formula determines a unique tangent vector $T_a f(v) \in T_b Y$. We show that it does not depend on the choice of h and k: if k is replaced by another chart l then

$$D(l \circ k^{-1}) \cdot T_a f(v)_k = D(l \circ k^{-1}) \cdot D(k \circ f \circ h^{-1}) \cdot v_h = D(l \circ f \circ h^{-1}) \cdot v_h = T_a f(v)_l\,,$$

and if g is another choice for h the identity

$$D(k \circ f \circ g^{-1}) \cdot v_g = D(k \circ f \circ h^{-1}) \cdot D(h \circ g^{-1}) \cdot v_g = D(k \circ f \circ h^{-1}) \cdot v_h$$

proves that we still get the same vector $T_a f(v) \in T_b Y$. Therefore $T_a f$ is well-defined.

Turning to the chain rule it is obvious that $T_a \operatorname{id} = \operatorname{id}_{T_a X}$, and for composable differentiable mappings f and g we calculate[2]

$$\begin{aligned}
\big((T_{f(a)} g \circ T_a f) \cdot v\big)_l &= D(l \circ g \circ k^{-1}) \cdot (T_a f \cdot v)_k \\
&= D(l \circ g \circ k^{-1})\, D(k \circ f \circ h^{-1}) \cdot v_h \\
&= D(l \circ g \circ f \circ h^{-1}) \cdot v_h \\
&= \big(T_a(g \circ f) \cdot v\big)_l
\end{aligned}$$

[2] In order to reduce the number of brackets we will henceforth usually indicate evaluation of a linear mapping, even an abstract one, by a central dot or mere juxtaposition.

using charts h, k, l at the points a, $f(a)$, and $(g\circ f)(a)$ respectively. □

Note It is clear that $\mathrm{rk}_a f$ is the rank of $T_a f$; in particular a more adequate way to express the rank hypothesis in the local inverse theorem 1.8 is to require that the differential $T_a f$ is a linear isomorphism.

2.4 Examples (1) Let $X \subset \mathbb{R}^n$ and $Y \subset \mathbb{R}^p$ be open subsets. Their tangent spaces reduce to $T_a X = \mathbb{R}^n$ and $T_b Y = \mathbb{R}^p$, and the differential to the Jacobian matrix $T_a f = Df(a) \in \mathrm{Mat}(p\times n, \mathbb{R})$.

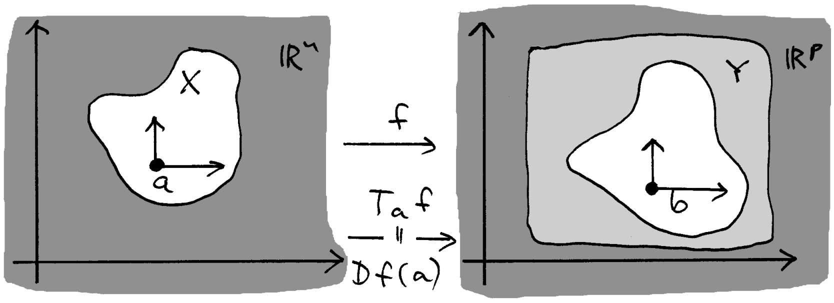

(2) Let X be a manifold of dimension $n+p$, and $S \subset X$ an n-dimensional submanifold. If (U,h) is a submanifold chart at the point $a \in S$ then, as we know, the restriction of h

$$S\cap U \xrightarrow{h'} h(S\cap U) \subset \mathbb{R}^n \times \{0\} = \mathbb{R}^n$$

is a chart for S. It yields a canonical identification of the tangent space $T_a S$ with a subspace of $T_a X$, sending $u \in T_a S$ to the tangent vector $v \in T_a X$ with $v_h = (u_{h'}, 0) \in \mathbb{R}^n \times \{0\} \subset \mathbb{R}^n \times \mathbb{R}^p$. More formally this identification is just the differential $T_a e$ of the inclusion mapping $e\colon S \subset X$, which of course is an embedding of differential manifolds.

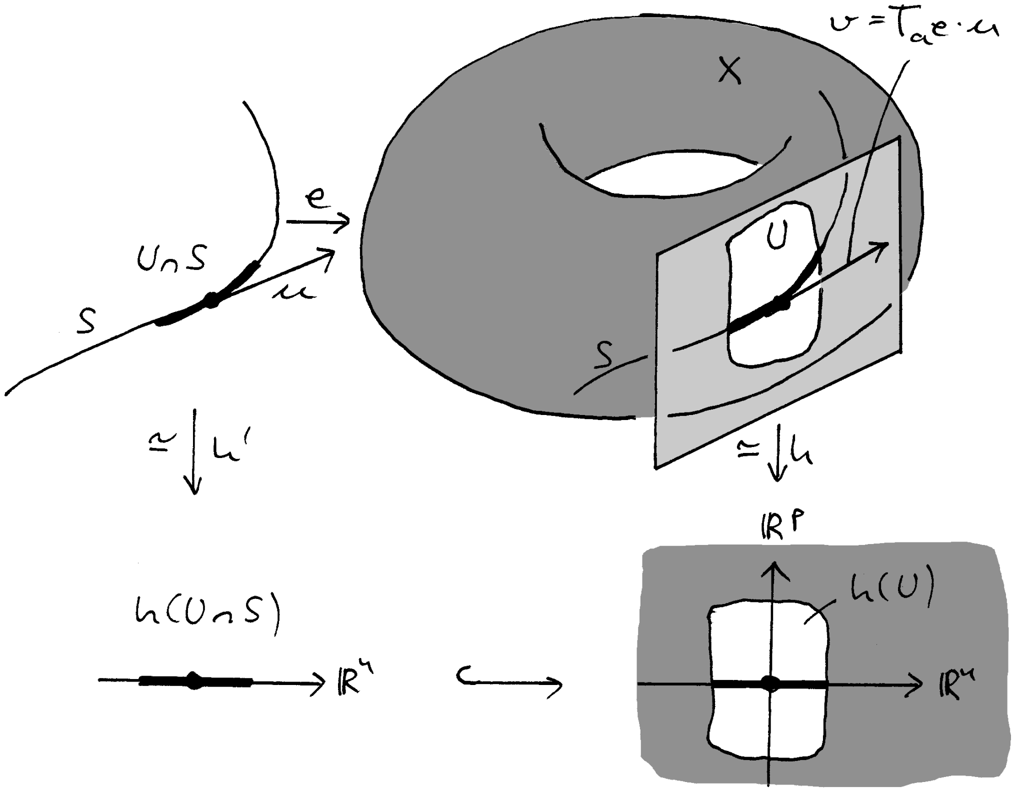

Note that by the chain rule we have

$$T_a(f|S) = T_a(f \circ e) = T_a f \circ T_a e = T_a f \,|\, T_a S$$

for differentiable maps $f: X \to Y$ into a further manifold Y.

Two special cases are worth singling out. One is that of an open subset $S \subset X$, that is $p = 0$: then the tangent spaces of S and X are just identified. The other is that of X an open subset of $\mathbb{R}^{n+p}$: in this case the identification of T_aX with $\mathbb{R}^{n+p}$ makes T_aS a subspace of $\mathbb{R}^{n+p}$. The figure

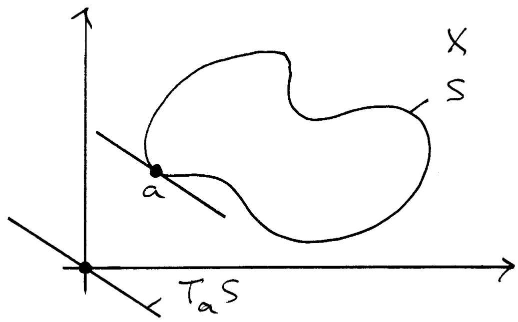

indicates this and includes the intuitive version of T_aS.

(3) Mappings $f: T \to X$ from an open interval T into a manifold X are known as (*parametrised*) *curves* on X. Together with any scalar function $g: X \to \mathbb{R}$ we obtain the differentiable function $g \circ f: T \to \mathbb{R}$ of one variable. Its differential is

$$T_t(g \circ f) = T_{f(t)} g \circ T_t f$$

or, in more classical notation,

$$(g \circ f)'(t) = T_{f(t)} g \cdot \dot{f}(t)$$

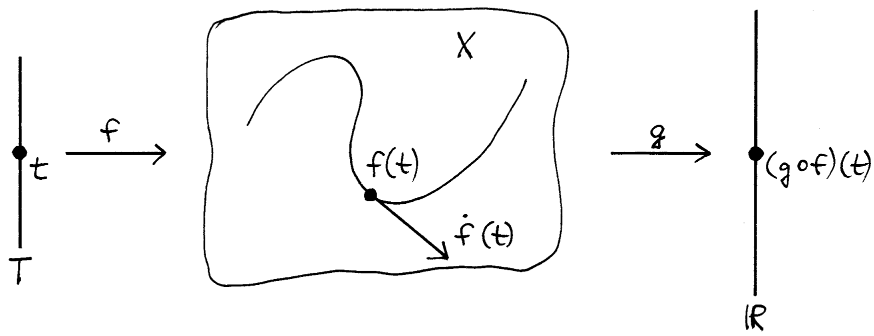

for all $t \in T$. Here $\dot{f}(t) = T_t(f) \cdot 1 \in T_{f(t)}X$ is the *velocity vector* of the curve f at the time t. Given a tangent vector $v \in T_bX$ at a point $b \in X$ the number $T_bg \cdot v$ is sometimes called the derivative of g *along* v: thus $(g \circ f)'(t)$ is the derivative of g along the vector $\dot{f}(t) \in T_{f(t)}X$.

Note that the notion of gradient plays no role in this picture. Indeed the differential of g is a linear form on $T_{f(t)}X$, and in the absence of a Euclidean structure on this space it is neither possible nor desirable to turn this linear form into a tangent vector. For the same reason we have refrained from calling $T_bg \cdot v$ a directional derivative, a terminology usually reserved to the case of a unitary tangent vector $v \in T_bX$.

(4) If $f: X \to Y$ and $b \in Y$ are as in the regular value theorem 1.13 then for any $a \in f^{-1}\{b\}$ the tangent space to the fibre is

$$T_a f^{-1}\{b\} = \operatorname{kernel} T_a f \subset T_a X.$$

Indeed the restriction $f|f^{-1}\{b\}$ is constant by definition, so that its differential at a vanishes: thus $T_af^{-1}\{b\} \subset \operatorname{kernel} T_af$ as observed in (2). On the other hand both $T_af^{-1}\{b\}$ and $\operatorname{kernel} T_af$ have the same dimension $\dim X - \dim Y$, so they must coincide.

In the special case of the map $f\colon \mathbb{R}^{n+1} \to \mathbb{R}$ with $f(x) = |x|^2$ and $b = 1$ we obtain the tangent space to the sphere

$$T_aS^n = T_af^{-1}\{1\} = \big\{x \in \mathbb{R}^{n+1} \,\big|\, 2a^{\mathrm{t}}x = 0\big\} = \{a\}^\perp \subset \mathbb{R}^{n+1},$$

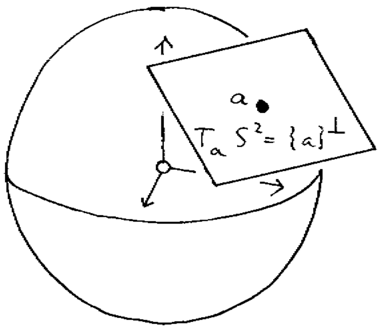

while in that of Example 1.14(2) — the determinant function on $\mathrm{Mat}(n\times n, \mathbb{R})$ — the tangent space at the unit matrix

$$T_1SL(n, \mathbb{R}) = \big\{x \in \mathrm{Mat}(n\times n, \mathbb{R}) \,\big|\, \operatorname{tr} x = 0\big\}$$

— the so-called *Lie algebra* $\mathfrak{sl}(n, \mathbb{R})$ of the Lie group $SL(n, \mathbb{R})$ — is the space of traceless matrices.

We turn to an important and beautiful extension of the local inverse theorem. It provides a local normal form for smooth mappings — albeit under a hypothesis which is quite restrictive.

2.5 Constant Rank Theorem Let X an n-dimensional and Y a p-dimensional manifold, and let $f\colon X \to Y$ be a smooth map of constant rank r: that is, $\mathrm{rk}_o f = r$ for all $o \in X$. Then for every $o \in X$ there exist centred charts (U, h) at o and (V, k) at $f(o) \in Y$ such that $f(U) \subset V$ and f expressed in these charts takes the normal form

$$h(U) \ni (x_1, \ldots, x_n) \xmapsto{k\circ f\circ h^{-1}} (x_1, \ldots, x_r, 0, \ldots, 0) \in k(V).$$

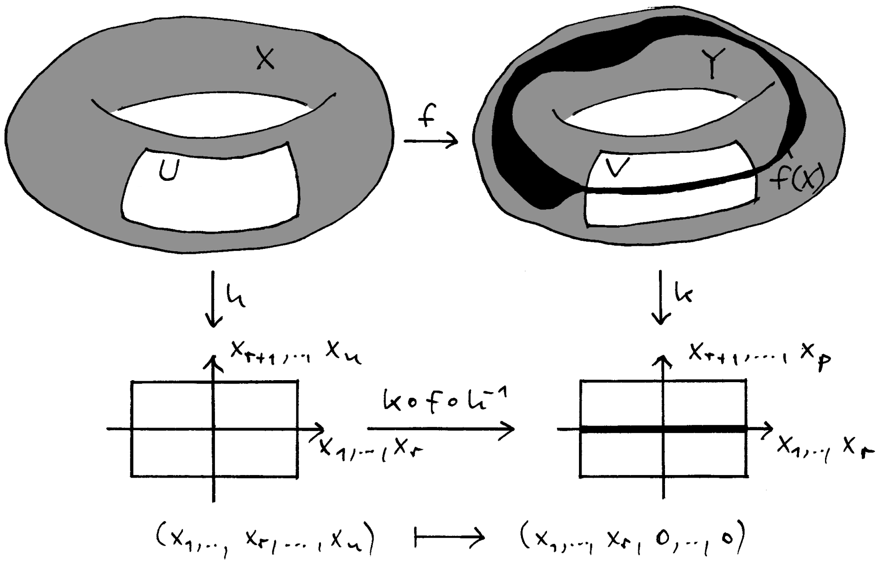

This conclusion may be strengthened in two special cases. In that of $r = p$ — surjective differential at o — the chart k may be prescribed arbitrarily, and in the case $r = n$, that is, injective T_of, the chart h may be obtained from an arbitrarily given one just by restriction.

Notes The conclusion of 2.5 being local, the rank of f need only be constant in some neighbourhood of o. Conversely if f admits such a local normal form its rank clearly is r in some neighbourhood of o. — Important special cases of maps with constant rank are immersions and submersions, which globally have ranks n and p by definition.

If $\mathrm{rk}_o\, f = r$ is maximal in the sense of $r = n$ or $r = p$ then $\mathrm{rk}_x\, f = r$ holds automatically for all x in some neighbourhood of o: working in charts the differential of f at o becomes a $p{\times}n$ matrix which contains an invertible $r{\times}r$ submatrix, and by continuity of the Jacobian determinant this submatrix remains invertible at all points sufficiently close to o.

It is interesting to compare 2.5 with the well-known theorem of linear algebra which states that every linear map of rank r takes the form $(x_1, \ldots, x_n) \mapsto (x_1, \ldots, x_r, 0, \ldots, 0)$ if suitable bases are chosen. Of course this is way easier to prove.

Proof At first we represent f in an arbitrary pair of centred charts at o and $f(o)$, and thus may assume that $X \subset \mathbb{R}^n$ and $Y \subset \mathbb{R}^p$ are open subsets while $o = 0$ and $f(o) = 0$. By assumption the Jacobian matrix $Df(0) \in \mathrm{Mat}(p{\times}n, \mathbb{R})$ has rank r, and suitably permuting the coordinates we further achieve that the submatrix

$$\left(\frac{\partial f_i}{\partial x_j}(0)\right)_{i,j=1}^{r} \in \mathrm{Mat}(r{\times}r; \mathbb{R})$$

is invertible. The smooth mapping

$$X \ni x \overset{h}{\longmapsto} \big(f_1(x), \ldots, f_r(x),\, x_{r+1}, \ldots, x_n\big) \in \mathbb{R}^n$$

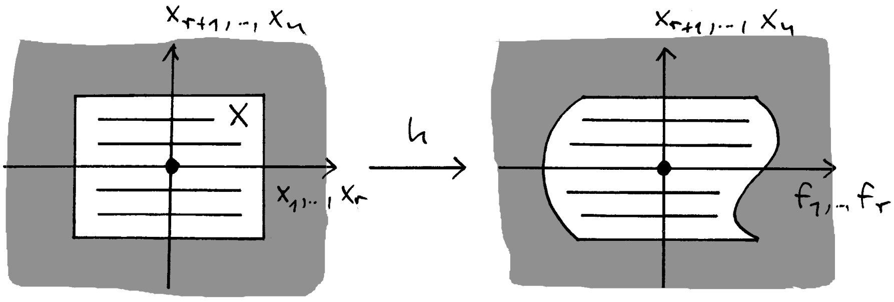

has the Jacobian

$$Dh = \left(\begin{array}{c|c} \frac{\partial f_i}{\partial x_j} & \frac{\partial f_i}{\partial x_j} \\ \hline 0 & 1 \end{array}\right)$$
$$\qquad\qquad \underbrace{}_{r}\ \underbrace{}_{n-r}$$

of format $n{\times}n$, which clearly is invertible at the point $0 \in X$. The local inverse theorem guarantees that we may shrink X so that the restricted mapping $h\colon X \to h(X)$ becomes a diffeomorphism and thereby a chart for X. Note that it is centred at $o = 0$ since $f(0) = 0$. This essentially completes the construction of the chart h.

We now represent the given map f in the chart h and the identical chart of Y, that is, we form the composition $g := f \circ h^{-1}\colon h(X) \to Y$. The commutative diagram

$$\begin{tikzcd}
(x_1, \ldots, x_n) \arrow[rr, "f"] \arrow[d, "h"', "\simeq"] & & \big(f_1(x), \ldots, f_p(x)\big) \\
\big(f_1(x), \ldots, f_r(x), x_{r+1}, \ldots, x_n\big) \arrow[urr, "g"'] & &
\end{tikzcd}$$

shows that g has the form

$$y = (y_1, \ldots, y_n) \longmapsto \big(y_1, \ldots, y_r, g_{r+1}(y), \ldots, g_p(y)\big).$$

Since g represents f in a chart the Jacobian of g

$$Dg = \left(\begin{array}{c|c} 1 & 0 \\ \hline * & \frac{\partial g_i}{\partial y_j} \end{array}\right),$$
$$\qquad\underbrace{\quad}_{r}\;\underbrace{\quad}_{n-r}$$

must have the same rank as f, that is rank r everywhere: this means that for all $i > r$ and $j > r$ the entries $\frac{\partial g_i}{\partial y_j}$ vanish identically. Shrinking X further we may assume that the open set $h(X) \subset \mathbb{R}^n$ is is a product $W \times Z$ of open balls $W \subset \mathbb{R}^r$ and $Z \subset \mathbb{R}^{n-r}$, and using convexity of the latter we conclude that for $i > r$ the value of g_i does not depend on $y_{r+1}, \ldots, y_n$:

$$g_i(y) = g_i(y_1, \ldots, y_r) \quad \text{for each } i > r$$

with differentiable functions $g_i: W \to \mathbb{R}$.

We now define the chart k. Its domain is $V := Y \cap (W \times \mathbb{R}^{p-r})$, and abbreviating $(y_1, \ldots, y_p)$ as $(y', y_{r+1}, \ldots, y_p)$ we put

$$k(y_1, \ldots, y_p) = \big(y_1, \ldots, y_r,\, y_{r+1} - g_{r+1}(y'),\, \ldots,\, y_p - g_p(y')\big).$$

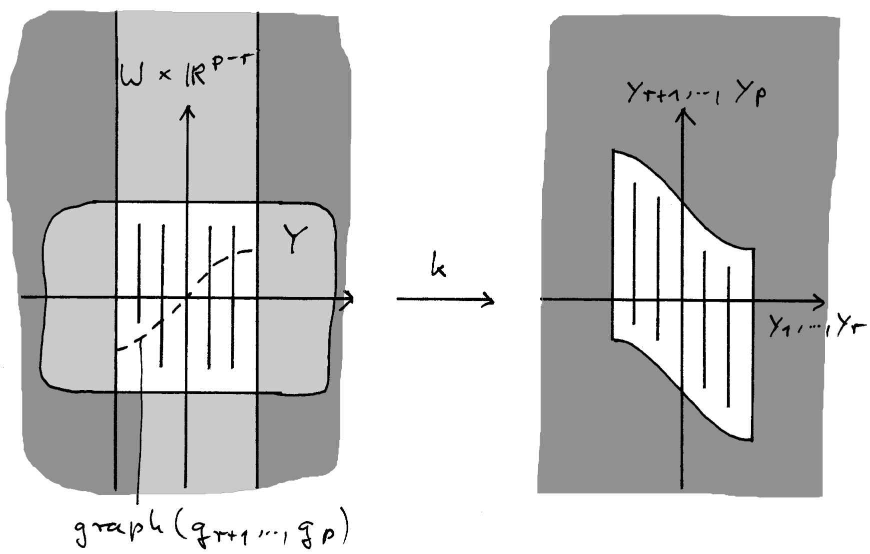

The inversion formula $(y_1, \ldots, y_p) \mapsto \big(y_1, \ldots, y_r,\, y_{r+1} + g_{r+1}(y'),\, \ldots,\, y_p + g_p(y')\big)$ shows that k is a diffeomorphism onto its image in $\mathbb{R}^p$, and as the g_i vanish at the origin k is a centred chart at $f(o) = 0 \in Y$. We finally restrict the chart $h: X \to h(X)$ to $U := X \cap f^{-1}V$ to ensure that $k \circ f \circ h^{-1} = k \circ g$ is defined. According to the diagram

$$\begin{tikzcd} & \big(y_1, \ldots, y_r, g_{r+1}(y'), \ldots, g_p(y')\big) \arrow[d, "k", "\simeq"'] \\ (y_1, \ldots, y_n) \arrow[ur, "g"] \arrow[r, "k\circ f\circ h^{-1}"'] & (y_1, \ldots, y_r, 0, \ldots, 0) \end{tikzcd}$$

the composition $k \circ f \circ h^{-1}$ now indeed acts as stated in the theorem.

As to the special cases we observe that for $r = p$ the k constructed in the last part of the proof is the identity: this means that the original choice of k, which was arbitrary, is preserved. — At the other extreme $r = n$, the constructed chart h does differ from the identity, but its effect is neutralised if the chart k is replaced by its composition with $h^{-1} \times \mathrm{id}_{\mathbb{R}^{p-n}}$:

$$\begin{tikzcd}
x \arrow[r, "f"] \arrow[d, "h"'] & f(x) \arrow[d, "k"] \\
y' \arrow[r, "k\circ f\circ h^{-1}"] \arrow[d, "h^{-1}"'] & (y', 0, \ldots, 0) \arrow[d, "h^{-1}\times\mathrm{id}_{\mathbb{R}^{p-n}}"] \\
x \arrow[r] & (x, 0, \ldots, 0)
\end{tikzcd}$$

This completes the proof of the constant rank theorem. □

Remark The regular value theorem 1.13 may be seen as a trivial application of the constant rank theorem: everything being local only a map f in normal form $(x_1, \ldots, x_{n+p}) \mapsto (x_1, \ldots, x_p)$ need be considered.

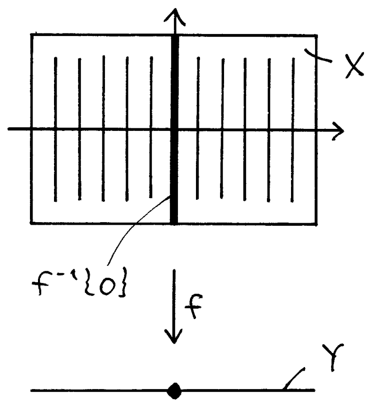

A similar application of the constant rank theorem in the dual situation of injective differential leads to a useful method to recognise embeddings of differential manifolds.

2.6 Theorem Let X and Y be manifolds. A differentiable map $e\colon X \to Y$ is an embedding of differential manifolds if and only if it is both a topological embedding and an immersion.

Proof Only one direction requires proof: We assume that e is an immersive topological embedding and will prove that it is a differentiable embedding. Let $\dim X = n$ and $\dim Y = n+p$, and let $o \in X$ be arbitrary. By the constant rank theorem we find centred charts (U, h) at o and (V, k) at $e(o) \in Y$ in which e is represented by

$$\mathbb{R}^n \supset h(U) \ni x \xmapsto{k\circ e\circ h^{-1}} (x, 0) \in k(V) \subset \mathbb{R}^n \times \mathbb{R}^p.$$

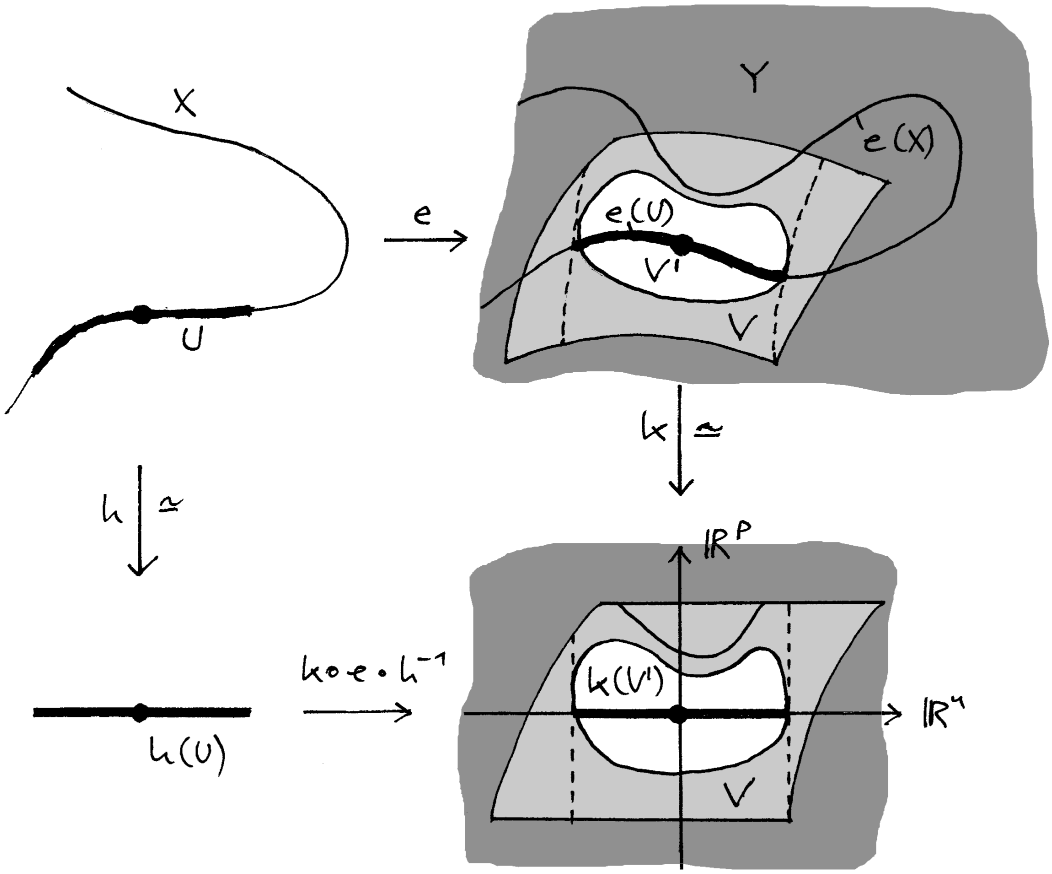

By the special form of $k \circ e \circ h^{-1}$ we have

$$k\big(e(U)\big) = (k \circ e \circ h^{-1})\big(h(U)\big) = h(U) \times \{0\} \subset h(U) \times \mathbb{R}^p$$

and therefore $e(U) \subset V \cap k^{-1}\big(h(U) \times \mathbb{R}^p\big)$. We thus may shrink the chart domain V to the smaller open neighbourhood $V \cap k^{-1}\big(h(U) \times \mathbb{R}^p\big)$ of $e(o)$, indicated by the dotted borders: by this step we achieve that $k(V) \subset h(U) \times \mathbb{R}^p$ and in particular

$$k(V) \cap \big(\mathbb{R}^n \times \{0\}\big) = h(U) \times \{0\}.$$

We now use that $X \xrightarrow{e} e(X)$ is a homeomorphism: the image $e(U) \subset e(X)$ is open, and we find an open subset $V' \subset V$ with $V' \cap e(X) = e(U)$. Then the restriction $V' \xrightarrow{k} k(V')$ is a submanifold chart for $e(X) \subset Y$. Indeed we have

$$k\big(V' \cap e(X)\big) = (k \circ e)(U) = (k \circ e \circ h^{-1})\big(h(U)\big) = h(U) \times \{0\},$$

so that the inclusion in

$$k(V') \cap \big(\mathbb{R}^n \times \{0\}\big) \subset k(V) \cap \big(\mathbb{R}^n \times \{0\}\big) = h(U) \times \{0\}$$

must be an equality. □

We have seen that all tangent spaces of an open submanifold $X \subset \mathbb{R}^n$ may be identified with $\mathbb{R}^n$ itself. Naturally this is no longer true for a general n-manifold X, where we have to work with a whole family $(T_xX)_{x \in X}$ of tangent spaces parametrised by X. An obstacle to further progress is that so far this is but a family in the sense of set theory, in particular the tangent spaces at two different points of X are entirely unrelated. Therefore our next project will be to endow the disjoint union of the tangent spaces with a geometric structure: it will become a new differentiable manifold, which has dimension $2n$ and projects onto X in such a way that the tangent spaces become the fibres. This is the so-called *tangent bundle* of X. — Let us first discuss an important topological tool.

2.7 Lemma — *'Topologies are Local'* Let X be a set, and $X = \bigcup_{\lambda\in\Lambda} X_\lambda$ a cover of X by topological spaces X_λ such that for all $\lambda, \mu \in \Lambda$

(a) $X_\lambda \cap X_\mu$ is open in X_λ, and

(b) X_λ and X_μ induce the same subspace topology on $X_\lambda \cap X_\mu$.

Then there is a unique topology on X such that for every $\lambda \in \Lambda$

(c) X_λ is open in X, and

(d) X induces the given topology on X_λ.

Proof Let $\mathcal{T}$ be any topology on X that satisfies (c) and (d). Then by (d), for every $U \in \mathcal{T}$ and every $\lambda \in \Lambda$ the intersection $U \cap X_\lambda \subset X_\lambda$ is open. Conversely, if U has this property for every $\lambda \in \Lambda$ then according to (d), $U \cap X_\lambda$ belongs to the subspace topology induced by $\mathcal{T}$ on X_λ, so that $U \cap X_\lambda \in \mathcal{T}$ by (c) and thus

$$U = \bigcup_{\lambda\in\Lambda} (U \cap X_\lambda) \in \mathcal{T}.$$

Therefore

$$\mathcal{T} := \big\{ U \subset X \,\big|\, U \cap X_\lambda \subset X_\lambda \text{ is open for each } \lambda \in \Lambda \big\}$$

is the only conceivable solution to the problem.

To see that the topology $\mathcal{T}$ thus defined indeed satisfies (c) and (d) we fix a $\lambda \in \Lambda$. Then by (a) the intersection $X_\lambda \cap X_\mu$ is open in X_μ for every $\mu \in \Lambda$, and therefore $X_\lambda \in \mathcal{T}$, which proves (c).

In view of the property (c) just established the subspace topology that $\mathcal{T}$ induces on X_λ is

$$\big\{ U \in \mathcal{T} \,\big|\, U \subset X_\lambda \big\} = \big\{ U \subset X_\lambda \,\big|\, U \cap X_\mu \subset X_\mu \text{ is open for each } \mu \in \Lambda \big\},$$

and it remains to see that this coincides with the given topology of X_λ. Thus let $U \subset X_\lambda$ be open in X_λ. The intersection $U \cap X_\mu$ is open in $X_\lambda \cap X_\mu$, which in turn is open in X_μ by (a) and (b), so that $U \cap X_\mu$ is open in X_μ for all $\mu \in \Lambda$: this means that U belongs to $\mathcal{T}$. The converse is clear, choosing $\mu = \lambda$. □

2.8 Example Let X be an n-manifold with atlas $\mathcal{A} = \{(U_\lambda, h_\lambda)\}_{\lambda\in\Lambda}$ — while strictly speaking we simply have $\Lambda = \mathcal{A}$ and $\lambda = (U_\lambda, h_\lambda)$ we prefer the more suggestive notation as a family. If we want we may completely forget the topology of X because it can be reconstructed from the atlas $\mathcal{A}$. Indeed for each λ we recover the topology of U_λ from the standard topology of $\mathbb{R}^n$ by declaring the bijection $h_\lambda \colon U_\lambda \to h_\lambda(U_\lambda)$ a homeomorphism with the open subspace $h_\lambda(U_\lambda) \subset \mathbb{R}^n$. Then the chart domains U_λ form a cover of the set X by topological spaces. For any two indices λ, μ

(a) $U_\lambda \cap U_\mu$ is open in U_λ since $h_\lambda(U_\lambda \cap U_\mu)$ is open in $h_\lambda(U_\lambda)$, and

(b) the topology induced on $U_\lambda \cap U_\mu$ from U_λ coincides with that induced from U_μ since the transition map $h_\lambda(U_\lambda \cap U_\mu) \xrightarrow{h_\mu \circ h_\lambda^{-1}} h_\mu(U_\lambda \cap U_\mu)$ is a homeomorphism.

Lemma 2.7 now returns the topology of X as the unique one that satisfies (c) and (d).

At first sight the example seems pointless since after all the topology of a manifold X — as perhaps its most basic ingredient — should be known before anything else. But the example does illustrate how an atlas can be used to take a given manifold to pieces and reassemble it from them.

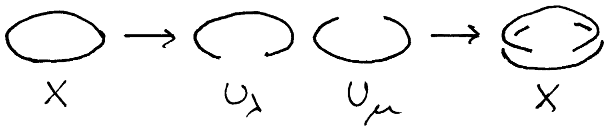

This process immediately becomes useful if between the two steps a third is performed which typically enhances the pieces in a way that cannot be directly described for the full manifold. One kind of enhancement consists in throwing in the tangent spaces, and this allows to construct the tangent bundle, as follows.

2.9 Theorem and Definition (1) Let X be a (smooth) n-manifold. The disjoint union

$$TX := \bigcup_{x\in X} \{x\}\times T_xX \xrightarrow{\ \pi\ } X$$

together with the projection $\pi\colon (x,v) \mapsto x$ is called the *tangent bundle* of X. The tangent bundle carries a natural structure of a smooth $2n$-manifold, which makes π a smooth map and is defined as follows. If $\mathcal{A} = \{(U_\lambda, h_\lambda)\}_{\lambda\in\Lambda}$ is the differentiable structure of X we give, for each $\lambda \in \Lambda$, the subset

$$\pi^{-1}U_\lambda = \bigcup_{x\in U_\lambda} \{x\}\times T_xX \subset TX$$

the topology that makes the bijection

$$\pi^{-1}U_\lambda \ni (x,v) \xmapsto{\ H_\lambda\ } \big(h_\lambda(x), v_{h_\lambda}\big) \in h_\lambda(U_\lambda)\times\mathbb{R}^n$$

a homeomorphism. The topological spaces $\pi^{-1}U_\lambda$ taken for all $\lambda\in\Lambda$, cover the set TX and define the topology of TX according to Lemma 2.7; even more, they provide a differentiable atlas

$$\tilde{\mathcal{A}} := \big\{(\pi^{-1}U_\lambda, H_\lambda)\big\}_{\lambda\in\Lambda}$$

and thereby the differentiable structure of TX.

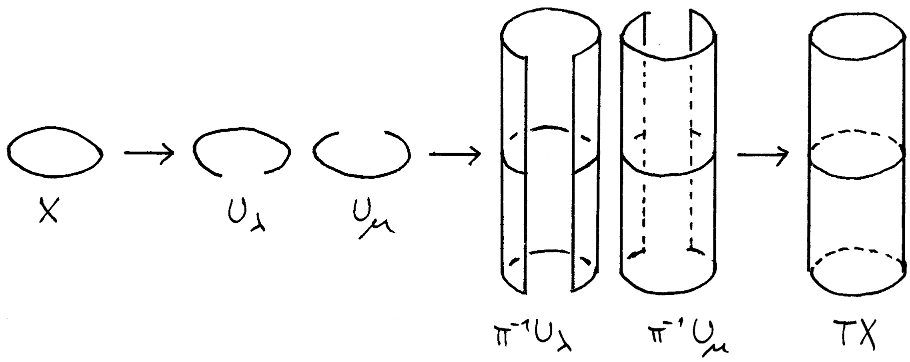

(2) Every smooth map $f\colon X\to Y$ induces a smooth mapping $Tf\colon TX\to TY$, called its *differential*, which acts on tangent vectors by[3]

$$T_xX \ni (x,v) \longmapsto \big(f(x), T_xf\cdot v\big) \in T_{f(x)}Y.$$

It thus renders the diagram

$$\begin{tikzcd} TX \arrow[r, "Tf"] \arrow[d] & TY \arrow[d] \\ X \arrow[r, "f"] & Y \end{tikzcd}$$

[3] We identify T_xX with the fibre $\pi^{-1}\{x\} = \{x\}\times T_xX$, and consequently write $v\in T_xX$ as $(x,v)\in T_xX$ when appropiate.

commutative and obeys the chain rule, which now simply reads

$$T\,\mathrm{id}_X = \mathrm{id}_{TX} \quad \text{and} \quad T(g\circ f) = Tg\circ Tf.$$

Proof In (1) we must verify conditions (a) and (b) of Lemma 2.7. This is analogous to Example 2.8: Condition (a) holds since

$$H_\lambda(\pi^{-1}U_\lambda\cap\pi^{-1}U_\mu) = h_\lambda(U_\lambda\cap U_\mu)\times\mathbb{R}^n \subset h_\lambda(U_\lambda)\times\mathbb{R}^n = H_\lambda(\pi^{-1}U_\lambda)$$

is an open subset, so that $\pi^{-1}U_\lambda\cap\pi^{-1}U_\mu$ is open in $\pi^{-1}U_\lambda$. Similarly (b) follows from the fact that the transition map

$$\begin{array}{ccccc} H_\lambda(\pi^{-1}U_\lambda\cap\pi^{-1}U_\mu) = h_\lambda(U_\lambda\cap U_\mu)\times\mathbb{R}^n & \xrightarrow{H_\mu\circ H_\lambda^{-1}} & h_\mu(U_\lambda\cap U_\mu)\times\mathbb{R}^n = H_\mu(\pi^{-1}U_\lambda\cap\pi^{-1}U_\mu) \\ (u,s) & \longmapsto & \big((h_\mu\circ h_\lambda^{-1})(u),\, D(h_\mu\circ h_\lambda^{-1})(u)\cdot s\big) \end{array}$$

is a homeomorphism. This is sufficient in order to apply Lemma 2.7 and conclude that $\tilde{\mathcal{A}}$ is a topological atlas for TX. Its differentiability follows from the fact that the homeomorphisms $H_\mu\circ H_\lambda^{-1}$ even are diffeomorphisms — note that in the C^∞ framework differentiation does not lower the degree of differentiability.

Before we proceed further it is interesting to note that the construction of the tangent bundle may equally well be based on any differentiable atlas $\mathcal{A}$ rather than the full differentiable structure of X. The uniqueness statement of Lemma 2.7 shows at once that the resulting topology is the same, and likewise the differentiable structures of TX coincide since all transition maps are diffeomorphisms.

TX is a Hausdorff space: Let $(a,v)\neq(b,w)$ be two distinct points of TX. In case $a\neq b$ we choose disjoint open neighbourhoods $U\subset X$ of a and $V\subset X$ of b; then $\pi^{-1}U$ and $\pi^{-1}V$ separate (a,v) and (b,w). If $a=b$ then we choose a $\lambda\in\Lambda$ with $a\in U_\lambda$. Then (a,v) and (a,w) are distinct points of the space $\pi^{-1}U_\lambda$ which under H_λ is homeomorphic to the Hausdorff space $h_\lambda(U_\lambda)\times\mathbb{R}^n$, so that again (a,v) and (a,w) can be separated.

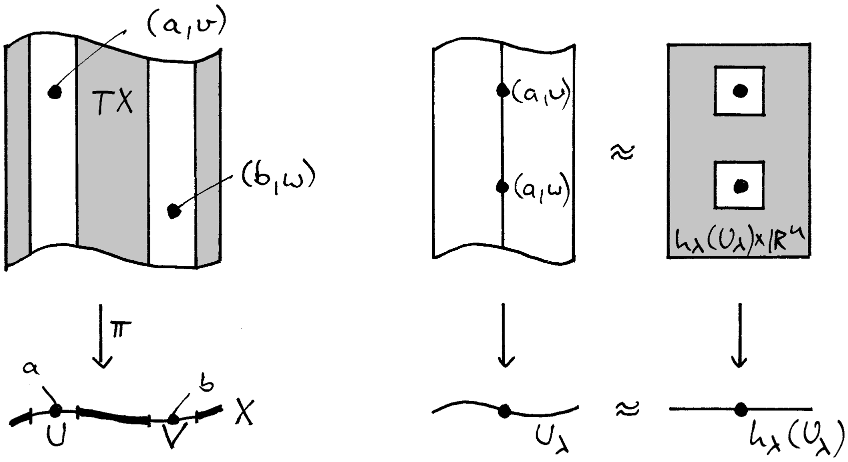

The topology of TX admits a countable base: By the result of Exercise 1.2 we find a countable smooth atlas $\mathcal{A}$ of X, and we use this atlas in the construction of TX. The resulting atlas $\tilde{\mathcal{A}}$ has the same index set as $\mathcal{A}$ and thus is countable too, and the result follows again from Exercise 1.2.

In order to prove differentiability of the projection $TX \xrightarrow{\pi} X$ it suffices to note that for any choice of a chart (U,h) the projection has a representation in the corresponding charts $(\pi^{-1}U, H)$ and (U,h), and that this representation is a Cartesian projection:

$$\begin{tikzcd} TX \arrow[d, "\pi"'] & \pi^{-1}U \arrow[l, hook'] \arrow[r, "H", "\simeq"'] \arrow[d, "\pi"'] & h(U)\times\mathbb{R}^n \arrow[d, "\mathrm{pr}"] \\ X & U \arrow[l, hook'] \arrow[r, "h", "\simeq"'] & h(U) \end{tikzcd}$$

We turn to part (2). The only statement that requires proof is the differentiability of Tf. To treat this local question near a point of TX over $o \in X$ we represent the map f in charts (U, h) at o and (V, k) at $f(o) \in Y$. In terms of the projections $TX \xrightarrow{\pi} X$ and $TY \xrightarrow{\rho} Y$ we write

$$\pi^{-1}U \ni (x, v) \xmapsto{H} \big(h(x), v_h\big) \in h(U) \times \mathbb{R}^n \quad \text{and} \quad \rho^{-1}V \ni (y, w) \xmapsto{K} \big(k(y), w_k\big) \in k(V) \times \mathbb{R}^p$$

for the corresponding charts of TX and TY. The differential Tf sends $\pi^{-1}U$ into $\rho^{-1}V$, and according to 2.3 its representation in the charts H and K

$$\begin{tikzcd} TX & \pi^{-1}U \arrow[l, hook'] \arrow[r, "Tf"] \arrow[d, "H"'] & \rho^{-1}V \arrow[r, hook] \arrow[d, "K"] & TY \\ & h(U) \times \mathbb{R}^n \arrow[r, "K \circ Tf \circ H^{-1}"] & k(V) \times \mathbb{R}^p & \end{tikzcd}$$

acts as

$$h(U) \times \mathbb{R}^n \ni (u, s) \longmapsto \big((k \circ f \circ h^{-1})(u), D(k \circ f \circ h^{-1})(u) \cdot s\big) \in k(V) \times \mathbb{R}^p.$$

This formula clearly defines a differentiable mapping, and this observation completes the proof. □

Exercises

2.1 Consider the point $a = \big(\frac{3}{4}\sqrt{2}, \frac{3}{4}\sqrt{2}, \frac{3}{2}\sqrt{3}\big) \in \mathbb{R}^3$ and let $v \in T_a\mathbb{R}^3$ be the tangent vector with Cartesian component $v_{\mathrm{id}} = (1, 1, -1)$. Compute the components $v_c \in \mathbb{R}^3$ and $v_s \in \mathbb{R}^3$ of v for each of the following charts at a:

- the chart c corresponding to cylinder coordinates $(r, \varphi, z) \mapsto (r\cos\varphi,\ r\sin\varphi,\ z)$
- the chart s corresponding to spherical coordinates $(R, \theta, \varphi) \mapsto (R\sin\theta\cos\varphi,\ R\sin\theta\sin\varphi,\ R\cos\theta)$

Do each case independently, then verify that v_c and v_s indeed are components of one and the same tangent vector in $T_a\mathbb{R}^3$.

Note that the given assignments in fact describe inverses of c and s, as in Exercise 1.3. Also remember that for standard matrix and differential calculus to work, vectors and even points of $\mathbb{R}^n$ should be written as columns — while in the formulation row notation is used for better readability.

2.2 Every real vector space X of finite dimension n has a canonical structure of a smooth n-manifold: explain why and how. Show furthermore that for every $a \in X$ the tangent space T_aX can be canonically identified with X itself. If $f: X \to Y$ is an affine mapping into another finite dimensional vector space, what is the differential $T_af: X \to Y$?

2.3 Let X be a non-empty compact manifold of positive dimension. Prove that for every differentiable function $f: X \to \mathbb{R}$ there exist two distinct points $a, b \in X$ with $T_af = 0$ and $T_bf = 0$.

2.4 Let $X \subset \mathbb{R}^n$ be a smooth submanifold, and let $a \in X$: the linear subspace $T_aX \subset T_a\mathbb{R}^n = \mathbb{R}^n$ is another (quite special) submanifold. Prove that the orthogonal projection $\pi: \mathbb{R}^n \to T_aX$ restricts to a map $\pi': X \to T_aX$ which is a local diffeomorphism at a.

2.5 Let X be a differentiable n-manifold, $o \in X$ a point. We consider smooth functions $f: V \to \mathbb{R}$ defined on open neighbourhoods $V \subset X$ of o and declare two such functions $f: V \to \mathbb{R}$ and $g: W \to \mathbb{R}$

equivalent if there exists a third neighbourhood U of o with $U \subset V \cap W$ and $f|U = g|U$. Briefly convince yourself that the equivalence classes form a real algebra — that is, a vector space with a compatible ring multiplication; this algebra is denoted by $\mathcal{E}_{X,o}$ or $\mathcal{E}_o$ for short, and its elements are called the *germs* of smooth functions at o (who is familiar with the notion will recognise $\mathcal{E}_{X,o}$ as a direct limit).

Set up a canonical isomorphism between the tangent space T_oX and the space of scalar *derivations* of $\mathcal{E}_o$, that is $\mathbb{R}$-linear functions $\eta\colon \mathcal{E}_o \to \mathbb{R}$ with the property

$$\eta(f \cdot g) = \eta(f) \cdot g(o) + f(o) \cdot \eta(g) \quad \text{for all } f, g \in \mathcal{E}_o\,.$$

Express the differential of a smooth mapping $X \xrightarrow{F} Y$ in terms of this alternative description of the tangent space.

Hint Taylor's formula — or a version of the mean value theorem — tells you that every germ $f \in \mathcal{E}_{\mathbb{R}^n,0}$ can be written in the form

$$f(x) = f(0) + \sum_{j=1}^{n} x_j \cdot f_j(x)$$

with germs $f_j \in \mathcal{E}_{\mathbb{R}^n,0}$.

2.6 Let $f\colon \mathbb{R}^{n+1} \to \mathbb{R}$ be given by

$$f(x) = \sum_{j=0}^{k} x_j^2 - \sum_{j=k+1}^{n} x_j^2$$

with $k \in \{-1, \ldots, n\}$. Prove that the affine quadric

$$Q := f^{-1}\{1\} \subset \mathbb{R}^{n+1}$$

is diffeomorphic to $S^k \times \mathbb{R}^{n-k}$. — Generalise to the case where $f\colon \mathbb{R}^{n+1} \to \mathbb{R}$ is an arbitrary quadratic form.

2.7 According to Exercise 1.10 the orthogonal group

$$O(n) = \big\{x \in \mathrm{Mat}(n\times n, \mathbb{R}) \,\big|\, x^{\mathrm{t}} x = 1\big\}$$

is a Lie group: what is its Lie algebra $\mathfrak{o}(n)$? Differentiate the following mappings at the point $1 \in O(n)$:

- the inversion $O(n) \ni x \mapsto i(x) := x^{-1} \in O(n)$,
- for fixed $u \in O(n)$ the conjugation mapping $O(n) \ni x \mapsto c_u(x) := uxu^{-1} \in O(n)$,
- the so-called *adjoint representation* $O(n) \ni u \mapsto \mathrm{Ad}(u) := Dc_u(1) \in \mathrm{End}\big(\mathfrak{o}(n)\big)$ with values in the vector space of linear endomorphisms of $\mathfrak{o}(n)$.

2.8 Prove that the *special unitary group*

$$SU(2) = \{u \in \mathrm{Mat}(2\times 2, \mathbb{C}) \,|\, \overline{u}^{\mathrm{t}} u = 1 \text{ and } \det u = 1\}$$

is a Lie group with Lie algebra

$$\mathfrak{su}(2) = \{x \in \mathrm{Mat}(2\times 2, \mathbb{C}) \,|\, \overline{x}^{\mathrm{t}} + x = 0 \text{ and } \operatorname{tr} x = 0\}.$$

Verify that $\mathfrak{su}(2)$ is a Euclidean vector space under the scalar product

$$\mathfrak{su}(2) \times \mathfrak{su}(2) \ni (x, y) \longmapsto \langle x, y \rangle := -\operatorname{tr} xy \in \mathbb{R}.$$

Prove that the adjoint representation of $SU(2)$,

$$SU(2) \ni u \longmapsto Dc_u(1) \in \operatorname{End}\big(\mathfrak{su}(2)\big) \quad \text{with } c_u(x) = uxu^{-1} = ux\overline{u}^{\mathrm{t}}$$

acts by orthogonal automorphisms of $\mathfrak{su}(2)$ and thus defines a homomorphism $f\colon SU(2) \to SO\big(\mathfrak{su}(2)\big)$ of groups.

Finally, express the differential $T_1 f$ in terms of the Pauli base

$$X = i \cdot \sigma_x = i \cdot \begin{pmatrix} & 1 \\ 1 & \end{pmatrix} \qquad Y = i \cdot \sigma_y = i \cdot \begin{pmatrix} & -i \\ i & \end{pmatrix} \qquad Z = i \cdot \sigma_z = i \cdot \begin{pmatrix} 1 & \\ & -1 \end{pmatrix}$$

of $\mathfrak{su}(2)$ and prove that f is a surjective local diffeomorphism with kernel $\{\pm 1\}$.

2.9 Let X be a connected n-manifold and $f\colon X \to X$ a smooth mapping with the property that $f \circ f = f$. Prove that the image set $S := f(X) \subset X$ is a closed differentiable submanifold of X.

Hints Note that S coincides with the fixed point set $\{x \in X \mid f(x) = x\}$ of f. — The problem is an application of the constant rank theorem.

2.10 Let $f\colon X \to Y$ be an injective immersion. Prove that if furthermore X is compact then f is a smooth embedding.

2.11 Describe an embedding $S^1 \times S^1 \to \mathbb{R}^3$ explicitly in terms of Cartesian and also in terms of angle coordinates θ and φ on S^1.

2.12 For arbitrary $n, p \in \mathbb{N}$ construct an embedding of $S^n \times S^p$ in $\mathbb{R}^{n+p+1}$.

Hint The manifolds $\mathbb{R} \times S^p$ and $\mathbb{R}^{p+1} \smallsetminus \{0\}$ are diffeomorphic.

2.13 Let $X := \mathbb{R}^2/\mathbb{Z}^2$ be the two-dimensional torus, and for fixed non-zero $\beta \in \mathbb{R}$ consider the parametrised curve

$$\mathbb{R} \ni t \overset{f}{\longmapsto} (t, \beta t) + \mathbb{Z}^2 \in X.$$

If $\beta \in \mathbb{Q}$ then f induces a smooth embedding of the circle S^1 as a smooth submanifold of X: make this precise and prove it.

Prove that if β is irrational then f is an injective immersion, and the image set $f(\mathbb{R}) \subset X$ is dense.

Hint The case of irrational β is one of the numerous manifestations of the following fact, which is half analysis and half algebra, and which you may simply use or wish to prove first: Every additive subgroup of the real line is either dense or cyclic.

2.14 Let $e\colon S \to X$ be an embedding of manifolds. Prove that $Te\colon TS \to TY$ also is an embedding.

3 Vector Bundles

The tangent bundle defined in the previous section is a particular instance of the notion of vector bundle. Vector bundles make precise the idea of linear algebra parametrised by non-linear geometry, a situation which arises naturally in a variety of topological, algebraic, and analytic contexts. It is the differentiable real version which is relevant here.

