Chapter 1
Vector Bundles

To motivate the definition of a vector bundle let us consider tangent vectors to the unit 2-sphere S^2 in $\mathbb{R}^3$. At each point $x \in S^2$ there is a tangent plane P_x. This is a 2-dimensional vector space with the point x as its zero vector 0_x. Vectors $v_x \in P_x$ are thought of as arrows with their tail at x. If we regard a vector v_x in P_x as a vector in $\mathbb{R}^3$, then the standard convention in linear algebra would be to identify v_x with all its parallel translates, and in particular with the unique translate $\tau(v_x)$ having its tail at the origin in $\mathbb{R}^3$. The association $v_x \mapsto \tau(v_x)$ defines a function $\tau: TS^2 \to \mathbb{R}^3$ where TS^2 is the set of all tangent vectors v_x as x ranges over S^2. This function τ is surjective but certainly not injective, as every nonzero vector in $\mathbb{R}^3$ occurs as $\tau(v_x)$ for infinitely many x, in fact for all x in a great circle in S^2. Moreover $\tau(0_x) = 0$ for all $x \in S^2$, so $\tau^{-1}(0)$ is a whole sphere. On the other hand, the function $TS^2 \to S^2 \times \mathbb{R}^3$, $v_x \mapsto (x, \tau(v_x))$, is injective, and can be used to topologize TS^2 as a subspace of $S^2 \times \mathbb{R}^3$, namely the subspace consisting of pairs (x, v) with v orthogonal to x.

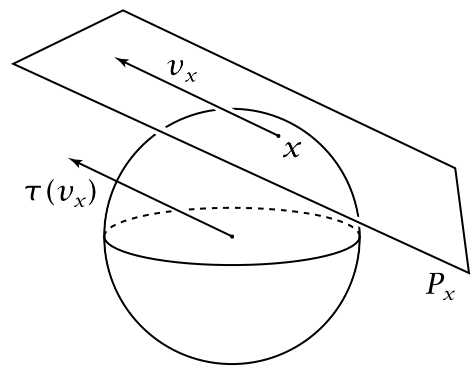

Thus TS^2 is first of all a topological space, and secondly it is the disjoint union of all the vector spaces P_x for $x \in S^2$. One can think of TS^2 as a continuous family of vector spaces parametrized by points of S^2.

The simplest continuous family of 2-dimensional vector spaces parametrized by points of S^2 is of course the product $S^2 \times \mathbb{R}^2$. Is that what TS^2 really is? More precisely we can ask whether there is a homeomorphism $h: TS^2 \to S^2 \times \mathbb{R}^2$ that takes each plane P_x to the plane $\{x\} \times \mathbb{R}^2$ by a vector space isomorphism. If we had such an h, then for each fixed nonzero vector $v \in \mathbb{R}^2$ the family of vectors $v_x = h^{-1}(x, v)$ would be a continuous field of nonzero tangent vectors to S^2. It is a classical theorem in algebraic topology that no such vector field exists. (See §2.2 for a proof using techniques from this book.) So TS^2 is genuinely twisted, and is not just a disguised form of the product $S^2 \times \mathbb{R}^2$.

Dropping down a dimension, one could consider in similar fashion the space TS^1 of tangent vectors to the unit circle S^1 in $\mathbb{R}^2$. In this case there is a continuous field v_x of nonzero tangent vectors to S^1, obtained by regarding points $x \in S^1$ as unit complex numbers and letting v_x be the translation of the vector ix that has its tail at x. This leads to a homeomorphism $S^1 \times \mathbb{R} \to TS^1$ taking (x,t) to tv_x, with $\{x\} \times \mathbb{R}$ going to the tangent line at x by a linear isomorphism. Thus TS^1 really is equivalent to the product $S^1 \times \mathbb{R}^1$.

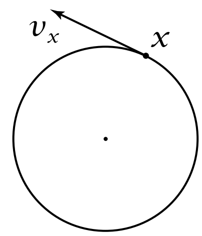

Moving up to S^3, the unit sphere in $\mathbb{R}^4$, the space TS^3 of tangent vectors is again equivalent to the product $S^3 \times \mathbb{R}^3$. Regarding $\mathbb{R}^4$ as the quaternions, an equivalence is the homeomorphism $S^3 \times \mathbb{R}^3 \to TS^3$ sending $(x,(t_1,t_2,t_3))$ to the translation of the vector $t_1 ix + t_2 jx + t_3 kx$ having its tail at x. A similar construction using Cayley octonions shows that TS^7 is equivalent to $S^7 \times \mathbb{R}^7$. It is a rather deep theorem, proved in §2.3, that S^1, S^3, and S^7 are the only spheres whose tangent bundle is equivalent to a product.

