Abstract Manifolds and Differential Equations

Including Algorithmic Aspects

Markus Rosenkranz Georg Regensburger

January 30, 2007

Contents

Chapter 0

Preliminaries

0.1 Vector Spaces

Unless stated otherwise, we assume that all vector spaces (including algebras) are *finite-dimensional* over the ground field $\mathbb{R}$. Rings and algebras are understood to be *unital* (meaning they have a distinguished multiplicative identity element). The *null space* $O = \{0\}$ is regarded as a subspace of dimension 0 within any space $\mathbb{R}^n$, which means 0 has to be understood as the corresponding null tuple $(0, \ldots, 0) \in \mathbb{R}^n$. Following [22[8]], we will also use the *subspace convention* of regarding $\mathbb{R}^n \subseteq \mathbb{R}^{n+s}$ in the sense of $\mathbb{R}^n \times O \subseteq \mathbb{R}^n \times \mathbb{R}^s$. Furthermore, we set $V^\times = V \setminus O$ for any vector space V.

Initially, one can think of $\mathbb{R}^n$ as an *affine space* (hence also having a natural topology), and this is all that we needed up to now when speaking about maps between open sets of affine spaces. When it comes to derivatives, however, we need n-dimensional vector spaces. As an example, look at the first formula of Subsection 1.1.1. Strictly speaking, the point x there comes from the affine space $\mathbb{R}^n$ while the vector h comes from its translation space. In order to avoid an overkill in notation, we view $\mathbb{R}^n$ both as an affine and a vector space (of columns), but we will usually use x, y, z when thinking of points in the affine space $\mathbb{R}^n$.

Sometimes—particularly in connection with co- and contravariant tensors—it is convenient to regard a vector space V together with its dual V^* on an equal footing, as achieved by following terminology. A bilinear map $\langle | \rangle \colon \bar{V} \times V \to K$ on two K-vector spaces $\bar{V}$ and V is called a *dual pairing* if it is non-degenerate in both arguments [37[239]], meaning $\langle \bar{v} | \_ \rangle \colon V \to K$ and $\langle \_ | v \rangle \colon \bar{V} \to K$ are injective (and hence bijective) for all $\bar{v} \in \bar{V}$ and $v \in V$. One says that $\bar{V}$ and V are dually paired by $\langle | \rangle$, and $(\bar{V}, \langle | \rangle, V)$ is a dual pair of vector spaces. For the sake of naming, we call $\bar{V}$ the dual space and V the primal space of the pair; see the next paragraph for a motivation of this naming convention.

If $(\bar{W}, \langle | \rangle, W)$ is another dual pair (overloading the notation for the pairing), a morphism of dual pairs is given by two maps $\bar{T} \colon \bar{W} \to \bar{V}$ and $T \colon V \to W$ with the conservation property $\langle \bar{T}\bar{w} | v \rangle = \langle \bar{w} | Tv \rangle$ for all $\bar{w} \in \bar{W}$ and $v \in V$. It is easily verified that the dual pairs form a category under this notion of morphism. It is then also clear what we understand by an isomorphism of dual pairs. For an isomorphism $(\bar{T}, T)$, one may also write $S = \bar{T}^{-1}$

such that the conservation property gets the more symmetric form $\langle \bar{v}|v\rangle = \langle S\bar{v}|Tv\rangle$ for all $\bar{v} \in \bar{V}$ and $v \in V$.

Given a vector space V, one can construct the *induced dual pair* $P_V = (V^*, \langle|\rangle, V)$, where the canonical pairing is defined by $\langle v^*|v\rangle = v^*(v)$ for all $v^* \in V^*$ and $v \in V$. Dually, one may also construct the dual pair P_V^*, where the roles of V and V^* are exchanged and the dual pairing has its arguments reversed. Note the P_V and P_V^* are anti-isomorphic to each other. Unless mentioned otherwise, we always use the construction $P_V = (V^*, V)$ rather than $P_V^* = (V, V^*)$, omitting reference to the pairing $\langle|\rangle$. (The covariant functor $V \mapsto P_V$ is an isomorphisms of categories, the contravariant functor $V \mapsto P_V^*$ a duality of categories; hence $P_V \mapsto P_V^*$ is a duality on the category of dual pairs.)

We write $\mathbb{R}_n$ for the vector space of *rows* ($1 \times n$ matrices) and $\mathbb{R}^n$ for the vector space of *columns* ($n \times 1$ matrices). Rows are also known as (linear) forms, and we often use the letters a, b for them; columns are sometimes addressed as vectors in the narrower sense (coming from the "primal space" rather than its dual), and we use h, k for them. Since both of these matrices act naturally to the left as well as to the right, we identify the actions with the corresponding matrices. Thus a row is seen as maps $\mathbb{R} \to \mathbb{R}_n$ and $\mathbb{R}^n \to \mathbb{R}$, a column as maps $\mathbb{R} \to \mathbb{R}^n$ and $\mathbb{R}_n \to \mathbb{R}$. We observe that $(\mathbb{R}_n, (|), \mathbb{R}^n)$ form a dual pair under the canonical scalar product $(|)\colon \mathbb{R}_n \times \mathbb{R}^n \to \mathbb{R}$ given by the matrix product $(a|h) = a \cdot h$. Furthermore, choosing a dual basis pair in a dual pair of vector spaces corresponds to an isomorphism between this dual pair and $(\mathbb{R}_n, (|), \mathbb{R}^n)$.

Following [37[268]], we call (β^*, β_*) a *dual basis pair* for $(V, \langle|\rangle, W)$ if $\beta^* = (\beta^1, \ldots, \beta^n)$ is a basis for V and $\beta_* = (\beta_1, \ldots, \beta_n)$ a basis for W such that $\langle \beta^i|\beta_j\rangle = \delta^i_j$ for all $i, j \in \{1, \ldots, n\}$. Here and elsewhere in these lecture notes we use δ^i_j for the Kronecker symbol and the following convention of placing the indices: The basis vectors from the "primal" space V are indexed above, the ones from the "dual" space W below; for the corresponding components $v = v_i\beta^i \in V$ and $w = w^i\beta_i \in W$, it is the other way round—such that we can employ the Einstein summation convention—summation over diagonal index pairs is implicit. If we consider a vector space V as a dual pair P_V, we can identify the bases β_* of V with the dual basis pairs (β^*, β_*) of P_V.

Writing $\delta^1, \ldots, \delta^n$ for the canonical basis in $\mathbb{R}_n$ and $\delta_1, \ldots, \delta_n$ for the one in $\mathbb{R}^n$, we obtain a dual basis pair for $(\mathbb{R}_n, (|), \mathbb{R}^n)$. With respect to this canonical dual basis pair, the scalar product of $a = a_i\delta^i \in \mathbb{R}_n$ and $h = h^i\delta_i \in \mathbb{R}^n$ takes on the familiar form $(a|h) = a_ih^i$. In the affine space $\mathbb{R}^n$, the basis columns appear as coordinate axes $\delta_k\colon \mathbb{R} \to \mathbb{R}^n, t \mapsto (0, \ldots, t, \ldots, 0)$ and the basis rows as coordinate projections $\delta^k\colon \mathbb{R}^n \to \mathbb{R}, (x_1, \ldots, x_n) \mapsto x_k$. Note that we can also view $\delta_k(t) = \delta_k t$ and $\delta^k(x) = \delta^k x$ as matrix products. A dual basis pair (β^*, β_*) in a dual pair of vector spaces $(V, \langle|\rangle, W)$ allows to extract the *components* with respect to the corresponding basis, so $v_i = \langle v|\beta_i\rangle$ and $w^i = \langle \beta^i|w\rangle$. We collect these components into rows and columns, respectively, writing $(v|_\beta = v_i\delta^i \in \mathbb{R}_n$ and $|w)_\beta = w^i\delta_i \in \mathbb{R}^n$.

The set of all $m \times n$ matrices is denoted by $\mathbb{R}^m_n$, so a matrix $A \in \mathbb{R}^m_n$ has m rows and n columns. We write A^i_j for its (i, j)-th element, with the row index i ranging over $\{1, \ldots, m\}$ and the column index j ranging over $\{1, \ldots, n\}$. Now a linear map F between an m-dimensional vector space V and an n-dimensional vector space W first has a *coordinate* *representation* $F_{\gamma\beta}\colon \mathbb{R}^m \to \mathbb{R}^n$, with respect to chosen bases β in V and γ in W. But the linear map $F_{\gamma\beta}$ can in turn be described by matrices, most canonically by the left representation matrix $L \in \mathbb{R}^n_m$ operating as $h \mapsto Lh$ and by the right representation matrix $R \in \mathbb{R}^m_n$ operating as $a \mapsto aR$. The matrix L has as columns $F(\beta_1), \ldots, F(\beta_n)$ expanded with respect to $\gamma_1, \ldots, \gamma_n$; the matrix R has as rows $F(\gamma^1), \ldots, F(\gamma^n)$ expanded with respect to $\beta^1, \ldots, \beta^n$. (The transformation $F \mapsto F_{\beta\alpha}$ is a functor from the category of vector spaces to its skeleton. The representation matrices L and R correspond to functors respectively from the skeleton and its opposite to the category of matrices.)

As an *example*, consider $\mathbb{C}$ as a two-dimensional real vector space V with the linear map $F\colon V \to V$ given by $z \mapsto iz$. Obviously F is a counterclockwise rotation by a right angle, whose coordinate representation is the map $F_{\beta\beta}\colon \mathbb{R}^2 \to \mathbb{R}^2$ given by $(x, y) \mapsto (-y, x)$, with respect to the canonical basis $\beta = (1, i)$. Its left and right matrix representations are

$$\begin{pmatrix} x \\ y \end{pmatrix} \mapsto \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix} \quad \text{and} \quad \begin{pmatrix} x & y \end{pmatrix} \mapsto \begin{pmatrix} x & y \end{pmatrix} \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix},$$

respectively. As one can see from this example, the left and right representation matrices are related to each other by transposition.

Since we shall normally work with dual pairs of vector spaces, we will adhere to the following *convention for representation matrices*: A coordinate representation is identified with its representation matrix—the left one operating on columns, the right one on rows. So if $F\colon V \to W$ is a linear map, we have $|Fv)_\gamma = F_{\gamma\beta}|v)_\beta$ for all $v \in V$, where β is a basis for V and γ a basis for W.

The *dual* of a linear map $F\colon V \to W$ is given by the pullback $F^*\colon W^* \to V^*$, defined by $a \mapsto a \circ F$. The matrix representation of F^* is reverse to that of F. In other words, if F had $A \in \mathbb{R}^m_n$ as its left representation matrix and hence $A^\top \in \mathbb{R}^m_n$ as its right representation matrix, its dual F^* has A as its right representation matrix and $A^\top$ as its left representation matrix (using in V^* and W^* the dual bases of the ones respectively used in V and W). As above we have now $(F^*d|_\beta = (d|_\gamma F_{\gamma\beta}$ for all $d \in V^*$, where β is again a basis for V and γ a basis for W.

0.2 Change of Bases

Changing from a basis $\beta = (\beta_1, \ldots, \beta_n)$ of a vector space V to a new basis $\bar\beta = (\bar\beta_1, \ldots, \bar\beta_n)$ can be expressed by the *transition matrix* $T \in \mathbb{R}^n_n$ from β to $\bar\beta$, formed by writing in columns the old basis vectors with respect to the new basis such that

$$\beta_j = \bar\beta_i T^i_j. \tag{1}$$

This means $T = 1_{\bar\beta\beta}$ is the (left) representation matrix of 1_V with respect to the bases β and $\bar\beta$.

A vector $v \in V$ may be expanded as $v = v^i\beta_i$ in the old basis and as $v = \bar v^i\bar\beta_i$ in the new basis. Hence we obtain

$$\bar v^i = T^i_j v^j \quad \text{or} \quad |v)_{\bar\beta} = T\,|v)_\beta \tag{2}$$

for the corresponding *component transformation*. Note that the matrix expressing $\bar{\beta}$ in terms of β is contragredient (i.e. inverse transpose) to the one expressing $|v)_{\bar{\beta}} \in \mathbb{R}^n$ in terms of $|v)_\beta \in \mathbb{R}^n$, the inverse coming from exchanging left/right sides of the equation and the transpose from exchanging row/column indices of the summation.

Now consider V as part of the *dual pair* (V^*, V). Then the above bases in V are written as $\beta_* = (\beta_1, \dots, \beta_n)$ and $\bar{\beta}_* = (\bar{\beta}_1, \dots, \bar{\beta}_n)$, and they come along with their dual bases $\beta^* = (\beta^1, \dots, \beta^n)$ and $\bar{\beta}^* = (\bar{\beta}^1, \dots, \bar{\beta}^n)$ in V^*. Since the component transformation law $\bar{v}^i = T^i_j v^j$ holds for all $v \in V$ and since $v^i = \beta^i(v)$ and $\bar{v}^i = \bar{\beta}^i(v)$, we obtain

$$\bar{\beta}^i = T^i_j \beta^j, \tag{3}$$

so the transition matrix from β^* to $\bar{\beta}^*$ in V^* is also contragredient to T. In analogy to before, this yields the component transformation law

$$d_j = \bar{d}_i T^i_j \quad \text{or} \quad (d|_\beta = (d|_{\bar{\beta}} T \tag{4}$$

for linear forms $d \in V^*$, whose matrix is this time just T itself.

The set $\mathfrak{Bas}(V)$ of bases of a vector space V form a so-called *torsor*. Since the theory of torsors—despite its simplicity—is not usually treated in standard texts of linear algebra, let us give a brief sketch of it at this point. A torsor over a group G, briefly called a G-torsor, is a free and transitive action $G \times X \to X$. We may regard it as an isomorphic view of G that does not have a distinguished neutral element ("a group that has forgotten which was its neutral element"). Consequently, one cannot "multiply" the elements of X, but one may "divide" them: For every $x, y \in X$ there is a unique $g =: x/y$ such that $x = g \cdot y$.

If X and Y are both torsors over the same group G, a map $h\colon X \to Y$ is called *equivariant* if $h(g \cdot x) = g \cdot h(x)$ for all $g \in G$ and $x \in X$. If one endows the G-actions with this notion of morphism, they form a category with the G-torsors being a full subcategory. Note also that we have only considered left actions so far, but everything can be extended to right actions as well: If both G and H act on the right, the condition for h to be equivariant is $h(x \cdot g) = h(x) \cdot g$. If G acts on the right but H on the left, one must require $h(x \cdot g) = g^{-1} \cdot h(x)$; and if G acts on the left but H on the right, we have accordingly $h(g \cdot x) = h(x) \cdot g^{-1}$.

Given a G_1-torsor X_1 and a G_2-torsor X_2, we can form the *torsor product*: We regard $X_1 \times X_2$ as a torsor over $G_1 \times G_2$ via the natural action $(g_1, g_2) \cdot (x_1, x_1) = (g_1 \cdot x_1, g_2 \cdot x_2)$. The division is now given by $(x_1, x_2)/(y_1, y_2) = (x_1/y_1, x_2/y_2)$.

Coming back to the *bases* of an n-dimensional vector space V, we see that $\mathfrak{Bas}(V)$ is a torsor over $GL_n(\mathbb{R})$. The action of a transition matrix $T \in GL_n(\mathbb{R})$ on a basis $\beta \in \mathfrak{Bas}(V)$ is given by $T \cdot \beta = \bar{\beta}$, where $\bar{\beta} \in \mathfrak{Bas}(V)$ is defined as in (1). In other words, the transition matrix from β to $\bar{\beta}$ is the quotient $\bar{\beta}/\beta$.

Every vector $v \in V$ induces a map $|v)\colon \mathfrak{Bas}(V) \to \mathbb{R}^n$ that associates to a given basis β its *components* $|v)_\beta$ with respect to β. The transformation law (2) can be written as

$$|v)_{T \cdot \beta} = T \cdot |v)_\beta \tag{5}$$

if we introduce the usual action of $GL_n(\mathbb{R})$ on the *component array* $\mathbb{R}^n$ as a matrix product (matrix times column). This means that $|v)$ is an equivariant map between the $GL_n(\mathbb{R})$-torsors $\mathfrak{Bas}(V)$ and $\mathbb{R}^n$.

One may also turn things upside-down and introduce n-dimensional vectors as equivariant maps. More precisely, choose an arbitrary torsor B over $GL_n(\mathbb{R})$, to be viewed as "abstract bases" or "basis labels". Then a *contravariant component vector* is defined as an equivariant map $v\colon B \to \mathbb{R}^n$. Transferring the linear structure from the tensor array $\mathbb{R}^n$, it is easy to see that the collection V of all component maps forms a vector space, which we might call the *contravariant component space*. Fixing any $\beta \in B$, it is clear from (5) that each component vector $v \in V$ is determined by its value $v(\beta) \in \mathbb{R}^n$; conversely, each column $h \in \mathbb{R}^n$ corresponds to the unique component vector v defined by $v(\beta) = h$. Thus every $\beta \in B$ induces a bijection $v \mapsto v(\beta)$ between V and $\mathbb{R}^n$, which is in fact a linear isomorphism $V \cong \mathbb{R}^n$. This implies in particular that $\dim V = n$.

Note also that every *basis* $(b_1, \ldots, b_n)$ of V corresponds to a unique "abstract basis" $\beta \in B$ in the following way: Choosing an arbitrary $\bar\beta \in B$, we obtain a basis $(b_1\,\bar\beta, \ldots, b_n\,\bar\beta)$ of $\mathbb{R}^n$. Letting T be the transition matrix from the canonical basis $(\delta_1, \ldots, \delta_n)$ to the basis $(b_1\,\bar\beta, \ldots, b_n\,\bar\beta)$ and setting $\beta = T \cdot \bar\beta$, we obtain $b_1(\beta) = \delta_1, \ldots, b_n(\beta) = \delta_n$. Identifying the concrete and abstract bases of V in this manner, we see that vector spaces and component spaces are in bijective correspondence with each other.

It becomes more interesting when we consider *dual spaces*. Just as $\mathbb{R}^n$ is a left torsor over $GL_n(\mathbb{R})$, its dual $\mathbb{R}_n$ is a right torsor over $GL_n(\mathbb{R})$. Every linear form $d \in V^*$ now induces a map $(d|\colon \mathfrak{Bas}(V) \to \mathbb{R}_n$ that assigns to d the row vector $(d|_\beta$ of components with respect to the unique V^*-basis dual to the V-basis β. The corresponding component transformation law (4) again says that $(d|$ is equivariant, this time meaning

$$(d|_{T\cdot\beta} = (d|_\beta \cdot T^{-1} \tag{6}$$

for all $T \in GL_n(\mathbb{R})$ and $\beta \in \mathfrak{Bas}(V)$.

If we view (V^*, V) as a dual pair of vector spaces with $\mathfrak{Bas}(V)$ consisting of dual basis pairs (β^*, β_*), a vector $v \in V$ induces an equivariant map $|v)\colon \mathfrak{Bas}(V) \to \mathbb{R}^n$ and a linear form $d \in V^*$ an equivariant map $(d|\colon \mathfrak{Bas}(V) \to \mathbb{R}_n$ as explained above. This can be done abstractly: If B is an arbitrary $GL_n(\mathbb{R})$-torsor, a *contravariant component vector* is again an equivariant map from B to $\mathbb{R}^n$ while a *covariant component vector* is an equivariant map from B to $\mathbb{R}_n$. As before, one may build the vector spaces V and $\bar V$ of contra- and covariant component vectors, respectively. Moreover, one can now introduce the bilinear form $\langle|\rangle\colon \bar V \times V \to \mathbb{R}$ by setting $\langle d|v\rangle = (a|h)$ where $a = (d|_\beta$ and $h = |v)_\beta$ for an arbitrary "abstract dual basis pair" $\beta \in B$; note that $\langle|\rangle$ is well-defined due to the transformation laws (5) and (6). Thus we obtain a dual pair of vector spaces $(\bar V, (|), V)$, whose dual basis pairs can again be identified with the "abstract dual basis pairs" in B. Under this identification, we obtain a bijective correspondence between dual pairs of vector spaces and co-/contravariant component spaces.

A final *remark on co- and contravariance* (to be investigated in a more general context in Chapter 2). For a single vector space, it does not make sense to speak of co- and contravariance, meaning to distinguish linear forms from vectors: Every vector space (recall that we always assume finite dimension!) is the dual of another space—namely of its dual space. Accordingly, an equivariant map from a set of "abstract bases" B to any n-dimensional vector space C (be it $\mathbb{R}_n$ or $\mathbb{R}^n$) is as good as any other—no matter whether we consider left or right actions on C: Writing "linear forms" as rows and "vectors" as columns is just an arbitrary (though useful) convention. The same is true about the action on B: We may write it on either side, depending on how we define the notion of transition matrix.

The crucial point in distinguishing co- from contravariance is to have a *dual pairing* between *two* vector spaces V and $\bar{V}$. Looking back at its definition, we see that it is not essential from which side $GL_n(\mathbb{R})$ acts on either B or C. But it is important that both the covariant and the contravariant maps are defined on the same torsor B with a fixed (left or right) action of $GL_n(\mathbb{R})$ and that its actions on $\mathbb{R}^n$ and $\mathbb{R}_n$ are from opposite sides (such that they cancel out). In such a case only, we may refer to V as the "primal space" (holding the "vectors" with their "contravariant components") and to $\bar{V}$ as the "dual space" (holding the "linear forms" with their "covariant components").

0.3 Maps in Topological Spaces

Let X and Y be sets. Then we call $\Phi \subseteq X \times Y$ a *graph* from X to Y if for every $x \in X$ there is a unique $y \in Y$ such that $(x, y) \in \Phi$. A *map* from X to Y is a triple $F = (X, Y, \Phi)$ such that Φ is a graph from X to Y; we write this as $F\colon X \to Y$, calling X the domain and Y the codomain of F. (If one does care for symmetry, one may as well drop X from the triple F; it can be obtained from Φ by collecting its left components.)

For reasons of clarity, we emphasize two points in the above definition: First, a map $F\colon X \to Y$ is not simply its graph; it "knows" its codomain (and of course also its domain). We write $\operatorname{dom} F$ and $\operatorname{cod} F$ for the *domain and codomain* of F, respectively. For example, if $X \subseteq Y$, the inclusion $\iota\colon X \to Y$ is to be distinguished from the identity $1_X\colon X \to X$; the graphs are in both cases $X \times X$, but we have $\operatorname{cod} \iota = Y$ whereas $\operatorname{cod} 1_X = X$. Second, we must be careful not to confuse the codomain Y of a map $F\colon X \to Y$ with its *image* $F(X)$; they coincide iff F is surjective.

If X and Y are topological spaces, a map $F\colon X \to Y$ is called open/continuous if the image/preimage of any open set in Y is open in Y. A bijection F that is both open and continuous is called a *homeomorphism*; this is equivalent to requiring that F and F^{-1} are both continuous (or both open). If $U \subseteq X$ and $V \subseteq Y$, we call a map $f\colon U \to V$ a C^0 *homeomorphism* if U is open in X and V open in Y and if F is a homeomorphism between the topological spaces U und V, considered with their induced topology. (We will extend this notion naturally to C^r diffeomorphisms between subsets of manifolds X and Y.)

A map $F\colon X \to Y$ will be called a *curve* in Y if X is $\mathbb{R}$ and a *function* on X if Y is $\mathbb{R}$. (The distinction between maps and functions is recommended in [2^{25}]. In the case of curves, Y will usually carry at least a topological structure, and we will require F to be continuous, differentiable or—most likely—even smooth. More about this later.) A curve or function is called *proper* iff it is not constant. Note that curves are normally defined on closed intervals like $[0, 1]$. For purposes of differentiation, however, open intervals like $]0, 1[$ are more appropriate; since they are diffeomorphic (as defined below) to the real line, a curve may also be defined on open intervals via reparametrizations.

0.4 Maps in Vector Spaces

Now let $U \subseteq \mathbb{R}^m$ and $V \subseteq \mathbb{R}^n$. Then the set of all *continuous* maps $F\colon U \to V$ is denoted by $C(U, V)$, the subset of all r-times *continuously differentiable* maps by $C^r(U, V)$, the subset of all *smooth* maps (meaning the union of all the sets $C^r(U, V)$ just introduced) by $C^\infty(U, V)$, and the subset of all *analytic* maps by $C^\omega(U, V)$. For reasons of convenience, one sets $C^0(U, V) = C(U, V)$, and one usually drops Y if it is the real line $\mathbb{R}$. Note that all these sets are real vector spaces (for $Y = \mathbb{R}$ even real algebras) and that U should be open when $r > 0$ in order for the definitions to make sense.

A bijection $F\colon U \to V$ is called a C^r *isomorphism* between U and V if U and V are open and if both F and F^{-1} are C^r maps. For $r > 0$ we prefer to use the speak of a C^r *diffeomorphism* and for $r = 0$ of a C^0 *homeomorphism*; the latter notion is clearly a special case of the definition given earlier. If r is not mentioned, it is generally understood as ∞. The derivative of a map $F\colon U \to V$ between open sets $U \subseteq \mathbb{R}^m$ and $V \subseteq \mathbb{R}^n$ at a point $x \in U$ is given by the Jacobian matrix $F'(x)$ if we use the standard bases (see Subsection 1.1.1).

The case of curves and functions deserves special mention. For a curve $c\colon \mathbb{R} \to \mathbb{R}^n$, its *tangent vector* (also known as the "velocity vector") for the parameter $t \in \mathbb{R}$ is to be written as a column $c'(t) \in \mathbb{R}^n$ and viewed as the linear curve $\mathbb{R} \to \mathbb{R}^n$ representing the tangent line (after the point $c(t)$ is translated to the origin). The analogous convention for a function $f\colon \mathbb{R}^n \to \mathbb{R}$ requires its *cotangent vector* (also known as the "gradient form") at the point $x \in \mathbb{R}^n$ to be written as a row $f'(x) \in \mathbb{R}_n$ and understood as the linear function $\mathbb{R}^n \to \mathbb{R}$ describing its tangent hyperplane (after the constant $f(x)$ is normalized to zero). For linear curves $h \in \mathbb{R}^n$ and linear functions $a \in \mathbb{R}_n$, these identifications lead to the uniformity relations $h'(t) = h$ for all $t \in \mathbb{R}$ and $a'(x) = a$ for all $x \in \mathbb{R}^n$.

Chapter 1

The Category of Manifolds

1.1 Embedded Manifolds

1.1.1 Differential Calculus in Vector Spaces

We first recall some notions and results from differential calculus. Let U be open in $\mathbb{R}^m$ and let $x \in U$. The linear approximation of a differentiable map $f\colon U \to \mathbb{R}^n$ locally at x is called the *differential* $d_x f\colon \mathbb{R}^m \to \mathbb{R}^n$ of f at x. It is the linear map uniquely defined by the property that

$$f(x+h) = f(x) + (d_x f)h + \|h\|\psi(h)$$

for a map ψ with $\lim_{h\to 0}\psi(h) = 0$ and h sufficiently small. For the standard basis, the differential is given by the *Jacobian matrix*

$$f'(x) = \left(\frac{\partial f_i}{\partial x_j}(x)\right)_{\substack{i=1,\ldots,n\\ j=1,\ldots,m}}$$

where $f_1, \ldots, f_n$ are the components of the map f.

Writing $y = f(x)$, the *chain rule*

$$d_x(g \circ f) = d_y g \circ d_x f$$

states that the linear approximation of the composite of two function is the composition of the linear approximations. Hence the Jacobian matrix of a composite function

$$(g \circ f)'(x) = g'(f(x))\, f'(x).$$

is the product of the Jacobian matrices of the two functions.

Let $V \subseteq \mathbb{R}^n$ be open and $f \in C^1(U, V)$. Suppose that there exists an inverse $g\colon V \to U$, which is differentiable. Applying the chain rule to

$$g \circ f = 1_U \quad \text{and} \quad f \circ g = 1_V,$$

we see that

$$d_y g \circ d_x f = 1_{\mathbb{R}^m} \quad \text{and} \quad d_x f \circ d_y g = 1_{\mathbb{R}^n}. \tag{1.1}$$

Hence $m = n$ since $d_x f$ is an isomorphism between vector spaces. Moreover, the Jacobian matrix of the inverse function is given by the inverse of the Jacobian, $g'(y) = f'(x)^{-1}$. Therefore g is also continuously differentiable and f is a C^1 *diffeomorphism.* So if a continuously differentiable function f has a differentiable inverse, the differential $d_x f$ is an isomorphism so that the Jacobian matrix $f'(x)$ is regular for all $x \in U$. The inverse mapping theorem, [9[36]] or [30[361]], tells us what happens if the differential is invertible at one point.

Now let $f \in C^r(U, \mathbb{R}^n)$. We call f a *local C^r isomorphism* at x if there exists an open neighborhood $U_1 \subseteq U$ of x such that $f|_{U_1}$ is a C^r isomorphism, which implies in particular that $f(U_1)$ is an open neighborhood of $f(x)$.

1.1 Theorem (Inverse Mapping Theorem) *Let $f \in C^r(U, \mathbb{R}^n)$ with $r \geq 1$ and $U \subseteq \mathbb{R}^n$ open, and choose a point $x \in U$. Then $d_x f$ is an isomorphism iff f is a local C^r diffeomorphism at x.*

We call a map $f \in C^r(U, \mathbb{R}^n)$ a *local C^r isomorphism* if it is one at all x in U. Note that a bijective local C^r isomorphism is a C^r diffeomorphism. So if f is bijective and its differential $d_x f$ is an isomorphism for all x in U, it is a C^r diffeomorphism.

1.1.2 Manifolds in Vector Spaces

We know from our experience that sitting on a sphere, the world around us looks flat. Mathematically speaking, we could also say that, at least locally, there exists a smooth coordinate change for the three dimensional space such that the sphere becomes an open subset of the plane.

Let us make this first more precise for the sphere. The two-dimensional *sphere* in $\mathbb{R}^3$ is the set of points

$$S^2 = \{(x_1, x_2, x_3) \in \mathbb{R}^3 \mid x_1^2 + x_2^2 + x_3^2 = 1\}.$$

We define an open subset

$$U_3^+ = \{(x_1, x_2, x_3) \in \mathbb{R}^3 \mid x_1^2 + x_2^2 < 1 \wedge x_3 > 0\}$$

and the map $\Phi_3^+ \colon U_3^+ \to \mathbb{R}^3$ by

$$(x_1, x_2, x_3) \mapsto (x_1, x_2, x_3 - \sqrt{1 - x_1^2 - x_2^2}).$$

Then Φ_3^+ is obviously smooth with the smooth inverse

$$(y_1, y_2, y_3) \mapsto (y_1, y_2, y_3 + \sqrt{1 - y_1^2 - y_2^2}).$$

defined on its image

$$A_3^+ = \{(y_1, y_2, y_3) \in \mathbb{R}^3 \mid y_1^2 + y_2^2 < 1 \wedge y_3 > -1\},$$

so it is a diffeomorphism. The image of Φ_3^+ restricted to the sphere S_2 is

$$\Phi_3^+(S_2 \cap U_3^+) = A_3^+ \cap (\mathbb{R}^2 \times \{0\}) = \{(y_1, y_2, 0) \mid y_1^2 + y_2^2 < 1\},$$

which is the open disk in the y_1y_2-plane of $\mathbb{R}^3$. Note that $\Phi_3^+|_{S_2}$ is just the projection on the first two coordinates. Clearly, we can define analogously diffeomorphisms Φ_3^-, $\Phi_1^\pm$ and $\Phi_2^\pm$ such that their corresponding domains cover the sphere and their images restricted to the sphere are the open disk in $\mathbb{R}^2$. We have thus interpreted the sphere S^2 as a geometric object embedded in $\mathbb{R}^3$.

More generally, we can think of an n-dimensional *embedded submanifold* of $\mathbb{R}^m$ (briefly an embedded n-submanifold) as a subset M of $\mathbb{R}^m$ that is locally diffeomorphic to an open subset of an n-dimensional subspace of $\mathbb{R}^m$. Recall our subspace convention of regarding $\mathbb{R}^n$ as a subset of $\mathbb{R}^n \times \mathbb{R}^{m-n} = \mathbb{R}^m$ in the sense that $\mathbb{R}^n \times \{0\} \subseteq \mathbb{R}^n \times \mathbb{R}^{m-n}$.

1.2 Definition *We call a subset M of $\mathbb{R}^m$ an embedded n-dimensional C^r submanifold of $\mathbb{R}^m$ if for every point $p \in M$ there exists an open neighborhood U of p and a C^r diffeomorphism $\Phi\colon U \to A \subseteq \mathbb{R}^m$ such that $\Phi(U \cap M) = A \cap \mathbb{R}^n$.*

Let M be a an embedded n-dimensional C^r submanifold of $\mathbb{R}^m$ and p a point on M. Let U be an open neighborhood of p and $\Phi\colon U \to A$ a diffeomorphism as above, with components $\Phi_1, \dots, \Phi_m$. We call Φ an *ambient chart* of M around the point p.

Locally, a manifold can be interpreted as the *zero set* of $k = m - n$ functions. We call $k = m - n$ the *codimension* of M. We can just take $g_i = \Phi_{i+n}$ for $i = 1, \dots, k$. Then $p \in M \cap U$ iff $g_i(p) = 0$ for $i = 1, \dots, k$. If we consider the map $g\colon U \to \mathbb{R}^k$ with the components $g_1, \dots, g_k$, we can can write this condition more compactly as

$$M \cap U = g^{-1}(0) \cap U.$$

Suppose the manifold is at least C^1. Since Φ is a diffeomorphism, we know that the Jacobian matrix of Φ is regular at p. Therefore the rank of the Jacobian matrix $g'(p)$ is k, and the *differential* d_pg is *surjective*. Using the Inverse Mapping Theorem 1.1, we will see that this gives us a condition to define manifolds as zero sets of functions.

Let U be open in $\mathbb{R}^m$ and $g\colon U \to \mathbb{R}^k$ be a differentiable map with $k \le m$. We call $x \in U$ a *regular point* of g if the differential d_xg at x is surjective so that the rank of the Jacobian matrix of $g'(x)$ is k; otherwise we call it a *critical* or *singular* point of g. We call $c \in \mathbb{R}^k$ a *regular value* of g if every $x \in g^{-1}(c)$ is a regular point of g. Note that this is in particular the case when $g^{-1}(c)$ is the empty set.

As an example consider the function $g\colon \mathbb{R}^3 \to \mathbb{R}, x \mapsto x_1^2 + x_2^2 + x_3^3$. The gradient $g'(x) = (2x_1, 2x_2, 2x_3)$ is nonzero and the differential d_xg surjective iff $x \ne 0$. So any nonzero vector is a regular point of g and the zero vector a singular point. Any $c \ne 0$ is a regular value, since $g^{-1}(0) = \emptyset$ for $c < 0$ and for any vector $x \in g^{-1}(c)$ with $c > 0$ at least one component $x_i \ne 0$, and 0 is a singular value.

1.3 Proposition *Let U be open in $\mathbb{R}^m$ and $g \in C^r(U, \mathbb{R}^k)$ with $r \ge 1$ and $k \le m$. Let $\tilde{x}$ be a regular point of g with value $c = g(\tilde{x}) \in \mathbb{R}^k$. Then locally at $\tilde{x}$, the level set $g^{-1}(c)$ is an embedded C^r submanifold of $\mathbb{R}^m$ of codimension k.*

Proof. After a translation by c we can assume that $g(\tilde{x}) = 0$. Observe that a translation does not change the differential. Since $d_{\tilde{x}}g$ is surjective we can assume, after a suitable reordering of coordinates, that the Jacobian matrix has the form

$$g'(\tilde{x}) = \begin{pmatrix} g_1'(\tilde{x}) & g_2'(\tilde{x}) \end{pmatrix},$$

where $g_2'(\tilde{x})$ is a regular $k \times k$ matrix. Let $n = m - k$. We consider the map $\Phi\colon U \to \mathbb{R}^m$ with

$$(x_1, \ldots, x_m) \mapsto (x_1, \ldots, x_n, g(x_1, \ldots, x_m)).$$

The Jacobian matrix of Φ at $\tilde{x}$ is

$$\Phi'(\tilde{x}) = \begin{pmatrix} I_n & 0 \\ g_1'(\tilde{x}) & g_2'(\tilde{x}) \end{pmatrix}.$$

It is a regular matrix since $g_2'(\tilde{x})$ is regular and so Φ is a local C^r isomorphism at $\tilde{x}$ by the Inverse Mapping Theorem 1.1. Therefore there exists an open neighborhood $U_1 \subseteq U$ of $\tilde{x}$ such that $A_1 = \Phi(U_1) \subseteq \mathbb{R}^k$ is open and $\Phi|_{U_1}$ is a C^r diffeomorphism from U_1 to A_1. Since

$$\Phi(U_1 \cap g^{-1}(0)) = A_1 \cap \mathbb{R}^n,$$

we see that $U_1 \cap g^{-1}(0)$ is indeed an embedded n-dimensional submanifold of $\mathbb{R}^m$. □

1.4 Corollary *If $c \in \mathbb{R}^k$ is a regular value of g, the level set $g^{-1}(c)$ is an embedded C^r submanifold of $\mathbb{R}^m$ of codimension k.*

Using this criterion, we can prove again that the sphere is a two-dimensional C^∞ submanifold of $\mathbb{R}^3$. It is the set $S^2 = g^{-1}(1)$ with $g(x) = x_1^2 + x_2^2 + x_3^3$, and we have seen before that 1 is a regular value of g. More generally, we see that the sphere

$$S^n = g^{-1}(1), \quad \text{with } g(x) = x_1^2 + \cdots + x_{n+1}^2$$

is an embedded smooth n-dimensional submanifold of $\mathbb{R}^{n+1}$.

1.1.3 From Ambient to Abstract Charts

In the definition of an embedded submanifold of $\mathbb{R}^m$ we used what we called "ambient charts", which are diffeomorphisms defined on some open subset of the ambient space $\mathbb{R}^m$ that map the manifold to a linear subspace. What happens if we *forget about the ambient space* and consider the "chart" obtained by restricting the ambient chart to the manifold and then projecting to the first n coordinates? From the definition we know that such a chart must be a bijection from a subset of the manifold to an open subset of $\mathbb{R}^n$. It will turn out that each chart is a homeomorphism, and we would actually like to say that it must also be a diffeomorphism—but this does not make sense since its domain fails to be open! Hence we will have to circumscribe this property in a roundabout way.

Let M be an embedded n-submanifold of $\mathbb{R}^m$ and $\Phi\colon U \to A$ be an ambient chart. Let

$$\hat{U} = U \cap M \quad \text{and} \quad \hat{A} = \pi(A \cap \mathbb{R}^n)$$

where $\pi\colon \mathbb{R}^m \to \mathbb{R}^n$ denotes the projection $\pi(x_1, \dots, x_m) = (x_1, \dots, x_n)$. Then $\hat{A}$ is open in $\mathbb{R}^n$ since every projection is an open map. We call the map

$$\varphi = \pi \circ \Phi|_{\hat{U}}\colon \hat{U} \to \hat{A}$$

an abstract chart, or briefly a *chart*. By Definition 1.2, every chart $\varphi\colon \hat{U} \to \hat{A}$ is then a bijection between a subset $\hat{U}$ of M and an open subset $\hat{A}$ of $\mathbb{R}^n$. (In our example of the sphere the charts were projections on two coordinates.) Since an embedded submanifold M is a subset of $\mathbb{R}^m$ it is naturally a topological space with the induced topology.

1.5 Proposition *Every chart is a C^0 homeomorphism.*

Proof. By definition of the induced topology, $\hat{U} = U \cap M$ is an open subset of M. The restriction of a continuous map to a subset with the induced topology is again continuous, and every projection is continuous. So a chart is a continuous bijective map and it remains to show that $\varphi = \pi \circ \Phi|_{\hat{U}}$ is open. Let $\hat{V} \subseteq \hat{U}$ be open, so that $\hat{V} = V \cap U \cap M$ with V open in $\mathbb{R}^m$. Since Φ is a diffeomorphism, $\Phi(V)$ is open in $\mathbb{R}^m$ and therefore also

$$\Phi(\hat{V}) = \Phi(V \cap U \cap M) = \Phi(V) \cap \Phi(U \cap M) = \Phi(V) \cap (A \cap \mathbb{R}^n)$$

by Definition 1.2. Hence the projection $\pi(\Phi(V) \cap A) = \varphi(\hat{V})$ is open in $\mathbb{R}^n$. $\square$

Now we have to face the problem of describing the differentiability of abstract charts—even though their domains fail to be open. The key to solve this problem is consider the "transition" of overlapping charts rather than single charts. Let $\varphi\colon \hat{U} \to \hat{A}$ and $\psi\colon \hat{V} \to \hat{B}$ be two charts. Then $\varphi(\hat{U} \cap \hat{V})$ and $\psi(\hat{U} \cap \hat{V})$ are open in $\mathbb{R}^n$ by the previous proposition. When $\hat{U} \cap \hat{V}$ is not empty, we call the map

$$\psi\varphi^{-1}\colon \varphi(\hat{U} \cap \hat{V}) \to \psi(\hat{U} \cap \hat{V})$$

the *transition* from the chart φ to the chart ψ. Note that the transition $\psi\varphi^{-1}$ is a bijection between open subsets of $\mathbb{R}^n$, so we can ask if it is differentiable.

1.6 Proposition *Every transition $\psi\varphi^{-1}$ is a C^r diffeomorphism.*

Proof. Let $\Phi\colon U \to A$ and $\Psi\colon V \to B$ be the corresponding ambient charts. The transition of the ambient charts $\Psi\Phi^{-1}$ is then a C^r diffeomorphism from $\Phi(U \cap V)$ to $\Psi(U \cap V)$. We denote by $\iota\colon \mathbb{R}^n \to \mathbb{R}^m$ the inclusion $\iota(x_1, \dots, x_n) = (x_1, \dots, x_n, 0, \dots, 0)$. Then

$$\psi\varphi^{-1} = \pi \circ \Psi\Phi^{-1} \circ \iota$$

is C^r since both injection and projection are C^∞. We see analogously that the inverse $\varphi\psi^{-1}$ is C^r, and so the transition is a diffeomorphism. $\square$

From now on we consider only abstract charts and we will denote their domains and codomains without hats to simplify the notation. By definition we have an ambient chart around every point of an embedded submanifold. Hence there exists a family of charts $\varphi_i\colon U_i \to A_i$ such that their domains cover M, meaning $M = \bigcup_{i\in I} U_i$. We call such a family an *atlas* of M. By the previous proposition, all charts in an atlas are mutually compatible in the sense that their transitions are C^r diffeomorphisms.