3.1 Definition A (*differentiable* or *smooth* real) *vector bundle* of *rank* d on, or over a manifold X consists of another manifold E, a smooth mapping

$$E \xrightarrow{\pi} X,$$

and a structure of a real d-dimensional vector space on each fibre

$$E_x := \pi^{-1}\{x\} \quad (x \in X).$$

These data are required to obey the *axiom of local triviality*: for each $x \in X$ there exists a *bundle chart* at x, consisting of an open neighbourhood $U \subset X$ of x and a diffeomorphism h that lets the diagram

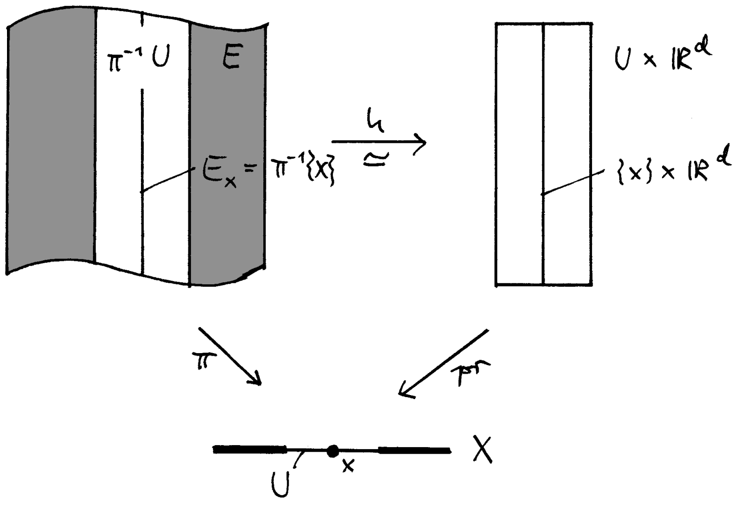

$$\begin{tikzcd} \pi^{-1}U \arrow[rr, "h"] \arrow[dr, "\pi"'] & & U \times \mathbb{R}^d \arrow[dl, "\mathrm{pr}"] \\ & U & \end{tikzcd}$$

commute and restricts to a — necessarily bijective — linear mapping

$$E_x \xrightarrow{h_x} \{x\} \times \mathbb{R}^d = \mathbb{R}^d$$

on each fibre. A *bundle atlas* is of course, a family of bundle charts $\pi^{-1}U \to U \times \mathbb{R}^d$ such that the open sets U cover X.

The manifolds X and E are called the *base* and the *total space* of the bundle respectively, while π is called the *bundle projection*. In practice often just the total space of a bundle is written down when the remaining data are clear from the context. — Quite generally, in the context of bundles the notion of locality always refers to the base rather than the total space.

3.2 Examples (1) Given a manifold X and an integer $d \in \mathbb{N}$ the *product bundle*

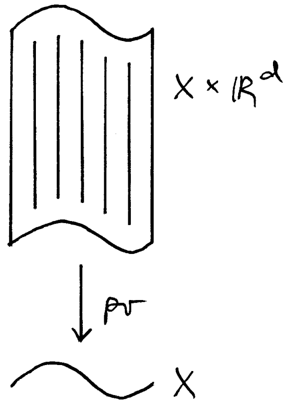

$$X \times \mathbb{R}^d \xrightarrow{\text{pr}} X$$

clearly is a vector bundle of rank d; in this case the identity mapping provides a single global bundle chart.

(2) The tangent bundle of a manifold X: if (U, h) is a chart then

$$TU \ni (x, v) \longmapsto \big(x, v_h\big) \in U \times \mathbb{R}^n$$

is a bundle chart. It may be expressed as the composition

$$TU \xrightarrow{H} h(U) \times \mathbb{R}^n \xrightarrow{h^{-1}\times\text{id}} U \times \mathbb{R}^n$$

where $H: (x, v) \to \big(h(x), v_h\big)$ is the chart derived from h and used in the construction of TX. The factor h^{-1} on the right simply undoes the action of H on the base, which here is undesired.

(3) The projective space $\mathbb{R}P^n$ carries the so-called *tautological bundle*

$$T = \big\{\big([x], v\big) \in \mathbb{R}P^n \times \mathbb{R}^{n+1} \,\big|\, v \in \mathbb{R}x\big\} \ni \big([x], v\big) \xrightarrow{\pi} [x] \in \mathbb{R}P^n.$$

It is a vector bundle of rank one or *line bundle*. Its name comes from the fact that the point $[x] = \{\pm x\} \in \mathbb{R}P^n$ may equivalently be thought of as the line $\mathbb{R}x \subset \mathbb{R}^{n+1}$, so that the fibre $T_{\mathbb{R}x} = \mathbb{R}x$ over it is just the point itself.

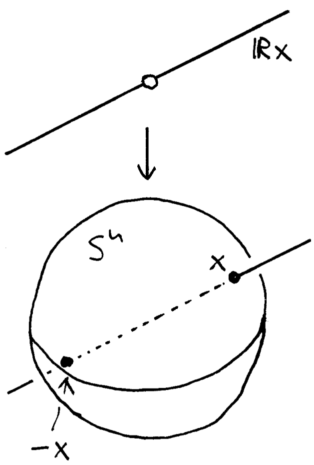

Recall that $\mathbb{R}P^n$ is covered by the charts

$$U_j = \{[x] \in \mathbb{R}P^n \mid x_j \neq 0\} \xrightarrow{h_j} U^n \quad \text{for } j = 0, \dots, n;$$

h_j sends $[x] \in U_j$ to the point $(x_0, \dots, \widehat{x_j}, \dots, x_n) \in U^n$ if the representative $x \in S^n$ is normalised by the condition $x_j > 0$. A suitable bundle chart for T over U_j is given by

$$\pi^{-1}U_j \ni ([x], v) \longmapsto \Big([x], \frac{v_j}{x_j}\Big) \in U_j \times \mathbb{R}$$

with inverse $([x], t \cdot x) \hookleftarrow ([x], t)$; here again x is normalised such that $x_j > 0$.

Several kinds of structure compatible mappings between vector bundles may be envisaged.

3.3 Definition (1) Let $E \to X$ and $F \to Y$ be vector bundles over manifolds X and Y. A *bundle mapping* from E to F is a pair (f, g) of smooth maps such that the diagram

$$\begin{tikzcd} E \arrow[r, "g"] \arrow[d, "\pi"'] & F \arrow[d, "\rho"] \\ X \arrow[r, "f"] & Y \end{tikzcd}$$

is commutative and $g\colon E \to F$ sends the vector space E_x isomorphically onto $F_{f(x)}$ for each $x \in X$.

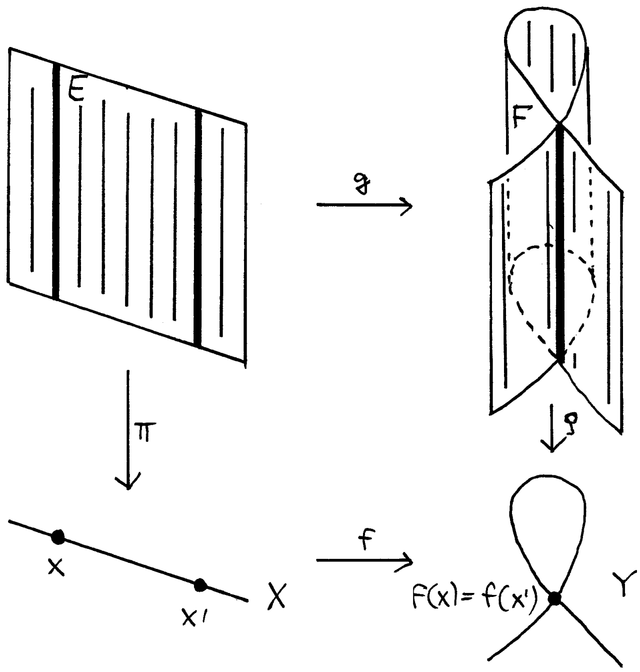

(For better readability the right hand side of the figure shows just the part of F that sits over the image $f(X)$.)

(2) Let $E \to X$ and $F \to X$ be vector bundles over one and the same manifold X. A *bundle homomorphism* from E to F is a smooth map $g\colon E \to F$ that for each $x \in X$ sends the vector space E_x linearly (but not necessarily bijectively) into F_x.

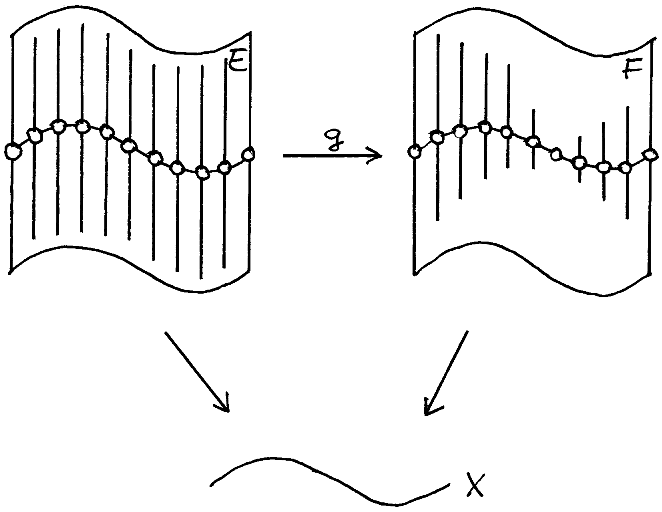

(3) A *bundle isomorphism* between bundles over the same base is a homomorphism which is invertible by another homomorphism. A bundle over X is called *trivial* if it is isomorphic to a product bundle $X\times\mathbb{R}^d \xrightarrow{\text{pr}} X$ (where d must be the rank of E of course).

Note The terminology introduced in (3) makes the axiom of local triviality self-explaining.

3.4 Examples (1) Bundle homomorphisms between product bundles $X\times\mathbb{R}^d \xrightarrow{g} X\times\mathbb{R}^e$ correspond to smooth functions $f\colon X \to \operatorname{Mat}(e\times d,\mathbb{R})$ via $g(x,v) = (x, f(x)\cdot v)$: note that by elementary analysis g is differentiable if and only if f is differentiable. It follows that g is an isomorphism of bundles if and only if $d = e$ and f takes values in the open subset $GL(d,\mathbb{R}) \subset \operatorname{Mat}(d\times d,\mathbb{R})$, for in that case g^{-1} corresponds to the equally smooth function $X \ni x \mapsto f(x)^{-1} \in GL(d,\mathbb{R}) \subset \operatorname{Mat}(d\times d,\mathbb{R})$. For a homomorphism $E \xrightarrow{g} F$ between general bundles these observations imply that g is an isomorphism if and only if it sends each fibre E_x isomorphically onto F_x, that is if and only if g is bijective as a map.

(2) The tangent bundle TS^1 is trivial[1], since

$$TS^1 = \big\{(x,v) \in S^1\times\mathbb{R}^2 \,\big|\, x^{\mathrm{t}}v = 0\big\} \ni (x,v) \longmapsto (x, -x_2v_1+x_1v_2) \in S^1 \times \mathbb{R}$$

is a bundle map: it measures v with respect to the tangent vector of length one pointing in anti-clockwise direction, and is inverted by $\big(x,\, t\cdot(-x_2,x_1)\big) \leftarrow\!\shortmid (x,t)$.

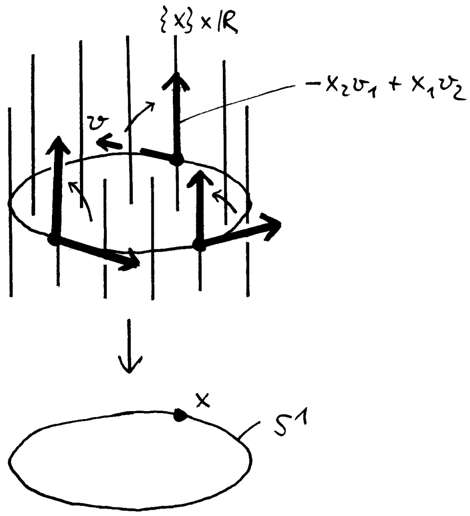

[1] Trying to visualise the tangent bundle invariably encounters a psychological barrier as it requires to think of the tangent directions as independent and transverse to the directions on the manifold — which on the other hand by definition are just one and the same.

This construction does not generalise to higher dimensions, and in fact TS^n is almost never trivial, the exceptions being $n = 0, 1, 3$, and 7.

3.5 Theorem and Definition Let $F \xrightarrow{\rho} Y$ be a rank d vector bundle and $f: X \to Y$ a map. Then

$$f^*F := \big\{(x, w) \in X \times F \,\big|\, f(x) = \rho(w)\big\} \subset X \times F$$

is a smooth submanifold, and the projection

$$f^*F \ni (x, w) \overset{\pi}{\longmapsto} x \in X$$

makes f^*F a vector bundle over X called the *induced* bundle or *pull-back* of F; it has the same rank d as F. The pair comprising f and the restricted projection $f^*F \ni (x, w) \overset{g}{\longmapsto} w \in F$ is a bundle map:

$$\begin{tikzcd} f^*F \arrow[r, "g"] \arrow[d, "\pi"'] & F \arrow[d, "\rho"] \\ X \arrow[r, "f"] & Y \end{tikzcd}$$

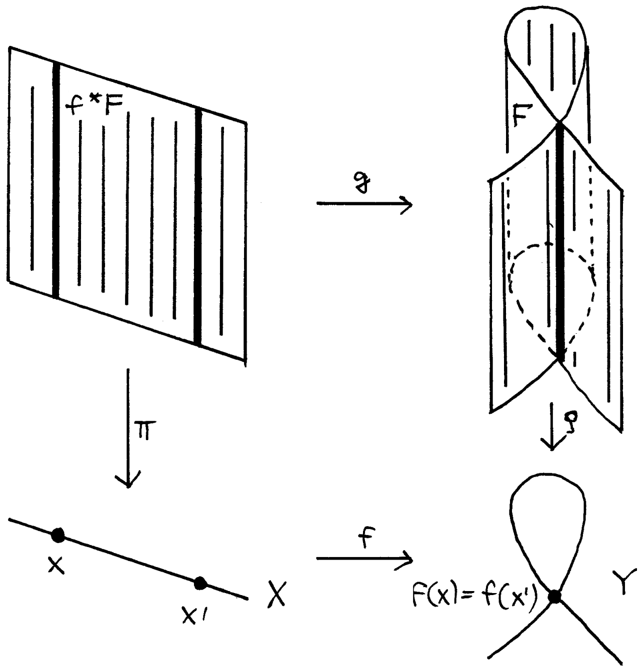

The pull-back bundle is natural in the sense that every bundle homomorphism $h: F \to F'$ pulls back to a homomorphism $f^*h: f^*F \to f^*F'$ such that $f^*\,\mathrm{id} = \mathrm{id}$ and $f^*(k \circ h) = f^*k \circ f^*h$ — in particular every bundle induced from a trivial bundle is trivial. The pull-back is also natural with respect to the inducing map: if e and f are composable maps then $(f \circ e)^*F$ is canonically isomorphic to e^*f^*F.

Proof The submanifold property of $f^*F \subset X \times F$ is local with respect to X and a fortiori with respect to Y, since f is continuous. To prove it we may thus assume that $Y \subset \mathbb{R}^p$ is an open subset and that $F = Y \times \mathbb{R}^d \to Y$ is the product bundle. Then

$$f^*F = \big\{(x, y, t) \in X \times Y \times \mathbb{R}^d \,\big|\, f(x) = y\big\} \subset X \times Y \times \mathbb{R}^d$$

is, up to order of the factors, the graph of the function $X \times \mathbb{R}^d \ni (x, t) \mapsto f(x) \in Y$, so that the formula

$$X \times Y \times \mathbb{R}^d \ni (x, y, t) \longmapsto \big(x, y - f(x), t\big) \in X \times \mathbb{R}^p \times \mathbb{R}^d$$

yields a flattening chart for f^*F. For the same reason the map

$$f^*F \ni (x,y,t) \overset{g}{\longmapsto} (x,t) \in X \times \mathbb{R}^d$$

is, in this special situation, a diffeomorphism; it makes the diagram

$$\begin{tikzcd} f^*F \arrow[rr, "g", "\simeq"'] \arrow[dr, "\pi"'] & & X \times \mathbb{R}^d \arrow[dl, "\mathrm{pr}"] \\ & X & \end{tikzcd}$$

commutative, restricts to the identity map between fibres

$$(f^*F)_x = \{x\} \times \{f(x)\} \times \mathbb{R}^d \to \{x\} \times \mathbb{R}^d$$

and therefore serves as a (global) bundle chart for f^*F.

The first naturality property stated is the subject of Exercise 3.1. As to the second, for given maps $U \xrightarrow{e} X \xrightarrow{f} Y$ the Cartesian projection induces a canonical bundle isomorphism between

$$e^*f^*F = \big\{(u,x,w) \in U \times X \times F \,\big|\, e(u) = x \text{ and } f(x) = \rho(w)\big\}$$

and

$$(f \circ e)^*F = \big\{(u,w) \in U \times F \,\big|\, (f \circ e)(u) = \rho(w)\big\}\,. \qquad \Box$$

3.6 Example In the case where $f\colon X \subset Y$ is the inclusion map of a submanifold the induced bundle f^*F is canonically isomorphic to the restriction $\rho^{-1}X \xrightarrow{\rho} X$, which in turn is simply written $F|X \to X$.

Proof The Cartesian projection

$$f^*F = \big\{(x,w) \in X \times F \,\big|\, x = \rho(w)\big\} \longrightarrow \big\{w \in F \,\big|\, \rho(w) \in X\big\} = F|X$$

does the job. $\Box$

The concepts of bundle mapping and bundle homomorphism suggest a common generalisation, which nevertheless can be reduced to these more restricted notions:

3.7 Factorisation Lemma Let $E \xrightarrow{\pi} X$ and $F \xrightarrow{\rho} Y$ be vector bundles (of possibly different ranks) over manifolds X and Y (of possibly different dimensions). Let

$$\begin{tikzcd} E \arrow[r, "g"] \arrow[d, "\pi"'] & F \arrow[d, "\rho"] \\ X \arrow[r, "f"] & Y \end{tikzcd}$$

be a commutative diagram of smooth maps such that for every $x \in X$ the restriction

$$E_x \ni v \longmapsto g(v) \in F_{f(x)}$$

is linear. Then there exists a unique bundle homomorphism $\tilde g\colon E \to f^*F$ that renders the diagram

$$\begin{tikzcd} E \arrow[rr, "g"] \arrow[d, "\pi"'] \arrow[dr, "\tilde g"] & & F \arrow[d, "\rho"] \\ X \arrow[dr, equal] & f^*F \arrow[ur] \arrow[d] & Y \\ & X \arrow[ur, "f"'] & \end{tikzcd}$$

commutative.

Proof There is no other choice for $\tilde{g}$ but to send $v \in E$ to $\big(\pi(v), g(v)\big) \in f^*F \subset X \times F$, and this clearly yields a smooth mapping with the desired properties. $\square$

3.8 Corollary If there exists a bundle map (g, f) from E to F then $E \simeq f^*F$. $\square$

Note While by the very definition bundle homomorphisms act on the fibres but not on the common base space the corollary shows that by contrast a bundle mapping acts on the base space but not essentially on the fibres: in the induced bundle the fibres of F are simply pulled back from points of Y over points of X. The factorisation lemma states that for the more general type of morphism between bundles the two effects can be neatly separated. In the following example the separation is between the differentiable mapping f on one side, and the linear action of its derivative Tf on the other.

3.9 Example The differential of a smooth map $f\colon X \to Y$ is a special case:

$$\begin{tikzcd}
TX \arrow[rr, "Tf"] \arrow[dr, "\tilde{g}"] \arrow[d] & & TY \arrow[d] \\
X \arrow[dr, equal] & f^*TY \arrow[ur] \arrow[d] & Y \\
& X \arrow[ur, "f"'] &
\end{tikzcd}$$

The resulting bundle homomorphism $\tilde{g}$ may be considered as an alternative version of the differential of f.

3.10 Definition Let $E \xrightarrow{\pi} X$ be a vector bundle. A *section* of E is a map $s\colon X \to E$ that satisfies $\pi \circ s = \mathrm{id}_X$. In view of Theorem 2.6 this condition forces s to be a smooth embedding, and sometimes the embedded submanifold $s(X) \subset E$ rather than s is referred to as a section — in any case the mapping s can be recovered from $s(X)$.

By the rules of differential calculus the sections of a given bundle $E \to X$ form a real vector space. Its zero vector is the section which assigns to each point $x \in X$ the zero vector of the fibre E_x, and is called the *zero section*; it provides a canonical embedding of X in E and is sometimes used to identify X with its image in E.

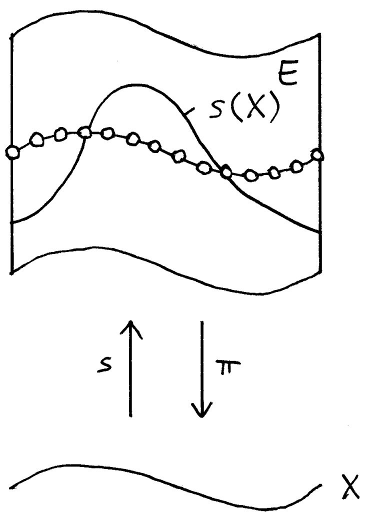

Linear algebra knows a variety of standard methods to construct new vector spaces from one or several given ones in an functorial way. These constructions readily extend to vector bundles. The technique is analogous to that which has led to the tangent bundle: the given vector bundles are taken apart using bundle atlases, the linear construction is applied to the pieces, and finally the results are glued back to form the new bundle.

3.11 Constructions with Vector Bundles (1) Let $E \xrightarrow{\pi} X$ and $F \xrightarrow{\rho} X$ be vector bundles over a manifold X. We define their *Whitney sum* $E \oplus F$, a new vector bundle over X, taking the direct sum fibre-wise:

$$(E \oplus F)_x := E_x \oplus F_x = E_x \times F_x.$$

While this defines the bundle $E \oplus F \xrightarrow{\sigma} X$ as a set over X with the required vector space structure on each fibre we must still declare its topology and differentiable structure. To this end pick bundle atlases

$$\Big\{\pi^{-1}U_\lambda \xrightarrow{h_\lambda} U_\lambda \times \mathbb{R}^d\Big\}_{\lambda\in\Lambda} \quad \text{and} \quad \Big\{\rho^{-1}V_\mu \xrightarrow{k_\mu} V_\mu \times \mathbb{R}^e\Big\}_{\mu\in\mathrm{M}}$$

for E and F; passing to intersections $U_\lambda \cap V_\mu$ we may assume a common index set $\Lambda = \mathrm{M}$ and identical covering sets $U_\lambda = V_\lambda$. For every $\lambda \in \Lambda$ the composition

$$\begin{array}{ccccc} \sigma^{-1}U_\lambda & \longrightarrow & U_\lambda \times \mathbb{R}^d \times U_\lambda \times \mathbb{R}^e & \xrightarrow{\text{pr}} & U_\lambda \times \mathbb{R}^d \times \mathbb{R}^e \\ (v,w) & \longmapsto & \big(h_\lambda(v), k_\lambda(w)\big) \quad (x,s,y,t) & \longmapsto & (x,s,t) \end{array}$$

is linear on fibres and bijective; indeed it is inverted by the assignment

$$\sigma^{-1}U_\lambda \ni \big(h_\lambda^{-1}(x,s), k_\lambda^{-1}(x,t)\big) \longleftarrow\!\!\shortmid (x,s,t) \in U_\lambda \times \mathbb{R}^d \times \mathbb{R}^e.$$

These mappings can be used in the usual way to put on $E \oplus F$ first a topology using Lemma 2.7, and then the manifold and vector bundle structures, noting that the arising transition maps are smooth and fibrewise linear. — A priori all these structures may depend on which bundle atlases are chosen for E and F. In order to compare two choices $\mathcal{B} = (U_\lambda, h_\lambda, k_\lambda)_{\lambda\in\Lambda}$ and $\mathcal{B}' = (U'_\lambda, h'_\lambda, k'_\lambda)_{\lambda\in\Lambda'}$ we may assume that the open cover $(U_\lambda)_{\lambda\in\Lambda}$ is a refinement of $(U'_\lambda)_{\lambda\in\Lambda'}$. Then the uniqueness statement of Lemma 2.7, applied to $(U_\lambda)_{\lambda\in\Lambda}$, shows that $\mathcal{B}'$ puts the same topology on $E \oplus F$ as $\mathcal{B}$ does. Likewise the smooth structures on $E \oplus F$ coincide since all transition maps are diffeomorphisms.

The Whitney construction is rightfully called a sum since it has the expected universal property: given bundle homomorphisms $e\colon E \to G$ and $f\colon F \to G$ to a third bundle over X there is a unique bundle homomorphism $E \oplus F \to G$ which restricts to e on $E \subset E \oplus F$ and to f on $F \subset E \oplus F$. With equal right $E \oplus F$ could be called the direct product of E and F since it also satisfies the dual universal property: given bundle homomorphisms $e\colon D \to E$ and $f\colon D \to F$ there is a unique bundle homomorphism $D \to E \oplus F$ which projects to e under $E \oplus F \to E$ and to f under $E \oplus F \to F$.

It is at once seen, and in fact a formal consequence of the universal property that the Whitney sum is functorial: every pair of bundle homomorphisms $E \xrightarrow{g} E'$ and $F \xrightarrow{h} F'$ induces a sum homomorphism of bundles

$$E \oplus F \xrightarrow{g \oplus h} E' \oplus F',$$

and apart from the triviality $\mathrm{id}_E \oplus \mathrm{id}_F = \mathrm{id}_{E\oplus F}$ the rule $(g \circ g') \oplus (h \circ h') = (g \oplus h) \circ (g' \oplus h')$ holds for pairs of composable homomorphisms.

(2) The Whitney sum serves as a model for the bundle versions of other constructions, including:

- Given bundles $E \to X$ and $F \to X$ there is a bundle $\mathrm{Hom}(E,F)$ such that for each $x \in X$ the fibre $\mathrm{Hom}(E,F)_x = \mathrm{Hom}(E_x, F_x)$ is the space of linear maps from E_x to F_x; in particular we have the dual bundle $E\check{} = \mathrm{Hom}(E, X \times \mathbb{R})$ where $E_x\check{} = \mathrm{Hom}(E_x, \mathbb{R})$ is the space of linear forms.
- Under the same conditions the tensor product of two bundles $E \otimes F$ may be formed.
- Given any bundle $E \to X$ we have its symmetric and alternating powers $S^kE \to X$ and $\Lambda^kE \to X$.
- Closely related with the latter but more widely known are symmetric bilinear forms and alternating multilinear k-forms; starting from any bundle $E \to X$ they give rise to the vector bundles $\mathrm{Sym}\,E \to X$ — isomorphic to $S^2E\check{} \to X$, and $\mathrm{Alt}^kE \to X$, which is isomorphic to $\Lambda^kE\check{}$.

3.12 Examples (1) For any $n \in \mathbb{N}$ let $N = S^n \times \mathbb{R} \longrightarrow S^n$ denote the trivial line bundle on S^n. Then the Whitney sum $N \oplus TS^n$ is a trivial vector bundle of rank $n{+}1$ over S^n. Indeed, the commutative diagram

$$\begin{tikzcd} N \oplus TS^n \arrow[rr, "h"] \arrow[dr] & & S^n \times \mathbb{R}^{n+1} \arrow[dl] \\ & S^n & \end{tikzcd}$$

with $h(x, t; x, v) = (x, tx + v)$ describes an isomorphism: in order to invert h just decompose a given vector of the fibre $\{x\} \times \mathbb{R}^{n+1}$ into its components which are parallel and orthogonal to x.

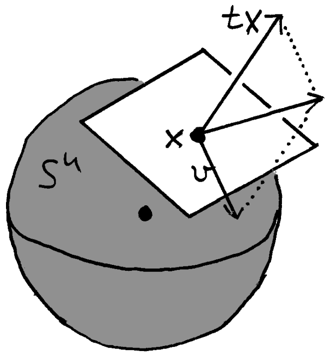

(2) For every line bundle $L \to X$ the bundle $\mathrm{Hom}(L, L) \to X$ is trivial, since the bundle homomorphism $X \times \mathbb{R} \to \mathrm{Hom}(L, L)$ that sends $1 \in \mathbb{R}_x$ to the identity map of L_x is bijective — its inverse is given by the trace.

(3) In Example 3.4(1) we observed the correspondence between homomorphisms of product bundles $X \times \mathbb{R}^n \to X \times \mathbb{R}^p$ and matrix valued functions $X \to \mathrm{Mat}(p \times n, \mathbb{R})$. We now can globalise: for any given vector bundles $E \xrightarrow{\pi} X$ and $F \xrightarrow{\rho} X$ the bundle homomorphisms $g\colon E \to F$ correspond to the sections $s\colon X \to \mathrm{Hom}(E, F)$ via the formula $g(v) = (s \circ \pi)(v) \cdot v$. The local description in 3.4(1) also shows that the points $x \in X$ such that $g_x\colon E_x \to F_x$ is an isomorphism form an open subset of X.

3.13 Definition Let $E \xrightarrow{\pi} X$ be a rank d vector bundle on a manifold X, and $s \le d$. A subset $S \subset E$ together with the restriction $S \xrightarrow{\pi|S} X$ is a *subbundle* of rank s if at every $x \in X$ there exists a bundle chart (U, h) for E

$$\begin{tikzcd} \pi^{-1}U \arrow[rr, "h"] \arrow[dr, "\pi"'] & & U \times \mathbb{R}^d \arrow[dl, "\mathrm{pr}"] \\ & U & \end{tikzcd}$$

with $h(\pi^{-1}U \cap S) = U \times \big(\mathbb{R}^s \times \{0\}^{d-s}\big)$ — the restricted projection $S \to X$ then clearly is a vector bundle in its own right.

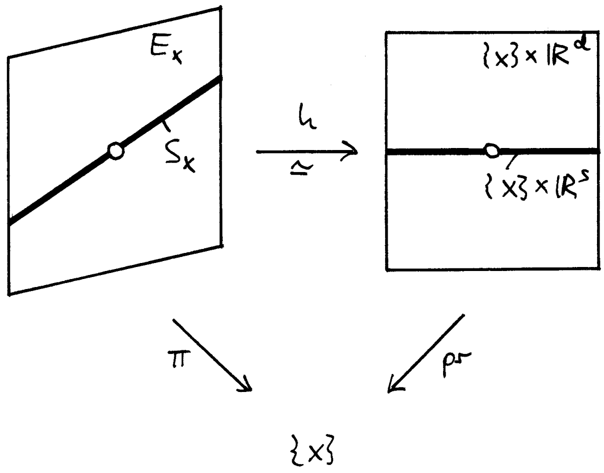

3.14 Examples (1) If $X \subset Y$ is a submanifold then $TX \subset TY|X$ is a subbundle of the restricted tangent bundle: assuming $\dim X = n$ and $\dim Y = n+p$ let $\pi\colon TY \to Y$ denote the bundle projection and let (U, h) be a flattening chart for $X \subset Y$, so that

$$h(U \cap X) = h(U) \cap \big(\mathbb{R}^n \times \{0\}\big) \subset \mathbb{R}^n \times \mathbb{R}^p.$$

Then the bundle chart for $TY|X$

$$\pi^{-1}(U \cap X) \ni (x, v) \longmapsto (x, T_x h \cdot v) \in (U \cap X) \times \mathbb{R}^{n+p}$$

sends $\pi^{-1}(U \cap X) \cap TX = T(U \cap X)$ onto $(U \cap X) \times (\mathbb{R}^n \times \{0\})$ and thus is a subbundle chart.

(2) The tautological bundle $T \to \mathbb{R}P^n$ was defined as the subset

$$T = \big\{([x], v) \in \mathbb{R}P^n \times \mathbb{R}^{n+1} \,\big|\, v \in \mathbb{R}x\big\}$$

of the total space of the trivial bundle $\mathbb{R}P^n \times \mathbb{R}^{n+1} \to \mathbb{R}P^n$, and one might ask whether it is a subbundle. Indeed it is: over the open set $U_j = \{[x] \in \mathbb{R}P^n \,|\, x_j \neq 0\}$ our bundle chart for T from Example 3.2(3)

$$\pi^{-1}U_j \ni ([x], v) \longmapsto \Big([x], \frac{v_j}{x_j}\Big) \in U_j \times \mathbb{R} \quad (x \text{ normalised so that } x_j > 0)$$

can be interpreted as assigning to the vector $v \in T_{[x]}$ its projection along $\mathbb{R}x$, and this suggests its extension to the subbundle chart

$$U_j \times \mathbb{R}^{n+1} \ni ([x], v) \longmapsto \Big([x], \frac{v_j}{x_j}, v - \frac{v_j}{x_j}\, x\Big) \in U_j \times \mathbb{R} \times \big\{w \in \mathbb{R}^{n+1} \,\big|\, w_j = 0\big\}$$

for $T \subset \mathbb{R}P^n \times \mathbb{R}^{n+1}$ — where for the sake of convenience the j-th coordinate exceptionally is listed first.

3.15 Lemma and Definition Let $E \xrightarrow{\pi} X$ be a vector bundle on a manifold X, and let $S \subset E$ be a subbundle. Then there is a natural vector bundle $E/S \to X$ with $(E/S)_x = E_x/S_x$ for all $x \in X$, called the *quotient bundle* of E by S. It fits into the equally natural exact sequence[2]

$$0 \longrightarrow S \longrightarrow E \longrightarrow E/S \longrightarrow 0$$

[2] Exactness at any position of a sequence requires that the image of the incoming arrow coincides with the kernel of the outgoing one. In the case at hand exactness is fibre-wise and simply restates the injectivity of the inclusion $S \subset E$, the surjectivity of the quotient homomorphism, and the fact that the latter's kernel is precisely S.

of vector bundles.

Proof As a set, of course, $E/S = \bigcup_{x\in X} E_x/S_x$ is the disjoint union of the quotient vector spaces, and we let $\rho\colon E/S \to X$ denote the projection. We choose an atlas indexed by $\lambda \in \Lambda$ say, of subbundle charts

$$\pi^{-1}U_\lambda \xrightarrow{h_\lambda} U_\lambda \times \mathbb{R}^d = U_\lambda \times \mathbb{R}^s \times \mathbb{R}^{d-s}$$

for the subbundle $S \subset E$. For each $\lambda \in \Lambda$ we define a chart of the quotient bundle by

$$\begin{array}{ccc} \rho^{-1}U_\lambda & \xrightarrow{k_\lambda} & U_\lambda \times \mathbb{R}^{d-s} \\ [v] & \longmapsto & \big(x, (\mathrm{pr}_3 \circ h_\lambda)\cdot v\big), \end{array}$$

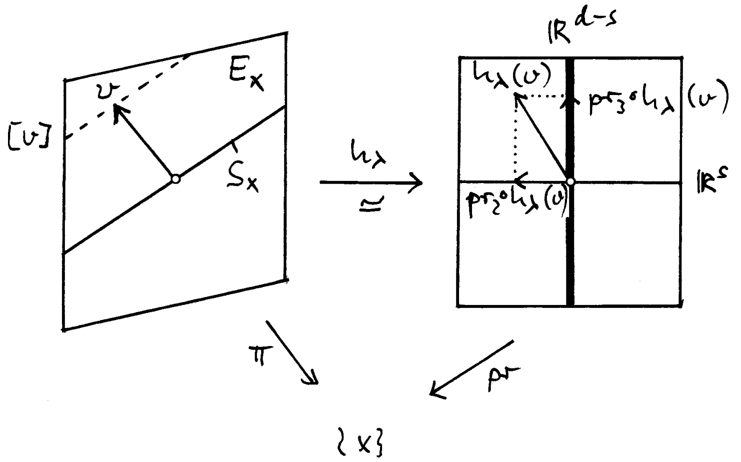

with $x := \pi([v])$, noting that the last component of h_λ does not depend on how the representative v is chosen in its congruence class modulo S_x. If $\mu \in \Lambda$ is another index then of course $(U_\lambda \cap U_\mu)\times\mathbb{R}^{d-s}$ is open in both $U_\lambda\times\mathbb{R}^{d-s}$ and $U_\mu\times\mathbb{R}^{d-s}$, and it inherits its natural topology from either. Since the transition map

$$\begin{array}{ccc} (U_\lambda \cap U_\mu) \times \mathbb{R}^{d-s} & \xrightarrow{k_\mu \circ k_\lambda^{-1}} & (U_\lambda \cap U_\mu) \times \mathbb{R}^{d-s} \\ (x, w) & \longmapsto & \big(x, (\mathrm{pr}_3 \circ h_\mu \circ h_\lambda^{-1})\cdot(x,0,w)\big) \end{array}$$

is a diffeomorphism the charts k_λ put well-defined topological and smooth bundle structures on E/S as desired. □

Given a homomorphism of vector bundles $E \xrightarrow{f} F$ on X one may form its kernel and its image fibrewise, taking kernel and image of the linear mapping $f_x\colon E_x \to F_x$ for each $x \in X$. In general the subsets of E and F thus defined will fail to be subbundles for the simple reason that the rank of f_x usually is a non-constant function of $x \in X$. Those familiar with the notion thus recognise that the category of vector bundles over X is not an abelian category. On the other hand the question of constant or non-constant rank turns out to be just the essential point:

3.16 Theorem Let E, F be vector bundles over X, and let $f\colon E \to F$ be a bundle homomorphism of constant rank r. Then the fibre-wise defined subsets

$$\mathrm{kernel}\, f \subset E \quad \text{and} \quad \mathrm{image}\, f \subset F$$

are subbundles.

Proof The question being local in X we may assume that $E = X\times\mathbb{R}^n$ and $F = X\times\mathbb{R}^p$ are product bundles, so that f has the form

$$X\times\mathbb{R}^n \ni (x,v) \longmapsto (x, f_x\cdot v) \in X\times\mathbb{R}^p$$

with a smooth matrix valued function $x \mapsto f_x$.

We first assume $n = p$ and consider a point $o \in X$. By elementary linear algebra we may assume that f_o has the standard form

$$f_o = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} \in \mathrm{Mat}(n\times n, \mathbb{R})$$

with an $r\times r$ unit matrix 1. We put

$$g = \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix} = \mathrm{id} - f_o \in \mathrm{Mat}(n\times n, \mathbb{R}).$$

The bundle homomorphism $h\colon (x, v) \longmapsto \big(x, (f_x + g) \cdot v\big)$ then acts as the identity at $x = o$, and reducing X to an open neighbourhood of o we may assume that h is a bundle automorphism of $X\times\mathbb{R}^n$. Even more, h is a subbundle chart for $\operatorname{kernel} f \subset X\times\mathbb{R}^n$ since for every $x \in X$ the inclusion

$$h_x(\operatorname{kernel} f_x) = (f_x+g)(\operatorname{kernel} f_x) = g(\operatorname{kernel} f_x) \subset \{0\}\times\mathbb{R}^{n-r}$$

must be an equality of vector spaces as both sides have the same dimension.

Likewise h^{-1} is a subbundle chart for $\operatorname{image} f \subset X\times\mathbb{R}^n$: we have

$$h_x\big(\mathbb{R}^r\times\{0\}\big) = (f_x+g)(\mathbb{R}^r\times\{0\}\big) = f_x(\mathbb{R}^r\times\{0\}\big) \subset \operatorname{image} f_x,$$

hence

$$\mathbb{R}^r\times\{0\} \subset h_x^{-1} \operatorname{image} f_x,$$

which again is an equality since the dimension on either side is the same for every $x \in X$. This completes the proof in the special case $n = p$.

We now assume that $n < p$. Recalling that $E = X\times\mathbb{R}^n$ is the product bundle we pick any linear surjection $\mathbb{R}^p \to \mathbb{R}^n$ and obtain a homomorphism between bundles of common rank p

$$X\times\mathbb{R}^p \longrightarrow X\times\mathbb{R}^n = E \xrightarrow{f} F$$

with the same image as f: by what we already know $\operatorname{image} f \subset F$ is a subbundle. Similarly for $n > p$ we embed $\mathbb{R}^p$ in $\mathbb{R}^n$ as a linear subspace; the composition

$$E \xrightarrow{f} F = X\times\mathbb{R}^p \longrightarrow X\times\mathbb{R}^n$$

has the same kernel as f, which therefore is a subbundle $\operatorname{kernel} f \subset E$.

Returning to the assumption $n < p$ we now rewrite f as the bundle homomorphism

$$E \xrightarrow{f} \operatorname{image} f$$

with the rank of E at least that of $\operatorname{image} f$, and conclude that its kernel, that is $\operatorname{kernel} f \subset E$ is a subbundle. Finally for $n > p$ we read the image of f as that of the injective homomorphism

$$E/\operatorname{kernel} f \longrightarrow F;$$

the latter has the same image as f, and this shows that $\operatorname{image} f \subset F$ is a subbundle in this case too.
□

Exercises

3.1 Verify the naturality of the pull-back with respect to bundle homomorphisms $h: F \to F'$ as stated in Theorem 3.5.

3.2 Let $E \xrightarrow{\pi} X$ be a vector bundle of rank $d \geq 2$. Prove that the complement of the zero section $E \setminus X$ is connected if and only if X is connected.

3.3 Prove that for $n > 0$ the tautological bundle $T \xrightarrow{\pi} \mathbb{R}P^n$ is not trivial.

3.4 Let X be a smooth manifold and $S \subset X$ a submanifold. Prove that $TS \subset TX|S$ is a subbundle.

3.5 Let X and Y be manifolds. Describe the tangent bundle $T(X \times Y)$ in terms of TX and TY using only vector bundle operations.

3.6 Let $E \xrightarrow{\pi} X$ be a vector bundle. Since the total space E is a manifold it has a tangent bundle $TE \to E$. Prove that there is a canonical exact sequence

$$0 \longrightarrow \pi^*E \longrightarrow TE \longrightarrow \pi^*TX \longrightarrow 0$$

of vector bundles on E.

3.7 Explain why the symmetric matrices with n distinct eigenvalues form an open subset X of the space $\mathrm{Sym}(n, \mathbb{R})$ of all symmetric real matrices of size n. Prove that the eigenspace decomposition represents the product bundle $X \times \mathbb{R}^n \to X$ as the Whitney sum $L_1 \oplus \cdots \oplus L_n$ of n line bundles.

3.8 Prove that for $n > 1$ the line bundles L_j of the previous problem are not trivial.

3.9 Let $T \xrightarrow{\pi} \mathbb{R}P^n$ be the tautological, and $V = \mathbb{R}P^n \times \mathbb{R}^{n+1} \xrightarrow{\mathrm{pr}} \mathbb{R}P^n$ the product bundle. Prove the isomorphy

$$T(\mathbb{R}P^n) \simeq \mathrm{Hom}\big(T, V/T\big)$$

of vector bundles on $\mathbb{R}P^n$.

4 Smooth Functions and Sections

4.1 Definition Let $s: X \to V$ be a function from a toplogical space to a finite dimensional real vector space, or alternatively a section of a vector bundle on X. The *support* of s

$$\operatorname{supp} s := \overline{\{x \in X \mid s(x) \neq 0\}} \subset X$$

is defined as the smallest closed subset with the property that s vanishes identically on the complement $X \setminus \operatorname{supp} s$.

4.2 Table Mountains From the smooth function $f: \mathbb{R} \to \mathbb{R}$,

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

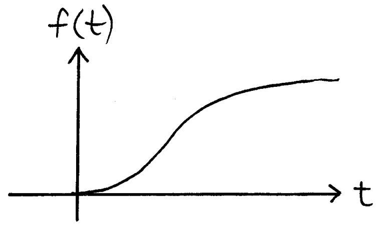

we construct functions $g: \mathbb{R} \to \mathbb{R}$ and $h: \mathbb{R} \to \mathbb{R}$,

$$g(t) = f(t) \cdot f(1-t) \quad \text{and} \quad h(t) = \int_0^t g(\tau) d\tau,$$

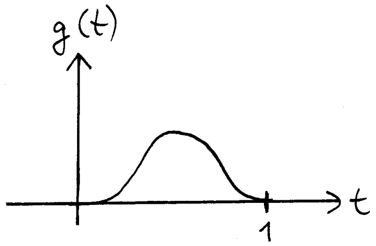

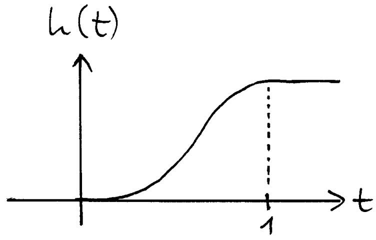

and normalise the latter to obtain $k: \mathbb{R} \longrightarrow [0, 1]$,

$$k(t) = \frac{1}{h(1)} \cdot h(t).$$

All these functions are smooth, and k vanishes identically on $(-\infty, 0]$, takes the constant value 1 on $[1, \infty)$, and strictly increases in between. Finally, depending on a dimension $n \in \mathbb{N}$ and real parameters $0 < r < R$ we define the *table mountain* $\tau: \mathbb{R}^n \to [0, 1]$ by

$$\tau(x) = h\left(\frac{R - |x|}{R - r}\right).$$

The table mountain likewise is smooth, has constant value 1 on the closed ball $D_r(0) \subset \mathbb{R}^n$ of radius r while it vanishes identically outside the open ball $U_R(0)$ of radius R.

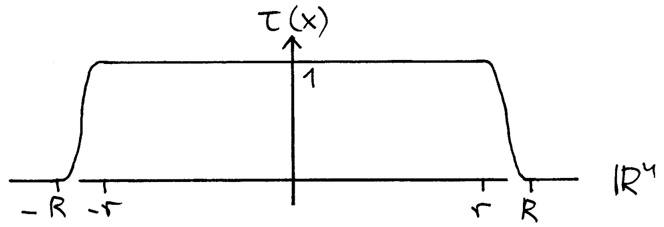

4.3 First Applications (1) Let X be an n-manifold and $o \in X$ a point. Choose any centred chart (U, h) at o and further choose $0 < r < R$ such that $D_R(0) \subset h(U)$. Then not only U but also $X \smallsetminus h^{-1}D_R(0)$ is an open subset of X since the compact set $h^{-1}D_R(0)$ must be closed in the Hausdorff space X. Therefore the formula

$$\sigma(x) = \begin{cases} \tau \circ h(x) & \text{for } x \in U, \text{ and} \\ 0 & \text{for } x \in X \smallsetminus h^{-1}D_R(0) \end{cases}$$

determines a well-defined C^∞-function $\sigma: X \to [0, 1]$. While σ has constant value 1 in a neighbourhood of o its support $h^{-1}D_R(0)$ is contained in U — in particular it may be forced to be arbitrarily small by a proper choice of U. We thus have implanted a table mountain on X which is concentrated near the point o.

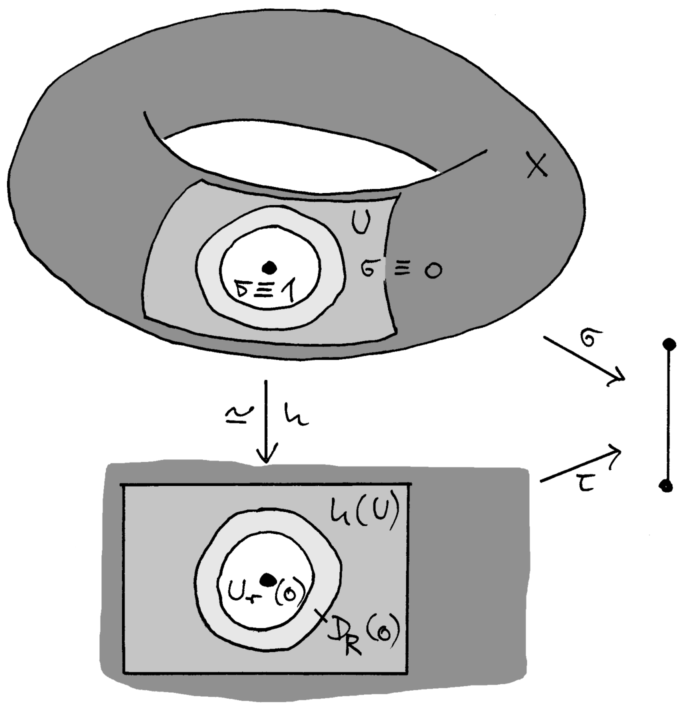

(2) To continue the first example let $f: U \to \mathbb{R}$ be a smooth function — which may, for instance, be obtained from any smooth real valued function on $h(U) \subset \mathbb{R}^n$ by composition with h. Then putting

$$F(x) = \begin{cases} \big(\tau \circ h(x)\big) \cdot f(x) & \text{for } x \in U, \text{ and} \\ 0 & \text{for } x \in X \smallsetminus h^{-1}D_R(0) \end{cases}$$

we obtain a global smooth function $F: X \to \mathbb{R}$ which coincides with f in a neighbourhood of o.

Note that it is certainly not possible in general to extend $f: U \to \mathbb{R}$ to a smooth or even continuous global function on X. But as this example shows extension is possible if you first allow U to be shrunk to a smaller neighbourhood of o.

Our first "serious" application of table mountains concerns the question of whether an abstract manifold can be realised as a submanifold of a Euclidean space.

4.4 Embedding Theorem Let X be a compact n-manifold. There exist a $p \in \mathbb{N}$ and an embedding $e: X \to \mathbb{R}^p$ of X as a smooth submanifold.

Proof As a first step we implant a table mountain at every point $o \in X$: we choose a chart (U_o, h_o) centred at o and compose it with a dilatation of $\mathbb{R}^n$ to make sure that $h_o(U_o) \subset \mathbb{R}^n$ is large enough so that $D_2(0) \subset h_o(U_o)$. As in 4.3(1) this allows to implant the table mountain function $\tau: \mathbb{R}^n \to [0,1]$ with uniform parameters $r = 1$ and $R = 2$ for all $o \in X$.

Since X is compact we find a finite set $\Lambda \subset X$ such that $X = \bigcup_{o\in\Lambda} h_o^{-1}U_1(0)$. We define the smooth mapping

$$e: X \longrightarrow \prod_{o\in\Lambda} (\mathbb{R} \oplus \mathbb{R}^n)$$

putting for every $x \in X$

$$e(x)_o = \begin{cases} (\tau \circ h_o)(x) \cdot \begin{pmatrix} 1 \\ h_o(x) \end{pmatrix} \in \mathbb{R} \oplus \mathbb{R}^n & \text{if } x \in U_o; \\ 0 & \text{if } x \in X \smallsetminus h_o^{-1}D_2(0). \end{cases}$$

We prove that e is injective. Thus let $x, y \in X$ be points with $e(x) = e(y)$. Since the sets $h_o^{-1}U_1(0)$ cover X there must exist an $o \in \Lambda$ with $(\tau \circ h_o)(y) = 1$. In particular we must have $y \in U_o$, and looking at the first component of e_o we see that then also

$$e(x)_o = e(y)_o = (\tau \circ h_o)(y) = 1$$

and therefore $x \in U_o$ and $(\tau \circ h_o)(x) = 1$. We now read off from the remaining components of e_o that

$$h_o(x) = (\tau \circ h_o)(x) \cdot h_o(x) = (\tau \circ h_o)(y) \cdot h_o(y) = h_o(y),$$

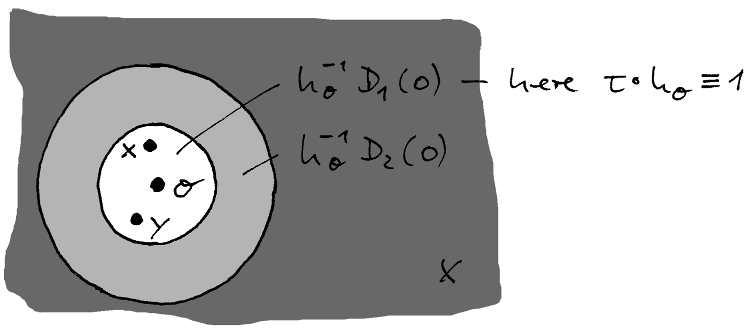

and since h_o is injective we conclude that $x = y$. Thus e is injective and furthermore a topological embedding since X is compact.

We now prove that e is an immersion, by showing that at every point $x \in X$ the rank of e is n — that is, no less than n. To this end we choose $o \in \Lambda$ such that $x \in h_o^{-1}U_1(0)$; the last n components of $e(x)_o$ reduce to $h_o(x)$, and we read off that

$$\mathrm{rk}_x\, e \geq \mathrm{rk}_x\, e(x)_o = \mathrm{rk}_x\, h_o = n.$$

It but remains to apply Theorem 2.6 in order to conclude the proof. □

Note Using finer techniques one can prove embedding theorems that give much more precise information than our rather crude version. To begin with it turns out that the compactness assumption is quite inessential. Other improvements concern the dimension of the embedding Euclidean space. Hassler Whitney proved in 1936 that every differentiable n-manifold admits an embedding as a closed differentiable submanifold of $\mathbb{R}^{2n+1}$ — in fact of $\mathbb{R}^{2n}$ by a later development. Of course special n-manifolds like the n-sphere may allow embeddings in $\mathbb{R}^p$ for p smaller than $2n$. On the other hand methods of algebraic topology provide non-embedding theorems that for certain classes of n-manifolds give interesting lower bounds for the possible embedding dimensions p.

To make available the full power of table mountains we must make a digression into point set topology. It concerns locally compact Hausdorff spaces that admit a countable base, and thus in particular applies to manifolds. It comprises two lemmata.

4.5 Lemma Let X be a locally compact Hausdorff space which admits a countable base. Then X is the union of a suitable sequence of compact subsets.

Proof We choose a countable base $(U_\lambda)_{\lambda\in\Lambda}$ of the topology of X and put

$$\Lambda' = \left\{\lambda \in \Lambda \,\middle|\, \overline{U_\lambda} \text{ is compact}\right\}.$$

Then $\bigcup_{\lambda\in\Lambda'} U_\lambda = X$. For let $x \in X$ be arbitrary; since X is locally compact we find an open neighbourhood U of x such that $\overline{U}$ is compact. The open set U must have a representation as a union of certain U_λ. As these must necessarily have compact closure only indices $\lambda \in \Lambda'$ are possible. In particular this implies $x \in U_\lambda$ for some $\lambda \in \Lambda'$.

Since Λ' is at most countable it but remains to arrange the at most countable family $(\overline{U_\lambda})_{\lambda\in\Lambda'}$ into a sequence (adding empty sets in the finite case). □

4.6 Lemma Let X be a locally compact Hausdorff space which admits a countable base. There exists a sequence $(K_j)_{j=0}^\infty$ of compact subspaces of X such that

$$K_j \subset K_{j+1}^\circ \text{ for all } j \in \mathbb{N}, \text{ and } \bigcup_{j=0}^\infty K_j = X.$$

Proof Let $(X_j)_{j=0}^\infty$ be a sequence as provided by Lemma 4.5. By induction on $j \in \mathbb{N}$ we will construct compact subsets $K_j \subset X$ with $K_j \subset K_{j+1}^\circ$ and $\bigcup_{i=0}^{j-1} X_i \subset K_j$.