Although TS^n is not usually equivalent to the product $S^n \times \mathbb{R}^n$, there is a sense in which this is true locally. Take the case of the 2-sphere for example. For a point $x \in S^2$ let P be the translate of the tangent plane P_x that passes through the origin. For points $y \in S^2$ that are sufficiently close to x the map $\pi_y : P_y \to P$ sending a tangent vector v_y to the orthogonal projection of $\tau(v_y)$ onto P is a linear isomorphism. This is true in fact for any y on the same side of P as x. Thus for y in a suitable neighborhood U of x in S^2 the map $(y, v_y) \mapsto (y, \pi_y(v_y))$ is a homeomorphism with domain the subspace of TS^2 consisting of tangent vectors at points of U and with range the product $U \times P$. Furthermore this homeomorphism has the key property of restricting to a linear isomorphism from P_y onto P for each $y \in U$. A convenient way of rephrasing this situation, having the virtue of easily generalizing, is to let $p : TS^2 \to S^2$ be the map $(x, v_x) \mapsto x$, and then we have a homeomorphism $p^{-1}(U) \to U \times P$ that restricts to a linear isomorphism $p^{-1}(y) \to \{y\} \times P$ for each $y \in U$. With this observation we are now ready to begin with the formal definition of a vector bundle.

1.1 Basic Definitions and Constructions

Throughout the book we use the word "map" to mean a continuous function.

An *n-dimensional vector bundle* is a map $p:E\to B$ together with a real vector space structure on $p^{-1}(b)$ for each $b\in B$, such that the following local triviality condition is satisfied: There is a cover of B by open sets U_α for each of which there exists a homeomorphism $h_\alpha : p^{-1}(U_\alpha)\to U_\alpha\times\mathbb{R}^n$ taking $p^{-1}(b)$ to $\{b\}\times\mathbb{R}^n$ by a vector space isomorphism for each $b\in U_\alpha$. Such an h_α is called a *local trivialization* of the vector bundle. The space B is called the *base space*, E is the *total space*, and the vector spaces $p^{-1}(b)$ are the *fibers*. Often one abbreviates terminology by just calling the vector bundle E, letting the rest of the data be implicit.

We could equally well take $\mathbb{C}$ in place of $\mathbb{R}$ as the scalar field, obtaining the notion of a *complex vector bundle*. We will focus on real vector bundles in this chapter. Usually the complex case is entirely analogous. In the next chapter complex vector bundles will play the larger role, however.

Here are some examples of vector bundles:

(1) The *product* or *trivial* bundle $E=B\times\mathbb{R}^n$ with p the projection onto the first factor.

(2) If we let E be the quotient space of $I\times\mathbb{R}$ under the identifications $(0,t)\sim(1,-t)$, then the projection $I\times\mathbb{R}\to I$ induces a map $p:E\to S^1$ which is a 1-dimensional vector bundle, or *line bundle*. Since E is homeomorphic to a Möbius band with its boundary circle deleted, we call this bundle the *Möbius bundle*.

(3) The tangent bundle of the unit sphere S^n in $\mathbb{R}^{n+1}$, a vector bundle $p:E\to S^n$ where $E=\{\,(x,v)\in S^n\times\mathbb{R}^{n+1}\mid x\perp v\,\}$ and we think of v as a tangent vector to S^n by translating it so that its tail is at the head of x, on S^n. The map $p:E\to S^n$ sends (x,v) to x. To construct local trivializations, choose any point $x\in S^n$ and let $U_x\subset S^n$ be the open hemisphere containing x and bounded by the hyperplane through the origin orthogonal to x. Define $h_x : p^{-1}(U_x)\to U_x\times p^{-1}(x)\approx U_x\times\mathbb{R}^n$ by $h_x(y,v)=(y,\pi_x(v))$ where π_x is orthogonal projection onto the hyperplane $p^{-1}(x)$. Then h_x is a local trivialization since π_x restricts to an isomorphism of $p^{-1}(y)$ onto $p^{-1}(x)$ for each $y\in U_b$.