Finally, we take the abstraction one step further and forget that we started with a subset of $\mathbb{R}^m$ and take instead an arbitrary set M. We consider this set together with a family of charts (bijections into open subsets of an $\mathbb{R}^n$) such that their domains cover M and that all charts are mutually compatible. This is the starting point for defining abstract manifolds in the next chapter.

1.2 Abstract Manifolds

The motivation for defining manifolds in a more abstract, invariant manner comes from the following observation: Even if we have a certain object as embedded in some $\mathbb{R}^N$, this embedding often suffers from a number of flaws: it may not be "natural", it may need the embedding dimension N to be significantly higher than the dimension n of the manifold. In fact, the famous Whitney Embedding Theorem [7[91]] tells us that one may always(!) embed an n-dimensional manifold into some $\mathbb{R}^N$ with $N \leq 2n$, but it also exposes the two flaws just mentioned: The upper bound of $2n$ for the embedding dimension may be reached (this is the case for the n-dimensional projective space) or it may at least be hard to find low dimensions N; and the embedding is not canonical, so to say adding bureaucratic ballast to the essential structure of the *manifold itself*.

So let M now be just a set and $n, r \in \mathbb{N}$. We will use A, B, C for open subsets of $\mathbb{R}^n$ and U, V, W for any subsets of M. Furthermore, k is to range over $\{1, \ldots, n\}$.

1.2.1 The Chart Topology

The crucial tools for abstracting from a surrounding vector space are the charts: they can be used for *pulling down* all kinds of objects and relations on the "abstract level" of M (usually pictured "above"), yielding numerical objects and relations on the "concrete level" of coordinates in $\mathbb{R}^n$ (usually pictured "below"). Conversely, they also allow us to *lift* objects and relations from the well-known Euclidean world of $\mathbb{R}^n$ to the more sophisticated domain of M. Such definitions of additional structures via the atlas must of course be independent of the choice of charts ("invariant under coordinate changes").

1.7 Definition *An n-dimensional* chart *is a bijection $\varphi\colon U \to A$. It is said to be* compatible *with another chart $\psi\colon V \to B$ if $\psi\varphi^{-1}$ is a C^r isomorphism between $\varphi(U \cap V)$ and $\psi(U \cap V)$. A family of mutually compatible charts is called an* atlas *if their domains cover M.*

To be more explicit, we should actually speak of charts, compatibility and atlases of *class* C^r. But r is to be fixed for the rest of this section, just like the dimension n and the base set M.

The *intuition* behind these concepts can be fetched from the one example that provided some of the standard terms: A good “atlas” partitions the Earth M into several overlapping “terrains” U, which are depicted by “charts” $\varphi\colon U \to A$ on various “pages” $A \subseteq \mathbb{R}^2$ of the atlas. So in the sense of the forthcoming definition, the Earth is indeed a two-dimensional manifold (meaning its surface is approximated by one).

Some *terminology on charts* will be useful in the future.

- Since $A \subseteq \mathbb{R}^n$, we can write the corresponding chart as $\varphi = (\varphi^1, \dots, \varphi^n)$ with n so-called *coordinate functions* $\varphi^1, \dots, \varphi^n\colon U \to \mathbb{R}$ defined via the projection as $\varphi^k = \delta^k \varphi$.

- Since every $p \in U$ has n coordinates $\varphi^1(p), \dots, \varphi^n(p)$, a chart is also known as a local coordinate system [7^{68}]. Its inverse $\varphi^{-1}\colon A \to U$ can be seen as a local *parametrization* of M.

- An element $p \in U$ is called a *point*, its image $\varphi(p) = (\varphi^1(p), \dots, \varphi^n(p)) \in A$ a *coordinate node*, consisting of the n coordinates of p.

- In [46^{129}] and [2^{52}], the domain U of a chart $\varphi\colon U \to A$ is called a coordinate neighborhood. In these notes, we will refer to U simply as a *chart domain*, while A will be called its *coordinate patch* embedded into the *coordinate space* $\mathbb{R}^n$.

- If $p \in U$, one calls $\varphi\colon U \to A$ a *chart around* the point p. If $\mathfrak{A}$ is an atlas, we write $\mathfrak{A}_p$ for the *local atlas* around p, containing all the charts $\varphi \in \mathfrak{A}$ around p.

- A chart $\varphi\colon U \to A$ around p with the special property $\varphi(p) = 0$ will be called a *chart centered at* p. We write $\mathfrak{A}_{p\bullet}$ for the centered atlas around p, consisting of all such charts.

- For centered charts, one can restrict the axes δ_k from $\mathbb{R} \to \mathbb{R}^n$ to $A_k \to A$ for some open set $A_k \subseteq \mathbb{R}^n$ and then build the so-called *coordinate curves* $\varphi_k = \varphi^{-1}\delta_k\colon A_k \to U$ through p.

- The map $\psi\varphi^{-1}$ from $\varphi(U \cap V)$ to $\psi(U \cap V)$ is variously called transition map [46^{114}], overlap map [1^{123}] or change of coordinates [7^{69}]; we will simply call it the *transition* from the chart φ to the chart ψ.

In order to define a topology on M, we apply the method of “pulling down” for the first time: In order to see whether a set is open, we probe its openness under every chart (restricting the set to the chart domain). The result will be called the *chart topology* on M.

1.8 Definition *Let $(\varphi_i \mid i \in I)$ be an atlas on M. Then we call a set $U \subseteq M$ open in M if $\varphi_i(U \cap U_i)$ is open in $\mathbb{R}^n$ for every chart φ_i with domain U_i.*

1.9 Lemma *The open sets form a topology on M.*

Proof. The topology of Definition 1.8 is just the final topology of the family $(\iota_i \circ \varphi_i^{-1} \colon A_i \to M)$, where $\iota_i \colon U_i \to M$ is the insertion of the corresponding chart domains. In order to see this, recall [6[32]] that a set U in such a final topology is open iff $(\iota_i \circ \varphi_i^{-1})^{-1}(U) = \varphi_i(U \cap U_i)$ is open in A_i and hence in $\mathbb{R}^n$; this is just what we have ein Definition 1.8. ☐

Since the charts translate between the abstract level of M and the concrete level of $\mathbb{R}^n$, we expect them to be *homeomorphisms* (they will even turn out to be diffeomorphisms once we have defined what that means). The following proposition verifies this important criterion for a reasonable notion of topology on M.

1.10 Proposition *Every chart is a C^0 homeomorphism.*

Proof. Before proceeding to the charts, let us first remark that each chart domain U_i is necessarily open since compatibility implies that $\varphi_j(U_i \cap U_j)$ is open for every chart $\varphi_j \colon U_j \to A_j$.

For showing that an arbitrary chart $\varphi_i \colon U_i \to A_i$ is a C^0 homeomorphism, we must prove that φ_i is both an open and a continuous map. Openness is easy: For an open $U \subseteq U_i$, the image $\varphi_i(U) = \varphi_i(U \cap U_i)$ is open in $\mathbb{R}^n$ by the definition of the topology on M. Note that since U_i is itself open, a subset of U_i is open in U_i iff it is open in M.

It remains to prove φ_i is continuous. For this, we choose $A \subseteq A_i$ open in $\mathbb{R}^n$ and prove $\varphi_i^{-1}(A)$ is open in M. So taking an arbitrary chart $\varphi_j \colon U_j \to A_j$, we must show that $B = \varphi_j(\varphi_i^{-1}(A) \cap U_j)$ is open. Since $\varphi_i^{-1}(A) \subseteq U_i$ and φ_i is bijective,

$$\varphi_i^{-1}(A) \cap U_j = \varphi_i^{-1}(A \cap \varphi_i(U_i \cap U_j)),$$

so $B = \varphi_j\varphi_i^{-1}(A \cap \varphi_i(U_i \cap U_j))$. Since A is open by hypothesis and $\varphi_i(U_i \cap U_j)$ by the the compatibility requirement of the atlas, $A \cap \varphi_i(U_i \cap U_j)$ is also open. But compatibility means that $\varphi_j\varphi_i^{-1}$ is diffeomorphic, and hence *a fortiori* homeomorphic, between the open sets $\varphi_i(U_i \cap U_j)$ and $\varphi_j(U_i \cap U_j)$. Therefore it maps $A \cap \varphi_i(U_i \cap U_j)$ to the open set B. The chart $\varphi_j \colon U_j \to A_j$ being arbitrary, this shows that $\varphi_i^{-1}(A)$ is open and thus concludes the proof that every chart $\varphi_i \colon U_i \to A_i$ is a C^0 homeomorphism. ☐

The chart property stated Proposition 1.10 is even *characteristic for the chart topology*, as we shall see now.

1.11 Proposition *The topology on M is uniquely characterized by the requirement that each chart is a C^0 homeomorphism.*

Proof. Note that the characteristic property of making every chart a C^0 homeomorphism includes that the chart domains are open sets. Now consider any topology fulfilling the characteristic property, and let $\mathcal{M}$ be the collection of its open sets.

We prove first that every set $U \in \mathcal{M}$ is open. For any chart $\varphi_i \colon U_i \to A_i$, we have $U_i \in \mathcal{M}$, hence also $U \cap U_i \in \mathcal{M}$. By its characteristic property, φ_i maps sets in $\mathcal{M}$ to open sets in $\mathbb{R}^n$, so each $\varphi_i(U \cap U_i)$ is open, which means that U is indeed open.

For the other direction, we have to prove $U \in \mathcal{M}$ for every open set $U \subseteq M$, so we assume $\varphi_i(U \cap U_i)$ is open in $\mathbb{R}^n$ for every chart $\varphi_i \colon U_i \to A_i$. Since the chart domains cover M, we have $U = \bigcup_{i \in I}(U_i \cap U)$, and it suffices to prove $U_i \cap U \in \mathcal{M}$ for each $i \in I$. But every $\varphi_i(U_i \cap U)$ is open, so φ_i^{-1} maps it to $U_i \cap U = \varphi_i^{-1}\varphi_i(U_i \cap U) \in \mathcal{M}$, due to its characteristic property. □

In order to check whether a set is open, we have to probe its images under all charts. Intuition tells us that one could be *more economic* than this, using only those charts whose domains touch the set in question. The following lemma confirms this expectation.

1.12 Lemma *Consider an atlas $(\varphi_i \mid i \in I)$ on M with charts $\varphi_i \colon U_i \to A_i$ and a set $U \subseteq M$ with $U = \bigcup_{j \in J} U_j$ for some $J \subseteq I$. Then U is open in M iff $\varphi_j(U \cap U_j)$ is open in $\mathbb{R}^n$ for every $j \in J$.*

Proof. The condition in the lemma is clearly necessary, so assume $U \subseteq M$ is such that $\varphi_j(U \cap U_j)$ is open in $\mathbb{R}^n$ for every $j \in J$. In order to prove that U is open in M, we have to take an arbitrary chart $\varphi_i \colon U_i \to A_i$ and show that $B_i = \varphi_i(U \cap U_i)$ is open in $\mathbb{R}^n$. If $B_i = \emptyset$, we are done; so assume now $B_i \neq \emptyset$. Then it suffices to prove that an arbitrary point $x \in B_i$ is an interior point.

Its preimage $p = \varphi_i^{-1}(x)$ lies in $U \cap U_i \subseteq U$, so there is a $j \in J$ with $p \in U_j$. Since the charts φ_j and φ_i are compatible, the coordinate patches $\varphi_i(U_i \cap U_j)$ and $\varphi_j(U_i \cap U_j)$ are open in $\mathbb{R}^n$, while $\varphi_j(U \cap U_j)$ is open by hypothesis. The chart φ_j being bijective, we see that

$$\varphi_j(U \cap U_i \cap U_j) = \varphi_j(U_i \cap U_j) \cap \varphi_j(U \cap U_j)$$

is then also open.

The transition $\varphi_i\varphi_j^{-1}$ is a diffeomorphism between $\varphi_j(U_i \cap U_j)$ and $\varphi_i(U_i \cap U_j)$, so it maps the open set $\varphi_j(U \cap U_i \cap U_j)$ into the open set $C_i = \varphi_i(U \cap U_i \cap U_j)$. Finally, we observe that $p \in U \cap U_i \cap U_j$, so $x \in C_i \subseteq B_i$ is indeed an interior point of B_i. □

The criterion of Lemma 1.12 could also be used for a definition of the topology on M, but we think that Definition 1.8 is of a more uniform character (given as the final topology of the parameterizations).

1.2.2 Differentiable Structures

An atlas is a specific way of "charting" the set M, which means mapping it somehow locally into $\mathbb{R}^n$ such that every point is reached by a chart. Usually there are many ways of charting the set M; as one sees already from using various coordinate systems in the plane (like orthogonal, oblique, polar coordinates). So an atlas is far from being unique on X; there are typically *many alternative coordinate systems* that one can use.

1.13 Definition *A chart* $\varphi\colon U \to A$ *is called* admissible *for an atlas* $(\varphi_i \mid i \in I)$ *if* φ *is compatible with every* φ_i.

For example, we may always replace a chart $\varphi\colon U \to A$ around a point p by a chart $\tilde{\varphi}$ centered at p; simply choose $\tilde{\varphi} = \varphi - \varphi(p)$. We could even construct an atlas consisting of charts $\varphi\colon U \to A$ such that A is the n-ball or n-cube centered at the origin [2[55]]; let us call such a chart φ a *ball chart* or *cube chart*, respectively. In any case, $\tilde{\varphi}$ will be an admissible chart, and replacing φ by $\tilde{\varphi}$ obviously does not change the atlas in any essential way—this is what leads us the concept of atlas equivalence.

1.14 Definition *Two atlases are said to be* equivalent *if each chart of one is admissible for the other.*

One can show that this is indeed an equivalence relation on atlases. Moreover, one sees that an atlas $\mathfrak{A}$ is equivalent to an atlas $\mathfrak{B}$ iff $\mathfrak{A} \cup \mathfrak{B}$ is still an atlas. As one would expect, the *topology* on M depends only on the equivalence class of an atlas.

1.15 Proposition *Let* $\mathfrak{A}$ *be an atlas for* M. *Then every chart* $\varphi\colon U \to A$ *admissible for* $\mathfrak{A}$ *is a* C^0 *homeomorphism.*

Proof. First of all, it is clear that U is open since $\varphi_i(U \cap U_i)$ is open for every $i \in I$. This follows from the fact that each transition $\varphi\varphi_i^{-1}$ is a diffeomorphism between the open sets $\varphi_i(U_i \cap U)$ and $\varphi(U_i \cap U)$.

Now we have to prove that $\varphi\colon U \to A$ is both open and continuous. Starting with openness, let us take an open set $V \subseteq U$ and prove that $\varphi(V)$ is open in $\mathbb{R}^n$. There is a chart domain U_j around every point in V, so we obtain an open cover $V = \bigcup_{j \in J}(U_j \cap V)$ for a suitable $J \subseteq I$. Then $\varphi(V) = \bigcup_{j\in J} \varphi(U_j \cap V)$, so it suffices to prove that $\varphi(U_j \cap V)$ is open for arbitrary $j \in J$. But we know that $\varphi_j(U_j \cap V)$ is open by the definition of the chart topology, while $\varphi\varphi_j^{-1}$ is a diffeomorphism between $\varphi_j(U_j \cap U)$ and $\varphi(U_j \cap U)$. Hence $\varphi(U_j \cap V) = (\varphi \circ \varphi_j^{-1})\, \varphi_j(U_j \cap V)$ is indeed open.

As for the continuity of $\varphi\colon U \to A$, we can repeat the proof of Proposition 1.10, replacing the chart φ_i by the admissible chart φ. □

1.16 Corollary *Let* $(\bar{\varphi}_j \mid j \in J)$ *be an atlas equivalent to the atlas* $(\varphi_i \mid i \in I)$. *Then* $U \subseteq M$ *is open in* M *iff* $\bar{\varphi}_j(U \cap \bar{U}_j)$ *is open in* $\mathbb{R}^n$ *for every chart* $\bar{\varphi}_j$ *with domain* $\bar{U}_j$.

Proof. Every chart $\bar{\varphi}_j$ is admissible for the atlas $(\varphi_i \mid i \in I)$ and hence a C^0 homeomorphism by Proposition 1.15, so $\bar{\varphi}_j(U \cap \bar{U}_j)$ is open as the image of the open set $U \cap \bar{U}_j$. □

Using all admissible charts, we can characterize the *neighborhood* of a point in a very intuitive manner: A chart domain around p is clearly a neighborhood of p, and using all admissible charts turns out to be enough to build up a neighborhood base. (The chart domains of the atlas itself will typically just be enough to cover M, it will not contain arbitrarily small neighborhoods of a point.)

1.17 Proposition *The admissible chart domains around a point p form a neighborhood base around p.*

Proof. We have to show that every neighborhood of p contains an admissible chart domain V around p. A neighborhood of p is any set that contains an open set U with $p \in U$. Since the chart domains cover M, there is a U_i with $p \in U_i$. Clearly U_i is open in M, hence also $V = U \cap U_i$, which is obviously an admissible chart domain, belonging to the chart $\varphi_i|_V$. □

We have seen that there are usually many different but mutually equivalent atlases for M; the only "invariant" one is the maximal atlas: the one containing *all admissible charts*. Every equivalence class contains a unique atlas with this property of being maximal with respect to $\subseteq$.

1.18 Definition *A* differentiable structure *on M is a maximal atlas on M.*

Some books like [1^{124}] take equivalence classes instead of maximal atlases, but this is clearly the same due to the bijective correspondence mapping an equivalence class $[\mathfrak{A}]$ to its maximal atlas $\bigcup[\mathfrak{A}]$ and a maximal atlas $\mathfrak{M}$ to its equivalence class $\{\mathfrak{A} \mid \mathfrak{A} \subseteq \mathfrak{M}\}$. The atlas $\mathfrak{A}_{\max} = \bigcup[\mathfrak{A}]$ is obtained by adding to $\mathfrak{A}$ all charts admissible for $\mathfrak{A}$, so it yields the unique maximal atlas containing $\mathfrak{A}$, called the *differentiable structure generated by* $\mathfrak{A}$. This is also how one uses differentiable structures in practice: by listing (typically finitely many) charts $\mathfrak{A} = \{\varphi_1, \ldots, \varphi_n\}$ and then passing to $\mathfrak{A}_{\max}$.

The set M together with a differentiable structure is called an n-dimensional C^r *manifold* if it fulfills certain topological constraints in order to avoid pathologies; the precise definition will be given in Subsection 1.2.4, together with a brief discussion of the topological constraints to be imposed. For the moment, however, we may ignore these complications and we rather give some first examples of manifolds (as we shall confirm later). The first two are the easiest at all. If we are already in $\mathbb{R}^n$, a chart need not do anything; if we are in an "abstract vector space", we can still do with a single chart!

1.19 Example *The canonical differentiable structure on the* Euclidean space $\mathbb{R}^n$ *is induced by the atlas $\{1_{\mathbb{R}^n}\}$ and comprises all C^r isomorphisms between open sets of $\mathbb{R}^n$.*

Proof. The chart domain of the bijection $1_{\mathbb{R}^n}$ trivially covers all of $\mathbb{R}^n$, and there is no compatibility relation; hence $\{1_{\mathbb{R}^n}\}$ is indeed an atlas. An admissible chart $\varphi\colon U \to A$ between open sets $U, A \subseteq \mathbb{R}^n$ has to be compatible with $1_{\mathbb{R}^n}$, which means that $\varphi \circ 1_{\mathbb{R}^n}^{-1} = \varphi$ must be a C^r isomorphism. □

1.20 Example *Let V be an n-dimensional* vector space *and choose a basis $(b_1, \ldots, b_n)$ for V. Then the canonical differentiable structure on V is induced by the atlas $\{\varphi\}$, with the* component chart *$\varphi\colon V \to \mathbb{R}^n$ given by $v \mapsto (\lambda^1, \ldots, \lambda^n)$ for every vector $v = \lambda^i\, b_i \in V$. Bases and component charts are in bijective correspondence with each other.*

Proof. The component charts are obviously linear isomorphisms and therefore also homeomorphisms; then we know from Proposition 1.11 the chart topology coincides with the canonical topology of V, which is the only one that makes V a Hausdorff space [23[32]] [4[I.13]]. The bijection between bases and component charts is evident [27[27]]. □

A slightly more general situation is given by the embedded manifolds treated in Section 1.1. (Note also that every open subset U of $\mathbb{R}^n$ is trivially an embedded n-dimensional submanifold of $\mathbb{R}^n$; as an atlas, one may simply take 1_U. The previous Example 1.19 is the case $U = \mathbb{R}^n$.)

1.21 Example *The charts of an embedded submanifold M of $\mathbb{R}^m$ form an atlas on M whose topology coincides with the induced topology inherited from $\mathbb{R}^n$.*

Proof. By Definition 1.2, the chart domains cover M. Furthermore, we have proved in Proposition 1.5 that all charts are bijections and in Proposition 1.6 that they are mutually compatible. So the charts of M indeed make up an atlas.

Finally, we have seen in Proposition 1.5 that all charts of M are C^0 homeomorphisms when M is regarded with its induced topology as a subset of $\mathbb{R}^n$. But Proposition 1.11 tells us that there is only one topology on M that fulfills this criterion, so it must coincide with the chart topology of M. □

1.22 Example *Let*

$$\mathbb{P}^n(\mathbb{R}) = (\mathbb{R}^{n+1} \setminus \{0\})/\sim$$

with the equivalence defined as

$$a \sim b \quad \textit{iff} \quad a = \lambda b \textit{ for some } \lambda \in \mathbb{R}^*$$

be the n-dimensional real projective space. So a point in $M = \mathbb{P}^n(\mathbb{R})$ is an equivalence class $[a]$ with $(a_0, \ldots, a_n) \in \mathbb{R}^{n+1} \setminus \{0\}$, usually denoted by $(a_0 : \ldots : a_n)$. With

$$U_i = \{(a_0 : \ldots : a_n) \mid a_i \neq 0\}$$

and $A_i = \mathbb{R}^n$, define maps

$$\begin{aligned} \varphi_i \colon U_i &\to A_i \\ (a_0 : \ldots : a_n) &\mapsto \left(\frac{a_0}{a_i}, \ldots, \frac{a_{i-1}}{a_i}, \frac{a_{i+1}}{a_i}, \ldots, \frac{a_n}{a_i}\right), \end{aligned}$$

for all $i = 0, \ldots, n$. Then the maps $(\varphi_0, \ldots, \varphi_n)$ are well-defined and constitute an n-dimensional smooth atlas.

Proof. Note that the condition $a_i \neq 0$ is independent of the representative. If $[a] = [b]$, then $a_k = \lambda b_k$ and so $\frac{b_k}{b_i} = \frac{\lambda a_k}{\lambda a_i} = \frac{a_k}{a_i}$. Hence the maps φ_i are well-defined. Their domains

U_i obviously cover M and their codomains A_i are open. The maps φ_i are bijective with inverses

$$(x_1, \ldots, x_n) \mapsto (x_1 : \ldots x_i : 1 : x_{i+1} : \ldots : x_n).$$

We have to show that the charts φ_i are mutually compatible. Let $i \neq j$. Then

$$U_i \cap U_j = \{(a_0 : \ldots : a_n) \mid a_i \neq 0 \wedge a_j \neq 0\}$$

and the images $\varphi_i(U_i \cap U_j) = \{x \in \mathbb{R}^n \mid x_j \neq 0\}$ and $\varphi_j(U_i \cap U_j) = \{x \in \mathbb{R}^n \mid x_i \neq 0\}$ are open in $\mathbb{R}^n$. We assume $i < j$. Let $x = (x_1, \ldots, x_n) \in \varphi_i(U_i \cap U_j)$. Then $\varphi_j \varphi_i^{-1}(x)$ equals

$$\varphi_j(x_1 : \ldots : x_i : 1 : x_{i+1} : \ldots : x_n) = \left(\frac{x_1}{x_j}, \ldots, \frac{x_i}{x_j}, \frac{1}{x_j}, \frac{x_{j+1}}{x_j}, \ldots, \frac{x_{j-1}}{x_j}, \frac{x_{i+1}}{x_j}, \ldots, \frac{x_n}{x_j}\right),$$

so the transitions $\varphi_j \varphi_i^{-1}$ are diffeomorphisms and $\varphi_0, \ldots, \varphi_n$ constitute an atlas. □

Finally, let us also mention another simple way of getting concrete manifolds: the open subsets of a given manifold! We will see in Subsection 1.4.1 how we can view this as a special case of forming submanifolds.

1.23 Example *If M has a differentiable structure $\mathfrak{A}$, every open set U of M has a canonical differentiable structure consisting of the restrictions to U of all the charts in $\mathfrak{A}$.*

Proof. If φ is any chart of $\mathfrak{A}$, it is compatible with any other chart ψ of $\mathfrak{A}$. But then $\varphi|_U$ is clearly also compatible with every other restricted chart $\psi|_U$. Moreover, their domains cover U, so they make up an atlas $\mathfrak{A}|_U$ that provides U with a differentiable structure. □

1.2.3 Manifolds as Patchwork

The above viewpoint—charts as a tool for locally exploring an intricate structure by pulling it down to a Euclidean space—might be called the *analytic* point of view. But one may also ascribe a more *synthetic* role to the charts of an atlas [8[3]]: In order to construct the manifold, one has to put the coordinate patches (hence their name!) together and glue them on their overlaps according to their transitions [42[15]]. Readers familiar with the language of schemes will notice a striking similarity with gluing schemes as explained in [20[80]].

In detail, if $(\varphi_i \colon U_i \to A_i \mid i \in I)$ is an atlas for M and if $\tilde{M}$ is the topological sum [6[33]] of all coordinate patches A_i, construct the quotient space $\tilde{M}/{\sim}$, where $\sim$ is defined as follows (all other points being equivalent only to themselves): For $x \in \varphi_i(U_i \cap U_j)$ and $y \in \varphi_j(U_i \cap U_j)$ we put

$$(x, i) \sim (y, j) \quad \text{iff} \quad y = \varphi_j \varphi_i^{-1}(x).$$

It turns out that this quotient space is essentially a reconstruction of M.

1.24 Proposition *Using the above construction based on an atlas $(\varphi_i \colon U_i \to A_i)$, the topological space $\tilde{M}/{\sim}$ is homeomorphic to M.*

Proof. Define $f\colon M \to \tilde{M}/{\sim}$ by $f(p) = [\varphi_i(p), i]_\sim$, where $\varphi_i\colon U_i \to A_i$ is any chart around p. Note that f is well-defined: If $\varphi_j\colon U_j \to A_j$ is another chart around p, we have $(\varphi_i(p), i) \sim (\varphi_j(p), j)$. But f is both continuous and open since it is the composition of three maps having this property $[6^{50}]$: the homeomorphism φ_i, the insertion $A_i \to \tilde{M}$, and the canonical map $\tilde{M} \to \tilde{M}/{\sim}$. (For the canonical map, openness follows $[6^{52}]$ because the equivalence is induced by a group of homeomorphisms acting on $\tilde{M}/{\sim}$.)

For obtaining an inverse, we define $g\colon \tilde{M}/{\sim} \to M$ by $g([x,i]_\sim) = \varphi_i^{-1}(x)$ if $x \in A_i$ and $\varphi_i\colon U_i \to A_i$ is the corresponding chart. Again it is clear that g is well-defined: If $(x,i) \sim (y,j)$ for $y \in U_j$ with another chart $\varphi_j\colon U_j \to A_j$, we have $\varphi_i^{-1}(x) = \varphi_j^{-1}(y)$ by the definition of $\sim$. It remains to show that g is the inverse of f, which is immediate: We have

$$g(f(p)) = g([\varphi_i(p), i]_\sim) = \varphi_i^{-1}(\varphi_i(p)) = p$$

together with

$$f(g([x,i]_\sim)) = f(\varphi_i^{-1}(x)) = [\varphi_i(\varphi_i^{-1}(x)), i]_\sim = [x,i]_\sim,$$

as required. □

One may also turn this idea upside-down [39], arriving at the historical roots of manifolds: A differentiable structure can be given by specifying a family $(A_i \mid i \in I)$ of open sets $A_i \subseteq \mathbb{R}^n$ together with *gluing diffeomorphisms* $\iota_{ji}\colon A_{ij} \to A_{ji}$ connecting the patches along their "fringes" $A_{ij} \subseteq A_i$ and $A_{ji} \subseteq A_j$. The ι_{ji} are to play the role of the transitions $\varphi_j\varphi_i^{-1}$, but we are now independent of any pre-given set M whose subsets U_i and U_j would serve as the domains of the corresponding charts φ_i and φ_j. In fact, we have no charts at all; only their "transitions" are given by the gluing diffeomorphisms ι_{ji}.

In order to glue corresponding points together, we need an *equivalence* $\sim$, to be defined on the topological $\tilde{M}$ of the coordinate spaces A_i in a completely analogous manner as before, using the gluing diffeomorphisms in place of the chart transitions: For $x \in A_{ij}$ and $y \in A_{ji}$ we set $(x,i) \sim (y,j)$ iff $y = \iota_{ji}(x)$, again viewing all other points equivalent only to themselves. The requirement that $\sim$ be an equivalence relation on $\tilde{M}$ implies three constraints on the ι_{ji}: We must have $\iota_{ii} = 1_{A_{ii}}$ for reflexivity, $\iota_{ji}^{-1} = \iota_{ij}$ for symmetry, and $\iota_{kj}\iota_{ji} = \iota_{ki}$ for transitivity. A short moment's reveals that the last condition (corresponding to the "cocycle condtion" for fiber bundles, explained in Subsection 2.3.1) already entails the other two. All this amounts is to guarantee a consistent specification of the gluing process.

This yields a parametrization $A_i \to \tilde{M}/{\sim}$ defined by $x \mapsto [x,i]_\sim$ and easily seen to be injective as well as continuous and open, using a similar argument as in the proof above. Writing U_i for the image of this parametrization, its inverse is then a homeomorphism $\varphi_i\colon U_i \to A_i$, acting by $[x,i]_\sim \mapsto x$. Using the criterion of Proposition 1.11, we see now that $\tilde{M}/{\sim}$ now bears the topology induced by the atlas $(\varphi_i \mid i \in I)$ or its generated differentiable structure. Let us reassure ourselves that the chart transitions $\varphi_j\varphi_i^{-1}$ are now given by the gluing diffeomorphisms ι_{ji}. Taking $x \in A_{ij}$, we compute

$$\varphi_j\varphi_i^{-1}(x) = \varphi_j[x,i]_\sim = \varphi_j[\iota_{ji}(x), j]_\sim = \iota_{ji}(x).$$

Hence the differentiability requirements on the ι_{ji} carry over to the differentiable structure on the set $\tilde{M}/{\sim}$, which is thus an n-dimensional C^r manifold whenever the gluing diffeomorphisms are C^r and the topological conditions are fulfilled (see the next subsection). We call this the *patchwork construction* of a differentiable manifold.

1.2.4 The Definition of Manifold

Finally now the official definition of a *manifold*, which imposes two crucial conditions on the induced topology of M.

1.25 Definition *The set M together with a differentiable structure of class C^r is said to be a C^r manifold of dimension n or an n-manifold of class C^r if M is a second-countable Hausdorff space when endowed with the chart topology.*

A C^r manifold with $r = 0$ is also called a *topological manifold*. In this case, an atlas just provides local homeomorphisms into $\mathbb{R}^n$; topological spaces admitting such an atlas are also called *locally Euclidean*. In other words, a topological manifold is a locally Euclidean, second countable Hausdorff space. By Proposition 1.11 the topology of M is uniquely determined by the condition that each chart is a C^0 homeomorphism, so the maximal atlas (the "differentiable structure" of class C^0) is uniquely determined by the topology of M. Hence it does not make sense to regard the charts as an additional structure on the topological space M.

The situation is drastically different as soon as we consider $r > 0$. As before, we speak of differentiable or smooth or analytic manifolds for $r \in \mathbb{N}$ or $r = \infty$ or $r = \omega$, respectively. (Similar conventions are in force for various related concepts like atlas or differentiable structure.) For example, even $\mathbb{R}$ admits—besides the standard differentiable structure generated by the identity—various other ones, like the one generated by the chart $\varphi\colon \mathbb{R} \to \mathbb{R}$ with $\varphi(x) = x^3$. For $\mathbb{R}$ it turns out that the resulting manifolds are still all diffeomorphic (in a sense that will soon be made precise); in fact, it is known $[17^2]$ that the only one-dimensional manifolds are the line $\mathbb{R}$ and the circle S^1. In contrast—just to mention Kervair and Milnor's famous example [25]—there are 27 exotic seven-dimensional spheres (a sphere is called exotic when endowed with a differentiable structure distinct from the standard one), while certain topological manifolds admit no differentiable structure at all $[7^{163}]$. So it does make sense to speak of *specific differential structures* imposed on a topological space.

This is why some authors like $[2^{53}]$, $[22^3]$, $[7^{68}]$, $[12^{161}]$ define manifolds as topological spaces with a differentiable structure imposed on them. As argued in $[46^{134}]$, this is not very convincing: Who introduces a metric space M as a topological space with a continuous map $M \times M \to \mathbb{R}$ fulfilling the metric axioms? What one usually does is better viewed as follows: If one wants to metrize a given topological space, the continuity of the metric is an additional constraint that guarantees that the metric structure is compatible with the topology. In the same way, we can view Proposition 1.11 as an additional constraint for *imposing a differentiable structure on a given topological space*; this is essentially what we did in showing that the patchwork construction can be used for specifying a differentiable structure.

For the degenerate dimension $n = 0$, each chart $\varphi\colon U \mapsto A$ around a point p maps into $A \subseteq \mathbb{R}^0 = O$, so $A = \{0\}$ and $U = \{p\}$. Being chart domains, all singletons $\{p\}$ are therefore open sets, so M carries the discrete topology. This means a 0-*dimensional* manifold is the same as a discrete topological space (of any smoothness class), and obviously there is only one "differentiable structure" providing this topology. (An even more degenerate case occurs in the case $M = \emptyset$; for convenience reasons, one regards this "manifold" as being of every dimension and every smoothness class.)

What happens if we allowed non-Hausdorff almost-manifolds? A famous example is the *drilled line* built over $\mathbb{R}^* = \mathbb{R} \setminus \{0\}$. Here $M = \mathbb{R}^* \cup \{0_+, 0_-\}$ with two separate (mutually distinct) zeros $0_+, 0_- \notin \mathbb{R}$. We introduce an atlas containing two charts

$$\varphi_\pm \colon U_\pm \to \mathbb{R} \quad \text{with domains} \quad U_\pm = \mathbb{R}^* \cup \{0_\pm\},$$

defined by $\varphi(p) = p$ if $p \in \mathbb{R}^*$ and $\varphi(0_\pm) = 0$. The charts φ_+ and φ_- obviously form an atlas, which generates a differentiable structure (even of class C^ω) on M. In this topology, any two neighborhoods of 0_+ and 0^- intersect, so M is not a Hausdorff space.

Our next remark concerns the *countability requirement.* Recall that we call a topological space first countable if it has a countable local base and second countable if it has a countable global base. (A local base is the same thing as a neighborhood base: around each point, any neighborhood contains a base neighborhood. A global base provides enough open sets to generate all open sets as union of base sets.) Since $\mathbb{R}^n$ is clearly first countable and each chart maps its domain homeomorphically into $\mathbb{R}^n$, any set M that admits an atlas is automatically first countable.

In order to qualify as a manifold, it must fulfill the stronger requirement of second countability. A weird example of a first but not second countable almost-manifold is the *repeated line*, the disjoint union of uncountably many copies of $\mathbb{R}$. (Note that the Hausdorff axiom is obviously fulfilled, so the only flaw of this example is the lack of second countability!) For another strange example (also fulfilling all the other properties and moreover being connected), see $[7^{59}]$.

Altogether, the purpose of the two topological conditions is to ensure that we can do *global analysis* on manifolds: Loosely speaking, the Hausdorff axiom guarantees unique limits (like derivatives) and thus accounts for the *analysis* part, whereas second countability employs partitions of unity—to be introduced soon—for pasting together local data (like solutions of a differential equation on a small chart domain) into global objects and so provides the *global* part. Note that while the Hausdorff axiom tends to increase the number of open sets, second countability restricts them; in this sense, the topological conditions force a certain equilibrium on the open sets.

There is also another way of interpreting the topological requirements: they are exactly what is needed to embed manifolds into vector spaces, as achieved by the Whitney Embedding Theorem $[7^{91}]$. More precisely, a connected Hausdorff space with a differentiable structure is *embeddable* into a vector space iff it is a second countable $[17^{6}]$. In this sense, we can clearly see the requirement of doing "global analysis": It is just an abstract (and therefore often more economic) way of doing calculus in $\mathbb{R}^n$.

Let us at least mention that one can develop substantial theory without some of these requirements: See for example [31], which starts out *without Hausdorff axiom or second countability* (even allowing charts to map into inifinte-dimensional Banach spaces instead of $\mathbb{R}^n$). Some authors drop just second countability as $[1^{126}]$ $[59^{12}]$ $[42^{6}]$, others just the Hausdorff axiom as $[19^{7}]$ or Olver in restricted parts $[44^{212}]$.

A less drastic weakening that is occasionally applied [24] is to replace second countability by *paracompactness*, which means that every open cover admits a locally finite open refinement $[6^{94}]$. It turns out that this is all one needs for partitions of unity $[59^{34}]$, but of course one loses embeddability. In the presence of Hausdorff axiom, there are various conditions equivalent to this one, most importantly metrizability, separability of the connected components, and the possibility of building Riemannian metrics $[28^{5}]$. Generalizing from second countability to paracompactness just allows uncountably many connected components each of which must be second countable [40].

This brings us to another condition that is often required in addition to the ones we stipulated: that M be *connected* as a topological space. In the presence of this third requirement, second countability and paracompactness actually coincide (the repeated line is paracompact, missing second countability only due to its failure of being-connected). In fact, the paper [17] lists 88(!) subtly different versions of paracompactness that all happen to coincide on connected manifolds. On a connected manifold, we could actually drop the equidimensionality requirement that all charts be n-dimensional for a fixed $n \in \mathbb{N}$ since the dimension is constant on each connected component. (If we dropped equidimensionality on non-connected manifolds, we would end up with aconnected components each of which is a manifold of possibly different dimension.)

In these notes, we keep only the two conditions (Hausdorff axiom and second countability) in Definition 1.25, and we state explicitly when a manifold is to be connected (but the connected components must have the same dimension.) For the future it is more practical to ensure these two conditions directly by a *suitable atlas*, as in [44[3]]. This is done by requiring a countable atlas $(\varphi_i \colon U_i \to A_i \mid i \in \mathbb{N})$ for M such that for any two distinct points $p, q \in M$ with $p \in U_i$ and $q \in U_j$ there are open sets $A, B \subseteq \mathbb{R}^n$ with $\varphi_i(p) \in A \subseteq A_i$ and $\varphi_j(q) \in B \subseteq A_j$ such that $\varphi_i^{-1}(A) \cap \varphi_j^{-1}(B) = \emptyset$. The latter condition is an obvious restatement of the Hausdorff axiom, while the subsequent lemma guarantess second countability when the charts are indexed by $\mathbb{N}$.

1.26 Lemma *A locally Euclidean space admits a countable atlas iff it is second countable.*

Proof. Assume first $(\varphi_i \colon U_i \to A_i \mid i \in \mathbb{N})$ is a countable atlas for a locally Euclidean space M. Then every chart domain U_i satisfies the second axiom of countability since it is homeomorphic to $\mathbb{R}^n$, so let $(V_{ij} \mid j \in \mathbb{N})$ be a base for U_i. Any open set $U \subseteq M$ can then be decomposed as

$$U = \bigcup_{i \in \mathbb{N}} (U \cap U_i) = \bigcup_{i \in \mathbb{N}} \bigcup_{j \in J_i} V_{ij}$$

if each $U \cap U_i = \bigcup_{j \in J_i} V_{ij}$ is the corresponding decomposition in U_i, with index set $J_i \subseteq \mathbb{N}$. Hence $(V_{ij} \mid i, j \in \mathbb{N})$ is a countable base for M.

Conversely, assume that M is second countable, so there is a base $(V_i \mid i \in \mathbb{N})$ of open sets for M. If $(\varphi_i \colon U_i \to A_i \mid i \in I)$ is any atlas for M, we can expand its chart domains in terms of the base as

$$U_i = \bigcup_{j \in J_i} V_j$$

for suitable index sets $J_i \subseteq \mathbb{N}$. Now we construct an atlas $(\psi_j \mid j \in J)$ with the countable index set

$$J = \bigcup_{i \in I} J_i \subseteq \mathbb{N}$$

as follows. For every $j \in J$, there is an $i \in I$ with $j \in J_i$, so that $V_j \subseteq U_i$. Now let ψ_j be the restriction of the chart $\varphi_i \colon U_i \to A_i$ to V_j. It is now clear that $(\psi_j \mid j \in J)$ forms an atlas since its chart domains V_j are sufficient for covering any base set U_i, which in turn cover M. $\square$

Just as second countability guarantees a countable atlas, the next lemma states that *compactness* guarantess a finite atlas (more suggestively for connected manifolds: paracompact $\Rightarrow$ countable atlas, compact $\Rightarrow$ finite atlas). Note that a manifold with a finite atlas need not be compact as we see in the example $\mathbb{R}$ with its standard differentiable structure (generated by the identity). By the converse of the lemma we can obtain non-compact manifolds if we can show that they do not admit any finite atlas—the "infinite doughnut" seems to be a case in point.

1.27 Lemma *A locally Euclidean space M admits a finite atlas if it is compact.*

Proof. Assume M is compact and choose any atlas $(\varphi_i \colon U_i \to A_i \mid i \in I)$. Then $(U_i \mid i \in I)$ is an open covering of M, so compactness yields a finite subcovering $(U_j \mid j \in J)$ with $J \subseteq I$. Obviously $(\varphi_j \colon U_j \to A_j)$ is a finite atlas for M. $\square$

1.3 Maps between Manifolds

1.3.1 Differentiability of a Map

In the sequel, we will have to consider more than one manifold (with its various structures like atlas and later also tangent spaces, tensor spaces etc) at a time. It would be quite cumbersome to always refer explicitly to the manifold they come from. Hence we will *omit type references* when the context makes it clear. For example, speaking of a chart around a point, we refer to a chart in the differentiable structure of the manifold containing this point.