Thus let $j \in \mathbb{N}$ be given and assume $K_0, \ldots, K_{j-1}$ constructed. Using that X is locally compact we choose a cover $(U_\lambda)_{\lambda\in\Lambda}$ of the subset $K_{j-1} \cup X_{j-1} \subset X$ consisting of open sets $U_\lambda \subset X$ with compact closure:

$$K_{j-1} \cup X_{j-1} \subset \bigcup_{\lambda\in\Lambda} U_\lambda \subset X.$$

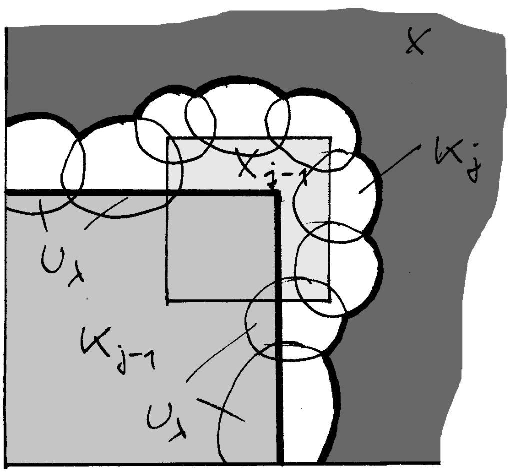

As $K_{j-1} \cup X_{j-1}$ is compact we can choose Λ as a finite set, and then $K_j := \bigcup_{\lambda\in\Lambda} \overline{U_\lambda}$ does the job. The lemma now follows at once. □

4.7 Definition Let X be a smooth manifold, and $(V_\mu)_{\mu\in\mathrm{M}}$ an open cover of X. A (*smooth*) *partition of unity subordinate* to $(V_\mu)_{\mu\in\mathrm{M}}$ is a sequence $(\rho_j)_{j=0}^\infty$ of smooth functions

$$\rho_j\colon X \longrightarrow [0,1]$$

such that

- for each $j\in\mathbb{N}$ the support $\operatorname{supp}\rho_j$ is compact,
- these supports form a locally finite family of subsets $(\operatorname{supp}\rho_j)_{j=0}^\infty$ of X, and
- the sum $\sum_{j=0}^\infty \rho_j = 1$ is the constant function;
- finally for each $j\in\mathbb{N}$ there exists some $\mu\in\mathrm{M}$ such that $\operatorname{supp}\rho_j \subset V_\mu$.

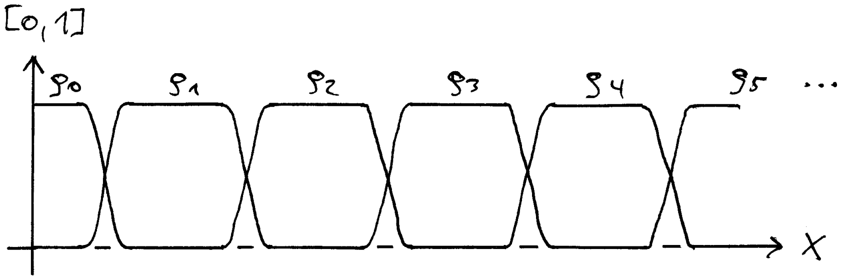

Notes Local finiteness means that every $x\in X$ has a neighbourhood which meets $\operatorname{supp}\rho_j$ for but finitely many $j\in\mathbb{N}$. This property guarantees that the infinite sum $\sum_j \rho_j$ not only makes sense but also that it defines a smooth function on X. Since the latter is positive everywhere the supports must cover X. — If X happens to be compact then ρ_j must vanish identically for all but finitely many j, and vice versa, see Exercise 4.1.

4.8 Theorem For every open cover of a manifold there exists a subordinate partition of unity.

Proof We let $(V_\mu)_{\mu\in\mathrm{M}}$ be the given open cover, fix a sequence $(K_j)_{j=0}^\infty$ as provided by Lemma 4.6, and put $K_{-1}=\emptyset$.

Consider any $j\in\mathbb{N}$. For every $o\in K_{j+1}\smallsetminus K_j^\circ$ we choose a centred chart

$$U_{jo} \xrightarrow{h_{jo}} U_3(0) \subset \mathbb{R}^{\dim X}$$

at o with $U_{jo}\subset K_{j+2}^\circ\smallsetminus K_{j-1}$ and $U_{jo}\subset V_\mu$ for some $\mu\in\mathrm{M}$:

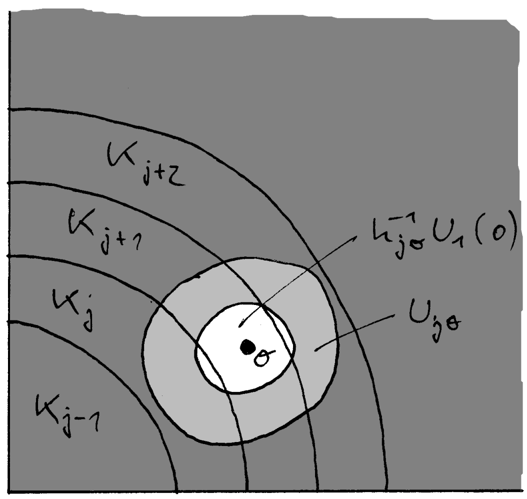

the condition $h_{jo}(U_{jo}) = U_3(0)$ is easily satisfied by blowing up the image of h_{jo} and restricting. The family

$$\big(h_{jo}^{-1} U_1(0)\big)_{o\in K_{j+1}\smallsetminus K_j^\circ}$$

is an open cover of the compact subset $K_{j+1}\smallsetminus K_j^\circ \subset X$, and we find a finite subcover indexed by $\Lambda_j \subset K_{j+1}\smallsetminus K_j^\circ$ say.

Throwing together all Λ_j we obtain the countable family of charts

$$\big(h_{jo}\colon U_{jo} \to U_3(0)\big)_{(j,o)\in\Lambda}$$

with $\Lambda := \sum_{j=0}^{\infty}\Lambda_j = \bigcup_{j=0}^{\infty}\{j\}\times\Lambda_j$. Its special properties include that

(1) even the smaller open sets $h_{jo}^{-1}U_1(0) \subset U_{jo}$ with $(j,o)\in\Lambda$ cover X,

(2) the U_{jo} themselves form a locally finite cover of X, and

(3) each U_{jo} is contained in V_μ for some $\mu\in\mathrm{M}$.

Properties (1) and (3) are clear by construction. To prove (2) consider any $x\in X$. Choose a $k\in\mathbb{N}$ with $x\in K_k^\circ$; then for $j>k$ all sets U_{jo} are disjoint from K_k, so that the open neighbourhood K_k° of x can meet U_{jo} only for the finitely many $(j,o)\in\Lambda$ with $j\le k$.

For each $\lambda = (j,o)\in\Lambda$ we now build the usual function $\sigma_\lambda\colon X\to[0,1]$ by

$$\sigma_\lambda(x) = \begin{cases} \tau\circ h_\lambda(x) & \text{for } x\in U_\lambda, \text{ and}\\ 0 & \text{for } x\in X\smallsetminus h_\lambda^{-1}D_2(0)\end{cases}$$

where τ is the table mountain with parameters $r=1$ and $R=2$. In view of property (2) the sum

$$\sigma := \sum_{\lambda\in\Lambda}\sigma_\lambda$$

is a smooth function on X, which by property (1) is positive everywhere. Dividing by σ we obtain functions

$$\rho_\lambda := \frac{\sigma_\lambda}{\sigma}\colon X \longrightarrow [0,\infty)$$

with $\sum_{\lambda\in\Lambda}\rho_\lambda = 1$. Finally replacing the countable index set Λ by the set $\mathbb{N}$ we obtain a sequence of functions which is a partition of unity subordinate to $(V_\mu)_{\mu\in\mathrm{M}}$, as follows at once form the fact that the support of ρ_λ is contained in (indeed equal to) the compact set $h_\lambda^{-1}D_2(0)\subset U_\lambda$. □

Partitions of unity are the key tool that allows to construct global objects from local ingredients. Let us illustrate the method by a number of applications.

4.9 Theorem Let A and B be disjoint closed subsets of a manifold X. Then there exists a smooth function $f\colon X\to[0,1]$ such that $f|A=0$ and $f|B=1$.

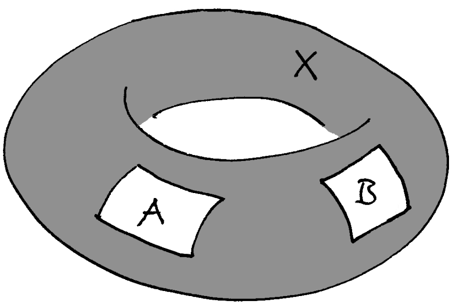

Proof The pair $(X\smallsetminus B, X\smallsetminus A)$ is an open cover of X, and we choose a subordinate partition of unity $(\rho_j)_{j=0}^{\infty}$. The function

$$f := \sum_{\operatorname{supp}\rho_j\subset X\smallsetminus A} \rho_j$$

solves the problem, for the ρ_j not included in the sum must have their support contained in $X \smallsetminus B$. □

One may note as a by-result that every manifold is a normal topological space.

4.10 Definition Let X and Y be manifolds and $A \subset X$ an arbitrary subset. A set mapping $f\colon A \to Y$ is said to be *smooth* at $a \in A$ if it can be locally extended to a smooth mapping, that is if there exist an open neighbourhood $U \subset X$ of a and a smooth map $g\colon U \to Y$ such that $f|(A \cap U) = g|(A \cap U)$.

4.11 Theorem Let A be a closed subset of a manifold X, and $f\colon A \to \mathbb{R}$ a smooth function. Then there exists a smooth extension $F\colon X \to \mathbb{R}$ of f.

Proof For every $a \in A$ we choose an open neighbourhood $V_a \subset X$ of a and a smooth function $g_a\colon V_a \to \mathbb{R}$ with $g_a|(A \cap V_a) = f|(A \cap V_a)$, while for every $a \in X\smallsetminus A$ we put $V_a = X\smallsetminus A$ and $g_a = 0$.

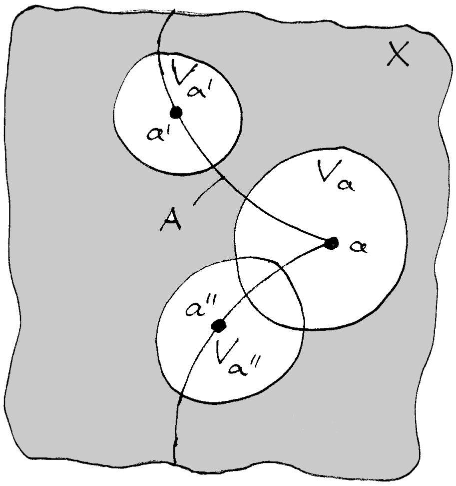

Then $(V_a)_{a\in X}$ is an open cover of X, and we find a partition of unity $(\rho_j)_{j=0}^{\infty}$ subordinate to it. Thus for every $j \in \mathbb{N}$ we can choose an index $a(j) \in X$ such that $\operatorname{supp}\rho_j \subset V_{a(j)}$, and we define the function $F\colon X \to \mathbb{R}$ by

$$F = \sum_{j=0}^{\infty} \rho_j \cdot g_{a(j)}$$

where $\rho_j \cdot g_{a(j)}$ is shorthand for

$$X \ni x \longmapsto \begin{cases} \rho_j(x) \cdot g_{a(j)}(x) & \text{if } x \in V_a \text{, and} \\ 0 & \text{if } x \in X \smallsetminus \operatorname{supp}\rho_j . \end{cases}$$

Let $x \in A$ be an arbitrary point. For every $j \in \mathbb{N}$ with $x \in \operatorname{supp}\rho_j$ we also have $x \in V_{a(j)}$ and $g_{a(j)}(x) = f(x)$. In view of $\sum_j \rho_j = 1$ this implies that $F(x) = f(x)$. Therefore F is an extension of f as claimed. □

The local objects that can be glued by a partition of unity need not be just scalar functions: the essential point is the presence of a linear or, more generally, a convex structure. Quite an important class of applications concerns sections of a vector bundle $E \to X$. Here a partition of unity on X allows to construct a global section $s\colon X \to E$ from a collection of *local sections*, that is sections of the restricted bundle $E|V$ over small open subsets V. The applications we give at this point are based on one particular kind of bundle.

4.12 Definition Let $E \xrightarrow{\pi} X$ be a vector bundle. A (smooth) section s of the bundle $\operatorname{Sym} E \to X$ of symmetric bilinear forms is called a *Riemannian metric* on E if for each $x \in X$ the symmetric bilinear form

$$s(x): E_x \times E_x \longrightarrow \mathbb{R}$$

is positive definite, thus turning each fibre E_x into a Euclidean vector space. If one particular Riemannian metric s is understood from the context the awkward notation $s(x)(v,w)$ for vectors $v, w \in E_x$ is usually replaced by the standard $\langle v, w\rangle$.

4.13 Theorem Let X be a manifold. Every vector bundle on X admits a Riemannian metric.

Proof In the case of a product bundle $X \times \mathbb{R}^d \to X$ the sections are the mappings

$$X \ni x \longmapsto \big(x, t(x)\big) \in X \times \operatorname{Sym}(d, \mathbb{R})$$

with a smooth map $t: X \to \operatorname{Sym}(d, \mathbb{R})$ into the space of symmetric matrices. Choosing for t the constant function with value $1 \in \operatorname{Sym}(d, \mathbb{R})$ we see that the conclusion certainly holds in this case, and thereby for every trivial bundle too.

Let now $E \to X$ be an arbitrary bundle. By local triviality every $x \in X$ has an open neighbourhood $U_x \subset X$ such that $E|U_x$ is trivial, and we can therefore choose a Riemannian metric

$$s_x: U_x \longrightarrow \operatorname{Sym} E|U_x.$$

Let $(\rho_j)_{j=0}^{\infty}$ be a partition of unity subordinate to the open cover $(U_x)_{x \in X}$, and choose for every $j \in \mathbb{N}$ an $x(j) \in X$ with $\operatorname{supp} \rho_j \subset U_{x(j)}$. The argument which by now is standard shows that

$$s := \sum_{j=0}^{\infty} \rho_j \cdot s_{x(j)}$$

is a smooth section of E. In view of $\sum_j \rho_j = 1$, for every $x \in X$ the value $s(x) \in \operatorname{Sym} E_x$ is a convex linear combination of positive definite bilinear forms $s_{x(j)}$, and therefore itself positive definite. This proves that s is a Riemannian metric. □

4.14 Applications (1) Every bundle $E \to X$ is (non-canonically) isomorphic to its dual $E\check{}$: any choice of a Riemannian metric s on E defines an isomorphism, which relates the fibres over $x \in X$ by

$$E_x \ni v \longmapsto \big(E_x \ni w \mapsto \langle v, w\rangle \in \mathbb{R}\big) \in E_x\check{}.$$

(2) Let $E \xrightarrow{\pi} X$ be a vector bundle and $S \subset E$ a subbundle. Given a Riemannian metric on E the fibre-wise orthogonal complement $S^\perp \subset E$ is a subbundle of E. Indeed it is the kernel of the surjective bundle homomorphism

$$E \ni v \longmapsto \big(w \mapsto \langle v, w\rangle\big) \in S\check{}.$$

This subbundle is a fibre-wise linear complement to S:

$$E_x = S_x \oplus S_x^\perp \quad \text{for all } x \in X,$$

and thus gives a Whitney sum decomposition $E = S \oplus S^\perp$.

(3) Let X be a differential manifold and $S \subset X$ a submanifold. We know from Example 3.14(1) that $TS \subset TX|S$ is a subbundle, and thus have the quotient bundle

$$N(S \subset X) := (TX|S)/TS \longrightarrow S,$$

called the *normal bundle* of S in X. But often one would prefer a normal bundle which is a subbundle rather than a quotient bundle of $TX|S$: this can be achieved choosing a Riemannian metric on $TX|S$ and taking the isomorphic subbundle $TS^{\perp} \subset TX|S$ instead. Unlike $N(S \subset X)$ it is not determined by S and X alone, but depends on the choice of the Riemannian metric, and should be referred to as *a* normal bundle rather than *the* normal bundle.

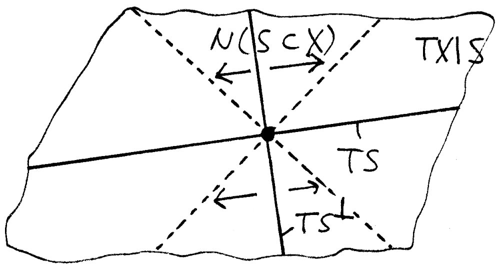

Whenever the context provides a natural choice of a Riemannian metric on TX — for instance if $X = \mathbb{R}^n$ — then one would of course prefer that one and dispose of an equally natural normal subbundle.

Exercises

4.1 Let X be a topological space and $(K_\lambda)_{\lambda\in\Lambda}$ a locally finite cover of X by compact subspaces. Prove the equivalence between the statements:

(1) For all but finitely many $\lambda \in \Lambda$ the set K_λ is empty.

(2) The space X ist compact.

4.2 Let X be a manifold. Show that every continuous function $f: X \to \mathbb{R}$ can be uniformly approximated by a smooth function: given f and $\varepsilon > 0$ there exists a smooth function $g: X \to \mathbb{R}$ such that $|f(x) - g(x)| < \varepsilon$ holds for all $x \in X$.

4.3 Let X be a manifold. Prove that there exists a proper smooth function $f: X \to \mathbb{R}$.

Recall or learn that a continuous map $f: X \to Y$ between locally compact Hausdorff spaces is called *proper* if for every compact $L \subset Y$ the pre-image $f^{-1}L \subset X$ also is compact.

4.4 Let $A \subset \mathbb{R}^n$ be a closed subset. Construct a smooth function $f: \mathbb{R}^n \to [0,1]$ with $A = f^{-1}\{0\}$.

4.5 Let $\sum_{k=0}^{\infty} a_k x^k$ be a power series with real cofficients. Prove that there exists a smooth function $f: \mathbb{R} \to \mathbb{R}$ whose Taylor series of f at 0 is the given series:

$$\frac{f^{(k)}(0)}{k!} = a_k \quad \text{for all } k \in \mathbb{N}.$$

Hint While the given series has no reason to converge you may force uniform convergence, even of the series of arbitrary order derivatives, by throwing in a factor $\tau(x/r_k)$ where τ is a fixed table mountain and $(r_k)_{k=0}^{\infty}$ a sequence of sufficiently small positive radii. To find such a sequence it is helpful to note that for every $k \in \mathbb{N}$ the set

$$\left\{ \left(\frac{d}{dx}\right)^j \Big(\tau(x) \cdot x^k\Big) \,\Big|\, j < k \text{ and } x \in \mathbb{R} \right\} \subset \mathbb{R}$$

is bounded.

4.6 Let $E \to X$ be a rank d vector bundle. Prove that the following statements are equivalent.

(1) There exists a section of E without zeros.

(2) There exists a trivial subbundle $L \subset E$ of rank one.

(3) There exists a Whitney sum decomposition $E = L \oplus E'$ of E into a trivial line bundle $L \subset E$ and a subbundle $E' \subset E$ of rank $d{-}1$.

5 Review of Linear Algebra

5.1 Definition Let V be a real vector space of finite dimension. A *multilinear form* of *degree* k on V, or multilinear *k-form* for short, is a function

$$\varphi: V^k = V\times\cdots\times V \longrightarrow \mathbb{R}$$

which becomes a linear function of each variable whenever the other $k-1$ variables are held fixed. Such a form is called *alternating* if it satisfies

$$\varphi(v_1,\ldots,v_k) = 0 \quad \text{if } v_i = v_j \text{ for some } i\neq j.$$

The sets of multilinear and of alternating forms clearly are real vector spaces $\operatorname{Alt}^k V \subset \operatorname{Mult}^k V$. Of course $\operatorname{Mult}^1 V = \operatorname{Alt}^1 V = V\check{}$ is just the dual vector space of linear forms on V.

5.2 Lemma Let $\varphi\in\operatorname{Alt}^k V$ be an alternating k-form, and let $v_1,\ldots,v_k\in V$ be vectors.

- Then $\varphi(v_1,\ldots,v_k) = 0$ whenever $(v_1,\ldots,v_k)$ is linearly dependent; in particular $\operatorname{Alt}^k V = \{0\}$ for all $k > \dim V$.
- If $f: \mathbb{R}v_1+\cdots+\mathbb{R}v_k \to \mathbb{R}v_1+\cdots+\mathbb{R}v_k$ is a linear endomorphism then

$$\varphi\big(f(v_1),\ldots,f(v_k)\big) = \det f\cdot\varphi(v_1,\ldots,v_k).$$

- For every permutation $\sigma\in\operatorname{Sym}_k$, if we write $(-1)^\sigma$ for the sign of σ we have

$$\varphi(v_{\sigma 1},\ldots,v_{\sigma k}) = (-1)^\sigma\varphi(v_1,\ldots,v_k).$$

Proof In the case of linear dependence we write one of the vectors, say v_j, as a linear combination of the others and apply multilinearity:

$$\varphi(v_1,\ldots,v_k) = \varphi\Big(v_1,\ldots,v_{j-1},\sum_{i\neq j}\lambda_i v_i, v_{j+1},\ldots,v_k\Big) = \sum_{i\neq j}\lambda_i\varphi(v_1,\ldots,v_{j-1},v_i,v_{j+1},\ldots,v_k).$$

The last expression has v_i at both the i-th and j-th positions, and therefore vanishes. This proves the first statement.

If in the second $(v_1,\ldots,v_k)$ happens to be linearly dependent, then this is equally true of the system of images $\big(f(v_1),\ldots,f(v_k)\big)$, and the claimed identity holds trivially since both sides vanish. We thus may assume that $(v_1,\ldots,v_k)$ is independent and therefore a base of the vector space $\mathbb{R}v_1+\cdots+\mathbb{R}v_k$. Via this base the endomorphisms of the latter correspond to matrices in $\operatorname{Mat}(k\times k,\mathbb{R})$, and if f_a denotes the endomorphism whose matrix is a then the function

$$\operatorname{Mat}(k\times k,\mathbb{R})\ni a\longmapsto\varphi\big(f_a(v_1),\ldots,f_a(v_k)\big)\in\mathbb{R}$$

is an alternating multilinear form in the columns of a. As is well known such a function must be proportional to the determinant function, thus there exists a $\lambda\in\mathbb{R}$ with

$$\varphi\big(f_a(v_1),\ldots,f_a(v_k)\big) = \lambda\cdot\det a \quad \text{for all } a\in\operatorname{Mat}(k\times k,\mathbb{R}).$$

Evaluating at $a = 1$, whence $f_a = \mathrm{id}$, we obtain $\lambda = \varphi(v_1,\ldots,v_k)$, and this concludes the proof of the second statement.

Applying it to the endomorphism f with $f(v_i) = v_{\sigma i}$ for all i we finally obtain the property stated as the third part of the lemma:

$$\varphi(v_{\sigma 1}, \ldots, v_{\sigma k}) = \varphi\big(f(v_1), \ldots, f(v_k)\big) = \det f \cdot \varphi(v_1, \ldots, v_k) = (-1)^\sigma \varphi(v_1, \ldots, v_k). \qquad \square$$

Note Conversely, if a multilinear form has this property for all permutations it clearly is alternating.

5.3 Notation Transformation of the multilinear form $\varphi \in \operatorname{Mult}^k V$ by $\sigma \in \operatorname{Sym}_k$ results in a new multilinear form $\sigma\varphi \in \operatorname{Mult}^k V$:

$$\sigma\varphi(v_1, \ldots, v_k) := (-1)^\sigma \varphi(v_{\sigma 1}, \ldots, v_{\sigma k}).$$

This defines a linear action of the symmetric group Sym_k on the vector space $\operatorname{Mult}^k V$:

$$1\varphi = \varphi \quad \text{and} \quad \tau(\sigma\varphi) = (\tau\sigma)\varphi \text{ for all } \sigma, \tau \in \operatorname{Sym}_k .$$

The alternating forms are just the fixed vectors of this action.

5.4 Definition and Lemma (1) Two multilinear forms $\varphi \in \operatorname{Mult}^k V$ and $\chi \in \operatorname{Mult}^l V$ may be multiplied putting

$$(\varphi\chi)(v_1, \ldots, v_k, w_1, \ldots, w_l) = \varphi(v_1, \ldots, v_k) \cdot \chi(w_1, \ldots, w_l);$$

then $\varphi\chi \in \operatorname{Mult}^{k+l} V$. Even if φ and χ are both alternating $\varphi\chi$ has no reason to be so.

(2) The *exterior* or *wedge* product of two alternating forms φ and χ is defined by

$$\varphi\wedge\chi = \frac{1}{k!\,l!} \sum_{\sigma\in\operatorname{Sym}_{k+l}} \sigma(\varphi\chi),$$

making it alternating by brute force; the bilinear products

$$\operatorname{Alt}^k V \times \operatorname{Alt}^l V \xrightarrow{\ \wedge\ } \operatorname{Alt}^{k+l} V$$

are associative and *graded commutative*:

$$\chi\wedge\varphi = (-1)^{kl} \varphi\wedge\chi \quad \text{if } \varphi \in \operatorname{Alt}^k V \text{ and } \chi \in \operatorname{Alt}^l V.$$

For $k = 0$, when $\operatorname{Alt}^k V = \mathbb{R}$, the product reduces to the scalar multiplication $\mathbb{R}\times\operatorname{Alt}^l V \to \operatorname{Alt}^l V$.

Note In the defining formula each term occurs $k!\,l!$ times in the sum — think of the scalar factor as eliminating this multiplicity rather than as a true denominator.

Proof In order to verify associativity we compute the products of three forms $\varphi \in \operatorname{Alt}^k V$, $\chi \in \operatorname{Alt}^l V$, and $\psi \in \operatorname{Alt}^m V$:

$$(\varphi\wedge\chi)\wedge\psi = \frac{1}{(k+l)!\,m!} \sum_\tau \tau\big((\varphi\wedge\chi)\psi\big) = \frac{1}{k!\,l!} \frac{1}{(k+l)!\,m!} \sum_{\sigma,\tau} \tau\big(\sigma(\varphi\chi)\psi\big)$$

where the sums are taken over all $\sigma \in \operatorname{Sym}_{k+l}$ and $\tau \in \operatorname{Sym}_{k+l+m}$. Considering $\operatorname{Sym}_{k+l} \subset \operatorname{Sym}_{k+l+m}$ as the subgroup that fixes the last m symbols and thereby does not affect ψ we re-write the equation as

$$(\varphi\wedge\chi)\wedge\psi = \frac{1}{k!\,l!} \frac{1}{(k+l)!\,m!} \sum_{\sigma,\tau} \tau\sigma(\varphi\chi\psi).$$

Now any fixed permutation $\rho \in \mathrm{Sym}_{k+l+m}$ is realised as a product $\tau\sigma$ in exactly $(k+l)!$ different ways, given by an arbitrary $\sigma \in \mathrm{Sym}_{k+l}$ and the $\tau = \rho\sigma^{-1} \in \mathrm{Sym}_{k+l+m}$ determined by it. Collecting equal summands we thus obtain the simplified equation

$$(\varphi\wedge\chi)\wedge\psi = \frac{1}{k!\,l!\,m!}\sum_{\rho}\rho(\varphi\chi\psi).$$

Evaluating $\varphi\wedge(\chi\wedge\psi)$ we would obtain exactly the same, and this proves the associative law.

The proof of commutativity relies on the special permutation $\rho \in \mathrm{Sym}_{k+l}$ which shifts the first l symbols past the last k without changing their order. Indeed, evaluating on vectors $v_i, w_j \in V$ we obtain

$$\begin{aligned}\rho(\chi\varphi)(v_1,\dots,v_k,w_1,\dots,w_l) &= (-1)^\rho(\chi\varphi)(w_1,\dots,w_l,v_1,\dots,v_k)\\ &= (-1)^\rho\chi(w_1,\dots,w_l)\cdot\varphi(v_1,\dots,v_k)\\ &= (-1)^\rho\varphi(v_1,\dots,v_k)\cdot\chi(w_1,\dots,w_l)\\ &= (-1)^\rho(\varphi\chi)(v_1,\dots,v_k,w_1,\dots,w_l),\end{aligned}$$

and, writing $\sigma = \tau\rho$,

$$\chi\wedge\varphi = \frac{1}{k!\,l!}\sum_{\sigma\in\mathrm{Sym}_{k+l}}\sigma(\chi\varphi) = (-1)^\rho\frac{1}{k!\,l!}\sum_{\tau\in\mathrm{Sym}_{k+l}}\tau(\varphi\chi) = (-1)^{kl}\varphi\wedge\chi. \qquad\square$$

The expression for the wedge product of three alternating forms established in the last proof easily generalises to an arbitrary number of factors, and is useful in its own right:

5.5 Formula The product of r alternating forms φ_j of degree k_j is

$$\varphi_1\wedge\varphi_2\wedge\cdots\wedge\varphi_r = \frac{1}{k_1!\,k_2!\cdots k_r!}\sum_{\sigma}\sigma(\varphi_1\varphi_2\cdots\varphi_r)$$

where σ runs over the group $\mathrm{Sym}_{k_1+\cdots+k_r}$. In particular

$$(\varphi_1\wedge\cdots\wedge\varphi_k)(v_1,\dots,v_k) = \det\Big(\varphi_i(v_j)\Big)_{i,j=1}^{k}$$

computes the product of k linear forms φ_i.

5.6 Example The first non-trivial case is $(\varphi\wedge\chi)(v,w) = \varphi(v)\chi(w) - \varphi(w)\chi(v)$.

5.7 Notation Every linear map $f\colon V\to W$ into another vector space of finite dimension induces linear maps

$$\mathrm{Mult}^k f\colon \mathrm{Mult}^k W \longrightarrow \mathrm{Mult}^k V \quad\text{and}\quad \mathrm{Alt}^k f\colon \mathrm{Alt}^k W \longrightarrow \mathrm{Alt}^k V$$

acting by

$$\chi \longmapsto \Big((v_1,\dots,v_k)\mapsto \chi\big(f(v_1),\dots,f(v_k)\big)\Big).$$

Both are usually abbreviated by f^*, where the upper position of the star reminds of the fact that the direction of f is reversed. In any case the chain rule

$$\mathrm{id}^* = \mathrm{id},\quad (g\circ f)^* = f^*\circ g^*$$

is evident, and f^* is compatible with the wedge product of alternating forms:

$$f^*(\varphi\wedge\chi) = f^*(\varphi)\wedge f^*(\chi).$$

5.8 Bases Let $\underline{b} = (\beta_1, \dots, \beta_n)$ be a base of $\mathrm{Alt}^1 V = V^\vee$. The wedge products

$$\beta_{j_1} \wedge \beta_{j_2} \wedge \dots \wedge \beta_{j_k}$$

with $1 \le j_1 < j_2 < \dots < j_k \le n$ form a base of the vector space $\mathrm{Alt}^k V$. In particular $\dim \mathrm{Alt}^k V = \binom{n}{k}$, and $\dim \bigoplus_{k=0}^{\infty} \mathrm{Alt}^k V = 2^n$.

Let $\underline{c} = (\gamma_1, \dots, \gamma_p)$ be a base of $W^\vee$ and $f: V \to W$ a linear mapping. In terms of the bases of V and W to which $\underline{b}$ and $\underline{c}$ are dual f is expressed by a matrix $a \in \mathrm{Mat}(p \times n, \mathbb{R})$. Then f^* acts on base forms by the determinants of submatrices of a:

$$f^*(\gamma_{i_1} \wedge \dots \wedge \gamma_{i_k}) = \sum \det \Big(a_{i_r j_s} \Big)_{r,s=1}^{k} \beta_{j_1} \wedge \dots \wedge \beta_{j_k}$$

where the sum is taken over all k-tuples $(j_1, \dots, j_k)$ with $1 \le j_1 < j_2 < \dots < j_k \le n$.

Proof Let $(v_1, \dots, v_n)$ denote the base of V dual to $\underline{b}$. A form $\varphi \in \mathrm{Mult}^k V$ clearly is determined by its values on all k-tuples of base vectors, and if φ alternates then even by its values on the k-tuples $(v_{j_1}, \dots, v_{j_k})$ with strictly increasing indices $j_1 < \dots < j_k$. This shows that the dimension of $\mathrm{Alt}^k V$ is at most $\binom{n}{k}$. On the other hand the k-fold wedge products $\beta_{j_1} \wedge \dots \wedge \beta_{j_k}$ for such j_s are at once seen to form a linearly independent system and therefore a base of $\mathrm{Alt}^k V$.

The matrix a of the linear map $f: V \to W$ has the coefficients $a_{ij} = (\gamma_i \circ f)(v_j) = (f^*\gamma_i)(v_j)$ by definition. Therefore we have the pull-back formula

$$f^*\gamma_i = \sum_{j=1}^{n} a_{ij}\, \beta_j,$$

which generalises to arbitrary degree via 5.5:

$$\begin{aligned} f^*(\gamma_{i_1} \wedge \dots \wedge \gamma_{i_k})(v_{j_1}, \dots, v_{j_k}) &= (f^*\gamma_{i_1} \wedge \dots \wedge f^*\gamma_{i_k})(v_{j_1}, \dots, v_{j_k}) \\ &= \det \Big((f^*\gamma_{i_r})(v_{j_s}) \Big)_{r,s=1}^{k} \\ &= \det \Big(a_{i_r, j_s} \Big)_{r,s=1}^{k} \end{aligned}$$

or

$$f^*(\gamma_{i_1} \wedge \dots \wedge \gamma_{i_k}) = \sum_{j_1 < \dots < j_k} \det \Big(a_{i_r j_s} \Big)_{r,s=1}^{k} \beta_{j_1} \wedge \dots \wedge \beta_{j_k}$$ □

5.9 Orientations Let V be a real vector space of dimension n. We consider the set of bases of V, and declare two bases equivalent if their transition map $\mathbb{R}^n \xrightarrow{\simeq} \mathbb{R}^n$ has positive determinant. Assuming $n > 0$ there are exactly two equivalence classes which are called the *orientations* of V. An *oriented vector space* is, of course, a finite dimensional real vector space together with one of its orientations. A linear isomorphism $f: V \simeq W$ between oriented vector spaces either *preserves* or *reverses* orientations. To determine which case occurs one can represent f by its matrix with respect to bases which are selected arbitrarily from the orientations, and read the sign of the determinant.

For the sake of completeness and consistency it is useful to extend the notion of orientation to the exceptional case $n = 0$: an orientation of the zero space $V = \{0\}$ is just one of the real numbers ± 1 assigned to it, and the isomorphism $\{0\} \simeq \{0\}$ preserves orientations if and only if these signs agree.

5.10 Definition The non-zero alternating n-forms on the n-dimensional vector space V are called *volume forms* since they provide V with a notion of oriented volume. A base $(v_1, \dots, v_n)$ of V belongs to the orientation defined by the volume form ω if and only if $\omega(v_1, \dots, v_n) > 0$ or, equivalently, if the representation of ω in terms of the dual base $\underline{b} = (\beta_1, \dots, \beta_n)$ is

$$\omega = \lambda \cdot \beta_1 \wedge \dots \wedge \beta_n \quad \text{with } \lambda > 0.$$

The intuitive meaning of the absolute value $\big|\omega(v_1,\dots,v_n)\big|$ is the n-dimensional volume of the parallelepiped

$$[0,1]v_1+\cdots+[0,1]v_n\subset V.$$

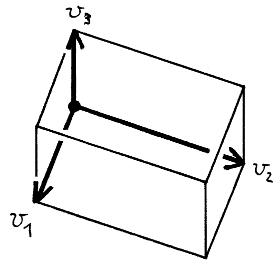

Exercises

5.1 Assume that (b_1,b_2,b_3,b_4) is a basis for the dual $V^\vee$ of some real vector space V. Do there exist linear forms $\varphi,\chi\in V^\vee$ such that

$$b_1\wedge b_2+b_3\wedge b_4=\varphi\wedge\chi\ ?$$

5.2 Let V be a (finite dimensional) real vector space and $k\in\mathbb{N}$. The formula

$$\sigma\varphi(v_1,\dots,v_k):=\varphi(v_{\sigma 1},\dots,v_{\sigma k})$$

(without the sign) defines another linear action of the symmetric group Sym_k on the vector space $\mathrm{Mult}^k V$, and the forms in the subspace

$$\mathrm{Sym}^k V:=\Big\{\varphi\in\mathrm{Mult}^k V\,\Big|\,\sigma\varphi=\varphi \text{ for all } \sigma\in\mathrm{Sym}_k\Big\}\subset\mathrm{Mult}^k V$$

are, of course, called the *symmetric* multilinear forms of degree k. Prove that the statement

$$\mathrm{Mult}^k V=\mathrm{Alt}^k V\oplus\mathrm{Sym}^k V$$

is true for $k=2$ but, assuming $\dim V\ge 2$, for no other $k\in\mathbb{N}$.

5.3 Let V be a real vector space of finite dimension, and let $\alpha\in V^\vee$ be a non-zero linear form. Prove that for every form $\chi\in\mathrm{Alt}^l V$ the following equivalence holds:

$$\alpha\wedge\chi=0\iff\text{there is a }\varphi\in\mathrm{Alt}^{l-1}V\text{ such that }\chi=\alpha\wedge\varphi$$

(put $\mathrm{Alt}^{-1}V:=\{0\}$ to include the case of $l=0$). If you are familiar with exact sequences you may prefer to read that the sequence

$$0\longrightarrow\mathrm{Alt}^0 V\xrightarrow{\alpha\wedge}\mathrm{Alt}^1 V\xrightarrow{\alpha\wedge}\ \cdots\ \xrightarrow{\alpha\wedge}\mathrm{Alt}^{l-1}V\xrightarrow{\alpha\wedge}\mathrm{Alt}^l V\xrightarrow{\alpha\wedge}\cdots$$

is exact, which means exactly the same.

5.4 Let V be an n-dimensional real vector space and fix a volume form $\omega \in \operatorname{Alt}^n V$. Prove that the assignment

$$h(v)(v_2, \ldots, v_n) := \omega(v, v_2, \ldots, v_n)$$

defines an isomorphism of vector spaces $h: V \simeq \operatorname{Alt}^{n-1} V$.

5.5 Let V be an oriented Euclidean vector space of dimension $n \in \mathbb{N}$.

- Prove that V carries a canonical volume form $\omega \in \operatorname{Alt}^n V$, characterized by the property that

$$\omega(u_1, \ldots, u_n) = 1$$

whenever $(u_1, \ldots, u_n)$ is a positively oriented orthonormal base of V.

- Let $(v_1, \ldots, v_n)$ be an arbitrary positively oriented base of V. Express ω in terms of the dual base $(\beta_1, \ldots, \beta_n)$ and the scalar products $\langle v_i, v_j \rangle$.

6 Differential Forms

We recall that natural constructions with vector spaces may be transferred to vector bundles, as described in 3.11. A case of particular interest is that of alternating forms on the tangent bundle of a smooth manifold.

6.1 Definition Let X be a differential manifold and $k \in \mathbb{N}$. A (smooth) section of the vector bundle $\operatorname{Alt}^k TX \xrightarrow{\pi} X$ is called a *differential form* of *degree* k on X, or simply a *k-form* on X. The set of all k-forms on X clearly is a real vector space denoted $\Omega^k X$ (usually of infinite dimension).

A k-form on X thus restricts, for every $x \in X$, to an alternating k-form $\varphi_x : T_x X \times \cdots \times T_x X \to \mathbb{R}$ on the tangent space $T_x X$. The exterior product, applied fibrewise, induces the exterior product of differential forms

$$\Omega^k X \times \Omega^l X \xrightarrow{\wedge} \Omega^{k+l} X,$$

which inherits its properties: bilinearity, associativity, and graded commutativity.

If $X \xrightarrow{f} Y$ is a smooth mapping then every differential form $\chi \in \Omega^k Y$ may be *pulled back* to the form $f^*\chi \in \Omega^k X$ with $(f^*\chi)_x = (T_x f)^* \chi_{f(x)}$. In the case that $f : X \subset Y$ is the inclusion of a submanifold this form is written $f^*\chi = \chi|X$ and called the *restriction* of χ to X.

Notes All these constructions are *pointwise* with respect to X: in order to know $(\varphi \wedge \chi)_x$ all you need know about φ and χ is $\varphi_x \in \operatorname{Alt}^k T_x X$ and $\chi_x \in \operatorname{Alt}^l T_x X$. Likewise the pull-back form $(f^*\chi)_x \in \operatorname{Alt}^k T_x X$ depends only on $\chi_{f(x)} \in \operatorname{Alt}^k T_{f(x)} Y$ rather than χ as a whole. — For formal completeness it is sometimes useful to put $\Omega^{-1} X = \{0\}$ and thus admit the trivial differential form of degree -1.

6.2 Basic Special Cases (0) A 0-form on X is the same as a smooth function $X \to \mathbb{R}$. Taking the wedge product with $f \in \Omega^0 X$ just multiplies the values of an arbitrary form $\chi \in \Omega^l X$ point by point with those of f.

(1) A 1-form on X is a smooth function $TX \to \mathbb{R}$ which is linear on each fibre. Every 0-form $f : X \to \mathbb{R}$ determines a particular 1-form df which composes Tf with the projection of the product bundle $T\mathbb{R}$ to its fibre:

$$df : TX \xrightarrow{Tf} T\mathbb{R} = \mathbb{R} \times \mathbb{R} \xrightarrow{\mathrm{pr}_2} \mathbb{R}.$$

As df carries the same information as Tf it is often taken as an alternative interpretation of the differential of f and referred to as such.

(2) Let (U, h) be a chart for the n-manifold X. The components $h_1, \ldots, h_n$ are smooth functions on the manifold U so that we may consider their differentials $dh_1, \ldots, dh_n \in \Omega^1 U$. The value of dh_j on the tangent vector $v \in T_x X$ is just the j-th component of $v_h \in \mathbb{R}^n$, so that at each point $x \in U$ the forms $dh_1, \ldots, dh_n$ restrict to a base of the cotangent space $T_x X\check{} = \operatorname{Hom}(T_x X, \mathbb{R})$. Therefore every 1-form $\varphi \in \Omega^1 U$ has a unique expression

$$\varphi = \sum_{j=1}^{n} \varphi_j \cdot dh_j$$

in terms of smooth coefficient functions $\varphi_j : U \to \mathbb{R}$, and more generally every differential form $\varphi \in \Omega^k U$ may be uniquely expressed as

$$\varphi = \sum \varphi_{j_1 \cdots j_k} \cdot dh_{j_1} \wedge \cdots \wedge dh_{j_k}$$

with smooth functions $\varphi_{j_1\cdots j_k}\colon U \to \mathbb{R}$ indexed by all k-tuples such that $1 \le j_1 < j_2 < \cdots < j_k \le n$.

(3) It does no harm and is quite convenient to use the names $h_1, \ldots, h_n$ of the component functions likewise for the coordinates on $h(U) \subset \mathbb{R}^n$. Thus if $f\colon U \to \mathbb{R}$ is a differentiable function the partial derivatives of f represented in the chart h with respect to h_j make sense; they are in turn interpreted as functions on U and, by abuse of language, abbreviated as

$$\frac{\partial f}{\partial h_j} := \frac{\partial (f \circ h^{-1})}{\partial h_j} \circ h.$$

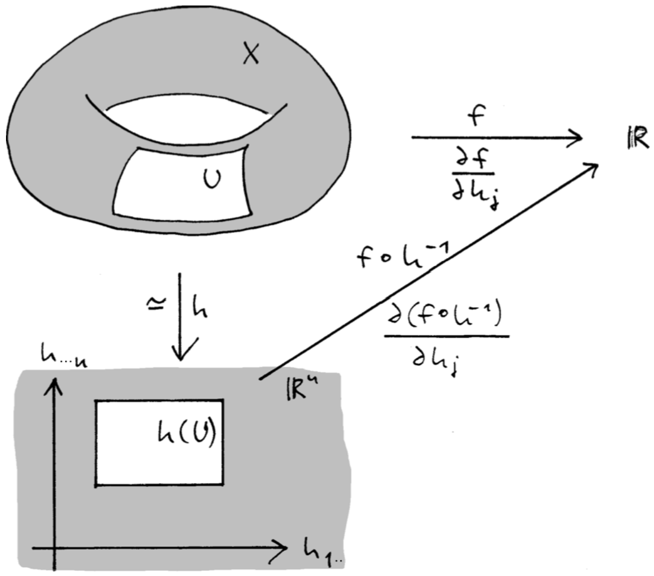

The differential of f then takes the easily memorised form

$$df = \sum_{j=1}^{n} \frac{\partial f}{\partial h_j} \cdot dh_j.$$

In particular f may be the i-th component of another chart (U, g), and we obtain the fool-proof formula

$$dg_i = \sum_{j=1}^{n} \frac{\partial g_i}{\partial h_j} \cdot dh_j$$

which governs the transition between representations of differential forms with respect to different charts. From the linear algebra of the previous section we infer that

$$dg_{i_1} \wedge \cdots \wedge dg_{i_k} = \sum_{j_1 < \cdots < j_k} \det \left(\frac{\partial g_{i_r}}{\partial h_{j_s}} \right)_{r,s=1}^{k} \cdot dh_{j_1} \wedge \cdots \wedge dh_{j_k}$$

is the transition formula for differential forms of arbitrary degree k.

6.3 Example The identical (global) chart of the 3-manifold $\mathbb{R}^3$ is traditionally written (x, y, z). Thus the differential forms on an open subset $X \subset \mathbb{R}^3$ are

$$\begin{aligned}
\Omega^0 X &: \quad f \cdot 1 = f \\
\Omega^1 X &: \quad f \cdot dx + g \cdot dy + h \cdot dz \\
\Omega^2 X &: \quad f \cdot dy \wedge dz + g \cdot dz \wedge dx + h \cdot dx \wedge dy \\
\Omega^3 X &: \quad f \cdot dx \wedge dy \wedge dz
\end{aligned}$$

with smooth functions $f, g, h\colon X \to \mathbb{R}$ — note how the base forms of degree 2 have been arranged according to the usual convention, based on the missing variable and cyclic permutation.

Spherical coordinates (R, θ, φ) give an alternative chart on the open subset, say,

$$X = \mathbb{R}^3 \smallsetminus \big([0,\infty)\times\{0\}\times\mathbb{R}\big)$$

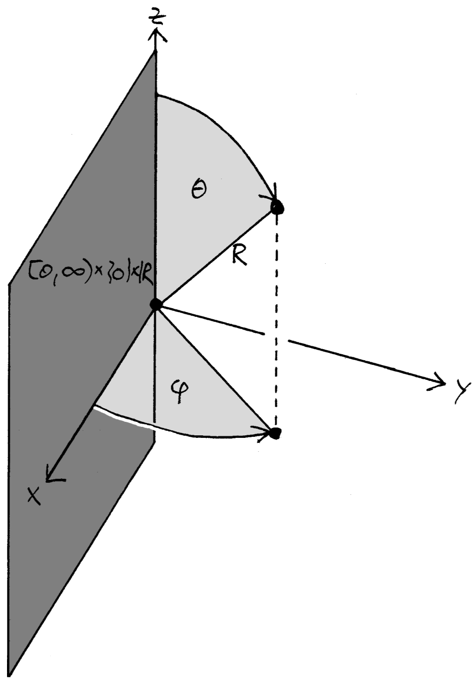

of $\mathbb{R}^3$ which corresponds to the restrictions of R positive, $\theta \in (0,\pi)$, and $\varphi \in (0, 2\pi)$. Differentiating the transition map

$$(R,\theta,\varphi) \longmapsto (x,y,z) = (R\sin\theta\cos\varphi, R\sin\theta\sin\varphi, R\cos\theta)$$

we obtain

$$\begin{aligned} dx &= \sin\theta\cos\varphi\, dR + R\cos\theta\cos\varphi\, d\theta - R\sin\theta\sin\varphi\, d\varphi \\ dy &= \sin\theta\sin\varphi\, dR + R\cos\theta\sin\varphi\, d\theta + R\sin\theta\cos\varphi\, d\varphi \\ dz &= \cos\theta\, dR - R\sin\theta\, d\theta, \end{aligned}$$

and thereby the transition formulae between the linear base forms of both charts. They directly imply those for forms of degree two and three like

$$\begin{aligned} dy\wedge dz &= \big(\sin\theta\sin\varphi\, dR + R\cos\theta\sin\varphi\, d\theta + R\sin\theta\cos\varphi\, d\varphi\big) \wedge \big(\cos\theta\, dR - R\sin\theta\, d\theta\big) \\ &= -R(\sin\theta)^2\sin\varphi\, dR\wedge d\theta + R(\cos\theta)^2\sin\varphi\, d\theta\wedge dR \\ &\quad + R\cos\theta\sin\theta\cos\varphi\, d\varphi\wedge dR - R^2(\sin\theta)^2\cos\varphi\, d\varphi\wedge d\theta \\ &= -R\sin\varphi\, dR\wedge d\theta - R\cos\theta\sin\theta\cos\varphi\, dR\wedge d\varphi + R^2(\sin\theta)^2\cos\varphi\, d\theta\wedge d\varphi\,, \end{aligned}$$

though

$$dx\wedge dy\wedge dz = R^2\sin\theta\cdot\, dR\wedge d\theta\wedge d\varphi$$

may be more readily obtained from the Jacobian determinant.

6.4 Orientations of Manifolds Let X be a manifold. An *oriented atlas* of X is a (differentiable) atlas for X which gives rise to orientation preserving transition maps: thus the Jacobian determinant of every transition map is required to be positive everywhere.

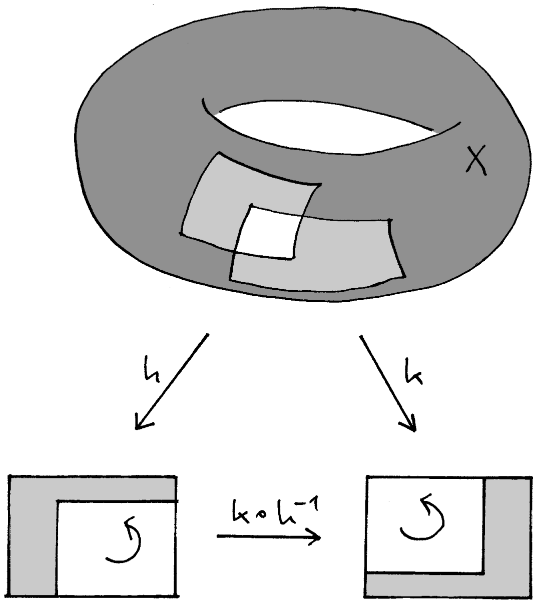

Every orientable atlas is contained in a unique maximal orientable one, and an *orientation* of X is defined as such a maximal orientable atlas for X; its charts are then called the oriented charts. If X is an oriented manifold then $-X$ is written to denote the same manifold with the opposite orientation.

A manifold may, but need not admit orientations and thus be *orientable*.

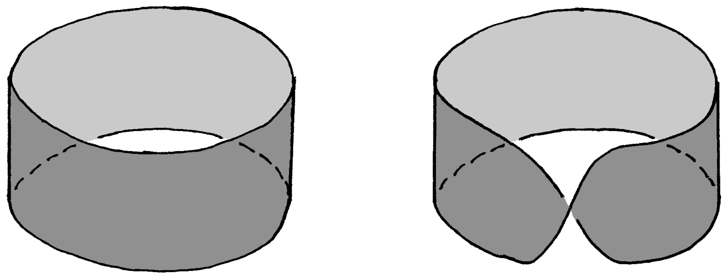

It is easy to see that a non-empty connected orientable manifold admits exactly two different orientations, see Exercise 6.3. A diffeomorphism between oriented manifolds X and Y is *orientation preserving* if all its local representations with respect to oriented charts have orientation preserving differentials — of course it is sufficient to verify this for suitably restricted charts taken from fixed oriented atlases of X and Y. All this remains perfectly true even for 0-manifolds X, where an orientation of X is just any function $X \to \{\pm 1\}$.

6.5 Integration Let X be an oriented n-manifold, and let $\varphi \in \Omega^n X$ be a differential form of the highest degree n, with compact support. The integral

$$\int_X \varphi$$

is defined in two steps:

- If the support of φ is contained in the domain of some oriented chart (U, h) then write

$$\varphi = f \cdot dh_1 \wedge \cdots \wedge dh_n$$

and put

$$\int_X \varphi := \int_{h(U)} f \circ h^{-1}.$$

The integral on the right hand side exists as it integrates a continuous function with compact support. It is clearly insensitive to shrinking the domain U as long as the latter contains the support of φ. We show that the integral does not change either if h is replaced by another oriented chart $k: U \to k(U)$ with the same domain. The form φ then reads

$$\varphi = f \cdot dh_1 \wedge \cdots \wedge dh_n = f \cdot \det \left(\frac{\partial h_i}{\partial k_j} \right)_{i,j=1}^{n} \cdot dk_1 \wedge \cdots \wedge dk_n,$$

and since the determinant is positive the classical transformation formula for integrals gives

$$\int_{k(U)} \Big(f \cdot \det \left(\frac{\partial h_i}{\partial k_j} \right)_{i,j=1}^{n} \Big) \circ k^{-1} = \int_{k(U)} (f \circ h^{-1}) \circ (h \circ k^{-1}) \left| \det D(h \circ k^{-1}) \right| = \int_{h(U)} f \circ h^{-1},$$

so that we obtain the same result indeed.

• For an arbitrary form $\varphi \in \Omega^n$ with compact support we cover X by oriented charts and choose a subordinate partition of unity $(\rho_j)_{j=0}^{\infty}$. The sum in

$$\int_X \varphi := \sum_{j=0}^{\infty} \int_X \rho_j \varphi$$

has but finitely many non-zero terms and is easily seen to be independent of the choice of the partition of unity.

Notes As to notation, if $X \subset Y$ happens to be a submanifold and φ an n-form on Y then the integral $\int_X (\varphi|X)$ is simply written $\int_X \varphi$.

Of course much more general n-forms may be considered for integration. Very much like in coordinate integration theory, the existence of the integral is no longer guaranteed beyond continuous forms with compact support, and the notion of integrable n-form on an oriented n-manifold arises naturally.

While the definition of the integral based on partitions of unity is perfectly satisfactory from a conceptual point of view, it is little suited to explicit calculus. Nor is it needed there since for the purpose of evaluating an integral any null sets may be happily ignored. On a smooth n-manifold X the notion of null set makes sense even though in the absence of further structure X does not carry a canonical measure: A subset $N \subset X$ is a *null set* if for every chart (U, h) of X the image $h(U \cap N) \subset \mathbb{R}^n$ is an n-dimensional Lebesgue null set. This condition only needs to be verified for the charts of some atlas of X rather than for all charts, since the transition functions, like all C^1 mappings, send null sets to null sets.