(4) The normal bundle to S^n in $\mathbb{R}^{n+1}$, a line bundle $p:E\to S^n$ with E consisting of pairs $(x,v)\in S^n\times\mathbb{R}^{n+1}$ such that v is perpendicular to the tangent plane to S^n at x, or in other words, $v=tx$ for some $t\in\mathbb{R}$. The map $p:E\to S^n$ is again given by $p(x,v)=x$. As in the previous example, local trivializations $h_x : p^{-1}(U_x)\to U_x\times\mathbb{R}$ can be obtained by orthogonal projection of the fibers $p^{-1}(y)$ onto $p^{-1}(x)$ for $y\in U_x$.

(5) Real projective n-space $\mathbb{R}\mathrm{P}^n$ is the space of lines in $\mathbb{R}^{n+1}$ through the origin. Since each such line intersects the unit sphere S^n in a pair of antipodal points, we can

also regard $\mathbb{RP}^n$ as the quotient space of S^n in which antipodal pairs of points are identified. The *canonical line bundle* $p:E\to\mathbb{RP}^n$ has as its total space E the subspace of $\mathbb{RP}^n\times\mathbb{R}^{n+1}$ consisting of pairs (ℓ,v) with $v\in\ell$, and $p(\ell,v)=\ell$. Again local trivializations can be defined by orthogonal projection.

There is also an infinite-dimensional projective space $\mathbb{RP}^\infty$ which is the union of the finite-dimensional projective spaces $\mathbb{RP}^n$ under the inclusions $RP^n\subset RP^{n+1}$ coming from the natural inclusions $\mathbb{R}^{n+1}\subset\mathbb{R}^{n+2}$. The topology we use on $\mathbb{RP}^\infty$ is the weak or direct limit topology, for which a set in $\mathbb{RP}^\infty$ is open iff it intersects each $\mathbb{RP}^n$ in an open set. The inclusions $RP^n\subset\mathbb{RP}^{n+1}$ induce corresponding inclusions of canonical line bundles, and the union of all these is a canonical line bundle over $\mathbb{RP}^\infty$, again with the direct limit topology. Local trivializations work just as in the finite-dimensional case.

(6) The canonical line bundle over $\mathbb{RP}^n$ has an orthogonal complement, the space $E^\perp=\{(\ell,v)\in\mathbb{RP}^n\times\mathbb{R}^{n+1}\mid v\perp\ell\}$. The projection $p:E^\perp\to\mathbb{RP}^n$, $p(\ell,v)=\ell$, is a vector bundle with fibers the orthogonal subspaces $\ell^\perp$, of dimension n. Local trivializations can be obtained once more by orthogonal projection.

A natural generalization of $\mathbb{RP}^n$ is the so-called Grassmann manifold $G_k(\mathbb{R}^n)$, the space of all k-dimensional planes through the origin in $\mathbb{R}^n$. The topology on this space will be defined precisely in §1.2, along with a canonical k-dimensional vector bundle over it consisting of pairs (ℓ,v) where ℓ is a point in $G_k(\mathbb{R}^n)$ and v is a vector in ℓ. This too has an orthogonal complement, an $(n-k)$-dimensional vector bundle consisting of pairs (ℓ,v) with v orthogonal to ℓ.

An *isomorphism* between vector bundles $p_1:E_1\to B$ and $p_2:E_2\to B$ over the same base space B is a homeomorphism $h:E_1\to E_2$ taking each fiber $p_1^{-1}(b)$ to the corresponding fiber $p_2^{-1}(b)$ by a linear isomorphism. Thus an isomorphism preserves all the structure of a vector bundle, so isomorphic bundles are often regarded as the same. We use the notation $E_1\approx E_2$ to indicate that E_1 and E_2 are isomorphic.

For example, the normal bundle of S^n in $\mathbb{R}^{n+1}$ is isomorphic to the product bundle $S^n\times\mathbb{R}$ by the map $(x,tx)\mapsto(x,t)$. The tangent bundle to S^1 is also isomorphic to the trivial bundle $S^1\times\mathbb{R}$, via $(e^{i\theta},ite^{i\theta})\mapsto(e^{i\theta},t)$, for $e^{i\theta}\in S^1$ and $t\in\mathbb{R}$.