Fix a differentiability order r, together with an m-manifold M and an n-manifold N, each at least of (possibly distinct) class C^r. In the sequel, we the term *differentiable* should be understood as being of class C^r; this obviously includes all classes C^s with $r \leq s$. In the same vein, an *isomorphism* is to be understood as a C^r isomorphism, and so on for all concepts that will be defined in dependence on r. The explicit reference C^r is only used when needed.

In order to deal with maps f from M to N, the crucial tool is once again the method *pulling down*, producing a kind of shadow of f.

1.28 Definition *Let $\varphi \colon U \to A$ be a chart around $p \in M$ and $\psi \colon V \to B$ a chart around $f(p) \in N$ with $f(U) \subseteq V$. Then $f_{\psi\varphi} = \psi f \varphi^{-1} \colon A \to B$ is called the* local representative *of f around p, with respect to the charts φ and ψ.*

The *terminology and notation* is taken from [1^{131}], except that we prefer to write $f_{\psi\varphi}$ rather than $f_{\varphi\psi}$ to reflect the traditional notation of functional composition. Note also that being a local representative with respect to some charts includes the above-stated condition of mapping one domain into the other.

We observe that a *transition* from a chart $\varphi \colon U \to A$ to another chart $\varphi' \colon U' \to A'$ of a manifold M is the local representative of $1 \colon M \to M$ with respect to φ and φ' since we have

$$1_{\varphi'\varphi} = \varphi'\varphi^{-1} \colon \varphi(U \cap U') \to \varphi'(U \cap U').$$

From now on we will employ this as a notation for transitions.

We can now apply the local representative of a map for *defining differentiability* at a given point.

1.29 Definition *We say $f \colon M \to N$ is* differentiable *at the point $p \in M$ if there is a differentiable local representative $f_{\psi\varphi}$ around $\varphi(p)$, using some suitable charts φ around p and ψ around $f(p)$.*

As one can see easily, the definition does not depend on the choice of the charts φ and ψ: If other charts φ' and ψ' are chosen instead of φ and ψ, we see that

$$f_{\psi'\varphi'} = 1_{\psi'\psi} f_{\psi\varphi} 1_{\varphi\varphi'}$$

is also differentiable at $\varphi'(p)$ since both transitions $1_{\psi'\psi}$ and $1_{\varphi\varphi'}$ are isomorphisms. Instead of saying that f is differentiable at p, we may also call f continuous or smooth or analytic at p, in the respective cases $s = 0$ or $s = \infty$ or $s = \omega$. For the first case, this actually needs some justification.

1.30 Lemma *A map $f \colon M \to N$ is continuous at $p \in M$ iff some local representative $f_{\psi\varphi}$ is continuous around $\varphi(p)$.*

Proof. If f is continuous around p, we can choose an arbitrary chart domain V around $f(p)$ and then find a sufficiently small neighborhood $\tilde{U}$ of p with $f(\tilde{U}) \subseteq V$. By Proposition 1.17 we can pick a chart domain $U \subseteq \tilde{U}$, and again we have $f(U) \subseteq V$. Hence we may form the local representative $f_{\psi\varphi} = \psi f \varphi^{-1}$, which is then continuous around $\varphi(p)$ since f is continuous around p and φ as well as ψ are isomorphisms.

Conversely, assume $f_{\psi\varphi}$ is continuous at $\varphi(p)$ for some charts $\varphi \colon U \to A$ and $\psi \colon V \to B$ with $f(U) \subseteq V$. Then ψ maps any neighborhood $V_0 \subseteq V$ of $f(p)$ to a neighborhood $\psi(V_0) \subseteq B$ of $\psi f(p) = f_{\psi\varphi}(\varphi(p))$ since it is open. By the continuity of $f_{\psi\varphi} \colon A \to B$, we can find a neighborhood $\varphi(U_0) \subseteq A$ of $\varphi(p)$, with $U_0 \subseteq U$ being a neighborhood of p, such that $f_{\psi\varphi}(\varphi(U_0)) \subseteq \psi(V_0)$ or equivalently $f(U_0) \subseteq V_0$. □

As to be expected, we call a map *differentiable* if it is differentiable at every point, so all its local representatives are.

1.31 Definition *A map $f \colon M \to N$ is called differentiable if its local representatives $f_{\psi\varphi}$ are, for all charts φ in M and ψ in N.*

If the differentiability order is to be mentioned explicitly, one speaks of an r-times differentiable or C^r map. The class of all C^r maps from the manifold M to the manifold N is denoted by $C^r(M, N)$. It is a trivial exercise to check that the identity is differentiable and that composition preserves differentiability. Hence we obtain a *category* with objects the manifolds of class at least C^r and morphisms the C^r maps. (The most important case is the so-called smooth category, consisting of smooth manifolds and smooth maps).

As with the test of openness, we should expect that it suffices to check differentiability of maps with respect to an *atlas* [1[132]].

1.32 Lemma *A map $f\colon M \to N$ is differentiable iff the local representatives of f relative to some atlases of M and N are so.*

Proof. If f is differentiable, the conclusion follows trivially. The converse follows from the observation that any admissible chart domain may be obtained by the following three operations: chart transition, union, restriction; all these operations preserve the differentiability of the corresponding local representatives. □

Lemma 1.32 is of particular importance if M or N is an open subset of some *vector space*, which can be charted by identity map alone (see Example 1.19). So if M is an open set $U \subseteq \mathbb{R}^m$ and N an open set $V \subseteq \mathbb{R}^n$, we regain the old definition $C^r(U, V)$ given in Chapter 0. The notion of diffeomorphism between manifolds is also inspired by the case of $\mathbb{R}^n$.

1.33 Definition *We call a bijection $f\colon M \to N$ a* diffeomorphism *if both f and f^{-1} are differentiable; in this case, M and N are said to be diffeomorphic.*

Since any *open sets* $U \subseteq M$ and $V \subseteq N$ may be regarded as submanifolds (see Example 1.23), this extends to maps $f\colon U \to V$. In accordance with our conventions of Chapter 0, we always imply that U and V are open when stating that $f\colon U \to V$ is an isomorphism (again preferring the term C^0 homeomorphism for $r = 0$ and C^r diffeomorphisms for $r > 0$). By Lemma 1.30, this is also consistent in the case $r = 0$.

We see now immediately that every *coordinate chart* $\varphi\colon U \to A$ is a diffeomorphism, and all *coordinate curves / functions* are differentiable curves / functions. Furthermore, the differentiable structure of a manifold consists exactly of all diffeomorphisms from open subsets into the coordinate space.

We say that $f\colon M \to N$ is *locally diffeomorphic at* a point p if it acts as a diffeomorphism in an open neighborhood of p. (This usage of the term "locally diffeomorphic" is nonstandard but seems practical. Postnikov [46[226]] uses the term "étale", alluding to certain analogies in algebraic geometry.) If f is diffeomorphic at every point $p \in M$, it is called a *local diffeomorphism.* Note that a local diffeomorphism is a fortiori also a local homeomorphism and hence an open map. The diffeomorphisms are just those local diffeomorphisms that are additionally bijective.

1.3.2 The Sheaf Structure on a Manifold

In the previous section, we have treated smooth maps defined on the whole manifold M. As we have seen in the case of the charts (or the coordinate functions they contain), it is often relevant to consider maps that are only locally defined. In fact, it is often crucial to cut down a global or local map to a narrower domain—see the proof (sketch) of Lemma 1.32 for a case in point. The underlying notion of *functional restriction* can best be analyzed if one focuses on differentiable functions on M; as mentioned before, this includes the coordinate functions of M.

The desire to capture the notion of functional restriction emerges in various other branches of mathematics, most notably in algebraic geometry [20[69]], and one usually codifies restriction by a standard structure called *sheaf*. As we do not intend to build up sheaf theory at this place, it will be sufficient if we review briefly how sheaves are defined. For our purposes, we are only interested in sheaves of algebras (one may of course supply any other category—like [20[61]] who uses mainly abelian groups—in place of algebras), and we prefer to enunciate the definition without explicit use of categories. By using algebras, we imply that "homomorphisms" should be understood in the sense of algebras.

Before going to sheaves, one usually starts with a preliminary structure, called a *presheaf*—it incorporates an abstract notion of functional restriction, but still lacks the power to use this notion effectively in pasting together functions from local patches: this will be the task of sheaves.

1.34 Definition *A* presheaf *on a topological space X is a map $\mathcal{F}$ that associates an algebra $\mathcal{F}(U)$ with any open set $U \subseteq X$, together with so-called restriction homomorphisms* $\mathrm{res}_{V,U}\colon \mathcal{F}(U) \to \mathcal{F}(V)$ *for any open sets $V \subseteq U \subseteq X$, such that* $\mathrm{res}_{U,U} = 1_{\mathcal{F}(U)}$ *and* $\mathrm{res}_{W,V}\,\mathrm{res}_{V,U} = \mathrm{res}_{W,U}$.

The elements of $F(U)$ are usually called *sections*; the reason for this terminology will become clear in Subsection 2.3.1. Section from $\mathcal{F}(X)$ are also called global, all other ones local. In many important cases, the sections $f \in \mathcal{F}(U)$ are *actual functions* on the open set U, and $\mathrm{res}_{U,V}$ is the set-theoretic restriction from the domain U to the smaller domain U. In this case, one can safely abbreviate $\mathrm{res}_{V,U}(f)$ by the familiar notation $f|_V$, since a map f encodes its domain U as explained in Chapter 0. Note that the presheaf axioms are then automatically fulfilled as long as the restriction is closed in the sense that $\mathrm{res}_{V,U}\,\mathcal{F}(U) \subseteq \mathcal{F}(V)$.

In order to get a *sheaf*, one has to ensure two further properties about "patching functions": Every open cover $(U_i \mid i \in I)$ must fulfill the following conditions: Two sections coincide whenever they have the same restriction on each U_i; and any family of sections on U_i with common restrictions comes from a single section. Some authors like [53[391]] subsume the first axiom into the second by requiring the unique existence of the single section $f \in \mathcal{F}(U)$. Let us now formulate the definition in a bit more detail.

1.35 Definition *A presheaf $\mathcal{F}$ is called a* sheaf *if every open cover $(U_i \mid i \in I)$ satisfies the following two axioms:*

1. *For $f, g \in \mathcal{F}(U)$ with $\mathrm{res}_{U_i,U} f = \mathrm{res}_{U_i,U}\, g$ for all $i \in I$, we have $f = g$.*
2. *If $f_i \in \mathcal{F}(U_i)$ is such that $\mathrm{res}_{U_i \cap U_j, U_i} f_i = \mathrm{res}_{U_i \cap U_j, U_j} f_j$ for all $i, j \in I$, there is an $f \in \mathcal{F}(U)$ with $\mathrm{res}_{U_i,U} f = f_i$ for all $i \in I$.*

We may remove axiom (1) if we stipulate unique existence in axiom (2).

One of the most important sheaves is given by the algebra of continuous functions on X. In this case, the restriction is of course meant in the set-theoretic sense, and one employs the notation $f|_U$ as explained above; the continuous functions defined on an open set U are accordingly denoted by $C_X(U)$. It is a trivial exercise to check that $U \mapsto C(U)$ is indeed a sheaf of algebras, written C_X and known as the *sheaf of continuous functions* (more compactly, the continuous sheaf) on X.

What is more important for our present purposes is the algebra of differentiable functions on a manifold M, introduced in Subsection 1.3.1. Since we may view any open set $U \subseteq M$ as a manifold it its own right (see Example 1.23), we can also introduce the algebra $C^r_M(U)$ of differentiable functions with domain U. Note that this is also consistent with the topological case discussed before since we have $C_M(U) = C^0_M(U)$ by Lemma 1.30. Using again set-theoretic restriction, this yields a sheaf: First of all, it is clear that we have a presheaf since we deal with actual functions. Then every $C^r_M(U)$ is a subalgebra of $C_M(U)$, so sheaf axiom (1) is inherited from $C_M(U)$. Finally, sheaf axiom (2) follows easily from Lemma 1.32. Let us summarize this now.

1.36 Definition *The* sheaf of differentiable functions *(more compactly, the differentiable sheaf) on M, denoted by C^r_M, assigns to each open set U of M the algebra $C^r_M(U)$ of differentiable functions on U. In the special case $M = \mathbb{R}^n$, we speak of the* differentiable Euclidean sheaf.

A *subsheaf* $\mathcal{F}'$ of the sheaf $\mathcal{F}$ is a sheaf that assigns to each open set $U \subseteq X$ a subalgebra $\mathcal{F}'(U)$ of the algebra $\mathcal{F}(U)$ such that the restriction homomorphisms of $\mathcal{F}'$ are induced by those of $\mathcal{F}$. Subsheaves typically arise by cutting down the ground space X to some subset X', regarded as a topological space with the induced topology. Then the *restriction sheaf* of $\mathcal{F}$ to X' is written as $\mathcal{F}|_{X'}$ and defined by

$$\mathcal{F}|_{X'}(U) = \mathcal{F}(U)$$

for all U open in X'. In other words, the restriction sheaf $\mathcal{F}_{X'}$ does exactly the same as $\mathcal{F}$ but ignores any open sets outside of X'.

A *sheaf morphism* between a sheaf $\mathcal{F}$ on the space X and a sheaf $\mathcal{G}$ on the space Y is given by a continuous map $\varphi \colon X \to Y$ together with homomorphisms $\varphi_V : \mathcal{G}(V) \to \mathcal{F}(\varphi^{-1}V)$ for open sets V in Y such that

$$\mathrm{res}_{U_1,U_2} \circ \varphi_{V_2} = \varphi_{V_1} \circ \mathrm{res}_{V_1,V_2}$$

for open sets $V_1 \subseteq V_2$ in Y and $U_1 = \varphi^{-1}V_1 \subseteq U_2 = \varphi^{-1}V_2$ in X. If the sections of $\mathcal{F}$ and $\mathcal{G}$ are all functions, the canonical choice for φ_V is the precomposition $g \in \mathcal{G}(V) \mapsto g \circ \varphi \in \mathcal{F}(\varphi^{-1}V)$, and we can think of φ alone as a sheaf morphism. In this case, one only has to make sure that $g \mapsto g \circ \varphi$ carries $\mathcal{G}(V)$ into $\mathcal{F}(\varphi^{-1}V)$.

As usual, a *sheaf isomorphism* is a sheaf morphism that has an inverse which is also a sheaf morphism (if all sections are functions, this is just a homeomorphism $\varphi \colon X \to Y$ such that both φ and φ^{-1} induce a sheaf morphism); in this case, we say that the sheaves $\mathcal{F}$ and $\mathcal{G}$ are isomorphic. This notion can be localized: The sheaves $\mathcal{F}$ and $\mathcal{G}$ are called *locally isomorphic* if every point has an open neighborhood $U \subseteq X$ such that $\mathcal{F}|_U$ is isomorphic to $\mathcal{G}|_V$ for some open set $V \subseteq Y$.

The *special role of the differentiable Euclidean sheaf* is that it models the differentiable sheaf on any manifold: Locally, they always look the same.

1.37 Proposition *The differentiable sheaf on M is locally isomorphic to the differentiable Euclidean sheaf.*

Proof. Given a point $p \in M$, we must find an open neighborhood $U \subseteq M$ of p and an open set $V \subseteq \mathbb{R}^n$ such that $C^r_M|_U$ is locally isomorphic to $C^r_{\mathbb{R}^n}|_V$. Choosing any chart $\varphi\colon U \to V$ around p, we show that φ gives the required sheaf isomorphism between $C^r_M|_U = C^r_U$ and $C^r_{\mathbb{R}^n}|_V = C^r_V$. As any other chart, φ is certainly a homeomorphism between the topological spaces U and V, so it remains to check that $\varphi_{V_0}\colon g \mapsto g \circ \varphi$ maps $C^r_{\mathbb{R}^n}(V_0)$ into $C^r_M(\varphi^{-1}V_0)$ for open sets $V_0 \subseteq V$ and $\varphi_{U_0}\colon f \mapsto f \circ \varphi^{-1}$ maps $C^r_M(U_0)$ into $C^r_{\mathbb{R}^n}(\varphi U_0)$ for open sets $U_0 \subseteq U$. But obviously $\{\varphi\}$ is an atlas for the submanifold $U \subseteq M$ while $\{1_{\mathbb{R}}\}$ is an atlas on $\mathbb{R}$, so Lemma 1.32 shows that a map $f\colon U_0 \to \mathbb{R}$ is differentiable iff $f \circ \varphi^{-1}$ is. □

The nature of sheaf isomorphisms between subsheaves of the differentiable Euclidean sheaf is particularly transparent: they are precisely the *differentiable maps*!

1.38 Lemma *For any open sets $A, B \subseteq \mathbb{R}^n$, a homeomorphism $\vartheta\colon A \to B$ is a sheaf isomorphism between C^r_A and C^r_B iff ϑ is differentiable.*

Proof. If $\vartheta\colon A \to B$ is differentiable, the precompositions $g \mapsto g \circ \vartheta$ and $f \mapsto f \circ \vartheta^{-1}$ clearly interchange differentiable functions defined on open subsets of B with those defined on open subsets of A. So assume conversely that ϑ is a sheaf isomorphism between C^r_A and C^r_B. Then $\delta^i \circ \vartheta\colon A \to \mathbb{R}$ is differentiable since each projection $\delta^i\colon B \to \mathbb{R}$ is differentiable for $i = 1, \dots, n$. A vector-valued function is differentiable iff its components are, so $\vartheta\colon A \to B \subseteq \mathbb{R}^n$ is also differentiable. □

The role of the Euclidean sheaf is that it models how differentiation works on a manifold M, so it can be expected that we may even characterize differentiable manifolds in this way—having differentiable sheaves that *look like* the differentiable Euclidean sheaf.

1.39 Definition *Given a topological space M, a subsheaf of C_M is called a* differentiable sheaf *on M if it is locally isomorphic to the differentiable Euclidean sheaf.*

We will now *reconstruct the differentiable structure* on M. Before doing so, observe that we have actually associated a differentiable sheaf $C^r_{\mathfrak{D}}$ to the differentiable structure $\mathfrak{D}$ of M since we have not used the topological conditions of Definition 1.25.

1.40 Proposition *For a topological space M, every differentiable sheaf $\mathcal{F}$ on M induces a differentiable structure $D^r_{\mathcal{F}}$ on M whose differentiable sheaf coincides with the original $\mathcal{F}$.*

Proof. Given a differentiable sheaf $\mathcal{F}$, we define the differentiable structure $D^r_{\mathcal{F}}$ by choosing as charts all those homeomorphisms between open sets $U \subseteq M$ and open sets $A \subseteq \mathbb{R}^n$ that induce sheaf isomorphisms between $\mathcal{F}|_U$ and $C^r_{\mathbb{R}^n}|_A$. Let us check that $D^r_{\mathcal{F}}$ is indeed a differentiable structure on M. First of all, it is clear that the chart domains cover M because $\mathcal{F}$ is locally isomorphic to $C^r_{\mathbb{R}^n}$. Next we have to ascertain differentiability of the transitions $\psi\varphi^{-1}\colon A' \to B'$

between charts $\varphi\colon U \to A$ and $\psi\colon V \to B$ with overlaps $A' = \varphi(U \cap V), B' = \psi(U \cap V) \subseteq \mathbb{R}^n$. But $\psi\varphi^{-1}$ is a sheaf isomorphism between $C^r_{A'}$ and $C^r_{B'}$, so differentiability follows from Lemma 1.38.

We have now proved that $D^r_{\mathcal{F}}$ is a differentiable atlas, and it only remains to show that it is maximal. So we must prove that any admissible chart $\varphi\colon U \to A$ induces a sheaf isomorphism between $\mathcal{F}|_U$ and $C^r_{\mathbb{R}^n}|_A$. For any point $p \in V$ there is a chart $\varphi\colon U' \to A'$ around p with $U_p = U \cap U' \neq \emptyset$. Write $\varphi_p\colon U_p \to A_p$ and $\psi_p\colon U_p \to A'_p$ with $A_p = \psi(U_p) \subseteq A$ and $A'_p = \varphi(U'_p) \subseteq A'$ for the corresponding restrictions. Since φ is compatible with ψ, the transition $\varphi_p\psi_p^{-1}\colon A'_p \to A_p$ is differentiable and thus a sheaf isomorphism between $C^r_{A'_p}$ and $C^r_{A_p}$ by Lemma 1.38. Being a chart, ψ also induces a sheaf isomorphism, and its restriction $\psi_p\colon U_p \to A'_p$ induces a sheaf isomorphism between $\mathcal{F}|_{U_p}$ and $C^r_{A'_p}$. Hence $\varphi_p = \varphi_p\psi_p^{-1} \circ \psi_p\colon U_p \to A_p$ is a sheaf isomorphism between $\mathcal{F}|_{U_p}$ and $C^r_{A_p}$.

Take any differentiable map $g\colon A_0 \to \mathbb{R}$ on an open set $A_0 \subseteq A$. Writing $U_0 = \varphi^{-1}(A_0)$, we obtain a differentiable map $g_0 \circ \varphi_p\colon U_0 \cap U_p \to \mathbb{R}$. Now the sets $U_0 \cap U_p$ form an open cover of U_0, while the $g_0 \circ \varphi_p$ clearly have common restrictions, so invoking sheaf axiom (2) for $\mathcal{F}|_U$ yields a section $f_0 \in \mathcal{F}(U_0)$. But this function $f_0\colon U_0 \to \mathbb{R}$ actually agrees with $g_0 \circ \varphi$ since we have $f_0(p) = g_0(\varphi_p(p)) = g_0(\varphi(p))$ for any $p \in U_0$. Hence we have proved $g_0 \circ \varphi \in \mathcal{F}(\varphi^{-1}A_0)$ for any function $g_0 \in C^r_A(A_0)$ defined on an open set $A_0 \subseteq A$, so φ induces a sheaf morphism from $\mathcal{F}|_U$ to $C^r_{\mathbb{R}^n}|_A$. In order to show that conversely $\varphi^{-1}\colon A \to U$ induces a sheaf morphism from $C^r_{\mathbb{R}^n}|_A$ to $\mathcal{F}|_U$, we can use the same technique with φ_p^{-1} in place of φ_p and sheaf axiom (2) for $C^r_{\mathbb{R}^n}|_A$ instead of $\mathcal{F}|_U$.

This concludes the proof that $\mathfrak{D} = D^r_{\mathcal{F}}$ is indeed a differentiable structure. Next we make sure that its differentiable sheaf $C^r_{\mathfrak{D}}$ actually coincides with the original differentiable sheaf $\mathcal{F}$. For an open set $V \subseteq M$, we have to prove that $f\colon V \to \mathbb{R}$ is differentiable iff $f \in \mathcal{F}(V)$. Again we can choose a chart around any point $p \in V$ and then restrict it to a small open neighborhood U_p within V. The restricted chart $\varphi_p\colon U_p \to A_p$ is clearly admissible, hence it induces a sheaf isomorphism between $\mathcal{F}|_{U_p}$ and $C^r_{\mathbb{R}^n}|_{A_p}$.

Assume now $f\colon V \to \mathbb{R}$ is differentiable. Then every local representative $f \circ \varphi_p^{-1}\colon A_p \to \mathbb{R}$ is also differentiable. Since φ_p induces a sheaf isomorphism between $\mathcal{F}|_{U_p}$ and $C^r_{\mathbb{R}^n}|_{A_p}$, it maps $f \circ \varphi_p^{-1} \in C^r_{\mathbb{R}^n}(A_p)$ to $f_p = (f \circ \varphi_p^{-1}) \circ \varphi_p \in \mathcal{F}(U_p)$. Again it is clear that the U_p form an open cover of V and the f_p agree on their common restrictions. Thus sheaf axiom (2) for $\mathcal{F}|_V$ yields $f \in \mathcal{F}(V)$ by an argumentation similar to the one above.

Now assume $f \in \mathcal{F}(V)$. By Lemma 1.32, it suffices to prove that the local representatives $f \circ \varphi_p^{-1}$ are differentiable. But this follows immediately from $f|_{U_p} \in \mathcal{F}(U_p)$ and the fact that φ_p^{-1} induces a sheaf morphism from $C^r_{\mathbb{R}^n}|_{A_p}$ to $\mathcal{F}|_{U_p}$. □

It is clear how we can go in the other direction: A differentiable structure provides a differentiable sheaf simply by taking all the differentiable functions—this is how it all started. But now we have seen how to build up a differentiable structure if we are just given the sheaf, and the question remains whether this leads us back to the *original differentiable structure*.

1.41 Proposition *For any differentiable structure $\mathfrak{D}$ on a topological space M, the corresponding differentiable sheaf $C^r_{\mathfrak{D}}$ induces the same differentiable structure $\mathfrak{D}$.*

Proof. We have to prove that a homeomorphisms φ between an open set $U \subseteq M$ and an open set $A \subseteq \mathbb{R}^n$ induces a sheaf isomorphism between $C^r_{\mathfrak{D}}|_U$ and $C^r_{\mathbb{R}^n}|_A$ iff φ is a chart of $\mathfrak{D}$.

Assume $\varphi\colon U \to A$ induces a sheaf isomorphism between $C^r_{\mathfrak{D}}|_U$ and $C^r_{\mathbb{R}^n}|_A$. For showing that φ is a chart of $\mathfrak{D}$, it is sufficient to prove that it is differentiable. A vector-valued function is differentiable iff all its components are, so we just have to show that $\delta^i \circ \varphi\colon U \to \mathbb{R}$ is differentiable for any $i \in \{1, \dots, n\}$. Clearly $\delta^i\colon \mathbb{R}^n \to \mathbb{R}$ is differentiable, so we have $\delta^i \in C^r_{\mathbb{R}^n}(A)$. But this implies $\delta^i \circ \varphi \in C^r_{\mathfrak{D}}(U)$ since φ induces a sheaf morphism between $C^r_{\mathfrak{D}}|_U$ and $C^r_{\mathbb{R}^n}|_A$, which means φ is indeed differentiable on its domain U.

Conversely, take a chart $\varphi\colon U \to A$ of $\mathfrak{D}$. In order to show that φ induces a sheaf morphism between $C^r_{\mathfrak{D}}|_U$ and $C^r_{\mathbb{R}^n}|_A$, take a $g \in C^r_{\mathbb{R}^n}(A)$. We have to show $g \circ \varphi \in C^r_{\mathfrak{D}}(U)$, which means that $g \circ \varphi\colon U \to \mathbb{R}$ is differentiable. Applying Lemma 1.32, it suffices to prove that $g \circ \varphi \circ \psi_p^{-1}\colon U_p \to \mathbb{R}$ is differentiable for every chart $\psi_p\colon U_p \to A_p$ around a point $p \in U$ with domain U_p. But φ itself is differentiable since it is a chart, hence also $\varphi \circ \psi_p^{-1}$ by the definition of differentiable maps on a manifold. Consequently, $g \circ \varphi \circ \psi_p^{-1}$ is differentiable as a composite of the differentiable maps g and $\varphi \circ \psi_p^{-1}$. Finally we must show that φ^{-1} is a sheaf morphism between $C^r_{\mathbb{R}^n}|_A$ and $C^r_{\mathfrak{D}}|_U$, so let $f \in C^r_{\mathfrak{D}}(U)$. Then $f\colon U \to \mathbb{R}$ is a differentiable map, so its local representative $f \circ \varphi^{-1}$ must also be differentiable, and therefore $f \circ \varphi^{-1} \in C^r_{\mathbb{R}^n}(A)$ as required. □

We can summarize our findings in the sense that differentiable structures and differentiable sheaves are *essentially the same*; see Exercise 2 in [7[71]]

1.42 Theorem *For a topological space M, differentiable structures and differentiable sheaves are in bijective correspondence.*

Proof. Write D^r for the construction $\mathcal{F} \mapsto D^r_{\mathcal{F}}$ associating a differentiable structure to a given differentiable sheaf, and correspondingly also C^r for the assignment $\mathfrak{D} \mapsto C^r_{\mathfrak{D}}$ extracting the differentiable sheaf from a given differentiable structure. If we denote the set of all differentiable structures on M by $\mathrm{Str}(M)$ and the set of all differentiable sheaves on M by $\mathrm{Shf}(M)$, we can view this as two function $D^r\colon \mathrm{Str}(M) \to \mathrm{Shf}(M)$ and $C^r\colon \mathrm{Shf}(M) \to \mathrm{Str}(M)$, and Proposition 1.40 states that $C^r \circ D^r = 1_{\mathrm{Shf}(M)}$ while Proposition 1.41 gives $D^r \circ C^r = 1_{\mathrm{Str}(M)}$. Hence C^r and D^r are inverse to each other, and the sets $\mathrm{Str}(M)$ and $\mathrm{Shf}(M)$ are indeed in bijective correspondence. □

Of course, this also provides an *alternative way of defining manifolds* in the first place [18[80]], emphasizing how "continuity structure" (the topology of M) is enriched by an additional "differentiability structure" (the differentiable sheaf on M). Besides this, the sheaf definition also makes certain relations with algebraic geometry more transparent [20].

Let us also remark that one can also define the notion of a *differentiable map* between manifolds in a "sheaf fashion" [18[80]]; see also Exercise 3 in [7[71]]. In the following proposition, we assume the same conventions about dimensions and differentiability orders as in Subsection 1.3.1.

1.43 Proposition *The differentiable maps between two manifolds M and N are exactly the sheaf morphisms between C^r_M and C^r_N.*

Proof. Assume first $F\colon M \to N$ is differentiable. For proving that F induces a sheaf morphism between C^r_M and C^r_N, we take $g \in C^r_N(V)$ for an open set $V \subseteq N$ and prove $g \circ F \in C^r_M(U)$ with $U = F^{-1}(V)$. Using Lemma 1.32, it is enough to prove that all local representatives $g \circ F \circ \psi^{-1}\colon \tilde{U} \to \mathbb{R}$ are differentiable, where $\psi\colon \tilde{U} \to \tilde{A}$ is a chart around an arbitrary point

$p \in U$ with domain $\tilde{U} \subseteq U$. Since F is differentiable, we know that $F \circ \psi^{-1} : \tilde{A} \to \mathbb{R}$ is as well. But then $g \circ F \circ \psi^{-1}$ is differentiable since it is the composition of the differentiable maps g and $F \circ \psi^{-1}$.

For the converse, assume $F : M \to N$ induces a sheaf morphism between C^r_M and C^r_N. For proving F to be differentiable, we take arbitrary charts $\varphi : U \to A$ on M and $\psi : V \to B$ on N with $F(U) \subseteq V$ and show that the local representative $\psi \circ F \circ \varphi^{-1} : A \to B$ is differentiable. A vector-valued function is differentiable iff all its components are, so it will be enough to prove that $\delta^i \circ \psi \circ F \circ \varphi^{-1} : A \to \mathbb{R}$ is differentiable for all $i \in \{1, \dots, n\}$. Note first that $\delta^i \circ \psi \circ F : U \to \mathbb{R}$ is differentiable: Since $\delta^i \circ \psi : V \to \mathbb{R}$ is differentiable, it belongs to $C^r_N(V)$, while precomposition with F carries $C^r_N(V)$ into $C^r_M(\varphi^{-1}V)$ because F is a sheaf morphism. Hence $\delta^i \circ \psi \circ F$ is differentiable even on $F^{-1}(V) \supseteq U$. Finally, $\delta^i \circ \psi \circ F \circ \varphi^{-1} : A \to \mathbb{R}$ is differentiable since it is the composition of the differentiable map $\delta^i \circ \psi \circ F : U \to \mathbb{R}$ and the diffeomorphism $\varphi^{-1} : A \to U$. □

We note finally that there is a sheaf analog of the *patchwork construction* described in Subsection 1.2.3: One may glue together sheaves given on an open cover of a common topological space [20[69]]. It seems that this provides again a fully analogous sheaf treatment of general patchworks in the following sense: The construction of Subsection 1.2.3 may also be carried out if we are given submanifolds $A_i \subseteq A$ of a given base manifold A, yielding a new "manifold" M modelled on A rather than $\mathbb{R}^n$. If A itself is an n-manifold, we may also regard the new manifold M as an n-manifold by composing the charts of A with the A-charts of M. This means we can construct an n-manifold by patching together various n-manifolds contained in a base manifold (which was $\mathbb{R}^n$ in Subsection 1.2.3). Under the sheaf-manifold correspondence of Theorem 1.42, this construction might correspond to the sheaf-gluing process as described in [20[69]].

1.3.3 Germs and Stalks

Assuming the same conventions as in Subsection 1.3.1, let $f : U \to N$ and $g : V \to N$ be maps defined on open neighborhoods of a point $p \in M$. Given a point $p \in M$, we put $f \sim_p g$ if there is an open neighborhood $W \subseteq U \cap V$ of p with $f|_W = g|_W$. It is easily verified that $\sim_p$ is an equivalence relation on the set of all functions defined on a neighborhood of p, and the resulting quotient forms a vector space, denoted by $G^r_{M,N}(p)$ and known as the differentiable *stalk* at p. Its elements are called differentiable *germs* at p, and the canonical projection is written as

$$C^r(U, N) \to G^r_{M,N}(p), \quad f \mapsto f_p.$$

This construction is based only on the sheaf of differentiable functions from M to N, and one may in general construct the stalks [20[62]] for any given sheaf (not necessarily consisting of actual functions with set-theoretic restriction) and for any category (not just the category of algebras used here).

In the special case $N = \mathbb{R}$, we obtain even an algebra that is briefly written as $G^r_M(p)$ or just $G^r(p)$ if M is understood from the context. For such germs $f_p \in G^r(p)$, one can also define evaluation by $f_p(p) = f(p)$. Since we will be working exclusively with such algebras, the objects $G^r(p)$ will be referred to by the more descriptive name *germ algebra* instead of the general term "stalk".

At any given point $p \in M$, a continuous map $f\colon M \to N$ with $q = f(p)$ induces a *germ transport*

$$f_{(p)}\colon G(q) \to G(p), \quad \psi_q \mapsto (\psi \circ f)_p;$$

note that $f_{(p)}$ is a homomorphism of algebras. Following [36[10]], we call $f\colon M \to N$ *loaded at* p if the germ transport $f_{(p)}$ restricts to $f_{(p)}\colon G^r(q) \to G^r(p)$, and a *loaded map* if it is loaded at all points $p \in M$. A bijection f is called a *loaded isomorphism* if both f and f^{-1} are loaded maps; consequently, f is a homeomorphism and the germ transports are isomorphisms $f_{(p)}\colon G^r(q) \to G^r(p)$.

Having differentiable structures on M and N, it turns out that *differentiability* can be characterized in a germ-theoretical way.

1.44 Proposition *Let $f\colon M \to N$ be a continuous map between two manifold M and N. Then f is differentiable at $p \in M$ iff it is loaded at p. Hence the differentiable maps coincide with the loaded ones.*

Proof. Assume first that f is differentiable at p with $q = f(p)$ and consider the germ transport $f_{(p)}\colon G(q) \to G(p)$. Take any $\psi_q \in G^r(q)$ and choose some representative $\psi\colon V \to \mathbb{R}$ of ψ_q. Since f is differentiable at p and ψ at $q = f(p)$, the composite ψf is also differentiable at p, so we have indeed $(\psi f)_p \in G^r(p)$.

Conversely, consider the germ transport $f_{(p)}\colon G^r(q) \to G^r(p)$. Choosing suitable charts $\varphi\colon U \to A$ around p and $\psi\colon V \to B$ around q, it is sufficient to prove that $f_{\psi\varphi} = \psi f \varphi^{-1}$ is differentiable at $\varphi(p)$. As usual, we employ the characterization of differentiable vector functions; so let us show that $\delta^i f_{\psi\varphi}$ is differentiable at $\varphi(p)$ for an arbitrary $i \in \{1, \ldots, n\}$. Since $(\delta^i\psi)_q \in G^r(q)$, we may infer that $(\delta^i\psi f)_p \in G^r(p)$. Hence we may choose some neighborhood $\tilde{U} \subseteq U$ of p such that $\delta^i\psi f\colon \tilde{U} \to \mathbb{R}$ is differentiable at p. Since we can choose any charts $\varphi\colon U \to A$ and $\psi\colon V \to B$ as long as $f(U) \subseteq V$, we may as well take $U = \tilde{U}$ and $V = f(U)$. But then $\delta^i\psi f\varphi$ is indeed differentiable at $\varphi(p)$. □

In the previous Subsection 1.3.2, we have seen that the differentiable structure of a manifold is already determined by its differentiable sheaf, and we have characterized which sheaves are suitable for this purpose. It turns out that we can even go down to the "infinitesimal" level of the germs: Knowing the germ algebra $G^r(p)$ at each point $p \in M$, we can *reconstruct the differentiable sheaf* and hence the differentiable structure of M, whose charts may also be characterized directly in terms of the germ algebras. The crucial point here is (as earlier!) to go back to the Euclidean space $\mathbb{R}^n$ as a reference object—now understood as a C^r manifold having its own germ algebras $G^r(x)$ for every coordinate node $x \in \mathbb{R}^n$.

1.45 Proposition *For any open set $U \subseteq M$, we have*

$$C^r(U) = \{f \in C(U) \mid \forall_{p \in U}\, f_p \in G^r(p)\}.$$

The charts of M are exactly the loaded isomorphisms between open sets of M and $\mathbb{R}^n$.

Proof. For $f \in C^r(U)$, it is clear that each germ f_p is a differentiable germ. So we take a continuous function $f\colon U \to \mathbb{R}$ with all germs $f_p \in G^r(p)$, and we must prove that $f \in C^r(U)$. Take a chart $\varphi\colon U \to A$ around p and let $g \in C^r(V)$ be a representative of f_p with $f|_V = g$ for a sufficiently small neighborhood $V \subseteq U$ of p. Restricting the chart to $\psi\colon V \to B$, we see that $g \circ \psi^{-1} = f \circ \psi^{-1}\colon B \to \mathbb{R}$ is differentiable at $\varphi(p)$, so f is differentiable at p by definition. Since $p \in U$ was arbitrary, we see that $f \in C^r(U)$ as claimed.

For proving the second statement of the proposition, one could probably evoke the correspondence of Theorem 1.42 and apply a suitable adaption of a general theorem in [20[63]] connecting sheaf morphisms and their induced operation on the stalks. As this seems to be rather awkward, we prefer to give an independent proof here.

As noted earlier, the charts of a manifold M are exactly those homeomorphisms between open sets of M and $\mathbb{R}^n$ which are differentiable. By Proposition 1.44, differentiability means being loaded everywhere, so the charts are indeed loaded isomorphisms and vice versa. ☐

Since the germ algebras of a manifold suffice to build its differentiable sheaf (equivalently, its differentiable structure), it is natural to characterize *what we need to describe a manifold* in terms of germ algebras.

1.46 Definition *A* loaded space *is a topological space M with an assignment $\mathcal{G}$ that associates a germ algebra $\mathcal{G}(p) \subseteq G(p)$ with each point $p \in M$; such an assignment $\mathcal{G}$ is then called a* load *on the topological space M. If specifically $M = \mathbb{R}^n$, the* Euclidean load $G^r_{\mathbb{R}^n}$ *associates the germ algebra $G^r_{\mathbb{R}^n}(x)$, briefly $G^r(x)$, to a point $x \in \mathbb{R}^n$.*

Any subset U of M is again a loaded space, carrying the induced topology as well as the obvious induced load $\mathcal{G}|_U$. The above *concepts about loads* (maps loaded at a point, loaded maps, loaded isomorphisms) generalize to loaded spaces.

A *loaded chart* is a homeomorphisms between open sets $U \subseteq M$ and $A \subseteq \mathbb{R}^n$ that acts as a loaded isomorphism between U and A, each carrying the induced load. (Here it is important which load we associate to $\mathbb{R}^n$ and hence to A, this depending on the differentiability order r inherent in the germ algebras $G^r(x)$ making up the Euclidean load!) The crucial role of the loaded charts is that they will allow us to characterize those loads that come from a differentiable structure (or sheaf).

1.47 Definition *A load on a topological space M is called a* differentiable load *on M if its loaded charts cover M.*

Any differentiable sheaf $\mathcal{F}$ on M provides a differentiable load $G^r_{\mathcal{F}}$ on M by associating

$$G^r_{\mathcal{F}}(p) = \{f_p \mid f \in \mathcal{F}(U) \wedge U \ni p\}$$

with a point $p \in M$; this works since Proposition 1.45 allows us to identify the charts of M with the loaded charts of $G^r_{\mathcal{F}}$. Let us next reassure ourselves that we can also go in the *other direction.*

1.48 Proposition *For a topological space M, every differentiable load $\mathcal{G}$ on M induces a differentiable sheaf $C^r_{\mathcal{G}}$ on M whose differentiable load coincides with the original $\mathcal{G}$.*

Proof. Given a load $\mathcal{G}$, we can proceed in a manner analogous to Proposition 1.45, defining

$$C^r_{\mathcal{G}}(U) = \{f \in C(U) \mid \forall_{p \in U}\, f_p \in \mathcal{G}(p)\}.$$

Let us first show that $C^r_{\mathcal{G}}$ is a differentiable sheaf on M. Obviously each $C^r_{\mathcal{G}}(U)$ is a subalgebra of of $C(U)$, and $U \mapsto C^r_{\mathcal{G}}(U)$ is at least a presheaf since we are now dealing with actual functions and set-theoretic restriction. In order to verify the sheaf axioms, consider an open cover $(U_i \mid i \in I)$ of U. Taking $f, g \in C^r_{\mathcal{G}}(U)$ with $f|_{U_i} = g|_{U_i}$ for all $i \in I$, we obtain $f = g$ since this is already true for $f, g \in C(U)$. For checking the other sheaf axiom, let $f_i \in C^r_{\mathcal{G}}(U_i)$ such that $f_i|_{U_i \cap U_j} = f_j|_{U_i \cap U_j}$ for all $i, j \in I$. From the continuous sheaf $U \mapsto C(U)$, we obtain at least an $f \in C(U)$ with $f|_{U_i} = f_i$ for all $i \in I$; it remains to prove that actually $f \in C^r_{\mathcal{G}}(U)$. Hence we must show $f_p \in \mathcal{G}(p)$ for an arbitrary $p \in U$. Since the U_i cover U, there is a U_i with $p \in U_i$ and so $f_p = (f_i)_p$. But $f_i \in C^r_{\mathcal{G}}(U_i)$ implies that $(f_i)_p \in \mathcal{G}(p)$, so the other sheaf axiom is also fulfilled.