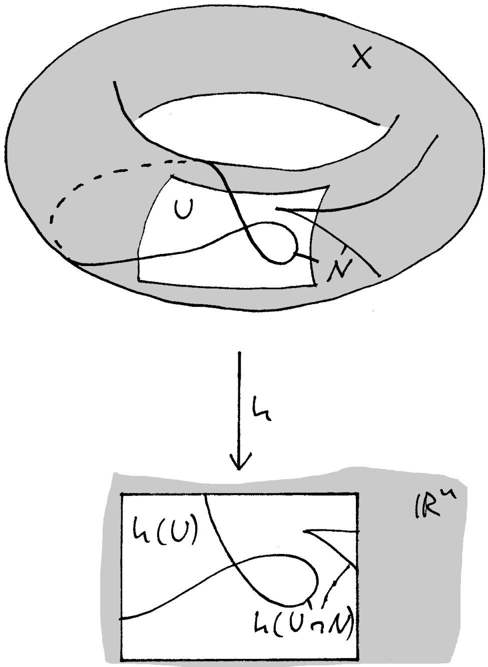

In practice removing suitable null sets usually allows to cover X by few pairwise disjoint chart domains, if not even a single chart suffices.

6.6 Example Spherical coordinates provide a chart

$$\mathbb{R}^3 \supset \mathbb{R}^3 \smallsetminus \big([0,\infty)\times\{0\}\times\mathbb{R}\big) = U \xrightarrow{\;h\;} h(U) = (0,\infty)\times(0,\pi)\times(0,2\pi) \subset \mathbb{R}^3$$

defined on the complement of a null set of $\mathbb{R}^3$. If the 3-form on $\mathbb{R}^3$ with compact support is given as $\psi = f(R,\theta,\varphi)\,dR\wedge d\theta\wedge d\varphi$ then its integral may be evaluated using Fubini's theorem:

$$\int_{\mathbb{R}^3}\psi = \int_U f(R,\theta,\varphi)\,dR\wedge d\theta\wedge d\varphi = \int_0^{2\pi}\int_0^{\pi}\int_0^{\infty} f(R,\theta,\varphi)\,dR\,d\theta\,d\varphi.$$

6.7 Volume Forms A nowhere-vanishing n-form $\omega\in\Omega^nX$ on an n-manifold X is called a *volume form*. It immediately defines an orientation of X, singling out those charts (U,h) in which ω is represented as

$$\omega = f\cdot dh_1\wedge\cdots\wedge dh_n \quad \text{with } f(x)>0 \text{ for all } x\in U.$$

While integrating a function on a manifold of positive dimension has no meaning in general it does make sense as soon as a volume form is selected, say ω, for then ω not only provides the necessary orientation of X but it also converts functions $f\colon X\to\mathbb{R}$ to n-forms by simple multiplication $f\mapsto f\cdot\omega$.

Let us return to the notion of differential form on a manifold in general. Apart from the possibility to integrate forms of highest degreee we have not yet seen compelling reasons to single out alternating forms rather than for instance symmetric, or general multilinear ones. Indeed quite analogous concepts exist which are based not on forms at all but rather on tensor products of tangent vectors, and even both ideas can be combined. Such objects, generally called tensor fields on manifolds are very interesting and important. But there is one feature that makes the alternating differential forms unique among all tensor fields: differential forms can be differentiated without reference to any additional structure on the underlying manifold.

6.8 Theorem and Definition Let X be a manifold. There exists a unique sequence

$$\Omega^kX \xrightarrow{\;d\;} \Omega^{k+1}X \quad \text{for all } k\in\mathbb{N}$$

of linear operators satisfying the following axioms. They are called *exterior derivatives* or the *Cartan differentials*.

(a) The 0-th differential $d\colon\Omega^0X\xrightarrow{d}\Omega^1X$ has the same meaning as before, explained in 6.2(1).

(b) The product rule $d(\varphi\wedge\chi) = d\varphi\wedge\chi + (-1)^k\varphi\wedge d\chi$ holds for all $\varphi\in\Omega^kX$ and $\chi\in\Omega^lX$.

(c) For every $k\in\mathbb{N}$ the composition $\Omega^kX\xrightarrow{d}\Omega^{k+1}X\xrightarrow{d}\Omega^{k+2}X$ is the zero operator.

Note that the sign in (b) is to be expected since d, which raises the degree of forms, may be thought of as an operator of degree 1.

Proof At first we make the additional assumption that the n-manifold X admits a global chart, say (X,h). Then every $\varphi\in\Omega^kX$ has a unique representation

$$\varphi = \sum_{j_1,\ldots,j_k}\varphi_{j_1\cdots j_k}\cdot dh_{j_1}\wedge\cdots\wedge dh_{j_k}$$

with indices $j_1<\cdots<j_k$ and coefficient functions $\varphi_{j_1\cdots j_k}$. Repeated application of the axioms shows that for the value $d\varphi\in\Omega^{k+1}X$ there is no choice but to put

$$d\varphi = \sum d\varphi_{j_1\cdots j_k}\wedge dh_{j_1}\wedge\cdots\wedge dh_{j_k}.$$

This formula proves uniqueness, and we now read it as the definition of $d\varphi$. It clearly complies with (a), and by linearity the axiom (b) need only be verified on differential forms $\varphi = f\cdot dh_{j_1}\wedge\cdots\wedge dh_{j_k}$ and $\chi = g\cdot dh_{j_1}\wedge\cdots\wedge dh_{j_k}$:

$$\begin{aligned}
d(\varphi\wedge\chi) &= d(f\,dh_{i_1}\wedge\cdots\wedge dh_{i_k}\wedge g\,dh_{j_1}\wedge\cdots\wedge dh_{j_l})\\
&= d(fg\,dh_{i_1}\wedge\cdots\wedge dh_{i_k}\wedge dh_{j_1}\wedge\cdots\wedge dh_{j_l})\\
&= (df\cdot g+f\cdot dg)\wedge dh_{i_1}\wedge\cdots\wedge dh_{i_k}\wedge dh_{j_1}\wedge\cdots\wedge dh_{j_l}\\
&= (df\wedge dh_{i_1}\wedge\cdots\wedge dh_{i_k})\wedge(g\,dh_{j_1}\wedge\cdots\wedge dh_{j_l})\\
&\qquad +(-1)^k(f\,dh_{i_1}\wedge\cdots\wedge dh_{i_k})\wedge(dg\wedge dh_{j_1}\wedge\cdots\wedge dh_{j_l})\\
&= d\varphi\wedge\chi+(-1)^k\varphi\wedge d\chi
\end{aligned}$$

As to axiom (c), applied to a form of degree zero it is a consequence of the fact that the partial differentiation operators commute with each other:

$$\begin{aligned}
df &= \sum_{j=1}^n \frac{\partial f}{\partial h_j}\,dh_j\\
ddf &= \sum_{i,j=1}^n \frac{\partial^2 f}{\partial h_i\partial h_j}\,dh_i\wedge dh_j = 0.
\end{aligned}$$

The product rule (b) then extends it to all degrees:

$$dd\varphi = dd(f\,dh_{i_1}\wedge\cdots\wedge dh_{i_k}) = d(df\wedge dh_{i_1}\wedge\cdots\wedge dh_{i_k}) = 0.$$

This completes the proof for manifolds X that admit a global chart.

We now remove this condition. We know that the theorem holds for every chart domain $U\subset X$, and denote by d_U the Cartan differentials of the manifold U. Note that since charts may be restricted the differentials are compatible in the sense that whenever V is a chart domain and $U\subset V$ an open subset then the diagram

$$\begin{tikzcd}
\Omega^k V \arrow[r, "d"] \arrow[d] & \Omega^{k+1}V \arrow[d]\\
\Omega^k U \arrow[r, "d"] & \Omega^{k+1}U
\end{tikzcd}$$

of differentials and restrictions commutes. — We now define a global Cartan differential d_X of X as follows. Let the form $\varphi\in\Omega^kX$ and a point $x\in X$ be given. We choose a chart domain $U\subset X$ with $x\in U$ and put

$$(d_X\varphi)_x = \big(d_U(\varphi|U)\big)_x \in \operatorname{Alt}^k T_xX.$$

This definition does not depend on the choice of U, and since on the other hand we clearly have $(d_X\varphi)|U = d_U(\varphi|U)$ the resulting Cartan operators d_X satisfy the three axioms.

The proof of uniqueness is more delicate. Let d denote any sequence of Cartan differentials for X, let φ be a k-form, and $x\in X$ a point. We choose a chart (U,h) and a table mountain $\sigma\colon X\to\mathbb{R}$ with $\operatorname{supp}\sigma\subset U$, both centred at x. If

$$\varphi|U = \sum \varphi_{j_1\cdots j_k}\cdot dh_{j_1}\wedge\cdots\wedge dh_{j_k}\in\Omega^kU$$

is the chart representation of φ — note that in view of axiom (a) the meaning of d is unambiguous — then we define a global form

$$\tilde\varphi = \sum \tilde\varphi_{j_1\cdots j_k}\cdot d\tilde h_{j_1}\wedge\cdots\wedge d\tilde h_{j_k}\in\Omega^kX$$

using the smooth functions $\tilde\varphi_{j_1\cdots j_k} = \sigma\cdot\varphi_{j_1\cdots j_k}$ and $\tilde h_j = \sigma\cdot h_j$ on X. Applying the axioms to d we obtain that

$$d\tilde\varphi = \sum d\tilde\varphi_{j_1\cdots j_k}\wedge d\tilde h_{j_1}\wedge\cdots\wedge d\tilde h_{j_k}$$

and, evaluating at x, that $(d\tilde{\varphi})_x = \big(d_U(\tilde{\varphi}|U)\big)_x = (d_X\tilde{\varphi})_x = (d_X\varphi)_x$ since φ and $\tilde{\varphi}$ coincide in some neighbourhood of x.

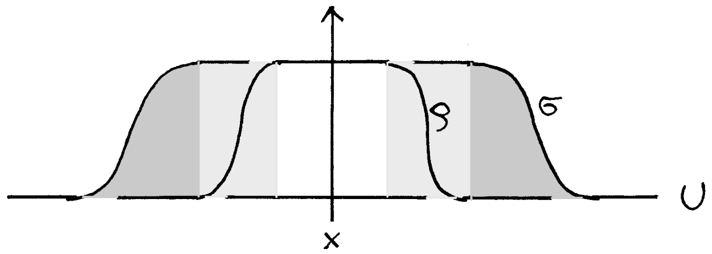

It only remains to prove that $(d\varphi)_x = (d\tilde{\varphi})_x$. To this end we choose a narrower table mountain $\rho\colon X \to \mathbb{R}$ at x such that $\operatorname{supp}\rho \subset \sigma^{-1}\{1\}$. Then $\rho\cdot\varphi = \rho\cdot\tilde{\varphi}$, and another application of the axioms to d yields

$$d\rho\wedge\varphi + \rho\cdot d\varphi = d(\rho\,\varphi) = d(\rho\,\tilde{\varphi}) = d\rho\wedge\tilde{\varphi} + \rho\cdot d\tilde{\varphi}.$$

Since ρ has constant value 1 near x evaluation just leaves $(d\varphi)_x = (d\tilde{\varphi})_x$ as desired: this concludes the proof. □

The last point of the proof is of independent interest and leads to the first part of the following proposition.

6.9 Proposition The Cartan differentials are local operators:

$$\operatorname{supp} d\varphi \subset \operatorname{supp}\varphi \quad \text{for every differential form } \varphi \in \Omega^k Y.$$

They also are functorial with respect to mappings: for every smooth mapping $f\colon X \to Y$ and every $k \in \mathbb{N}$ the diagram

$$\begin{tikzcd}
\Omega^k Y \arrow[r, "d"] \arrow[d, "f^*"] & \Omega^{k+1} Y \arrow[d, "f^*"] \\
\Omega^k X \arrow[r, "d"] & \Omega^{k+1} X
\end{tikzcd}$$

commutes.

Proof Locality is clear since the value of $d\varphi$ at a point $x \in X$ is defined by first restricting φ to a neighbourhood of x that we may prescribe to be small. Functoriality for forms of degree zero, that is functions $\varphi\colon Y \to \mathbb{R}$, follows from the chain rule:

$$d(f^*\varphi) = d(\varphi\circ f) = d\varphi\circ Tf = f^*d\varphi.$$

It also holds for 1-forms of type dh, for the simple reason that $d(f^*dh) = ddf^*h$ and $f^*d(dh)$ both vanish. Locally, in terms of a chart, every form of higher degree is obtained from these two types as a linear combination of wedge products. Therefore it is sufficient to prove that f^* and d commute on the wedge product $\varphi\wedge\chi$ whenever they commute on both $\varphi \in \Omega^k X$ and $\chi \in \Omega^l X$. The direct calculation

$$\begin{aligned}
f^*d(\varphi\wedge\chi) &= f^*\big(d\varphi\wedge\chi + (-1)^k\varphi\wedge d\chi\big)\\
&= f^*d\varphi\wedge f^*\chi + (-1)^k f^*\varphi\wedge f^*d\chi\\
&= d(f^*\varphi)\wedge f^*\chi + (-1)^k f^*\varphi\wedge d(f^*\chi)\\
&= d(f^*\varphi\wedge f^*\chi)\\
&= d\big(f^*(\varphi\wedge\chi)\big)
\end{aligned}$$

achieves this and thus completes the proof of the proposition. □

6.10 Summary of Differential Form Calculus The triangle comprises the most prominent laws that govern exterior multiplication, Cartan differentials, and pull-backs of differential forms. Rules concerning one of these operations are placed at the corners, and compatibility formulae between operations at the corresponding edges.

$$\begin{tikzcd}
 & \chi\wedge\varphi = (-1)^{\deg\varphi\cdot\deg\chi}\,\varphi\wedge\chi \arrow[dl, dash, "f^*(\varphi\wedge\chi) = f^*\varphi\wedge f^*\chi"'] \arrow[dr, dash, "d(\varphi\wedge\chi) = d\varphi\wedge\chi + (-1)^{\deg\varphi}\varphi\wedge d\chi"] & \\
(gf)^* = f^*g^* \arrow[rr, dash, "f^*d = d\,f^*"'] & & dd = 0
\end{tikzcd}$$

Note that these laws are universal in the sense that they do not depend on, in fact do not even refer to any manifold chart. If on the other hand charts must be used for an explicit calculation the differential form calculus does not impose a particular one: a chart may be arbitrarily selected to suit the particular application.

6.11 Classical Vector Analysis Let $X \subset \mathbb{R}^3$ be an open subset. Cartesian coordinates x, y, z allow to identify every differential form on X with the system of its coefficient functions as in 6.3: for instance $f\,dx + g\,dy + h\,dz \in \Omega^1 X$ with $(f, g, h) \in (\Omega^0 X)^3$. The basic operations $\wedge$ and d then translate according to the diagrams

$$\begin{tikzcd}
\Omega^1 X \times \Omega^1 X \arrow[r, "\wedge"] \arrow[d, "\simeq"] & \Omega^2 X \arrow[d, "\simeq"] \\
(\Omega^0 X)^3 \times (\Omega^0 X)^3 \arrow[r, "\times"] & (\Omega^0 X)^3
\end{tikzcd}
\qquad
\begin{tikzcd}
\Omega^1 X \times \Omega^2 X \arrow[r, "\wedge"] \arrow[d, "\simeq"] & \Omega^3 X \arrow[d, "\simeq"] \\
(\Omega^0 X)^3 \times (\Omega^0 X)^3 \arrow[r, "\bullet"] & \Omega^0 X
\end{tikzcd}$$

and

$$\begin{tikzcd}
\Omega^0 X \arrow[r, "d"] \arrow[d, "\simeq"] & \Omega^1 X \arrow[r, "d"] \arrow[d, "\simeq"] & \Omega^2 X \arrow[r, "d"] \arrow[d, "\simeq"] & \Omega^3 X \arrow[d, "\simeq"] \\
\Omega^0 X \arrow[r, "\mathrm{grad}"] & (\Omega^0 X)^3 \arrow[r, "\mathrm{rot}"] & (\Omega^0 X)^3 \arrow[r, "\mathrm{div}"] & \Omega^0 X
\end{tikzcd}$$

to operations on the coefficient functions respectively vectors. The algebraic operations, which of course already make sense on the level of linear algebra rather than differential forms, are known as the vector and scalar products

$$\begin{pmatrix} f \\ g \\ h \end{pmatrix} \times \begin{pmatrix} f' \\ g' \\ h' \end{pmatrix} = \begin{pmatrix} gh' - hg' \\ hf' - fh' \\ fg' - gf' \end{pmatrix} \quad \text{and} \quad \begin{pmatrix} f \\ g \\ h \end{pmatrix} \bullet \begin{pmatrix} f' \\ g' \\ h' \end{pmatrix} = ff' + gg' + hh'.$$

As to the analytic operators, they are

$$\mathrm{grad}\, f = \begin{pmatrix} \partial f/\partial x \\ \partial f/\partial y \\ \partial f/\partial z \end{pmatrix}, \quad \mathrm{rot} \begin{pmatrix} f \\ g \\ h \end{pmatrix} = \begin{pmatrix} \partial h/\partial y - \partial g/\partial z \\ \partial f/\partial z - \partial h/\partial x \\ \partial g/\partial x - \partial f/\partial y \end{pmatrix}, \quad \text{and} \quad \mathrm{div} \begin{pmatrix} f \\ g \\ h \end{pmatrix} = \frac{\partial f}{\partial x} + \frac{\partial g}{\partial y} + \frac{\partial h}{\partial z}$$

with alternative fancy symbols $\mathrm{grad} = \nabla$, $\mathrm{rot} = \mathrm{curl} = \nabla\times$, and $\mathrm{div} = \nabla\bullet$. These traditional vector operations, known and in common use by the end of the 19th century, were definitely established in 1901 in a monograph titled “Vector Analysis” by the mathematician Edwin Bidwell Wilson, based on lectures by the physicist Josiah Willard Gibbs. In terms of these classical operations the concise laws 6.10 of vector analysis disintegrate into an extensive zoo of seemingly unrelated identities. Physicists and engineers live with them to these days (enhancing the zoo by as many versions of the integral theorem 7.4 to be discussed in the next section). This may be due to the authority exercised by Gibbs through Wilson's book[1], but also to the fact that the smooth calculus of differential forms as we know it today took its time to emerge. While the notion of differential form is generally attributed to a 1899 paper by Elie Cartan, it was not until 1946 that it was clarified in a mathematically satisfactory way by Claude Chevalley. Nevertheless it seems peculiar if not absurd that scientists and engineers even today largely disdain differential forms and continue to bother with an awkward formalism in presence of a much simpler alternative.

6.12 Definition A differential form φ is *closed*[2] if $d\varphi = 0$. It is called *exact* if there exists a form χ such that $d\chi = \varphi$.

Note In view of $d \circ d = 0$ every exact form is closed. The converse does not hold in general but is true under special conditions:

6.13 Poincaré's Lemma Let $X \subset \mathbb{R}^n$ be an open set which is star-shaped with respect to one of its points a, in the sense that for every $x \in X$ the segment

$$\big\{(1-t)a + tx \,\big|\, t \in [0,1]\big\}$$

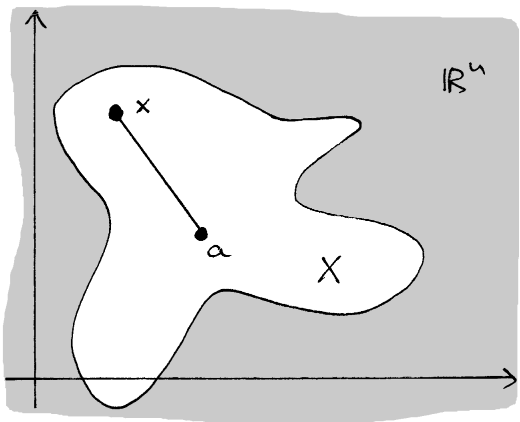

is completely contained in X. Then for $k > 0$ every closed differential form $\varphi \in \Omega^k X$ is exact. A so-called *homotopy operator* $H\colon \Omega^k X \to \Omega^{k-1} X$ allows to explicitly compute a form $H\varphi \in \Omega^{k-1} X$ with $d(H\varphi) = \varphi$:

$$(H\varphi)_x(v_2, \ldots, v_k) = \int_0^1 t^{k-1} \varphi_{(1-t)a+tx}(x, v_2, \ldots, v_k)\, dt \quad \text{for } \varphi \in \Omega^k X$$

Proof Note how the geometric hypothesis ensures that the integral makes sense at all. Patient calculation gives the identity

$$H \circ d + d \circ H = \mathrm{id}$$

of endomorphisms of $\Omega^k X$. Applied to a closed form φ there remains the equation $d(H\varphi) = \varphi$, which was the claim. □

As charts may always be chosen to map onto an open ball we obtain at once:

6.14 Corollary A differential form φ of positive degree is closed if and only it is *locally* exact, that is if every point admits an open neighbourhood on which φ is exact. □

[1] This is a well-written book, and Gibbs' notions were eminently useful in their time.
[2] In German: geschlossen, not abgeschlossen.

Notes A 0-form on any manifold X, that is a function $\varphi: X \to \mathbb{R}$ is closed if and only if it is locally constant. On the other hand the only exact 0-form is the zero form. — Physicists and engineers tend to ignore the difference between closedness and exactness when they state that "every conservative vector field admits a scalar potential", or that "every incompressible vector field admits a vector potential" — in mathematical terms they would erroneously claim that every closed form of degree 1 respectively 2 is exact. This negligence is at the root of "unexpected" phenomena like non-uniqueness of solutions in certain instances of Maxwell's equations.

6.15 Example The polar angle of a point $(x, y) \in S^1$

$$\begin{array}{ll} \arctan\,(y/x) \text{ where } x > 0 & \arctan\,(y/x) + \pi \text{ where } x < 0 \\ \operatorname{arccot}\,(x/y) \text{ where } y > 0 & \operatorname{arccot}\,(x/y) + \pi \text{ where } y < 0 \end{array}$$

is not a well-defined function φ on the manifold S^1; nevertheless it is often treated as a smooth scalar function which is locally defined up to the choice of an additive constant (a multiple of 2π). By contrast the inverse $\varphi^{-1}: \mathbb{R} \to S^1$ *is* well-defined, sending $\varphi \in \mathbb{R}$ to $(\cos\varphi, \sin\varphi) \in S^1$.

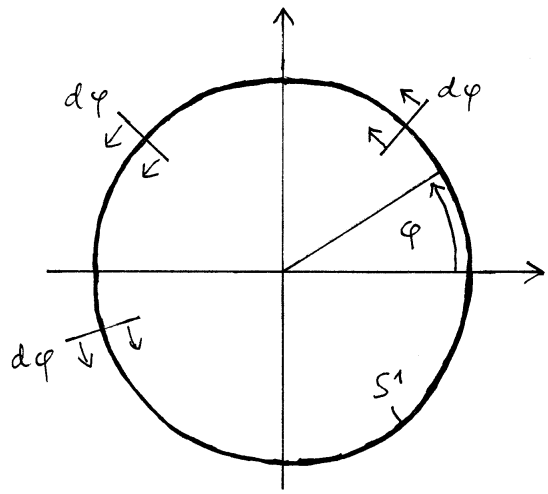

Also well-defined is the differential $d\varphi \in \Omega^1 S^1$ since differentiation kills additive constants. Nevertheless the notation $d\varphi$ is an abuse of language as it suggests an exact form. In fact $d\varphi$ is closed since there are no non-trivial 2-forms on S^1, but it fails to be exact, as we now see. First note that the three formulae for the polar angle with $x < 0$ or $y \neq 0$ do yield a well-defined function φ, which is a chart for S^1 defined on the complement of the point $(1, 0) \in S^1$ and with values in the interval $(0, 2\pi)$. The integral of $d\varphi$ is easily evaluated in terms of this chart:

$$\int_{S^1} d\varphi = \int_0^{2\pi} 1 \cdot d\varphi = \varphi\Big|_{\varphi=0}^{2\pi} = 2\pi.$$

If on the other hand there were a true function $f \in \Omega^0 S^1$ with $df = d\varphi$ then the integral would vanish:

$$\int_{S^1} d\varphi = \int_{S^1} df = \int_{S^1} \frac{\partial f}{\partial \varphi} d\varphi = \int_0^{2\pi} \frac{d}{d\varphi}(f \circ \varphi^{-1}) d\varphi = f \circ \varphi^{-1}\Big|_{\varphi=0}^{2\pi} = 0.$$

Thus $d\varphi$ cannot be exact.

The discrepancy between closedness and exactness has in itself been an interesting subject for investigation:

6.16 Definition Let $k \in \mathbb{N}$. For every differential manifold X the quotient vector space

$$H^k X := \frac{\operatorname{kernel}\big(\Omega^k X \xrightarrow{d} \Omega^{k+1} X\big)}{\operatorname{image}\big(\Omega^{k-1} X \xrightarrow{d} \Omega^k X\big)}$$

is called the k-th *de Rham cohomology* of X. It is functorial in the sense that every smooth map $f\colon X \to Y$ induces linear mappings $f^*\colon H^kY \to H^kX$ which satisfy $\mathrm{id}^* = \mathrm{id}$ and $(g \circ f)^* = f^* \circ g^*$.

Notes While spaces of differential forms nearly always have infinite dimension it can be shown that the de Rham cohomology spaces of, for instance, a compact manifold are finite dimensional. Their functoriality generalises from smooth to all continuous mappings, so that the cohomology spaces of a manifold X are in fact invariants of the topological space that underlies X. They turn out to admit several other interpretations in the framework of algebraic topology, and can often be explicitly determined. Example 6.14 for instance shows that H^1S^1 is non-trivial; in fact the congruence class of $d\varphi$ turns out to span this vector space, so that $H^1S^1 = \mathbb{R} \cdot [d\varphi] \simeq \mathbb{R}$ — see Exercise 6.9.

Exercises

6.1 The assignment $\chi_x(v, w) := \det \begin{pmatrix} x & v & w \end{pmatrix}$ defines a differential form $\chi \in \Omega^2\mathbb{R}^3$.

- Represent χ in Cartesian coordinates.
- Represent the restriction $\chi|S^2$ in spherical coordinates (θ, φ). What is the intuitive meaning of this form?
- Giving S^2 the "standard" orientation — that for which (θ, φ) is an oriented chart — compute the area of the sphere $\int_{S^2} \chi$.

6.2 Let X be a differentiable manifold of dimension n. Show that the following are equivalent:

(1) X is orientable.

(2) There exists a volume form on X.

(3) The line bundle $\mathrm{Alt}^n\, TX \to X$ ist trivial.

6.3 Prove that a non-empty connected orientable manifold admits exactly two orientations.

6.4 Every integral matrix $a \in \mathrm{Mat}(n{\times}n, \mathbb{Z})$ induces a smooth self-mapping $f\colon X \to X$ of the n-dimensional torus $X = \mathbb{R}^n/\mathbb{Z}^n$. Explain how to read off from a whether f is

- a local diffeomorphism,
- a (global) diffeomorphism,
- an orientation preserving local diffeomorphism.

Find the relation between $\int_X f^*\chi$ and$\int_X \chi$ for differential forms $\chi \in \Omega^nX$.

Hint You may find it useful to represent a in Smith normal form.

6.5 Compute the exterior derivative of the following differential forms:

- an arbitrary $(n-1)$-form $\alpha = \sum_{j=1}^n (-1)^{j-1}\, \alpha_j(x)\, dx_1 \wedge \cdots \wedge \widehat{dx_j} \wedge \cdots \wedge dx_n$ on an open subset $X \subset \mathbb{R}^n$;
- the form $\beta = -\frac{y}{\sqrt{x^2+y^2}}\, dx + \frac{x}{\sqrt{x^2+y^2}}\, dy$ on $\mathbb{R}^2 \smallsetminus \{0\}$;
- the form $\gamma = r\, d\varphi$ (in polar coordinates) on $\mathbb{R}^2 \smallsetminus \{0\}$.

Is there a connection between β und γ?

6.6 Prove that $\beta = 2xz\, dy\wedge dz - dz\wedge dx - (e^x + z^2)\, dx\wedge dy$ is a closed form on $\mathbb{R}^3$. Compute a form α on $\mathbb{R}^3$ with $d\alpha = \beta$.

6.7 A popular formula of classical vector analysis states that

$$\nabla \bullet (\vec{F} \times \vec{G}) = (\nabla \times \vec{F}) \bullet \vec{G} - \vec{F} \bullet (\nabla \times \vec{G})$$

for "vector fields" $\vec{F}$ and $\vec{G}$, that is, functions on an open subset of $\mathbb{R}^3$ with values in $\mathbb{R}^3$. Which identity is hidden behind it?

6.8 Let $\varphi \in \Omega^k X$ und $\chi \in \Omega^l X$ be differential forms on the manifold X. Prove the following:

- If φ and χ are closed then so is $\varphi\wedge\chi$.
- If φ is a closed, and χ an exact form then $\varphi\wedge\chi$ is exact.

Explain how these facts imply that the total de Rham cohomology $\bigoplus_{k=0}^{\infty} H^k X$ carries the structure of a graded commutative algebra over $\mathbb{R}$.

6.9 Prove that the congruence class of $d\varphi$ spans $H^1 S^1$ as a real vector space.

7 Boundaries and Integrals

For many purposes the concept of manifold that we have developed is too narrow. One reason is that one would often like to think of a connected manifold as glued from smaller pieces of the same dimension without much overlap: based on the present notion alone this is clearly impossible since pieces which themselves are manifolds would have to intersect in non-empty open subsets.

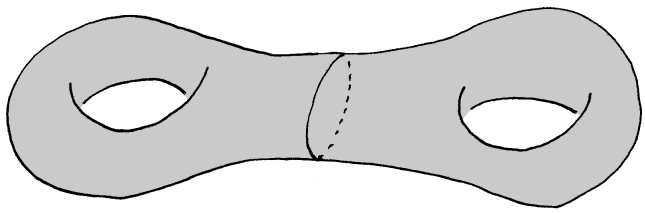

7.1 Definition The concept of n-dimensional *manifold with boundary* extends that of differential n-manifold by allowing charts (U, h) such that $h(U)$ is an open subset of the closed half-space

$$\mathbb{R}^n_- = (-\infty, 0] \times \mathbb{R}^{n-1} = \big\{(x_1, \ldots, x_n) \in \mathbb{R}^n \,\big|\, x_1 \leq 0\big\}$$

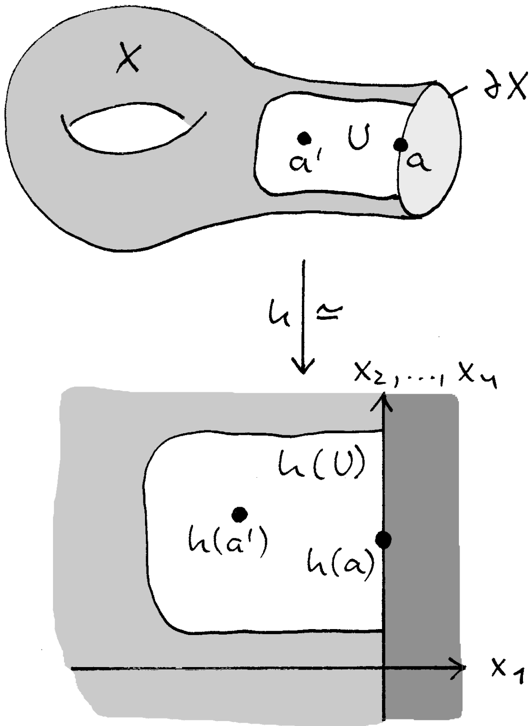

rather than $\mathbb{R}^n$. A point $a \in X$ is called an *inner* respectively a *boundary point*[1] of X according to whether such a chart at a throws a into $(-\infty, 0) \times \mathbb{R}^{n-1}$ or into $\{0\} \times \mathbb{R}^{n-1}$. The set of all boundary points of the manifold X is called its *boundary* $\partial X \subset X$.

[1] These notions are not to be confused with those pertaining to subsets of a topological space.

Explanation This definition presupposes the notion of differentiability of a mapping

$$\mathbb{R}^n_- \supset S \xrightarrow{f} T \subset \mathbb{R}^p_-$$

between open subsets $S \subset \mathbb{R}^n_-$ and $T \subset \mathbb{R}^p_-$ via the existence of local smooth extensions, which in a special case has already been considered in Section 4. More precisely f is considered differentiable at $a \in S$ if and only if there exist an open neighbourhood $V \subset \mathbb{R}^n$ of a and a smooth mapping $F\colon V \to \mathbb{R}^p$ such that

$$f|(S\cap V) = F|(S\cap V).$$

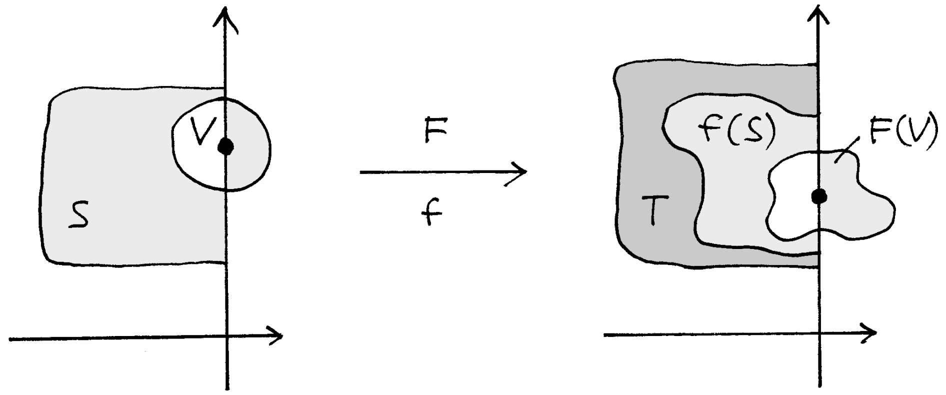

In contrast to Section 4, where arbitrary domains of f were allowed, the now relevant domain S is sufficiently "thick" to make the differential $DF(a) \in \mathrm{Mat}(p\times n, \mathbb{R})$ independent of the choice of the local extension F, and therefore a well-defined feature $Df(a)$ of f. As a consequence differentiable mappings between manifolds with boundary have well-defined differentials at all points, including boundary points.

By the local inverse theorem any transition map

$$\mathbb{R}^n_- \supset S \xrightarrow{k\circ h^{-1}} T \subset \mathbb{R}^n_-$$

must send S°, the interior of S with respect to $\mathbb{R}^n$, into T°: therefore the distinction between inner and boundary points of X does not depend on the choice of a particular chart. It also follows that $\partial X \subset X$ is a closed subset.

From a mere logical point of view the charts of 7.1 might supersede the old style ones of Definition 1.1; yet it does no harm, and is useful to retain the latter as additional charts. In any case the term of manifold as such will continue to mean a manifold without boundary; it is of course included in the extended notion as that of a manifold with empty boundary.

7.2 Example The closed disk (or synonymously, ball) $D^n \subset \mathbb{R}^n$ is an n-dimensional manifold with boundary $\partial D^n = S^{n-1}$. In fact the charts we used in 1.11(1) to show that the sphere is a submanifold of $\mathbb{R}^n$ require a but slight adjustment to become charts for D^n, for instance

$$\big\{x \in D^n \,\big|\, x_1 > 0\big\} \ni (x_1,\ldots,x_n) \longmapsto \Big(x_1 - \sqrt{1-\textstyle\sum_{j>1} x_j^2},\, x_2,\ldots,x_n\Big) \in U^n \cap \mathbb{R}^n_-.$$

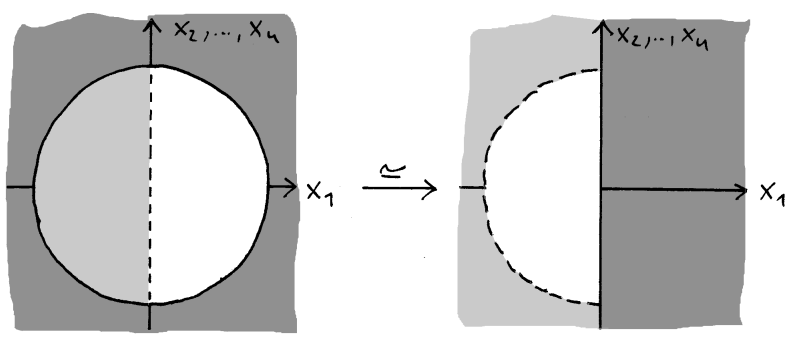

7.3 Analogies and No-Analogies (1) At boundary points the local inverse theorem does not hold, as is testified by the smooth map

$$\mathbb{R}^2_- \ni (x, y) \mapsto (x-y^2, y) \in \mathbb{R}^2_-.$$

The differential of this map at the origin is the identity, yet there is no neighbourhood of the origin on the left that is sent onto a full neighbourhood of the origin on the right.

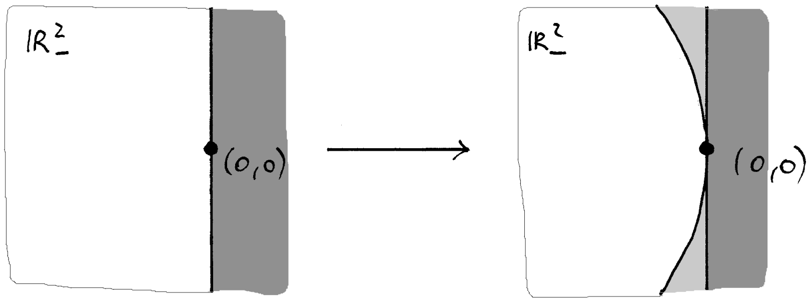

(2) The Cartesian product of a manifold X with boundary and one without boundary Y is a new manifold with boundary $\partial(X \times Y) = \partial X \times Y$, but the Cartesian product of two manifolds with non-empty boundary is not — rather one with "corners". Depending on the context one may wish to elaborate and systematically work with this more general notion, or avoid it applying a technique called "straightening the angle" which turns such products back into ordinary manifolds with boundary.

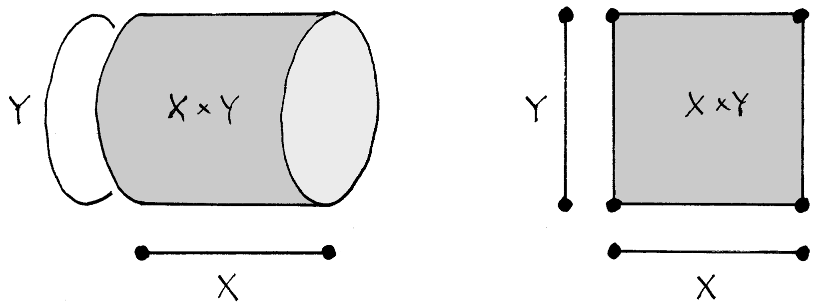

(3) The regular value theorem has a variant which is an abundant source of manifolds with boundary: Let X be an n-manifold without boundary, and $f: X \to \mathbb{R}$ a smooth scalar function with regular value $b \in \mathbb{R}$. Then $S := f^{-1}(-\infty, b]$ is in a natural way an n-manifold with boundary $\partial S = f^{-1}\{b\}$.

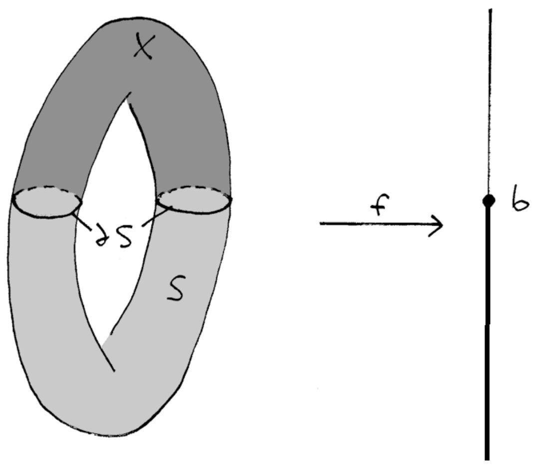

(4) There is not one universal notion of submanifold but several, pertaining to different situations. A first case is that of equidimensional submanifolds, with $S \subset X$ from (3) as an example.

Another is that of the boundary $\partial X \subset X$, which itself can be considered as a closed submanifold of codimension one: simply restrict charts (U, h) for X to $U \cap \partial X$ and note that the transition map to another chart (V, k) restricts to a diffeomorphism

$$h(U \cap V) \cap \big(\{0\} \times \mathbb{R}^{n-1}\big) \xrightarrow{\simeq} k(U \cap V) \cap \big(\{0\} \times \mathbb{R}^{n-1}\big).$$

In particular ∂X is a manifold (without boundary) in its own right.

The probably most useful class of submanifold comprises those $S \subset X$ which at boundary points are transverse to the boundary in the following sense. Assuming some positive dimension $s \leq n = \dim X$ it is required that at each point $a \in S$ there exists a chart (U, h) for X such that

$$h(U \cap S) = h(U) \cap \big(\mathbb{R}^s_- \times \{0\}^{n-s}\big).$$

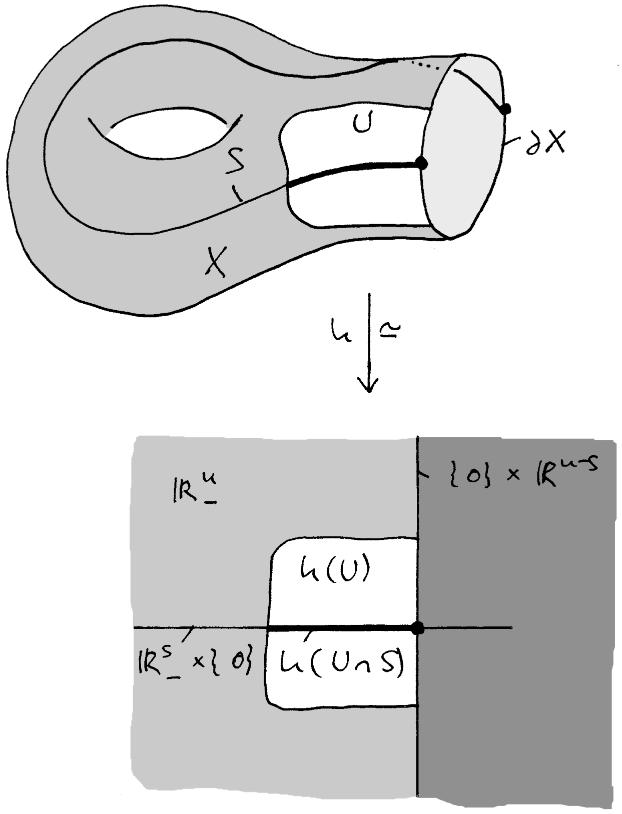

(5) The tangent space of an n-manifold X at a boundary point $o \in \partial X$ is defined exactly as for interior points o, and likewise is an n-dimensional real vector space (not a half-space). Let (U, h) and (V, k) be charts for X at o. A pointed out in (4) the transition map $k \circ h^{-1}$ may be restricted to a diffeomorphism

$$h(U \cap V) \cap (\{0\} \times \mathbb{R}^{n-1}) \xrightarrow{\simeq} k(U \cap V) \cap (\{0\} \times \mathbb{R}^{n-1});$$

therefore its Jacobian has the form

$$D(k \circ h^{-1})(h(o)) = \left(\begin{array}{c|c} + & 0 \;\cdots\; 0 \\ \hline * & \\ \vdots & D' \\ * & \end{array}\right)$$

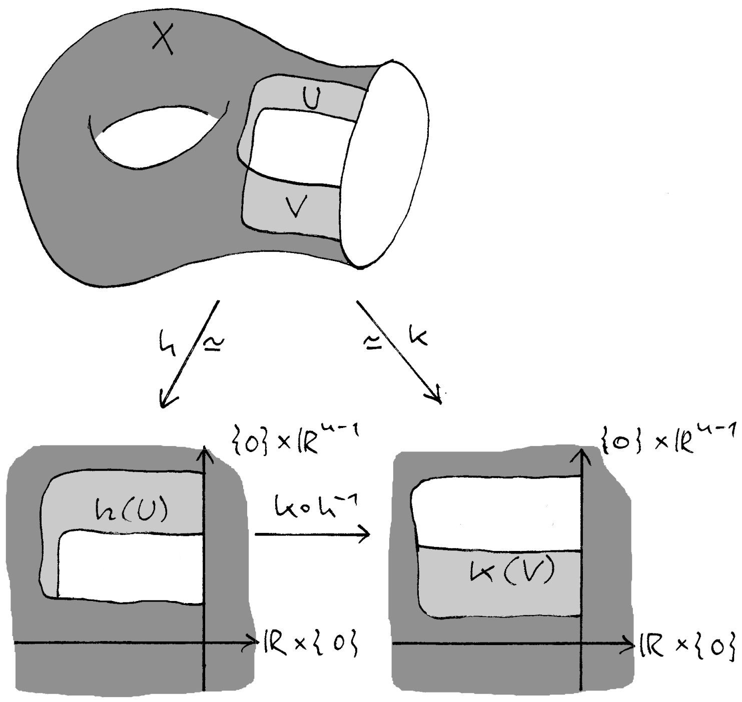

and contains as a submatrix the Jacobian $D' \in GL(n-1, \mathbb{R})$ of the transition map for ∂X. Furthermore the first component of the transition map is non-positive throughout, which implies that the top left entry of $D(k \circ h^{-1})(h(o))$ is positive. From these special properties of the Jacobian we infer that the tangent space T_oX carries two additional structures:

- the hyperplane $T_o(\partial X) \subset T_oX$, and
- a preferred connected component of $T_oX \setminus T_o(\partial X)$ comprising the *outward pointing* vectors.

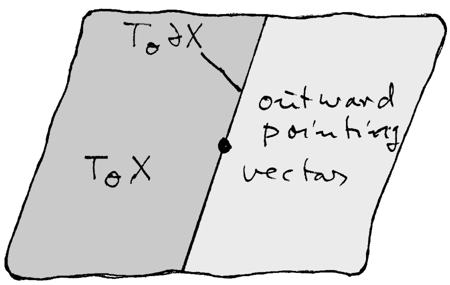

The latter are just those tangent vectors whose representation with respect to one, and thus every chart at o, has a positive first component.

Globally speaking we have the codimension one subbundle $T(\partial X) \subset TX \,|\, \partial X$. It is easy to construct a section of the bundle $TX \,|\, \partial X$ which points outward everywhere: such sections clearly exist locally, and they may be glued using a partition of unity on ∂X. The point here is that every convex linear combination of outward pointing vectors again points outward.

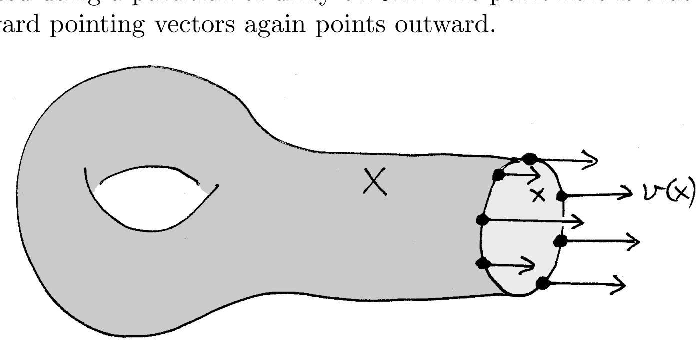

In particular such a section $\nu\colon \partial X \to TX \,|\, \partial X$ vanishes nowhere, and thus defines an embedding

$$\partial X \times \mathbb{R} \ni (x,t) \longmapsto t \cdot \nu(x) \in TX \,|\, \partial X$$

of the product line bundle. By Theorem 3.16 its image is a trivial rank one subbundle of $TX \,|\, \partial X$ and furthermore a complement to $T(\partial X) \subset TX \,|\, \partial X$. In other words the section ν determines a bundle isomorphism

$$TX \,|\, \partial X \simeq (\partial X \times \mathbb{R}) \oplus T(\partial X).$$

(6) The notion of orientation carries over to manifolds with boundary. Just the one-dimensional case requires a minor adjustment since there need not exist an oriented chart at a boundary point: thus either charts that map into open subsets of $[0,\infty)$ rather than $\mathbb{R}_- = (-\infty, 0]$ must also be allowed, or both orientation-preserving and -reversing charts must be used, keeping track of orientation signs. Given an orientation of the n-manifold X a boundary orientation is induced on ∂X simply by restricting all oriented charts of X to the boundary: these restrictions form an oriented atlas for ∂X.

The calculus of differential forms carries over to manifolds with boundary without any problem, as does the connection between orientations and volume forms. A volume form $\omega \in \Omega^n X$ does not in a canonical way induce a volume form on ∂X, but it does so once an everywhere outward pointing section ν of the bundle $TX \,|\, \partial X$ has been selected: the formula

$$\omega'_x(v_2, \ldots, v_n) = \omega_x(\nu(x), v_2, \ldots, v_n) \quad \text{for all } x \in \partial X \text{ and } v_2, \ldots, v_n \in T_x\partial X$$

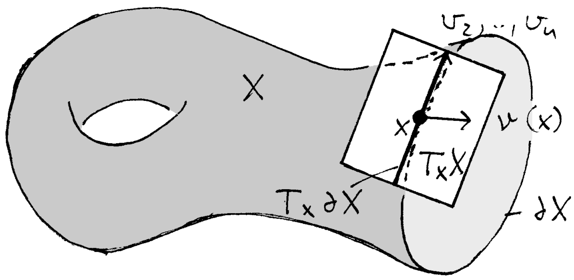

then defines a volume form $\omega' \in \Omega^{n-1}(\partial X)$ on ∂X. Evaluation on an oriented base $(v_2, \ldots, v_n)$ in terms of a chart at x at once verifies not only that $\omega'_x \neq 0$, but also that ω' induces on ∂X the boundary orientation inherited from X (oriented by the form ω).

The following theorem is generally regarded as the culmination of integral calculus on manifolds.

7.4 Stokes' Theorem Let X be a smooth oriented n-manifold with boundary. Then the integral formula

$$\int_X d\varphi = \int_{\partial X} \varphi$$

holds for all differential forms $\varphi \in \Omega^{n-1}X$ with compact support.

Proof In nearly all cases[2] the manifold X admits an atlas consisting of oriented charts (U, h) such that $h(U)$ either is the open unit cube $(-1,0) \times (0,1)^{n-1} \subset \mathbb{R}^n$ — let us call such charts of interior type — or the relatively open cube $(-1,0] \times (0,1)^{n-1} \subset \mathbb{R}^n_-$: call these of boundary type. We first make the additional assumption that the support of φ is completely contained in the domain U of one of these charts, so that we can express

$$\varphi = \sum_{j=1}^{n} f_j \cdot dh_1 \wedge \cdots \wedge \widehat{dh_j} \wedge \cdots \wedge dh_n$$

[2] One-dimensional oriented manifolds do not admit an oriented chart at a negatively oriented boundary point: a minor modification of the argument will take care of this.

in terms of h. Since φ has compact support the functions $f_j \circ h^{-1}$ have compact support contained in $h(U)$, and therefore extend by zero to C^∞ functions on all $\mathbb{R}^n_-$. We use these functions in order to independently evaluate both integrals of Stokes' formula.

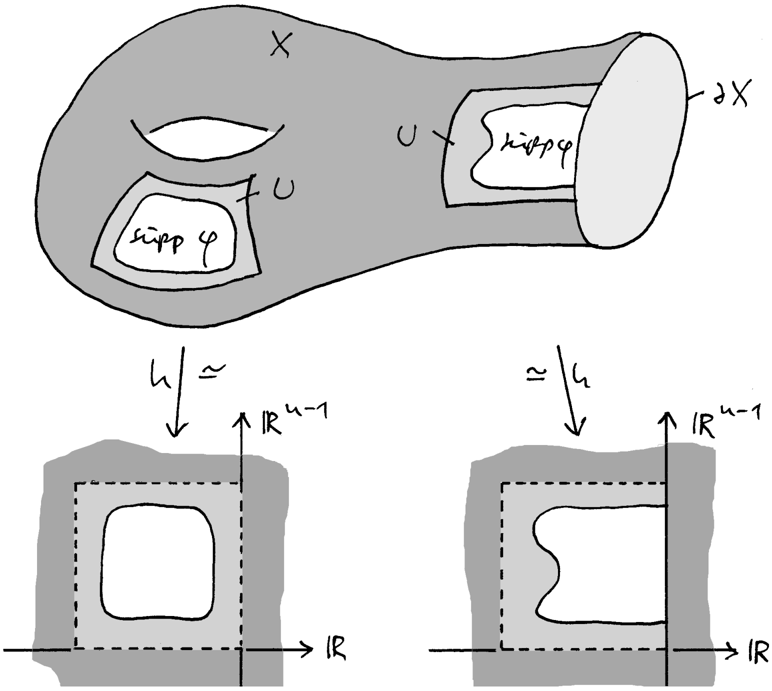

The form

$$d\varphi = \sum_{j=1}^{n} (-1)^{j-1} \frac{\partial f_j}{\partial h_j} \cdot dh_1 \wedge \cdots \wedge dh_n$$

by definition integrates to

$$\int_X d\varphi = \int_U d\varphi = \sum_{j=1}^{n} (-1)^{j-1} \int_{h(U)} \frac{\partial f_j}{\partial h_j} \cdot dh_1 \wedge \cdots \wedge dh_n.$$

Fubini's theorem converts the j-th integral into the multiple integral

$$\int_{h(U)} \frac{\partial f_j}{\partial h_j} \cdot dh_1 \wedge \cdots \wedge dh_n = \int_0^1 \cdots \int_0^1 \int_{-1}^0 \frac{\partial f_j}{\partial h_j}\, dh_1 \, \ldots \, dh_n.$$

Performing the j-th integration first we distinguish the cases of $j > 1$ from that of $j = 1$. In the former

$$\int_0^1 \frac{\partial f_j}{\partial h_j}\, dh_j = f_j(h_1, \ldots, h_n)\Big|_{h_j=0}^{1} = 0$$

vanishes since the integrand has compact support contained in the open interval $(0, 1)$. The same conclusion holds for $j = 1$ if the chart h is of interior type. If not then the support is contained in $(-1,\, 0\,]$, and there remains the contribution

$$\int_{-1}^0 \frac{\partial f_1}{\partial h_1}\, dh_1 = f_1(h_1, \ldots, h_n)\Big|_{h_j=-1}^{0} = f_1(0, h_2, \ldots, h_n).$$

Thus the overall result is

$$\int_X d\varphi = \int_0^1 \cdots \int_0^1 f_1(0, h_2, \ldots, h_n)\, dh_2 \, \ldots \, dh_n.$$

The boundary integral is immediately evaluated: Note first that the domain of integration is empty unless h is of boundary type. In that case the function h_1 is constant on $U\cap\partial X$, so that $dh_1|(U\cap\partial X)$ is the zero form. Therefore we have

$$\int_{\partial X}\varphi=\int_{U\cap\partial X}\varphi=\int_{h(U\cap\partial X)}f_1\cdot dh_2\wedge\cdots\wedge dh_n=\int_0^1\cdots\int_0^1 f_1(0,h_2,\ldots,h_n)\,dh_2\,\ldots\,dh_n.$$

This completes the proof for forms with support contained in a chart of interior or boundary type. Since these charts constitute an open cover of X we find a subordinate partition of unity $(\rho_j)_{j=0}^{\infty}$. If now φ is an arbitrary form with compact support then what we have shown applies to each form $\rho_j\cdot\varphi$. Only finitely many of them are not the zero form, therefore we may simply sum up, and conclude that Stokes' integral formula holds for φ as well. □

7.5 Examples (1) Consider the 1-manifold $X=[0,1]$ with the standard orientation (which makes the translation $X\supset(0,1]\ni t\mapsto t-1\in(-1,0]\subset\mathbb{R}^1_-$ an oriented chart). A differential form $\varphi\in\Omega^0X$ is just a smooth function on $[0,1]$, and since the oriented boundary of X is $\partial[0,1]=\{1\}-\{0\}$ Stokes' formula reads

$$\int_0^1\varphi'(t)\,dt=\int_{[0,1]}d\varphi=\int_{\partial[0,1]}\varphi=\varphi(1)-\varphi(0).$$

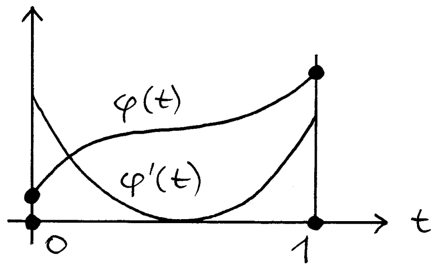

Thus the case of dimension 1 reduces to the fundamental theorem of calculus. This does not surprise since the latter is at the heart of the proof of Stokes' formula.