As a further example, the Möbius bundle in (2) above is isomorphic to the canonical line bundle over $\mathbb{RP}^1\approx S^1$. Namely, $\mathbb{RP}^1$ is swept out by a line rotating through an angle of π, so the vectors in these lines sweep out a rectangle $[0,\pi]\times\mathbb{R}$ with the two ends $\{0\}\times\mathbb{R}$ and $\{\pi\}\times\mathbb{R}$ identified. The identification is $(0,x)\sim(\pi,-x)$ since rotating a vector through an angle of π produces its negative.

Sections

A *section* of a vector bundle $p:E\to B$ is a map $s:B\to E$ assigning to each $b\in B$ a vector $s(b)$ in the fiber $p^{-1}(b)$. The condition $s(b)\in p^{-1}(b)$ can also be written as

$ps = \mathbb{1}$, the identity map of B. Every vector bundle has a canonical section, the *zero section* whose value is the zero vector in each fiber. We often identity the zero section with its image, a subspace of E which projects homeomorphically onto B by p.

One can sometimes distinguish nonisomorphic bundles by looking at the complement of the zero section since any vector bundle isomorphism $h : E_1 \to E_2$ must take the zero section of E_1 onto the zero section of E_2, so the complements of the zero sections in E_1 and E_2 must be homeomorphic. For example, we can see that the Möbius bundle is not isomorphic to the product bundle $S^1 \times \mathbb{R}$ since the complement of the zero section is connected for the Möbius bundle but not for the product bundle.

At the other extreme from the zero section would be a section whose values are all nonzero. Not all vector bundles have such a section. Consider for example the tangent bundle to S^n. Here a section is just a tangent vector field to S^n. As we shall show in §2.2, S^n has a nonvanishing vector field iff n is odd. From this it follows that the tangent bundle of S^n is not isomorphic to the trivial bundle if n is even and nonzero, since the trivial bundle obviously has a nonvanishing section, and an isomorphism between vector bundles takes nonvanishing sections to nonvanishing sections.

In fact, an n-dimensional bundle $p : E \to B$ is isomorphic to the trivial bundle iff it has n sections $s_1, \cdots, s_n$ such that the vectors $s_1(b), \cdots, s_n(b)$ are linearly independent in each fiber $p^{-1}(b)$. In one direction this is evident since the trivial bundle certainly has such sections and an isomorphism of vector bundles takes linearly independent sections to linearly independent sections. Conversely, if one has n linearly independent sections s_i, the map $h : B \times \mathbb{R}^n \to E$ given by $h(b, t_1, \cdots, t_n) = \sum_i t_i s_i(b)$ is a linear isomorphism in each fiber, and is continuous since its composition with a local trivialization $p^{-1}(U) \to U \times \mathbb{R}^n$ is continuous. Hence h is an isomorphism by the following useful technical result:

Lemma 1.1. *A continuous map $h : E_1 \to E_2$ between vector bundles over the same base space B is an isomorphism if it takes each fiber $p_1^{-1}(b)$ to the corresponding fiber $p_2^{-1}(b)$ by a linear isomorphism.*

Proof: The hypothesis implies that h is one-to-one and onto. What must be checked is that h^{-1} is continuous. This is a local question, so we may restrict to an open set $U \subset B$ over which E_1 and E_2 are trivial. Composing with local trivializations reduces to the case of an isomorphism $h : U \times \mathbb{R}^n \to U \times \mathbb{R}^n$ of the form $h(x, v) = (x, g_x(v))$. Here g_x is an element of the group $GL_n(\mathbb{R})$ of invertible linear transformations of $\mathbb{R}^n$, and g_x depends continuously on x. This means that if g_x is regarded as an $n \times n$ matrix, its n^2 entries depend continuously on x. The inverse matrix g_x^{-1} also depends continuously on x since its entries can be expressed algebraically in terms of the entries of g_x, namely, g_x^{-1} is $1/(\det g_x)$ times the classical adjoint matrix of g_x. Therefore $h^{-1}(x, v) = (x, g_x^{-1}(v))$ is continuous. $\square$

As an example, the tangent bundle to S^1 is trivial because it has the section $(x_1, x_2) \mapsto (-x_2, x_1)$ for $(x_1, x_2) \in S^1$. In terms of complex numbers, if we set $z = x_1 + ix_2$ then this section is $z \mapsto iz$ since $iz = -x_2 + ix_1$.