Now we construct the differentiable load $G^r_{\mathcal{F}}$ induced by $\mathcal{F} = C^r_{\mathcal{G}}$, and we will show that actually $G^r_{\mathcal{F}} = \mathcal{G}$. By definition, $G^r_{\mathcal{F}}(p)$ is the algebra of all germs f_p coming from a continuous function f on a neighborhood U of p such that $f_q \in \mathcal{G}(q)$ for all $q \in U$. In particular, we have then $f_p \in \mathcal{G}(p)$, which immediately yields $G^r_{\mathcal{F}}(p) \subseteq \mathcal{G}(p)$.

For the other inclusion, take an $f_p \in \mathcal{G}(p) \subseteq G_p(M)$. Since $f_p \in G_p(M)$, there is a neighborhood U of p and a function $f \in C(U)$ such that f_p is represented by f, incidentally justifying our choice of notation for f_p. At this point we make use of the fact that the load $\mathcal{G}$ is a differentiable one, meaning the loaded charts of $\mathcal{G}$ cover M. Hence there is also a loaded chart defined on a neighborhood of p, and we may assume that the above set U was chosen so small that the loaded chart restricts to a smaller loaded chart $\varphi\colon U \to A$. Then $f_p \in \mathcal{G}(p)$ implies $(f \circ \varphi^{-1})_x \in C^r_x(\mathbb{R}^n)$ for $x = \varphi(p)$, so $f \circ \varphi^{-1}$ restricts to a differentiable function defined on a neighborhood $A' \subseteq A$ of x. Writing $U' = \varphi^{-1}(A') \subseteq U$ for the corresponding neighborhood of p, we obtain a restricted loaded chart $\varphi\colon U' \to A'$. In order to prove $f_p \in G^r_{\mathcal{F}}(p)$, it suffices now to show that $f_q \in \mathcal{G}(q)$ for all $q \in U'$. But we know already that $(f \circ \varphi^{-1})_y \in C^r_y(\mathbb{R}^n)$ for all $y \in A'$ since $f \circ \varphi^{-1}$ is differentiable on A'. Therefore we obtain $f_q \in \mathcal{G}(q)$ for the corresponding $q = \varphi^{-1}(y)$ by φ, acting as a loaded isomorphism between $\mathcal{G}|_{U'}$ and $C^r_{\mathbb{R}^n}|_{A'}$. This ends the proof of the inclusion $\mathcal{G}(p) \subseteq G^r_{\mathcal{F}}(p)$. □

1.49 Theorem *For a topological space M, differentiable loads and differentiable sheaves are in bijective correspondence.*

Proof. Write C^r for the construction $\mathcal{G} \mapsto C^r_{\mathcal{G}}$ associating a differentiable sheaf to a given differentiable load, as before

$$C^r_{\mathcal{G}}(U) = \{f \in C(U) \mid \forall_{p \in U}\, f_p \in \mathcal{G}(U)\}.$$

Correspondingly, we write also G^r for the assignment $\mathcal{F} \mapsto G^r_{\mathcal{F}}$ extracting the differentiable load from a given differentiable sheaf, meaning

$$G^r_{\mathcal{F}}(p) = \{f_p \mid f \in \mathcal{F}(U) \wedge U \ni p\}.$$

If we denote the set of all differentiable loads on M by $\mathrm{Grm}(M)$ and the set of all differentiable sheaves on M by $\mathrm{Shf}(M)$, we can view this as two function $C^r\colon \mathrm{Grm}(M) \to \mathrm{Shf}(M)$ and $G^r\colon \mathrm{Shf}(M) \to \mathrm{Grm}(M)$. Proposition 1.48 states that $G^r \circ C^r = 1_{\mathrm{Grm}(M)}$ while Proposition 1.45 gives $C^r \circ G^r = 1_{\mathrm{Shf}(M)}$. Hence C^r and G^r are inverse to each other, and the sets $\mathrm{Shf}(M)$ and $\mathrm{Grm}(M)$ are indeed in bijective correspondence. □

By Theorem 1.49, we have now a fourth possibility of *defining a differentiable manifold* (the others being: atlas, patchwork, sheaf), with an appropriate characterization of differentiable maps given in Proposition 1.43. This framework is used in [36] to build up the theory of differentiable manifolds. There is one peculiar point in this approach that should however be noted: Even though one can dispense with charts in specifying a differentiable manifold (all we need is to specify at each point a subalgebra of the algebra of continuous germs), they are again needed (under the guise of "loaded charts") for characterizing which loads are differentiable—and hence make up a differentiable manifold.

We conclude this section by pointing out an important property of all except the analytic germs. So let M be a C^r manifold with $r \in \mathbb{N} \cup \{\infty\}$ and $p \in M$. Then every germ in $G^r_M(p)$ may be realized as f_p coming from a *global section* $f \in C^r(M)$. The proof [36[15]] uses a partition of unity for constructing a bump function that connects a local representative smoothly (in as much as required by the order r) to the rest of M; the proof fails for $r = \omega$ since the bump function cannot be made analytic. As a consequence, one obtains an obvious isomorphism of algebras

$$G^r(p) \cong C^r(M)/C^r_{[p]}(M),$$

where $C^r_{[p]}(M)$ is the ideal of all global sections vanishing around p.

1.4 Constructions on Manifolds

1.4.1 Submanifolds

Recall Definition 1.2 of an n-dimensional embedded submanifold M of $\mathbb{R}^m$. We have around every point a diffeomorphism—an ambient chart—that maps an open set of M to an open set of $\mathbb{R}^n$. From Example 1.19 we know that every diffeomorphism is a chart of the canonical differentiable structure on $\mathbb{R}^m$. This observation motivates the following notion of an embedded submanifold of an arbitrary manifold.

1.50 Definition *Let M be an m-dimensional manifold and N a subset of M. We say that N is an n-dimensional* embedded submanifold *of M, if for every point $p \in N$ there exists a chart $\varphi\colon U \to A$ around p such that $\varphi(U \cap N) = A \cap \mathbb{R}^n$. We call $m-n$ the* codimension *of N.*

To justify our terminology we have to show that N is a manifold in its own right. Indeed, we obtain an atlas for N analogously to the construction in Section 1.1.3, by restricting a chart to the submanifold and then projecting to the first n coordinates.

Let $\varphi\colon U \to A$ be a chart as in the above definition (an "ambient chart" in our earlier terminology). Let

$$\hat{U} = U \cap N \quad \text{and} \quad \hat{A} = \pi(A \cap \mathbb{R}^n)$$

where $\pi\colon \mathbb{R}^m \to \mathbb{R}^n$ denotes the projection $\pi(x_1, \dots, x_m) = (x_1, \dots, x_n)$. Then $\hat{A}$ is open in $\mathbb{R}^n$ and the map

$$\hat{\varphi} = \pi \circ \varphi|_{\hat{U}}\colon \hat{U} \to \hat{A}$$

is an ("abstract") chart for N. Since an embedded submanifold is a subset of M it is naturally a topological space with the induced topology. Analogously to Proposition 1.5 and 1.6, we obtain the following properties for the induced charts $\hat{\varphi}$.

1.51 Proposition *Every chart $\hat{\varphi}$ is a C^0 homeomorphism.*

1.52 Proposition *Every transition is a C^r diffeomorphism.*

So the induced charts constitute an atlas for the embedded submanifold, and the corresponding chart topology is the *induced topology* by Propositions 1.11. Note moreover that the induced topology is again Hausdorff and second-countable. It is also obvious that embedded submanifolds of $\mathbb{R}^m$ as in Definition 1.2 are embedded submanifolds of $\mathbb{R}^m$, the latter seen as a manifold with the canonical differentiable structure. Open subsets of of a manifold are 0-codimensional embedded submanifolds by Example 1.23.

We have the following characterization of differentiable maps for submanifolds.

1.53 Proposition *Let $\hat{M}$ be an embedded submanifold of M. Let $f\colon N \to M$ be a map from a manifold to N such that $f(N) \subseteq \hat{M}$, and consider a point $p \in N$. Then f is differentiable at p iff the induced map $\hat{f}\colon N \to \hat{M}$ is differentiable at p.*

Proof. Let f be differentiable at p. Let $\varphi\colon U \to A$ be a chart around $f(p)$ as in Definition 1.50. By Lemma 1.30, f is continuous at p, so there exists a neighborhood $\tilde{V}$ of p with $f(\tilde{V}) \subseteq U$. We can choose a chart domain $V \subseteq \tilde{V}$ with the corresponding chart $\psi\colon V \to B$ by Proposition 1.17, and since $f(N) \subseteq \hat{M}$ we also have $f(V) \subseteq U \cap \hat{M}$. The local representative $f_{\varphi\psi}$ is differentiable around $\psi(p)$, and therefore also the local representative

$$\hat{f}_{\hat{\varphi}\psi} = \pi \circ \varphi f \psi^{-1}.$$

for the induced chart $\hat{\varphi}$ is differentiable since π is C^ω. So the induced map $\hat{f}$ is indeed differentiable at p.

Suppose conversely that $\hat{f}$ is differentiable at p. Let $\varphi\colon U \to A$ be a chart around $f(p)$ as in Definition 1.50 and $\hat{\varphi}$ the corresponding induced chart. Since $\hat{f}$ is continuous at p, we can find as before a chart $\psi\colon V \to B$ around p such that $f(V) \subseteq U \cap \hat{M}$. The local representative $\hat{f}_{\hat{\varphi}\psi}$ is differentiable around $\psi(p)$, and therefore also the local representative

$$f_{\varphi\psi} = \iota \circ \hat{\varphi} \hat{f} \psi^{-1}.$$

for the chart φ is differentiable since the inclusion $\iota(x_1, \ldots, x_n) = (x_1, \ldots, x_n, 0, \ldots, 0)$ is C^ω. So the map f is differentiable at p, as claimed. □

Several authors consider also a more general notion of submanifolds often called immersed submanifolds, see for example [7[79]], [46[234]] and [2[75]]. We will discuss immersions and immersed submanifolds in a later section.

1.4.2 Products and Sums of Manifolds

Given an m-manifold M and and n-manifold N, we can canonically construct a differentiable structure on the cartesian product $M \times N$ such that it becomes an $(m+n)$-manifold.

Let $\mathfrak{A} = (\varphi_i \mid i \in I)$ and $\mathfrak{B} = (\psi_j \mid j \in J)$ be atlases on M and N, respectively. Then one sees that

$$\mathfrak{A} \times \mathfrak{B} = (\varphi_i \times \psi_j \mid i \in I, j \in J)$$

is an $(m+n)$-dimensional atlas on $M \times N$. The corresponding chart topology is the product topology, hence it is again Hausdorff and second-countable. We call $M \times N$ with the differentiable structure $(\mathfrak{A} \times \mathfrak{B})_{\max}$ the *product manifold* of M and N.

1.54 Example *A classical example of a product manifold is the* torus

$$S^1 \times S^1 \subseteq \mathbb{R}^2 \times \mathbb{R}^2.$$

To illustrate the torus, we usually use a diffeomorphic embedded submanifold in $\mathbb{R}^3$ *that looks like a doughnut.*

We know that a map $f = (f_1, f_2)\colon U \to \mathbb{R}^m \times \mathbb{R}^n$ for some open subset $U \subseteq \mathbb{R}^l$ is differentiable iff its components f_1 and f_2 are differentiable. Hence we obtain the following characterization of differentiable maps for product manifolds.

1.55 Proposition *Let L, M, N be manifolds and $f = (f_1, f_2)\colon L \to M \times N$. Then f is differentiable at p iff its components f_1 and f_2 are differentiable at p.*

Let now M and N be two n-dimensional manifolds. Then we can canonically define an n-dimensional atlas on the disjoint union $M + N$ by taking the union $\mathfrak{A} + \mathfrak{B}$ (automatically disjoint) of two atlases for M and N, respectively. This is indeed an atlas since we do not have any additional transition between charts. The corresponding chart topology is the sum topology, hence it is again Hausdorff and second-countable. We call $M + N$ with the differentiable structure $(\mathfrak{A} + \mathfrak{B})_{\max}$ the *sum* of the manifolds M and N.

Note that we can even construct the sum of countably many manifolds, but not uncountably many since we want to have second countability. Obviously, we obtain an analogous characterization of differentiable maps for the sum of manifolds with the canonical inclusions.

Chapter 2

The Tangent Space

Differential calculus on manifolds begins with the tangent space: If we want to *approximate* a differentiable map between two manifolds near a given point by a linear map (and this is what the calculus is all about!), we need a vector space "at" the given point: the tangent space. But how can we create a vector space structure on an abstract manifold? Obviously we cannot add their points—unless the manifold is embedded in some Euclidean space. But this is just what we wanted to avoid in the present setting. Hence we must seek other means of setting up a vector space structure at a point.

2.1 Cotangent and Tangent Vectors

2.1.1 Abstract Setting

We fix an n-dimensional manifold M of class C^r with $r \geq 1$, and choose a point $p \in M$. Let $\mathfrak{D}$ be the differentiable structure of M. The variables i, j range over $\{1, \ldots, n\}$.

We will introduce tangent and cotangent vectors as geometrical objects, with special emphasis on their duality. As a visual expression of this co/contra duality, portions of the text will be layed out in two columns.

We write $\mathcal{F}(p)$ for the set of *function germs* through p, consisting of all C^1 germs of maps $f\colon M \to \mathbb{R}$. From now on, the variable f and its embellishments range over $\mathcal{F}(p)$.	We write $\mathcal{C}(p)$ for the set of *curve germs* through p, consisting of all C^1 germs of maps $c\colon \mathbb{R} \to M$ with $c(0) = p$. From now on, the variable c and its embellishments range over $\mathcal{C}(p)$.

We can measure the rate of change of a function germ f along a curve germ c at the point p by the directional derivative

$$\langle f|c\rangle = (f \circ c)'(0).$$

We simply call $\langle f|c\rangle$ the *rate* of f along c. The notation $\langle|\rangle$ should remind of an inner product—and we shall soon understand why!

We call f and $\bar{f}$ cotangent to each other, denoted by $f \sim \bar{f}$, if

$$\langle f|c\rangle = \langle \bar{f}|c\rangle$$

for all c. If in particular $f \sim 0$, we call the function germ f *stationary*.
The cotangent space is given by $T_p^*M = \mathcal{F}(p)/{\sim}$, whose elements are accordingly known as the *cotangent vectors* (or the "covectors") at p. The variables d, e and their embellishments range over T_p^*.

We call c and $\bar{c}$ tangent to each other, denoted by $c \sim \bar{c}$, if

$$\langle f|c\rangle = \langle f|\bar{c}\rangle$$

for all f. If in particular $c \sim 0$, we call the curve germ c *singular*.
The tangent space is given by $T_pM = \mathcal{C}(p)/{\sim}$, whose elements are accordingly known as the *tangent vectors* (or the "vectors") at p. The variables v, w and their embellishments range over T_p.

Since we have fixed the manifold M from the outset, we will subsequently abbreviate T_p^*M and T_pM by T_p^* and T_p, respectively. Note that our *terminology and notation* anticipates already some results that we will prove presently: Both T_p^* and T_p are n-dimensional vector spaces, and T_p^* is indeed dual to T_p.

In order to establish these dimensions, we will have to go back to the defining property of the manifold: Locally, it must be diffeomorphic to $\mathbb{R}^n$ via any *local chart*. Hence it is reasonable to expect that we can reduce the "testing germs" (function germs / curve germs) in the equivalence relation (cotangency / tangency) to those coming from a single local chart (coordinate functions / coordinate lines). In order to make for a smoother notation, we extend the bracket to centered charts by writing respectively

$$\begin{aligned}\langle f|\varphi^{-1}\rangle &= (f\circ\varphi^{-1})'(0) = \langle f|\varphi_i\rangle\,\delta^i \in \mathbb{R}_n,\\ \langle \varphi|c\rangle &= (\varphi\circ c)'(0) = \langle \varphi^i|c\rangle\,\delta_i \in \mathbb{R}^n\end{aligned}$$

for the rate of a chart (meaning its coordinate functions) along a curve and the rate of a function along a parametrization (meaning its coordinate curves).

2.1 Lemma *We have $f \sim \bar{f}$ iff $\langle f|\varphi^{-1}\rangle = \langle \bar{f}|\varphi^{-1}\rangle$ and $c \sim \bar{c}$ iff $\langle\varphi|c\rangle = \langle\varphi|\bar{c}\rangle$, where φ is some chart centered at p.*

Proof. The direction from left to right is obvious. For the other direction, we can compute $\langle f|c\rangle$ as

$$(f\circ c)'(0) = ((f\circ\varphi^{-1})\circ(\varphi\circ c))'(0) = (f\circ\varphi^{-1})'(0)\,(\varphi\circ c)'(0) = \langle f|\varphi^{-1}\rangle\,\langle\varphi|c\rangle,$$

which obviously coincides with $\langle f|\bar{c}\rangle$ or $\langle\bar{f}|c\rangle$ whenever either $\langle f|\varphi^{-1}\rangle = \langle\bar{f}|\varphi^{-1}\rangle$ or $\langle\varphi|c\rangle = \langle\varphi|\bar{c}\rangle$, respectively. □

In order to address T_p^* and T_p as vector spaces, we have to make sure that there is a linear structure on them. For T_p^*, this is obviously the case already, and for T_p, we define the *linear structure* by putting

$$\lambda'[c']_\sim + \lambda''[c'']_\sim = \{c \in \mathcal{C}(p) \mid \forall_{f\in\mathcal{F}(p)}\ \langle f|c\rangle = \lambda'\,\langle f|c'\rangle + \lambda''\,\langle f|c''\rangle\} \tag{2.1}$$

for all $\lambda', \lambda'' \in \mathbb{R}$ and $c', c'' \in \mathcal{C}(p)$. Note that the linear structure is well-defined since the result is indeed an equivalence class under tangency, and the vector-space axioms are easily verified. Hence we can indeed view both T_p and T_p^* as vector spaces. Moreover, we can see that the rate factors through these spaces as a bilinear form

$$\begin{aligned} \langle | \rangle \colon \qquad T_p^* \times T_p &\to \mathbb{R}, \\ ([f]_\sim, [c]_\sim) &\mapsto \langle f | c \rangle. \end{aligned}$$

Note also that $\langle | \rangle \colon T_p^* \times T_p \to \mathbb{R}$ is non-degenerate, due to the definition of cotangency and tangency.

2.2 Lemma *For any chart φ centered at p, we have $\langle \varphi^i | \varphi_j \rangle = \delta^i_j$.*

Proof. We compute

$$\langle \varphi^i | \varphi_j \rangle = (\varphi^i \circ \varphi_j)'(0) = ((\delta^i \circ \varphi) \circ (\varphi^{-1} \circ \delta_j))'(0) = (\delta^i \circ \delta_j)'(0) = \delta_i^j,$$

the latter identity since $\delta^i \circ \delta_j \colon \mathbb{R} \to \mathbb{R}$ acts by $t \mapsto \delta_i^j\, t$. □

Our introduction of the cotangent and tangent space follows along the lines of [35], but it can actually be done in a *general construction* similar to [32[145,523]] and in the spirit of [43[1]]. We will briefly outline this idea because it gives us more insight into the symmetry between the cotangent and tangent space. Suppose we have a "bisurjective" function $\langle | \rangle \colon F \times C \to K$ with values in a field K, meaning $\langle f | \_ \rangle$ and $\langle \_ | c \rangle$ are surjective for all nonstationary f and nonsingular c. Here one defines the relations of cotangency and tangency in a completely analogous manner, and f is again called stationary if $f \sim 0$ and c singular if $c \sim 0$. One can then proceed to impose a K-vector space structure on $\tilde{F} = F/{\sim}$ and $\tilde{C} = C/{\sim}$ by defining linear combinations on $\tilde{C}$ as we did in (2.1), and dually also on $\tilde{F}$. This makes $\langle | \rangle$ a non-degenerate bilinear form on $\tilde{F} \times \tilde{C}$.

Bisurjectivity can be guaranteed by having a *scalar multiplication* on either C or F. Taking F for definiteness, this means we have a map $\cdot \colon K \times F \to K$ such that $\langle \lambda \cdot f | c \rangle = \lambda \langle f | c \rangle$ for all $\lambda \in K$ and $(f, c) \in F \times C$. This is what we used in our example where we had $F = \mathcal{F}(p)$, $C = \mathcal{C}(p)$ and $K = \mathbb{R}$.

But in general, one will not obtain the dualities $\tilde{F} \cong \tilde{C}^*$ and $\tilde{C} \cong \tilde{F}^*$ from this bilinear form—in fact, this happens iff we know that one (and therefore also the other) of these spaces is finite-dimensional. *Banach manifolds*, as treated e.g. in [31] and [11], are a case in point here: The tangent and cotangent space (in the sense of our definition) are no longer dual to each other since each of them is of the same (small) dimension as the underlying Banach space playing the role of $\mathbb{R}^n$. Hence there are now four spaces involved: tangent, cotangent, tangent dual (usually called "cotangent"!), and cotangent dual.

In the finite-dimensional case of our cotangent and tangent spaces, we have a *special situation* that introduces an apparent asymmetry into the construction: There is already a linear structure on F, but not on C. But when we look at the geometric picture and our actual definitions, we realize that the linear structure on $\mathbb{R}$ is never really used—neither for curves $c \colon \mathbb{R} \to M$ nor for functions $f \colon M \to \mathbb{R}$. The crucial point is just that the *rate* of f along c produces an element of $\mathbb{R}$; and it is from them that we obtain linear structures on the equivalence classes (and we could introduce them in the symmetric fashion of the abstract construction sketched above).

In fact, we could consider *generalized curves* parametrized by an arbitrary set and *generalized functions* with values in an arbitrary set (think of replacing $\mathbb{R}$ by a suitable algebraic curve); all we need is a way of measuring their rate of change! Another possibility would be to keep the curves and functions as they are, defining a *generalized rate* e.g. by $\langle f|g\rangle = (f \circ \Psi \circ c)'(0)$ with an arbitrary differentiable map $\Psi\colon M \to M$ fixing p.

2.1.2 Representation through Components

In order to arrive at the desired dimension and duality results, we will set up an isomorphism between the (co)tangent space and the real vector space. This isomorphism will depend on the choice of any chart φ centered at p; it maps the abstract (co)tangent vectors to their *components* in the chosen coordinate system. Note that the component isomorphisms below are well-defined because of Lemma 2.1.

We define *cotangent component isomorphism* by

$$\begin{aligned} (|_\varphi\colon \quad T_p^* &\to \mathbb{R}_n, \\ [f]_\sim &\mapsto \langle f|\varphi^{-1}\rangle \end{aligned}$$

and the *cotangent representation isomorphism* by

$$\begin{aligned} \langle|_\varphi\colon \quad \mathbb{R}_n &\to T_p^*, \\ a &\mapsto [a \circ \varphi]_\sim. \end{aligned}$$

As usual, we view the rows $a \in \mathbb{R}_n$ here as linear functions $\mathbb{R}^n \to \mathbb{R}$.

We define the *tangent component isomorphism* by

$$\begin{aligned} |)_\varphi\colon \quad T_p &\to \mathbb{R}^n, \\ [c]_\sim &\mapsto \langle \varphi|c\rangle \end{aligned}$$

and the *tangent representation isomorphism* by

$$\begin{aligned} |\rangle_\varphi\colon \quad \mathbb{R}^n &\to T_p, \\ h &\mapsto [\varphi^{-1} \circ h]_\sim. \end{aligned}$$

As usual, we view the columns $h \in \mathbb{R}^n$ here as linear curves $\mathbb{R} \to \mathbb{R}^n$.

As mentioned before, we can also write the component isomorphisms in the more *explicit forms*

$$[f]_\sim \mapsto (\langle f|\varphi_1\rangle, \ldots, \langle f|\varphi_n\rangle) \quad \text{and} \quad [c]_\sim \mapsto (\langle \varphi^1|c\rangle, \ldots, \langle \varphi^n|c\rangle)^\top. \tag{2.2}$$

Our immediate aim is now to put the component isomorphisms into use: We prove that they establish the desired isomorphism between the (co)tangent spaces and their Euclidean counterparts (thus justifying their "isomorphism" labels).

2.3 Proposition *The component and representation isomorphisms are inverse to each other as linear maps. Hence $T_p^* \cong \mathbb{R}_n$ and $T_p \cong \mathbb{R}^n$, implying in particular $\dim T_p^* = \dim T_p = n$.*

Proof. On the vector space $\mathcal{F}(p)$, we can see the cotangency relation as a linear congruence that corresponds to the subspace of stationary germs

$$\mathcal{S}_p = \{f \mid \forall_{c \in \mathcal{C}(p)}\, \langle f|c\rangle = 0\} = \{f \mid \langle f|\varphi^{-1}\rangle\},$$

where the latter identity follows from Lemma 2.1. Observe now that the gradient operator

$$\gamma_\varphi \colon \mathcal{F}(p) \to \mathbb{R}_n,\ f \mapsto \langle f|\varphi^{-1}\rangle$$

is as a linear map that factors through the cotangency projection to $(|_\varphi \colon T_p^* \to \mathbb{R}_n$. Since $T_p^* = \mathcal{F}(p)/\!\sim\, = \mathcal{F}(p)/\mathcal{S}_p$, the homomorphism theorem implies that $(|_\varphi$ is a linear map, while Lemma 2.2 tells us that γ_φ and hence $(|_\varphi$ is surjective. Moreover, $(|_\varphi$ is injective since $\operatorname{Ker}\gamma_\varphi = \mathcal{S}_p$. Hence $(|_\varphi$ is indeed a linear isomorphism.

Let us now check that $\langle|_\varphi$ is the inverse of $(|_\varphi$, which immediately guarantees that $\langle|_\varphi$ is also a linear map (and thus a linear isomorphism as claimed). For a linear isomorphism, any left or right inverse is automatically its (unique) inverse. So it suffices to check that $\langle|_\varphi$ is, say, a left inverse of $(|_\varphi$. For $d = [f]_\sim \in T_p^*$ we obtain

$$\langle(d|_\varphi|_\varphi = \langle\gamma_\varphi(f)|_\varphi = [\gamma_\varphi(f)\circ\varphi]_\sim = [f]_\sim = d$$

if we can show $f \sim \gamma_\varphi(f)\circ\varphi$. Using Lemma 2.1, this is equivalent to

$$\gamma_\varphi(f) = ((\gamma_\varphi(f)\circ\varphi)\circ\varphi^{-1})'(0),$$

which is evidently true since the right-hand side simplifies to $\gamma_\varphi(f)'(0) = \gamma_\varphi(f)$.

The tangency relation is defined on the set $\mathcal{C}(p)$, and it is obviously compatible with the tangent operator

$$\tau_\varphi \colon \mathcal{C}(p) \to \mathbb{R},\ c \mapsto \langle\varphi|c\rangle$$

Hence τ_φ factors through the tangency projection to the map $|)_\varphi \colon T_p \to \mathbb{R}$. We must now check that the map $|)_\varphi$ is linear on $T_p = \mathcal{C}(p)/\!\sim$ with respect to the linear structure we imposed there. Taking $c', c'' \in \mathcal{C}(p)$ and $\lambda', \lambda'' \in \mathbb{R}$, we pick out any representative $c \in \mathcal{C}(p)$ of the equivalence class $[c]_\sim = \lambda'\,[c']_\sim + \lambda''\,[c'']_\sim$, computing $\tau_\varphi(c) = (\varphi\circ c)'(0)$ as

$$\begin{pmatrix}\langle\varphi^1|c\rangle\\ \vdots\\ \langle\varphi^n|c\rangle\end{pmatrix} = \begin{pmatrix}\lambda'\,\langle\varphi^1|c'\rangle + \lambda''\,\langle\varphi^1|c''\rangle\\ \vdots\\ \lambda'\,\langle\varphi^n|c'\rangle + \lambda''\,\langle\varphi^n|c''\rangle\end{pmatrix} = \lambda'\begin{pmatrix}\langle\varphi^1|c'\rangle\\ \vdots\\ \langle\varphi^n|c'\rangle\end{pmatrix} + \lambda''\begin{pmatrix}\langle\varphi^1|c''\rangle\\ \vdots\\ \langle\varphi^n|c''\rangle\end{pmatrix},$$

where the first equality uses the linear structure on T_p and the second the one on $\mathbb{R}^n$. Using now $|[c]_\sim)_\varphi = \tau_\varphi(c)$ and the analogous relations for c' and c'', this yields indeed $|[c]_\sim)_\varphi = \lambda'\,|[c']_\sim)_\varphi + \lambda''\,|[c'']_\sim)_\varphi$ for $[c]_\sim = \lambda'\,[c']_\sim + \lambda''\,[c'']_\sim$, so $|)_\varphi \colon T_p \to \mathbb{R}^n$ is a linear map. Since $\tau_\varphi(c) = (\langle\varphi^1|c\rangle, \ldots, \langle\varphi^n|c\rangle)^\top$, we see from Lemma 2.2 that τ_φ and hence $|)_\varphi$ is surjective. It is injective and thus a linear isomorphism because the equivalence kernel is just the tangency relation.

It remains to be proved now that $|)_\varphi$ has $|\rangle_\varphi$ as its, say, left inverse. For $v = [c]_\sim \in T_p$ we have

$$||v)_\varphi\rangle_\varphi = |\tau_\varphi(c)\rangle_\varphi = [\varphi^{-1}\circ\tau_\varphi(c)]_\sim = [c]_\sim = v,$$

provided we can prove $c \sim \varphi^{-1}\circ\tau_\varphi(c)$. Again employing Lemma 2.1, this is equivalent to

$$(\varphi\circ c)'(0) = (\varphi\circ(\varphi^{-1}\circ\tau_\varphi(c)))'(0)$$

which is satisfied since the right-hand side simplifies to $\tau_\varphi(c)'(0) = \tau_\varphi(c)$. □

We can now prove the announced *duality* of cotangent and tangent space (thus justifying the terminology).

2.4 Proposition *The cotangent and tangent spaces form a dual pair of vector spaces* $(T_p^*, \langle|\rangle, T_p)$, *hence we will identify* $(T_p)^*$ *with* T_p^*, *and* $(T_p^*)^*$ *with* T_p.

Proof. Each cotangent vector $d \in T_p^*$ determines a map $d^*\colon T_p \to \mathbb{R}$ given by $v \mapsto \langle d|v\rangle$; since this map is obviously linear, we have indeed $d^* \in (T_p)^*$. Analogously, a tangent vector $v \in T_p$ determines a map $v^*\colon T_p^* \to \mathbb{R}$ given by $d \mapsto \langle d|v\rangle$; this map is again linear, so we have also $v^* \in (T_p^*)^*$.

The maps $d \mapsto d^*$ and $v \mapsto v^*$ are obviously linear, and the construction of T_p and T_p^* ensures their injectivity. Injective linear maps between equidimensional vector spaces are automatically surjective (alternatively we could apply Lemma 2.2), so we have indeed $(T_p)^* \cong T_p^*$ and $(T_p^*)^* \cong T_p$. □

For each centered chart φ around p, an "abstract" tangent vector $v \in T_p$ is now represented by the "concrete" component vector $h = |v)_\varphi \in \mathbb{R}^n$. Choosing another chart ψ centered at p, we get a different component vector $k = |v)_\psi \in \mathbb{R}^n$. There must be a tight relationship between these two component vectors, though, since we have the composite isomorphism

$$\mathbb{R}^n \overset{|\rangle_\varphi}{\rightsquigarrow} T_p \overset{|)_\psi}{\rightsquigarrow} \mathbb{R}^n \tag{2.3}$$

effecting $h \mapsto k$. Similarly, if a cotangent vector $d \in T_p^*$ is represented by component forms $a = (d|_\varphi \in \mathbb{R}_n$ and $b = (d|_\psi \in \mathbb{R}_n$, the composite isomorphism

$$\mathbb{R}_n \overset{\langle|_\varphi}{\rightsquigarrow} T_p^* \overset{(|_\psi}{\rightsquigarrow} \mathbb{R}_n \tag{2.4}$$

realizes $a \mapsto b$. It turns out that we can characterize these isomorphisms without recourse to the component and representation isomorphisms (recall that an asterisk on a linear map $f\colon \mathbb{R}^n \to \mathbb{R}^n$ denotes its dual $f^*\colon \mathbb{R}_n \to \mathbb{R}_n$ which in canonical bases operates by multiplying vectors from the left instead of the right).

2.5 Proposition *The linear isomorphisms of* (2.3) *and* (2.4) *are given by the vector differential* $d_0\, 1_{\psi\varphi}\colon \mathbb{R}^n \to \mathbb{R}^n$ *and codifferential* $d_0^*\, 1_{\varphi\psi}\colon \mathbb{R}_n \to \mathbb{R}_n$, *respectively.*

Proof. For $h \in \mathbb{R}^n$, we compute

$$||h\rangle_\varphi)_\psi = |[\varphi^{-1}h]_\sim)_\psi = (\psi \circ \varphi^{-1}h)'(0) = (\psi\varphi^{-1} \circ h)'(0) = (\psi\varphi^{-1})'(0)\, h'(0),$$

and this is indeed the same as $(d_0\, 1_{\psi\varphi})\, h$, written in canonical coordinates of $\mathbb{R}^n$, since $\psi\varphi^{-1} = 1_{\psi\varphi}$ and $h'(0) = h$. With $a \in \mathbb{R}_n$, we perform a similar computation

$$(\langle h|_\varphi|_\psi = ([a\varphi]_\sim|_\psi = (a\varphi \circ \psi^{-1})'(0) = (a \circ \varphi\psi^{-1})'(0) = a'(0)\, (\varphi\psi^{-1})'(0),$$

which equals this time $(d_0^*\, 1_{\varphi\psi})\, a$, written in the canonical coordinates of $\mathbb{R}_n$, since $\varphi\psi^{-1} = 1_{\varphi\psi}$ and $a'(0) = a$. □

Note that Proposition 2.5 is very useful since it allows us to change between different component vectors without "actually knowing" the abstract cotangent/tangent spaces: Given the components in one coordinate system, we can compute its components in any other coordinate system by virtue of the transition Jacobian. This is crucial because we never "actually know" the abstract spaces (whose elements are equivalence classes consisting of infinitely many germs)! In this sense, the component isomorphisms translate the abstract entities of the original definition (sometimes called the "geometric" view) into a simple isomorphic object that is amenable to *practical computations* (also called the "physical" view since computations are crucial in physical applications).

We should understand, however, that cotangent and tangent vectors are not simply elements of $\mathbb{R}_n$ and $\mathbb{R}^n$, respectively. Of course, if we are only speaking about "vectors" (meaning elements of an algebraic structure that satisfies the vector space axioms), there is nothing more to it, and even the distinction between $\mathbb{R}_n$ and $\mathbb{R}^n$ disappears. But a cotangent/tangent vector *of* a certain manifold M *at* a certain point $p \in M$ is more than that: it must *behave* in the way function/curve germs behave (of course modulo cotangency/tangency). And exactly how they must behave is specified in Proposition 2.5; we will come back to this question at the end of the next subsection.

2.1.3 Co- and Contravariance

In order to understand this behavior in some depth, let us restate the transformation formulae of Proposition 2.5 in a more classical notation: In this subsection we will use x, y, z for coordinate systems (which is the same as "charts" in our earlier terminology). Writing $\partial\bar{x}/\partial x = 1'_{\bar{x}x}(0)$ for the transition Jacobian from x to $\bar{x}$, its inverse $\partial x/\partial\bar{x} = 1'_{x\bar{x}}(0)$ is transition Jacobian from $\bar{x}$ to x by (1.1), and the transformation formulae can be expressed by

$$(d|_{\bar{x}} = (d|_x \,\frac{\partial x}{\partial \bar{x}} \quad \text{and} \quad |v)_{\bar{x}} = \frac{\partial \bar{x}}{\partial x}\,|v)_x. \tag{2.5}$$

We may view these formulae as giving a generalization of the Linear Algebra notion of *base change*, as explained in Subsection 2.2.3.

In the so-called *Ricci calculus*, relations of the form (2.5) are written entirely in components with lower/upper indices for extracting covector/vector components (recall that we think of lower/upper indices as ranging along rows/columns). As mentioned in Chapter 0, summation over diagonal index pairs is always implicit ("Einstein summation convention") in order to avoid excessive summation signs. In order to further simplify our notation, we leave out the component isomorphisms when using indices. This is harmless since d_j makes sense only if we think of d as $(d|_x \in \mathbb{R}_n$ and v^j only if we now interpret v as $|v)_x \in \mathbb{R}^n$. Of course, this notation loses the information about the coordinate systems in use, but this is usually solved by embellishments on d and v: If we use charts x and $\bar{x}$, the corresponding components would be called (d_j) and $(\bar{d}_i)$ for a cotangent vector d or (v^j) and $(\bar{v}^i)$ for a tangent vector v. With these conventions, the transformation laws can succinctly be written as

$$\bar{d}_i = \frac{\partial x^j}{\partial \bar{x}_i}\, d_j \quad \text{and} \quad \bar{v}^i = \frac{\partial \bar{x}^i}{\partial x_j}\, v^j, \tag{2.6}$$

all indices (both free and bound) ranging over $\{1, \ldots, n\}$. This index convention is another important standard ingredient of the Ricci calculus. Note also that in the transition Jacobian, the upper index appears in the numerator and the lower index in the denominator, just as one would expect.

Another crucial aspect is the *expansion* of (co)tangent vectors in the certain natural bases, called (co)frames.

The *coframe* for a chart x around p is given by the n cotangent vectors

$$[x^1]_\sim, \ldots, [x^n]_\sim \in T_p^*,$$

abbreviated as $dx^1, \ldots, dx^n$.

The *frame* for a chart x around p is given by the n tangent vectors

$$[x_1]_\sim, \ldots, [x_n]_\sim \in T_p,$$

abbreviated as $dx_1, \ldots, dx_n$.

The significance of the (co)frames is that they are all that one needs: Having chosen a fixed coordinate system x, they result from the cotangent representation isomorphism $\langle|_x \colon \mathbb{R}_n \to T_p^*$ as $\delta^i \mapsto dx^i$ and from the tangent representation isomorphism $|\rangle_x \colon \mathbb{R}^n \to T_p$ as $\delta_i \mapsto dx_i$. Of course, one could choose different (co)bases in T_p^* and T_p. But somehow this would be pointless because any such "unnatural" basis corresponds to a (co)frame for a different coordinate system $\bar{x}$, and it is much more convenient to work with this (co)frame; below we will derive the corresponding transformation laws. In view of these facts, we will henceforth use *only (co)frames*.

The *explicit component formulae* (2.2) can now be restated in terms of (co)frames: For a cotangent vector $d = [f]_\sim \in T_p^*$ and tangent vector $v = [c]_\sim \in T_p$, the components can be extracted by

$$d_i = \langle f|x_i\rangle = \langle d|dx_i\rangle \quad \text{and} \quad v^i = \langle x^i|c\rangle = \langle dx^i|v\rangle \tag{2.7}$$

since the rate factors through the canonical projections $\mathcal{F}(p) \to T_p^*$ and $\mathcal{C}(p) \to T_p$.

From Lemma 2.2 we infer immediately that $\langle dx^i|dx_j\rangle = \delta^i_j$, which implies that both the dx^i and the dx_j are linearly independent. In order to see that they also span the respective spaces (and hence constitute mutually dual bases for the cotangent/tangent space), we claim the *expansions*

$$d = d_i\, dx^i \quad \text{and} \quad v = v^i\, dx_i, \tag{2.8}$$

now making use of the conventions explained above.

2.6 Proposition *For any cotangent vector $d \in T_p^*$ and tangent vector $v \in T_p$ we have the expansions* (2.8) *and the formula*

$$\langle d|v\rangle = (d|_x\, |v)_x = d_i\, v^i \tag{2.9}$$

for the rate of d along v.

Proof. Choose arbitrary representatives to write $d = [f]_\sim$ and $v = [c]_\sim$. We start by proving the formula (2.9) for the rate

$$\langle d|v\rangle = \langle f|c\rangle = (f \circ c)'(0) = (fx^{-1} \circ xc)'(0) = (fx^{-1})'(0)\,(xc)'(0) = (d|_x|v)_x,$$

unfolding the definition of the component isomorphism in the last step. The second equality of (2.9) is just a restatement using the component convention introduced above. Note also that we can write the rate as $(d|\,|v)$ if we omit reference to the chart x. Running together the components, this could then be written as $(d|v)$ and parsed as the canonical scalar product of $(d_1, \ldots, d_n)$ and $(v^1, \ldots, v^n)^\top$.

In order to prove the identities (2.8), it suffices to verify the relations $f \sim d_i\, x^i$ and $c \sim v^i\, x_i$. Choosing an arbitrary $\bar{c} \in \mathcal{C}(p)$ with associated $\bar{v} = [\bar{c}]_\sim \in T_p$, we compute

$$\langle f|\bar{c}\rangle = \langle d|\bar{v}\rangle = d_i\, \bar{v}^i = \langle d_i\, x^i|\bar{c}\rangle$$

by the explicit component formulae (2.7). Dually, we take $\bar{f} \in \mathcal{F}(p)$ with $\bar{d} = [\bar{f}]_\sim \in T_p^*$ to compute

$$\langle \bar{f}|c\rangle = \langle \bar{d}|v\rangle = \bar{d}_i v^i = \langle \bar{f}|v^i x_i\rangle,$$

again using (2.7). □

Using frame and coframe, we can also express the *transition Jacobian* $\partial\bar{x}/\partial x = 1'_{\bar{x}x}(0) = (\bar{x}x^{-1})'(0)$ for the coordinate change from x to $\bar{x}$. Its i-th row is obviously $(\bar{x}^i x^{-1})'(0) = (d\bar{x}^i|_x$ by the definition of the cotangent component map. Using the explicit formulae (2.2), this yields

$$\frac{\partial \bar{x}}{\partial x} = \begin{pmatrix} \langle \bar{x}^1|x_1\rangle & \cdots & \langle \bar{x}^1|x_n\rangle \\ \vdots & \ddots & \vdots \\ \langle \bar{x}^n|x_1\rangle & \cdots & \langle \bar{x}^n|x_n\rangle \end{pmatrix}$$

for the Jacobian matrix and $\partial\bar{x}^i/\partial x_j = \langle \bar{x}^i|x_j\rangle$ for its generic entry.