(2) If more generally $f\colon[0,1]\to X$ is a smooth curve on any manifold X, and $\varphi\in\Omega^1X$ a differential form of degree 1 then the *curve*, *path*, or *contour integral*

$$\int_0^1 f^*\varphi=\int_{[0,1]}f^*\varphi$$

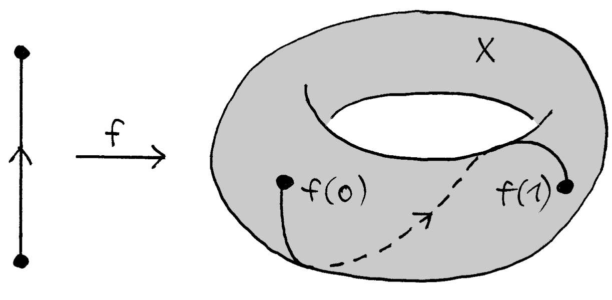

is defined. If φ happens to be exact, say $\varphi=d\chi$, then the integral may be evaluated as

$$\int_0^1 f^*d\chi=\int_{[0,1]}d\,f^*\chi=\int_{\partial[0,1]}f^*\chi=\chi\circ f\,\Big|_0^1=\chi\Big|_{f(0)}^{f(1)}$$

and thus depends only on the values of χ at the initial and end points of the curve f.

(3) Let X be a compact n-manifold without boundary — for conciseness such manifolds are often called *closed*[3] manifolds. Let an orientation of X be fixed, and let $f, g\colon X \to Y$ be smooth mappings into another manifold Y. Assume there exists a *homotopy* from f to g, that is a smooth mapping $F\colon [0,1]\times X \to Y$ such that $F(0,x) = f(x)$ and $F(1,x) = g(x)$ for all $x \in X$. Finally let $\varphi \in \Omega^n Y$ be a closed differential form — in many interesting cases Y has the same dimension $n = \dim X$, so that closedness becomes automatic.

Under these assumptions $\int_X f^*\varphi = \int_X g^*\varphi$.

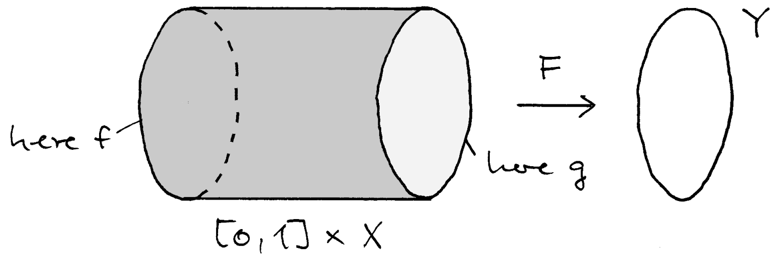

Indeed we have $d\varphi = 0$ by assumption, and if we apply Stokes' theorem to the pull-back form $F^*\varphi \in \Omega^n\big([0,1]\times X\big)$ we obtain

$$\begin{aligned} 0 &= \int_{[0,1]\times X} F^* d\varphi = \int_{[0,1]\times X} d(F^*\varphi) \\ &= \int_{\partial[0,1]\times X} F^*\varphi = \int_{\{1\}\times X} F^*\varphi - \int_{\{0\}\times X} F^*\varphi = \int_X g^*\varphi - \int_X f^*\varphi. \end{aligned}$$

(4) Let X be an oriented $(n+1)$-manifold with boundary, and let $S \subset X$ be a finite subset with $S \cap \partial X = \emptyset$. Let further φ be a closed n-form on X with "singularities" in S, which simply means that it is not defined there: $\varphi \in \Omega^n(X \setminus S)$. Then the boundary integral $\int_{\partial X} \varphi$ turns out to be completely determined by data which is concentrated near S. To make this precise, pick for each $s \in S$ a compact neighbourhood $D_s \subset X \setminus \partial X$ (arbitrarily small if desired) which is diffeomorphic to a ball. Then

$$\int_{\partial X} \varphi = \sum_{s\in S} \int_{\partial D_s} \varphi.$$

This follows at once, applying Stokes' Theorem to the "Emmental cheese" $X' := X \setminus \bigcup_{s\in S} D_s^\circ$: note that the integrals on the right refer to the orientation of the sphere ∂D_s as the boundary of D_s, while the orientation of ∂D_s as a boundary component of X' is the opposite one.

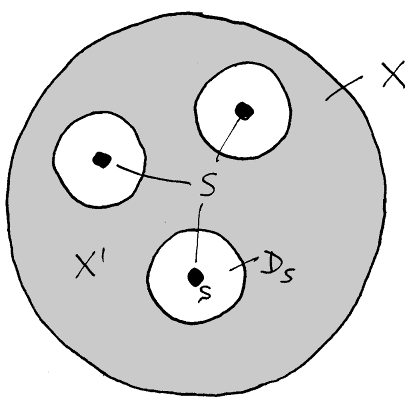

[3] Make sure to note that this terminology has nothing to do with the notion of closed subset of a topological space. Unfortunately it introduces some ambiguity into statements involving closed submanifolds.

As the importance of Stokes' theorem reaches far beyond pure mathematics one would wonder whether the C^∞ smoothness assumptions can be relaxed. In fact we have been working with C^∞ for mere convenience, and closer inspection would reveal that much less is sufficient:

7.6 Addendum Stokes' theorem holds for C^1 differential forms on a C^2 manifold rather than C^∞ data. □

It is also possible to allow differential forms with singularities, that is, forms defined on the complement of a certain subset of the base manifold. The essential points are that this singular subset should have codimension at least two and that the growth of the differential form and its Cartan derivative near the singular points be sufficiently moderate. Of course this would have to be made precise, but we do not go into the details.

In yet another interesting direction one may allow for singularities of the manifold boundary. One instance would be manifolds with corners as they occur as Cartesian products of manifolds with boundary, and it turns out that Stokes' theorem extends entirely to this and similar cases. Rather than formulate a general theorem let us investigate one example.

7.7 Example The quadrant $X = [0,\infty)^2 \subset \mathbb{R}^2$ is not a smooth manifold with boundary. Removing the origin we do obtain a manifold with boundary X', but restricting a given 1-form φ to X' in general will destroy compactness of its support. To re-establish it we choose a table mountain $\tau\colon \mathbb{R}^2 \to [0,1]$ at the origin, say of inner radius 1 and outer radius 2, put $\tau_r(x) = \tau(x/r)$ for radii $r > 0$ yet to be determined, and apply Stokes' theorem to the form $(1-\tau_r)\cdot\varphi$:

$$\int_{X'} d\big((1-\tau_r)\,\varphi\big) = \int_{\partial X'} (1-\tau_r)\,\varphi$$

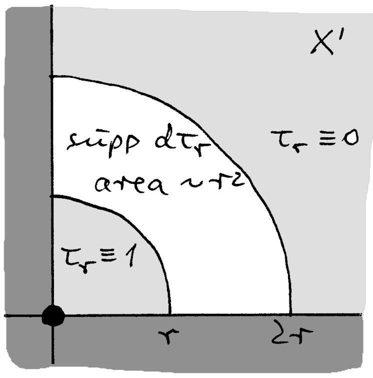

The surface integral expands into

$$\int_{X'} d\big((1-\tau_r)\,\varphi\big) = -\int_{X'} d\tau_r\wedge\varphi + \int_{X'} (1-\tau_r)\,d\varphi\,,$$

and we observe that if we let run r through a sequence with limit zero the first integrand also tends to zero while the second converges to $d\varphi$. We show that in each case the pointwise convergence is dominated by an integrable function. Indeed the coefficients of the C^∞ forms φ and $d\varphi$ are bounded, so that the second integrand is dominated by a constant multiple of the characteristic function of $\operatorname{supp}\varphi$. The same incidentally is true of the line integral $\int_{\partial X'}(1-\tau_r)\,\varphi$ on the right hand side of Stokes' formula, which therefore converges to $\int_{\partial X'}\varphi$. The case of the first surface integral is slightly more delicate, since

$$(d\tau_r)_x = \frac{1}{r}(d\tau)_{x/r}$$

is not bounded as r tends to zero. On the other hand the support of this expression is contained in the quarter annulus $X' \cap D_{2r}(0) \setminus U_r(0)$, and we do find a constant c such that the 2-form $d\tau_r\wedge\varphi$ is bounded by c/r. If we now let run r through the set $\{2^{-k} \mid k \in \mathbb{N}\}$ the function $g\colon X' \cap D_2(0) \to \mathbb{R}$ defined by

$$g(x) = 2^k \cdot c \quad \text{if } 2^{-k} < |x| \le 2^{-k+1}$$

dominates the sequence of integrands $\big(d\tau_r\wedge\varphi \,|\, r=2^{-k}\big)_{k=0}^{\infty}$. Since the area of $X'\cap D_{2r}(0)\smallsetminus U_r(0)$ is proportional to r^2 evaluating the integral of g leads to a convergent geometric series. Therefore g is integrable, the integrals of $d\tau_r\wedge\varphi$ converge to zero, and Stokes' formula holds for φ.

In low-dimensional situations in Euclidean spaces Stokes's theorem obviously can be expressed in the terms of classical vector analysis in order to recover the nineteenth century integral formulae named after Gauß, Green, and others — see Exercise 7.8. But nowadays this would seem like doing number theory using Roman numerals ...

Exercises

7.1 Let X be a manifold with boundary, and let $f\colon X\to\mathbb{R}$ a smooth function without critical values. Prove that any point where f takes its smallest value must lie on the boundary ∂X.

7.2 Let X be a compact n-manifold without boundary that admits a fixed point free involution, that is, a self-diffeomorphism $h\colon X\to X$ with $h\circ h=\mathrm{id}$ and $h(x)\neq x$ for all $x\in X$. Prove that there exist a compact $(n{+}1)$-manifold W with boundary and a diffeomorphism $X\simeq\partial W$.

7.3 Let $X\subset\mathbb{R}^3$ be a 3-dimensional submanifold with boundary.

- Explain why the surface ∂X carries a canonical area form $\omega'\in\Omega^2(\partial X)$, which generalises that of Exercise 6.1.
- Let $W\xrightarrow{f}\mathbb{R}^3$ be an orientation preserving embedding of an open subset $W\subset\mathbb{R}^2$ in $\mathbb{R}^3$ with image $f(W)=\partial X$ — such an f often is called a *parametrization* of ∂X. Show that

$$f^*\omega'=\left|\frac{\partial f}{\partial u}\times\frac{\partial f}{\partial v}\right|\cdot du\wedge dv$$

where $\mathbb{R}^3\times\mathbb{R}^3\xrightarrow{\times}\mathbb{R}^3$ denotes the classical vector product.
- How does this simplify if f is a graph mapping $(u,v)\mapsto f(u,v)=\big(u,v,\varphi(u,v)\big)$?
- For $0<r<R$ compute the area of the embedded torus

$$\big\{\big((R{+}r\cos\theta)\cos\varphi,(R{+}r\cos\theta)\sin\varphi,r\sin\theta\big)\in\mathbb{R}^3\;\big|\;\varphi,\theta\in\mathbb{R}\big\}\,.$$

7.4 For fixed $n\in\mathbb{N}$ and real $R>0$ let $D:=D_R(0)\subset\mathbb{R}^n$ be the closed n-dimensional ball of radius R. Describe the area form — that is, the $(n-1)$-dimensional volume form induced by the canonical outward pointing unit field — on the boundary $S:=S_R(0)$, and prove the identity

$$n\cdot\mathrm{volume}(D)=R\cdot\mathrm{area}(S).$$

7.5 Let X be a non-empty orientable compact n-manifold (without boundary). Prove that $H^nX\neq\{0\}$.

7.6 Let $X=S^1\times S^1$ be the torus. Prove that the multiplication $H^1X\times H^1X\xrightarrow{\wedge}H^2X$ introduced in Exercise 6.8 is non-trivial.

7.7 Let X be a non-empty orientable compact manifold with boundary. Prove that there cannot exist a smooth mapping $f\colon X \to \partial X$ such that $f|\partial X = \mathrm{id}$.

7.8 Let $Y \subset \mathbb{R}^3$ be a compact oriented surface, that is, a two-dimensional smooth submanifold. Note that Y carries the unit normal field $\nu\colon Y \to T\mathbb{R}^3$ determined by the property that for every $y \in Y$ the vector $\nu(y) \in T_y\mathbb{R}^3 = \mathbb{R}^3$ makes $\big(\nu(y), v, w\big)$ a positively oriented orthonormal base if (v, w) is a positively oriented orthonormal base of T_yY, as in Exercise 7.3. Also recall that using ν the canonical volume form ω of $\mathbb{R}^3$ induces a volume (area) form ω' on Y. In nineteenth century terminology $\omega = dV$ is called the *volume element*, and $\omega' = dS$ the *surface* or *area element* of Y. They occur in the classical integral theorems:

• *Divergence theorem* Here $X \subset \mathbb{R}^3$ is a smooth compact three-dimensional submanifold and $Y = \partial X$. The theorem states that for every smooth $F\colon X \to \mathbb{R}^3$ — a "vector field" in that terminology — the formula

$$\int_X \mathrm{div} F\, dV = \int_{\partial X} F \bullet \nu\, dS$$

holds. Show that this is just the result of applying Theorem 7.4 to the situation.

• *Theorem of Kelvin-Stokes* Forget X and allow the surface Y to have a boundary. The latter is an oriented closed curve which carries not only an induced volume form — the *line element* ds — but also the unit vector field $\tau\colon \partial Y \to T\partial Y$ which assigns to $y \in \partial Y$ the positively oriented tangent vector of unit length in $T_y\partial Y$. Then for every "vector field" F which is defined on some open neighbourhood of Y in $\mathbb{R}^3$

$$\int_Y \mathrm{rot} F \bullet \nu\, dS = \int_{\partial Y} F \bullet \tau\, ds.$$

Again show that this is what results from Theorem 7.4.

8 Vector Fields and Flows

The aim of this section is to translate the basic theory of ordinary differential equations to the setting of smooth manifolds. Recall that manifolds have no boundary unless the contrary is stated explicitly.

8.1 Definition Let X be a manifold. A smooth section $\xi\colon X \to TX$ of the tangent bundle is called a *vector field* on X.

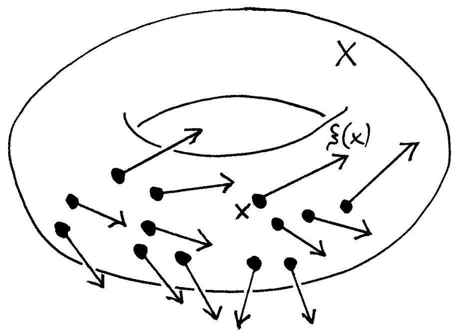

8.2 Definition Let X be a manifold. A *global flow* on X is a (smooth) mapping

$$\Phi\colon \mathbb{R} \times X \longrightarrow X$$

that satisfies the flow axioms

- $\Phi(0, x) = x$ for all $x \in X$;
- $\Phi\big(s, \Phi(t, x)\big) = \Phi(s+t, x)$ for all $s, t \in \mathbb{R}$ and $x \in X$.

For a given point $x \in X$ the partial map

$$\mathbb{R} \ni t \stackrel{\Phi_x}{\longmapsto} \Phi(t, x) \in X$$

is a smooth curve called the *flow line through* x (or having x as its *initial point*). Via these flow lines every global flow Φ on X produces a vector field $\xi\colon X \to TX$ by

$$\xi(x) = \dot{\Phi}_x(0) = \frac{d}{dt}\Phi(t, x)\bigg|_{t=0} \in T_x X,$$

which is called the *velocity field* of Φ.

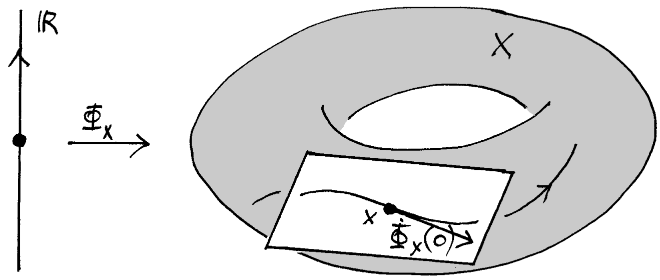

Remarks To aid intuition think of $t \in \mathbb{R}$ as a time parameter. — In different terminology the flow axioms state that Φ is an *action* of the additive group $\mathbb{R}$ on the set X. In terms of the partial maps

$$\Phi_t \colon X \longrightarrow X \quad \text{defined by } \Phi_t(x) = \Phi(t,x)$$

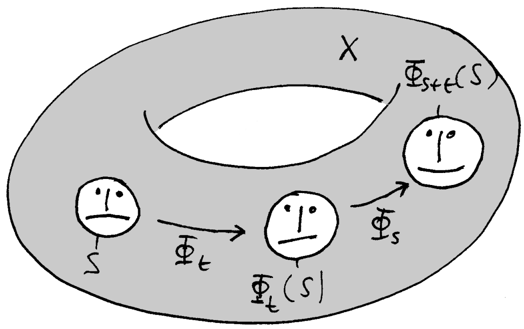

they may be restated as $\Phi_0 = \mathrm{id}_X$ and $\Phi_s \circ \Phi_t = \Phi_{s+t}$, and we see in particular that for every $t \in \mathbb{R}$ the map Φ_t is a diffeomorphism with inverse Φ_{-t}. — The theoretically conflicting notations Φ_x and Φ_t pose no problem in practice if one avoids time-like names as s or t for “space” variables.

8.3 Examples (1) For every fixed vector $b \in \mathbb{R}^n$ the assignment $(t,x) \mapsto t\,b + x$ defines a global flow on the manifold $\mathbb{R}^n$; in terms of the identification $T\mathbb{R}^n = \mathbb{R}^n \times \mathbb{R}^n$ its velocity field has the constant value b.

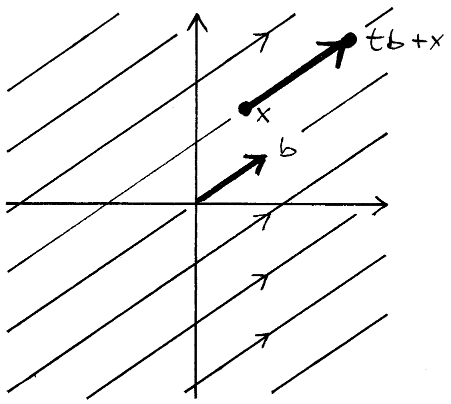

(2) The same formula defines a global flow on the n-dimensional torus $\mathbb{R}^n/\mathbb{Z}^n$.

(3) If $a \in \mathrm{Mat}(n\times n, \mathbb{R})$ is any matrix then

$$\Phi(t,x) := e^{ta} \cdot x = \sum_{k=0}^{\infty} \frac{t^k a^k}{k!} \cdot x$$

defines a global flow on $\mathbb{R}^n$. If a happens to be skew-symmetric then the matrices e^{ta} are orthogonal:

$$(e^{ta})^{\mathrm{t}} \cdot e^{ta} = e^{ta^{\mathrm{t}}} \cdot e^{ta} = e^{-ta} \cdot e^{ta} = e^0 = 1.$$

In this case Φ may be restricted to give a global flow $\mathbb{R}\times S^{n-1} \to S^{n-1}$ on the sphere. In either case the corresponding velocity field assigns the tangent vector $a \cdot x \in \mathbb{R}^n = T_x\mathbb{R}^n$ to the point x.

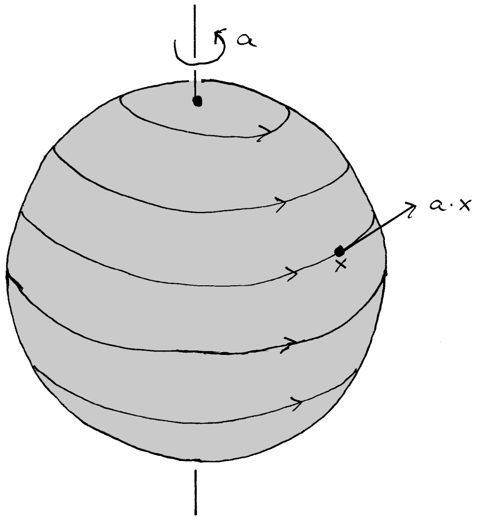

(4) In terms of a chart (U,h) a vector field ξ on an n-manifold X reduces to a smooth function $v\colon U \to \mathbb{R}^n$ via

$$v(x) = \big(\xi(x)\big)_h = T_x h \cdot \xi(x).$$

By contrast any attempt to represent a global flow Φ in terms of (U,h) encounters the problem that restricting a flow to the open subset $U \subset X$ will not result in a global flow on U but just a smooth mapping

$$\mathbb{R}\times U \supset (\mathbb{R}\times U)\cap \Phi^{-1}U \xrightarrow{\Phi|\ldots} U.$$

The problem already occurs if U is formed by removing a single point x from the manifold X which does not happen to be a *fixed point*, that is one with $\Phi(t,x) = x$ for all $t \in \mathbb{R}$.

This last example motivates to complement the notion of global flow by a weaker version which does allow restriction to arbitrary open subsets.

8.4 Definition Let X be a manifold. A (not necessarily global) *flow* on X is a smooth mapping

$$\mathbb{R}\times X \supset D \xrightarrow{\Phi} X$$

where

- $D \subset \mathbb{R}\times X$ is open,
- for each $x \in X$ the open set $\big\{t \in \mathbb{R} \,\big|\, (t,x) \in D\big\}$ is an interval containing $0 \in \mathbb{R}$, and
- the flow axioms $\Phi(0,x) = x$ and $\Phi\big(s, \Phi(t,x)\big) = \Phi(s+t,x)$ hold, the latter for all $s,t \in \mathbb{R}$ such that both sides of the identity are defined.

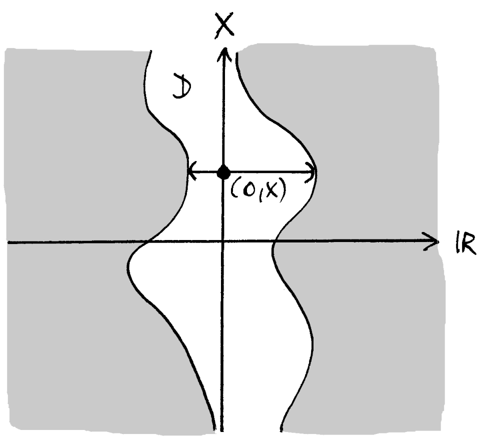

Notes (1) For fixed $x \in X$ the pairs (s,t) such that $\Phi\big(s,\Phi(t,x)\big)$ and $\Phi(s{+}t,x)$ are defined are determined by the conditions

$$(t,x) \in D\,,\ \big(s,\Phi(t,x)\big) \in D\,, \text{ and } (s{+}t,x) \in D$$

and thus form an open neighbourhood of $(0,0)$ in $\mathbb{R}^2$.

(2) It now makes sense to restrict a flow $\Phi\colon D \to X$ to an open subset $U \subset X$: the subset

$$D' := \big\{(t,x) \in D \,\big|\, [\,0,t] \times \{x\} \subset \Phi^{-1}U\big\} \subset \mathbb{R} \times U$$

— for $t < 0$ read $[\,0,t]$ as $[t,0\,]$ — is open and has the required geometry, and Φ restricts to a mapping $\Phi'\colon D' \to U$ which is a flow on U.

(3) If $\Phi\colon D \to X$ is a flow then its flow line through a point $x \in X$ still is a smooth curve

$$(\alpha_x,\omega_x) \xrightarrow{\Phi_x} X$$

with $(\alpha_x,\omega_x) := \big\{t \in \mathbb{R} \,\big|\, (t,x) \in D\big\}$, but its life span (into the past, the future, or both) may be limited if Φ is not global.

(4) While for a non-global flow the diffeomorphisms Φ_t need not be defined the velocity field of Φ still is, since for every $x \in X$ the domain D is a neighbourhood of $(0,x)$ in $\mathbb{R}{\times}X$.

8.5 Lemma Let $\Phi\colon D \to X$ be a flow on X with velocity field $\xi\colon X \to TX$, and let $\Phi_x\colon(\alpha_x,\omega_x) \to X$ be the flow line through a point $x \in X$. Then

$$\dot\Phi_x(t) = \xi\big(\Phi_x(t)\big) \quad \text{for all } t \in (\alpha_x,\omega_x).$$

In other words Φ_x is a solution of the *ordinary differential equation* $\dot\varphi = \xi \circ \varphi$ for φ with initial value $\varphi(0) = x$. Therefore the flow lines are also called *integral curves* of the velocity field.

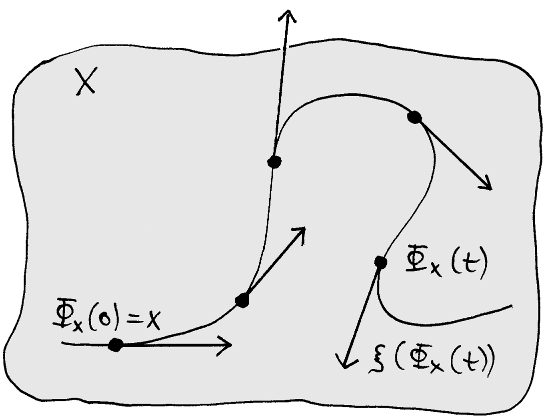

Proof For fixed t and x differentiate the flow identity $\Phi(\tau{+}t,x) = \Phi\big(\tau,\Phi(t,x)\big)$ at $\tau = 0$. □

We will take the local theory of ordinary differential equations for granted. All that we need is contained in the following theorem, which is a standard result and which we state without proof as our starting-point.

8.6 Fundamental Theorem on Ordinary Differential Equations Let $X \subset \mathbb{R}^n$ be open and $v\colon X \to \mathbb{R}^n$ a C^1 map. Then we have:

- *Uniqueness of solutions* — If the C^1 curves $\varphi\colon I \to X$ and $\chi\colon J \to X$ defined on open intervals containing $0 \in \mathbb{R}$ satisfy $\dot\varphi(t) = v(\varphi(t))$ and $\dot\chi(t) = v(\chi(t))$ for all t, as well as $\varphi(0) = \chi(0)$ then φ and χ agree on some neighbourhood of 0 in $I \cap J$.

• *Existence of solutions and their dependence on the initial condition* — For every $x \in X$ there exist a $\delta > 0$, an open neighbourhood $U \subset X$ of x, and a C^1 mapping

$$\Phi\colon (-\delta, \delta) \times U \to X$$

such that $\frac{d}{dt}\Phi(t,x) = v\big(\Phi(t,x)\big)$ and $\Phi(0,x) = x$ for all $t \in (-\delta,\delta)$ and $x \in U$. If v is even C^k differentiable for some larger k then so is Φ. □

8.7 Maximal Integral Curves of a Vector Field Let X be a manifold and ξ a vector field on X, and let $x \in X$ be a point. A (weak) application of the existence part of the fundamental theorem shows that there exists at least one curve $\varphi\colon (-\delta,\delta) \to X$ solving the differential equation $\dot\varphi = \xi \circ \varphi$ with initial condition $\varphi(0) = x$: expressing ξ in any chart at x takes us straight to the situation of 8.6.

Now consider two integral curves $\varphi\colon I \to X$ and $\chi\colon J \to X$, both defined on open intervals containing $0 \in \mathbb{R}$. The subset

$$T := \big\{t \in I \cap J \,\big|\, \varphi(t) = \chi(t)\big\} \subset I \cap J$$

is closed by continuity. It is also open: given $t \in T$ both

$$\tau \mapsto \varphi(\tau + t) \quad \text{and} \quad \tau \mapsto \chi(\tau + t)$$

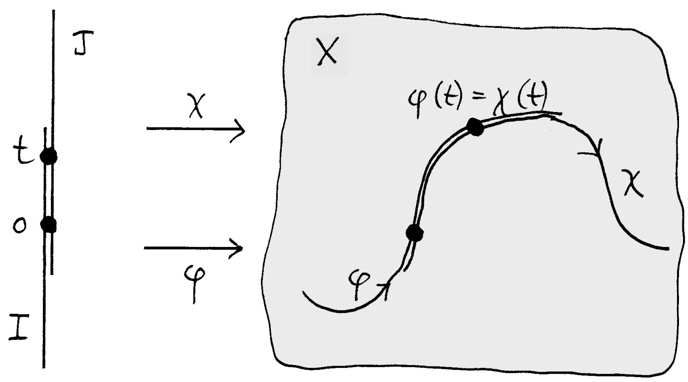

are integral curves starting at the point $\varphi(t) = \chi(t)$; thus they coincide near 0. Since the interval $I \cap J$ is connected we obtain $T = I \cap J$, so that φ and χ must coincide on $I \cap J$.

Now the definition of the *maximal*[1] *integral curve* through x

$$\varphi_x\colon (\alpha_x, \omega_x) \longrightarrow X$$

makes sense: simply take for (α_x, ω_x) the union of the domains of all integral curves through x and define $\varphi_x(t)$ as the common value of those curves which are defined at t.

We come to the central result of the section: assigning to a flow its velocity field is reversed by solving a differential equation.

8.8 Theorem Let X be a manifold with a vector field ξ. For each $x \in X$ let $\varphi_x\colon (\alpha_x, \omega_x) \to X$ denote the maximal integral curve of the differential equation $\dot\varphi = \xi \circ \varphi$ through x. Then the union

$$D := \bigcup_{x \in X} (\alpha_x, \omega_x) \times \{x\} \subset \mathbb{R} \times X$$

is open, and

$$\Phi\colon D \longrightarrow X \quad \text{defined by } \Phi(t,x) = \varphi_x(t)$$

[1] This is the usual name even though from a logical point of view it rather is the *largest* integral curve through x.

is a flow on X with velocity field ξ. Given $(t,x)\in D$ the domains of the curves φ_x and $\varphi_{\Phi(t,x)}$ are related by

$$(\alpha_{\Phi(t,x)},\omega_{\Phi(t,x)}) = (\alpha_x - t,\ \omega_x - t)$$

(with the convention that $\pm\infty - t = \pm\infty$).

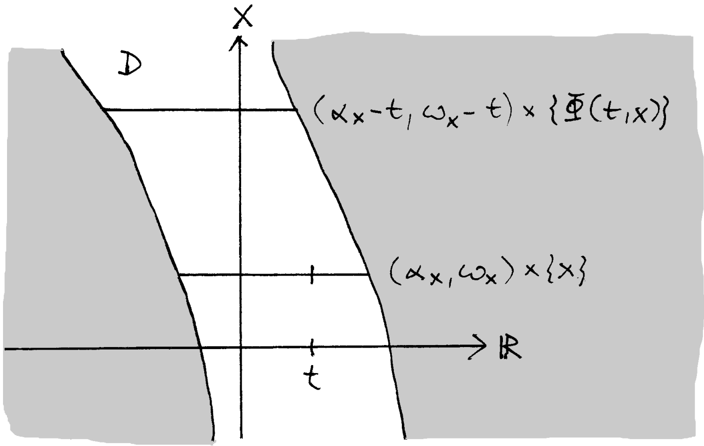

Proof The essential point is to show that D is open and that Φ is smooth. We fix a point $a\in X$ and consider the set

$$J := \left\{ t\in [\,0,\omega_a) \;\middle|\; \begin{array}{c} D \text{ contains an open neighbourhood of} \\ [\,0,t]\times\{a\}\subset \mathbb{R}\times X \text{ on which } \Phi \text{ is smooth}\end{array} \right\} \subset \mathbb{R}\,.$$

Our main goal is to prove that $J=[\,0,\omega_a)$: of course the analogous reasoning works for $(\alpha_a,0\,]$, and since the choice of $a\in X$ was arbitrary we will have shown that every point of D is an interior point and that Φ is smooth near it.

From Theorem 8.6 we know that at least $0\in J$: using a chart at a the existence part supplies a $\delta>0$, an open neighbourhood $U\subset X$ of a, and a smooth map

$$\Psi\colon (-\delta,\delta)\times U \longrightarrow X$$

with $\frac{d}{dt}\Psi = \xi\circ\Psi$ and $\Psi(0,x)=x$ for all $x\in U$; by maximality and uniqueness this further implies $(-\delta,\delta)\times U\subset D$ and $\Psi = \Phi|(-\delta,\delta)\times U$.

It is clear from the definition that $J\subset[\,0,\omega_a)$ is an open subset. We will prove that it is also closed. Thus let $t\in\overline{J}\cap[\,0,\omega_a)$ be any point of the closure: we must show that $t\in J$, and to this end may assume $t>0$. We put

$$c=\Phi(t,a)\in X$$

and apply the existence part of 8.6 there. We thus find a $\delta>0$, an open neighbourhood $V\subset X$ of c, and a smooth mapping

$$\Psi\colon(-\delta,3\delta)\times V\longrightarrow X$$

such that $\frac{d}{dt}\Psi=\xi\circ\Psi$ and $\Psi(0,y)=y$ for all $y\in V$. Since the maximal half-curve $\varphi_a\colon[\,0,\omega_a)\to X$ is continuous at t with value $\varphi_a(t)=c$ we may make $\delta>0$ smaller to achieve that $t-3\delta>0$ and

$$b:=\varphi_a(t-2\delta)\in V.$$

Now by definition of J there exists an open neighbourhood $U\subset X$ of a such that $[\,0,t-\delta)\times U\subset D$ and $\Phi\colon[\,0,t-\delta)\times U\to X$ is smooth. In particular the map

$$U\ni x \xmapsto{\Phi_{t-2\delta}} \Phi(t-2\delta,x)\in X$$

is continuous, and we may replace U by the smaller open set

$$U \cap \Phi_{t-2\delta}^{-1} V \subset X,$$

which in view of $\Phi_{t-2\delta}(a) = \varphi_a(t-2\delta) = b \in V$ still is a neighbourhood of a.

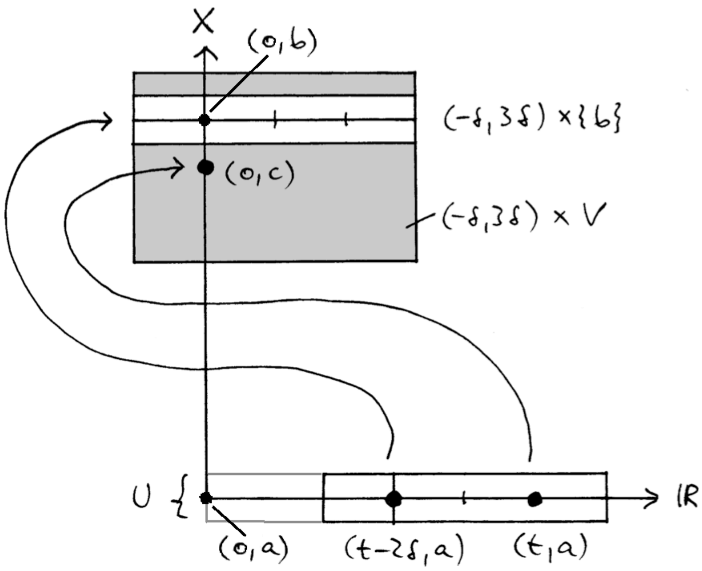

Next we define a smooth map $\tilde{\Phi}\colon [\,0,\, t+\delta) \times U \to X$ by

$$\tilde{\Phi}(\tau, x) = \begin{cases} \Phi(\tau, x) & \text{if } \tau \in [\,0,\, t-\delta), \\ \Psi\big(\tau - t + 2\delta,\, \Phi(t-2\delta, x)\big) & \text{if } \tau \in (t-3\delta,\, t+\delta). \end{cases}$$

Why do both cases give the same value on the intersection $(t-3\delta,\, t-\delta) \times U$? Consider the curves $\varphi, \chi\colon (-\delta, \delta) \to X$ given by

$$\varphi(\sigma) = \Phi(\sigma + t - 2\delta, x) \quad \text{and} \quad \chi(\sigma) = \Psi\big(\sigma,\, \Phi(t-2\delta, x)\big)$$

— they are translates of the terms defining $\tilde{\Phi}$. Both are integral curves: $\dot{\varphi} = \xi \circ \varphi$ and $\dot{\chi} = \xi \circ \chi$, with $\varphi(0) = \Phi(t-2\delta) = \chi(0)$. Hence we have $\varphi = \chi$ throughout, and are sure that $\tilde{\Phi}$ is well-defined.

By maximality and uniqueness we conclude that $[\,0,\, t+\delta) \subset D$ and $\tilde{\Phi} = \Phi|[\,0, t+\delta) \times U$, and in particular that $t \in J$ holds as desired. We thus have shown that J is an open and closed subset of the interval $[\,0, \omega_a)$. Since the latter is connected we must have $J = [\,0, \omega_a)$, and this concludes the proof of the main assertions.

We now know that $D \subset \mathbb{R} \times X$ is open and that Φ is smooth. As to the flow identities the first, $\Phi(0, x) = x$ for all $x \in X$, is obvious. For the second we fix a pair $(t, x) \in D$ and compare the curves

$$s \mapsto \varphi_{\Phi(t,x)}(s) \quad \text{and} \quad s \mapsto \varphi_x(s+t).$$

Both solve the differential equation with initial value $\Phi(t, x)$: thus by uniqueness they coincide wherever both are defined, and by maximality they also have the same interval of definition. This proves the second flow axiom, as well as the equality $(\alpha_{\Phi(t,x)}, \omega_{\Phi(t,x)}) = (\alpha_x - t,\, \omega_x - t)$.

Finally the identity $\frac{d}{dt}\Phi(t, x) = \xi\big(\Phi(t, x)\big)$, which holds by definition of Φ, reduces for $t = 0$ to the equation that establishes ξ as the velocity field of Φ. This completes the proof of the theorem. □

8.9 Definition Let X be a manifold. A flow $\mathbb{R} \times X \supset D \xrightarrow{\Phi} X$ is called *maximal*, if it allows no extension to a flow $\mathbb{R} \times X \supset \tilde{D} \xrightarrow{\tilde{\Phi}} X$ such that $\tilde{D}$ strictly contains D.

8.10 Summary Let X be a manifold. Assigning to a flow its velocity field sets up a one-to-one correspondence between maximal flows and vector fields on X; the inverse assignment sends the vector field ξ to the flow assembled from the maximal integral curves of the differential equation $\dot\varphi = \xi \circ \varphi$ — by definition this flow is maximal even in the stronger sense that none of its flow lines can be extended.

8.11 Corollaries A flow Φ is maximal if and only if every one of its flow lines is a maximal integral curve of the velocity field of Φ. — Every flow Φ extends to a unique maximal flow, namely that corresponding to the velocity field of Φ.

Remark The existence of the maximal flow corresponding to a given vector field ξ is expressed in more classical terminology by saying that every smooth vector field ξ is *integrable*. It is called *globally integrable* in case the solution flow is global.

There is a striking and extremely useful criterion that characterizes the maximality of flow lines.

8.12 Maximality Criterion Given a manifold X and a maximal flow Φ let

$$\Phi_x\colon (\alpha_x, \omega_x) \longrightarrow X$$

be the flow line through x, and for some finite positive $\omega \le \omega_x$ consider the restriction $\varphi := \Phi_x|[0,\omega)$. Then the following statements are equivalent:

- $\omega = \omega_x$;
- the mapping $[\,0,\omega) \xrightarrow{\varphi} X$ is proper[2].

Proof Assume $\omega < \omega_x$: then $L := \Phi_x\big([\,0,\omega]\big)$ is a compact subset of X while its preimage $\varphi^{-1}L = [0,\omega)$ fails to be compact, so that φ is not proper.

We now fix an arbitrary finite positive $\omega \le \omega_x$ and assume that φ is not proper. We select a compact set $L \subset X$ such that $\varphi^{-1}L$ is not compact, and therefore not contained in any compact subinterval of $[\,0,\omega)$. We thus find a sequence $(t_k)_{k=0}^{\infty}$ in $[\,0,\omega)$ such that $\lim_k t_k = \omega$ and $\varphi(t_k) \in L$ for all k.

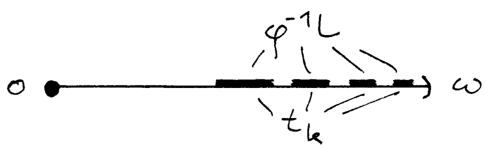

Since L is compact we may replace this sequence by a subsequence such that $y := \lim_k \varphi(t_k) \in L$ also exists.

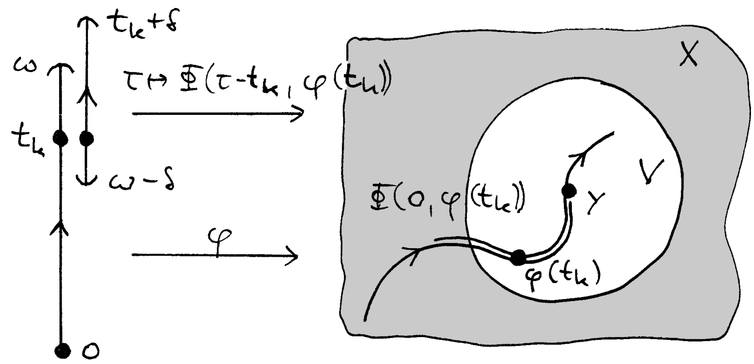

[2] A continuous map $f\colon X \to Y$ between locally compact Hausdorff spaces is *proper* if for every compact set $L \subset Y$ the pre-image $f^{-1}L \subset X$ also is compact.

We choose a number $\delta \in (0,\omega)$ and an open neighbourhood $V \subset X$ of y such that $(-\delta,\delta)\times V$ is contained in the domain of Φ, and select one $k \in \mathbb{N}$ so large that $\omega-\delta < t_k < \omega$ and $\varphi(t_k) \in V$. Then the curve $\tilde{\varphi}\colon [0, t_k+\delta) \to X$ with

$$\tilde{\varphi}(\tau) = \begin{cases} \varphi(\tau) & \text{if } \tau \in [0,\, \omega), \\ \Phi\bigl(\tau-t_k, \varphi(t_k)\bigr) & \text{if } \tau \in (\omega-\delta,\, t_k+\delta) \end{cases}$$

is well-defined since the curves

$$\sigma \mapsto \varphi(\sigma+t_k) \quad \text{and} \quad \sigma \mapsto \Phi\bigl(\sigma, \varphi(t_k)\bigr)$$

both are integral curves through $\varphi(t_k)$. In view of $t_k+\delta > \omega$ we have thus constructed a strict extension of φ as a flow line, and thereby shown that $\omega < \omega_x$. □

8.13 Corollary Every maximal flow on a compact manifold must be global.

Proof If X is compact the half-curve φ of the criterion can never be proper, thus no ω_x can be finite. Of course the same argument applies to α_x. □

The correspondence 8.10 between flows and vector fields is routinely used to construct flows with prescribed properties. The strategy is to translate such properties to equivalent ones of the corresponding vector field, where they are easier to deal with since vector fields, being sections of the tangent bundle, are *linear* objects. Going back to the solution flow will then solve the original problem. In the end this is yet another instance of the central idea of differentiation: to solve non-linear problems by treating them in linear approximations. — The following examples are held rather simple, but we will study quite a sophisticated application in the next section.

8.14 Examples (1) Let X be a manifold and $a \in X$ be a point. If you wish to construct flows Φ that leave a fixed you must look for vector fields $\xi\colon X \to TX$ with $\xi(a) = 0$: For if a is a fixed point of Φ then Φ_a is constant and the value of the velocity field $\xi(a) = \dot{\Phi}_a(0)$ certainly is zero. Conversely if $\xi(a) = 0$ then the constant curve $\mathbb{R} \ni t \xmapsto{\varphi} a \in X$ is a solution of the differential equation $\dot{\varphi} = \xi \circ \varphi$, and therefore coincides with the flow line of Φ through a. Incidentally it follows that the flow line through a fixed point always has infinite life span. The figure shows a selection of such fixed points $a \in X$.

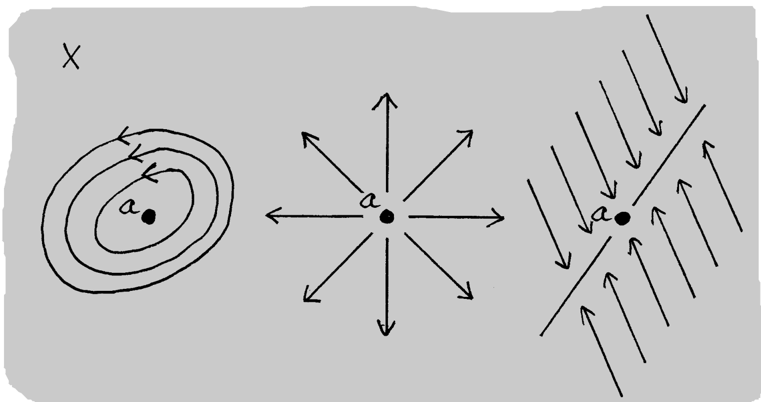

(2) Suppose a smooth function $f\colon X \to \mathbb{R}$ is given on the manifold X. If you are interested in flows Φ that leave f invariant in the sense that f is constant along each flow line Φ_x then you must look for vector fields ξ with the property

$$Tf \circ \xi = 0.$$

For if $f \circ \Phi_x$ is constant then

$$(Tf \circ \xi)(x) = T_x f \cdot \dot{\Phi}_x(0) = \frac{d}{dt}(f \circ \Phi_x)(0) = 0.$$

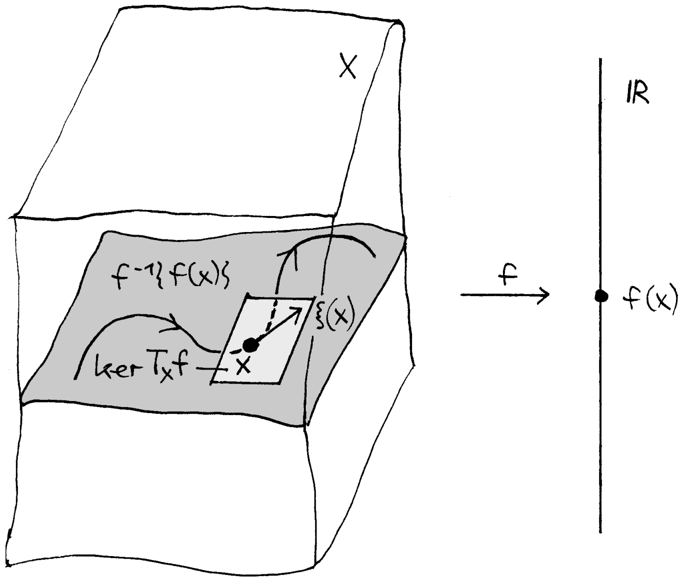

Conversely, if we assume $Tf \circ \xi = 0$ then for every $x \in X$ the composition $f \circ \Phi_x$ is a scalar function on an interval such that

$$\frac{d}{dt}(f \circ \Phi_x)(t) = T_{\Phi(t,x)} f \cdot \dot{\Phi}_x(t) = (Tf \circ \xi)\big(\Phi_x(t)\big) = 0 \quad \text{for all } t,$$

so that $f \circ \Phi_x$ solves the differential equation $\dot{\varphi} = 0$ for the scalar function φ on the real line: such a function must be constant by uniqueness of solutions (or common knowledge). This being true for every $x \in X$ it follows that Φ leaves f invariant.

We now look into the question of what can be said about the geometry of the individual flow lines. The following definition generalizes the notion of orbit as it is used in the context of group actions.

8.15 Definition Let X be a manifold, $\mathbb{R} \times X \supset D \xrightarrow{\Phi} X$ a maximal flow, and $x \in X$. The image set

$$\Phi_x(\alpha_x, \omega_x) \subset X$$

of the flow line through x is called the *orbit* of x.

8.16 Lemma Let X be a manifold, $\mathbb{R} \times X \supset D \xrightarrow{\Phi} X$ a maximal flow. Two points $x, y \in X$ have the same orbit if and only if there exists a $t \in (\alpha_x, \omega_x) \in D$ with $\Phi(t, x) = y$. Otherwise the orbits of x and y are disjoint. In particular the distinct orbits form a partition of the set X.

Proof Assume $\Phi(t, x) = y$. By Theorem 8.8 we have $(\alpha_y, \omega_y) = (\alpha_x - t,\ \omega_x - t)$ and

$$\Phi_y\big((\alpha_y, \omega_y)\big) = \Phi_{\Phi(t,x)}\big((\alpha_x - t,\ \omega_x - t)\big) = \Phi_x\big((\alpha_x, \omega_x)\big),$$

so that x and y have the same orbit. Conversely assume that the orbits of x and y intersect: we must show that $y = \Phi(t, x)$ for some t. But by what we have just seen both orbits coincide with that of the intersection point; therefore the orbits are equal. In particular y belongs to the orbit of x, and the claim follows. Finally the orbits form a partition since clearly every point of X lies in its own orbit. □

8.17 Theorem Let X be a manifold, $\mathbb{R} \times X \supset D \xrightarrow{\Phi} X$ a maximal flow, and $x \in X$ a point. Then exactly one of the following statements is true.

- The domain (α_x, ω_x) of the integral curve Φ_x is all $\mathbb{R}$, and x a fixed point of Φ.

- The point x is a not a fixed point of Φ, but $(\alpha_x, \omega_x) = \mathbb{R}$ and there exists a time $p > 0$ such that $\Phi(p, x) = x$; the flow line Φ_x induces an embedding of the circle $\mathbb{R}/p\mathbb{Z} \xrightarrow{\simeq} \Phi_x(\mathbb{R}) \subset X$ as a so-called *periodic* orbit.
- The flow line $\Phi_x: (\alpha_x, \omega_x) \to X$ is an injective immersion.

Proof Let $\xi: X \to TX$ be the velocity field of Φ. We have seen in 8.14(1) that the case $\xi(x) = 0$ corresponds to the first statement, where x is a fixed point. We now assume $\xi(x) \neq 0$. If there were a time $t \in (\alpha_x, \omega_x)$ with $\xi\big(\Phi(t, x)\big) = 0$ then $\Phi(t, x)$ would be a fixed point, and by Lemma 8.16 also x would be a fixed point, in contradiction to $\xi(x) \neq 0$. Thus applying Lemma 8.5 we have

$$\dot{\Phi}_x(t) = \xi\big(\Phi(t, x)\big) \neq 0 \quad \text{for all } t \in (\alpha_x, \omega_x),$$

and the curve Φ_x is an immersion.

Continuing with the case $\xi(x) \neq 0$ we now assume that Φ_x is not injective and prove that it is periodic. If we pick distinct times $p, q \in (\alpha_x, \omega_x)$ such that $\Phi_x(p) = \Phi_x(q)$ the intervals of definition satisfy

$$(\alpha_x - p,\ \omega_x - p) = (\alpha_{\Phi(p,x)}, \omega_{\Phi(p,x)}) = (\alpha_{\Phi(q,x)}, \omega_{\Phi(q,x)}) = (\alpha_x - q,\ \omega_x - q),$$

so that necessarily $(\alpha_x, \omega_x) = \mathbb{R}$. Therefore the restricted map

$$\mathbb{R} \times \Phi_x(\mathbb{R}) \ni (t, x) \longmapsto \Phi(t, x) \in \Phi_x(\mathbb{R})$$

is a continuous action of the real line on the orbit $\Phi_x(\mathbb{R})$, and a standard argument from the theory of group actions applies: the isotropy subgroup $\big\{t \in \mathbb{R} \,\big|\, \Phi(t, x) = x\big\}$ is a closed non-zero proper subgroup of $\mathbb{R}$, therefore of the form $p\mathbb{Z} \subset \mathbb{R}$ for a unique positive period p, and the induced map $\mathbb{R}/p\mathbb{Z} \to \Phi_x(\mathbb{R})$ is bijective. It is even a homeomorphism since $\mathbb{R}/p\mathbb{Z}$ is compact, and composing with the inclusion $\Phi_x(\mathbb{R}) \to X$ we obtain a mapping which is both a topological embedding and an immersion. By Theorem 2.6 it is an embedding of the circle as a smooth submanifold of X, and this completes the proof. □

Our examples from 8.3 nicely illustrate the three kinds of orbits. The flow of (3) — rotation of the sphere say about the polar axis — has the poles as two fixed points, while the circles of latitude, which comprise all other orbits, are periodic (incidentally all with the same period).