There is an analogous construction using quaternions instead of complex numbers. Quaternions have the form $z = x_1 + ix_2 + jx_3 + kx_4$, and form a division algebra $\mathbb{H}$ via the multiplication rules $i^2 = j^2 = k^2 = -1$, $ij = k$, $jk = i$, $ki = j$, $ji = -k$, $kj = -i$, and $ik = -j$. If we identify $\mathbb{H}$ with $\mathbb{R}^4$ via the coordinates (x_1, x_2, x_3, x_4), then the unit sphere is S^3 and we can define three sections of its tangent bundle by the formulas

$$\begin{aligned} z \mapsto iz \quad &\text{or} \quad (x_1, x_2, x_3, x_4) \mapsto (-x_2, x_1, -x_4, x_3) \\ z \mapsto jz \quad &\text{or} \quad (x_1, x_2, x_3, x_4) \mapsto (-x_3, x_4, x_1, -x_2) \\ z \mapsto kz \quad &\text{or} \quad (x_1, x_2, x_3, x_4) \mapsto (-x_4, -x_3, x_2, x_1) \end{aligned}$$

It is easy to check that the three vectors in the last column are orthogonal to each other and to (x_1, x_2, x_3, x_4), so we have three linearly independent nonvanishing tangent vector fields on S^3, and hence the tangent bundle to S^3 is trivial.

The underlying reason why this works is that quaternion multiplication satisfies $|zw| = |z||w|$, where $|\cdot|$ is the usual norm of vectors in $\mathbb{R}^4$. Thus multiplication by a quaternion in the unit sphere S^3 is an isometry of $\mathbb{H}$. The quaternions $1, i, j, k$ form the standard orthonormal basis for $\mathbb{R}^4$, so when we multiply them by an arbitrary unit quaternion $z \in S^3$ we get a new orthonormal basis z, iz, jz, kz.

The same constructions work for the Cayley octonions, a division algebra structure on $\mathbb{R}^8$. Thinking of $\mathbb{R}^8$ as $\mathbb{H} \times \mathbb{H}$, multiplication of octonions is defined by $(z_1, z_2)(w_1, w_2) = (z_1 w_1 - \overline{w}_2 z_2, z_2 \overline{w}_1 + w_2 z_1)$ and satisfies the key property $|zw| = |z||w|$. This leads to the construction of seven orthogonal tangent vector fields on the unit sphere S^7, so the tangent bundle to S^7 is also trivial. As we shall show in §2.3, the only spheres with trivial tangent bundle are S^1, S^3, and S^7.

Another way of characterizing the trivial bundle $E \approx B \times \mathbb{R}^n$ is to say that there is a continuous projection map $E \to \mathbb{R}^n$ which is a linear isomorphism on each fiber, since such a projection together with the bundle projection $E \to B$ gives an isomorphism $E \approx B \times \mathbb{R}^n$, by Lemma 1.1.

Direct Sums

Given two vector bundles $p_1 : E_1 \to B$ and $p_2 : E_2 \to B$ over the same base space B, we would like to create a third vector bundle over B whose fiber over each point of B is the direct sum of the fibers of E_1 and E_2 over this point. This leads us to define the *direct sum* of E_1 and E_2 as the space

$$E_1 \oplus E_2 = \{ (v_1, v_2) \in E_1 \times E_2 \mid p_1(v_1) = p_2(v_2) \}$$

There is then a projection $E_1 \oplus E_2 \to B$ sending (v_1, v_2) to the point $p_1(v_1) = p_2(v_2)$. The fibers of this projection are the direct sums of the fibers of E_1 and E_2, as we wanted. For a relatively painless verification of the local triviality condition we make two preliminary observations:

(a) Given a vector bundle $p:E\to B$ and a subspace $A\subset B$, then $p:p^{-1}(A)\to A$ is clearly a vector bundle. We call this the *restriction of E over* A.

(b) Given vector bundles $p_1:E_1\to B_1$ and $p_2:E_2\to B_2$, then $p_1\times p_2:E_1\times E_2\to B_1\times B_2$ is also a vector bundle, with fibers the products $p_1^{-1}(b_1)\times p_2^{-1}(b_2)$. For if we have local trivializations $h_\alpha:p_1^{-1}(U_\alpha)\to U_\alpha\times\mathbb{R}^n$ and $h_\beta:p_2^{-1}(U_\beta)\to U_\beta\times\mathbb{R}^m$ for E_1 and E_2, then $h_\alpha\times h_\beta$ is a local trivialization for $E_1\times E_2$.