Having bases in T_p^* and T_p, we can also ask how they change when we pass from one coordinate system to another. We have already derived the transformation formulae (2.5) for the components, and from Linear Algebra we would expect those for the *basis change* to be contragredient to them. This is indeed the case as we shall see now.

2.7 Proposition *Passing from a coordinate system x to $\bar{x}$, the coframe $dx^1, \ldots, dx^n$ and the frame $dx_1, \ldots, dx_n$ change to $d\bar{x}^1, \ldots, d\bar{x}^n$ and $d\bar{x}_1, \ldots, d\bar{x}_n$ such that the transformation laws*

$$d\bar{x}^i = \frac{\partial \bar{x}^i}{\partial x_j}\, dx^j \quad \textit{and} \quad d\bar{x}_i = \frac{\partial x^j}{\partial \bar{x}_i}\, dx_j \tag{2.10}$$

are satisfied.

Proof. For expressing "new" coframe vectors $d\bar{x}^i$ and frame vectors $d\bar{x}_i$ in terms of their "old" cognates dx^i and dx_i, we expand them in the old (co)frame, computing the corresponding component vectors according to formula (2.5). The latter are given by

$$|d\bar{x}^i)_x = |dx^i)_x\, \frac{\partial \bar{x}}{\partial x} \quad \text{and} \quad (d\bar{x}_i|_x = \frac{\partial x}{\partial \bar{x}}\, (dx_i|_x.$$

Since $|dx^i)_x = \delta^i$ and $(dx_i|_x = \delta_i$, this yields

$$(d\bar{x}^i)_j = \frac{\partial \bar{x}^i}{\partial x_j} \quad \text{and} \quad (d\bar{x}_i)^j = \frac{\partial x^j}{\partial \bar{x}_i}$$

for the components (all referring to the coordinate system x). Inserting them now into the expansions, we obtain the claimed relations

$$d\bar{x}^i = (d\bar{x}^i)_j \, dx^j = \frac{\partial \bar{x}^i}{\partial x_j} \, dx^j \quad \text{and} \quad d\bar{x}_i = (d\bar{x}_i)^j \, dx_j = \frac{\partial x^j}{\partial \bar{x}_i} \, dx_j$$

for the new (co)frame vectors. □

The formulae (2.10) for the basis change are clearly *contragredient* to the formulae (2.6) describing the coefficient transformation. In detail this means the following: The matrix describing the transformation of the coframe is contragredient (i.e. inverse-transpose) to the matrix describing the corresponding coefficient transformation for cotangent vectors, and also the transformation matrix of the frame is contragredient to the coefficient transformation matrix for tangent vectors. The combined effect of contragredience is that cotangent and tangent vectors remain invariant under all coordinate transformations (which we could have used as an alternative starting point for deriving the formulae of basis change). From the abstract point of view, this is of course trivial; let us nevertheless see how this manifests in components: Taking a cotangent vector $d = d_i \, dx^i \in T_p^*$ and passing to a new coordinate system $\bar{x}$, we regain the original cotangent vector since

$$\bar{d}_i \, d\bar{x}^i = \Big(\frac{\partial x^j}{\partial \bar{x}_i} \, d_j\Big)\Big(\frac{\partial \bar{x}^i}{\partial x_l} \, dx^l\Big) = \delta_l^j \, d_j \, dx^l = d_j \, dx^j = d;$$

similarly, a tangent vector $v = v^i \, dx_i \in T_p$ is replicated as we see from

$$\bar{v}^i \, d\bar{x}_i = \Big(\frac{\partial \bar{x}^i}{\partial x_j} \, v^j\Big)\Big(\frac{\partial x^l}{\partial \bar{x}_i} \, dx_l\Big) = \delta_j^l \, v^j \, dx_l = v^j \, dx_j = v.$$

This shows how the contragredient transformation of the frames and coefficients combines to keep cotangent/tangent vectors invariant.

The classification of the transformation laws seems to create considerable confusion as one can see from [52]. We will try to clarify the issue from a geometric point of view. The key to understanding is to consider a uniform expansion by a factor $\kappa > 1$. This can be expressed by passing from the coordinate system x to a new coordinate system $\bar{x}$ such that $x(p) = \kappa \, \bar{x}(p)$. (Note that in $\bar{x}$ all coordinates are smaller by κ^{-1}, thus making everything appear larger by κ. In other words, measuring in smaller units leads to larger results!) The transition Jacobians are then given by $\partial x / \partial \bar{x} = \kappa$ and $\partial \bar{x} / \partial x = \kappa^{-1}$. Hence everything that transforms with $\partial x / \partial \bar{x}$, expands along *with* the coordinate space: its scale varies in harmony or *covariantly.* Everything that transforms with $\partial \bar{x} / \partial x = \kappa^{-1}$, however, contracts *against* the coordinate space: its scale varies in opposition or *contravariantly.* We must now distinguish three different object types in discussing "how they change under coordinate transformations":

Vectors: As we have seen, cotangent and tangent vectors are *invariant*. This is just the point in defining "geometric objects" in the first place: they are independent of any coordinate systems, and that is why we have already introduced them in a coordinate-free manner. (In another terminology, such objects—along with points, curves, functions and the like—are called intrinsic, as opposed to extrinsic ones like charts and everything coming through them—in particular, bases and components.)

Bases: From Proposition 2.7 we can see that the cobasis transforms contravariantly, whereas the basis transforms covariantly. This should be geometrically obvious: Expanding the space, the basis vectors grow along with it, but the cobasis vectors (think of evenly spaced level planes) shrink.

Components: Equation (2.6) tells us that the cotangent components transform covariantly, while the tangent components transform contravariantly. We can understand this geometrically as follows: Since the basis vectors are now bigger, fewer are needed for building up a composite vector, but more level planes fit in a reference volume.

We can now see the source of *confusion* in the terminology: Tangent vectors are sometimes called contravariant (because their components are) but at other times covariant (since their bases are), while they are in fact invariant! In the same spirit, cotangent vectors are frequently called covariant vectors (after their components) and at times also contravariant (like their cobases), but they are of course again invariant! Therefore we do well to remain with the terminology "cotangent vector" versus "tangent vector", reserving the variance terms for the corresponding transformations of bases and components (since only they are actually transformed).

A last word on the terminology (in case this is not already obvious): The conventions imposed by the Ricci calculus on the *position of indices* is chosen in such a way that all superscripts (hence row indices) mark contravariant objects while subscripts (hence column indices) mark covariant objects. The whole calculus is marvellously consistent in this respect, and that can be a great help for keeping order in one's calculations.

At this point one could have the impression that the distinction between covariant and contravariant transformations is completely artificial since the "actual geometric objects" are anyway invariant. But there is really more to it: Even though a cotangent vector remains invariant under a coordinate transformation, we know that its components transform covariantly—and not contravariantly or whatever! And similarly for the invariant tangent vectors, whose components transform contravariantly—and not covariantly or whatever! The way in which the components transform turns out to be characteristic for (co)tangent vectors—this is the *traditional physicists' view* of cotangent and tangent vectors: "Anything that transforms covariantly is a cotangent vector. Anything that transforms contravariantly is a tangent vector".

Here "anything" means any vector (in a finite-dimensional vector space). Up to isomorphism, every vector space is given by $\mathbb{R}^n$, but it turns out to be more convenient to use $\mathbb{R}_n$ in the first case. So we can rephrase the above statement as: "The cotangent space is $\mathbb{R}_n$ with a covariant transformation law. The tangent space is $\mathbb{R}^n$ with a contravariant transformation law." We will refer to the former as the *covariant component space* and to the latter as the *contravariant component space*; their elements will accordingly be called co-/contravariant component vectors.

The idea of characterizing geometric objects by their *transformation laws* is a very fruitful one. (More precisely, we should say that one specifies geometric objects by some representations that are subject to certain constraints on how they transform under coordinate changes). This was at the heart of Felix Klein's "Erlangen program" of classifying geometries through their invariance properties. In his geometry book [26], we can find numerous transformation laws at prominent places, for all kinds of geometric objects (of course including our cotangent and tangent vectors, albeit under different guises).

Let us now make this precise. For *imposing a transformation law*, we associate a row/column of components to each coordinate system in such a way that the component rows/columns of different coordinate systems are linked to each other by the appropriate transformation law. (According to Proposition 2.5, this means they come from the same "abstract" cotangent or tangent vector.)

The *covariant component space* at the point p consists of all maps

$$D\colon \mathfrak{D}_{p\bullet} \to \mathbb{R}_n$$

with the property

$$D_{\bar{x}} = D_x \frac{\partial x}{\partial \bar{x}}$$

for all $x, \bar{x} \in \mathfrak{D}_{p\bullet}$.

The *contravariant component space* at the point p consists of all maps

$$V\colon \mathfrak{D}_{p\bullet} \to \mathbb{R}^n$$

with the property

$$V_{\bar{x}} = \frac{\partial \bar{x}}{\partial x} V_x$$

for all $x, \bar{x} \in \mathfrak{D}_{p\bullet}$.

Note that using $\mathbb{R}_n$ instead of $\mathbb{R}^n$ for defining the covariant component space is not essential. The point is just that the transformation law in the definition of $\equiv$ looks nicer by using rows for the component vectors. If we insisted on *columns*, the transition Jacobian would have to be transposed, and the component vectors would appear to the right of it.

Our definition of the component spaces follows [22]. Other authors [12][38] use the following *alternative but equivalent definition*: The covariant component space is given by $(\mathfrak{D}_{p\bullet} \times \mathbb{R}_n)/{\equiv}$ with an equivalence $\equiv$ defined by $(x, a) \equiv (\bar{x}, \bar{a})$ iff $\bar{a} = a\,(\partial x/\partial \bar{x})$; the contravariant component space by $(\mathfrak{D}_{p\bullet} \times \mathbb{R}^n)/{\equiv}$ with another equivalence defined by $(x, h) \equiv (\bar{x}, \bar{h})$ iff $\bar{h} = (\partial \bar{x}/\partial x)\,h$.

The *connection* is the following: In the original definition, a co- or contravariant component vector is specified by a certain map, in the alternative definition by its (partitioned) graph. Conversely, one can obtain maps from the equivalence class because the latter are actually maps in disguise (which suggests that the definition by maps is more natural): Whenever both (x, a) and $(x, \bar{a})$ are in the same equivalence class of $(\mathfrak{D}_{p\bullet} \times \mathbb{R}_n)/{\equiv}$, we obtain $\bar{a} = a\partial x/\partial x = a$; similarly, (x, h) and $(x, \bar{h})$ being in a single equivalence class of $(\mathfrak{D}_{p\bullet} \times \mathbb{R}^n)/{\equiv}$ implies $\bar{h} = \partial x/\partial x\; h = h$. Hence we see that the two definitions of component spaces are indeed equivalent.

2.8 Theorem *The covariant/contravariant component space at the point p is isomorphic to the cotangent/tangent space at p.*

Proof. Fixing $d \in T_p^*$ and $v \in T_p$, the component isomorphisms induce a covariant component vector $x \mapsto (d|_x$ and a contravariant component vector $x \mapsto |v)_x$. Thus we obtain a map ι^* from

T_p^* to the covariant component space at p and a map ι from T_p to the contravariant component space at p.

These maps are bijective since every covariant component vector D and every contravariant component vector V gives rise to abstract counterparts in the following way: Picking any coordinate system $x \in \mathfrak{D}_{p\bullet}$, we obtain the corresponding $(\iota^*)^{-1}(D) = \langle D_x|_x \in T_p^*$ and $\iota^{-1}(V) = |V_x\rangle_x \in T_p$. Proposition 2.3 implies immediately that $(\iota^*)^{-1}$ and ι^* as well as ι^{-1} and ι are actually linear isomorphisms and inverse to each other.

The only thing left to check is that $(\iota^*)^{-1}$ and ι^{-1} are well-defined. Picking another coordinate system $\bar{x} \in \mathfrak{D}_{p\bullet}$, we must verify that $\langle D_x|_x = \langle D_{\bar{x}}|_{\bar{x}}$ and $|V_x\rangle_x = |V_{\bar{x}}\rangle_{\bar{x}}$. Set $d = \langle D_x|_x$ and $v = |V_x\rangle_x$. Then we obtain

$$\langle D_{\bar{x}}|_{\bar{x}} = \langle D_x \frac{\partial x}{\partial \bar{x}}|_{\bar{x}} = \langle (d|_x \frac{\partial x}{\partial \bar{x}}|_{\bar{x}} = \langle (d|_{\bar{x}}|_{\bar{x}} = d$$

and

$$|V_{\bar{x}}\rangle_{\bar{x}} = |\frac{\partial \bar{x}}{\partial x} V_x\rangle_{\bar{x}} = |\frac{\partial \bar{x}}{\partial x} |v)_x\rangle_{\bar{x}} = ||v)_{\bar{x}}\rangle_{\bar{x}} = v$$

from the definition of the co- and contravariant component isomorphism, the representation/component isomorphisms of Proposition 2.3, the component transformation laws (2.5), and the representation/component isomorphisms again (this time reversed). □

The characterization of T_p^* and T_p described in Theorem 2.8 is also useful for defining the (co)tangent space in a *patchwork* directly (i.e. without first constructing a differentiable atlas and then using our earlier definition); actually, this is where the idea of the component spaces has apparently come from [38]. The point is that we just need the transition Jacobians—and these are of course available in a patchwork as the derivatives of the gluing diffeomorphisms.

Since the (co)tangent vectors are now attached to equivalence classes of glued points in $\mathbb{R}^n$, it is now more advantageous to apply the above-mentioned alternative definition: The equivalences for gluing points and for identifying components can be combined into a *single equivalence* defined on triples made up of label/coordinates/components. (The label identifies the coordinate system in use, the coordinates specify the point, and the components refer to the actual (co)tangent vector at this point with respect to this coordinate system.) Including the point actually anticipates a construction that we will come across later: the (co)tangent bundle.

2.1.4 Coderivations and Derivations

Up to now, we have treated cotangent and tangent vectors in such a way that we can perceive the symmetry between them in a clear light. In some sense, however, this symmetry is not perfect: Although one *need not* use any asymmetric features for buildung up the theory (and we have indeed avoided them in our treatment), one *can* also take advantage of such asymmetric features. In this subsection, we will give a brief outline of such an approach. Let M again be a fixed n-manifold and choose a point $p \in M$.

The main point is that the set $\mathcal{F}(p)$ is an algebra whereas $\mathcal{C}(p)$ is not. We can add functions, but we *cannot add curves*. This also why it was trivial to introduce a vector space structure on $T_p^* = \mathcal{F}(p)/{\sim}$ while the definition (2.1) of linear combinations is much more indirect for $T_p = \mathcal{C}(p)/\sim$. In fact, it amounts to taking a detour via the derivations, as we shall soon realize.

A *coderivation* at $p \in M$ is a map $$\eta\colon \mathcal{C}(p) \to \mathbb{R}$$ that is compatible with the tangency relation $\sim$ and has a linear quotient map $\eta_\sim\colon T_p \to \mathbb{R}$. We write $\Delta_p^* M$ or briefly Δ_p^* for the vector space of all coderviations at p, and we use η, ζ and its embellishments for elements of Δ_p^*.	A *derivation* at $p \in M$ is a map $$\delta\colon \mathcal{F}(p) \to \mathbb{R}$$ that is compatible with the cotangency relation $\sim$ and has a linear quotient map $\delta_\sim\colon T_p^* \to \mathbb{R}$. We write $\Delta_p M$ or briefly Δ_p for the vector space of all derivations at p, and we use δ, ε and its embellishments for elements of Δ_p.

Obviously every cotangent vector $d = [f]_\sim$ gives rise to a coderivation $\eta\colon c \mapsto \langle f|c\rangle$, and every tangent vector $v = [c]_\sim$ to a derivation $\delta\colon f \mapsto \langle f|c\rangle$. We call η the *coderivation onto* d, written as $\langle f|\_\_\rangle$ or $\langle d|\_\_\rangle$, while δ is called the *derivation along* v, denoted analogously by $\langle \_\_|c\rangle$ or $\langle \_\_|v\rangle$. (Imagine $\langle d|\_\_\rangle$ as a velocity component projected onto f and $\langle \_\_|v\rangle$ as a directional derivative directed along v.)

Passing from a (co)derivation to its (co)tangent vectors is achieved —as we shall see immediately—by the quotient maps $\eta_\sim$ and $\delta_\sim$ occurring in the definition above. We call $\eta_\sim$ and $\delta_\sim$ the *direction* of η and δ, respectively. They can be directly characterized by suitable equivalence classes.

2.9 Lemma *We have*

$$\eta_\sim = \{f \in \mathcal{F}(p) \mid \forall_{c\in\mathcal{C}(p)}\ \langle f|c\rangle = \eta(c)\}$$

and

$$\delta_\sim = \{c \in \mathcal{C}(p) \mid \forall_{f\in\mathcal{F}(p)}\ \langle f|c\rangle = \delta(f)\}$$

for every $\eta \in \Delta_p^*$ and $\delta \in \Delta_p$.

Proof. By definition of Δ_p^* and Δ_p, we obtain from η and δ yield linear forms $\eta_\sim\colon T_p \to \mathbb{R}$ and $\delta_\sim\colon T_p^* \to \mathbb{R}$. Using the identifications of Proposition 2.4, this means $\eta_\sim \in T_p^*$ and $\delta_\sim \in T_p$. For every $v = [c]_\sim \in T_p$ and $d = [f]_\sim \in T_p^*$ we have

$$\eta_\sim(v) = \eta(c) = \langle \{f \in \mathcal{F}(p) \mid \forall_{c\in\mathcal{C}(p)}\ \langle f|c\rangle = \eta(c)\}|v\rangle$$

and

$$\delta_\sim(d) = \eta(f) = \langle \{c \in \mathcal{C}(p) \mid \forall_{f\in\mathcal{F}(p)}\ \langle f|c\rangle = \delta(c)\}|d\rangle,$$

so the claim follows by the duality of Proposition 2.4. $\square$

2.10 Proposition *The maps $d \mapsto \langle d|\_\_\rangle$ and $\eta \mapsto \eta_\sim$ are inverse to each other, defining a linear isomorphism $T_p^* \cong \Delta_p^*$. Similarly, the maps $v \mapsto \langle \_\_|v\rangle$ and $\delta \mapsto \delta_\sim$ are also inverse to each other, defining a linear isomorphism $T_p \cong \Delta_p$.*

Proof. By the bilinearity of $\langle|\rangle$, the maps $d \mapsto \langle d|\_\_\rangle$ and $v \mapsto \langle \_\_|v\rangle$ are linear. Being projections onto quotients modulo $\sim$, the maps $\eta \mapsto \eta_\sim$ and $\delta \mapsto \delta_\sim$ are linear as well. It follows immediately that $\langle d|\_\_\rangle_\sim = d$ and $\langle \_\_|v\rangle_\sim = v$ for all $d \in T_p^*$ and $v \in T_p$ as well as $\langle \eta_\sim|\_\_\rangle = \eta$ and $\langle \_\_|\delta_\sim\rangle = \delta$ for all $\eta \in \Delta_p^*$ and $\delta \in \Delta_p$. $\square$

Now we can understand the definition (2.1) of *linear combinations of tangent vectors* in a new light: Paraphrasing it as $\lambda' v' + \lambda'' v'' = (\lambda' \langle \_\_|v'\rangle + \lambda'' \langle \_\_|v''\rangle)_\sim$, its effect is to push the linear structure from Δ_p to T_p. As mentioned after Lemma 2.2, one could do the same for the cotangent space—but the point is that one need not!

Next we consider *component view* of (co)derivations. By the above proposition, it is clear that $\langle dx^i|\_\_\rangle$ and $\langle \_\_|dx_i\rangle$ constitute bases for Δ_p^* and Δ_p, which we shall denote respectively by ∂x^i and ∂x_i or even—if the coordinate system is assumed to be known—by ∂^i and ∂_i. Using these bases, a coderivation $\eta = \eta_i \, \partial x^i$ operates on $c \in \mathcal{C}(p)$ by

$$\eta(c) = \eta_i \, \partial^i c = \eta_i \, \langle dx^i|c\rangle = \eta_i \, (x^i \circ c)'(0) = \eta_i \, \frac{\partial c^i}{\partial t}(0)$$

and a derivation $\delta = \delta^i \, \partial x_i$ on $f \in \mathcal{F}(p)$ by

$$\delta(f) = \delta^i \, \partial_i f = \delta^i \, \langle f|dx_i\rangle = \delta^i (f \circ x_i)'(0) = \delta^i \, \frac{\partial f}{\partial x_i}(0).$$

Roughly speaking, a coderivation is a vector differential operator on curves and a derivation a partial differential operator on functions. (Note the last expression in both lines, exemplifying some common "abuse of notation": One writes c^i for the i-th coordinate projection of c, and one reuses the name x_i of the i-th coordinate function for the corresponding formal differentiation parameter of f.)

Up to this point, out treatment was completely symmetric between coderivations and derivations. From now on, however, we will consider some *special features of derivations* since the coderivations are apparently of a rather meagre structural substance (this is why they are not explicitly introduced in the standard texts on manifolds). Our first remark serves to justify the term "derivation".

2.11 Definition *A* derivation *on $G^r(p)$ is a linear functional δ satisfying the Leibniz law in the form $\delta(fg) = f(p)\,\delta(g) + g(p)\,\delta(f)$.*

This notion of derivation can be seen as an instance of the *general concept of derivations* $\delta\colon A \to M$ from an R-algebra A into an A-module M; see for example $[32^{746}]$. In this sense, a derivation δ is an R-linear map satisfying the Leibniz law in its "classical" form $\delta(fg) = f\,\delta(g) + g\,\delta(f)$. In our case, the ring R is given by the reals $\mathbb{R}$, the algebra A by the germs $G^r(p)$, and the module M again by $G^r(p)$, but now equipped with the module operation $A \times M \to M$ defined by $(f, g) \mapsto f(p)\, g$.

Let us also remark that one obtains a *differential algebra* if one takes $M = A$, viewed as a module over itself. A typical case would be $C^\infty(I)$, the smooth functions on some open interval $I \subseteq \mathbb{R}$ with the usual differentiation as a derivation (and the product rule as the Leibniz law).

Since $\delta(1) = \delta(1 \cdot 1) = 1 \cdot \delta(1) + \delta(1) \cdot 1$, it is clear that δ *vanishes on constants*. Every derivation $\delta \colon \mathcal{F}(p) \to \mathbb{R}$ has a canonical extension $\tilde{\delta} \colon G^r(p) \to \mathbb{R}$ via $\tilde{\delta}(f) = \delta(f - f(0))$ for all $f \in G^r(p)$. Therefore we will from now on identify the derivations on $G^r(p)$ with those on $\mathcal{F}(p)$. Now for the announced result on derivations.

2.12 Proposition *A linear functional δ on $G^r(p)$ is a derivation iff $\delta|_{\mathcal{F}(p)}$ factors through the canonical projection $\lambda \colon \mathcal{F}(p) \to \mathcal{F}(p)/\mathcal{F}(p)^2$. Then we have in particular $\delta\left(\mathcal{F}(p)^2\right) = 0$.*

Proof. For defining the quotient map $\tilde{\delta} \colon \mathcal{F}(p)/\mathcal{F}(p)^2 \to \mathbb{R}$ we have to make sure that $\delta(\operatorname{Ker} \lambda) = O$. Taking $f \in \mathcal{F}(p)$ with $\lambda f = 0$, there are $f_1, f_2 \in \mathcal{F}(p)$ with $f = f_1 f_2$. But then $\delta(f) = f_1(p)\, \delta f_2 + f_2(p)\, \delta f_1 = 0$ since $f_1(p) = f_2(p) = 0$, so $\delta|_{\mathcal{F}(p)} = \tilde{\delta} \circ \lambda$.

Conversely, assume $\delta|_{\mathcal{F}(p)} = \tilde{\delta} \circ \lambda$ for a linear map $\tilde{\delta} \colon \mathcal{F}(p)/\mathcal{F}(p)^2 \to \mathbb{R}$ and take $f, g \in G^r(p)$. Using the linearity of $\tilde{\delta}$ and λ, we obtain

$$\begin{aligned} \delta(fg) = \tilde{\delta}\lambda\big(fg - f(p)\, g(p)\big) &= \tilde{\delta}\lambda\left(f - f(p)\right)\big(g - g(p)\big) \\ &\quad + f(p)\, \tilde{\delta}\lambda\big(g - g(p)\big) + g(p)\, \tilde{\delta}\lambda\big(f - f(p)\big), \end{aligned}$$

where the first summand of the right-hand side vanishes since $(f - f(p)(g - g(p)) \in \mathcal{F}(p)^2$. The Leibniz rule now follows since

$$\delta(fg) = f(p)\, \tilde{\delta}\lambda\big(g - g(p)\big) + g(p)\, \tilde{\delta}\lambda\big(f - f(p)\big) = f(p)\, \delta(g) + g(p)\, \delta(f),$$

so δ is indeed a derivation on $G^r(p)$. □

We can now prove that the linear functionals of Δ_p are indeed derivations on $G^r(p)$. In fact, we can even prove a little bit more. For that, let us call a derivation C^r *compatible* if it annihilates all stationary function germs.

2.13 Corollary *Every $\delta \in \Delta_p$ is a C^r compatible derivation on $G^r(p)$.*

Proof. By definition, every $\delta \in \Delta_p$ is compatible with the cotangency relation; in particular, we have $\delta(f) = 0$ for all stationary function germs $f \in \mathcal{F}(p)$, so δ is C^r compatible. But every $f \in \mathcal{F}(p)^2$ is *a fortiori* stationary since $f = f_1 f_2$ with $f_1, f_2 \in \mathcal{F}(p)$ implies

$$\langle f|c\rangle = (f_1 f_2 \circ c)'(0) = ((f_1 \circ c)(f_2 \circ c))'(0) = f_1(0)\, \langle f_2|c\rangle + f_2(0)\, \langle f_1|c\rangle = 0$$

for all $c \in \mathcal{C}(p)$. Therefore $\delta(\operatorname{Ker} \lambda) = O$ and $\delta|_{\mathcal{F}(p)}$ factors through λ, so Proposition 2.12 ensures that δ is a derivation on $G^r(p)$. □

So every $\delta \in \Delta_p$ is a derivation on $G^r(p)$, but one should not think that Δ_p provides *all derivations*. This is not true for $r < \infty$; see the remarks in [55[14]]. For the cases $r < \infty$, one needs the additional hypothesis of C^r compatibility [41[56]], which can also be achieved implicitly [7[76]] by considering only those derivations that can be written as $\langle \_ | c \rangle$ for some $c \in \mathcal{C}(p)$.

2.14 Proposition *The algebra of C^r-compatible derivations on $G^r(p)$ coincides with Δ_p.*

Proof. By Corollary 2.13, every $\delta \in \Delta_p$ is a C^r compatible derivation on $G^r(p)$. Conversely, let δ be a C^r compatible derivation on $G^r(p)$. Then δ is a linear functional that is also compatible with the cotangency relation, and Proposition 2.12 yields a factorization $\delta|_{\mathcal{F}(p)} = \tilde{\delta} \circ \lambda$ with a linear map $\tilde{\delta}\colon \mathcal{F}(p)/\mathcal{F}(p)^2 \to \mathbb{R}$. This means that $\delta \in \Delta_p$. □

In order to understand the special case of C^∞ manifolds, we have to characterize the *structure of stationary function germs* [41[58]] in a suitable way. Assuming $r > 0$, we write now more explictly $\mathcal{F}^r(p)$ for the r-times differentiable function germs vanishing at p.

2.15 Proposition *Every stationary $f \in \mathcal{F}^r(p)$ can be written as $f = g_1h_1 + \ldots + g_nh_n$ with $g_1, \ldots, g_n, h_1, \ldots, h_n \in \mathcal{F}^{r-1}(p)$.*

Proof. Using a ball chart, we can reduce the statement to the case $M = \mathbb{R}^n$ and $p = 0$. By the Fundamental Theorem of Calculus, we have

$$f(x) = \int_0^1 \frac{d}{dt} f(tx)\,dt = \sum_{i=0}^{n} x_i g_i(x) \qquad \text{with} \qquad g_i(x) = \int_0^1 \frac{\partial f}{\partial x_j}(tx)\,dt,$$

and $g_i \in G^{r-1}(p)$. Since f is stationary, we have also $g_i(0) = \partial f/\partial x_j(0) = 0$, so $g_i \in \mathcal{F}^{r-1}(p)$. Setting $h_i(x) = x_i$, this yields the representation $f = g_1h_1 + \ldots + g_nh_n$. □

2.16 Corollary *Every derivation on $G^\infty(p)$ is C^∞ compatible.*

Proof. Take a derivation δ on $G^\infty(p)$ and a stationary $f \in G^\infty(p)$. By Proposition 2.15, we can write $f = g_1h_1 + \ldots + g_nh_n$ for some $g_1, \ldots, g_n, h_1, \ldots, h_n \in \mathcal{F}^\infty(p)$. But then $\delta(f) = 0$ follows since $\delta(g_ih_i) = 0$ by the Leibniz law. □

This result explains the *special status of smooth manifolds*: We do not need any additional compatibility condition, so the tangent vectors can be identified with *all* derivations. Restricting themselves from the outset to smooth manifolds (clearly the most important for appliations), some textbooks therefore introduce the tangent space simply as the algebra of all C^∞ derivations.

One can see this point clearly for the *Euclidean manifold* $M = \mathbb{R}^n$. In this case, it is clear that all tangent spaces V can be identified canonically (see Subsection 2.2.3 for more details). Choosing a V-basis $e_1, \ldots, e_n$ for V, its dual V^*-basis consists of the coordinate functionals $x_1, \ldots, x_n$. (Note that for once we do not adhere to the usual convention of index placement.) The bidual V^{**} consists of linear functionals on V^*, and the V^{**}-basis dual to $x_1, \dots, x_n$ are just the partial derivatives $\partial_1, \dots, \partial_n$. Hence V^{**} may be described as the space of linear partial differential operators with constant coefficients. In these bases, the *canonical isomorphism* $V \to V^{**}$ is realized by $e_i \mapsto \partial_i$, and we have $\partial_i(f) = f(e_i)$ for every $f \in V^*$. Hence the dual pairing is either realized by evaluation of linear functionals as

$$\begin{aligned} V^* \times V &\to \mathbb{R} \\ (f, v) &\mapsto f(v) \end{aligned}$$

or by differentiation of linear functionals as

$$\begin{aligned} V^* \times V^{**} &\to \mathbb{R} \\ (f, \partial) &\mapsto \partial(f). \end{aligned}$$

The isomorphism $V \to V^{**}$ is packed up in the isomorphism $T_pM \to \Delta_pM$ of Proposition 2.10.

Going back to the C^r case with $0 < r < \infty$, we provide an alternative characterization of Δ_p in terms of *lower-order derivations*. Choose any s with $0 < s < r$. Then every C^r compatible derivation on $G^r(p)$ can be extended to a C^s compatible derivation on $G^s(p)$, as pointed out in [41[58]]. The converse is also true, yielding the following characterization theorem [41[59]].

2.17 Proposition *We have $\delta \in \Delta_p$ iff $\delta = \tilde{\delta}|_{\mathcal{F}^r(p)}$ for some derivation $\tilde{\delta}$ on $G^s(p)$.*

Proof. Take a derivation $\tilde{\delta}$ on $G^s(p)$ and a stationary $f \in \mathcal{F}^r(p)$. Then by Proposition 2.15, we have $f = g_1h_1 + \ldots + g_nh_n$ for suitable $g_1, \dots, g_n, h_1, \dots, h_n \in \mathcal{F}^s(p)$. Since $\tilde{\delta}$ is a derivation on $G^s(p)$, we have $\tilde{\delta}(g_ih_i) = 0$ by the Leibniz law. Conversely, every $\delta \in \Delta_p$ represents a C^r compatible derivation $\delta\colon G^r(p) \to \mathbb{R}$ that can be extended to an (even C^s compatible) derivation $\tilde{\delta}\colon G^s(p) \to \mathbb{R}$ by the above remark. □

In concluding, we turn now briefly to the *cotangent space*. From the definition, it is immediately clear that $T_p^* = \mathcal{F}^r(p)/{\sim}$ is the quotient space of $\mathcal{F}(p)$ modulo the stationary function germs. Dually to the tangent space, this can be formulated in a "nicer" (more algebraic) way in the special case of smooth manifolds.

2.18 Proposition *A function germ $f \in \mathcal{F}^\infty(p)$ is stationary iff $f \in \mathcal{F}^\infty(p)^2$.*

Proof. From Proposition 2.15 we know that every stationary function germ is in $\mathcal{F}^\infty(p)^2$. Conversely, let $f \in \mathcal{F}^\infty(p)^2$. Then the Leibniz law implies $\delta(f) = 0$ for every derivation δ on $G^\infty(p)$. By Corollary 2.16, we have $\delta \in \Delta_p$ for every such derivation. By the definition of Δ_p and Lemma 2.9, this means $\langle f|c\rangle = 0$ for all $c \in \mathcal{C}(p)$ and hence $f \sim 0$. □

In view of this restult, one can define the *cotangent space of a smooth manifold* as $T_p^* = \mathcal{F}^\infty(p)/\mathcal{F}^\infty(p)^2$, and one can interpret the duality between T_p and T_p^* in the following way [55[13]]. Every cotangent vector represents a linear functional $\tilde{\delta}\colon \mathcal{F}^\infty(p)/\mathcal{F}^\infty(p)^2 \to \mathbb{R}$, which by Proposition 2.12 gives a derivation when composed with the canonical projection $\lambda\colon \mathcal{F}^\infty(p) \to \mathcal{F}^\infty(p)/\mathcal{F}^\infty(p)^2$, and vice versa. Hence in the smooth case, we have a natural isomorphism between Δ_p and $\mathcal{F}^\infty(p)/\mathcal{F}^\infty(p)^2$. For non-smooth manifolds this approach fails because $\mathcal{F}^r(p)/\mathcal{F}^r(p)^2$ is infinite-dimensional for $r < \infty$; see [55[14]].

The construction of this isomorphism can be generalized [32[748]][18[98]] [5[569]] to derivations $\delta\colon A \to M$ between an R-algebra A and an A-module M. In this case, it turns out that the A-module of derivations is isomorphic to the module of all A-homorphisms between the A-modules J/J^2 and M, where the ideal J is given as the kernel of the tensor product $A \otimes A \to A, a_1 \otimes a_2 \mapsto a_1 a_2$. The module J/J^2 is referred to as the *universal differential module* of A, and its elements are often called Kähler differentials [48[199]]. This construction is of some importance is algebraic geometry since it provides a kind of substitute for differential forms on manifolds (which we will introduce soon). In our case, $\mathcal{F}^\infty(p)$ plays the role of J, so we can think of the cotangent vectors in $\mathcal{F}^\infty(p)/\mathcal{F}^\infty(p)^2$ as the Kähler differentials at $p \in M$.

2.2 The Differential

Now that we have defined the tangent space at each point of a manifold, we can define the *linear approximation*—the differential—of a differentiable map between manifolds as a linear map between the tangent spaces. In the following, we discuss the definition of the differential and its basic properties. Let M and N be two manifolds as in Subsection 1.3.1, with a map $\Phi\colon M \to N$, differentiable at a point $p \in M$ mapping to $q = \Phi(p)$.

2.2.1 Abstract Setting

By the usual pushforward and pullback constructions, every differentiable curve $c\colon \mathbb{R} \to M$ through p is "pushed forward" to a differentiable curve $\Phi \circ c\colon \mathbb{R} \to N$ through q and every differentiable function $\bar{f}\colon N \to \mathbb{R}$ at q is "pulled back" to a differentiable function $\bar{f} \circ \Phi\colon M \to \mathbb{R}$ at p. The transformations $c \mapsto \Phi \circ c$ and $\bar{f} \mapsto f \circ \Phi$ obviously factor through the germ equivalences $\sim_p$ and $\sim_q$, yielding the corresponding *germ transports*

$$\begin{aligned} \mathcal{F}(q) &\to \mathcal{F}(p) \\ \bar{f}_q &\mapsto (\bar{f} \circ \Phi)_p \end{aligned} \qquad \text{and} \qquad \begin{aligned} \mathcal{C}(p) &\to \mathcal{C}(q) \\ c_p &\mapsto (\Phi \circ c)_q. \end{aligned}$$

For simplicity, we will suppress the germ subscripts p and q and use f for the function germs of $\mathcal{F}(p)$ and c for the curve germs of $\mathcal{C}(p)$, just as in Section 2.1. The corresponding objects in N will be denoted by bars, so $\bar{f}$ ranges over $\mathcal{F}(q)$ and $\bar{c}$ over $\mathcal{C}(q)$.

Now the (co)differential arises simply by *factoring the germ transports* through the (co)tangency relation as made explicit in the following definition.

The *codifferential* of Φ at the point p is defined as the map

$$\begin{aligned} T_q^*N &\to T_p^*M \\ [\bar{f}]_\sim &\mapsto [\bar{f}\circ\Phi]_\sim, \end{aligned}$$

denoted by $d_p^*\Phi$.

The *differential* of Φ at the point p is defined as the map

$$\begin{aligned} T_pM &\to T_qN \\ [c]_\sim &\mapsto [\Phi\circ c]_\sim, \end{aligned}$$

denoted by $d_p\Phi$.

2.19 Proposition *The codifferential $d_p^*\Phi\colon T_q^*N \to T_p^*M$ and the differential $d_p\Phi\colon T_pM \to T_qN$ are well-defined linear maps.*

Proof. For showing $d_p^*\Phi\colon T_q^*N \to T_p^*M$ well-defined, we take $\bar{f}_1, \bar{f}_2 \in \mathcal{F}(q)$ with $\bar{f}_1 \sim \bar{f}_2$ and prove $\bar{f}_1\circ\Phi \sim \bar{f}_2\circ\Phi$. Hence we fix $c \in \mathcal{C}(p)$ and show $\langle \bar{f}_1\circ\Phi|c\rangle = \langle \bar{f}_2\circ\Phi|c\rangle$, which is equivalent to $\langle \bar{f}_1|\Phi\circ c\rangle = \langle \bar{f}_2|\Phi\circ c\rangle$. But this follows from $\bar{f}_1 \sim \bar{f}_2$ since $\Phi\circ c \in \mathcal{C}(q)$. In a completely analogous fashion, one can also prove that $d_p\Phi\colon T_pM \to T_qN$ is well-defined.

The linearity of $d_p^*\Phi$ is immediate since

$$\begin{aligned} d_p^*\Phi(\lambda_1\,[\bar{f}_1]_\sim + \lambda_2\,[\bar{f}_2]_\sim) &= d_p^*\Phi\,[\lambda_1\bar{f}_1 + \lambda_2\bar{f}_2]_\sim = [(\lambda_1\bar{f}_1 + \lambda_2\bar{f}_2)\circ\Phi]_\sim \\ &= [\lambda_1\,\bar{f}_1\circ\Phi + \lambda_2\,\bar{f}_2\circ\Phi]_\sim = \lambda_1\,[\bar{f}_1\circ\Phi]_\sim + \lambda_2\,[\bar{f}_2\circ\Phi]_\sim = \lambda_1\,d_p^*\Phi\bar{f}_1 + \lambda_2\,d_p^*\Phi\bar{f}_2. \end{aligned}$$

In order to prove $d_p\Phi$ linear, however, we have to invest a little more work (because linear combinations are more involved for tangent vectors than for cotangent vectors). We must prove that

$$\begin{aligned} L = d_p\Phi(\lambda_1\,[c_1]_\sim + \lambda_2\,[c_2]_\sim) &= d_p\Phi\,\{c \in \mathcal{C}(p) \mid \forall_{f\in\mathcal{F}(p)}\,\langle f|c\rangle = \lambda_1\,\langle f|c_1\rangle + \lambda_2\,\langle f|c_2\rangle\} \\ &= [\Phi\circ c \mid c \in \mathcal{C}(p) \wedge \forall_{f\in\mathcal{F}(p)}\,\langle f|c\rangle = \lambda_1\,\langle f|c_1\rangle + \lambda_2\,\langle f|c_2\rangle]_\sim \end{aligned}$$

and

$$\begin{aligned} R = \lambda_1\,d_p\Phi[c_1]_\sim + \lambda_2\,d_p\Phi[c_2]_\sim &= \lambda_1\,[\Phi\circ c_1]_\sim + \lambda_2\,[\Phi\circ c_2]_\sim \\ &= \{\bar{c} \in \mathcal{C}(q) \mid \forall_{\bar{f}\in\mathcal{F}(q)}\,\langle \bar{f}|\bar{c}\rangle = \lambda_1\,\langle \bar{f}|\Phi\circ c_1\rangle + \lambda_2\,\langle \bar{f}|\Phi\circ c_2\rangle\} \end{aligned}$$

are the same. For showing $L \subseteq R$, we take $c \in \mathcal{C}(p)$ with $\langle f|c\rangle = \lambda_1\,\langle f|c_1\rangle + \lambda_2\,\langle f|c_2\rangle$ for all $f \in \mathcal{F}(p)$ and prove $\langle \bar{f}|\Phi\circ c\rangle = \lambda_1\,\langle \bar{f}|\Phi\circ c_1\rangle + \lambda_2\,\langle \bar{f}|\Phi\circ c_2\rangle$ for all $\bar{f} \in \mathcal{F}(q)$. But the latter is equivalent to $\langle \bar{f}\circ\Phi|c\rangle = \lambda_1\,\langle \bar{f}\circ\Phi|c_1\rangle + \lambda_2\,\langle \bar{f}\circ\Phi|c_2\rangle$, and this follows from the choice of c since $\bar{f}\circ\Phi \in \mathcal{F}(p)$. Now for the converse $L \supseteq R$, let $\bar{c} \in \mathcal{C}(q)$ be such that $\langle \bar{f}|\bar{c}\rangle = \lambda_1\,\langle \bar{f}|\Phi\circ c_1\rangle + \lambda_2\,\langle \bar{f}|\Phi\circ c_2\rangle$ for all $\bar{f} \in \mathcal{F}(q)$; we must find $c \in \mathcal{C}(p)$ with $\Phi\circ c \sim \bar{c}$ such that $\langle f|c\rangle = \lambda_1\,\langle f|c_1\rangle + \lambda_2\,\langle f|c_2\rangle$ for all $f \in \mathcal{F}(p)$. Choosing c by $[c] = \lambda_1\,[c_1]_\sim + \lambda_2\,[c_2]_\sim$, the last condition is trivially fulfilled, so it remains to prove $\langle \bar{f}|\Phi\circ c\rangle = \langle \bar{f}|\bar{c}\rangle$ for all $\bar{f} \in \mathcal{F}(q)$. By the definition of c, we can compute

$$\langle \bar{f}|\Phi\circ c\rangle = \langle \bar{f}\circ\Phi|c\rangle = \lambda_1\,\langle \bar{f}\circ\Phi|c_1\rangle + \lambda_2\,\langle \bar{f}\circ\Phi|c_2\rangle = \lambda_1\,\langle \bar{f}|\Phi\circ c_1\rangle + \lambda_2\,\langle \bar{f}|\Phi\circ c_2\rangle;$$

but this is $\langle \bar{f}|\bar{c}\rangle$ by the choice of $\bar{c}$. □

For emphasizing the symmetry of $\langle f \circ \Phi | c\rangle = \langle f | \Phi \circ c\rangle$, we may also introduce the *generalized rate* $\langle f | \Phi | c\rangle = (f \circ \Phi \circ c)'(0)$ for it. We can think of this expression in the following way: It measures the rate of change of the f-value of Φ along the c-direction.