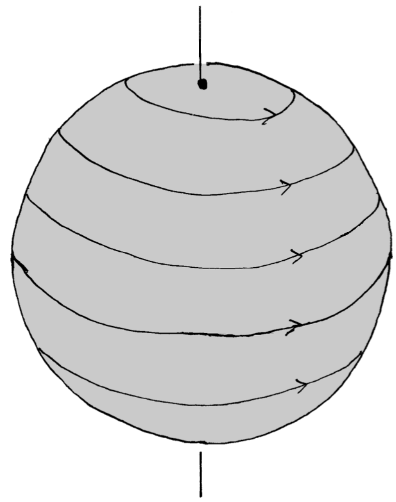

While example (1) shows the case of flow lines that embed injectively the most interesting flows are those of (2) on the torus, say with $n = 2$ and $b = (1, \beta) \in \mathbb{R}^2$. Given a rational $\beta = m/p$ every orbit is periodic, and if the denominator $p > 0$ is chosen minimal then this is the common period. But if β is irrational then every flow line is an injective immersion of the real line in the torus $\mathbb{R}^2/\mathbb{Z}^2$. Its image turns out to be a dense subset of $\mathbb{R}^2/\mathbb{Z}^2$ even though, being a countable union of 1-dimensional submanifolds it must be a null set, in particular a proper subset of $\mathbb{R}^2/\mathbb{Z}^2$.

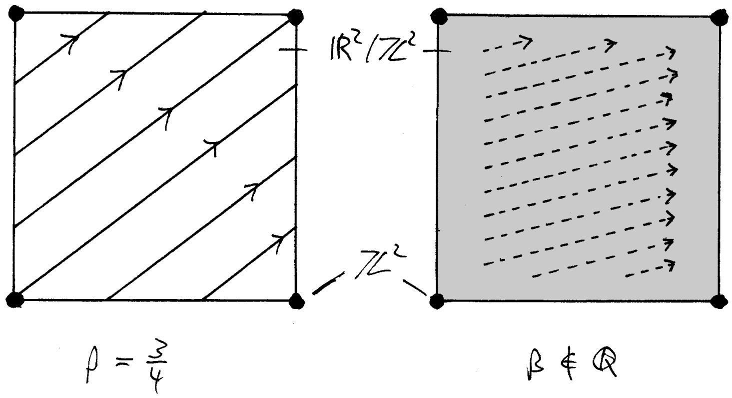

We turn to a result that bears a certain resemblance to the constant rank theorem. Rather than to the local classification of smooth mappings it pertains to that of vector fields, and states that vector fields are locally described by a single normal form wherever they do not vanish: intuitively speaking, ξ can locally be "combed". By contrast local classification of vector fields near their zeros is quite a difficult task, indeed a hopeless one unless strong simplifying conditions are imposed — quite like the task of classifying smooth maps near points of non-constant rank. For this reason the zeros of a vector field are often called its *singular points* or *singularities*.

8.18 Theorem Let ξ be a vector field on the n-manifold X, and let $o \in X$ be a point with $\xi(o) \neq 0 \in T_oX$. There exists a centred chart (U, h) at o such that ξ expressed in this chart becomes the constant mapping

$$U \ni x \longmapsto \xi(x)_h = e_1 \in \mathbb{R}^n$$

whose value is the first standard base vector $e_1 = (1, 0, \ldots, 0)$. Thus the flow that corresponds to ξ is of the type considered in Example 8.3(1), acting in terms of h by

$$(t, h) \longmapsto (t+h_1, h_2, \ldots, h_n)$$

on its domain.

Proof Since the conclusion is of local nature we may assume that $X \subset \mathbb{R}^n$ is open with $o = 0$, and that $\xi(0) = (0, e_1) \in X \times \mathbb{R}^n = TX$. Let $\mathbb{R} \times X \supset D \xrightarrow{\Phi} X$ be the maximal flow corresponding to ξ. We abbreviate

$$X' := X \cap \big(\{0\} \times \mathbb{R}^{n-1}\big)$$

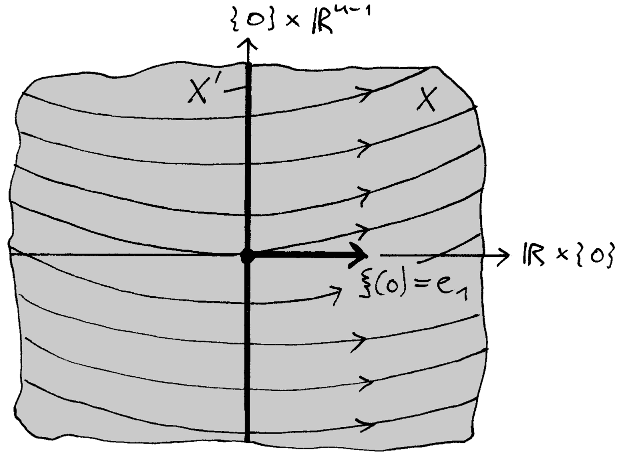

and let g denote the restriction of Φ to $D \cap (\mathbb{R} \times X')$. Then $g|(\{0\} \times X')$ sends $(0, h') \in \{0\} \times X'$ to $h' \in X' \subset X$, and in view of $\frac{\partial g}{\partial t}(0) = \xi(0) = e_1$ the Jacobian matrix of g at the origin is the unit matrix $1 \in \mathrm{Mat}(n\times n, \mathbb{R})$. By the local inverse theorem g is a local diffeomorphism at the origin, and we define the chart (U, h) as a local inverse.

We summarize the situation in the following diagram. For better readability we do not name the precise domains of the maps involved, thus the arrows stand for smooth mappings that are defined near the origins of $\mathbb{R}$, $\mathbb{R}^{n-1}$, and their Cartesian products. Points of $h(U)$ are written as $h = (h_1, \ldots, h_n) = (h_1, h')$:

$$\begin{tikzcd} \mathbb{R}\times U \arrow[r, "\Phi"] & U \\ \mathbb{R}\times h(U) \arrow[r, "\Psi"] \arrow[u, "\mathrm{id}\times g", "\simeq"'] & h(U) \arrow[u, "\simeq", "g"'] \end{tikzcd} \qquad \begin{tikzcd} \big(t, \Phi(h_1, h')\big) \arrow[r, mapsto] & \Phi(t+h_1, h') \\ (t, h) \arrow[r, mapsto] \arrow[u, mapsto] & (t+h_1, h') \arrow[u, mapsto] \end{tikzcd}$$

The definition of the bottom arrow Ψ is $(t, h_1, h')) \mapsto (t+h_1, h')$, so that the diagram commutes as indicated on the right. This shows that in terms of the chart h the flow of ξ does act as claimed by the theorem, and this in turn implies the stated normal form of ξ itself. $\square$

So far in the section we have treated what in the language of ordinary differential equations would be called *autonomous* equations. By contrast, a *non-autonomous* differential equation on the manifold X is defined by a time-dependent vector field: such a field assigns to each point $x \in X$ a tangent vector $\xi(t, x) \in T_xX$ which also depends on a time parameter t that varies in an open interval $J \subset \mathbb{R}$. More precisely ξ is a smooth section of the pull-back bundle

$$\begin{tikzcd} \mathrm{pr}_X^* TX \arrow[r] \arrow[d, shift left] & TX \arrow[d] \\ J\times X \arrow[r, "\mathrm{pr}_X"] \arrow[u, shift left, "\xi"] & X \end{tikzcd}$$

induced from the tangent bundle. A solution of the corresponding differential equation is a curve φ defined on a subinterval $J' \subset J$ such that $\dot\varphi(t) = \xi\big(t, \varphi(t)\big)$ holds for all $t \in J'$.

While such a time-dependent vector field ξ does not give rise to a flow on X it is at once accommodated in our set-up if the time parameter is given the status of an independent variable: in terms of the identification $TJ = J\times\mathbb{R}$ the assignment

$$J\times X \ni (t, x) \longmapsto \tilde\xi(t, x) := (t, x; 1) \oplus \xi(t, x) \in \mathrm{pr}_J^* TJ \oplus \mathrm{pr}_X^* TX = T(J\times X)$$

defines a vector field on $J\times X$, and a curve $\varphi\colon J' \to X$ is a solution curve of ξ with initial point $\varphi(a) = x$ if and only if the curve $\tilde\varphi\colon J' \to J\times X$ with

$$\tilde\varphi(t) = \big(t, \varphi(t)\big) \in J\times X$$

is the solution curve of $\tilde\xi$ with initial point $\tilde\varphi(a) = (a, x)$.

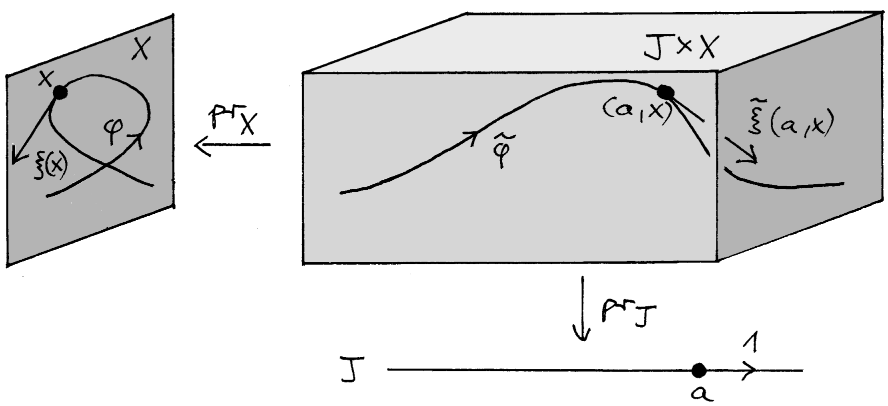

Note that now the curve φ may well fail to be injective without being periodic, as indicated in the figure. By contrast, $\tilde{\varphi}$ always is a smooth embedding of course.

In a similar way every second order differential equation on X may be written as a first order equation on the total space TX of the tangent bundle $TX \xrightarrow{\pi} X$, treating the first order derivative as an independent variable. A second order differential equation on X thus becomes the first order equation given by a vector field $\xi: TX \to T(TX)$ on TX with the special property $T\pi \circ \xi = \mathrm{id}_{TX}$.

A final comment on flows and vector fields on manifolds with non-empty boundary. Clearly most propositions made in this section then no longer can apply literally: for instance there will certainly not exist an integral curve with a positive life span through a point of the boundary where the vector field points outward. As so often the precise way of transfer to the case with boundary depends on the intended application. Nevertheless it is worth singling out one situation of quite general interest where everything we have said remains *literally* true: here the flows $\mathbb{R}\times X \supset D \xrightarrow{\Phi} X$ are required to preserve the boundary in the sense of

$$\Phi\big(D \cap (\mathbb{R} \times \partial X)\big) \subset \partial X,$$

and on the other side only vector fields $\xi: X \to TX$ are admitted that are tangent to the boundary:

$$\xi(x) \in T(\partial X) \quad \text{for all } x \in \partial X.$$

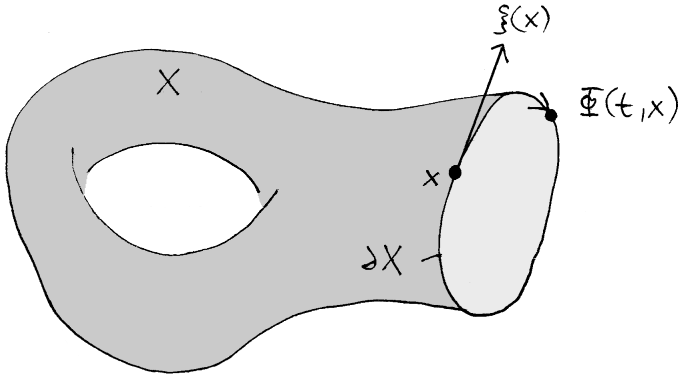

Exercises

8.1 For fixed $n \in \mathbb{N}$ let $f: S^n \to S^n$ be the antipodal map that sends x to $-x$.

- Prove the formula
$$\int_{S^n} f^*\varphi = (-1)^{n+1} \int_{S^n} \varphi$$
for all differential forms $\varphi \in \Omega^n S^n$.

- Assume that there exists a vector field $\xi: S^n \to TS^n$ without zeros. Prove that there exists a (smooth) homotopy from id_{S^n} to f.

- Conclude that there exists a vector field without zeros on S^n if and only if n is odd.

Note This exercise would belong entirely to the previous section if it did not involve the notion of vector field. This is but used in an elementary way, and the relation to flows is not relevant.

8.2 Let ξ be the vector field on the manifold $X = \mathbb{R}$ that sends x to $(x, x^2) \in TX = X \times \mathbb{R}$. Write out the corresponding ordinary differential equation, solve it by explicit calculation, and determine the corresponding maximal flow on Φ, including its precise domain $D \subset \mathbb{R} \times X$.

8.3 Prove that every vector field with compact support is globally integrable.

8.4 Let X be a manifold and $\Phi: D \to X$ a maximal flow on it. Prove that if there exists a $\delta > 0$ with $(-\delta, \delta) \times X \subset D$ then Φ is global.

8.5 Let ξ be a vector field on $\mathbb{R}^n$ such that in the representation

$$\xi(x) = \bigl(x, v(x)\bigr) \in \mathbb{R}^n \times \mathbb{R}^n = T\mathbb{R}^n$$

the smooth function $v: \mathbb{R}^n \to \mathbb{R}^n$ is bounded. Prove that v is globally integrable.

8.6 Let $X \subset \mathbb{R}^n$ be a non-empty star-shaped open subset. Prove that X is diffeomorphic to $\mathbb{R}^n$.

Hint A useful ingredient is a proper smooth function $f: X \to [0, \infty)$ — such a function exists by Exercise 4.3.

8.7 Let X be a connected manifold. Prove that for any two points $a, b \in X$ there exist a diffeomorphism $h: X \to X$ and a compact subset $K \subset X$ such that $h(a) = b$ and $f(x) = x$ for all $x \in X \setminus K$.

8.8 Let Φ be a flow on the manifold X with velocity field ξ, and let $S \subset X$ be a closed submanifold. Prove that S is stable under Φ if and only if ξ is everywhere tangent to S. The first of these statements means

$$\Phi(t, x) \in S \quad \text{for all } x \in S \text{ and all } t \text{ such that } \Phi(t, x) \text{ is defined,}$$

and the second that $\xi(x) \in T_x S$ for all $x \in S$.

8.9 Let Φ be a flow on the manifold X with velocity field ξ, and let $S \subset X$ be a closed equidimensional submanifold with boundary: $\dim S = n = \dim X$. Prove that if $\xi(x) \in T_x X = T_x S$ points inward for every $x \in \partial S$ then flow lines that start in S will stay in S:

$$\Phi(t, x) \in S \quad \text{for all } x \in S \text{ and all } t \geq 0 \text{ such that } \Phi(t, x) \text{ is defined.}$$

8.10 Prove the conclusion of the previous exercise under the weaker assumption that $\xi(x)$ does not point outward at any point $x \in \partial S$.

9 Fibre Bundles

One purpose of the section is to illustrate the method of flows and vector fields in an interesting context. As proper mappings will be of particular importance we collect some of the most important aspects of properness.

9.1 Proper Mappings (1) A continuous map $p: X \to Y$ between locally compact Hausdorff spaces is called *proper* if for every compact $L \subset Y$ the inverse image $p^{-1}L \subset X$ is compact: that is the definition[1].

(2) Methods to establish properness of a map include the following simple observations:

- If X is a compact space (and, as always, Y locally compact) then every map $X \to Y$ is proper.
- If K is a compact space then the cartesian projection $\mathrm{pr}: Y \times K \longrightarrow Y$ is a proper mapping.
- If $p: X \to Y$ is proper and if Z is an arbitrary locally compact space then the product mapping

$$X \times Z \xrightarrow{p \times \mathrm{id}} Y \times Z$$

is proper.

- Compositions of proper maps are proper.
- The restriction of a proper mapping $p: X \to Y$ to a closed subspace $F \subset X$ is proper.
- Let $p: X \to Y$ is a proper mapping and $Y' \subset Y$ a locally closed subspace (that is, the intersection of any open with any closed subset of Y). This ensures that Y' and $p^{-1}Y'$ also are locally compact Hausdorff spaces, and the mapping

$$p': p^{-1}Y' \longrightarrow Y'$$

obtained by restricting p is proper.

- If $p: X \to Y$ is any mapping between locally compact spaces, and if Y admits a locally finite closed cover $(F_\lambda)_{\lambda \in \Lambda}$ such that for each $\lambda \in \Lambda$ the restriction of p

$$p_\lambda: p^{-1}F_\lambda \longrightarrow F_\lambda$$

is proper, then p is proper: given a compact $L \subset Y$, write $p^{-1}L = \bigcup_{\lambda \in \Lambda} p^{-1}(L \cap F_\lambda)$.

(3) The last observation would establish properness of p as a local property with respect to Y if the locally finite cover could be replaced by an arbitrary open one. This holds true by definition if Y is a so-called paracompact space: a Hausdorff space Y is *paracompact* if for every open cover $(Y_\lambda)_{\lambda \in \Lambda}$ of Y there is a locally finite closed cover $(F_\mu)_{\mu \in \mathrm{M}}$ which refines it in the sense that each F_μ is contained in Y_λ for some $\lambda \in \Lambda$.

It can be shown that every compact Hausdorff space is paracompact. On the other hand the class of paracompact spaces is much wider and for instance includes all manifolds (with or without boundary): given the open cover $(Y_\lambda)_{\lambda \in \Lambda}$ choose a subordinate partition of unity $(\rho_j)_{j=0}^{\infty}$ and, with $\mathrm{M} = \mathbb{N}$, put $F_j = \operatorname{supp} \rho_j$ for each $j \in \mathbb{N}$.

[1] The definition for a continuous map $p: X \to Y$ between general topological spaces requires that for every topological space Z the map $p \times \mathrm{id}: X \times Z \to Y \times Z$ sends closed subsets of $X \times Z$ to closed subsets of $Y \times Z$. While, for instance, algebraic geometers must have recourse to this definition when working with Zariski topologies, in the context of locally compact Hausdorff spaces it is equivalent to the simpler version we use.

The figure suggests how some of these rules may be applied to a typical situation of local analytic geometry, where a variety X defined by analytic equations is exhibited as a so-called branched cover of an open subset Y of a Euclidean space.

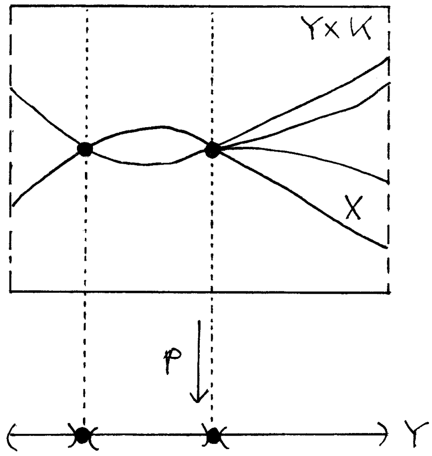

The notion of fibre bundle is more general and weaker than that of vector bundle, from which it takes up the geometric but not the algebraic aspect.

9.2 Definition Let F and X be manifolds. A (*smooth*) *fibre bundle* over X with fibre F is a smooth mapping called the bundle *projection*

$$E \xrightarrow{\pi} X$$

from another manifold E to X which satisfies the axiom of *local triviality*: for every point $x \in X$ there exist an open neighbourhood $U \subset X$ of x and a diffeomorphism h such that the diagram

$$\begin{tikzcd} \pi^{-1}U \arrow[rr, "h", "\simeq"'] \arrow[dr, "\pi"'] & & U \times F \arrow[dl, "\mathrm{pr}"] \\ & U & \end{tikzcd}$$

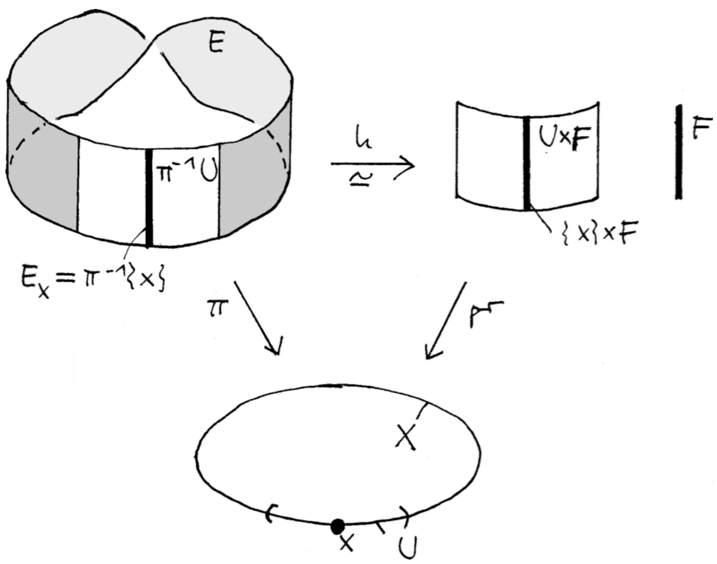

commutes. The spaces X and E are called the *base* and the *total space* of the bundle. Fibre bundles are alternatively called *locally trivial fibrations*.

Much further terminology concerning fibre bundles is identical with that for vector bundles, and includes the notions of product and trivial bundles, bundle mappings and isomorphisms, induced bundles, and sections.

Notes Since $\{x\}\times F \subset U\times F$ is a submanifold and the diffeomorphism h^{-1} sends $\{x\}\times F$ onto $E_x = \pi^{-1}\{x\}$, every fibre of π is a submanifold of E diffeomorphic to F: in this abstract sense F is "the" fibre of π. In case of a connected base it is possible to rewrite the definition without explicit reference to F, which then is replaced by E_x with an arbitrary fixed choice of $x \in U$. For the axiom of local triviality then implies that the partition of X according to diffeomorphism type of the fibre is a partition into open sets and therefore the trivial one, so that all fibres are diffeomorphic and F may be selected as any one of them.

It is clear from the definition that the bundle projection π always is a submersion. Of course we must have $\dim E = \dim X + \dim F$ whenever E is not empty.

The definition would allow either X or F to have a boundary; in both cases E will have a boundary, and in the former π may be restricted to another fibre bundle $\partial E \to \partial X$ with fibre F while in the latter $\pi|\partial E\colon \partial E \to X$ is a fibre bundle with fibre ∂F.

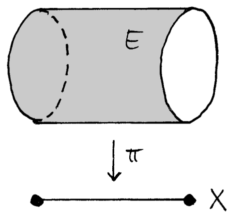

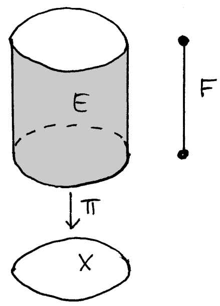

9.3 Lemma The bundle projection of a fibre bundle $E \xrightarrow{\pi} X$ with fibre F over a non-empty base space X is proper if and only if F is compact.

Proof Let π be proper. Choose any $x \in X$, then $\pi^{-1}\{x\}$ is compact and homeomorphic to F, so that F is compact.

Conversely assume that F is compact. Let $(X_\lambda)_{\lambda\in\Lambda}$ be an open cover of X such that π is trivial over each X_λ, and let $(F_\mu)_{\mu\in\mathrm{M}}$ be a locally finite closed refinement. Then each restriction $\pi^{-1}F_\mu \to F_\mu$ is proper, and therefore π is proper by the last point of 9.1(2). $\square$

9.4 Examples (1) Of course every rank d vector bundle $E \xrightarrow{\pi} X$ reduces to a fibre bundle with fibre $\mathbb{R}^d$ when the linear structure of the fibres is forgotten. Its projection π is not a proper mapping, with two trivial exceptions, $X = \emptyset$ and $\operatorname{rk} E = 0$.

(2) If the rank d vector bundle $E \xrightarrow{\pi} X$ carries a Riemannian metric $\rho\colon X \to \operatorname{Sym} E$ then the disk and sphere bundles

$$DE := \big\{v \in E \,\big|\, \rho_{\pi(v)}(v,v) \le 1\big\} \xrightarrow{\pi|DE} X \quad\text{and}\quad SE := \big\{v \in E \,\big|\, \rho_{\pi(v)}(v,v) = 1\big\} \xrightarrow{\pi|SE} X$$

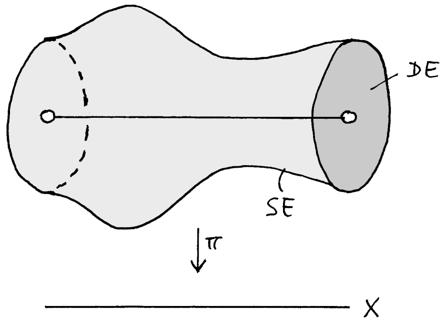

are fibre bundles with fibre D^d and S^{d-1} respectively. The projections of both bundles are proper.

(3) Fibre bundles with a discrete fibre F are known as *covering projections*[2]. Their projections are proper if and only if F is finite. Well-known explicit examples with base space $S^1 = \{z \in \mathbb{C} \mid |z| = 1\}$ include the exponential map

$$\mathbb{R} \ni t \longmapsto e^{2\pi i t} \in S^1 \quad \text{with infinite fibre } \mathbb{Z}$$

and, for non-zero $n \in \mathbb{Z}$, the n-th power mapping

$$S^1 \ni z \longmapsto z^n \in S^1 \quad \text{with finite fibre } \mathbb{Z}/|n|\,.$$

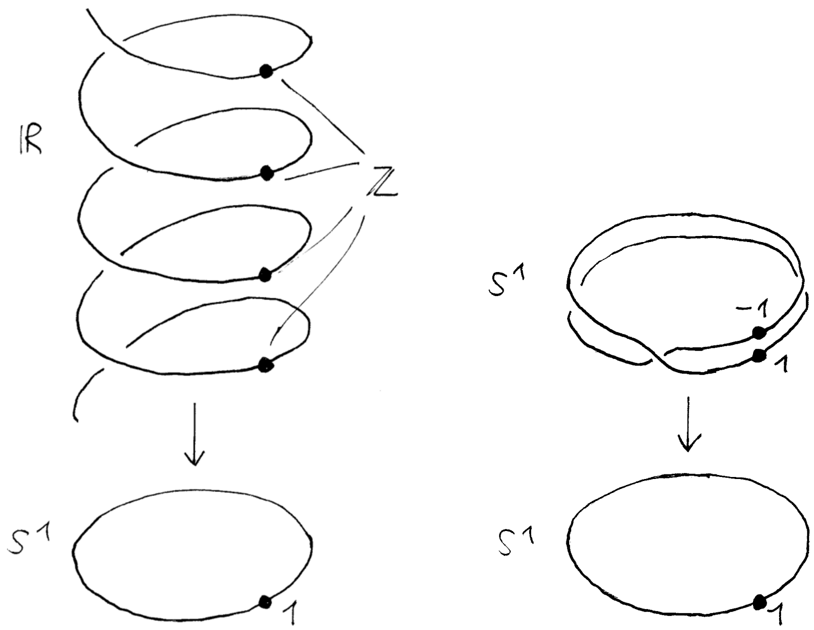

The two-to-one quotient map $S^n \ni x \mapsto [x] \in \mathbb{R}P^n$ provides a further example.

When a concrete smooth map $E \xrightarrow{\pi} X$ is given it may be quite cumbersome to verify the definition in order to recognize π as the projection of a fibre bundle. In certain cases this task is greatly facilitated by the following theorem, which is central to this section. Indeed when only proper maps π are considered it is a criterion that decides whether such a π is a fibre bundle.

9.5 Ehresmann's Fibration Theorem[3] Let X be a connected, and E an arbitary manifold. If the smooth mapping $\pi: E \to X$ is a proper submersion then $E \xrightarrow{\pi} X$ is a fibre bundle.

Proof The conclusion is local in X, and we may therefore assume that $X = \mathbb{R}^p$. To move around in $\mathbb{R}^p$ we will use p particularly simple flows: denoting the standard base vectors by $e_1, \ldots, e_p$ we put

$$\Psi_i: \mathbb{R} \times \mathbb{R}^p \longrightarrow \mathbb{R}^p; \quad (\tau, y) \mapsto y + \tau \cdot e_i \quad \text{for } i = 1, \ldots, p.$$

These are in fact global flows, and given an arbitrary point $t = (t_1, \ldots, t_p) \in \mathbb{R}^p$ the image of $0 \in \mathbb{R}^p$ under the composition

$$(\Psi_1)_{t_1} \circ (\Psi_2)_{t_2} \circ \cdots \circ (\Psi_p)_{t_p}: \mathbb{R}^p \longrightarrow \mathbb{R}^p$$

[2] More precisely they are *smooth* covering projections. Virtually all bundle notions are easily adapted to a purely topological context, replacing manifolds by mere topological spaces, and smoothness by mere continuity.

[3] The French mathematician Charles Ehresmann was a pioneer of differential topology and geometry. He used his fibration theorem from the 1940s onward as one of his standard tools.

is just t.

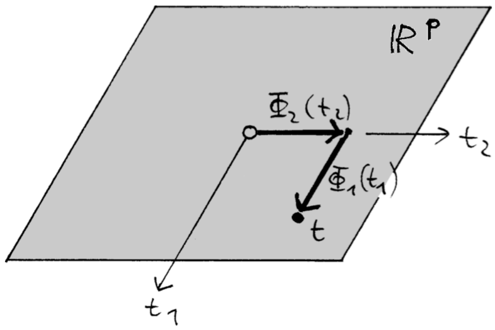

In order to find the desired trivialisation of $E \xrightarrow{\pi} \mathbb{R}^p$ we shall try and "lift" each of these flows to E, that is construct equally global flows

$$\Phi_i : \mathbb{R} \times E \longrightarrow E \quad \text{for } i = 1, \dots, p$$

such that all diagrams

$$\begin{tikzcd}
\mathbb{R} \times E \arrow[r, "\Phi_i"] \arrow[d, "\mathrm{id} \times \pi"'] & E \arrow[d, "\pi"] \\
\mathbb{R} \times \mathbb{R}^p \arrow[r, "\Psi_i"] & \mathbb{R}^p
\end{tikzcd}$$

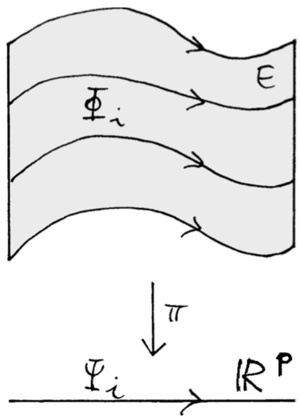

commute. Then, putting $E_0 = \pi^{-1}\{0\}$ as usual, the mapping

$$\begin{array}{ccc}
\mathbb{R}^p \times E_0 & \xrightarrow{g} & E \\
(t, x) & \longmapsto & \Big((\Phi_1)_{t_1} \circ \cdots \circ (\Phi_p)_{t_p} \Big)(x)
\end{array}$$

is a diffeomorphism — it is inverted by sending the point $y \in E$ with $\pi(y) = s$ to the pair

$$\Big(s, \Big((\Phi_p)_{-s_p} \circ \cdots \circ (\Phi_1)_{-s_1} \Big)(y) \Big) \in \mathbb{R}^p \times E_0 \,.$$

It is clear from the definition that g will send the fibre $\mathrm{pr}^{-1}\{t\} = \{t\} \times E_0$ onto $\pi^{-1}\{t\} = E_t$ for every $t \in \mathbb{R}^p$: thus the diagram

$$\begin{tikzcd}
E \arrow[rr, "g^{-1}", "\simeq"'] \arrow[dr, "\pi"'] & & \mathbb{R}^p \times E_0 \arrow[dl, "\mathrm{pr}"] \\
& \mathbb{R}^p &
\end{tikzcd}$$

commutes, and g^{-1} is a (global) trivialisation of the bundle.

It is quite unfeasible to construct the flows Φ_i directly, and at this point the dictionary translating flows to vector fields becomes crucial. We fix one index i and put $\Phi := \Phi_i$ and $\Psi := \Psi_i$; note that the velocity field of Ψ has the constant value $e_i \in \mathbb{R}^p$. If we had found Φ then its velocity field $\xi\colon E \to TE$ would satisfy

$$T\pi \circ \xi = T\pi \circ \dot{\Phi}_t\Big|_{t=0} = \frac{d}{dt}(\pi \circ \Phi_t)\Big|_{t=0} = \frac{d}{dt}(\Psi_t \circ \pi)\Big|_{t=0} = e_i$$

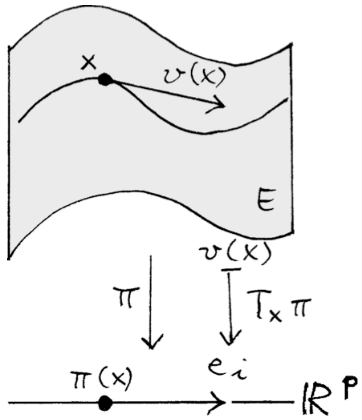

and in this sense be a lifting of e_i to E. Assume conversely that we have constructed a maximal flow Φ on E whose velocity field ξ lifts e_i. Then if $\varphi\colon (\alpha_x, \omega_x) \to E$ is the flow line through x the composition

$$\pi\circ\varphi\colon (\alpha_x, \omega_x) \to \mathbb{R}^p$$

is a flow line of e_i through $\pi(x)$:

$$\frac{d}{dt}(\pi\circ\varphi) = T\pi \circ \xi \circ \varphi = e_i \quad \text{and} \quad (\pi\circ\varphi)(0) = \pi(x)\,.$$

By the uniqueness of flow lines we must have

$$(\pi\circ\varphi)(t) = \pi(x) + t \cdot e_i \quad \text{for all } t \in (\alpha_x, \omega_x).$$

We claim that $\alpha_x = -\infty$ and $\omega_x = \infty$, so that in fact $\pi \circ \varphi = \Psi_{\pi(x)}$. Assuming $\omega_x < \infty$ the curve $\varphi|[0, \omega_x)$ is proper by the maximality criterion 8.12. Since π is proper the curve

$$[0, \omega_x) \ni t \xmapsto{\pi\circ\varphi} \pi(x) + t \cdot e_i \in \mathbb{R}^p$$

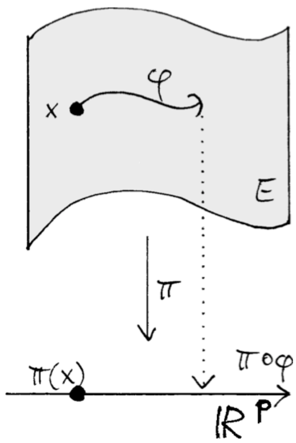

also is proper — which clearly is untrue, so that we have arrived at a contradiction. Therefore $\omega_x = \infty$, and symmetrically $\alpha_x = -\infty$. This completes the proof that the flow $\Phi = \Phi_i$ is a global lifting of Ψ_i.

In view of Theorem 8.8 on the equivalence of maximal flows and vector fields our task has by now been reduced to that of constructing a vector field ξ on E which lifts the constant vector field e_i on $\mathbb{R}^p$. We first solve this problem locally near an arbitrary point $a \in E$: We know that $\mathrm{rk}_a\,\pi = p$ and find, by the constant rank theorem 2.5, a centred chart (U, h) at a such that $\pi \circ h^{-1}$ becomes

$$h(U) \ni (s_1, \ldots, s_n, t_1 \ldots, t_p) \longmapsto \pi(a) + (t_1, \ldots, t_p) \in \mathbb{R}^p.$$

Therefore the assignment $U \ni x \mapsto (T_x h)^{-1} \cdot e_{n+i} \in T_x E$ defines a vector field ξ on U with $T\pi \circ \xi = e_i$ as required. Applying this construction to every point a we obtain an open cover $(U_a)_{a \in E}$ of E, and for every $a \in E$ a lifting vector field $\xi_a : U_a \to TU_a = TE|U_a$. Let $(\rho_j)_{j=0}^{\infty}$ be a partition of unity subordinate to this cover, and choose for each $j \in \mathbb{N}$ a point $a(j) \in E$ with $\operatorname{supp} \rho_j \subset U_{a(j)}$. Then the sum

$$\xi := \sum_{j=0} \rho_j \cdot \xi_{a(j)}$$

is a well-defined global vector field on E. It lifts e_i since at every point $x \in E$ the lifting condition $T_x \pi \cdot \xi(x) = e_i$ is a convex condition on the vector $\xi(x)$.

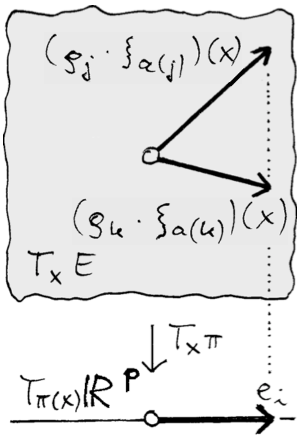

This completes the construction of ξ and thereby the proof of the fibration theorem. □

In a special case we have in fact proved a stronger result:

9.6 Corollary Every proper submersion $E \to \mathbb{R}^p$ is a trivial fibre bundle. □

The method that has allowed to control the life spans of the flow lines, on the other hand, is worth being recorded for general use.

9.7 Proposition Let $X \xrightarrow{f} Y$ be a proper smooth mapping between manifolds, and let $\xi : X \to TX$ and $\eta : Y \to TY$ be vector fields such that ξ lifts η:

$$Tf \circ \xi = \eta \circ f$$

Let $\mathbb{R} \times Y \supset D \xrightarrow{\Psi} Y$ be the maximal flow of η. Then the maximal flow of ξ has the domain

$$(\mathrm{id} \times f)^{-1} D \subset \mathbb{R} \times X.$$

Proof Let $x \in X$ be arbitrary, and let $\varphi : (\alpha_x, \omega_x) \to X$ be the maximal integral curve through x. Then $f \circ \varphi$ is an integral curve of η, and we must show that it is maximal. By symmetry it is sufficient to prove that the restriction $f \circ \varphi|[0, \omega_x)$ has maximal life span as a flow line of η. This is certainly true if $\omega_x = \infty$. If not, then we know from the criterion 8.12 that $\varphi|[0, \omega_x)$ is proper, and since f is proper the composition $f\circ\varphi|[0,\omega_x)$ also is proper. This composition is an integral curve of η, and again by the criterion we conclude its maximality. □

Since fibrations with closed fibres — compact manifolds without boundary — are relatively special let us briefly treat a version of the theorem with boundary.

9.8 Theorem Let X be a connected manifold, and E a manifold with boundary. Let $\pi: E\to X$ be a proper smooth mapping such that not only π itself but also the restriction $\pi|\partial E$ is submersive. Then $E\xrightarrow{\pi} X$ is a fibre bundle.

Proof The proof from 9.5 requires but little change. In view of the remarks at the end of the previous section we only must make sure that the vector field $\xi: E\to TE$ is everywhere tangent to the boundary ∂E. As this is a linear property concerning the values of ξ on ∂E it can be treated locally. At interior points $a\in E$ there is no new condition at all, and the local vector fields ξ_a are chosen as before.

Thus consider a boundary point $a\in\partial E$. First using any chart at a we may assume that E is an equidimensional submanifold of $\mathbb{R}^{n+p}$ and that π is defined on an open neighbourhood $\tilde E$ of a in $\mathbb{R}^{n+p}$.

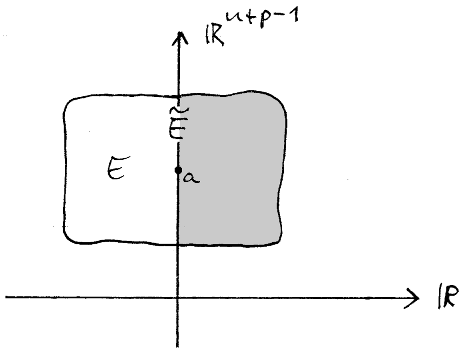

There the constant rank theorem applies, and we find a centred chart (W,h) at a such that $\pi\circ h^{-1}$ takes the form

$$h(W)\ni(s_1,\dots,s_n,t_1,\dots,t_p)\longmapsto\pi(a)+(t_1,\dots,t_p)\in\mathbb{R}^p.$$

While this chart has no reason to send $W\cap\partial E$ to $\{0\}\times\mathbb{R}^{n+p-1}$ from the assumption that $\pi|\partial E$ is submersive we do know that the projection $(s_1,\dots,s_n,t_1,\dots,t_p)\mapsto(t_1,\dots,t_p)$ maps the tangent space $T_0\big(h(W\cap\partial E)\big)$ onto $\mathbb{R}^p$:

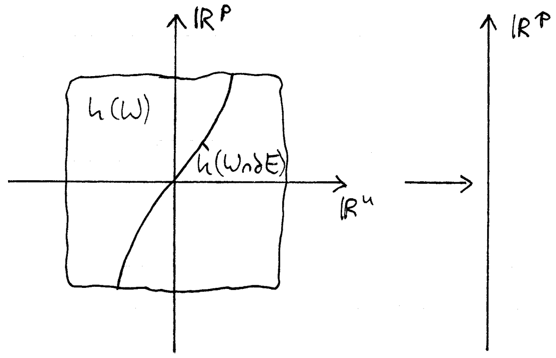

In particular this $(n{+}p{-}1)$-dimensional space cannot contain $\mathbb{R}^n\times\{0\}$, and after a rotation of $\mathbb{R}^n$ — which does not affect $\pi\circ h^{-1}$ — we may assume that $T_0\big(h(W\cap\partial E)\big)$ does not contain the first coordinate axis $\mathbb{R}\times\{0\}\times\{0\}\subset\mathbb{R}\times\mathbb{R}^{n-1}\times\mathbb{R}^p$. By the implicit function theorem we now may shrink W such that $h(W)$ becomes an open box $U\times V\subset\mathbb{R}\times\mathbb{R}^{n+p-1}$, and $h(W\cap\partial E)\subset U\times V$ the graph of a smooth function $\varphi: V\to U$. Correcting h by composition with the local diffeomorphism

$$(h_1,\dots,h_{n+p})=(h_1,h')\longmapsto\big(h_1-\varphi(h'),h'\big)$$

and possibly flipping h_1 we obtain a boundary chart for E.

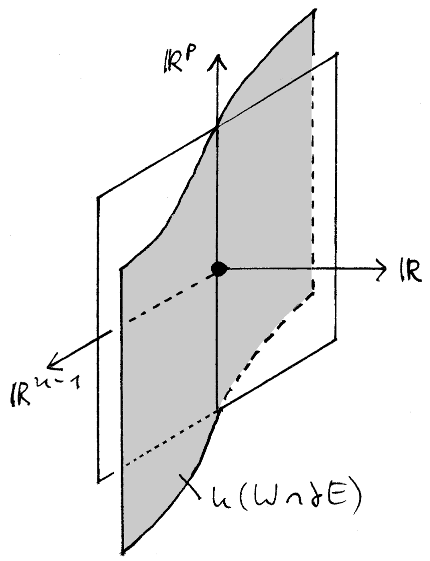

None of these corrections has an effect on the map $\pi \circ h^{-1}$, and we can now lift the constant vector field e_i from $\mathbb{R}^p$ to e_{n+i} on $h(W)$, which corresponds to a vector field on W which is tangential to $W \cap \partial E$ as required. □

9.9 Examples (1) For fixed radii $0 < r < R$ consider the embedded torus

$$T = \big\{\big((R+r\cos\theta)\cos\varphi, (R+r\cos\theta)\sin\varphi, r\sin\theta\big) \in \mathbb{R}^3 \bigm| \varphi, \theta \in \mathbb{R}\big\} \subset \mathbb{R}^3$$

and let $f\colon T \to \mathbb{R}$ be the restriction of the coordinate function $x\colon \mathbb{R}^3 \to \mathbb{R}$. Since T is compact f is proper; on the other hand the differential of f vanishes exactly at the four points $(\pm R \pm r, 0, 0) \in T$, so that the numbers $\pm R \pm r \in \mathbb{R}$ are the critical values of f. Thus for any choice of the interval X among

$$(-\infty, -R-r),\ (-R-r, -R+r),\ (-R+r, R-r),\ (R-r, R+r)\ , \text{and } (R+r, \infty)$$

the restricted mapping $f^{-1}X \longrightarrow X$ is a proper submersion and therefore, by Ehresmann's theorem, a locally trivial fibration — even a trivial one by Corollary 9.6.

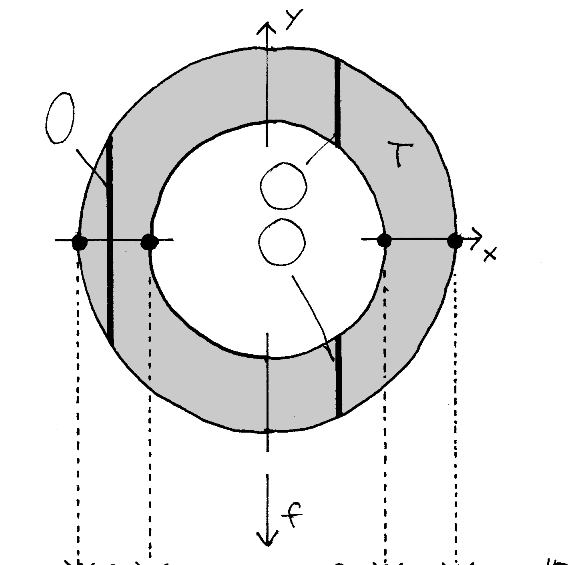

The fibrations over the two unbounded intervals with empty fibre are of no interest. The fibre over the middle interval is S^1+S^1, and that over the lateral ones, a single copy of S^1. Thus a dense part of T is built of cylinders over the fibres of f. One might think of scanning T starting at low values of f where the fibres are empty, then increasing the value of f to observe a change in the type of fibre every time a critical value is crossed — on the other hand there is no such change between two consecutive critical values. This idea of analyzing a manifold via the fibres of a cleverly chosen scalar function has led to a surprisingly powerful method — named after the American mathematician Marston Morse — to investigate the manifold's topology.

(2) One particular application of the boundary version 9.8 has been much used in order to study so-called isolated singular points of complex polynomials. For the sake of simplicity we consider a complex polynomial

$$f(z) = \sum_{|j|=d} a_j z^j \in \mathbb{C}[z] = \mathbb{C}[z_1, \ldots, z_n]$$

which is homogeneous of positive degree $d \in \mathbb{N}$; the summation is over n-dimensional indices $j \in \mathbb{N}^n$ of constant total degree $|j| := \sum_{k=1}^n j_k$, and z^j is shorthand for the monomial $z_1^{j_1} \cdots z_n^{j_n}$. We interpret f as a holomorphic function

$$f\colon \mathbb{C}^n \longrightarrow \mathbb{C}$$

and assume that $Df(z) \neq 0$ for all $z \in \mathbb{C}^n \setminus \{0\}$: thus by the regular value theorem all fibres $f^{-1}\{t\}$ with $t \neq 0$ are submanifolds of $\mathbb{C}^n$ of codimension 2, while for the zero fibre $f^{-1}\{0\}$ this becomes true when the origin of $\mathbb{C}^n$ is removed.

We wish to show that for sufficiently small real $\delta > 0$ the mapping

$$E := D^{2n} \cap f^{-1} S_\delta(0) \xrightarrow{\pi} S_\delta(0)$$

obtained by restricting f is a smooth fibre bundle. First note that the function f has complex rank one — thus real rank 2 — at every point $z \in S^{2n-1} \cap f^{-1}\{0\}$. Since f is homogeneous the differential $Df(z)$ annihilates the radial vector $z \in \mathbb{C}^n = T_z D^{2n}$, and this further implies that the auxiliary mapping

$$\mathbb{C}^n \supset D_2(0) \ni z \overset{F}{\longmapsto} \big(|z|^2, f(z)\big) \in \mathbb{R} \times \mathbb{C}$$

has rank 3 along $S^{2n-1} \cap f^{-1}\{0\}$: in other words the pair $(1,0) \in \mathbb{R}\times\mathbb{C}$ is a regular value of F.

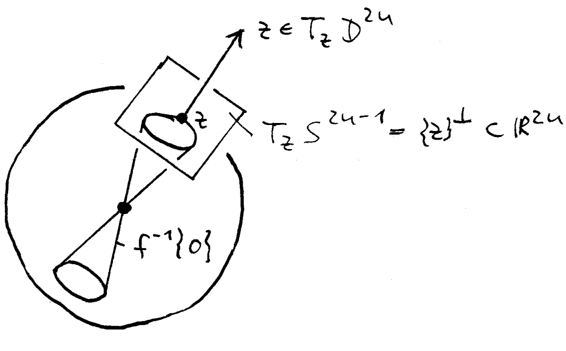

Using the fact that the set of critical values of $F|D_2(0)$ must be compact we choose $\delta > 0$ sufficiently small so that all points of $\{1\}\times D_{\delta^2} \subset \mathbb{R}\times\mathbb{C}$ are regular values of $F|U_2(0)$. A fortiori the composition

$$U_2(0) \ni z \longmapsto \big(|z|^2, |f(z)|^2\big) \in \mathbb{R} \times \mathbb{R}$$

obtained from F has $(1, \delta^2) \in \mathbb{R}\times\mathbb{R}$ as a regular value, and applying the regular value theorem first in the original version 1.13 and then in the boundary version 7.3(3) we obtain that

$$U_2(0) \cap f^{-1} S_\delta(0) = \big\{ z \in U_2(0) \,\big|\, |f(z)|^2 = \delta^2 \big\}$$

and

$$E = D^{2n} \cap f^{-1}S_\delta(0) = \{z \in U_2(0) \cap f^{-1}S_\delta(0) \mid |z|^2 \le 1\}$$

are $(2n-1)$-dimensional submanifolds of $\mathbb{C}^n$, the latter with boundary $\partial E = S^{2n-1} \cap f^{-1}S_\delta(0)$.

Thus $E \xrightarrow{\pi} S_\delta(0)$ is a smooth mapping between manifolds. At any given point of E the rank of π is 1 since the original function f has rank 2. At a point of ∂E the rank of the auxiliary mapping F is 3, and since $\pi|\partial E$ essentially is the restriction of F to the pre-image of $\{1\}\times S_\delta(0) \subset \mathbb{R}\times\mathbb{C}$ the rank of $\pi|\partial E$ must be at least $3-2$, and thus also equal to 1.

We now know that

$$E \xrightarrow{\pi} S_\delta(0) \quad \text{and} \quad \partial E \xrightarrow{\pi} S_\delta(0)$$

are submersions, and by Ehresmann's theorem conclude that $E \xrightarrow{\pi} S_\delta(0)$ is a fibre bundle over the circle as claimed.

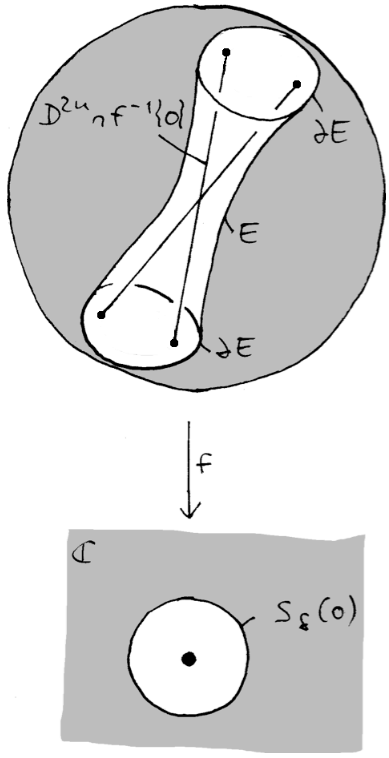

This locally trivial fibration was first constructed in the 1960s by John Milnor, and is named after him. Recall that in the previous decade Milnor had been the first to describe exotic spheres, and his interest was catched by the discovery that some quite simple choices of the polynomial f lead to Milnor fibres whose boundaries are such exotic spheres. The idea has since been partially reversed and put in a much more general context. Already Milnor had shown that given a homogeneous[4] polynomial f with an isolated singular point, the fibration does not essentially depend on the choice of the number δ and thus is an invariant of f. Its fibre E_δ — the *Milnor fibre* — turns out to reveal a lot of information about the singular point, and is closely related to other invariants of purely algebraic nature. A first instance of such relations may be expressed in terms of the de Rham cohomology of this $(2n-2)$-dimensional manifold with boundary: all de Rham spaces but $H^0E_\delta = \mathbb{R}$ and $H^{n-1}E_\delta = \mathbb{R}^{\mu(f)}$ vanish, and the dimension of the latter, the so-called *Milnor number* $\mu(f)$ has an algebraic interpretation as the dimension of the *Jacobian algebra*

$$\mathbb{C}[z_1,\dots,z_n]\Big/\Big(\frac{\partial f}{\partial z_1},\dots,\frac{\partial f}{\partial z_n}\Big)$$

[4] More generally inhomogeneous polynomials and even holomorphic functions defined near the origin of $\mathbb{C}^n$ may be considered if D^{2n} is replaced by a ball of sufficiently small radius, and if the Jacobian algebra referred to below is suitably localized.

as a complex vector space.

(3) We let $H := \{\tau \in \mathbb{C} \mid \operatorname{Im} \tau > 0\}$ denote the complex upper half plane (the name τ of the complex coordinate is traditional here), and consider the product bundle

$$H \times \mathbb{C} \xrightarrow{\mathrm{pr}} H.$$

The group $\mathbb{Z}^2$ acts on the total space $H \times \mathbb{C}$ via

$$\mathbb{Z}^2 \times (H \times \mathbb{C}) \longrightarrow H \times \mathbb{C}; \quad (k,l) \cdot (\tau, z) = (\tau,\, k + l\,\tau + z)$$

preserving the fibres of the bundle, and this action gives rise to the orbit space

$$E := (H \times \mathbb{C})/\mathbb{Z}^2$$

obtained by identifying (τ, w) with (τ, z) if and only if the difference $w - z$ belongs to the lattice $\Lambda_\tau := \mathbb{Z} + \mathbb{Z}\tau \subset \mathbb{C}$. By the discreteness of this lattice sufficiently small restrictions of the identical mapping $(\tau, z)\colon H \times \mathbb{C} \to H \times \mathbb{C}$ yield an atlas on E whose transition maps locally are affine linear. Therefore the quotient space E is a smooth 4-manifold, and, even more, the projection induces a smooth mapping $E \ni [\tau, z] \xmapsto{\pi} \tau \in H$ which is a proper submersion. Note that the fibre over $\tau \in H$ is the torus $E_\tau = \mathbb{C}/\Lambda_\tau$.