Then if E_1 and E_2 both have the same base space B, the restriction of the product $E_1\times E_2$ over the diagonal $B=\{(b,b)\in B\times B\}$ is exactly $E_1\oplus E_2$.

The direct sum of two trivial bundles is again a trivial bundle, clearly, but the direct sum of nontrivial bundles can also be trivial. For example, the direct sum of the tangent and normal bundles to S^n in $\mathbb{R}^{n+1}$ is the trivial bundle $S^n\times\mathbb{R}^{n+1}$ since elements of the direct sum are triples $(x,v,tx)\in S^n\times\mathbb{R}^{n+1}\times\mathbb{R}^{n+1}$ with $x\perp v$, and the map $(x,v,tx)\mapsto(x,v+tx)$ gives an isomorphism of the direct sum bundle with $S^n\times\mathbb{R}^{n+1}$. So the tangent bundle to S^n is *stably trivial*: it becomes trivial after taking the direct sum with a trivial bundle.

As another example, the direct sum $E\oplus E^\perp$ of the canonical line bundle $E\to\mathbb{R}P^n$ with its orthogonal complement, defined in example (6) above, is isomorphic to the trivial bundle $\mathbb{R}P^n\times\mathbb{R}^{n+1}$ via the map $(\ell,v,w)\mapsto(\ell,v+w)$ for $v\in\ell$ and $w\perp\ell$. Specializing to the case $n=1$, the bundle $E^\perp$ is isomorphic to E itself by the map that rotates each vector in the plane by 90 degrees. We noted earlier that E is isomorphic to the Möbius bundle over $S^1=\mathbb{R}P^1$, so it follows that the direct sum of the Möbius bundle with itself is the trivial bundle. To see this geometrically, embed the Möbius bundle in the product bundle $S^1\times\mathbb{R}^2$ by taking the line in the fiber $\{\theta\}\times\mathbb{R}^2$ that makes an angle of $\theta/2$ with the x-axis, and then the orthogonal lines in the fibers form a second copy of the Möbius bundle, giving a decomposition of the product $S^1\times\mathbb{R}^2$ as the direct sum of two Möbius bundles.

Example: The tangent bundle of real projective space. Starting with the isomorphism $S^n\times\mathbb{R}^{n+1}\approx TS^n\oplus NS^n$, where NS^n is the normal bundle of S^n in $\mathbb{R}^{n+1}$, suppose we factor out by the identifications $(x,v)\sim(-x,-v)$ on both sides of this isomorphism. Applied to TS^n this identification yields $T\mathbb{R}P^n$, the tangent bundle to $\mathbb{R}P^n$. This is saying that a tangent vector to $\mathbb{R}P^n$ is equivalent to a pair of antipodal tangent vectors to S^n. A moment's reflection shows this to be entirely reasonable, although a formal proof would require a significant digression on what precisely tangent vectors to a smooth manifold are, a digression we shall skip here. What we will show is that even though the direct sum of $T\mathbb{R}P^n$ with a trivial line bundle may not be trivial as it is for a sphere, it does split in an interesting way as a direct sum of nontrivial line bundles.

In the normal bundle NS^n the identification $(x,v) \sim (-x,-v)$ can be written as $(x,tx) \sim (-x,t(-x))$. This identification yields the product bundle $\mathbb{R}P^n \times \mathbb{R}$ since the section $x \mapsto (-x,-x)$ is well-defined in the quotient. Now let us consider the identification $(x,v) \sim (-x,-v)$ in $S^n \times \mathbb{R}^{n+1}$. This identification respects the coordinate factors of $\mathbb{R}^{n+1}$, so the quotient is the direct sum of $n+1$ copies of the line bundle E over $\mathbb{R}P^n$ obtained by making the identifications $(x,t) \sim (-x,-t)$ in $S^n \times \mathbb{R}$. The claim is that E is just the canonical line bundle over $\mathbb{R}P^n$. To see this, let us identify $S^n \times \mathbb{R}$ with NS^n by the isomorphism $(x,t) \mapsto (x,tx)$. hence $(-x,-t) \mapsto ((-x,(-t)(-x)) = (-x,tx)$. Thus we have the identification $(x,tx) \sim (-x,tx)$ in NS^n. The quotient is the canonical line bundle over $\mathbb{R}P^n$ since the identifications $x \sim -x$ in the first coordinate give lines through the origin in $\mathbb{R}^{n+1}$, and in the second coordinate there are no identifications so we have well-defined vectors tx in these lines.