As we may expect from the symmetry in the definition of the differential and the codifferential, the latter is the *dual* of the former. In other words, the categorical duality (pullback versus pushforward) is respected by the linear structure (operating on the dual versus primal space).

2.20 Proposition *The codifferential is the dual of the differential, $(d_p\Phi)^* = d_p^*\Phi$. In detail, this means*

$$\langle (d_p^*\Phi)\, \bar{d} | v\rangle = \langle \bar{d} | (d_p\Phi)\, v\rangle \tag{2.11}$$

for all $\bar{d} \in T_q^*N$ and $v \in T_pM$.

Proof. In the sense of the identification installed by Proposition 2.4, the identity $(d_p\Phi)^* = d_p^*\Phi$ is nothing more than a reformulation of the defining property (2.11), so we just have to prove the latter. For arbitrary $\bar{d} = [\bar{f}]_\sim \in T_q^*N$ and $v = [c]_\sim \in T_pM$, we obtain

$$\langle (d_p^*\Phi)\, \bar{d} | v\rangle = \langle \bar{f} | \Phi | c\rangle = \langle \bar{d} | (d_p\Phi)\, v\rangle$$

as claimed. □

The language of Subsection 2.1.4 provides another way of understanding this pullback/pushforward action. Let us use the explicit notation Φ^* for the pullback of Φ and Φ_* for its pushforward; this notation will be "globalized" in Subsection 2.3.4. Thus Φ^* is defined by $f \mapsto f \circ \Phi$ for any function f defined in a neighborhood of $q \in N$ while Φ_* operates on curves c through p via $c \mapsto \Phi \circ c$. Factoring through germ equivalence, this yields the maps $\Phi^* \colon \mathcal{F}(q) \to \mathcal{F}(p)$ and $\Phi_* \colon \mathcal{C}(p) \to \mathcal{C}(q)$. In the sense of Proposition 2.10, we may interpret the codifferential as a map $d_p^*\Phi \colon \Delta_q^*N \to \Delta_p^*M$ and accordingly the differential as a map $d_p\Phi \colon \Delta_pM \to \Delta_qN$. Then one sees immediately that $d_p^*\Phi$ acts by $\bar{\eta} \mapsto \bar{\eta} \circ \Phi_*$ and $d_p\Phi$ by $\delta \mapsto \delta \circ \Phi^*$.

As we can immediately read off from the definition, both the codifferential and the differential respects identity and composition in the sense of the generalized uniformity relation

$$d_p^*\, 1_M = 1_{T_p^*M} \quad \text{and} \quad d_p\, 1_M = 1_{T_pM}$$

and the generalized chain rule

$$d_p^*(\Psi \circ \Phi) = d_p^*\Phi \circ d_q^*\Psi \quad \text{and} \quad d_p(\Psi \circ \Phi) = d_q\Psi \circ d_p\Phi$$

if Ψ is a map from N to another manifold, differentiable at q. Hence we can view the codifferential and the differential respectively as a *cofunctor and functor*

$$d_{\dots}^* \colon \mathfrak{PtMan} \to \mathfrak{Vec} \quad \text{and} \quad d_{\dots} \colon \mathfrak{PtMan} \to \mathfrak{Vec}$$

from the category $\mathfrak{PtMan}$ of pointed manifolds to the category $\mathfrak{Vec}$ of vector spaces. (Following [13], we use "functor" as a synonym for "covariant functor" and "cofunctor" as a synonym for "contravariant functor".)

Let us make this a bit more explicit: A *pointed manifold* is a manifold with a distinguished point called the "base point". Clearly, a morphism Φ between pointed manifolds (M, p) and (N, q) is then given by a continuous map $\Phi\colon M \to N$ with $\Phi(p) = q$, differentiable at p. The functors $d^*_{\dots}$ and $d_{\dots}$ are thought to operate with the basepoint p substituted for the slots $_{\dots}$; they map a pointed manifold $(M, p) \in \mathfrak{PtMan}$ respectively to the vector spaces $T^*_p M \in \mathfrak{Vec}$ and $T_p M \in \mathfrak{Vec}$, and a morphism of pointed manifolds $\Phi\colon (M, p) \to (N, q)$ to the codifferential $d^*_p\Phi\colon T^*_q N \to T^*_p M$ and the differential $d_p\Phi\colon T_p M \to T_q N$.

2.2.2 Representation through Components

In Subsection 2.1.2 we have seen how to represent any cotangent and tangent vector by its list of *components*. These components depend on the choice of a coordinate system around the point p in question, and they refer to an expansion with respect to the (co)frames that we have introduced inSubsection 2.1.3.

Now the question is near at hand: Fixing a coordinate system x around p and $\bar{x}$ around q, how can we compute the *matrix representation* of the (co)differential with respect to the (co)frame in the (co)tangent spaces? The answer is strikingly simple: Extracting the local representative with respect to x and $\bar{x}$, its Jacobian operates from the left as the differential and from the right as the codifferential!

2.21 Proposition *Let x and $\bar{x}$ be coordinate systems around p and q, respectively. If $d = (d^*_p\Phi)\,\bar{d}$ and $\bar{v} = (d_p\Phi)\,v$ for a cotangent vector $\bar{d} \in T^*_q N$ and tangent vector $v \in T_p M$, we have the matrix relations*

$$(d|_x = (\bar{d}|_{\bar{x}}\, \Phi'_{\bar{x}x}(0) \quad \textit{and} \quad |\bar{v})_{\bar{x}} = \Phi'_{\bar{x}x}(0)\, |v)_x$$

for the corresponding component vectors (see the diagrams below), yielding

$$d_i = \frac{\partial \bar{x}^j(\Phi)}{\partial x_i}\, \bar{d}_j \quad \textit{and} \quad \bar{v}^i = \frac{\partial \bar{x}^i(\Phi)}{\partial x_j}\, v^j \tag{2.12}$$

when written in the Ricci calculus.

$$\begin{tikzcd} T^*_p M \arrow[d, "(|_x"'] & T^*_q N \arrow[l, "d^*_p\Phi"'] \arrow[d, "(|_{\bar{x}}"'] \\ \mathbb{R}_m & \mathbb{R}_n \arrow[l, "d^*_0\Phi_{\bar{x}x}"] \end{tikzcd} \qquad \begin{tikzcd} T_p M \arrow[r, "d_p\Phi"] \arrow[d, "|)_x"] & T_q N \arrow[d, "|)_{\bar{x}}"] \\ \mathbb{R}^m \arrow[r, "d_0\Phi_{\bar{x}x}"'] & \mathbb{R}^n \end{tikzcd}$$

Proof. Using the explicit component formulae (2.2), we can compute

$$d_i = \langle d|dx_i\rangle = \langle (d_p^*\Phi)\,\bar{d}|dx_i\rangle = \langle (d_p^*\Phi)\,\bar{d}_j\,d\bar{x}^j|dx_i\rangle = \langle \bar{x}^j \circ \Phi|x_i\rangle\,\bar{d}_j$$

for the cotangent components and analogously

$$\bar{v}^i = \langle d\bar{x}^i|\bar{v}\rangle = \langle d\bar{x}^i|(d_p\Phi)\,v\rangle = \langle d\bar{x}^i|(d_p\Phi)\,v^j dx_j\rangle = \langle \bar{x}^i|\Phi \circ x_j\rangle\,v^j$$

for the tangent components. But the Jacobian $\Phi'_{\bar{x}x}(0) = (\bar{x} \circ \Phi \circ x^{-1})'(0)$ obviously has

$$\langle \bar{x}^i|\Phi|x_j\rangle = \langle \bar{x}^i|\Phi \circ x_j\rangle = \langle \bar{x}^i \circ \Phi|x_j\rangle$$

for its (i, j) entry, hence we obtain (2.12). □

The transformation formulae (2.12) can be seen as a linear approximation of Φ in the row space $\mathbb{R}_n$ or column space $\mathbb{R}^n$. The notation $\partial\bar{x}\,(\Phi)/\partial x$ for the transformation Jacobian is quite similar to our notation for the *transition Jacobian* introduced in Subsection 2.1.3. In fact, the corresponding transition formulae (2.5) are a special case when we have one manifold $M = N$ and the identity map $\Phi = 1_M$. In this case, we suppress the latter in the notation for the Jacobian, just writing $\partial\bar{x}/\partial x$ then.

Sometimes one also *suppresses reference to the charts.* If the target chart $\bar{x}$ is inessential, one can write $\partial\Phi/\partial x$ as in [22]. If both the target chart $\bar{x}$ and the source chart x are clear from the context, the notatation $\partial\Phi$ can be used. (In components, these notations would be $\partial\Phi^i/\partial x_j$ and $\partial_j\Phi^i$, respectively.)

In the language of the *component spaces* explained in Subsection 2.1.3, the formulae (2.12) take on the form

$$D_x = \bar{D}_{\bar{x}}\,\frac{\partial\bar{x}(\Phi)}{\partial x} \quad \text{and} \quad \bar{V}_{\bar{x}} = \frac{\partial\bar{x}(\Phi)}{\partial x}\,V_x,$$

which should now be read as a definition: The codifferential is a transition-compatible linear map between the cotangent vectors $\bar{D}$ and D, the differential between the tangent vectors V and $\bar{V}$.

2.2.3 The Special Case of Vector Spaces

If the manifold in question is a *vector space* V, we know from Example 1.20 that each choice of basis $(b_1, \ldots, b_n)$ corresponds to a canonical atlas $\{x\}$ for V, where $x\colon V \to \mathbb{R}^n$ is the component chart [27[27]] , given by $v \mapsto (\lambda^1, \ldots, \lambda^n)$ for every vector $v = \lambda^i\, b_i$. In the special case $V = \mathbb{R}^n$, we have the canonical basis $(\delta_1, \ldots, \delta_n)$, corresponding to the component chart 1_V and the atlas $\{1_V\}$ treated in Example 1.19.

An "abstract" linear map $\Phi\colon V \to W$ between "abstract" vector spaces V and W can be described by its coordinate representation if we choose bases in V and W. If x and $\bar{x}$ are the corresponding component charts, the coordinate representation of Φ with respect to these bases is nothing else than the *local representative* $\Phi_{\bar{x}x}\colon \mathbb{R}^m \to \mathbb{R}^n$.

A change of basis in V corresponds to replacing x by another component chart $\bar{x}$. In Linear Algebra, this is usually described by a *transition matrix*, defined as the matrix of the identity map with respect to these bases [3[116]]. Hence it is just the transition Jacobian $\partial\bar{x}/\partial x = 1'_{\bar{x}x}(0)$, and the usual formula for "coordinate transformations" can be seen in the right-hand formula of (2.5); the left-hand formula of (2.5) describes the corresponding basis change in the dual space.

Let us now turn to *cotangent and tangent spaces*. Any component chart $x\colon V \to \mathbb{R}^n$ gives rise to the charts $x - x(p)$ centered at a point $p \in V$; these charts are again denoted by x if there is no ambiguity. Thus we get a coframe $dx^1, \ldots, dx^n$ for T_p^*V and a frame $dx_1, \ldots, dx_n$ for T_pV, and these spaces are all canonically isomorphic to each other (once we have fixed the component chart x). Observe that $dx^1, \ldots, dx^n$ are just the coordinate planes, while $dx_1, \ldots, dx_n$ are the coordinate axes. As opposed to general manifolds, we can thus identify all cotangent spaces with each other (and also all tangent spaces with each other).

For a *differentiable map* $\Phi\colon V \to W$ with $q = \Phi(p)$, these canonical isomorphisms allow us to interpret the differential $d_p\Phi\colon T_pV \to T_qW$ with respect to fixed bases (and thus component charts) in V and W as a linear map $d_p\Phi\colon \mathbb{R}^m \to \mathbb{R}^n$ and the codifferential $d_p^*\Phi\colon T_q^*W \to T_p^*V$ as a linear map $d_p^*\Phi\colon \mathbb{R}_n \to \mathbb{R}_m$. This is essentially what we did in Subsection 1.1.1, using the canonical bases in $V = \mathbb{R}^m$ and $W = \mathbb{R}^n$ corresponding to the component charts 1_V and 1_W, respectively.

If $\Phi\colon V \to W$ is a *linear map*, its codifferentials coincide for all points $p \in M$ (and of course also its differentials coincide for all points); hence we may suppress p in this case. We see here a generalized picture of the three levels described in Chapter 0: First the abstract map $\Phi\colon M \to N$ along with its dual $\Phi^*\colon N^* \to M^*$. Second their coordinate representations $d\Phi\colon \mathbb{R}^n \to \mathbb{R}^n$ and $d^*\Phi\colon \mathbb{R}_n \to \mathbb{R}_n$, which are identical with the local representatives of Φ and Φ^*, respectively. Third the Jacobians $\partial\Phi$ and $\partial\Phi^\top$ as the corresponding representation matrices.

Next we describe *hybrid cases*—those where either the domain or the codomain of a map is a vector space. If this vector space is just $\mathbb{R}$, we are led to functions and curves. For functions on $\mathbb{R}^n$ and curves in $\mathbb{R}^n$, the derivative has the intuitive meaning of cotangent vector and tangent vectors as explained in Chapter 0. We can now generalize this to functions on M and curves in M.

2.22 Definition *The* cotangent vector of a function $f\colon M \to \mathbb{R}$ *at a point* $p \in M$ *is the codifferential* d_p^*f. *The* tangent vector of a curve $c\colon \mathbb{R} \to M$ *at* p *is the differential* d_pc.

As a codifferential or differential, these concepts are linear functions and not just covectors and vectors as in the special case $M = \mathbb{R}^n$. But they can be naturally *identified with covectors and vectors* as we shall see now. For writing this succinctly, let us agree that $[f]_\sim$ for a function $f\colon M \to \mathbb{R}$ means $[\tilde{f}]_\sim$, where $\tilde{f}$ is the function germ of $f(\_\_) - \alpha$ for $\alpha = f(p)$; similarly, $[c]_\sim$ for a curve $c\colon \mathbb{R} \to M$ stands for $[\tilde{c}]_\sim$, where $\tilde{c}$ is the curve germ of $c(\_\_ + \tau)$ for $\tau \in \mathbb{R}$ such that $c(\tau) = p$. (These conventions reflect the intuition that differentiation is only concered with the local behavior and that we can furthermore discard constants.)

2.23 Proposition *There are natural isomorphisms carrying the cotangent vector of a function f at a point $p \in M$ to the covector $[f]_\sim \in T_p^*M$ and the tangent vector of a curve c at p to the vector $[c]_\sim \in T_pM$.*

Proof. The codifferential $d_p^* f\colon T_\alpha^*\mathbb{R} \to T_p^*M$ is defined on a one-dimensional vector space with canonical coframe $\{[1_\mathbb{R} - \alpha]_\sim\}$, so we can identify $d_p^* f$ with its value

$$(d_p^* f)\,[1_\mathbb{R} - \alpha]_\sim = [(1_\mathbb{R} - \alpha) \circ f]_\sim = [\tilde{f}]_\sim = [f]_\sim,$$

as claimed.

Also the differential $d_p c\colon T_\tau\mathbb{R} \to T_pM$ is defined on a one-dimensional vector space, now having $\{[1_\mathbb{R} + \tau]_\sim\}$ as its canonical frame, and we may identify

$$(d_p c)\,[1_\mathbb{R} + \tau]_\sim = [c \circ (1_\mathbb{R} + \tau)]_\sim = [\tilde{c}]_\sim = [c]_\sim$$

with the differential $d_p c$. □

Somewhat loosely, we may *summarize* this as follows: A cotangent vector $[f]_\sim \in T_p^*M$ is the cotangent vector of any of its representatives f, and a tangent vector $[c]_\sim \in T_pM$ is the tangent vector of any of its representatives c. This confirms the terminology of "cotangent space" and "tangent space", and it justifies the notation $(dx^1, \ldots, dx^n)$ for the coframe and $(dx_1, \ldots, dx_n)$ for the frame in the following sense: By Proposition 2.23, we can identify $dx^i = [x^i]_\sim \in T_p^*M$ with $d_p^* x^i$ and $dx_i = [x_i]_\sim \in T_pM$ with $d_p x_i$. Formally speaking, we are just suppressing the point $p \in M$ and the asterisk in the coframe (the latter is clearly motivated by Definition 2.22).

The next interesting hybrid map is exemplified by a *chart* $x\colon M \to \mathbb{R}^n$ and its *parametrization* $x^{-1}\colon \mathbb{R}^n \to M$. It turns out that we obtain well-known maps when we pass to their (co)differential—the component maps for the tangent and cotangent space, respectively! This is not surprising since the idea of the tangent space is to provide a linear approximation of the manifold (locally given through a chart or a parametrizations).

2.24 Proposition *For a chart x centered at $p \in M$, we have $d_p x = |)_x$ and $d_0^* x^{-1} = (|_x$ as well as $d_0 x^{-1} = |\rangle_x$ and $d_p^* x = \langle|_x$.*

Proof. The differential $d_p x\colon T_pM \to T_0\mathbb{R}^n$ may be interpreted as a linear map $d_p x\colon T_pM \to \mathbb{R}^n$. This can be done by identifying $T_0\mathbb{R}^n$ with $\mathbb{R}^n$, mapping the frame $(d\delta_1, \ldots, d\delta_n)$ of the former to the frame $(\delta_1, \ldots, \delta_n)$ of the latter. For proving $d_p x = |)_x$, we consider an arbitrary frame vector $dx_i = [x^{-1} \circ \delta_i]_\sim \in T_pM$ and compute

$$(d_p x)\, dx_i = [x \circ (x^{-1} \circ \delta_i)]_\sim = [\delta_i]_\sim = d\delta_i$$

which can be identified with $|dx_i)_x = \delta_i \in \mathbb{R}^n$.

Analogously, the codifferential $d_0^* x^{-1}\colon T_p^*M \to T_0^*\mathbb{R}^n$ is interpreted as a linear map $d_0^* x^{-1}\colon T_p^*M \to \mathbb{R}_n$ by identifying $T_0^*\mathbb{R}^n$ with $\mathbb{R}_n$, this time mapping the frame $(d\delta^1, \ldots, d\delta^n)$

of the former to the frame $(\delta^1, \dots, \delta^n)$ of the latter. For proving $d_0^* x^{-1} = (|_x$, we take now a coframe vector $dx^i = [\delta^i \circ x]_\sim \in T_p^* M$ and compute

$$(d_0^* x^{-1})\, dx^i = [(\delta^i \circ x) \circ x^{-1}]_\sim = [\delta^i]_\sim = d\delta^i,$$

which corresponds of course to $\delta^i \in \mathbb{R}_n$.

The other two formulae are an immediate consequence of the bijectivity of x and the functoriality of d_p as well as d_p^*. □

Recall that we called the image of an "abstract point" under a chart x its *coordinates*, whereas the image of an "abstract cotangent / tangent vector" under the component isomorphisms are of course called its *components*. In this sense, we could say that the chart differential passes from coordinates (nonlinear labels) to components (linearized labels), and this is also true for the differential of an arbitrary map between manifolds.

2.2.4 Co- and Contravariance Revisited

In Subsection 2.1.3 we have investigated the notion of co- and contravariant transformation laws from a *geometric viewpoint*: Changing from a coordinate system x to another $\bar{x}$, an object transforms by rescaling its representation by $\partial x/\partial \bar{x}$ for the covariant and $\partial \bar{x}/\partial x$ for the contravariant case.

On the other hand, we know from *category theory* that contravariant functors (also called "cofunctors") reverse the order of composition while covariant functors (otherwise just "functors") preserve it. We have seen an example of this in Subsection 2.2.1: the differential $d_{\dots} \colon \mathfrak{PtMan} \to \mathfrak{Vec}$ is a functor, the codifferential $d^*_{\dots} \colon \mathfrak{PtMan} \to \mathfrak{Vec}$ a cofunctor.

As one would suspect, there is a relation between the categorical and the geometric notions of variance. But we must be careful because, as pointed out in Subsection 2.1.3, a coordinate change (also known as a chart transition) manifests itself in an *inverse* manner: Expanding by a factor $\kappa > 1$ is realized by a coordinate change $1_{\bar{x}x} \colon A \to \bar{A}$ with $x(p) = \kappa\, \bar{x}(p)$, which appears to be a compression by a factor κ^{-1}, however!

For any vector space V, a linear endomorphism $A \colon V \to V$ may be interpreted as an *active or passive transformation*: as a coordinate change or as an actual movement. These two interpretations are exactly inverse to each other, in the following sense. On the one hand, a passive transformation $1_{\bar{x}x}$ assigns to a point p new coordinates, which can be interpreted as the coordinates of another point $\bar{p}$ in the old system. On the other hand, the point p is moved to $\bar{p}$, by the active transformation $1_{x\bar{x}}$. More precisely: The transition matrix $1_{\tilde{x}x}$ from a basis $x = (x_1, \dots, x_n)$ to its A-image $\tilde{\beta} = (Ax_1, \dots, Ax_n)$ is the inverse of the representation matrix A_{xx}.

From a geometric point of view, we must therefore regard a transition $1_{\bar{x}x} \colon A \to \bar{A}$ as a morphism from $\bar{A}$ to A. This makes the differential into a contravariant functor $d_{\dots} \colon \mathfrak{PtMan}^* \to \mathfrak{Vec}$ and the codifferential into a covariant functor $d^*_{\dots} \colon \mathfrak{PtMan}^* \to \mathfrak{Vec}$. In order to avoid the notational encumbrance of dealing with chart domain overlaps, let us introduce the quotient category [34[51]] of $\mathfrak{PtMan}$ with respect to *germ equivalence* at the base points. Writing $\mathfrak{PtMan}_\sim$ for this category, it is obvious that the differential and codifferential factor over this quotient category, so we obtain in our setting a contravariant functor $d_{\dots}\colon \mathfrak{PtMan}_\sim^* \to \mathfrak{Vec}$ and a covariant functor $d^*_{\dots}\colon \mathfrak{PtMan}_\sim^* \to \mathfrak{Vec}$.

The co- and contravariance of the transformation laws comes from the corresponding *Hom-functors*. Fix a manifold M with base point p as an object in $\mathfrak{PtMan}_\sim$ and consider the two covariant Hom-functors $\mathrm{Hom}(T_pM, \_\_), \mathrm{Hom}(T_p^*M, \_\_)\colon \mathfrak{Vec} \to \mathfrak{Set}$ and the two contravariant Hom-functors $\mathrm{Hom}(\_\_, T_pM), \mathrm{Hom}(\_\_, T_p^*M)\colon \mathfrak{Vec} \to \mathfrak{Set}$. Composing them with the contravariant functor $d_{\dots}\colon \mathfrak{PtMan}_\sim^* \to \mathfrak{Vec}$ and the covariant functor $d^*_{\dots}\colon \mathfrak{PtMan}_\sim \to \mathfrak{Vec}$, we obtain two covariant functors

$$\mathbf{B}_* = \mathrm{Hom}(d_p\,\_\_, T_pM), \mathbf{C}_* = \mathrm{Hom}(T_p^*M, d_p^*\,\_\_)\colon \mathfrak{PtMan}_\sim^* \to \mathfrak{Set}$$

as well as two contravariant functors

$$\mathbf{B}^* = \mathrm{Hom}(d_p^*\,\_\_, T_p^*M), \mathbf{C}^* = \mathrm{Hom}(T_pM, d_p\,\_\_)\colon \mathfrak{PtMan}_\sim^* \to \mathfrak{Set}$$

As we shall see in a moment, the covariant / contravariant functors give rise to covariant / contravariant transformation laws, namely $\mathbf{B}_*$ and $\mathbf{B}^*$ for the bases, $\mathbf{C}_*$ and $\mathbf{C}^*$ for the components.

$$\begin{tikzcd}
& T_p^*M & & & & M \arrow[dl, "x"'] \arrow[dr, "\bar x"] & & & & T_pM \arrow[dl, "|)_x"'] \arrow[dr, "|)_{\bar x}"] & \\
\mathbb{R}_n \arrow[ur, "\langle|_x"] & & \mathbb{R}_n \arrow[ll, "\cdot\,\partial\bar x/\partial x"'] \arrow[ul, "\langle|_{\bar x}"'] & \overset{d^*_{\dots}}{\Longleftarrow} & A \arrow[rr, "1_{\bar x x}"] \arrow[dr, "x^{-1}"'] & & \bar A \arrow[dl, "\bar x^{-1}"] & \overset{d_{\dots}}{\Longleftarrow} & \mathbb{R}^n \arrow[rr, "\partial\bar x/\partial x\,\cdot"] \arrow[dr, "|\rangle_x"'] & & \mathbb{R}^n \arrow[dl, "|\rangle_{\bar x}"] \\
& T_p^*M \arrow[ul, "(|_x"] \arrow[ur, "(|_{\bar x}"'] & & & & M & & & & T_pM &
\end{tikzcd}$$

The above *diagram* visualizes the action of these functors, showing (from left to right): in the upper layer the contravariant functors $\mathbf{B}^*$ and $\mathbf{C}^*$, in the lower layer the covariant functors $\mathbf{C}_*$ and $\mathbf{B}_*$. (In order to avoid confusion, we have displayed the morphisms in $\mathfrak{PtMan}_\sim$ rather than $\mathfrak{PtMan}_\sim^*$, namely by way of representatives $x\colon U \to A$ and $\bar x\colon \bar U \to \bar A$ but with the chart domains U and $\bar U$ suppressed. Note also that the functors $\mathbf{B}^*, \mathbf{C}^*, \mathbf{C}_*, \mathbf{B}_*$ are all to be applied to the middle morphism $1_{\bar x x}$, leading to a vertical reflection of the diagonal arrows in the two outer diagrams.)

Let us now see what these diagrams mean in detail. Applying the functor $\mathbf{C}^*$ and $\mathbf{C}_*$ respectively to the identities $\bar x = 1_{\bar x x} \circ x$ and $x^{-1} = \bar x^{-1} \circ 1_{\bar x x}$, Proposition 2.24 yields $|)_{\bar x} = (\partial\bar x/\partial x) \cdot |)_x$ and $(|_x = (|_{\bar x} \cdot (\partial\bar x/\partial x)$. If we multiply the latter identity by the inverse Jacobian $\partial x/\partial\bar x$ and substitute an arbitrary cotangent vector $d \in T_p^*M$ and tangent vector $v \in T_pM$, we obtain $|v)_{\bar x} = (\partial\bar x/\partial x) \cdot |v)_x$ and $(d|_{\bar x} = (\partial x/\partial\bar x) \cdot (d|_x$; these are exactly the formulae (2.5) of Subsection 2.1.3, describing the co- and contravariant *transformations of the components*.

In a similar fasion, we can now derive the corresponding *transformation of the bases* by applying the functors $\mathbf{B}^*$ and $\mathbf{B}_*$ to the identities $\bar{x} = 1_{\bar{x}x} \circ x$ and $x^{-1} = \bar{x}^{-1} \circ 1_{\bar{x}x}$. Using again Proposition 2.24, this gives $\langle|_{\bar{x}} = \langle\cdot(\partial\bar{x}/\partial x)|_x$ and $|\rangle_x = |(\partial\bar{x}/\partial x)\cdot\rangle_{\bar{x}}$. Precomposing the latter identity by the inverse Jacobian $\partial x/\partial\bar{x}$ and substituting an arbitrary coframe vector δ^i and frame vector δ_i, we obtain respectively

$$d\bar{x}^i = \langle\delta^i \frac{\partial\bar{x}}{\partial x}|_x = \langle\delta^i_k \frac{\partial\bar{x}^k}{\partial x_j}\,\delta^j|_x = \langle\frac{\partial\bar{x}^i}{\partial x_j}\,\delta^j|_x = \frac{\partial\bar{x}^i}{\partial x_j}\,\langle\delta^j|_x = \frac{\partial\bar{x}^i}{\partial x_j}\,dx^j$$

and

$$d\bar{x}_i = |\frac{\partial x}{\partial\bar{x}}\,\delta_i\rangle_x = |\delta_j\,\frac{\partial x^j}{\partial\bar{x}_k}\,\delta^k_i\rangle_x = |\delta_j\,\frac{\partial x^j}{\partial\bar{x}_i}\rangle_x = \frac{\partial x^j}{\partial\bar{x}_i}\,|\delta_j\rangle_x = \frac{\partial x^j}{\partial\bar{x}_i}\,dx_j,$$

which are exactly the basis transformation laws (2.10) of Subsection 2.1.3.

The functorial view of cotangent and tangent vectors also enables a better understanding of the *component spaces* introduced at the end of Subsection 2.1.3.

2.25 Proposition *Up to isomorphism, the (co)differential is determined uniquely as (co)functor $\mathfrak{PtMan}_\sim \to \mathfrak{Vec}$ that maps open sets $A, \bar{A}$ of $\mathbb{R}^n$ to the vector space $\mathbb{R}^n$ and differentiable maps $f\colon A \to \bar{A}$ to their (co)differential.*

Proof. Let $\mathcal{T}\colon \mathfrak{PtMan}_\sim \to \mathfrak{Vec}$ be any functor with the required differential properties and M a fixed manifold with differentiable structure $\mathfrak{D}$ and base point p. Then every $v \in \mathcal{T}M$ induces a map $V\colon \mathfrak{D}_{p\bullet} \to \mathbb{R}^n$ by sending $x \in \mathfrak{D}_{p\bullet}$ to $V_x = (\mathcal{T}x)v$. Since $\mathcal{T}$ is a functor, we have $V_{\bar{x}} = (\partial\bar{x}/\partial x)\cdot V_x$, which means that V is a contravariant component vector in the sense of Theorem 2.8. Conversely, any contravariant component vector $V\colon \mathfrak{D}_{p\bullet} \to \mathbb{R}^n$ comes from a unique $v \in \mathcal{T}M$ since we have $v = (\mathcal{T}x)^{-1}\,V_x$ for an arbitrary $x \in \mathfrak{D}_{p\bullet}$. This means that $\mathcal{T}M$ is isomorphic to the contravariant component space and hence by Theorem 2.8 also to T_pM.

Now consider a map $\Phi\colon M \to \bar{M}$ between an n-manifold M and an m-manifold $\bar{M}$ mapping $p \in M$ to $q \in \bar{M}$. The morphism $\mathcal{T}\Phi\colon \mathcal{T}M \to \mathcal{T}\bar{M}$ can be interpreted as a linear map between the corresponding component spaces the contravariant component vector determined by $(\mathcal{T}\Phi)v \in \mathcal{T}\bar{M}$ to the contravariant component vector determined by $v \in \mathcal{T}M$. Applying the functor $\mathcal{T}$ to the identity $\Phi_{\bar{x}x} = \bar{x} \circ \Phi \circ x^{-1}$ yields $\bar{V} = (\mathcal{T}\Phi)\,V$ with

$$\bar{V}_{\bar{x}} = \frac{\partial\bar{x}(\Phi)}{\partial x}\,V_x,$$

which is exactly the component representation 2.12 of the differential $d_p\Phi$ under the above isomorphisms $\mathcal{T}M \cong T_pM$ and $\mathcal{T}\bar{M} \cong T_q\bar{M}$.

The proof for the codifferential is completely analogous, using covariant instead of contravariant component vectors. □

Note that a differential $\mathcal{T}^*\colon \mathfrak{PtMan}_\sim \to \mathfrak{Vec}$ in the sense of Proposition 2.25 is a functor, a codifferential $\mathcal{T}_*\colon \mathfrak{PtMan}_\sim \to \mathfrak{Vec}$ a cofunctor. Looking at it "geometrically", however, the differential appears as a contravariant functor $\mathcal{T}^*\colon \mathfrak{PtMan}^*_\sim \to \mathfrak{Vec}$ and the codifferential as a contravariant functor $\mathcal{T}_*\colon \mathfrak{PtMan}^*_\sim \to \mathfrak{Vec}$. This is consistent with the terminology of Theorem 2.8: As we have seen in the proof of Proposition 2.25, the object function of $\mathcal{T}^*$ provides the *contravariant component space*, the object function of $\mathcal{T}_*$ the *covariant component space.*

2.2.5 The Rank Theorem

We know from differential calculus in $\mathbb{R}^n$ that several (local) properties of a differentiable map can be characterized by properties of its linear approximation. Recall for example the Inverse Mapping Theorem 1.1. In the following, we first discuss the Rank Theorem, which is a generalization of the Inverse Mapping Theorem and the Implicit Function Theorem. Using charts we translate the result to differentiable maps between manifolds and their differentials.

We first look at the situation in Linear Algebra. Let $A\colon V \to W$ be a linear map of rank r between finite dimensional vector spaces. Then we know that there are linear isomorphisms $\varphi\colon V \to \mathbb{R}^m$ and $\psi\colon W \to \mathbb{R}^n$ such that $\psi A \varphi^{-1}\colon \mathbb{R}^m \to \mathbb{R}^n$ is given by

$$(x_1, \ldots, x_m) \mapsto (x_1, \ldots, x_r, 0, \ldots, 0).$$

Let now $f\colon U \to \mathbb{R}^n$ be a differentiable map and $x \in U \subseteq \mathbb{R}^m$. We define the *rank* of f at x, denoted by $\operatorname{rnk}_x f$, as the rank of the differential $d_x f$ or the Jacobian matrix $f'(x)$. The Rank Theorem [9[109]] tells us that every differentiable map that of locally constant rank looks like the linear counterpart in suitable coordinates.

2.26 Theorem (Rank Theorem for Vector Spaces) *Let $f \in C^s(\tilde{U}, \mathbb{R}^n)$ with $s \geq 1$, $\tilde{U} \subseteq \mathbb{R}^m$, and consider a point $p \in \tilde{U}$ with image $q = f(p)$. If the map f has locally constant rank r around p, there exist open neighborhoods U and V of p and q, respectively, and C^s diffeomorphisms $\varphi\colon U \to A$ and $\psi\colon V \to B$ with $\varphi(p) = 0$, $\psi(q) = 0$ such that $\psi f \varphi^{-1}\colon A \to B$ is given by*

$$(x_1, \ldots, x_m) \mapsto (x_1, \ldots, x_r, 0, \ldots, 0),$$

locally around 0.

Notice that if we consider f as a differentiable map between the manifolds $\tilde{U}$ and $\mathbb{R}^n$, the conclusion of this theorem says that there exists a local representative $f_{\psi\varphi}$ of the indicated form.

Locally the rank of a differentiable map *cannot decrease*: For $\operatorname{rnk}_x f = r$, the Jacobian matrix $f'(x)$ has an $r \times r$ submatrix with nonzero determinant. Since the Jacobian depends continuously on x and the determinant is continuous, this subdeterminant is also nonzero in a neighborhood of x, and the rank there is at least r. In fact, the rank can *increase* as we can see from considering $x = 0$ in the example $f(x) = x^2$.

There are two cases where the rank cannot increase and is thus locally constant: If the rank of f at x is either m or n hence maximal for the given m and n. The condition $\operatorname{rnk}_x f = m$ is of course only possible if $m \leq n$ and equivalent to $d_x f$ being injective (by the Rank-Nullity Theorem of Linear Algebra). Then we call f an *immersion* at x. Similarly, we call f a *submersion* at x if $d_x f$ is surjective so $\operatorname{rnk}_x f = n$ and hence in particular $m \geq n$. We call f an immersion / submersion if it is one at every point.

We are in the situation of the *Inverse Mapping Theorem* if f is both an immersion and a submersion at x, so $\operatorname{rnk}_x f = m = n$. If f is a submersion at x, then x is also called a regular point of f, see Subsection 1.1.2.

By the Rank Theorem, we know that an immersion (assuming $m \leq n$) can be represented locally by an embedding

$$\begin{aligned} \mathbb{R}^m &\to \mathbb{R}^n \\ (x_1, \ldots, x_m) &\mapsto (x_1, \ldots, x_m, 0, \ldots, 0), \end{aligned}$$

and a submersion (assuming $m \geq n$) by a projection

$$\begin{aligned} \mathbb{R}^m &\to \mathbb{R}^n \\ (x_1, \ldots, x_n, \ldots, x_m) &\mapsto (x_1, \ldots, x_n). \end{aligned}$$

Therefore an immersion is locally injective while a submersion is open.

Let now M and N be two manifolds with a map $\Phi\colon M \to N$, differentiable at a point $p \in M$ mapping to $q = \Phi(p) \in N$. We define the *rank* of Φ at p, denoted by $\operatorname{rnk}_p \Phi$, as the rank of the differential $d_p\Phi\colon T_pM \to T_qN$. From Section 2.2.2 we know that a matrix representation of the differential is given by the Jacobian of a local representative of Φ. So the rank of a differentiable map is also the rank of the Jacobian matrix of any of its local representative, $\operatorname{rnk}_p \Phi = \operatorname{rnk}_x \Phi_{\psi\varphi}$ with $x = \varphi(p)$.

From the Rank Theorem 2.26 for vector spaces, we immediately obtain the following version for manifolds.

2.27 Theorem (Rank Theorem for Manifolds) *Let $\Phi\colon M \to N$ be differentiable with locally constant rank $r = \operatorname{rnk}_p \Phi$ around p. Then there exist charts φ and ψ centered at p and q, respectively, such that the* local representative *$\Phi_{\psi\varphi}$ is given by*

$$(x_1, \ldots, x_m) \mapsto (x_1, \ldots, x_r, 0, \ldots, 0),$$

locally at 0.

Immersions and submersions are defined as before. Then the Rank theorem implies in particular that for *immersions* there is a local representative $\Phi_{\psi\varphi}$ given by

$$(x_1, \ldots, x_m) \mapsto (x_1, \ldots, x_m, 0, \ldots, 0)$$

and for *submersions* by

$$(x_1, \ldots, x_n, \ldots, x_m) \mapsto (x_1, \ldots, x_n).$$

We obtain also the following version of the Inverse Mapping Theorem for differentiable maps between manifolds: The differential $d_p\Phi$ at p is an isomorphism iff Φ is a local diffeomorphism at p.

Immersions can be used to define a more general notion of submanifolds than we considered in Subsection 1.4.1, see for example [7[79]], [46[234]] and [2[75]]. Let us briefly discuss this notion and its relation to our definition. A manifold N is called a *submanifold* of a manifold M if $N \subseteq M$ and the embedding

$$\begin{aligned} \iota\colon N &\to M \\ p &\mapsto p \end{aligned}$$

is an immersion. Note that in the definition we require a differentiable structure on N, so N already has a topology. Since N is a subset of M, we can also consider the induced topology on N. In general, these two topologies are *different*. We only know that the induced topology is contained in the topology as a manifold (this is simply the continuity of the map ι).

The manifold N is an *embedded submanifold* (also called a *regular submanifold*) if the embedding ι is a homeomorphism onto its image, that is, the induced topology and the topology as a manifold coincide. Otherwise N is sometimes called an *immersed submanifold*, to emphasize that it is not an embedded submanifold.

Of course we should verify that this definition of embedded submanifold is equivalent to Definition 1.50. From Subsection 1.4.1 we know that the topology as a manifold and the induced topology are equal, and from Proposition 1.53 that the embedding ι is differentiable. Conversely, one can show that the above definition implies the existence of charts as required in Definition 1.50 of embedded submanifolds by applying the Rank Theorem 2.27 to the given immersion.

Submanifold in the above sense can also be interpreted as parametrizations with arbitrary manifolds as parameter set. Let $\Phi\colon N \to M$ be an *injective immersion*. We can decompose $\Phi = \iota \circ \tilde{\Phi}$ into a bijective map $\tilde{\Phi}\colon N \to \Phi(N)$ and the insertion ι. Since $\tilde{\Phi}$ is a bijection, it induces a differentiable structure on $N' = \Phi(N)$. Then N' becomes a manifold, $\tilde{\Phi}$ a diffeomorphism, and $\iota = \Phi \circ \tilde{\Phi}^{-1}$ an immersion. But this means that N' is a submanifold of M.

Submersions can be used to implicitly define manifolds. Compare to Subsection 1.1.2, where we discussed this construction for embedded submanifolds of $\mathbb{R}^m$. Let M and L be two manifolds of dimension m and k, respectively, with $k \leq m$. Let $\Phi\colon M \to L$ a differentiable map, $q \in L$ and $N = \Phi^{-1}(q)$. We call q a *regular value* of Φ if Φ is a submersion at each point in N.

2.28 Theorem (Regular Value Theorem) *If $q \in L$ is a regular value of $\Phi\colon M \to L$, then N is an embedded submanifold of M of codimension k.*

Proof. Let $p \in N$ and $n = m - k$. Since Φ is a submersion at p, the Rank Theorem 2.27 implies that there exist charts $\varphi\colon U \to A$ and ψ centered at p and q, respectively, such that the local representative $\Phi_{\psi\varphi}$ is given by

$$(x_1, \dots, x_n, x_{n+1}, \dots, x_{n+k}) \mapsto (x_{n+1}, \dots, x_{n+k}).$$

We show that for the chart φ, we have

$$p' \in U \cap N \Leftrightarrow (\varphi^{n+1}(p') \dots, \varphi^{n+k}(p')) = 0.$$

Note that for $p' \in U$, we have $\Phi(p') = q$ iff

$$\psi(\Phi(p')) = \psi\Phi\varphi^{-1}(\varphi(p')) = \psi(q).$$

Since $\psi(q) = 0$, and the local representative is given by

$$\Phi_{\psi\varphi}(\varphi(p')) = (\varphi^{n+1}(p'), \ldots, \varphi^{n+k}(p')),$$

the stated equivalence is clear. Hence

$$\varphi(U \cap N) = A \cap \mathbb{R}^n,$$

and we have proved for every $p \in N$ the existence of charts as required in Definition 1.50 of embedded submanifolds. $\square$

Just like injective immersions can be compared with a parametric description of varieties (see above), submersions correspond to an implicit description of varieties.