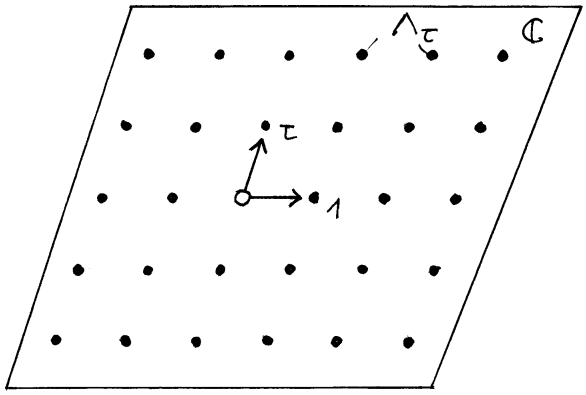

We now might apply Ehresmann's theorem and conclude that $E \xrightarrow{\pi} H$ is a trivial fibre bundle — but this would be overkill since a global trivialisation can be easily written down explicitly:

$$\begin{tikzcd} E \arrow[dr, "\pi"'] & & H \times \mathbb{R}^2/\mathbb{Z}^2 \arrow[ll, "{[\tau, x+y\tau] \,\leftarrow\!\shortmid\, (\tau,x,y)}"', "\simeq"] \arrow[dl, "\mathrm{pr}"] \\ & H & \end{tikzcd}$$

In any case it is true that for all choices of $\tau \in H$ the fibres E_τ are diffeomorphic to each other, for instance to the standard torus $E_i = \mathbb{C}/(\mathbb{Z}+\mathbb{Z}i) = \mathbb{R}^2/\mathbb{Z}^2$.

The true interest of this example lies in the fact that it carries a richer than just a smooth structure. The charts that define the differentiable structure of the manifold E not only take complex rather than real values but also give rise to holomorphic transition functions. Therefore E in fact carries the structure of a complex (holomorphic) manifold — a notion whose local theory is quite analogous to ours; simply the notion of differentiable must be read as holomorphic throughout. Of course this is also true for the upper half plane H, and our bundle projection $\pi\colon E \to H$ is a proper holomorphic submersion. The fibres of π thus also are complex manifolds; they are rightfully called 1-dimensional *complex tori*[5]. If Ehresmann's theorem held for complex manifolds then in particular all these tori

[5] They also turn out to be the elliptic curves of complex algebraic geometry. — Of course complex dimension 1 means real — that is topological — dimension 2, a fact reflected in the traditional name of *Riemann surface* for a 1-dimensional complex manifold.

would be *biholomorphically equivalent* to each other — this is the accepted term for isomorphy in the category of complex manifolds and holomorphic maps.

In reality the picture is quite different. The E_τ form a continuous family of generically biholomorphically *inequivalent* tori; in order to make a simple precise statement let us say that there are uncountably many $\tau \in H$ such that the corresponding complex tori E_τ are pairwise biholomorphically inequivalent.

To see this, we define an invariant that distinguishes inequivalent fibres. Fix any $\tau \in H$. On a complex manifold like E_τ the notion of holomorphic differential form makes sense, and here we have the particular holomorphic 1-form $dz\colon TE_\tau \to \mathbb{C}$ at hand: the global coordinate z of $\mathbb{C}$ provides local coordinates on E_τ whose transition functions locally are translations, and thus have no effect on the differential dz. Note that the complex rank of the bundle TE_τ is just one, and since dz clearly has no zeros, every further holomorphic 1-form must be a scalar multiple $\lambda(z)\cdot dz$ of it. But then $\lambda\colon E_\tau \to \mathbb{C}$ is a holomorphic function on the compact torus E_τ, and by the well-known maximum principle for holomorphic functions must be a constant.

Let $f\colon[0,1] \to E_\tau$ be a C^∞ loop (closed curve) in E_τ. If it happens to be of the form

$$[0,1] \xrightarrow{\tilde f} \mathbb{C} \xrightarrow{\pi} E_\tau$$

for a smooth curve $\tilde f\colon[0,1] \to \mathbb{C}$ then the difference $\tilde f(1)-\tilde f(0)$ must belong to the lattice Λ_τ. On the other hand this difference is the value of the path integral

$$\int_{[0,1]} f^*dz = \int_{[0,1]} d(\tilde f^*\pi^*z) = \int_0^1 d\tilde f = \tilde f\Big|_0^1 = \tilde f(1) - \tilde f(0)$$

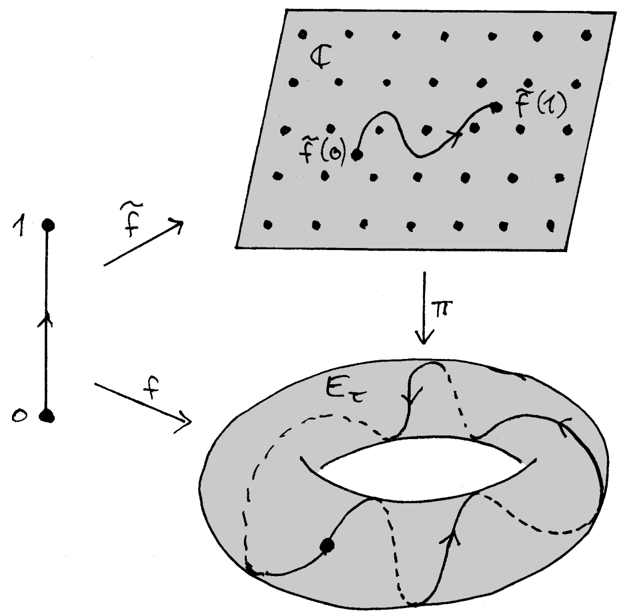

along f. It is not difficult to see that π is a covering projection and that every loop f lifts to a curve $\tilde f$, and we conclude that $\Lambda_\tau \subset \mathbb{C}$ is the set of all path integrals of dz along arbitrary loops f in E_τ.

Let now $\tau' \in H$ be a second number and assume that there exists a biholomorphic equivalence between E_τ and $E_{\tau'}$.

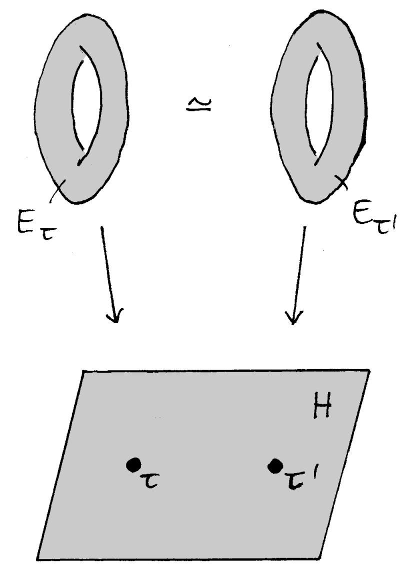

It must take the 1-form dz from E_τ to some holomorphic 1-form on $E_{\tau'}$, that is to some constant multiple $\lambda \cdot dz$. Therefore the sets of path integrals are related by

$$\Lambda_\tau = \lambda \cdot \Lambda_{\tau'} \subset \mathbb{C}.$$

Explicitly this means that there exists an invertible matrix $g = \begin{pmatrix} a & b \\ c & d \end{pmatrix} \in GL(2,\mathbb{Z})$ with

$$g \cdot \begin{pmatrix} \tau \\ 1 \end{pmatrix} = \lambda \cdot \begin{pmatrix} \tau' \\ 1 \end{pmatrix}$$

or, explicitly, $\tau' = \dfrac{a\tau + b}{c\tau + d}$. Since $GL(2,\mathbb{Z})$ is countable this proves our claim: there must exist uncountably many $\tau \in H$ such that the corresponding complex tori E_τ are pairwise biholomorphically inequivalent.

The picture is easily completed. The formula relating τ and τ' defines an action of the special linear group $SL(2,\mathbb{Z})$ on H, and this action has a so-called *fundamental domain*

$$F := \big\{\tau \in H \,\big|\, |\mathrm{Re}\,\tau| \le \frac{1}{2} \text{ and } |\tau| \ge 1\big\}.$$

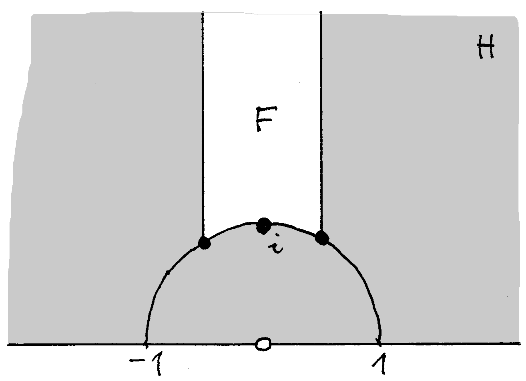

Roughly speaking this statement means that every orbit in H intersects F in exactly one point of the interior F°, or just finitely many points of the boundary $F \setminus F^\circ$. In particular any two distinct $\tau \in F^\circ$ yield holomorphically inequivalent tori E_τ. In complex analysis as well as algebraic geometry parameters like τ are called *moduli*, and parameter spaces like H *moduli spaces* — they describe continuous variations of the holomorphic or algebraic structure. They constitute a fundamental contrast with smooth manifold theory, since Ehresmann's theorem implies that from the smooth point of view such variations, at least in the compact case, must be trivial.

If we attempt to adapt the proof of Ehresmann's theorem to the holomorphic setting, something of course must break down, and it is easily identified: the identity theorem for holomorphic function precludes the existence of holomorphic table mountains and thereby of partitions of unity.

Exercises

9.1 Let $p\colon X \to Y$ be a proper mapping between locally compact Hausdorff spaces. Prove that p sends closed subsets of X to closed subsets of Y.

9.2 Let X be a smooth compact 2-manifold and $f\colon X \to \mathbb{R}$ a function such that the differential of f vanishes at exactly two points of X. Prove that X is homeomorphic to S^2.

9.3 Putting $a = -3/\sqrt[3]{4}$ and $X = (-\infty, a) \cup (a, \infty)$ consider the projection

$$E := \big\{ \big(s, [x:y:z]\big) \in X \times \mathbb{R}P^2 \,\big|\, y^2 z = x^3 + sxz + z^3 \big\} \xrightarrow{\pi} X$$

that sends $\big(s, [x:y:z]\big)$ to s. Prove that E is a smooth manifold and that π restricts to a fibre bundle over each of the two connected components of X.

Note $[x:y:z]$ is the customary notation for the equivalence class of $(x, y, z) \in S^2$ in $\mathbb{R}P^2$ — it emphasises the fact that what counts in a projective space is exactly the *ratios* between the numbers x, y, and z.

9.4 Let X and E be manifolds, and $E \xrightarrow{\pi} \mathbb{R} \times X$ be a fibre bundle with compact fibre. Prove that the bundle is isomorphic to the bundle

$$\mathbb{R} \times E' \xrightarrow{\mathrm{id} \times \pi'} \mathbb{R} \times X$$

where $E' \xrightarrow{\pi'} X$ is obtained from π by restriction over $\{0\} \times X$:

$$\begin{tikzcd}
E' \arrow[r, equal] \arrow[d, "\pi'"] & \pi^{-1}\big(\{0\} \times X\big) \arrow[r, hook] \arrow[d] & E \arrow[d, "\pi"] \\
X \arrow[r, "\simeq"] & \{0\} \times X \arrow[r, hook] & \mathbb{R} \times X
\end{tikzcd}$$

9.5 Let $E \xrightarrow{\pi} Y$ be a fibre bundle with compact fibre, and let $f_0\colon X \to Y$ and $f_1\colon X \to Y$ be two smoothly homotopic maps. Prove that the induced bundles $f_0^* E$ and $f_1^* E$ on X are isomorphic.

9.6 Let $f \in \mathbb{C}[z_1, \ldots, z_n]$ be the quadratic polynomial $f(z) = \sum_{j=1}^{n} z_j^2$. Show that any choice of δ with $0 < \delta < 1$ is good in the sense that the restriction

$$E := D^{2n} \cap f^{-1} S_\delta(0) \longrightarrow S_\delta(0)$$

becomes a fibre bundle (the Milnor fibration). Prove that the Milnor fibre E_δ is diffeomorphic to the total space of the unit disk bundle

$$\big\{(u,v)\in TS^{n-1}\,\big|\,|v|\le 1\big\}\subset S^{n-1}\times\mathbb{R}^n$$

of the tangent bundle of S^{n-1}.

9.7 Verify that the formula

$$\begin{pmatrix} a & b\\ c & d\end{pmatrix}\cdot\tau=\frac{a\tau+b}{c\tau+d}$$

defines an action of the *modular group* $\Gamma := PSL(2,\mathbb{Z}) = SL(2,\mathbb{Z})/\{\pm1\}$ on the upper half plane H. Prove that the set

$$F=\Big\{\tau\in H\,\Big|\,|\mathsf{Re}\,\tau|\le\frac{1}{2}\text{ and }|\tau|\ge 1\Big\}$$

is a fundamental domain of the action. This means precisely

- that every orbit intersects F,
- that this intersection consists either of just one point of the interior F°, or a finite subset of the boundary $F\smallsetminus F^\circ$, and
- that the isotropy group is trivial in the first case, and finite else.

What is the quotient H/Γ as a topological space?

10 Lie Derivatives

The equivalence between flows and vector fields can be used to introduce a new kind of differentiation on an arbitrary manifold.

10.1 Definition Let $h: X \to Y$ be a smooth map between manifolds. The process of pulling back differential forms

$$\Omega^k Y \ni \chi \longmapsto h^*\chi \in \Omega^k X \quad \text{with } (h^*\chi)_x(v_1, \dots, v_k) = \chi\big(T_x h \cdot v_1, \dots, T_x h \cdot v_k\big)$$

has no analogue for vector fields in general, but it does in the case that h is a local diffeomorphism: a vector field $\eta: Y \to TY$ is sent to

$$h^*\eta := Th^{-1} \circ \eta \circ h, \text{ explicitly } (h^*\eta)(x) = (T_x h)^{-1} \cdot \eta\big(h(x)\big).$$

Remarks As in previous sections we often indicate evaluation of a linear map on a vector by a dot rather than use brackets. — Sections of any tensor product bundle $E \to X$ with (possibly repeated) factors TX, $(TX)\check{}$, and their symmetric and alternating powers are called *tensors* or *tensor fields* on X. The construction of h^* extends from the vector fields and differential forms that we are mainly considering, to such tensors $\alpha: X \to E$ in general:

$$(h^*\alpha)_x = l_h \cdot \alpha_{h(x)}$$

where the linear mapping $l_h: E_{h(x)} \to E_x$ is induced from one copy of $(T_x h)^{-1}$ for each tensor factor TX, and one of $(T_x h)\check{}$ for each factor $TX\check{}$ of E.

10.2 Definition We denote by $\operatorname{Vect} X$ the real vector space of (smooth) vector fields on the manifold X. For every $k \in \mathbb{N}$ there is a partial evaluation map

$$\operatorname{Vect} X \times \Omega^k X \ni (\xi, \varphi) \longmapsto \xi \lrcorner \varphi \in \Omega^{k-1} X$$

defined by

$$(\xi \lrcorner \varphi)_x(v_2, \dots, v_k) = \varphi\big(\xi(x), v_2, \dots, v_k\big) \quad \text{(to be interpreted as zero for } k = 0\text{)}.$$

Note that this bilinear operation is purely algebraic like the wedge product of differential forms: in order to know $(\xi \lrcorner \varphi)_x \in \operatorname{Alt}^{k-1} X$ for fixed $x \in X$ it is sufficient to know $\xi(x) \in T_x X$ and $\varphi_x \in \operatorname{Alt}^k T_x X$. For $k = 1$ the operation reduces to the evaluation of linear forms on vectors, and is usually written as such:

$$\langle \xi, \varphi \rangle := \xi \lrcorner \varphi \in \Omega^0 X \quad \text{for } \varphi \in \Omega^1 X.$$

10.3 Definition Let X be a manifold, $\xi \in \operatorname{Vect} X$, and α another vector field or a differential form — or, quite generally, any tensor on X. Then the *Lie derivative* of α with respect to ξ is defined as

$$\mathcal{L}_\xi \alpha = \frac{d}{dt}\, \Phi_t^* \alpha \,\bigg|_{t=0}$$

where Φ is the maximal flow on X that integrates ξ.

Explanation While Φ_t need not make global sense, in order to determine the derivative on the right hand side near a point $o \in X$ we only need $\Phi_t(x) = \Phi(t,x)$ for t close to $0 \in \mathbb{R}$ and x in some neighbourhood of o: the flow will supply these values since its domain is open and includes $\{0\}\times X$. — If the tensor α is, say, a section of $E \to X$ then for fixed $o \in X$ the assignment $t \mapsto (\Phi_t^*\alpha)(o)$ is a curve in the vector space E_o. Hence also $(\mathcal{L}_\xi\alpha)(x) \in E_o$, and therefore $\mathcal{L}_\xi\alpha$ is a tensor of the same type as α.

10.4 Lemma Let $\Phi\colon D \to X$ denote the maximal flow generated by ξ. The formula defining $\mathcal{L}_\xi\alpha$ generalises to the identities

$$\frac{d}{dt}\,\Phi_t^*\alpha\,(x) = \mathcal{L}_\xi\Phi_t^*\alpha\,(x) = \Phi_t^*\mathcal{L}_\xi\alpha\,(x)\,,$$

valid for all times t and points $x \in X$ with $(t,x) \in D$.

Proof For any fixed $(t,x) \in D$ we have, using a time parameter τ close to 0,

$$\frac{d}{dt}\,(\Phi_t^*\alpha)(x) = \frac{d}{d\tau}\,(\Phi_{\tau+t}^*\alpha)(x)\,\Big|_{\tau=0} = \frac{d}{d\tau}\,\big(\Phi_\tau^*(\Phi_t^*\alpha)\big)(x)\,\Big|_{\tau=0} = (\mathcal{L}_\xi\Phi_t^*\alpha)(x)$$

and

$$\frac{d}{dt}\,(\Phi_t^*\alpha)(x) = \frac{d}{d\tau}\,(\Phi_{t+\tau}^*\alpha)(x)\,\Big|_{\tau=0} = \frac{d}{d\tau}\,\big(\Phi_t^*(\Phi_\tau^*\alpha)\big)(x)\,\Big|_{\tau=0} = \Phi_t^*\,\frac{d}{d\tau}\,(\Phi_\tau^*\alpha)(x)\,\Big|_{\tau=0} = (\Phi_t^*\mathcal{L}_\xi\alpha)(x)\,.$$

In the second case we have used that the action of Φ_t^* is by a linear mapping $E_{\Phi(t,x)} \to E_x$. □

Intuitively speaking $\mathcal{L}_\xi\alpha$ is the infinitesimal change that α undergoes when ξ, more precisely the flow generated by ξ, is applied.

10.5 Example Let $f \in \Omega^0X$ be a differential form of degree 0, in other words a smooth function on X. Then

$$\mathcal{L}_\xi f = \frac{d}{dt}(\Phi_t^*f)\,\Big|_{t=0} = \frac{d}{dt}(f\circ\Phi_t)\,\Big|_{t=0} = df\circ\xi = \langle\xi, df\rangle$$

is a new function on X, its value at x is the directional derivative of f with respect to the tangent vector $\xi(x)$. In particular the value of $\mathcal{L}_\xi f$ at a point x depends only on $\xi(x)$, not on other values of ξ; this is not a property of Lie derivatives in general. — Often the Lie action of the vector field ξ on the smooth function f is written even shorter as ξf rather than $\mathcal{L}_\xi f$.

10.6 Notation Let (U,h) be a chart for the n-dimensional manifold X and fix some $i \in \{1,\ldots,n\}$. The vector field $\xi\colon U \to TU = TX|U$ whose h-component has the constant value

$$\xi(x)_h = T_xh\cdot\xi(x) = e_i \quad \text{for all } x \in U,$$

acts on functions on U by

$$\xi f = d(f\circ h^{-1})\circ Th\circ\xi = \big\langle e_i, d(f\circ h^{-1})\big\rangle = \frac{\partial f}{\partial h_i}\,.$$

This property, which uniquely determines ξ, explains the classical and seemingly weird notation $\dfrac{\partial}{\partial h_i} \in \operatorname{Vect}U$ for this particular vector field. The fields

$$\frac{\partial}{\partial h_1},\ \ldots,\ \frac{\partial}{\partial h_n} \in \operatorname{Vect}U$$

clearly are dual to the differential 1-forms $dh_1, \ldots, dh_n$, so that $\Big\langle\dfrac{\partial}{\partial h_i}, dh_j\Big\rangle = \delta_{ij}$. The notation suggests the correct transition formula

$$\frac{\partial}{\partial h_i} = \sum_{j=1}^{n}\frac{\partial k_j}{\partial h_i}\frac{\partial}{\partial k_j}$$

when the chart is changed to k. Unfortunately it also suggests that $\dfrac{\partial}{\partial h_i}$ might just depend on the function h_i and not the other components of h — this is not true, and unlike df which makes sense for any scalar function f on X the expression $\dfrac{\partial}{\partial f}$ has no meaning for general f.

10.7 Definition Let X be a manifold. A linear mapping $D\colon \Omega^0 X \to \Omega^0 X$ is called a *derivation* if it satisfies the product rule

$$D(f\cdot g) = Df\cdot g + f\cdot Dg \quad \text{for all functions } f, g \in \Omega^0 X\,.$$

Similarly a linear function $D\colon \Omega^0 X \to \mathbb{R}$ is called a *derivation at the point* $o \in X$ if

$$D(f\cdot g) = Df\cdot g(o) + f(o)\cdot Dg \quad \text{for all functions } f, g \in \Omega^0 X\,.$$

10.8 Proposition Let X be a manifold. Every derivation D is a local operator:

$$\operatorname{supp} Df \subset \operatorname{supp} f \quad \text{for all } f \in \Omega^0 X\,,$$

and if D is a derivation at $o \in X$ then $Df = 0$ unless $o \in \operatorname{supp} f$.

Proof If D is a derivation and $o \in X$ then the assignment $\Omega^0 X \ni f \mapsto Df(o) \in \mathbb{R}$ is a derivation at o; thus it is sufficient to prove the second statement, concerning a derivation D at the point o. To this end we assume $o \notin \operatorname{supp} f$. We choose a table mountain $\tau\colon X \to [0,1]$ at o with $\tau\cdot f = 0$ and obtain

$$Df = D\big((1-\tau)\cdot f\big) = D(1-\tau)\cdot f(o) + (1-\tau)(o)\cdot Df = 0\,. \qquad \square$$

It is clear that for every vector field $\xi \in \operatorname{Vect} X$ the map $\mathcal{L}_\xi\colon f \mapsto \xi f$ is a derivation, and that for a given point $o \in X$ and tangent vector v the assignment $f \mapsto (df)_o \cdot v$ defines a derivation at o. It turns out that in fact all derivations at o are of this type:

10.9 Theorem Let X be a manifold and $o \in X$. The map that assigns to each vector $v \in T_o X$ the derivation $f \mapsto (df)_o \cdot v$ is an isomorphism between $T_o X$ and the vector space of derivations at the point o.

Proof In view of 10.8 and the possibility of extending smooth functions defined near o by a table mountain the question is local near o, and we do not restrict the generality if we assume that X is covered by the domain of a single chart h at o. Then, putting $n = \dim X$, the derivation assigned to

$$v = \sum_{i=0}^{n} v_i \frac{\partial}{\partial h_i} \in T_o X$$

sends the j-th component function $h_j \in \Omega^0 X$ to $v_j \in \mathbb{R}$, so that the coefficients of v can be recovered from the derivation as $v_i = (dh_i)_o \cdot v$: this shows that the assignment is injective. To prove it is also surjective it suffices to show that the vector space of derivations at o has dimension at most n.

We may work in a centred chart at o and thus assume that $X \subset \mathbb{R}^n$ is a convex open neighbourhood of $o = 0$. By Taylor's formula every smooth function $f \in \Omega^0 X$ may be written as

$$f(x) = f(0) + \sum_{i=0}^{n} x_i \cdot f_i(x) \quad \text{with} \quad f_i(x) = \int_0^1 \frac{\partial f}{\partial x_i}(tx)\, dt\,.$$

If now D is a derivation at 0 then firstly

$$D1 = D(1\cdot 1) = D1\cdot 1 + 1\cdot D1 = 2\cdot D1$$

so that $D1 = 0$, and D annihilates constant functions. Secondly

$$D(x_i \cdot f_i) = Dx_i \cdot f_i(0) + x_i(0) \cdot Df_i = Dx_i \cdot f_i(0)\,,$$

so that

$$Df = \sum_{i=0}^{n} Dx_i \cdot f_i(0) = \sum_{i=0}^{n} Dx_i \cdot \frac{\partial f}{\partial x_i}(0) \quad \text{for all } f \in \Omega^0 X\,.$$

This expresses D as a linear combination of the n partial derivation operators at 0, and thereby completes the proof of the lemma. □

10.10 Lemma The Lie derivative with respect to a fixed vector field ξ is a derivation with respect to both the wedge product of differential forms, and the partial evaluation map between vector fields and differential forms:

$$\mathcal{L}_\xi(\varphi\wedge\chi) = \mathcal{L}_\xi\varphi \wedge \chi + \varphi \wedge \mathcal{L}_\xi\chi \quad \text{and} \quad \mathcal{L}_\xi(\eta \lrcorner \chi) = \mathcal{L}_\xi\eta \lrcorner \chi + \eta \lrcorner \mathcal{L}_\xi\chi$$

— in fact for any bilinear algebraic tensor operation the analogue would be true. Furthermore the Lie derivative commutes with the Cartan differential:

$$\mathcal{L}_\xi(d\varphi) = d(\mathcal{L}_\xi\varphi).$$

Proof Let $\Phi\colon D \to X$ be the maximal flow generated by ξ. Then

$$\frac{d}{dt}\,\Phi_t^*(\varphi\wedge\chi) = \frac{d}{dt}\,(\Phi_t^*\varphi\wedge\Phi_t^*\chi) = \frac{d}{dt}\Phi_t^*\varphi \,\wedge\, \Phi_t^*\chi + \Phi_t^*\varphi \,\wedge\, \frac{d}{dt}\Phi_t^*\chi\,,$$

and

$$\frac{d}{dt}\,\Phi_t^*(\eta \lrcorner \chi) = \frac{d}{dt}\,(\Phi_t^*\eta \lrcorner \Phi_t^*\chi) = \frac{d}{dt}\Phi_t^*\eta \,\lrcorner\, \Phi_t^*\chi + \Phi_t^*\eta \,\lrcorner\, \frac{d}{dt}\,\Phi_t^*\chi\,.$$

Similarly we have

$$\frac{d}{dt}\,\Phi_t^*(d\varphi) = \frac{d}{dt}\,d(\Phi_t^*\varphi) = d\Big(\frac{d}{dt}\Phi_t^*\varphi\Big)$$

since time and space derivatives commute with each other. Evaluation at $t = 0$ now yields the desired identities. □

10.11 Example Let (U, h) be a chart for the n-dimensional manifold X. By Lemma 10.10 the vector field $\xi = \sum_j v_j \frac{\partial}{\partial h_j}$ acts on a general 1-form by

$$\begin{aligned}
\mathcal{L}_\xi \sum_{i=1}^{n} \varphi_i\, dh_i &= \sum_{i=1}^{n} (\mathcal{L}_\xi\varphi_i\, dh_i + \varphi_i\, d\mathcal{L}_\xi h_i) \\
&= \sum_{i=1}^{n} \left(\sum_{j=1}^{n} v_j\, \frac{\partial\varphi_i}{\partial h_j}\, dh_i + \varphi_i\, dv_i \right) \\
&= \sum_{i,j=1}^{n} v_j\, \frac{\partial\varphi_i}{\partial h_j}\, dh_i + \sum_{i,j=1}^{n} \varphi_i\, \frac{\partial v_i}{\partial h_j}\, dh_j \\
&= \sum_{i,j=1}^{n} \left(v_j\, \frac{\partial\varphi_i}{\partial h_j} + \frac{\partial v_j}{\partial h_i}\, \varphi_j \right) dh_i\,.
\end{aligned}$$

Note that the expression involves first derivatives of the coefficient functions v_i, so that the Lie derivative no longer is a pointwise operator with respect to ξ.

The Lie derivative in 10.11 is the coordinate version of the $\deg\varphi = 1$ case of the following more general conceptual formula.

10.12 Theorem Let X be a manifold and ξ a vector field on X. The Lie derivative of a differential form may be expressed as

$$\mathcal{L}_\xi\varphi = \xi \lrcorner\, d\varphi + d(\xi \lrcorner\, \varphi).$$

Proof We fix $\xi \in \operatorname{Vect} X$. From 10.10 we know that the left hand side of the formula behaves like an (ungraded) derivation with respect to φ and the wedge product. If we write $D\psi = \xi \lrcorner\, \psi$ then the right hand side of the formula becomes $(Dd+dD)\,\varphi$. We know the Cartan differential d also is a derivation, though a graded one[2], and we now show that the latter is equally true for the operator D. Since the value $(D\psi)_x$ only depends on $\xi(x)$ and ψ_x this is a mere question of linear algebra involving tangent spaces at a fixed point x. It is therefore sufficient to verify the equation $D(\psi\wedge\chi) = D\psi\wedge\chi\pm\psi\wedge D\chi$ for the case of $\xi = \dfrac{\partial}{\partial h_1}$ with respect to some chart h, and the differential forms

$$\psi = dh_{i_1}\wedge\cdots\wedge dh_{i_k} \quad\text{and}\quad \chi = dh_{j_1}\wedge\cdots\wedge dh_{j_l}$$

with $i_1 < \cdots < i_k$ and $j_1 < \cdots < j_l$. Inserting $\xi = \dfrac{\partial}{\partial h_1}$ into the wedge product we obtain

$$\begin{aligned} &D\big((dh_{i_1}\wedge\cdots\wedge dh_{i_k})\wedge(dh_{j_1}\wedge\cdots\wedge dh_{j_l})\big) \\ &= \begin{cases} 0 & \text{if } i_1 > 1 \text{ and } j_1 > 1\,, \\ (dh_{i_2}\wedge\cdots\wedge dh_{i_k})\wedge(dh_{j_1}\wedge\cdots\wedge dh_{j_l}) & \text{if } i_1 = 1 \text{ and } j_1 > 1\,, \\ (-1)^k(dh_{i_1}\wedge\cdots\wedge dh_{i_k})\wedge(dh_{j_2}\wedge\cdots\wedge dh_{j_l}) & \text{if } i_1 > 1 \text{ and } j_1 = 1\,, \\ 0 & \text{if } i_1 = 1 \text{ and } j_1 = 1\,, \end{cases} \end{aligned}$$

which coincides with

$$\begin{aligned} &D(dh_{i_1}\wedge\cdots\wedge dh_{i_k})\wedge(dh_{j_1}\wedge\cdots\wedge dh_{j_l}) + (-1)^k(dh_{i_1}\wedge\cdots\wedge dh_{i_k})\wedge D(dh_{j_1}\wedge\cdots\wedge dh_{j_l}) \\ &= \begin{cases} 0+0 & \text{if } i_1 > 1 \text{ and } j_1 > 1\,, \\ (dh_{i_2}\wedge\cdots\wedge dh_{i_k})\wedge(dh_{j_1}\wedge\cdots\wedge dh_{j_l}) + 0 & \text{if } i_1 = 1 \text{ and } j_1 > 1\,, \\ 0 + (-1)^k(dh_{i_1}\wedge\cdots\wedge dh_{i_k})\wedge(dh_{j_2}\wedge\cdots\wedge dh_{j_l}) & \text{if } i_1 > 1 \text{ and } j_1 = 1\,, \\ (dh_{i_2}\wedge\cdots\wedge dh_{i_k})\wedge(dh_{j_1}\wedge\cdots\wedge dh_{j_l}) & \\ +(-1)^k(dh_{i_1}\wedge\cdots\wedge dh_{i_k})\wedge(dh_{j_2}\wedge\cdots\wedge dh_{j_l}) & \text{if } i_1 = 1 \text{ and } j_1 = 1 \end{cases} \end{aligned}$$

(to make the two terms of the last expression cancel the factor $dh_{i_1} = dh_1$ is shifted past $k-1$ other degree 1 factors to become dh_{j_1}).

We now know that d and D are graded derivations with respect to the wedge product, and it is readily shown that $Dd+dD$ then is an ungraded one. Thus both sides of $\mathcal{L}_\xi\varphi = \xi \lrcorner\, d\varphi + d(\xi \lrcorner\, \varphi)$ are local operators and derivations with respect to φ, and this reduces the proof to the two cases of a 0-form $\varphi = f$ and of the special 1-forms $\varphi = dh_k$. Again writing $\xi = \sum_j v_j \dfrac{\partial}{\partial h_j}$, the first reads

$$\mathcal{L}_\xi f = \langle \xi, df\rangle + d\,0$$

and was observed in Example 10.5 while the second,

$$\mathcal{L}_\xi dh_k = \sum_{i=1}^n \frac{\partial v_k}{\partial h_i}\, dh_i = \xi \lrcorner\, 0 + dv_k$$

is included in Example 10.11. □

[2] We are here working with derivations that do not fit into the formal framework of Definition 10.7 because they act on the algebra of differential forms rather than that of scalar functions.

Note that the formula of Theorem 10.12 reduces to $\mathcal{L}_\xi\varphi = d(\xi \lrcorner \varphi)$ whenever φ is a closed differential form.

10.13 Application Let X be an n-manifold and $K \subset X$ a compact oriented n-dimensional submanifold with boundary. Let further ξ be a vector field on X and $\mathbb{R}\times X \supset D \xrightarrow{\Phi} X$ be the maximal flow generated by ξ: since K is compact D contains a set of the form $(-\delta,\delta)\times K$ for some positive δ, so that $\Phi_t\colon K \to X$ is a well-defined smooth embedding for all times t close to zero. Under these conditions the identity

$$\frac{d}{dt}\int_K \Phi_t^*\omega\,\bigg|_{t=0} = \int_{\partial K}(\xi \lrcorner \omega)$$

holds for every differential form $\omega \in \Omega^n X$.

Interpretation If ω is the density of some physical quantity, say charge, then the formula equates the variation of total charge in K with the flow of charge across the boundary of K.

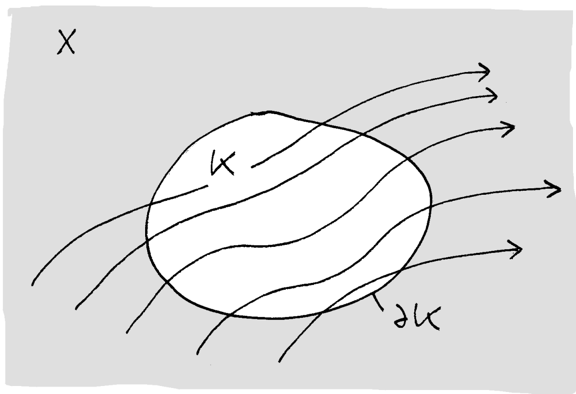

Proof Of course $\omega \in \Omega^n X$ is closed, so that Theorem 10.12 and Stokes' formula yield

$$\int_K \mathcal{L}_\xi\omega = \int_K d(\xi \lrcorner \omega) = \int_{\partial K}(\xi \lrcorner \omega)\,.$$

We substitute the definition of the Lie derivative; since differentiation with respect to t commutes with integration we obtain

$$\frac{d}{dt}\int_K \Phi_t^*\omega\,\bigg|_{t=0} = \int_K \frac{d}{dt}\Phi_t^*\omega\,\bigg|_{t=0} = \int_K \mathcal{L}_\xi\omega = \int_{\partial K}(\xi \lrcorner \omega)\,. \qquad \square$$

10.14 Definition Let X be an n-manifold, possibly with boundary, and assume that a volume form $\omega \in \Omega^n X$ is given. Then multiplication by ω defines canonical isomorphisms

$$\Omega^0 X \ni f \longmapsto f\cdot\omega \in \Omega^n X \quad\text{and}\quad \operatorname{Vect} X \ni \xi \longmapsto \xi \lrcorner \omega \in \Omega^{n-1}X.$$

The new mapping that makes the diagram

$$\begin{tikzcd} \operatorname{Vect} X \arrow[r, "\operatorname{div}"] \arrow[d, "\simeq"'] & \Omega^0 X \arrow[d, "\simeq"] \\ \Omega^{n-1}X \arrow[r, "d"] & \Omega^n X \end{tikzcd}$$

commutative and assigns a scalar function to each vector field is called the *divergence*.

The intuitive meaning of the divergence is easily read off from the formula

$$\mathcal{L}_\xi\omega = d(\xi \lrcorner \omega) = \operatorname{div}\xi \cdot \omega \; :$$

at every point $x \in X$ the number $\operatorname{div}\xi(x)$ is the infinitesimal change that the local volume suffers under the flow generated by ξ. In particular ξ expands or compresses volumes near x if $\operatorname{div}\xi(x)$ is positive or negative, respectively — note in passing that this sign has nothing to do with orientations since the divergence remains the same when ω is replaced by $-\omega$ and thus the orientation of X is flipped.

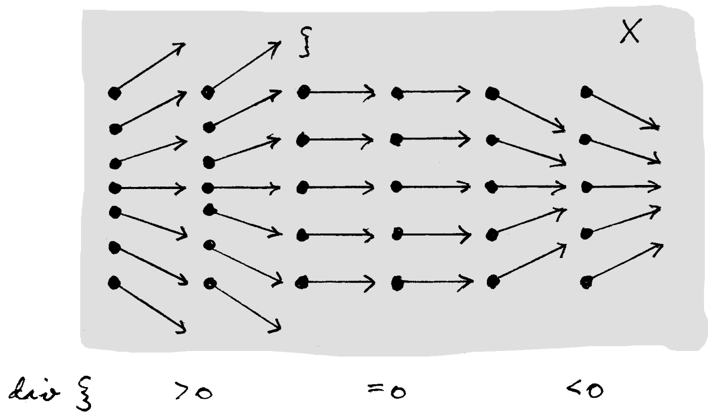

We turn from the Lie derivative of differential forms to that of vector fields as another basic instance of tensor fields.

10.15 Example Let X be a manifold and $\xi, \eta \in \operatorname{Vect} X$ two vector fields. In order to understand the Lie derivative of η we calculate that of the evaluation bracket between η and an exact form $df \in \Omega^1 X$:

$$\mathcal{L}_\xi\langle\eta, df\rangle = \langle\mathcal{L}_\xi\eta, df\rangle + \langle\eta, \mathcal{L}_\xi df\rangle$$

By 10.5 the left hand side and the first term of the right are respectively $\mathcal{L}_\xi\mathcal{L}_\eta f$ and $(\mathcal{L}_\xi\eta)f$, while 10.10 and again 10.5 show that the remaining term is $\langle\eta, \mathcal{L}_\xi df\rangle = \big\langle\eta, d(\mathcal{L}_\xi f)\big\rangle = \mathcal{L}_\eta\mathcal{L}_\xi f$. The overall result is

$$(\mathcal{L}_\xi\eta)f = \mathcal{L}_\xi\mathcal{L}_\eta f - \mathcal{L}_\eta\mathcal{L}_\xi f \quad \text{for all smooth } f\colon X \to \mathbb{R}\,,$$

or

$$\mathcal{L}_\xi\eta = \xi \circ \eta - \eta \circ \xi$$

if we write the operator $f \mapsto \mathcal{L}_\xi f = \xi f$ simply as ξ. Since the vector field $\mathcal{L}_\xi\eta \in \operatorname{Vect} X$ is uniquely determined by its action on scalar functions we thus have found a new representation of it as the commutator of the first order differential operators ξ and η — note that the composition $\xi\circ\eta = \mathcal{L}_\xi\mathcal{L}_\eta$ is of second order but that its second order terms must cancel with those of $\mathcal{L}_\eta\mathcal{L}_\xi$. Written out in a chart (U, h) the formula reads

$$\mathcal{L}_\xi\eta = \sum_{j,k=1}^{n} \left(v_k \frac{\partial w_j}{\partial h_k} - w_k \frac{\partial v_j}{\partial h_k} \right) \frac{\partial}{\partial h_j}$$

with $n = \dim X$ and vector fields $\xi = \sum_j v_j \frac{\partial}{\partial h_j}$ and $\eta = \sum_j w_j \frac{\partial}{\partial h_j}$ on U.

10.16 Definition The bilinear mapping

$$\operatorname{Vect} X \times \operatorname{Vect} X \ni (\xi, \eta) \longmapsto [\xi, \eta] := \mathcal{L}_\xi\eta \in \operatorname{Vect} X$$

is called the *Lie bracket*, and it gives the real vector space $\operatorname{Vect} X$ the structure of a *Lie algebra*. Apart from bilinearity the Lie algebra axioms require the now obvious skew-symmetry $[\eta, \xi] = -[\xi, \eta]$ and the so-called *Jacobi identity*

$$\big[\xi, [\eta, \zeta]\big] + \big[\eta, [\zeta, \xi]\big] + \big[\zeta, [\xi, \eta]\big] = 0 \quad \text{for all } \xi, \eta, \zeta \in \operatorname{Vect} X\,;$$

the latter is immediately verified by writing out the action of each term on scalar functions.

Lie algebras form an interesting class of algebraic objets, and an extensive theory is available, though much more complete for algebras which have finite vector space dimension. By contrast the Lie algebra $\operatorname{Vect} X$ has infinite dimension up to the obvious exception of 0-manifolds X.

Exercises

10.1 Let X be an n-manifold with volume form $\omega \in \Omega^n X$, and let Φ be a flow on X with velocity field ξ. Prove for every compact equidimensional submanifold[1] $K \subset X$ the formula

$$\frac{d}{dt} \int_{\Phi_t(K)} \omega \,\bigg|_{t=0} = \int_K \operatorname{div} \xi \cdot \omega$$

which expresses the rate of change of volume as K is carried along by Φ.

10.2 Let X be an n-manifold, $f\colon X \to \mathbb{R}$ a smooth function, and $o \in X$ a critical point of f. Let further $\xi, \eta \in \operatorname{Vect} X$ be two vector fields. Prove that the values of

$$\mathcal{L}_\xi \mathcal{L}_\eta f \quad \text{and} \quad \mathcal{L}_\eta \mathcal{L}_\xi f$$

at the point o coincide and that they only depend on the tangent vectors $\xi(o)$ and $\eta(o)$ rather than on the fields ξ and η. Express the number $\mathcal{L}_\xi \mathcal{L}_\eta f(o)$ in terms of a chart (U, h) at o.

10.3 Let $\eta = \sum_{j=1}^n w_j \frac{\partial}{\partial x_j}$ be a vector field on the open ball U^n. Prove that the Lie bracket

$$\Big[\frac{\partial}{\partial x_1}, \eta\Big] \in \operatorname{Vect} U^n$$

vanishes identically if and only if the values of the coefficient functions $w_j(x_1, \ldots, x_n)$ do not depend on x_1.

10.4 Let Φ and Ψ be flows on the manifold X. Prove that the flows commute in the sense of

$$\Phi\big(s, \Psi(t, x)\big) = \Psi\big(t, \Phi(s, x)\big) \quad \text{for all } x \in X \text{ and all } s, t \in \mathbb{R} \text{ sufficiently close to } 0$$

if and only if their velocity fields ξ and η satisfy $[\xi, \eta] = 0$.

10.5 Let $\xi, \eta \in \operatorname{Vect} X$ be vector fields on a manifold X such that $[\xi, \eta] = 0$. Prove that at every point $o \in X$ where $\xi(o)$ and $\eta(o)$ are linearly independent in $T_o X$ there exists a chart (U, h) for X such that

$$\xi|U = \frac{\partial}{\partial h_1} \quad \text{and} \quad \eta|U = \frac{\partial}{\partial h_2}\,.$$

[1] In fact any compact subset of X would do.

11 Lie Groups

We already have informally met Lie groups in the context of the regular value theorem. The aim of this section is to take a first systematic look at Lie groups — but we will not even get near to the heart of their theory, since the latter by itself would be a subject for a full course.

11.1 Definition A *Lie group* G is, simultaneously, a smooth manifold and a group, such that the group multiplication

$$G \times G \ni (x, y) \longmapsto xy \in G$$

is differentiable[1]. A *homomorphism of Lie groups* by definition is a differentiable group homomorphism.

A homomorphism of Lie groups $\mathbb{R} \to G$ is called a *one parameter group of* G (or *in* G)[2].

11.2 Notation Let G be a Lie group. We will write $G \times G \xrightarrow{\mu} G$ whenever a letter for the group multiplication map is required. For any fixed $g \in G$ we let

$$G \ni y \overset{\lambda_g}{\longmapsto} gy \in G \quad \text{and} \quad G \ni x \overset{\rho_g}{\longmapsto} xg \in G$$

denote the left, respectively right *translation mapping*. All these maps are smooth, and the translation maps are diffeomorphisms as translation by g^{-1} inverts them. While for given elements $g, h \in G$ the translation λ_g need not commute with λ_h, the associative law guarantees that λ_g and ρ_h always commute with each other.

11.3 Proposition Let G be a Lie group. Then the inverting map

$$G \ni x \longmapsto x^{-1} \in G$$

also is smooth and therefore a self-diffeomorphism of G.

Proof Fix a point $g \in G$. The partial differential of $\mu: G \times G \to G$ with respect to the second variable at the point (g, g^{-1}) is just the differential $T_{g^{-1}}\lambda_g$, and therefore invertible. By the implicit function theorem there exist open neighbourhoods $U \subset G$ of g and $V \subset G$ of g^{-1} such that

$$\mu^{-1}\{1\} \cap (U \times V) = \big\{(x, y) \in U \times V \,\big|\, \varphi(x) = y\big\}$$

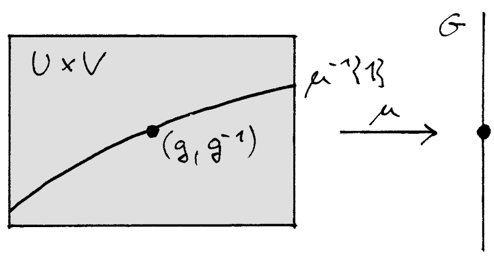

[1] It is not unusual to require the differentiability of the inverting map $x \mapsto x^{-1}$ as well, in order to keep the definition analogous to that of a topological group. Proposition 11.3 below will show that this property is a consequence of the definition as stated.

[2] Of course these are not groups but mappings. Yet in the literature it is not uncommon even to call them one parameter *subgroups* of G

is the graph of a smooth function $\varphi\colon U \to V$. This function necessarily sends x to x^{-1}, and therefore the inverting map is differentiable at g. □

11.4 Definition Let G be a Lie group. The tangent space at the unit element

$$\operatorname{Lie} G := T_1 G$$

is called the *Lie algebra* of G. If $f\colon G \to H$ is a homomorphism of Lie groups we will use $\operatorname{Lie} f$ as an alternative notation for the differential $T_1 f$ in order to emphasize functoriality of the Lie algebra.

Let us briefly recall examples of Lie groups that we have met before. The euclidean spaces $\mathbb{R}^n$ are abelian Lie groups with respect to vector addition, and so are all finite dimensional real or complex vector spaces V. Their Lie algebra $\operatorname{Lie} V$ is canonically identified with V itself.

If such a V is fixed any collection of linearly independent vectors span a lattice $\Lambda \subset V$, and the quotient V/Λ is an abelian Lie group with Lie algebra V. This class of Lie groups includes the explicit examples of $\Lambda = \mathbb{Z}^n \subset \mathbb{R}^n = V$ as well as the complex tori $\mathbb{C}/\Lambda_\tau$ — which in fact carry an even richer structure as *holomorphic* Lie groups.

In the context of the regular value theorem we encountered some of the linear groups, namely the general and special linear groups $GL(n,\mathbb{R})$ and $SL(n,\mathbb{R})$, and the orthogonal group $O(n)$. Their Lie algebras are the space of all $n\times n$-matrices

$$\operatorname{Lie} GL(n,\mathbb{R}) = \mathfrak{gl}(n,\mathbb{R}) = \operatorname{Mat}(n\times n,\mathbb{R}),$$

that of the traceless matrices

$$\operatorname{Lie} SL(n,\mathbb{R}) = \mathfrak{sl}(n,\mathbb{R}) = \big\{x \in \operatorname{Mat}(n\times n,\mathbb{R}) \,\big|\, \operatorname{tr} x = 0\big\},$$

and that of the skew-symmetric matrices

$$\operatorname{Lie} O(n) = \mathfrak{o}(n) = \big\{x \in \operatorname{Mat}(n\times n,\mathbb{R}) \,\big|\, x^{\mathrm{t}} + x = 0\big\}$$

respectively.

11.5 Observation Let G be a Lie group. The tangent bundle of $TG \xrightarrow{\pi} G$ is canonically trivial: the triangle

$$\begin{tikzcd} TG \arrow[dr, "\pi"'] & & G \times \operatorname{Lie} G \arrow[ll, "T_1\lambda_x\cdot\xi \,\leftarrow\!\shortmid\, (x,\xi)"'] \arrow[dl, "\mathrm{pr}"] \\ & G & \end{tikzcd}$$

describes a global trivialisation.

Proof The inverse sends $\eta \in TG$ to $T_{\pi(\eta)}\lambda_{\pi(\eta)^{-1}} \cdot \eta$. □

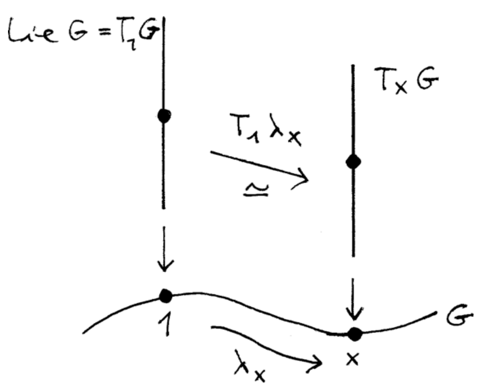

11.6 Definition A vector field $\xi\colon G \to TG$ on the Lie group G is called *left invariant* if $\lambda_x^*\xi = \xi$ for every $x \in G$, or explicitly if for every $x \in G$ the diagram

$$\begin{array}{ccc} TG & \xrightarrow{T\lambda_x} & TG \\ \uparrow{\scriptstyle\xi} & & \uparrow{\scriptstyle\xi} \\ G & \xrightarrow{\lambda_x} & G \end{array}$$

commutes.

11.7 Lemma If ξ is a left invariant vector field on the Lie group G then the corresponding maximal flow $\mathbb{R}\times G \ni D \xrightarrow{\Phi} G$ is global and left invariant too:

$$\Phi(t, gx) = g \cdot \Phi(t, x) \quad \text{for all } t \in \mathbb{R} \text{ and } g, x \in G.$$

In particular the flow line through the unit element is a one parameter group of G.

Proof Fix a point $x \in G$ and let $\Phi_x\colon (\alpha_x, \omega_x) \to G$ be the maximal flow line through x. For any choice of $g \in G$ the curve

$$\varphi_g := \lambda_g \circ \Phi_x$$

starts at $\lambda_g(x) = gx$, satisfies

$$\dot\varphi_g(\tau) = \frac{d}{d\tau}\,(\lambda_g \circ \Phi)(\tau, x) = T_{\Phi(\tau,x)}\lambda_g \cdot \dot\Phi(\tau, x) = T_{\Phi(\tau,x)}\lambda_g \cdot \xi\big(\Phi(\tau, x)\big) = \xi\big((\lambda_g \circ \Phi)(\tau, x)\big) = (\xi \circ \varphi_g)(\tau)$$

and thus is an integral curve of ξ through gx.

We apply this fact in two ways. Firstly, if t is chosen such that $(t, x) \in D$ the particular choice of $g := \Phi(t, x)\, x^{-1}$ shows that

$$(\alpha_x, \omega_x) \subset (\alpha_{\Phi(t,x)}, \omega_{\Phi(t,x)}) = (\alpha_x - t, \omega_x - t),$$

where the last identity is taken from Theorem 8.8. Since both negative and positive choices of t are possible this implies that $\alpha_x = -\infty$ and $\omega_x = \infty$. Therefore Φ is global. — Secondly, we read from the identity of curves $\lambda_g \circ \Phi_x = \varphi_g = \Phi_{gx}$ the left invariance of Φ. It further implies for all $s, t \in \mathbb{R}$ that

$$\Phi(s+t, 1) = \Phi\big(t, \Phi(s, 1) \cdot 1\big) = \Phi(s, 1) \cdot \Phi(t, 1)\,,$$

so that the curve $t \mapsto \Phi(t, 1)$ is a homomorphism. □

11.8 Corollary If ξ and η are two left invariant vector fields on the Lie group G then their Lie bracket $[\xi, \eta] \in \operatorname{Vect} G$ again is left invariant. Therefore the left invariant vector fields form a Lie subalgebra of $\operatorname{Vect} G$.

Proof Let $\mathbb{R}\times G \xrightarrow{\Phi} G$ be the flow generated by ξ, and $g \in G$ an arbitrary element. Then

$$\lambda_g^*[\xi, \eta] = \lambda_g^*\, \frac{d}{dt}\, \Phi_t^*\eta\, \Big|_{t=0} = \frac{d}{dt}\, \lambda_g^*\Phi_t^*\eta\, \Big|_{t=0} = \frac{d}{dt}\, \Phi_t^*\lambda_g^*\eta\, \Big|_{t=0} = \frac{d}{dt}\, \Phi_t^*\eta\, \Big|_{t=0} = [\xi, \eta]\,;$$

note that λ_g^* commutes with differentiation since the factor $T_{gx}\lambda_g^{-1}$ contained in it is a linear mapping for every $x \in G$. □

11.9 Lemma and Definition Let G be a Lie group. Evaluation at the unit element $\xi \mapsto \xi(1)$ defines an isomorphism from the vector space of left invariant vector fields on G to the Lie algebra $\operatorname{Lie} G$. This isomorphism gives the vector space $\operatorname{Lie} G$ a real Lie algebra structure and thereby justifies its name. If $f\colon G \to H$ is a Lie group homomorphism then the induced linear mapping

$$\operatorname{Lie} G \xrightarrow{\operatorname{Lie} f} \operatorname{Lie} H$$

is a homomorphism of Lie algebras.