2.3 Tensor Fields on a Manifold

2.3.1 Fiber Bundles

In Subsection 2.2.1 we have viewed the codifferential and the differential respectively as a *cofunctor and functor*

$$d^*_{\ldots}\colon \mathfrak{PtMan} \to \mathfrak{Vec} \quad \text{and} \quad d_{\ldots}\colon \mathfrak{PtMan} \to \mathfrak{Vec}$$

from the category $\mathfrak{PtMan}$ of *pointed manifolds* to the category $\mathfrak{Vec}$ of vector spaces. Specifically, if $\Phi\colon M \to N$ is a continous map between manifolds M and N, differentiable at the point $p \in M$ with image $q = \Phi(p)$, we obtain linear maps $d^*_p\Phi\colon T^*_qN \to T^*_pM$ and $d_p\Phi\colon T_pM \to T_qN$.

Suppose now that Φ is differentiable at every point $p \in M$. Then we can construct the tangent spaces and differentials for all these points (we need not mention that everything works just as well in the co-style). It would be nicer, though, to compress all these tangent spaces into just two super-objects (one for M and one for N), and to paste together all the differentials into one super-morphism between these two super-objects. This is what *vector bundles* can do for us: they provide the kind of super-objects that we need here, and the corresponding bundle maps are the required super-morphisms. Writing $\mathfrak{VecBnd}$ for this category, the differential comes out as a functor $d\colon \mathfrak{Man} \to \mathfrak{VecBnd}$ and the codifferential as a cofunctor $d^*\colon \mathfrak{Man} \to \mathfrak{VecBnd}$. But before introducing vector bundles, we will briefly have a look at the more general concept of *fiber bundles*.

Let B be a fixed topological space. A *bundle* over B is a topological space E together with a surjective continuous map $\pi\colon E \to B$. We call E the "total space" (containing the "bundle points") and B the "basis space" or just "basis" (containing the "basis points"), whereas π is referred to as the "bundle projection" or just "projection". A bundle over B will be called a B-bundle for short. Now for some basic vocabulary about bundles.

- The *fiber* over $p \in B$ is defined as the closed set $\pi^{-1}(p)$, denoted by E_p.

- For an open set U in B, the *slice* over U, here denoted by E_U, is the bundle $\pi^{-1}(U)$ over B, of course using the restricted projection $\pi|_{E_U}$. (A more common notation for E_U is actually $E|_U$, but we prefer to make the relation to fibers more explicit: We view E_p as abbreviation for $E_{\{p\}}$, except that $\{p\}$ is not open.)

- Let $\tilde{E}$ be another B-bundle. A *bundle morphism* (over B) is a continuous map $h\colon E \to \tilde{E}$ such that $\pi = \tilde{\pi} \circ h$, where π and $\tilde{\pi}$ are respectively the projections of E and $\tilde{E}$. This means we have $h(E_p) \subseteq \tilde{E}_p$ for every $p \in B$.

- If $h\colon E \to \tilde{E}$ is moreover a homeomorphism, we speak of a *bundle isomorphism* (over B), and we call the bundles E and $\tilde{E}$ isomorphic. Note that h swaps E_p and $\tilde{E}_p$ for every $p \in B$.

- The bundles over a fixed basis B form a category with morphisms the bundle morphisms over B, a subcategory of the slice slice category $\mathfrak{Top} \downarrow B$, where $\mathfrak{Top}$ denotes the usual category of topological spaces with morphisms the continuous maps.

- More generally, one may also consider the *bundles over arbitrary bases* as a category, this time a subcategory of the comma category $\mathfrak{Top} \downarrow \mathfrak{Top}$. Given a B-bundle E with projection π and a $\tilde{B}$-bundle $\tilde{E}$ with projection $\tilde{\pi}$, a bundle morphism is given by a continous map $h\colon E \to \tilde{E}$ together with a continuous map $h_0\colon B \to \tilde{B}$ between the bases such that $h_0 \circ \pi = \tilde{\pi} \circ h$; in this case, one speaks of a map h "over" a map h_0.

Having a map h over a map h_0, there is for each basis point $p \in B$ a point $\tilde{p} = h_0(p) \in \tilde{B}$ yielding the fiber inclusion $h(E_p) \subseteq \tilde{E}_{\tilde{p}}$. Conversely, every continous map $h\colon E \to \tilde{E}$ with this fiber inclusion property determines the basis map $h_0\colon B \to \tilde{B}$ uniquely via $h_0(p) = \tilde{\pi}h(E_p)$, but note that h_0 need not be continuous in general. At any rate, we may describe a morphism between a B-bundle E and a $\tilde{B}$-bundle $\tilde{E}$ as a continuous map $h\colon E \to \tilde{E}$ that fulfills the fiber inclusion property such that the induced basis map is continuous.

A *section* (sometimes also called "cross section") of a B-bundle E is a continuous map $\sigma\colon B \to E$ such that $\pi \circ \sigma = 1_B$. The geometric motivation for the name (meaning "cut" in Latin) becomes evident in simple cases like a cylinder $E = S^1 \times [0,1]$, considered as a bundle over the basis $B = S^1$. (The terminology is also consistent with the generic case: In any category, if one has two morphisms $\pi\colon E \to B$ and $\sigma\colon B \to E$ with $\pi \circ \sigma = 1_B$, one calls π a left inverse or a retraction for σ and dually σ a right inverse or a section for π.)

In a general B-bundle E, it is not always possible—or at least not easy—to find a section $\sigma\colon B \to E$. One is therefore often content with a *local section*, meaning a map $\sigma\colon U \to E$ that is a section of E_U. Note that such sections form a sheaf—and this is actually the motivation for the name "section" used in sheaf theory (see at the beginning of Subsection 1.3.2). Certain powerful topological tools (like partitions of unity) allow the construction of a global section (meaning a section $\sigma\colon B \to E$) from a collection of local sections covering B.

One important example of a bundle is the following. For a given topological space F, we consider $E = B \times F$ as a B-bundle, using the canonical projection $\pi\colon B \times F \to B$ acting by $\pi(p, v) = p$. We call E the *straight* B-bundle with *typical fiber* F.

2.29 Definition *The B-bundle E is called* locally straight *if every basis point $p \in B$ has a neighborhood U such that E_U is isomorphic to a straight B-bundle.*

In detail, there is for each such neighborhood U of the basis point p a homeomorphism $\varphi\colon E_U \to U \times F$ with $\pi(\varphi^{-1}(p, v)) = p$ for all $p \in U$ and $v \in F$; such a map is called a *bundle chart* of E over U. Obviously we can write $\varphi = (\pi, \varphi_F)$ with a continuous *fiber chart* $\varphi_F\colon E_U \to F$. The fibers F need not be the same for different bundle points p, but we have always $F \cong E_p$. If all fibers are indeed homeomorphic to a fixed F, we call E locally straight with *typical fiber* F. (This terminology is consistent with the case of a globally straight bundle $E = B \times F$, where each fiber is given by $\{p\} \times F \cong F$.)

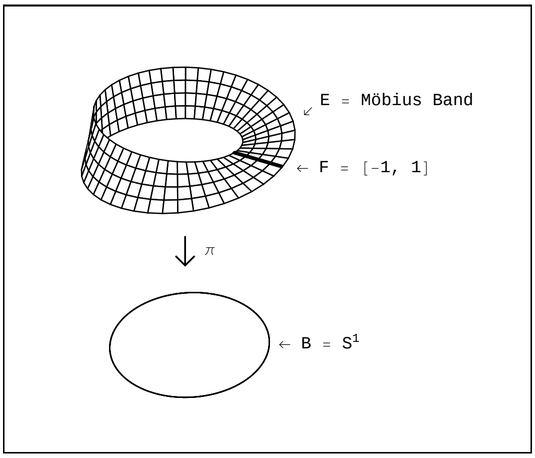

Figure 2.1: The Möbius Band

Note the analogy to manifolds here: Again we have maps (bundle charts) for locally "pulling down" a complicated structured (twisting fibers in a B-bundle E) to a simple reference object (a straight B-bundle). The point is that even though E is *locally straight*, it may still be *globally twisted*. Think of the Möbius band with basis S^1 and fiber $[-1, 1]$, as visualized in Figure 2.1.

In analogy to manifolds, we can also introduce a *bundle atlas* for a locally straight bundle as a collection of bundle charts whose domains cover the total space (or equivalently their projections cover the basis space). Hence we can read Definition 2.29 as characterizing the locally straight bundles as the ones having an atlas. Again similar to manifolds, a *bundle structure* is a maximal bundle atlas (with respect to set inclusion), containing all admissible bundle charts. In fact, we have not yet restricted the transitions—which we will do by the structure group to be introduced below—so that all bundle charts are admissible at this point.

On a locally straight B-bundle E, *local sections* $\sigma\colon U \to E$ always exist. Moreover, they are in bijective correspondence with the continuous maps $\tilde{\sigma}\colon U \to F$, which we might call *fiber maps*: Given an admissible bundle chart $\varphi\colon E_U \to U \times F$, we obtain fiber maps from local sections via $\tilde{\sigma}(p) = \varphi_F(\sigma(p))$ and local sections from fiber maps via $\sigma(p) = \varphi^{-1}(p, \tilde{\sigma}(p))$.

Note that the homeomorphism type of E_p is *locally constant* in every B-bundle E, so we have a typical fiber at least for every connected component of B. In particular, all bundles over a connected basis B do have a typical fiber $F \cong E_p$ for all $p \in B$. Hence it is only a small step from locally straight bundles to fiber bundles.

2.30 Definition *A* fiber bundle *is a locally straight bundle with a typical fiber.*

Let us return to the analogy with manifolds. The crucial idea of a C^r manifold is that every point has a neighborhood that is locally mapped to an open set of $\mathbb{R}^n$ in such a way that all the transitions are C^r. One can impose similar restrictions on the *bundle charts* of a fiber bundle.

Recall that every B-bundle E comes with a bundle atlas, which can be extended to the unique bundle structure consisting of all admissible bundle charts—just like in the manifold case. If E is a fiber bundle with typical fiber F, every admissible bundle chart can be cast into a map $\varphi\colon E_U \to U \times F$. Now let $\psi\colon E_V \to V \times F$ be another admissible bundle chart with non-empty overlap $W = U \cap V$. Then the *bundle transition* from φ to ψ is given by the homeomorphism

$$1_{\psi\varphi} = \psi\varphi^{-1}\colon W \times F \to W \times F.$$

Since $1_{\psi\varphi}$ is a bundle isomorphism, we have

$$1_{\psi\varphi}(p, v) = (p, 1'_{\psi\varphi}(p)\, v) \quad \text{with} \quad 1'_{\psi\varphi}(p)\colon F \to F$$

being a unique homeomorphism depending on the basis point p; let us call $1'_{\psi\varphi}(p)$ the *fiber transition* at p from φ to ψ. (Note that, for the moment, $1'_{\psi\varphi}(p)$ is just a convenient notation for a certain map depending on the chosen charts φ and ψ; it has nothing to do with derivatives, which are not even defined when the fiber is a plain topological space! But in the special case of tangent bundles on a manifold, $1'_{\psi\varphi}(p)$ will indeed be the good old transition Jacobian: the derivative of the coordinate change $1_{\psi\varphi}$ around p.)

The fiber transitions obviously form a group. Taking any homeomorphism γ of F and any open set W of B, one can build admissible bundle transitions $(1_W, \gamma)$, from any chart φ with domain W to its corresponding chart $(1_W, \gamma) \circ \varphi$ with the same domain W. Hence the *group of fiber transitions* is given by all homeomorphisms of F. In order to restrict the possible transitions in the bundle, we must prescribe a smaller subgroup G. Usually one thinks of G not as a subset of the full homeomorphism group of F, but as a separate group acting effectively on F. Moreover, one also specifies a topology on G such that G can be required to act continuously on F; this is sometimes convenient for fine-tuning the admissible transitions. Let us summarize these ideas.

2.31 Definition *Let E be a fiber bundle over B with typical fiber F and with a topological group G acting effectively on F as a group of homeomorphisms. Then E is said to admit the* structure group G *if there is a continuous map $\Theta_{\psi\varphi}\colon W \to G$ for any two admissible charts φ and ψ with overlap $W \subseteq B$ such that $1'_{\psi\varphi}(p)\, v = \Theta_{\psi\varphi}(p) \cdot v$ for all $p \in W$ and $v \in F$.*

For the sake of a readability, we identify the homeomorphism $1'_{\psi\varphi}(p)\colon F \to F$ and the group element $\Theta_{\psi\varphi}(p) \in G$, using the former for the latter. A fiber bundle with structure group G will briefly be called a G-structured bundle. Just as with manifolds constructed via patchwork (see Subsection 1.2.3), the fiber transitions from G satisfy the three *transition conditions*, thus making the gluing relation an equivalence (for all base points in the appropriate overlaps): reflexivity $\Theta_{\varphi\varphi} = 1_F$, symmetry $\Theta_{\varphi\psi}\Theta_{\psi\varphi} = 1_F$, and transitivity $\Theta_{\chi\varphi} = \Theta_{\chi\psi}\Theta_{\psi\varphi}$. As in Subsection 1.2.3, the third condition—also named the "cocycle condition" in Čech cohomology—implies the other two.

The geometric significance of the structure group is to impose various *symmetry constraints* on how the fibers can by glued together. As an example, consider the Möbius band (Figure 2.1), an S^1-bundle with typical fiber $F = [-1, 1]$. One can impose $\mathbb{Z}_2$ as a structure group on it because running 360° around the band one will end up on the same fiber with opposite orientation, corresponding to a fiber transition $F \to F$ that flips the interval F via $v \mapsto -v$.

Everything which we have done up to now (bundles, bundle morphisms, sections, fiber bundles, structure groups) can be specialized to the category of manifolds: Just replace all topological spaces by manifolds and all continuous functions by differentiable ones. So a *differentiable fiber bundle* is a manifold E based on a manifold B with a manifold F for its fiber, and the projection $\pi\colon E \to B$ would be a differentiable map commuting with the (differentiable!) canonical projection $B \times F \to B$. A structure group for E would be required to be a Lie group with a differentiable action on F, inducing differentiable fiber transitions.

Before proceeding to the next subsection on vector bundles (bundles with linear fibers and linear structure groups—as we shall see soon), let us mention a few other special cases of fiber bundles.

- In complex analysis, *covering spaces* are an important tool for classifying analytic continuations. We can view them as fiber bundles with discrete fibers.

- If the fibers have a free and transitive group action, one speaks of a *principal bundle*. In this case, the fiber itself serves as a structure group.

- The n-dimensional *disk bundle* has as fiber a disk $D^n = \{x \in \mathbb{R}^n \mid \|x\| \leq 1\}$. Its structure group is often required to be a subgroup of the orthogonal group on $\mathbb{R}^n$.

- Analogously, the n-dimensional *sphere bundle* has a sphere $S^n = \{x \in \mathbb{R}^n \mid \|x\| = 1\}$ for its fiber, and again the transitions are often restricted to rotations. Disk bundles and sphere bundles are used for constructing the Thom space of a vector bundle; see [23[53]].

A few words on the *terminology* in the theory of fiber bundles and vector bundles. First of all, we must confess that to some degree we have introduced our own terms here because there seems to be no global agreement on how to name various key concepts of the theory.

For example in $[23^{157}]$, a *bundle* E over B is called a "topological space over" B, without mentioning the projection π; whereas in $[7^{106}]$, they call π a "bundle projection", ignoring the total space E. Our own usage of the "bundle" is similar to [58].

Speaking of a bundle E over B suggests a rough *parallel to vector spaces* V over a field K. In this comparision, V corresponds to E and K to B, so "vectors" are like "bundle points" and "scalars" like "basis points", while the heterogeneous operation $\cdot\colon V \times K \to V$ has a vague counterpart in the projection $\pi\colon E \to B$. Moreover, the brief name "B-bundle" corresponds to the common term "K-vector space". Generalizing the analogy from vector spaces over a field K to modules over a commutative ring K, a straight B-bundle $B \times F$ with typical fiber F corresponds to a free K-module K^F with generators F.

The term "(locally) straight bundle" is also nonstandard; the usual name is "(locally) trivial bundle". We prefer the more explicit term that refers to the *geometric contents* of this concept: Calling something trivial just provokes the question: trivial with respect to what? The expressions "fiber chart" and "fiber map" are also new (there seem to be no words for them in the literature). Note also that the pictorial terminology of "slices" is uncommon, the usual term being "restriction".

For the *transition of bundle charts*, the terminology in the literature is again undecided. Some people use the name "transition function" for what we called "bundle transition", some use it for our "fiber transition". We find it more useful to distinguish them in a natural way by referring to what they act on (either on the "bundle" in its locally isomorphic manifestation or on the fiber).

Regarding *structure groups*, we have a small problem with our terminology of B-bundles: In the literature, a fiber bundle with structure group G is sometimes called a "G-bundle". This would conflict with our notion of B-bundles since it may not always be clear from the context (at least a check would be distracting) that B is a basis space and G a group. But we may instead speak of a "G-structured B-bundle", roughly analogous to speaking of an "n-dimensional K-vector space".

2.3.2 Vector Bundles

As announced in the previous section, we are especially interested in those fiber bundles that have linear fibers and linear structure groups (subgroups of the corresponding general linear group).

2.32 Definition *An n-dimensional* vector bundle *is a fiber bundle with typical fiber an n-dimensional vector space V and linear fiber transitions $V \to V$.*

Obviously (recall we are only dealing with real vector spaces), we may replace V by $\mathbb{R}^n$ if we choose a *basis*. In particular, a one-dimensional vector bundle (also known as a line bundle) may be fibered by $\mathbb{R}$. For general vector bundles (like a Lie algebra or a space of algebraic tensors), however, there may not always be a natural basis.

If E and $\tilde{E}$ are vector bundles over the same basis B, a *vector bundle morphism* is a bundle morphism $h\colon E \to \tilde{E}$ restricting to linear maps on all fibers (then h is sometimes called a "fiberwise linear" map). This means the restriction $h\colon E_p \to \tilde{E}_p$ is a linear map for every $p \in B$. We observe that, with these morphisms, the vector bundles form a (non-full) subcategory of the fiber bundles. If $\tilde{E}$ is a vector bundle over a different basis $\tilde{B}$, a morphism h between E and $\tilde{E}$ is a fiberwise linear morphism of fiber bundles. This makes the vector bundles (over arbitrary bases) a subcategory $\mathfrak{VecBnd}$ of the fiber bundles (over arbitrary bases).

Just as with fiber bundles, one may specialize to the category of C^r manifolds, defining a *differentiable vector bundle* as a vector bundle that is at the same time a differentiable fiber bundle (and hence all the more a manifold). So we have a manifold with a differentiable projection to a basis which is also a manifold, and the structure group is a Lie group. The corresponding morphisms are of course also assumed to be differentiable maps; let us write this category as $\mathfrak{VecBnd}_r$, with the understanding $\mathfrak{VecBnd} = \mathfrak{VecBnd}_0$.

Before we proceed to our most important example—the cotangent and tangent bundle of a manifold—let us mention that one can use the basis space for transferring its topology (in case the basis is a manifold, also its differentiable structure) to the total space, thus *building a vector bundle* (in case the basis is a manifold, even a differentiable vector bundle) on what is called a "vector prebundle" in [8[29]].

More specifically, a *vector prebundle* is given as follows: We have a set E, a basis space B, a surjection $\pi\colon E \to B$, an n-dimensional vector space V, and an atlas $\mathfrak{A}$ of prebundle charts $\varphi\colon E_U \to U \times V$. The latter condition has the expected meaning:

- The $U \subseteq B$ are open sets such that the slices E_U cover E.
- Each φ is a bijection with $\pi(\varphi^{-1}(p, v)) = p$ for all $p \in U$ and $v \in V$.
- The bundle transitions $1_{\psi\varphi}\colon U \times V \to U \times V$ are homeomorphisms.
- The induced fiber transitions $1'_{\psi\varphi}(p)\colon V \to V$ linear.

Note that $\mathfrak{A}$ is then a bundle atlas for E except that E has no topology, so we cannot speak of φ as a homeomorphism. But we can endow E with a topology in the same manner as in Proposition 1.11, namely by requiring all prebundle charts to be homeomorphisms (and thus charts). Then E becomes indeed a vector bundle with continous projection π and bundle atlas $\mathfrak{A}$.

Now assume that E is what could be called a *differentiable prevector bundle*, meaning that—in addition to the above—the bundle transitions $1_{\psi\varphi}\colon U \times V \to U \times V$ are C^r with B an m-dimensional C^r manifold (this includes of course the case when the bundle transitions are smoother than B). The first condition can also be formulated [31[42]] by requiring that all group operations

$$U \to GL(V),\ p \mapsto 1'_{\psi\varphi}(p)$$

are C^r. Now each bundle chart $\varphi\colon E_U \to U \times V$ may be followed by a B-chart $U \to A$ going into an open set $A \subseteq \mathbb{R}^m$ and a component chart $V \to \mathbb{R}^n$ corresponding to a fixed basis in V. Combining these, one obtains an E-chart $\Phi\colon E_U \to A \times \mathbb{R}^n$ making E into an $(m+n)$-dimensional C^r manifold. Hence E is a differentiable vector bundle.

One may even go one step further [41[157]], starting *without the total space* E and the prebundle charts $\varphi\colon E_U \to U \times V$. One starts from an open cover $(U_i \mid i \in I)$ of a C^r basis manifold B having at each point $p \in B$ a group of fiber transitions $\Theta_{ji}(p)\colon V \to V$ that fulfills the cocycle condition $\Theta_{kj}\Theta_{ji} = \Theta_{ki}$ and induces a C^r map $\Theta\colon U_i \cap U_j \to GL(V)$. Then one can construct the total space E as a C^r manifold, along with a C^r bundle atlas, in a manner similar to the patchwork construction of Subsection 1.2.3.

As announced at the beginning of this section, we can now summarize the *cotangent and tangent spaces* in one convenient data structure—a vector bundle, utilizing the above construction based on a vector prebundle.

2.33 Definition *The* cotangent and tangent bundle *of an n-dimensional C^r manifold M with $r > 0$ are respectively given by the disjoint unions*

$$T^*M = \biguplus_{p\in M} T_p^*M \quad \textit{and} \quad TM = \biguplus_{p\in M} T_pM$$

over the basis M with the natural bundle projections.

Note that T^*M and TM are the *total spaces* of the vector bundle. By the definition of disjoint union, it consists of pairs (p, d) and (p, v) with basis point $p \in M$ and fiber vectors $d \in T_M^*$ and $v \in T_M$. Then the natural bundle projections $T^*M \to M$ and $TM \to M$ operate by $(p, d) \mapsto p$ and $(p, v) \mapsto p$, respectively.

Next we introduce the *bundle charts* of the cotangent and tangent bundle. Let $\mathfrak{A}$ be a differentiable structure on M and take a chart $x \in \mathfrak{A}$ with domain $U \subseteq M$ and coordinate patch $A \subseteq \mathbb{R}^n$.

One slight technicality has to be addressed at this point. Up to now we have always used charts centered at a fixed point p; for bundle charts we have to relax this restriction since we need the components for the (co)vectors in a whole neighborhood of a point. Since the components are defined in terms of the *rate* $\langle|\rangle$ at a fixed point p, it suffices to generalize $\langle|\rangle$ to non-centered charts. This is done in the obvious way: Let f be any function germ at p, and let c be any curve germ through $p = c(t)$. Then we set $\langle f|c\rangle = \langle f|\bar{c}\rangle$, where $\bar{c}$ is the centered curve germ through p defined by $\bar{c}(\tau) = c(\tau - t)$. And if x is any chart with $x(p) = x_0$, we define

$$\begin{aligned} \langle f|x^{-1}\rangle &= (f \circ x^{-1})'(x_0) = \langle f|x_i\rangle\, \delta^i, \\ \langle x|c\rangle &= (x \circ c)'(t) = \langle x^i|c\rangle\, \delta_i \end{aligned}$$

for the rate of a chart along a curve and the rate of a function along a parametrization. Using this generalized rate, a cotangent vector $d = [f]_\sim \in T_p^*M$ has a again the components $(d|_x = \langle f|x^{-1}\rangle$ and a tangent vector $v = [c]_\sim \in T_pM$ the components $\langle x|c\rangle$. Now we can define the bundle charts.

The map

$$\begin{aligned}(x|\colon\quad T^*M_U &\to U \times \mathbb{R}_n\\ (p,d) &\mapsto (p,(d|_x\ )\end{aligned}$$

is a bundle chart of T^*M. We write

$$(\mathfrak{A}| = \{(x| \mid x \in \mathfrak{A}\}$$

for the collection of bundle charts.

The map

$$\begin{aligned}|x)\colon\quad TM_U &\to U \times \mathbb{R}^n\\ (p,v) &\mapsto (p,|v)_x\ )\end{aligned}$$

is a bundle chart of TM. We write

$$|\mathfrak{A}) = \{|x) \mid x \in \mathfrak{A}\}$$

for the collection of bundle charts.

As explained above, $(\mathfrak{A}|$ and $|\mathfrak{A})$ turn T^*M and TM into differentiable vector bundles, if we can ensure certain conditions: If $\bar{x}\colon \bar{U} \to \bar{A}$ is another chart with overlap $W = U \cap \bar{U}$, the *bundle transitions* come out as follows. For any basis point $p \in M$, the transformation formulae (2.5) yield

$$1_{(\bar{x}|(x|}(p,a) = (p, a \cdot \frac{\partial x}{\partial \bar{x}}(\bar{x}_0)) \quad \text{and} \quad 1_{|\bar{x})|x)}(p,h) = (p, \frac{\partial \bar{x}}{\partial x}(x_0) \cdot h)$$

for all $a \in \mathbb{R}_n$ and $h \in \mathbb{R}^n$; here we have written $\bar{x}_0 = \bar{x}(p)$ and $x_0 = x(p)$. In other words, we have

$$1'_{(\bar{x}|(x|}(p)\, a = a\, 1'_{x\bar{x}}(\bar{x}_0) \quad \text{and} \quad 1'_{|\bar{x})|x)}(p)\, h = 1'_{\bar{x}x}(x_0)\, h$$

for the corresponding *fiber transitions*. Note that these maps are indeed linear and the group operations

$$W \to GL(T^*_pM),\ p \mapsto 1'_{(\bar{x}|(x|}(p) \quad \text{and} \quad W \to GL(T_pM),\ p \mapsto 1'_{|\bar{x})|x)}(p)$$

are C^{r-1}. Hence we may conclude that T^*M and TM are $2n$-dimensional differentiable vector bundles of class C^{r-1}.

Finally, let us interpret the *codifferential and differential* as appropriate vector bundle morphisms. Let N be another manifold and let $\Phi\colon M \to N$ be a sufficiently smooth map (and a bijection in case of the codifferential).

The map

$$\begin{aligned}d^*\Phi\colon\qquad T^*N &\to T^*M\\ (\Phi(p),\bar{d}) &\mapsto (p,(d^*_p\Phi)\,\bar{d})\end{aligned}$$

is called the codifferential of Φ.

The map

$$\begin{aligned}d\Phi\colon\quad TM &\to TN\\ (p,v) &\mapsto (\Phi(p),(d_p\Phi)\,v)\end{aligned}$$

is called the differential of Φ.

Both $d^*\Phi$ and $d\Phi$ are obviously vector bundle morphisms. Moreover, the observations at the end of Subsection 2.2.1 can now be reformulated as follows. We have the generalized uniformity relation

$$d^*\,1_M = 1_{T^*M} \quad \text{and} \quad d\,1_M = 1_{TM}$$

and the generalized chain rule

$$d^*(\Phi \circ \Psi) = d^*\Psi \circ d^*\Phi \quad \text{and} \quad d(\Psi \circ \Phi) = d\Psi \circ d\Phi$$

if Ψ is a map (of the same type) from N into another manifold. Writing $\mathfrak{Man}_r$ for the category of C^r manifolds (with $r > 0$), we may thus view the codifferential and the differential respectively as a *cofunctor and functor*

$$d^*\colon \mathfrak{Man}_r' \to \mathfrak{VecBnd}_{r-1} \quad \text{and} \quad d\colon \mathfrak{Man}_r \to \mathfrak{VecBnd}_{r-1},$$

just as promised at the outset of this section. There is just one catch in this apparent symmetry (a situation that will reappear when we consider pushforward and pullback of tensors in Subsection 2.3.4): As mentioned above, the codifferential works only for bijections, so the category $\mathfrak{Man}_r$ must be restricted to C^r diffeomorphisms; this is what the notation $\mathfrak{Man}_r'$ tries to suggest.

2.3.3 The Tensor Bundles

A *differential form* and a *vector field* assign respectively a cotangent and a tangent vector to each point of a manifold. In this subsection, we make this more precise by using the construction of the (co)tangent bundle from the previous subsection. After recalling some facts of the tensor product from Linear Algebra, we discuss the generalization to tensor fields, which assign to each point a tensor.

Let M be an n-dimensional C^r manifold with $r \geq 1$. Then we know from the previous subsection that the (co)tangent bundle is a C^{r-1} manifold with a C^{r-1} projection.

A *differential form* on M is a section

$$\omega\colon M \to T^*M$$

of the contangent bundle, meaning $\omega(p) \in T_p^*M$ for all $p \in M$.

A *vector field* on M is a section

$$\xi\colon M \to TM$$

of the tangent bundle, that is, meaning $\xi(p) \in T_pM$ for all $p \in M$.

More precisely, we speak of C^s differential form or vector field if the section is a C^s map where $s \leq r - 1$. For the sake of symmetry, differential forms are sometimes also called covector fields.

Now we discuss how differential forms and vector fields can be represented using local coordinates. Let $\mathfrak{A}$ be a differentiable structure on M and take a chart $x \in \mathfrak{A}$ with domain $U \subseteq M$ and coordinate patch $A \subseteq \mathbb{R}^n$. Then for every point $p \in U$, we have respectively the natural

$$\text{cobasis } dx^1(p), \ldots, dx^n(p) \text{ of } T_p^*M \quad \text{and} \quad \text{basis } dx_1(p), \ldots, dx_n(p) \text{ of } T_pM.$$

So for any differential form and vector field there exist respectively *component functions* $v_1, \ldots, v_n$ and $v^1, \ldots, v^n$ on U such that we obtain a local representation

$$\omega(p) = (p, v_i(p)\, dx^i(p)) \quad \text{and} \quad \xi(p) = (p, v^i(p)\, dx_i(p)).$$

It is clear that the differential form or vector field is C^s if and only if the corresponding component functions are C^s.

Let $\bar{x}\colon \bar{U} \to \bar{A}$ be another chart with overlap $W = U \cap \bar{U}$. Then we know from (2.6) the transformation laws

$$\bar{v}_i = \frac{\partial x^j}{\partial \bar{x}_i}\, v_j \quad \text{and} \quad \bar{v}^i = \frac{\partial \bar{x}^i}{\partial x_j}\, v^j,$$

for the components of a (co)tangent vector. Recall that here $\partial\bar{x}/\partial x$ denotes the transition Jacobian from x to $\bar{x}$ and $\partial x/\partial\bar{x}$ its inverse. From these equations we immediately obtain the *transformation laws*

$$\bar{v}_i(p) = \frac{\partial x^j}{\partial \bar{x}_i}(\bar{x}_0)\, v_j(p) \quad \text{and} \quad \bar{v}^i(p) = \frac{\partial \bar{x}^i}{\partial x_j}(x_0)\, v^j(p)$$

for the component functions, writing again $\bar{x}_0 = \bar{x}(p)$ and $x_0 = x(p)$.

We recall some facts on the tensor product. Let U and V be K-vector spaces. Then we know from Linear Algebra [32[602]] that there exists a K-vector spaces $U \otimes V$, called a *tensor product* of U and V, together with a bilinear map

$$\otimes\colon U \times V \to U \otimes V$$

having the following universal property: For any bilinear map $f\colon U \times V \to W$ with W being a K-vector space, there exists a unique K-linear map $\tilde{f}\colon U \otimes V \to W$ such that

$$f = \tilde{f} \circ \otimes.$$

The tensor product is unique up to isomorphism. Moreover, if U and V are finite-dimensional with bases

$$(e_i)_{1\le i\le m} \quad \text{and} \quad (f_j)_{1\le j\le n},$$

then

$$(e_i \otimes f_j)_{1\le i\le m, 1\le j\le n}$$

is a basis of $U \otimes V$, so that $\dim U \otimes V = \dim U \cdot \dim V$. Considering multilinear instead of bilinear maps, one can analogously define a tensor product of finitely many K-vector spaces.

2.34 Proposition *For finite-dimensional vector spaces, a tensor product of U and V is given by the K-vector space $L^2(U^*, V^*; K)$ of bilinear maps from $U^* \times V^*$ to K.*

Proof. To show that it is a tensor product, we have to construct a bilinear map

$$\otimes\colon U \times V \to L^2(U^*, V^*; K)$$

satisfying the universal property. For any $(u, v) \in U \times V$, we define a bilinear map

$$u \otimes v\colon U^* \times V^* \to K$$

by

$$u \otimes v(u^*, v^*) = u^*(u)\, v^*(v).$$

It is easy to see that $\otimes$ is a bilinear map. Let (e_i) and (f_j) be respectively bases for U and V with dual bases (e^k) and (f^l). Then

$$e_i \otimes f_j(e^k, f^l) = e_i(e^k)\, f_j(f^l) = \delta_i^k\, \delta_j^l,$$

so $(e_i \otimes f_j)$ is a basis of $L^2(U^*, V^*; K)$. Let now $f\colon U \times V \to W$ be a bilinear map, given by its values

$$f(e_i, f_j) = w_{ij} \in W$$

on the tuples (e_i, f_j) of basis vectors. Then $\tilde{f}$ defined by

$$\tilde{f}(e_i \otimes f_j) = w_{ij}$$

obviously satisfies $f = \tilde{f} \circ t$, and we have shown that $L^2(U^*, V^*; K)$ is indeed a tensor product of U and V. □

Analogously one can show that a tensor product for finite-dimensional K-vector spaces $V_1, \dots, V_\alpha$ is given by the K-vector space of *multilinear maps* from $V_1^* \times \cdots \times V_\alpha^*$ to K, denoted by

$$L^\alpha(V_1^*, \dots, V_\alpha^*; K).$$

Moreover, by identifying a vector space V with its bidual V^{**}, we see that a tensor product of U^* and V^* is given by $L^2(U, V; K)$. Thus we have

$$U^* \otimes V^* = L^2(U, V; K) \quad \text{and} \quad U \otimes V = L^2(U^*, V^*; K).$$

Let V be a finite dimensional K-vector space of dimension n, and α, β positive integers. The K-vector space $T^{\alpha,\beta}(V)$ of *tensors* of *valence* $\left[\begin{smallmatrix}\alpha\\ \beta\end{smallmatrix}\right]$ and *rank* $\alpha+\beta$ on V is a tensor product of β copies of V^* and α copies of V, so that

$$T^{\alpha,\beta}(V) = \underbrace{V^* \otimes \cdots \otimes V^*}_{\beta} \otimes \underbrace{V \otimes \cdots \otimes V}_{\alpha} = L^{\alpha+\beta}(\underbrace{V, \dots, V}_{\beta}, \underbrace{V^*, \dots, V^*}_{\alpha}; K).$$

Elements of $T^{\alpha,\beta}(V)$ are also called (α,β)-valent tensor or *contravariant* of order α and *covariant* of order β.

Let (e_i) be a basis of V and (e^j) the corresponding dual basis. Then we know that a basis of $T^{\alpha,\beta}(V)$ is given by

$$e^{j_1} \otimes \cdots \otimes e^{j_\beta} \otimes e_{i_1} \otimes \cdots \otimes e_{i_\alpha},$$

so $\dim T^{\alpha,\beta}(V) = n^{\alpha+\beta}$ and any tensor $t \in T^{\alpha,\beta}(V)$ can be written uniquely as

$$t = t^{i_1,\dots,i_\alpha}_{j_1,\dots,j_\beta}\, e^{j_1} \otimes \cdots \otimes e^{j_\beta} \otimes e_{i_1} \otimes \cdots \otimes e_{i_\alpha}$$

with an array $t^{i_1,\dots,i_\alpha}_{j_1,\dots,j_\beta}$ of *components* in K.

If $(\bar{e}_{\bar\imath})$ is another basis of V with the *transition matrix* $(c^{\bar\imath}_i)$ from (e_i) to $(\bar{e}_{\bar\imath})$, we have

$$e_i = c^{\bar\imath}_i \, \bar{e}_{\bar\imath}.$$

Let $(c^j_{\bar\jmath})$ denote the inverse of the transition matrix. (Precisely speaking, the transition matrix itself is an object c^- and its inverse an object c_-. In connection with these objects, one imposes the index convention that barred indices are always to be used under the barred positions of these matrices, while unbarred indices must only be used in their unbarred positions. Hence one may leave out the overbars on the indices in barred positions—which would anyway look ugly from the typographic perspective.) Then we know that the corresponding dual basis $(\bar{e}^{\bar\jmath})$ for $(\bar{e}_{\bar\imath})$ transforms with the inverse transpose of the transition matrix, so that the relation

$$e^j = c^j_{\bar\jmath} \, \bar{e}^{\bar\jmath}.$$

is fulfilled.

By multilinearity of the tensor product we obtain the *transformation law*

$$\bar{t}^{\bar\imath_1,\dots,\bar\imath_\alpha}_{\bar\jmath_1,\dots,\bar\jmath_\beta} = c^{j_1}_{\bar\jmath_1} \cdots c^{j_\beta}_{\bar\jmath_\beta} \, c^{\bar\imath_1}_{i_1} \cdots c^{\bar\imath_\alpha}_{i_\alpha} \, t^{i_1,\dots,i_\alpha}_{j_1,\dots,j_\beta} \tag{2.13}$$

for the components of a tensor.

Interpreting t as *multilinear function*

$$t\colon \underbrace{V \times \cdots \times V}_{\beta} \times \underbrace{V^* \times \cdots \times V^*}_{\alpha} \to K,$$

we have

$$t(e_{k_1}, \dots, e_{k_\beta}, e^{l_1}, \dots, e^{l_\alpha}) = t^{l_1,\dots,l_\alpha}_{k_1,\dots,k_\beta},$$

and for

$$v_1 = v_1^{k_1} e_{k_1}, \dots, v_\beta = v_\beta^{k_\beta} e_{k_\beta} \quad \text{and} \quad d^1 = d^1_{l_1} e^{l_1}, \dots, d^\alpha = d^\alpha_{l_\alpha} e^{l_\alpha}$$

we obtain by multilinearity

$$t(v_1, \dots, v_\beta, d^1, \dots, d^\alpha) = t^{l_1,\dots,l_\alpha}_{k_1,\dots,k_\beta} \, v_1^{k_1} \cdots v_\beta^{k_\beta} \, d^1_{l_1} \cdots d^\alpha_{l_\alpha}.$$

Let us collect here some additional notes about the *tensor product* of K-vector spaces $V \cong K^m$ and $W \cong K^n$. Besides the "official" way of introducing the tensor product $V \otimes W$ by way of its universal property (leading to $V \otimes W$ as a suitable quotient of the space $K^{(V \times W)}$, a construction that generalizes to arbitrary commutative rings K), there are three alternative definitions:

Basis-free: As discussed in the text, one can set $V \otimes W = L^2(V^*, W^*; K)$

Basis-invariant: One can introduce $V \otimes W$ in a fasion similar to the component spaces introduced at the end of Subsection 2.1.3.

Basis-depenent: If $x_1, \dots, x_m$ and $y_1, \dots, y_n$ are bases respectively for V and W, we can also interpret $V \otimes W$ as a suitable space of the noncommutative polynomials.

Each of these constructions of $V \otimes W$ must also provide an appropriate definition for the universal map $\otimes\colon V \times W \to V \otimes W$, and we shall see that this is indeed the case.

Depending on how the spaces V and W are given (and what one wants to do with them), one will give *preference* to one or the other. The basis-free definition is best suited if V and W are given in the usual way, namely as sets with operations like $+\colon V \times V \to V$ and $\cdot\colon \mathbb{R} \times V \to V$, subject to certain well-known axioms. If V and W are given as torsors (see Section 0.2), however, one would naturally use the basis-invariant definition. And if V and W come with specific bases, the bases-dependent definition is of course most convenient.

We have already discussed the *basis-free* definition. We define $V \otimes W = L^2(V^*, W^*; K)$ with the universal map $\otimes\colon V \times W \to V \otimes W$ is obviously by $v \otimes w\,(v^*, w^*) = v^*(v)\,w^*(w)$. See the text above for a more detailed treatment.

The other extreme is the *basis-dependent* definition. Assume V has the basis $(x_1, \dots, x_m)$ and W the basis $(y_1, \dots, y_n)$. Then we may view both V and W as subspaces of the noncommutative polynomial ring $K\langle x_1, \dots, x_m, y_1, \dots, y_n\rangle$, namely as the corresponding linear spans $V = K\langle x_1, \dots, x_m\rangle$ and $W = K\langle y_1, \dots, y_n\rangle$. Then we may set

$$V \otimes W = V \cdot W \subseteq K\langle x_1, \dots, x_m, y_1, \dots, y_n\rangle,$$

where $\cdot$ denotes the usual product of noncommutative polynomials, which sereves also as the universal map $\cdot\colon V \times W \to V \otimes W$.

This sheds more light on the intuitive meaning of tensors: They are just the noncommutative analogs of mixed homogeneous polynomials of degree two (also called bilinear forms!), and this can be continued. For any vector space V, its *tensor algebra* is given by the direct sum of arbitrarily high tensor products $V^{\otimes i} = \underbrace{V \otimes \ldots \otimes V}_{i}$ as

$$\mathfrak{T}(V) = \bigoplus_{n=0}^{\infty} V^{\otimes i},$$

so we have $\mathfrak{T}(V) = K\langle x_1, \dots, x_m\rangle$ where $(x_1, \dots, x_m)$ is again a basis for $V = K\langle x_1, \dots, x_m\rangle$. If $(y_1, \dots, y_n)$ is a basis for the vector space $W = K\langle y_1, \dots, y_n$, we see that

$$V \oplus W = K\langle x_1, \dots, x_m, y_1, \dots, y_n\rangle$$

so $V \otimes W \leq \mathfrak{T}(V \oplus W)$ is just that part of the tensor algebra that combines exactly one "indeterminate" from V with one "indeterminate" from W. As we said before: Mixed homogeneous polynomials of degree two.

For the *basis-invariant definition*, we use the language of torsors (see Section 0.2). So let V and W be contravariant component spaces with "abstract bases" B and C with dimensions m and n, respectively. Then every $v \in V$ is an equivariant map $v\colon B \to \mathbb{R}^n$ and equally every $w \in W$ an equivariant map $w\colon C \to \mathbb{R}^n$. Now we consider the "abstract bases" $B \times C$ and the component array $\mathbb{R}^{m \times n}$ as torsors over $GL_m(\mathbb{R}) \times GL_n(\mathbb{R})$. For the component array, we have to view $m = \{0, 1, \dots, m-1\}$ and $n = \{0, 1, \dots, n-1\}$ in their ordinal character. On $B \times C$, the action of $GL_m(\mathbb{R}) \times GL_n(\mathbb{R})$ is given by the torsor product. For $M \in \mathbb{R}^{m \times n}$ we define the action of $(S, T) \in GL_m(\mathbb{R}) \times GL_n(\mathbb{R})$ by

$$(S, T) \cdot M = S^{\top} M T;$$

note that this is just the usual transformation rule for the components of a bilinear form [32[528]]. We can now define $V \otimes W$ to be the set of all equivariant maps between the $GL_m(\mathbb{R}) \times GL_n(\mathbb{R})$ torsors $B \times C$ and $\mathbb{R}^{m \times n}$. This gives a vector space of dimension mn, as one may easily check.

The universal map $\otimes\colon V \times W \to V \otimes W$ is constructed as follows. Given two equivariant maps $v\colon B \to \mathbb{R}^m$ and $w\colon C \to \mathbb{R}^n$, we construct the map

$$\begin{aligned} v \otimes w\colon \quad & B \times C \to \mathbb{R}^{m \times n}, \\ & (\beta, \gamma) \mapsto \Big((i,j) \mapsto v_i(\beta)\, w_j(\gamma)\Big) \end{aligned}$$

so that $v \otimes w$ behaves just like a bilinear form should behave. Hence $v \otimes w$ is indeed an equivariant map $v \otimes w \in V \otimes W$.

Of course, one need not stop at tensoring two vector spaces; the same approach works for defining $V_1 \otimes \cdots \otimes V_\alpha$, yielding tensors of valence $\left[\begin{smallmatrix}0\\ \alpha\end{smallmatrix}\right]$. In fact, one can immediately introduce tensors of arbitrary mixed valence $\left[\begin{smallmatrix}\alpha\\ \beta\end{smallmatrix}\right]$, and this is where the torsor construction shows its benefits (assuming $V_1 = \ldots = V_\alpha$ for simplicity): The construction is completely symmetric in co- and contravariant components. Indeed, assume a dual pair of vector spaces $(V^*, (|), V)$ is given as the space of equivariant maps from a set B of "abstract bases" to $\mathbb{R}_n$ and to $\mathbb{R}^n$, describing the co- and contravariant vectors, respectively. Then a tensor of valence $\left[\begin{smallmatrix}\alpha\\ \beta\end{smallmatrix}\right]$ is just an equivariant map

$$\tau\colon B \to \mathbb{R}^{n^\alpha \times n^\beta},$$

where

$$n^\alpha = \Big\{\kappa \mid \kappa\colon \{0, 1, \ldots, \alpha - 1\} \to \{0, 1, \ldots, n-1\}\Big\}$$

and likewise n^β are to be understood in the sense of ordinals. The action of a transition matrix $c = c^- \in GL_n(\mathbb{R})$ on a component array $t \in \mathbb{R}^{n^\alpha + n^\beta}$ results in another component array $c \cdot t = \bar{t} \in \mathbb{R}^{n^\alpha + n^\beta}$ given by

$$\bar{t}^{\bar{\imath}_1, \ldots, \bar{\imath}_\alpha}_{\bar{\jmath}_1, \ldots, \bar{\jmath}_\beta} = c^{\bar{\imath}_1}_{i_1} \cdots c^{\bar{\imath}_\alpha}_{i_\alpha}\, c^{j_1}_{\bar{\jmath}_1} \cdots c^{j_\beta}_{\bar{\jmath}_\beta}\, t^{i_1, \ldots, i_\alpha}_{j_1, \ldots, j_\beta},$$

which is just the tensor transformation law (2.13), again using c_- for the inverse of c^- in the way described above.

We can now apply the construction of the tensor product to the tangent space of a manifold and define tensor fields, which assign to each point of the manifold a tensor. We first have to generalize the definition of the (co)tangent bundle to tensor bundles. Recall that

$$T_p^{\alpha,\beta} M := \underbrace{T_p^* M \otimes \cdots \otimes T_p^* M}_{\beta} \otimes \underbrace{T_p M \otimes \cdots \otimes T_p M}_{\alpha}$$

is an $n^{\alpha+\beta}$-dimensional vector space.

2.35 Definition *The* tensor bundle *of valence* $\left[\begin{smallmatrix}\alpha\\ \beta\end{smallmatrix}\right]$ *of an n-dimensional C^r manifold M with $r > 0$ is given by the disjoint union*

$$T^{\alpha,\beta} M = \biguplus_{p \in M} T_p^{\alpha,\beta} M$$

over the basis M with the natural bundle projection.

Similar to Subsection 2.3.2, we can turn $T^{\alpha,\beta}M$ into $(n+n^{\alpha+\beta})$-dimensional *differentiable vector bundle* over M of class C^{r-1}. Note also that

$$T^{1,0}M = TM \quad \text{and} \quad T^{0,1}M = T^*M$$

recovers the (co)tangent bundle.

A *tensor field* of valence $\begin{bmatrix}\alpha\\ \beta\end{bmatrix}$ on M is a section

$$\tau\colon M \to T^{\alpha,\beta}M$$

of the tensor bundle $T^{\alpha,\beta}M$, meaning $\tau(p) \in T_p^{\alpha,\beta}M$ for all $p \in M$. More precisely, we speak of C^s tensor field if the section is a C^s map where $s \leq r-1$.

We can represent tensor fields using local coordinates, as before differential forms and vector fields. Let $x\colon U \to A$ be a chart. Then there exist *component functions* $t^{i_1,\dots,i_\alpha}_{j_1,\dots,j_\beta}$ on U such that we obtain a local representation

$$\tau(p) = (p, t^{i_1,\dots,i_\alpha}_{j_1,\dots,j_\beta}(p)\, dx^{j_1} \otimes \cdots \otimes dx^{j_\beta} \otimes dx_{i_1} \otimes \cdots \otimes dx_{i_\alpha}),$$

where we omit the dependence on p in the basis vectors. It is clear that the tensor field is C^s if and only if the corresponding component functions are C^s.

Let $\bar{x}\colon \bar{U} \to \bar{A}$ be another chart with overlap $W = U \cap \bar{U}$. Then from the transformation laws (2.6) for the (co)tangent vectors and (2.13) for the components of a tensor, we obtain the *transformation laws*

$$\bar{t}^{\bar{\imath}_1,\dots,\bar{\imath}_\alpha}_{\bar{\jmath}_1,\dots,\bar{\jmath}_\beta} = \frac{\partial x^{j_1}}{\partial \bar{x}_{\bar{\jmath}_1}} \cdots \frac{\partial x^{j_\beta}}{\partial \bar{x}_{\bar{\jmath}_\beta}} \frac{\partial \bar{x}^{\bar{\imath}_1}}{\partial x_{i_1}} \cdots \frac{\partial \bar{x}^{\bar{\imath}_\alpha}}{\partial x_{i_\alpha}}\, t^{i_1,\dots,i_\alpha}_{j_1,\dots,j_\beta} \tag{2.14}$$

for the component functions of a tensor, where everything depends on p.

The set of all sections of the tensor bundle of valence $\begin{bmatrix}\alpha\\ \beta\end{bmatrix}$ is denoted by $\mathcal{T}^{\alpha,\beta}M$. Note that $\mathcal{T}^{0,0}M$ is just the sheaf of differentiable functions on M; one may check that every $\mathcal{T}^{\alpha,\beta}M$ is also a sheaf, accordingly called the *sheaf of tensor fields* of valence $\begin{bmatrix}\alpha\\ \beta\end{bmatrix}$. The special cases of vector fields and differential forms will be denoted by $\Xi M = \mathcal{T}^{1,0}M$ and $\Omega M = \mathcal{T}^{0,1}M$, respectively.

2.3.4 Pullback and Pushforward

We now discuss the pullback and pushforward of tensor fields for a differentiable map. Again we first consider the situation in Linear Algebra, so let $A\colon V \to W$ be a linear map between finite dimensional K-vector spaces. The dual map

$$A^*\colon W^* \to V^*$$

is defined by $w^* \mapsto w^* \circ A$. Note that $W^* = T^{0,1}(W)$ and $V^* = T^{0,1}(V)$. We can generalize this idea. The dual of a linear map induces the *pullback* on covariant tensors of order β, which is again denoted by A^*. More precisely, we have the linear map

$$A^*\colon T^{0,\beta}(W) \to T^{0,\beta}(V)$$

defined by

$$e^{j_1} \otimes \cdots \otimes e^{j_\beta} \mapsto A^* e^{j_1} \otimes \cdots \otimes A^* e^{j_\beta}.$$

Analogously, A also induces a map

$$A_* \colon T^{\alpha,0}(V) \to T^{\alpha,0}(W),$$

defined by

$$e_{i_1} \otimes \cdots \otimes e_{i_\alpha} \mapsto A e_{i_1} \otimes \cdots \otimes A e_{i_\alpha}$$

and known as the *pushforward* on contravariant tensors of order α.

If A is an isomorphism, we can define respectively the pushforward of covariant tensors of order β as the pullback of A^{-1} and the pullback of contravariant tensors of order α as the pushforward of A^{-1}. So for an isomorphism A, we can define the *pullback*

$$A^* \colon T^{\alpha,\beta}(W) \to T^{\alpha,\beta}(V)$$

of mixed tensors by

$$e^{j_1} \otimes \cdots \otimes e^{j_\beta} \otimes e_{i_1} \otimes \cdots \otimes e_{i_\alpha} \mapsto A^* e^{j_1} \otimes \cdots \otimes A^* e^{j_\beta} \otimes A^{-1} e_{i_1} \otimes \cdots \otimes A^{-1} e_{i_\alpha}$$

and the *pushforward*

$$A_* \colon T^{\alpha,\beta}(V) \to T^{\alpha,\beta}(W)$$

by

$$e^{j_1} \otimes \cdots \otimes e^{j_\beta} \otimes e_{i_1} \otimes \cdots \otimes e_{i_\alpha} \mapsto (A^{-1})^* e^{j_1} \otimes \cdots \otimes (A^{-1})^* e^{j_\beta} \otimes A e_{i_1} \otimes \cdots \otimes A e_{i_\alpha},$$

containing the previous definitions as the special cases $\alpha = 0$ and $\beta = 0$.

We can also apply the pullback and the pushforward pointwise to *tensor fields* along a diffeomorphism $\Phi \colon M \to N$, where the differential $d\Phi \colon TM \to TN$ plays the role of the linear map A and accordingly the codifferential $d^*\Phi \colon T^*N \to T^*M$ the role of the dual map A^*. Writing the arguments of Φ as subscripts for readability, this yields

$$\begin{aligned} \Phi^*\tau \colon \quad & T^{\alpha,\beta}N \to T^{\alpha,\beta}M \\ & (\Phi_p, \tau) \mapsto (p, d_p^* \tau(p)) \end{aligned}$$

for the pullback of Φ and

$$\begin{aligned} \Phi_*\tau \colon \quad & T^{\alpha,\beta}M \to T^{\alpha,\beta}N \\ & (p, \tau) \mapsto (p, d_p \tau(\Phi_p)) \end{aligned}$$

for the pushforward of Φ. If one takes a chart/parametrization for Φ, one may use the pullback and pushout of tensors for extracting their components with respect to Φ in complete analogy to Proposition 2.24.

Let us analyze the important special case of *differential forms* $\omega \colon M \to T^*M$ and *vector fields* $\xi \colon M \to TM$ a bit closer. The diagram below illustrates the definition of the pushforwards $\Phi_*\xi$ and $\Phi^*\omega$ as well as the pullbacks $\Phi^*\xi$ and $\Phi^*\omega$. Up to now we have always assumed that Φ is bijective. In the diagrams we see that this is actually needed—except in the one case of $\Phi^*\omega$, because we can compute $\Phi^*\omega = d^*_\Phi \circ \omega \circ \Phi$ without inverting either the map itself or its linearization. In the other three cases, we need the inverse of either one of them or both. Note that the failure for $\Phi^*\xi$ and $\Phi_*\omega$ is already apparent for linear maps A as discussed above; but the failure for $\Phi_*\xi$ and $\Phi_*\omega$ only comes from not being able to "transport" the tensors to the right points of M, a phenomenon that we have already encountered in the definition of the codifferential on the cotangent bundle.

$$\begin{tikzcd}
M \arrow[r, "\Phi"] \arrow[d, "\xi"'] & N \arrow[d, dotted, "\Phi_*\xi"] \\
TM \arrow[r, dotted, "d\Phi"'] & TN
\end{tikzcd}
\qquad
\begin{tikzcd}
M \arrow[r, dotted, "\Phi"] \arrow[d, dotted, "\Phi^*\xi"'] & N \arrow[d, "\xi"] \\
TM \arrow[r, "d\Phi"'] & TN
\end{tikzcd}$$

$$\begin{tikzcd}
M \arrow[r, "\Phi"] \arrow[d, "\omega"'] & N \arrow[d, dotted, "\Phi_*\omega"] \\
T^*M \arrow[r, dotted, "(d^*\Phi)^{-1}"'] & T^*N
\end{tikzcd}
\qquad
\begin{tikzcd}
M \arrow[r, dotted, "\Phi"] \arrow[d, dotted, "\Phi^*\omega"'] & N \arrow[d, "\omega"] \\
T^*M \arrow[r, "(d^*\Phi)^{-1}"'] & T^*N
\end{tikzcd}$$

The above diagrams have been drawn in the style of *categorical pullbacks and pushforwards* (actually the latter is then usually termed "pushout"). The idea is that the morphisms represented by the solid arrows are regarded as given (together with the three objects at their vertices) while the ones represented by dotted arrows are to be found (together with the missing vertex). For example, the upper left diagram would describe the pushout square for $\Phi\colon M \to N$ and $\xi\colon M \to TM$, yielding the "fibered sum" TN, and $d\Phi\colon TM \to TN$ for the pushout of Φ along ξ as well as $\Phi_*\xi\colon N \to TN$ for the pushout of ξ along Φ. Likewise, the lower right diagram would appear as the pullback square for $\Phi\colon M \to N$ and $(d^*\Phi)^{-1}\colon T^*M \to T^*N$, yielding the "fibered product" M, and $\Phi\colon M \to N$ for the pullback of $(d^*\Phi)^{-1}$ along ω as well as $\Phi^*\omega$ for the pullback of ω along $(d^*\Phi)^{-1}$.

The *situation in manifolds* is somewhat degenerate as compared with general categories because the horizontal arrows determine each other: Given a diffeomorphism $\Phi\colon M \to N$, its differential $d\Phi\colon TM \to TN$ and codifferential $d^*\Phi\colon T^*N \to T^*N$ are fixed (and bijective again); conversely, one can of course extract Φ from the fiber maps $d\Phi$ and $d^*\Phi$. Hence the missing vertical morphism is the only interesting thing (and thus the only one named "pullback" or "pushforward"), and it is obtained by the appropriate composition of Φ, its linearization, and the given (co)vector field.

We recall from Subsection 2.2.1 that the differential and codifferential were constructed by "pushing forward" curve germs and codifferential by "pulling back" function germs, respectively.

As the names suggest, there is some connection between these notions of pushforward/pullback and the ones we have been studying now. As before, we consider manifolds M and N, a diffeomorphism $\Phi\colon M \to N$, a vector field $\xi\colon M \to TM$, and a differential form $\omega\colon N \to T^*N$. Fixing a point $p \in M$ with image $q = \Phi(p) \in N$, the pushforward is defined by

$$(\Phi_*\xi)_q = (d_p\Phi)\,\xi_p$$

and the pullback by

$$(\Phi^*\omega)_p = (d_p^*\Phi)\,\omega_q,$$

where we have used the conventions $\xi(p) = (p, \xi_p)$ and $\omega(q) = (q, \omega_q)$, with their analogs on the left-hand side. As we can see from this, the pushforward of a vector field is a pointwise application of the differential (a pushforward of a single vector), and the pullback of a covector field accordingly a pointwise application of the codifferential (a pullback of a single covector). In other words, the pushforward/pullback of vector/covector fields is a *fiberwise pushforward/pullback* of single vectors/covectors.

In order to make these relations more explicit, let us first consider the pushforward. We know [31[88]] that there is an *integral curve* $c\colon \mathbb{R} \to M$ locally through $p = c(0)$, meaning

$$\xi(c(t)) = d_t c = (c(t), c'(t))$$

or briefer $\xi_{c(t)} = c'(t)$ for all $t \in \mathbb{R}$. (Recall the convention of identifying tangent vectors with curve differentials as explained after Proposition 2.23.) For a curve point $p = c(t)$, we can then compute

$$(\Phi_*\xi)_q = (d_p\Phi)\,[c_p]_\sim = [(\Phi \circ c)_p]_\sim = (\Phi \circ c)'(t) \tag{2.15}$$

for the pushforward of the tangent vectors. Note here the different meanings of the subscripts: For vector fields, the subscript is defined in the previous paragraph, for the differential it denotes the base point of the fiber, and for curves it passes to the curve germ. In equation 2.15 we can see that the effect of the pushforward of a vector field is to "push forward" its integral curves in the sense of Subsection 2.2.1.

We can express (2.15) more succinctly by introducing the *tangent field* of a curve $c\colon \mathbb{R} \to M$ as the vector field

$$\begin{aligned} C &\to TC \\ c(t) &\mapsto d_t c = (c(t), c'(t)) \end{aligned}$$

on the submanifold $C = c(\mathbb{R}) \subseteq M$. We will write c' for this vector field (the context will remove any ambiguity). Then c is an integral curve iff $c' = \xi|_C$, and equation (2.15) may be restated as $\Phi_*\, c' = (\Phi \circ c)'$. Since every curve may be seen as part of some vector field (its restriction), we can forget about ξ. Then we can summarize equation 2.15 as follows: The tangent field of a curve is pushed forward to the tangent field of its image curve.

We can do analogous things for the pullback. Given a differential form $\omega\colon N \to T^*N$, there is a local *potential function*, meaning

$$\omega(q) = d_q f = (q, f'(q))$$

or briefer $\omega_q = f'(q)$ for all $q \in N$ in a sufficiently small open region. (Recall the conventions of identifying cotangent vectors with function codifferentials, also discussed after Proposition 2.23) Writing the preimage points again $p = \Phi^{-1}(q)$, we can now compute

$$(\Phi^*\omega)_p = d_p^*\Phi\,[f_q]_\sim = [(f \circ \Phi)_q]_\sim = (f \circ \Phi)'(q) \tag{2.16}$$

for the pullback of the cotangent vectors. In other words, the pullback of a covector field "pulls back" its potential functions in the sense of Section 2.2.1.

In order to express equation (2.16) in a compact manner, we introduce the *gradient form* of a function $f\colon N \to \mathbb{R}$ as the differential form

$$\begin{aligned} N &\to T^*N \\ q &\mapsto d_q^* f = (q, f'(q)). \end{aligned}$$

Again we will write f' for this differential form (relying on the context for resolving ambiguity), so f is a potential function iff $f' = \omega$. Now equation (2.16) can be stated as $\Phi^* f' = (f \circ \Phi)'$ and verbalized thus: The gradient form of a function is pulled back to the gradient form of its image function.

2.3.5 Various Tensor-like Quantities

Tensors are an incredibly flexible tool for describing many *geometrical and physical structures* occurring in practice; this is why one finds a stupendous array of tensor-like objects with all sorts of (all too often inconsistent) terminology. Here are some examples from geometry: length, area, volume, angle, curvature, torsion. Here are some examples from physics: electro-magnetic fields and fluxes, stress and strain, stress-energy in relativity.

For some of these objects, one needs certain variations of tensors. We can try to put some order into them by looking at their *behavior under coordinate changes* in an n-dimensional manifold M. Locally at a point $p \in M$, such a coordinate change is always linear (realized by the transition Jacobian), so we are really studying $GL(V)$ with $V = T_pM$. Roughly speaking, we can distinguish the following features in overview (more on them later):

Translations: It is maybe stupid to mentions this here, but just to make sure: The fact that we linearize chart transitions by their Jacobian includes a hidden transportation to the origin of the tangent space. Hence translations are preserved ab initio—the universe is *isotropic*.

Rotations: Every tensor-like quantity (even its most fancy variations) should respect rotations. This is the core idea of all tensors, expressed in their transformation law: We may violate this law for various special transformations, but never for rotations—the universe is *isotropic*.

Dilations: Passing to *tensors densities*, we may get extra determinant factors when rescaling some of its components. Geometrically, this means that the tensor spreads out transversally.

Reflections: Working with *pseudo-tensors*, we have to flip sign if orientation is reverses; combined with densities, this brings in determinant factors with absolute value. Geometrically, the easiest example is the cross product of vectors.

Permutations: For quantitites—think of areas, volumes, angles—that extend in more than one direction (more precisely: tensors with rank above one), the order of these directions (more precisely: the order of the V and V^* factors in the definition of the tensor bundle) may have a certain influence: The two extreme cases are: *symmetry* (nothing happens) and *antisymmetry* (the sign flips).

Note that the first four operations are the basic building blocks of $GL(V)$, while the last one allows to let $GL(V)$ act on more than just single vectors. All these choices are independent of each other, and they are also exhaustive in the following sense: Any combination leads to some variation $\tilde{T}$ of the plain tensor bundle T, and $GL(V)$ acts on $\tilde{T}$ by way of a *linear representation.* Conversely, any representation of $GL(V)$ can be obtained in this fashion. (This relation seems to be folklore among mathematical physicists, but unfortunately I have not found any references.) The whole menagerie is seldom layed out, although some physics book [50[36]] do mention all of the relevant "geometric objects."

Let us start with the action of permuations. A tensor T of rank called pure if it is either fully contravariant (having valence $\left[\begin{smallmatrix}\alpha\\0\end{smallmatrix}\right]$ then) or fully covariant (having valence $\left[\begin{smallmatrix}0\\\beta\end{smallmatrix}\right]$ then). For reasons of simplicity, we will restrict ourselves to these cases, viewing all tensors as multilinear maps into $\mathbb{R}$. Then we call T *symmetric* if T remains invariant under all permutation of its arguments and *alternating* (also called skew- or anti-symmetric) if T changes sign for alternating permutations (one may obviously reduce these requirements to transpositions). As one sees immediately, this gives subspaces of the corresponding plain tensor spaces. In analogy to the usual contravariant and covariant tensors, this gives respectively symmetric tensor bundles $\Sigma^\alpha M$ and $\Sigma^\beta_* M$ as well as alternating tensor bundles $\Lambda^\alpha M$ and $\Lambda^\beta_* M$. In their component arrays, symmetric/alternating tensors display an analogous behavior, e.g. (skew)symmetric matrices in rank two; in fact, this was is the old way of defining symmetric/alternating tensors—together with the remark that the (skew)symmetry remains invariant under coordinate changes.

An arbitrary tensor may be *symmetrized* or *alternated* in the obvious manner (see [2[202]] for details). This means we have projection maps Sym and Alt, both defined on $\mathcal{T}^{\alpha,0}M$ as well as $\mathcal{T}^{0,\beta}M$, such that

$$\begin{aligned} \operatorname{Sym}(\mathcal{T}^{\alpha,0}M) &= \Sigma^\alpha M \quad \text{and} \quad \operatorname{Sym}(\mathcal{T}^{0,\beta}M) = \Sigma^\beta_* M, \\ \operatorname{Alt}(\mathcal{T}^{\alpha,0}M) &= \Lambda^\alpha M \quad \text{and} \quad \operatorname{Alt}(\mathcal{T}^{0,\beta}M) = \Lambda^\beta_* M. \end{aligned}$$

A tensor τ is symmetric iff $\operatorname{Sym}(\tau) = \tau$ and alternating iff $\operatorname{Alt}(\tau) = \tau$. Furthermore, Sym and Alt commute with pullback/pushforward (this means we can pull them back and push them forward just as if they were plain tensors). In components, one indicates symmetrization/alternation by putting round/square brackets around the indices [45[238]]. One may also multiply symmetric/alternating tensors by usual tensor multiplication with subsequent symmetrization/alternation, resulting in a tensor with rank the sum of ranks.

An alternative approach to symmetric/alternating tensors is via *quotient spaces* of the plain tensors: One simply collapses the subspace generated by all commutators/anticommutators of the basis elements. This is not essentially different from the selective approach outlined above; in fact, we have just selected canonical representatives from the equivalence classes of the quotient.

Of course, the symmetric/alternating tensor bundles are of a lower dimension than the plain ones. In fact, one may restrict the *tensor basis vectors* (being tensor products of basis vectors) for example to weakly/strictly ascending ones (meaning the indices of the contributing basis vectors form a monotonically increasing / strictly increasing sequence). Hence we have

$$\dim \Sigma^k M = \dim \Sigma^k_* M = \binom{n+k-1}{k}, \qquad \dim \Lambda^k M = \dim \Lambda^k_* M = \binom{n}{k}$$

basis vectors in the respective cases.

There are two important *examples* that equip the manifold M with additional structure: Metrics are elements of $\Sigma^2_* M$, while symplectic forms are from $\Lambda^2_* M$. In physics, we may quote stress and strain or the moment of inertia, all coming from $\Sigma^2_* M$ as well.

The alternating structures are more important than the symmetric ones. An alternating tensor of rank α is called a *multivector* or rank α or briefly an α-vector (for $\alpha = 1$ we get back the old vectors); an alternating form of rank β a *multiform* of rank β or briefly a β-form (for $\beta = 1$ we get back the old linear forms). In the vector case, we may also speak of monovectors, bivectors, trivectors and the like for $\alpha = 1, 2, 3, \ldots$; the anlog for multiforms seems to be uncommon—and the name "multiform" as well. The sections of $\Lambda^\alpha M$ are accordingly called *α-vector fields*, and their corresponding sheaf is denoted by $\Omega^\alpha M$; the sections of $\Lambda^\beta_* M$ are then named the *differential β-forms*, and their sheaf is denoted by $\Xi^\beta M$.

As an aid for *visualization* one may take the models given for $M = \mathbb{R}^3$ in the beautiful booklet [56]. Then we may picture vectors of course as "arrows", forms as "stacks" of equally spaced sheets with a big arrow for orienting them, bivectors as "thumbtacks" with designated tip direction, and biforms as an "grid" of equally spaces wires with a specific current direction. Note that all these four objects have three components (since they are tensors of rank 1 on a manifold of dimension 3), while the behavior of their components under transformation is very different! Indeed, expanding the space uniformly shrinks an arrow shaft but condenses the sheets of a stack. Expanding the space along a thumbtack or a grid has no effect, but expanding it transversally shrinks the thumbtack head and condenses the grid lines.

As explained above, multivectors and multiforms can be multiplied: an α_1-vector times an α_2-vector gives an $(\alpha_1+\alpha_2)$-vector, and a β_1-form times a β_2-form gives a $(\beta_1+\beta_2)$-form. Hence the direct sum of all multivectors/multiforms gives an algebra, called the *exterior algebra* of vectors/forms. This is important for the integration theory behind the general Stokes theorem, where k-chains serve as the global analogs for k-vectors and k-integrals for k-forms.

For *example*, assume we want to integrate a “vector field” (doing a biintegral) over a “surface patch” (a bichain) in space. An infinitesimal piece of the patch is spanned by two vectors; since only are and orientation counts for integrating whatever quantity on the surface patch, we may think of the path as the bivector formed by the product of the two vectors. The integrand is a scalar $f(p)$ depending on the point p of the surface patch and the product of two differentials dx and dy; the precise formulation is that both dx and dy are actually forms (stacks of sheets), whose product gives a biform (the grid of wires at the sheet intersections); then the biform may evaluate the bivector (counting the wires traversing the surface patch).

The n-forms $\omega \in \Lambda^n_* M$ have a special role to play. If the manifold is orientable, they can be used for determining the volume—this is why they are also called *volume forms*. A manifold is orientable iff it has volume form that vanishes nowhere. Note that $\Lambda^n_* M$ is a line bundle (also known as the “determinant bundle”), just like $\Lambda^0_* M \cong C^r(M)$, so a volume form is somehow like a scalar field. But the difference is that—unlike an honest scalar field—a volume form expands if the basis vectors expand: after all, it is supposed to measure out the volume spanned by the basis vectors! There is also an analog for n-vectors, forming another line bundle (let us call it the “codeterminant bundle”).

Now let us go to *tensor densities*, also called “densitized” or “weighted” tensors. For a basis-free definition of the corresponding bundles, one proceeds most conveniently by forming the tensor product with certain line bundles of *scalar densitites*, also called “densitized” or “weighted” scalars. For any $\rho \in \mathbb{R}$, a scalar density of weight ρ is a map $\mu\colon \Lambda^n_* M \to \mathbb{R}$ with the property that for all $\lambda \in \mathbb{R}^\times$ one has

$$\mu(\lambda\,\omega) = \lambda^{-\rho/n}\,\mu(\omega) \quad \text{or} \quad \mu(\lambda\,\omega) = |\lambda|^{-\rho/n}\,\mu(\omega);$$

such a densitiy is then called *relative* or *absolute*, respectively. In the literature, the relative densities are mostly just called densities, whereas the absolute ones are often called pseudo-tensors.

We write ${}^\pm\Pi^\rho M$ for the space of all ρ-weighted scalar densities at p with $+$ designating the relative and $-$ the absolute ones. Besides the usual laws for the signs, one obtains the isomorphisms $\Pi^{\rho_1} M \otimes \Pi^{\rho_2} M = \Pi^{\rho_1+\rho_2} M$ and $(\Pi^\rho M)^* = \Pi^{-\rho} M$. The scalar densities of weight 0 are obviously the old scalars. The product $\Pi^n M \otimes {}^-\Pi^{-n} M$ is called the pseudo-bundle; its elements, accordingly called *pseudo-scalars*, are like ordinary scalars except that they flip sign when the orientation of the basis is reversed. In forming tensor products, of relative/absolute tensor densities, they act like a minus sign (hence one needs only relative or absolute tensor densitites once the pseudo-scalars are available).

For the common case of weigths $\rho = nl$ with $l \in \mathbb{Z}$, there is an additional interpretation of tensor densities: twisting l times with the determinant bundle for $l \geq 0$ and with the codeterminant bundle for $l \leq 0$. (Forming the tensor product with a line bundle is often called “twisting with such-and-such”.) Another important relation is given by the *Weinreich duality* [52[23]], which associates k-forms with k'-vector densities where $k' = n-k$. (One often leaves out the specification of the weight if $l = 1$.) Note that we have silently made this identification above when we described the visualization of bivectors and biforms.

In *coordinates*, the effect of these twists is just what one expects: The tensor transformation laws (2.14) gets an additional factor $(\partial\bar{x}/\partial x)^\rho$ for relative and $|\partial\bar{x}/\partial x|^\rho$ for absolute tensor densities of weight ρ. Tensor densities are seldom treated in the literature—if at all, mostly by this coordinate characterization [47], [57].

Pseudo-tensors are important for *integration* on non-oriented manifolds. This is in particular true for rank zero: pseudo-scalar densities (the "function" that one wants to integrate), called "densities" in [31[304]]. Instead of a volume form, one needs a *volume pseudo-form* in order to determine the volume of a non-orientable manifold (for example the area of the Möbius band). On an orientable manifold, there is a bijective correspondence between volume pseudo-forms and volume forms, so one does not notice the need of pseudo-scalars there.

In *physics*, there is one famous group of quantities, namely the fundamental fields of electro-magnetism as they appear in the Maxwell equations (everything alternating): The electric field E is a monoform, the magnetic field H a biform, the electric flux D a pseudo-biform, and the magnetic flux B a pseudo-monoform. (Note that in many texts, all these quantities are sometimes just termed "vectors"—just because they can be described by real numbers. Or because one may identify all four quantities if one has additional structure like metrics. Or because one restricts oneself to orthogonal coordinate systems so that the transition Jacobians are always rotation matrices.) Another case in point is the physicists' distinction between *axial and polar* vectors: White the latter are plain old vectors, the former ones are pseudo-vectors: vectors twisted by the pseudo-bundle.

In the presence of additional structure, there are numerous other ways of identifying some of these objects. For example, in the presence of a Riemannian metric, one may apply *Hodge duality* for associating k-forms with k'-pseudo-forms or—in case of an oriented manifold—even with k-forms, where $k' = n - k$. This is where the theory of manifolds opens itself into the endless mathematical ocean of ever new structures ...

Bibliography

[1] Ralph Abraham, Jerrold E. Marsden, and Tudor Raţiu. *Manifolds, tensor analysis, and applications*, volume 2 of *Global Analysis Pure and Applied: Series B*. Addison-Wesley, Reading (Massachussettes), 1983.

[2] William M. Boothby. *An introduction to differentiable manifolds and Riemannian geometry*, volume 63 of *Pure and Applied Mathematics*. Academic Press, New York, 1975.

[3] Siegfried Bosch. *Lineare Algebra*. Springer Lehrbuch. Springer, second edition, 2003.

[4] N. Bourbaki. *Topological vector spaces. Chapters 1–5*. Elements of Mathematics (Berlin). Springer, Berlin, 1987. Translated from the French by H. G. Eggleston and S. Madan.

[5] Nicolas Bourbaki. *Algebra I. Chapters 1–3*. Elements of Mathematics (Berlin). Springer-Verlag, Berlin, 1998. Translated from the French, Reprint of the 1989 English translation [MR0979982 (90d:00002)].

[6] Nicolas Bourbaki. *General topology. Chapters 1–4*. Elements of Mathematics (Berlin). Springer-Verlag, Berlin, 1998. Translated from the French, Reprint of the 1989 English translation.

[7] Glen E. Bredon. *Topology and geometry*, volume 139 of *Graduate Texts in Mathematics*. Springer, New York, 1993.

[8] T. Bröcker and K. Jänich. *Einführung in die Differentialtopologie*, volume 143 of *Heidelberger Taschenbücher*. Springer, reprinted 1990 edition, 1973.

[9] Theodor Bröcker. *Analysis II*. Bibliographisches Institut, Mannheim, 1992.

[10] Theodor Bröcker. *Analysis III*. Bibliographisches Institut, Mannheim, 1992.

[11] J. Dieudonné. *Foundations of modern analysis Vol. 3*, volume 10 of *Pure and Applied Mathematics*. Academic Press, New York, 1969.

[12] C. T. J. Dodson and T. Poston. *Tensor geometry. The geometric viewpoint and its uses*, volume 130 of *Graduate Texts in Mathematics*. Springer, Berlin, second edition, 1991.

[13] Albrecht Dold. *Lectures on Algebraic Topology.* Classics in Mathematics. Springer, Berlin, 1995. Reprint of the 1972 edition.

[14] Harley Flanders. *Differential forms with applications to the physical sciences.* Academic Press, New York, 1963.

[15] Otto Forster. *Riemannsche Flächen*, volume 184 of *Heidelberger Taschenbücher.* Springer, Berlin, 1977.

[16] Jean Gallier. Notes on group actions, manifolds, lie groups and lie algebras. Book in progress.

[17] David Gauld. Metrisability of manifolds. Technical report, University of Auckland, New Zealand, October 2003. Electronic Notes.

[18] Lothar Gerritzen. *Grundbegriffe der Algebra. Eine Einführung unter Berücksichtigung funktorieller Aspekte.* Vieweg, Braunschweig, 1994.

[19] G. W. Gibbons. Applications of differential geometry to physics. Technical report, Cambridge University, Great Britain, January 2006. Lecture Notes.

[20] Robin Hartshorne. *Algebraic geometry.* Springer-Verlag, New York, 1977. Graduate Texts in Mathematics, No. 52.

[21] Nigel Hitchin. Differentiable manifolds. Lecture notes, Oxford University, 2003.

[22] Klaus Jänich. *Vektoranalysis.* Springer, Berlin, fourth edition, 2003.

[23] Klaus Jänich. *Topologie.* Springer-Lehrbuch. [Springer Textbook]. Springer, Berlin, eighth edition, 2005. With a chapter by Theodor Bröcker.

[24] Nikolaos I. Katzourakis. The riemann surface of the logarithm constructed in a geometrical framework. *Rhit. Math. Journal*, 5, 2005.

[25] Michel A. Kervaire and John W. Milnor. Groups of homotopy spheres. I. *Ann. of Math. (2)*, 77:504–537, 1963.

[26] Felix Klein. *Elementarmathematik vom höheren Standpunkte aus. Zweiter Band: Geometrie.* Dritte Auflage. Ausgearbeitet von E. Hellinger. Für den Druck fertig gemacht und mit Zusätzen versehen von Fr. Seyfarth. Die Grundlehren der mathematischen Wissenschaften, Band 15. Springer-Verlag, Berlin, 1968.

[27] W. Klingenberg. *Lineare Algebra und Geometrie.* Hochschultext. [University Textbooks]. Springer-Verlag, Berlin, 1984.

[28] Ivan Kolár, Peter W. Michor, and Jan Slovák. *Natural Operations in Differential Geometry.* Springer, Berlin, 1993.

[29] Andreas Kriegl and Peter W. Michor. *The convenient setting of global analysis*, volume 53 of *Mathematical Surveys and Monographs*. American Mathematical Society, Providence, RI, 1997.

[30] Serge Lang. *Real and functional analysis*, volume 142 of *Graduate Texts in Mathematics*. Springer, New York, third edition, 1993.

[31] Serge Lang. *Differential and Riemannian manifolds*, volume 160 of *Graduate Texts in Mathematics*. Springer, New York, third edition, 1995.

[32] Serge Lang. *Algebra*, volume 211 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, third edition, 2002.

[33] Jeffrey M. Lee. Manifolds and differential geometry. Lecture notes (draft version of book), Texas Tech University, 2006.

[34] Saunders Mac Lane. *Categories for the working mathematician*, volume 5 of *Graduate Texts in Mathematics*. Springer, New York, second edition, 1998.

[35] Saunders MacLane. Geometrical mechanics I. Lecture notes, University of Chicago, Winter 1968. Call Number 2583 at IBM Linz.

[36] Saunders MacLane. Geometrical mechanics II. Lecture notes, University of Chicago, Spring 1968. Call Number 2584 at IBM Linz.

[37] Saunders MacLane and Garrett Birkhoff. *Algebra*. MacMillan, New York, 1968.

[38] Robert Milson. Tangent space. PlanetMath, 16 February 2002. Version 2.

[39] Robert Milson. Notes on the classical definition of a manifold. PlanetMath, 3 December 2003. Version 7.

[40] James Munkres. Introduction to topology. Online Lecture Notes, 2004. MIT Open Courseware.

[41] Raghavan Narasimhan. *Analysis on Real and Complex Manifolds*, volume 1 of *Advanced Studies in Pure Mathematics*. Masson & Cie Editeur, Paris and North-Holland, Amsterdam, 1968.

[42] Karl-Hermann Neeb. Manifolds and transformation groups. Lecture notes, Technische Universität Darmstadt, 2005.

[43] Ulrich Oberst. Lineare algebra. Lecture notes, $\sim$ 2000.

[44] Peter J. Olver. *Applications of Lie groups to differential equations*, volume 107 of *Graduate Texts in Mathematics*. Springer, New York, second edition, 1993.

[45] Roger Penrose. *The Road to Reality. A Complete Guide to the Laws of the Universe.* Vintage Books, London, 2005.

[46] M. Postnikov. *Smooth manifolds.* Nauka, Moscow, 1989. Lectures in geometry. Semester III, translated from the Russian by Vladimir Shokurov.

[47] J. A. Schouten. *Ricci-calculus. An introduction to tensor analysis and its geometrical applications.* Die Grundlehren der mathematischen Wissenschaften in Einzeldarstellungen mit besonderer Berücksichtigung der Anwendungsgebiete, Bd X. Springer-Verlag, Berlin, second edition, 1954. IBM Call Number 1859.

[48] Igor R. Shafarevich. *Basic algebraic geometry. 1.* Springer, second edition, 1994. Varieties in projective space, Translated from the 1988 Russian edition and with notes by Miles Reid.

[49] Reyer Sjamaar. Manifolds and differential forms. Lecture notes, Cornell University, 2006.

[50] Hans Stephani. *Allgemeine Relativitätstheorie.* Hochschulbücher für Physik [University Books for Physics], 43. VEB Deutscher Verlag der Wissenschaften, Berlin, second edition, 1980. Eine Einführung in die Theorie des Gravitationsfeldes. [An introduction to the theory of the gravitational field].

[51] Walter Thirring. *A course in mathematical physics. Vol. I. Classical dynamical systems.* Springer-Verlag, New York, 1978. Translated from the German by Evans M. Harrell.

[52] Chris Tiee. Contravariance, covariance, densities, and all that: An informal discussion on tensor calculus. Web document, 2006.

[53] Marius van der Put and Michael F. Singer. *Galois theory of linear differential equations*, volume 328 of *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences].* Springer, Berlin, 2003.

[54] Matt Visser. Notes on differential geometry. Technical report, Victoria University of Wellington, New Zealand, 2006.

[55] Frank W. Warner. *Foundations of differentiable manifolds and Lie groups.* Scott, Foresman and Co., Glenview, Ill.-London, 1971.

[56] Gabriel Weinreich. *Geometrical vectors.* Chicago Lectures in Physics. University of Chicago Press, Chicago, IL, 1998.

[57] Hermann Weyl. *Raum, Zeit, Materie. Vorlesungen über allgemeine Relativitätstheorie.* Sechste, unveränderte Auflage. Springer, Berlin, eighth edition, 1993. Main Library (Magazine).

[58] Wikipedia. Fiber bundle — wikipedia, the free encyclopedia, 2006. [Online; accessed 13-December-2006].

[59] Jie Wu. Differentiable manifolds. Technical report, National University of Singapore, 2004.