Proof In order to invert the evaluation map $\xi \mapsto \xi(1)$ we assign to each vector $\xi_1 \in \operatorname{Lie} G$ the constant section $G \ni x \mapsto (x, \xi_1) \in G \times \operatorname{Lie} G$ and then compose with the canonical trivialisation 11.5: the resulting vector field $G \ni x \mapsto T_1\lambda_x \cdot \xi_1 \in TG$ is left invariant by definition.

As to functoriality with respect to $f\colon G \to H$ we note that the differential $\operatorname{Lie} f = T_1 f$ may now be re-interpreted as a map which assigns to every left invariant vector field $\xi\colon G \to TG$ a left invariant vector field $f_*\xi\colon H \to TH$ such that $(f_*\xi)(1) = T_1 f \cdot \xi(1)$. This assignment makes the diagram

$$\begin{tikzcd} TG \arrow[r, "Tf"] & TH \\ G \arrow[u, "\xi"] \arrow[r, "f"] & H \arrow[u, "f_*\xi"'] \end{tikzcd}$$

commutative since f is a homomorphism: given a fixed $x \in G$ we calculate

$$T_x f \cdot \xi(x) = T_x f \cdot T_1\lambda_x \cdot \xi(1) = T_1\lambda_{f(x)} \cdot T_1 f \cdot \xi(1) = T_1\lambda_{f(x)} \cdot (f_*\xi)(1) = (f_*\xi)\big(f(x)\big).$$

Turning to the action of vector fields on a scalar function $h\colon H \to \mathbb{R}$ we now obtain the identity

$$\mathcal{L}_{f_*\xi} h \circ f = \langle f_*\xi, dh\rangle \circ f = \langle Tf \cdot \xi, dh \circ f\rangle = \langle \xi, f^* dh\rangle = \mathcal{L}_\xi (h \circ f)\,,$$

and if η is another left invariant vector field on G a second application of this identity yields

$$\mathcal{L}_{f_*\xi}\mathcal{L}_{f_*\eta} h \circ f = \mathcal{L}_\xi(\mathcal{L}_{f_*\eta} h \circ f) = \mathcal{L}_\xi\mathcal{L}_\eta(h \circ f)\,.$$

Using the representation 10.15 of the Lie bracket as a commutator we conclude that

$$\mathcal{L}_{[f_*\xi, f_*\eta]} h \circ f = \mathcal{L}_{[\xi,\eta]}(h \circ f) = \mathcal{L}_{f_*[\xi,\eta]} h \circ f\,.$$

As this holds for all scalar functions h the vector fields $[f_*\xi, f_*\eta]$ and $f_*[\xi,\eta]$ take the same value at least at the point $1 \in H$. Since both are left invariant fields on H they coincide altogether. □

Besides the interpretations of the Lie algebra as a tangent space, and the space of left invariant vector fields there are others, which we now summarize.

11.10 Theorem Let G be a Lie group. There are canonical bijections between the following sets, and they respect the vector space structures of the first and second:

- vectors in $\operatorname{Lie} G = T_1 G$, interpreted as tangent vectors at the unit element,
- elements of $\operatorname{Lie} G$, interpreted as left invariant vector fields on G,
- left invariant global flows on G, and
- one parameter groups of G.

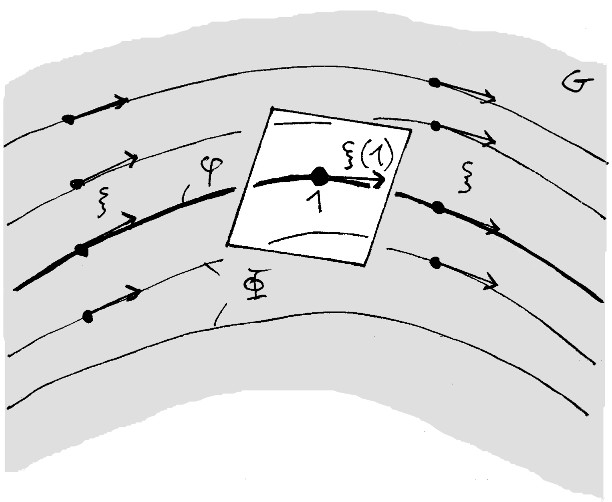

Proof The first two items have been linearly related in 11.9, and in 11.7 we have seen that every left invariant vector field defines a left invariant global flow, which in turn restricts to a one parameter group of G. If conversely $\varphi\colon \mathbb{R} \to G$ is a one parameter group with, say, $\dot\varphi(0) = \xi \in \operatorname{Lie} G$ then the formula

$$\mathbb{R} \times G \ni (t, x) \longmapsto x \cdot \varphi(t) \in G$$

defines a global flow on G which is clearly left invariant, and whose time derivative is

$$\frac{d}{dt}\, x \cdot \varphi(t)\Big|_{t=0} = \frac{d}{dt}\,(\lambda_x \circ \varphi)(t)\Big|_{t=0} = T_{\varphi(0)}\lambda_x \cdot \dot\varphi(0) = T_1\lambda_x \cdot \xi(1) = \xi(x)\ :$$

this must be the left invariant flow generated by ξ.

It is immediately verified that the various assignments we have described are compatible with each other and thus also provide inverses. □

In accordance with the basic idea of differential calculus we should expect the tangent space $\operatorname{Lie} G = T_1 G$ to hold at least local information about the Lie group G near the unit element. A first glance seems to disappoint that hope: since the multiplication mapping μ satisfies $\mu(x, 1) = x = \mu(1, x)$ for all $x \in G$ its differential $T_1\mu\colon T_1G \times T_1G \to T_1G$ must be just vector addition; therefore the differential $T_1\mu$ reveals no information whatsoever about μ itself. This may be interpreted saying that the group laws determine the multiplication of a Lie group to first — that is, linear — order. Particular features of the multiplication, including a failure to be commutative, are of higher order and not reflected in $\operatorname{Lie} G$ as a mere vector space. We will now see that the Lie algebra structure recovers the seemingly lost information. We begin with an alternative description of the Lie bracket.

11.11 Definition and Theorem Let G be a Lie group. The group homomorphism

$$\operatorname{Ad}\colon G \longrightarrow GL(\operatorname{Lie} G)$$

which assigns to $x \in G$ the differential at $1 \in G$ of the conjugation map $G \ni y \mapsto xyx^{-1} \in G$ is called the *adjoint representation of the group* G. Differentiating this homomorphism at $1 \in G$ and using the canonical identification of the tangent space $T_1 GL(\operatorname{Lie} G)$ with $\operatorname{End}(\operatorname{Lie} G)$ we obtain the linear mapping

$$\operatorname{ad}\colon \operatorname{Lie} G \longrightarrow \operatorname{End}(\operatorname{Lie} G)\,,$$

the *adjoint representation of the algebra* $\operatorname{Lie} G$. Its action is given by the Lie bracket:

$$\operatorname{ad}(\xi) \cdot \eta = [\xi, \eta] \quad \text{for all } \xi, \eta \in \operatorname{Lie} G.$$

Proof Given left invariant vector fields $\xi, \eta \in \operatorname{Lie} G$ we let $\varphi, \chi\colon \mathbb{R} \to G$ be the corresponding one parameter groups of G. For every $s \in \mathbb{R}$ we have

$$\operatorname{Ad}\big(\varphi(s)\big) \cdot \eta(1) = \operatorname{Ad}\big(\varphi(s)\big) \cdot \dot\chi(0) = \frac{\partial}{\partial t}\, \varphi(s)\, \chi(t)\, \varphi(s)^{-1}\ \Big|_{t=0}$$

by the chain rule, and further

$$\operatorname{ad}(\xi) \cdot \eta(1) = \operatorname{ad}\big(\dot\varphi(0)\big) \cdot \eta(1) = \frac{\partial}{\partial s}\frac{\partial}{\partial t}\, \varphi(s)\, \chi(t)\, \varphi(s)^{-1}\ \Big|_{s=t=0} \in T_1G = \operatorname{Lie} G\,.$$

This formula implies that the vector field $\operatorname{ad}(\xi) \cdot \eta$ acts on scalar functions $f\colon G \to \mathbb{R}$ in such a way that

$$(\mathcal{L}_{\operatorname{ad}(\xi)\cdot\eta} f)(1) = \frac{\partial^2}{\partial s\,\partial t}\, f\big(\varphi(s)\, \chi(t)\, \varphi(s)^{-1}\big)\ \Big|_{s=t=0}.$$

To see this we first show that the right hand side defines a derivation

$$D\colon f \longmapsto \frac{\partial^2}{\partial s\,\partial t}\, f\big(\varphi(s)\, \chi(t)\, \varphi(s)^{-1}\big)\ \Big|_{s=t=0}$$

at the unit element $1 \in G$, satisfying $D(f \cdot g) = Df \cdot g(1) + f(1) \cdot Dg$ for all functions $f, g\colon G \to \mathbb{R}$. Indeed, abbreviating $e(s,t) := \varphi(s)\,\chi(t)\,\varphi(s)^{-1}$ we observe $e(s,0) = 1$ for all s, and calculate

$$D(f \cdot g) = \frac{\partial}{\partial s}\left(\left\langle \frac{\partial e}{\partial t}, df \circ e\right\rangle \cdot (g \circ e) + (f \circ e) \cdot \left\langle \frac{\partial e}{\partial t}, dg \circ e\right\rangle\right)\Bigg|_{s=t=0} = Df \cdot g(1) + f(1) \cdot Dg.$$

By Theorem 10.9 the derivation D must correspond to some tangent vector in T_1G. To identify this tangent vector we choose a chart h centred at the point $1 \in G$: our previous calculation shows that $(\mathcal{L}_{\mathrm{ad}(\xi)\cdot\eta} f)(1) = Df$ certainly holds if f is any of the cartesian components of h. Therefore $\mathrm{ad}(\xi) \cdot \eta(1) \in T_1G$ is the tangent vector that corresponds to D, and the formula therefore holds for all functions f.

Since the flows generated by ξ and η are $(s,x) \mapsto x \cdot \varphi(s)$ and $(t,y) \mapsto y \cdot \chi(t)$, we can compute

$$\frac{\partial}{\partial t} f\big(\varphi(s)\chi(t)\big)\Big|_{t=0} = \mathcal{L}_\eta f\big(\varphi(s)\big)$$

and further

$$\frac{\partial^2}{\partial s\,\partial t} f\big(\varphi(s)\chi(t)\big)\Big|_{s=t=0} = \frac{\partial}{\partial s} \mathcal{L}_\eta f\big(\varphi(s)\big)\Big|_{s=0} = \mathcal{L}_\xi \mathcal{L}_\eta f\,(1)$$

as well as, switching the roles of φ and χ,

$$\frac{\partial^2}{\partial s\,\partial t} f\big(\chi(t)\,\varphi(s)\big)\Big|_{s=t=0} = \mathcal{L}_\eta \mathcal{L}_\xi f\,(1).$$

The fact that the differential $T_1\mu(1,1)$ is just vector addition in T_1G implies that products in G are differentiated like usual bilinear products: the differential of a product $z \mapsto a(z) \cdot b(z) = \mu\big(a(z), b(z)\big)$ at a point z with $a(z) = b(z) = 1 \in G$ is

$$\zeta \mapsto \begin{pmatrix} 1 & 1 \end{pmatrix} \begin{pmatrix} T_z a \cdot \zeta \\ T_z b \cdot \zeta \end{pmatrix} = (T_z a + T_z b) \cdot \zeta.$$

This allows us to finally evaluate

$$\begin{aligned}
(\mathcal{L}_{\mathrm{ad}(\xi)\cdot\eta} f)(1) &= \frac{\partial^2}{\partial s\,\partial t} f\big(\varphi(s) \cdot \chi(t) \cdot \varphi(-s)\big)\Big|_{s=t=0} \\
&= \frac{\partial^2}{\partial s\,\partial t} f\big(\varphi(s) \cdot \chi(t)\big)\Big|_{s=t=0} - \frac{\partial^2}{\partial s\,\partial t} f\big(\chi(t) \cdot \varphi(s)\big)\Big|_{s=t=0} \\
&= \mathcal{L}_\xi \mathcal{L}_\eta f\,(1) - \mathcal{L}_\eta \mathcal{L}_\xi f\,(1) \\
&= (\mathcal{L}_{[\xi,\eta]} f)(1).
\end{aligned}$$

Thus the vector fields $\mathrm{ad}(\xi) \cdot \eta$ and $[\xi, \eta]$ agree at the unit element. Since both are left invariant they agree everywhere. □

11.12 Examples (1) For an abelian Lie group the adjoint representation of the group is constant, and the Lie bracket vanishes identically: Lie algebras with this property are likewise called *abelian* or *commutative* Lie algebras.

(2) The adjoint representation of the group $GL(n,\mathbb{R})$ with Lie algebra $\mathfrak{gl}(n,\mathbb{R}) = \mathrm{Mat}(n \times n, \mathbb{R})$ is given by

$$\mathfrak{gl}(n,\mathbb{R}) \ni \eta \xmapsto{\mathrm{Ad}(x)} x\,\eta\,x^{-1} \in \mathfrak{gl}(n,\mathbb{R})\,,$$

inducing that of the Lie algebra

$$\mathfrak{gl}(n,\mathbb{R}) \ni \eta \xmapsto{\mathrm{ad}(\xi)} \xi\,\eta - \eta\,\xi \in \mathfrak{gl}(n,\mathbb{R})\,.$$

Therefore the Lie bracket $[\xi,\eta]$ of two matrices has the usual meaning as the commutator $\xi\,\eta-\eta\,\xi$. By naturality of the Lie algebra structure this property is inherited by all Lie subgroups of $GL(n,\mathbb{R})$.

Like curves on any manifold, one parameter groups of a Lie group have velocity vectors and therefore induce infinitesimal data in tangent spaces. A special feature of Lie groups is that there also is a correspondence in the opposite direction.

11.13 Definition and Lemma Let G be a Lie group. The *exponential map*

$$\operatorname{Lie} G \xrightarrow{\exp} G$$

assigns to the vector $\xi \in \operatorname{Lie} G$ the value $\varphi(1) \in G$ of the one parameter group $\varphi\colon \mathbb{R} \to G$ which corresponds to ξ.

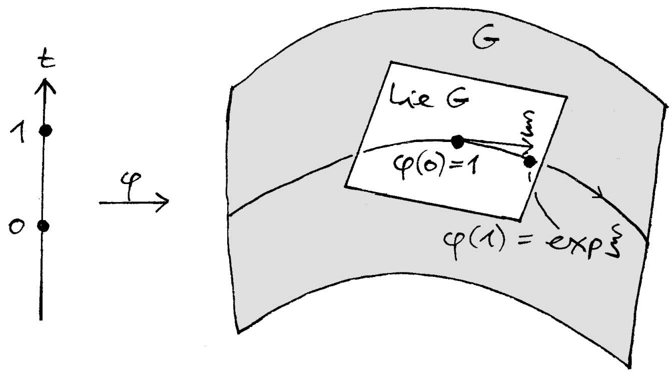

The exponential map is smooth, and its differential at $0 \in \operatorname{Lie} G$ is the identity mapping. The exponential map also is natural: for every homomorphism of Lie groups $f\colon G \to H$ the diagram

$$\begin{tikzcd}
\operatorname{Lie} G \arrow[r, "\operatorname{Lie} f"] \arrow[d, "\exp"'] & \operatorname{Lie} H \arrow[d, "\exp"] \\
G \arrow[r, "f"] & H
\end{tikzcd}$$

commutes.

Proof Throughout the proof we treat elements of the Lie algebra consistently as tangent vectors $\xi \in T_1G$. We denote the corresponding one parameter group by φ_ξ and recall that

$$\mathbb{R} \times G \ni (t,x) \longmapsto x \cdot \varphi_\xi(t) \in G$$

is the global flow of ξ, whose (left invariant) velocity field sends $x \in G$ to $T_1\lambda_x \cdot \xi \in T_xG$.

The first point in question is the smooth dependence of φ_ξ on ξ. To this end we observe that the tangent bundle of the manifold $T_1G \times G$ contains the subbundle

$$\begin{tikzcd}
\operatorname{pr}^* TG \arrow[rr, hook] \arrow[dr] & & T(T_1G \times G) \arrow[dl] \\
& T_1G \times G &
\end{tikzcd}$$

induced from TG under the cartesian projection $T_1G \times G \xrightarrow{\operatorname{pr}} G$. The section

$$T_1G \times G \ni (\xi,x) \longmapsto T_1\lambda_x \cdot \xi \in \operatorname{pr}^* TG$$

of this bundle therefore defines a smooth vector field $\tilde{\xi}$ on $T_1G \times G$. The maximal flow determined by $\tilde{\xi}$ acts as

$$\mathbb{R} \times T_1G \times G \ni (t, \xi, x) \longmapsto \tilde{\Phi}(t, \xi, x) := \big(\xi,\, x \cdot \varphi_\xi(t)\big) \in T_1G \times G\,.$$

Indeed the values of this formula depend smoothly on t and satisfy the differential equation

$$\frac{d}{dt}\tilde{\Phi}(t,\xi,x) = \big(0, T_{\varphi_\xi(t)}\lambda_x \cdot T_1\lambda_{\varphi_\xi(t)} \cdot \xi\big) = \big(0, T_1\lambda_{x\cdot\varphi_\xi(t)} \cdot \xi\big) = \tilde{\xi}\big(\xi, x \cdot \varphi_\xi(t)\big) = (\tilde{\xi} \circ \tilde{\Phi})(t,\xi,x)$$

as well as the initial condition $\tilde{\Phi}(0,\xi,x) = (\xi, x)$ for all $(\xi, x) \in T_1G \times G$. In particular we now know that $\tilde{\Phi}$ is smooth, and since the exponential map can be written as the composition

$$\operatorname{Lie} G \ni \xi \mapsto (1, \xi, 1) \in \mathbb{R} \times \operatorname{Lie} G \times G \xrightarrow{\tilde{\Phi}} \operatorname{Lie} G \times G \xrightarrow{\mathrm{pr}} G$$

it is smooth too.

In order to compute the differential

$$\operatorname{Lie} G = T_0 \operatorname{Lie} G \xrightarrow{T_0 \exp} \operatorname{Lie} G$$

we consider a tangent vector $\xi \in \operatorname{Lie} G$. For every $t \in \mathbb{R}$ the curve

$$\tau \longmapsto \varphi_\xi(t \cdot \tau)$$

is a one parameter group with initial velocity vector $t \cdot \xi$, and therefore coincides with $\varphi_{t\xi}$. This implies that

$$\frac{d}{dt}\, \exp(t \cdot \xi)\,\bigg|_{t=0} = \frac{d}{dt}\, \varphi_{t\xi}(1)\,\bigg|_{t=0} = \frac{d}{dt}\, \varphi_\xi(t)\,\bigg|_{t=0} = \xi,$$

so that $T_0 \exp$ is the identity as claimed.

Finally the naturality of the exponential map follows from the observation that for given $\xi \in T_1G$ the curve $f \circ \varphi_\xi$ is a one parameter group of H with initial velocity $T_1f \cdot \dot{\varphi}_\xi(0) = T_1f \cdot \xi$. □

11.14 Examples (1) Let V be a finite dimensional real vector space, considered as a Lie group. To the tangent vector $\xi \in \operatorname{Lie} V = V$ there corresponds the one parameter group $\mathbb{R} \ni t \mapsto t\xi \in V$, and the exponential map is the identical mapping. If $\Lambda \subset V$ is a lattice and we pass from V to the quotient group V/Λ then the exponential map essentially becomes the quotient homomorphism:

$$\operatorname{Lie}(V/\Lambda) = \operatorname{Lie} V = V \longrightarrow V/\Lambda$$

(2) The group $GL(n,\mathbb{R})$ has Lie algebra $\mathfrak{gl}(n,\mathbb{R}) = \operatorname{Mat}(n \times n, \mathbb{R})$, and the one parameter group corresponding to the matrix $\xi \in \mathfrak{gl}(n,\mathbb{R})$ is

$$\mathbb{R} \ni t \longmapsto \exp(t\xi) = \sum_{j=0}^{\infty} \frac{1}{j!}\, t^j \xi^j \in GL(n,\mathbb{R})\,,$$

so that the exponential map is

$$\mathfrak{gl}(n,\mathbb{R}) \ni \xi \longmapsto \exp \xi = \sum_{j=0}^{\infty} \frac{1}{j!}\, \xi^j \in GL(n,\mathbb{R})\,.$$

The name of exponential map goes back to this fundamental example. That we are truly dealing with a one parameter group in $GL(n,\mathbb{R})$ follows from the observation $\exp 0 = 1$ and the functional relation

$$\exp(\xi+\eta) = \sum_{m=0}^{\infty} \frac{1}{m!}\, (\xi+\eta)^j = \sum_{j,k=0}^{\infty} \frac{1}{j!}\frac{1}{k!}\, \xi^j \eta^k = \exp \xi \cdot \exp \eta$$

valid for all *commuting* matrices ξ and η: the proof relies on absolute convergence and may be copied word by word from the well-known scalar case.

By naturality of the exponential map its description as a power series passes to all Lie subgroups of $GL(n,\mathbb{R})$. The simplest example is the special linear group $SL(n,\mathbb{R}) \subset GL(n,\mathbb{R})$ of Example 2.4(4) with Lie algebra

$$\mathfrak{sl}(n,\mathbb{R}) = \left\{\xi \in \mathrm{Mat}(n\times n,\mathbb{R}) \,\middle|\, \operatorname{tr}\xi = 0\right\},$$

and it shows in particular that the condition $\operatorname{tr}\xi = 0$ implies $\det\exp\xi = 1$.

The fact that the exponential map is a local diffeomorphism at the origin of the Lie algebra — as follows from 11.13 by the local inverse theorem — makes the relation between a Lie group and its tangent space at the unit a particularly tight one. A first striking consequence is this:

11.15 Theorem Let G and H be Lie groups, G connected. Homomorphisms $G \xrightarrow{f} H$ are uniquely determined by their differentials $\mathrm{Lie}\,G \xrightarrow{\mathrm{Lie}\,f} \mathrm{Lie}\,H$.

Proof The image of $\mathrm{Lie}\,G \xrightarrow{\exp} G$ contains a neighbourhood of $1 \in G$; therefore the subgroup generated by this image is an open subgroup of G. All its cosets in G likewise are open, and since G is connected this subgroup must be all of G.

By naturality of the exponential map two homorphisms f and f' with $\mathrm{Lie}\,f = \mathrm{Lie}\,f'$ must coincide on the image set $\exp(\mathrm{Lie}\,G) \subset G$ and therefore on the subgroup generated by it, which we have just seen is G itself. □

The exponential map need not be a group homomorphism, nor need it be surjective even for a connected Lie group. The following lemma gives more precise information on the first point.

11.16 Theorem Let be a connected Lie group. The exponential map $\mathrm{Lie}\,G \xrightarrow{\exp} G$ is a group homomorphism if and only if G is abelian. In this case exp also surjective.

Proof Assume that $\mathrm{Lie}\,G \xrightarrow{\exp} G$ is a homomorphism. Then its image is a subgroup which — as we saw in the previous proof — must be all G. Thus the exponential map is surjective, and since $\mathrm{Lie}\,G$ is commutative under addition so is the group multiplication of G.

Assume conversely that G is abelian. Then the multiplication map $G\times G \xrightarrow{\mu} G$ is a homomorphism of Lie groups, and by naturality of the exponential map the diagram

$$\begin{tikzcd}
\mathrm{Lie}\,G \times \mathrm{Lie}\,G \arrow[r, "\mathrm{Lie}\,\mu"] \arrow[d, "\exp_{G\times G}"'] & \mathrm{Lie}\,G \arrow[d, "\exp_G"] \\
G\times G \arrow[r, "\mu"] & G
\end{tikzcd}$$

commutes. The exponential map of $G\times G$ is $\exp_G\times\exp_G$, and the differential of μ is vector addition in $\mathrm{Lie}\,G$: thus the commutativity of the diagram just means that $\exp_G$ is a group homomorphism. □

11.17 Corollary Every connected abelian Lie group is isomorphic to $(S^1)^r \times \mathbb{R}^s$ for integers $r, s \in \mathbb{N}$.

Proof Let G be a connected abelian Lie group. By Theorem 11.16 the exponential map is a surjective homomorphism of Lie groups $\mathrm{Lie}\,G \xrightarrow{\exp} G$. We know from 11.13 that $\mathrm{Lie}\,G \xrightarrow{\exp} G$ is a local diffeomorphism at $0 \in \mathrm{Lie}\,G$; therefore $\Lambda := \operatorname{kernel}\exp \subset \mathrm{Lie}\,G$ is a discrete subgroup, and the induced homomorphism $(\mathrm{Lie}\,G)/\Lambda \to G$ is a diffeomorphism and thus an isomorphism of Lie groups. On the other hand, as a discrete subgroup of a finite dimensional real vector space Λ is known[3] to be a lattice, that is, the additive subgroup spanned by a set of linearly independent vectors in $\mathrm{Lie}\,G$. Up to isomorphism we may assume that $\mathrm{Lie}\,G = \mathbb{R}^{r+s}$ is a Euclidean space and that $\Lambda \subset \mathrm{Lie}\,G$ is

[3] If you wish to prove this the hint for Exercise 2.13 may be helpful.

spanned by the first r standard base vectors $e_1, \dots, e_r \in \mathbb{R}^{r+s} = \operatorname{Lie} G$. This makes $(\operatorname{Lie} G)/\Lambda$ and thereby G isomorphic to $\mathbb{R}^r/\mathbb{Z}^r \times \mathbb{R}^s = \left(S^1\right)^r \times \mathbb{R}^s$. □

In view of this corollary the theory of Lie groups is mainly about Lie groups which are not commutative. A fact that unexpectedly carries over from the commutative case is that the algebraic structure of a Lie group is quite faithfully reflected in the — now non-trivial — multiplicative structure of its Lie algebra: this allows to study Lie groups by methods of linear algebra and thus once more perfectly fits into the general concept of differential calculus.

Exercises

11.1 Let G be a Lie group. Show that the connected component of the unit element is a normal Lie subgroup of G.

11.2 Prove that every bijective homomorphism of Lie groups is an isomorphism.

11.3 Show that that every connected Lie group of dimension 1 is abelian.

11.4 Let $\varphi: \mathbb{R} \to SO(3)$ be a non-trivial one parameter group. Prove that the image of φ is a Lie subgroup isomorphic to S^1.

11.5 Let G be a Lie group. Explain why the right invariant vector fields on G — those $\xi \in \operatorname{Vect} G$ with $\rho_y^* \xi = \xi$ for all $y \in G$ — form a Lie subalgebra of $\operatorname{Vect} G$ and prove that this subalgebra is canonically isomorphic to $\operatorname{Lie} G$.

11.6 Let G be a Lie group. Prove the formula

$$\exp\bigl(\operatorname{Ad}(g) \cdot \xi\bigr) = g \cdot \exp \xi \cdot g^{-1}$$

for all $g \in G$ and $\xi \in \operatorname{Lie} G$.

11.7 Let G be a holomorphic Lie group. Prove that if G is compact and connected, then G is abelian.

11.8 Prove that the exponential map of the Lie group $SO(n)$ is surjective.

11.9 Determine the image of the exponential map of the Lie group $SL(2, \mathbb{R})$.

Hint Study the trace of $\exp \xi$ for $\xi \in \mathfrak{sl}(2, \mathbb{R})$.

12 Symplectic Manifolds

12.1 Symplectic Linear Algebra A *symplectic structure* on a finite dimensional real vector space V is a form $\omega \in \mathrm{Alt}^2 V$ which has full rank in the sense that the linear mapping

$$V \ni v \longmapsto \omega(v, ?) \in V^{\check{}}$$

is an isomorphism. A *symplectic vector space* is a vector space with a symplectic form on it, and a linear isomorphism $f\colon (V,\omega) \to (V',\omega')$ between such spaces is *symplectic* if $\omega = f^*\omega'$.

Every symplectic vector space (V,ω) admits *symplectic bases*, with respect to which the skew-symmetric matrix of ω is

$$\begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix} \in \mathrm{Mat}(2n \times 2n, \mathbb{R}),$$

built from the unit matrix $1 \in \mathrm{Mat}(n \times n, \mathbb{R})$; in particular $\dim V = 2n$ must be even. If the base $(b_1, \ldots, b_n, c_1, \ldots, c_n)$ of $V^{\check{}}$ is dual to a symplectic base we have

$$\omega = \sum_{j=1}^{n} b_j \wedge c_j \in \mathrm{Alt}^2 V,$$

and therefore

$$\frac{1}{n!}\,\omega^n = \frac{1}{n!}\,\omega \wedge \cdots \wedge \omega = (b_1 \wedge c_1) \wedge (b_2 \wedge c_2) \wedge \cdots \wedge (b_n \wedge c_n) \in \mathrm{Alt}^{2n} V$$

is a volume form on V.

Note that a linear subspace $L \subset V$ does not in general inherit a symplectic structure from (V,ω). While its orthogonal complement

$$L^{\perp} := \{v \in V \mid \omega(v,w) = 0 \text{ for all } w \in L\} \subset V$$

always has the complementary dimension $2n - \dim L$ it need not be a linear complement to L in V. If it is then ω does restrict to a symplectic structure on L. At the other extreme subspaces $L \subset V$ with $L \subset L^{\perp}$ are called *isotropic*; among them there are all one-dimensional subspaces. The maximal isotropic subspaces have dimension n and are called *Lagrangian*. For instance in terms of any symplectic base $\mathbb{R}^n \times \{0\}^n$ and $\{0\}^n \times \mathbb{R}^n$ correspond to Lagrangian subspaces.

12.2 Definition A *symplectic manifold* is a differentiable manifold X with a *symplectic form* on it, which in turn means a differential form $\omega \in \Omega^2 X$ which is closed and restricts for each $x \in X$ to a symplectic structure $\omega_x \in \mathrm{Alt}^2 T_x X$ on the tangent space. A local diffeomorphism $f\colon (X,\omega) \to (X',\omega')$ is called *symplectic* if $\omega = f^*\omega'$.

12.3 Examples (1) Every open subset $X \subset \mathbb{R}^{2n} = \{(p_1, \ldots, p_n, q_1, \ldots, q_n)\}$ carries the standard symplectic form

$$\omega = \sum_{j=1}^{n} dp_j \wedge dq_j$$

and thus is a symplectic manifold.

(2) A symplectic structure on a surface — a smooth manifold X of dimension two — is the same as a volume form (or rather a surface form) on X.

(3) Let X be an arbitrary n-manifold. The total space of the cotangent bundle $T^\vee X \xrightarrow{\pi} X$ carries a natural 1-form $\psi \in \Omega^1 T^\vee X$ which is defined as follows. Given a point $\xi \in (T_x X)^\vee \subset T^\vee X$ and a tangent vector $v \in T_\xi(T^\vee X)$ the value of the linear form ψ_ξ on v is

$$\psi_\xi(v) = \big\langle v, (T_\xi \pi)^* \xi \big\rangle_{T_\xi(T^\vee X)} = \big\langle T_\xi \pi \cdot v, \xi \big\rangle_{T_x X} \in \mathbb{R}.$$

Every chart (U, q) for X induces a chart

$$T^\vee q \colon T^\vee U \longrightarrow q(U) \times \mathbb{R}^n$$

for the manifold $T^\vee X$ — essentially the corresponding bundle chart for the cotangent bundle; it sends $\xi \in (T_x U)^\vee$ to the pair $\big(q(x), p(\xi)\big) \in q(U) \times \mathbb{R}^n$ where the component p_j of $p(\xi) = (p_1, \ldots, p_n)$ is determined by evaluation on the j-th standard vector $e_j \in \mathbb{R}^n$:

$$p_j = \xi \cdot (T_x q)^{-1} \cdot e_j = \Big\langle \frac{\partial}{\partial q_j}, \xi \Big\rangle$$

Note that the diagram

$$\begin{tikzcd} T^\vee U \arrow[r, "T^\vee q = (q,p)"] \arrow[d, "\pi"'] & q(U) \times \mathbb{R}^n \arrow[d, "\mathrm{pr}"] \\ U \arrow[r, "q"] & q(U) \end{tikzcd}$$

commutes by definition. Returning to the natural 1-form ψ defined above we see that its value on the tangent vector $\dfrac{\partial}{\partial q_j} \in T_\xi(T^\vee U)$ is just p_j. Since ψ clearly annihilates all tangent vectors to the fibres of π this means that

$$\psi = \sum_{j=1}^{n} p_j \cdot dq_j$$

is the expression of ψ in terms of the chart $(T^\vee U, T^\vee q) = (T^\vee U, q, p)$. It is remarkable that the expression of this differential form is the same for all charts q. It further results that the 2-form

$$\omega := d\psi = \sum_{j=1}^{n} dp_j \wedge dq_j \in \Omega^2 T^\vee X$$

is a symplectic form: thus the total space of every cotangent bundle is a symplectic manifold in a canonical way.

12.4 Darboux's Theorem Let X be a symplectic $2n$-manifold, and let $o \in X$ be a point. Then there exists a *symplectic chart* (U, p, q) at o: one which is a symplectic diffeomorphism onto an open subset $V \subset \mathbb{R}^{2n}$ with its standard symplectic form $\sum_j dp_j \wedge dq_j$.

Notes We have recorded in 12.1 that every symplectic vector space admits a symplectic base. Stated in another way, any two $2n$-dimensional symplectic vector spaces are symplectically isomorphic to each other. In a similar vein Darboux's theorem states that any two symplectic $2n$-manifolds are locally related by symplectic diffeomorphisms. This theorem thus stands in the same relation to the existence of symplectic bases as the constant rank theorem 2.5 to the existence of normal forms for linear mappings between two vector spaces.

As a consequence of Darboux's theorem there can be no local invariant that would distinguish between two different symplectic manifolds.[1]

[1] The situation is quite different in differential geometry, which studies smooth manifolds equipped with a Riemannian metric on the tangent bundle. In differential geometry the curvature is just such an invariant.

Proof It will be shown that locally any given symplectic form can be smoothly deformed into a standard form. We begin with a few simple preparations.

The statement being local we need but consider the case where $X \subset \mathbb{R}^{2n}$ is an open subset and $o = 0$. Also after a linear transformation we may assume that the value of the symplectic form ω at $0 \in X$ is the standard form

$$\omega_0 = \sum_{j=1}^{n} dp_j \wedge dq_j \in \operatorname{Alt}^2 \mathbb{R}^{2n}.$$

As a final preparation we shrink $X \subset \mathbb{R}^{2n}$ to an open ball centred at the origin.

On the $(2n+1)$-manifold $\mathbb{R}\times X$ we denote the extra coordinate by t and consider the differential forms

$$\tilde{\omega} := (1-t)\cdot\omega + t\cdot\omega_0 \in \Omega^2(\mathbb{R}\times X) \quad\text{and}\quad \Omega := dt\wedge\tilde{\omega} \in \Omega^3(\mathbb{R}\times X)$$

where $\omega \in \Omega^2(\mathbb{R}\times X)$ is shorthand for the pull-back $\operatorname{pr}^*\omega$ under the projection $\operatorname{pr}\colon \mathbb{R}\times X \to X$, while ω_0 is now re-interpreted as a differential form with constant coefficients in $\Omega^2 X$ or $\Omega^2(\mathbb{R}\times X)$. We define the open subset $Y \subset \mathbb{R}\times X$ as

$$Y = \big\{(t,x) \in \mathbb{R}\times X \,\big|\, \tilde{\omega}_{t,x}|(\{0\}\times\mathbb{R}^{2n}) \in \operatorname{Alt}^2\mathbb{R}^{2n} \text{ has full rank}\big\}$$

and note that $\mathbb{R}\times\{0\} \subset Y$.

The central tool in the proof Darboux's theorem is a maximal flow

$$\mathbb{R}\times Y \supset D \xrightarrow{\Phi} Y$$

that we will construct on the manifold Y. It will have the following properties:

(1) $\Phi(\tau,t,0)$ is defined for all $\tau, t \in \mathbb{R}$, and its value is $\Phi(\tau,t,0) = (\tau+t,0)$.

(2) The diagram

$$\begin{tikzcd}
\mathbb{R}\times Y \arrow[d, "\mathrm{id}\times\mathrm{pr}"'] & D \arrow[l, hook'] \arrow[r, "\Phi"] & Y \arrow[d, "\mathrm{pr}"] \\
\mathbb{R}\times\mathbb{R} \arrow[rr, "(\tau,t)\longmapsto\tau+t"] & & \mathbb{R}
\end{tikzcd}$$

commutes.

(3) At every point (t,x) the identity $\Phi_\tau^*\Omega = \Omega$ holds for all $\tau \in \mathbb{R}$ such that $(\tau,t,x) \in D$.

Once we have such a flow the conclusion of the theorem will follow. For (1) implies that there is an open neighbourhood $U \subset X$ of the origin such that $[0,1]\times\{0\}\times U \subset D$, so that in particular the restriction of Φ

$$\Phi_1\colon \{0\}\times U \longrightarrow Y$$

is defined. By (2) we have $\Phi_\tau(t,x) \in \{\tau+t\}\times X$ whenever the left hand side is defined, in particular Φ_1 sends $\{0\}\times U$ into $\{1\}\times X$ and thus restricts to a diffeomorphism

$$\Phi_1\colon \{0\}\times U \xrightarrow{\cong} \{1\}\times V$$

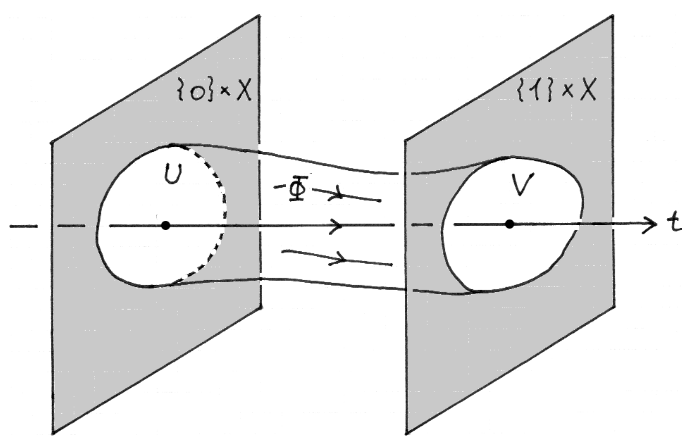

for a suitable open subset $V \subset X$ which by (1) contains the origin. Lastly from (3) we know that $\Omega = \Phi_1^*\Omega$, or

$$dt \wedge \tilde\omega = \Phi_1^*(dt \wedge \tilde\omega) = d\Phi_1^* t \wedge \Phi_1^*\tilde\omega = dt \wedge \Phi_1^*\tilde\omega.$$

Partial evaluation on the vector field $\frac{\partial}{\partial t}$ yields

$$\tilde\omega = \Phi_1^*\tilde\omega - dt \wedge \Big(\frac{\partial}{\partial t} \lrcorner \Phi_1^*\tilde\omega\Big)$$

— recall from the proof of 10.12 that partial evaluation is a graded derivation and note that $\tilde\omega$ contains no dt. Restricting this identity of 2-forms to $\{0\}\times U$ we finally obtain

$$\omega = \tilde\omega|_{t=0} = (\Phi_1^*\tilde\omega)|_{t=0} = \Phi_1^*\left(\tilde\omega|_{t=1}\right) = \Phi_1^*\omega_0.$$

Therefore Φ_1, read as a mapping $U \xrightarrow{\simeq} V$, is a symplectic chart for X.

We now turn to the construction of the maximal flow Φ. We rely, of course, on the standard technique of constructing the corresponding vector field ξ on Y instead. In order to learn which properties we must give ξ we differentiate the requirements listed under (1–3) and thereby translate them into properties of ξ.

(1) This becomes $\xi(t,0) = \dfrac{d}{d\tau}\Phi(\tau,t,0)\Big|_{\tau=0} = \dfrac{d}{d\tau}(\tau+t,0)\Big|_{\tau=0} = \dfrac{\partial}{\partial t}$ for all $t \in \mathbb{R}$.

(2) Here differentiation yields

$$T_{t,x}\,\mathrm{pr}\cdot\xi(t,x) = \frac{d}{d\tau}\big(\mathrm{pr}\circ\Phi(\tau,t,x)\big)\Big|_{\tau=0} = \frac{d}{d\tau}(\tau+t)\Big|_{\tau=0} = \frac{\partial}{\partial t} \quad \text{for all } (t,x) \in Y.$$

In other words ξ must have the form

$$\xi = \frac{\partial}{\partial t} + \xi' = \frac{\partial}{\partial t} + \sum_{j=1}^{n} \xi_j \frac{\partial}{\partial p_j} + \sum_{j=1}^{n} \eta_j \frac{\partial}{\partial q_j}$$

with coefficient functions $\xi_j, \eta_j \colon Y \to \mathbb{R}$ which by (1) must vanish identically along $\mathbb{R}\times\{0\}$.

(3) Translation of the last condition simply gives

$$\mathcal{L}_\xi\Omega = \frac{d}{d\tau}(\Phi_\tau^*\Omega)\Big|_{\tau=0} = \frac{d}{d\tau}\Omega\Big|_{\tau=0} = 0.$$

Conversely, if ξ is any vector field on Y with these three properties we integrate it to obtain the maximal flow $\mathbb{R}\times Y \supset D \xrightarrow{\Phi} Y$. By uniqueness of integral curves we recover the original conditions imposed on Φ:

(1) For every $t \in \mathbb{R}$ the assignment $\tau \mapsto (\tau+t,0)$ defines an integral curve of ξ, which must coincide with $\Phi_{t,0}$.

(2) For all $(t,x) \in Y$ the curve $\varphi\colon \tau \mapsto \mathrm{pr}\circ\Phi_{t,x}$ is an integral curve of the vector field $\frac{\partial}{\partial t}$ on $\mathbb{R}$, thus we must have $\varphi(\tau) = \tau+t$ for all τ.

(3) Fix any $(t,x) \in Y$. The time derivative of the curve

$$\tau \mapsto \left(\Phi_\tau^*\Omega\right)_{t,x} \in \mathrm{Alt}^3\,\mathbb{R}^{2n+1}$$

vanishes identically, so it is the curve with constant value $\Omega_{t,x}$.

It remains to construct a vector field ξ with properties (1–3). Let us first work out the third condition $\mathcal{L}_\xi\Omega = 0$ in more detail. We begin with

$$\mathcal{L}_\xi\Omega = \mathcal{L}_\xi(dt \wedge \tilde\omega) = d\mathcal{L}_\xi t \wedge \tilde\omega + dt \wedge \mathcal{L}_\xi\tilde\omega = dt \wedge \mathcal{L}_\xi\tilde\omega$$

using that $\mathcal{L}_\xi t = \langle \xi, dt\rangle = 1$ in view of (2). Expanding the Lie derivative

$$\begin{aligned}\mathcal{L}_\xi \tilde\omega &= d(\xi \lrcorner \tilde\omega) + \xi \lrcorner d\big((1-t)\cdot\omega + t\cdot\omega_0\big)\\ &= d(\xi \lrcorner \tilde\omega) - \xi \lrcorner \big(dt\wedge(\omega-\omega_0)\big)\\ &= d(\xi \lrcorner \tilde\omega) - (\omega-\omega_0) + dt\wedge\big(\xi \lrcorner (\omega-\omega_0)\big)\end{aligned}$$

we see that condition (3) becomes equivalent to

$$dt\wedge\big(d(\xi \lrcorner \tilde\omega) - (\omega-\omega_0)\big) = 0.$$

This suggests how to construct ξ: The symplectic forms ω and ω_0 are closed forms on X, and so is their difference $\omega-\omega_0 \in \Omega^2 X$. By Poincaré's lemma the latter even is exact, say $\omega-\omega_0 = d\psi$ with $\psi\in\Omega^1 X$. We adjust ψ so that it vanishes at the point $0\in X$; this is possible simply because the base forms

$$dh_{j_1}\wedge\cdots\wedge dh_{j_k}$$

with respect to any chart h whatever are closed.

Let

$$p\colon Y\subset \mathbb{R}\times X \xrightarrow{\mathrm{pr}} X$$

denote the restriction of the Cartesian projection. The assignment $\xi'\mapsto \xi'\lrcorner\tilde\omega$ defines a homomorphism

$$\begin{tikzcd} Y\times T_0X \arrow[r, equal] \arrow[dr, "\mathrm{pr}"'] & p^*TX \arrow[r, "\simeq"] \arrow[d] & p^*T^\vee X \arrow[r, equal] \arrow[d] & Y\times \mathrm{Alt}^1 T_0X \arrow[dl, "\mathrm{pr}"] \\ & Y \arrow[r, equal] & Y & \end{tikzcd}$$

of (trivial) vector bundles, in fact an isomorphism since by the choice of Y the alternating form $\tilde\omega_{t,x}|(\{0\}\times\mathbb{R}^{2n})\in \mathrm{Alt}^2\,\mathbb{R}^{2n}$ has full rank at all points $(t,x)\in Y$. Therefore there is a unique $\frac{\partial}{\partial t}$-free vector field

$$\xi' = \sum_{j=1}^n \xi_j\frac{\partial}{\partial p_j} + \sum_{j=1}^n \eta_j\frac{\partial}{\partial q_j}$$

on Y with $\xi'\lrcorner\tilde\omega = p^*\psi$. Since ψ vanishes at $0\in X$ the field ξ' vanishes along $\mathbb{R}\times\{0\}\subset Y$. We finally put $\xi = \frac{\partial}{\partial t}+\xi'$; then obviously (1) and (2) are true, and so is (3):

$$\mathcal{L}_\xi\Omega = dt\wedge\big(d(\xi\lrcorner\tilde\omega)-(\omega-\omega_0)\big) = dt\wedge\big(d(\xi'\lrcorner\tilde\omega)-d\psi\big)=0.$$

This completes the proof of Darboux's theorem.

12.5 Definition A vector field ξ on a symplectic manifold (X,ω) is called *symplectic* if

$$\mathcal{L}_\xi\omega = 0$$

or equivalently if the flow it generates preserves the symplectic form ω. As ω is closed the condition may be rephrased as $d(\xi\lrcorner\omega)=0$.

The easiest way to ensure this is to make $\xi\lrcorner\omega$ an exact form: let $H\colon X\to\mathbb{R}$ be any smooth function, and let the vector field ξ correspond to $-dH\in\Omega^1X$ under the isomorphism

$$\mathrm{Vect}\,X\ni\xi\longmapsto\xi\lrcorner\omega\in\Omega^1X.$$

In this case ξ is called a *Hamiltonian vector field*, and the function H, which is well-defined up to addition of a locally constant function, its *Hamiltonian*.

Notes Every symplectic vector field ξ on a $2n$-manifold X must preserve the symplectic volume of X:

$$\mathcal{L}_\xi\Big(\frac{1}{n!}\,\omega^n\Big) = \frac{1}{n!}\,n\cdot\mathcal{L}_\xi\omega\wedge\omega^{n-1}=0.$$

This fact was first observed studying the equations of motion of certain mechanical systems; it is known as Liouville's theorem though in the context of symplectic geometry it is a mere observation.

The Hamiltonian H of a vector field ξ is automatically invariant under the flow of ξ, since

$$\mathcal{L}_\xi H = \xi \lrcorner dH = \xi \lrcorner (-\xi \lrcorner \omega) = 0$$

— recall $\omega(v,v) = 0$ for all tangent vectors v. In classical language H is said to be a *first integral* of ξ. In particular cases there may exist other first integrals, but not more than n independent ones, in a sense that would need to be made precise.

12.6 Example The 1-dimensional pendulum consists of a rod suspended at one end with a mass fixed to the other; its movement under the earth's gravitational field is restricted to a vertical plane that passes through the point of suspension. The mass is idealised as concentrated in a single point while the rod is massless, and the gravitational field is assumed a constant parallel field that points downwards.

Using suitable units the *configuration space* of the pendulum — that is, the manifold of its possible positions — is the unit circle S^1 with a coordinate q (defined up to addition of multiples of 2π) which measures the counter-clockwise angle from the downward vertical to the rod.

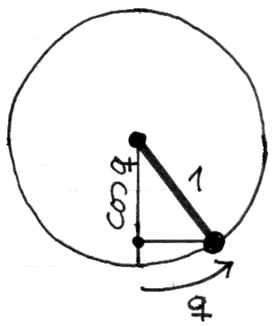

As in abundantly many physical systems the canonical fibre coordinate p on the cotangent bundle $T^\vee S^1 = S^1 \times \mathbb{R}$ is naturally identified with momentum, making $T^\vee S^1$ itself the *phase space* of the system. The total energy is a function $H: T^\vee S^1 \to \mathbb{R}$ called the *Hamiltonian*, and is easily determined explicitly as the sum of kinetic and potential energy:

$$H(p,q) = \frac{1}{2}p^2 - \cos q$$

The corresponding Hamiltonian vector field ξ gives the equations of motion of the pendulum. The possible trajectories of the system in phase space are the flow lines obtained by integrating ξ. Each flow line is uniquely determined by its starting point, and this is one advantage of phase over configuration space: while the flow lines project from $T^\vee S^1$ to trajectories in S^1 there are many of the latter that pass through a given starting point, corresponding to different initial velocities.

From the formula for H and the symplectic form $dp \wedge dq$ the field ξ is immediately worked out:

$$\xi = -\frac{\partial H}{\partial q} \cdot \frac{\partial}{\partial p} + \frac{\partial H}{\partial p} \cdot \frac{\partial}{\partial q} = -\sin q \cdot \frac{\partial}{\partial p} + p \cdot \frac{\partial}{\partial q}$$

Thus the equation of motion is the differential equation

$$\begin{aligned} \dot p &= -\sin q \\ \dot q &= p \end{aligned}$$

which would become the second order equation $\ddot q = -\sin q$ on the configuration space S^1.

The vector field ξ is globally integrable. As remarked above, the Hamiltonian H takes a constant value E — the total energy — along each flow line, and in this example energy furthermore turns out to provide a perfect classification of the orbits.

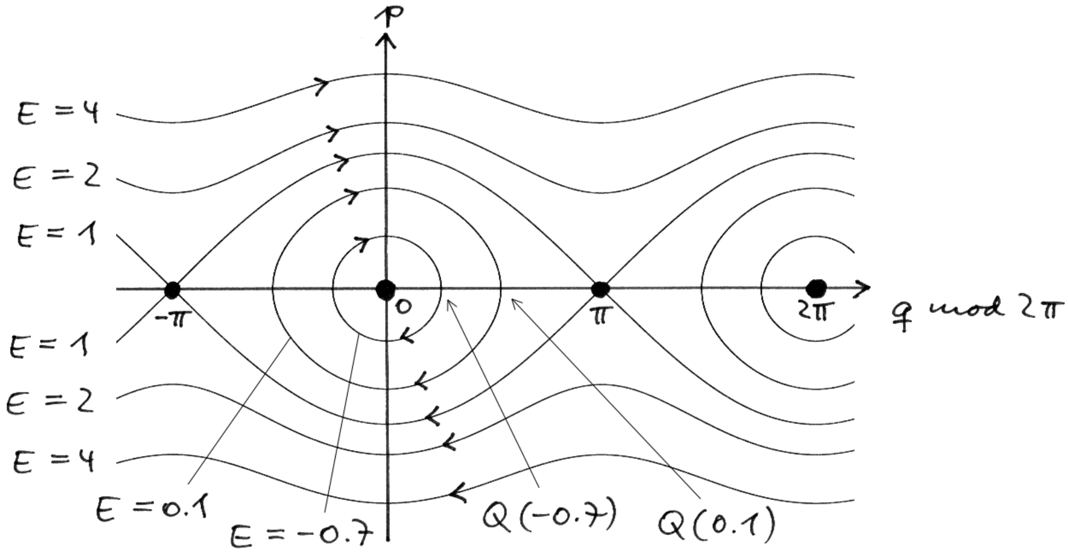

- The lowest possible value $E = -1$ is taken at $(q, p) = (0, 0) \in T^\vee S^1$, which is a fixed point and corresponds to the pendulum at rest in the stable equilibrium position.

- The other critical point of H is $(\pi, 0)$ with value $E = 1$. We immediately recognise the unstable equilibrium position of the pendulum. The remainder of the fibre $H^{-1}\{1\} \setminus \{(\pi, 0)\}$ is made up of two embedded orbits attached to $(\pi, 0)$ in the form of loops. They represent a pendulum that has just the energy necessary to rise to the top but never reaches it (in finite time as physicists would add).

- All other orbits are periodic. A fibre $H^{-1}\{E\}$ with $-1 < E < 1$ consists of a single orbit which is periodic and displays the typical pendulum movement with turning points at $\pm Q(E) := \arccos(-E)$. At the turning points all the energy is potential.

- Finally for $E > 1$ the kinetic energy is positive throughout, and the pendulum rotates continuously. The fibre $H^{-1}\{E\}$ accommodates two embedded orbits of which the sense of rotation selects one.

Exercises

12.1 Verify the facts of symplectic linear algebra stated at the beginning of the section.

12.2 The area form $\omega := \chi|S^2$ constructed in Exercise 6.1 turns the 2-sphere into a symplectic manifold. Compute the vector field ξ on S^2 derived from the height function $H = z$ as a Hamiltonian, and determine the flow generated by it.

12.3 Let X be a smooth n-manifold with a Riemannian metric s on its tangent bundle. According to 4.14(1) the metric determines a bundle isomorphism $TX \xrightarrow{\simeq} T^\vee X$, which may be used to pull back the canonical form ω of Example 12.3(3) to a symplectic form on the total space TX. Given a chart q of X, express this form in terms of the corresponding chart

$$Tq\colon (x,v) \longmapsto \big(q(x), T_xq\cdot v\big)$$

of TX and the scalar products $s_{ij} = s\left(\frac{\partial}{\partial q_i}, \frac{\partial}{\partial q_j}\right)$ for $i,j = 1,\ldots,n$.

Note Physicists have no scruples to write the i-th component of $T_xq\cdot v$ as $\dot q_i$ even though this is a free coordinate which does not depend in any way on the values of q. You may wish to do the exercise using his notation.

12.4 Prove what is stated in Example 12.6 about the trajectories of the pendulum.

12.5 Prove that in Example 12.6 the orbit of energy $E \in (-1,1)$ has period

$$T(E) = 2\sqrt{2}\int_0^{Q(E)} \frac{dq}{\sqrt{\cos q - \cos Q(E)}}$$

while those with $E > 1$ have period

$$T(E) = \sqrt{2}\int_0^{\pi} \frac{dq}{\sqrt{\cos q + E}}\,.$$