Thus we have shown that the tangent bundle $T\mathbb{R}P^n$ is stably isomorphic to the direct sum of $n+1$ copies of the canonical line bundle over $\mathbb{R}P^n$. When $n = 3$, for example, $T\mathbb{R}P^3$ is trivial since the three linearly independent tangent vector fields on S^3 defined earlier in terms of quaternions pass down to linearly independent tangent vector fields on the quotient $\mathbb{R}P^3$. Hence the direct sum of four copies of the canonical line bundle over $\mathbb{R}P^3$ is trivial. Similarly using octonions we can see that the direct sum of eight copies of the canonical line bundle over $\mathbb{R}P^7$ is trivial. In §2.5 we will determine when the sum of k copies of the canonical line bundle over $\mathbb{R}P^n$ is at least stably trivial. The answer turns out to be rather subtle: This happens exactly when k is a multiple of $2^{\varphi(n)}$ where $\varphi(n)$ is the number of integers i in the range $0 < i \le n$ with i congruent to 0, 1, 2, or 4 modulo 8. For $n = 3,7$ this gives $2^{\varphi(n)} = 4,8$, the same numbers we obtained from the triviality of $T\mathbb{R}P^3$ and $T\mathbb{R}P^7$. If there were a 16-dimensional division algebra after the octonions then one might expect the sum of 16 copies of the canonical line bundle over $\mathbb{R}P^{15}$ to be trivial. However this is not the case since $2^{\varphi(15)} = 2^7 = 128$.

Inner Products

An *inner product* on a vector bundle $p: E \to B$ is a map $\langle\,,\rangle : E \oplus E \to \mathbb{R}$ which restricts in each fiber to an inner product, a positive definite symmetric bilinear form.

Proposition 1.2. *An inner product exists for a vector bundle $p: E \to B$ if B is compact Hausdorff or more generally paracompact.*

The definition of paracompactness we are using is that a space X is paracompact if it is Hausdorff and every open cover has a subordinate partition of unity, a collection of maps $\varphi_\beta : X \to [0,1]$ each supported in some set of the open cover, and with $\sum_\beta \varphi_\beta = 1$, only finitely many of the φ_β being nonzero near each point of X. Constructing such functions is easy when X is compact Hausdorff, using Urysohn's

Lemma. This is done in the appendix to this chapter, where we also show that certain classes of noncompact spaces are paracompact. Most spaces that arise naturally in algebraic topology are paracompact.

Proof: An inner product for $p:E \to B$ can be constructed by first using local trivializations $h_\alpha : p^{-1}(U_\alpha) \to U_\alpha \times \mathbb{R}^n$, to pull back the standard inner product in $\mathbb{R}^n$ to an inner product $\langle \cdot, \cdot \rangle_\alpha$ on $p^{-1}(U_\alpha)$, then setting $\langle v, w \rangle = \sum_\beta \varphi_\beta p(v) \langle v, w \rangle_{\alpha(\beta)}$ where $\{\varphi_\beta\}$ is a partition of unity with the support of φ_β contained in $U_{\alpha(\beta)}$. □

In the case of complex vector bundles one can construct Hermitian inner products in the same way.

In linear algebra one can show that a vector subspace is always a direct summand by taking its orthogonal complement. We will show now that the corresponding result holds for vector bundles over a paracompact base. A *vector subbundle* of a vector bundle $p:E \to B$ has the natural definition: a subspace $E_0 \subset E$ intersecting each fiber of E in a vector subspace, such that the restriction $p:E_0 \to B$ is a vector bundle.

Proposition 1.3. *If $E \to B$ is a vector bundle over a paracompact base B and $E_0 \subset E$ is a vector subbundle, then there is a vector subbundle $E_0^\perp \subset E$ such that $E_0 \oplus E_0^\perp \approx E$.*

Proof: With respect to a chosen inner product on E, let $E_0^\perp$ be the subspace of E which in each fiber consists of all vectors orthogonal to vectors in E_0. We claim that the natural projection $E_0^\perp \to B$ is a vector bundle. If this is so, then $E_0 \oplus E_0^\perp$ is isomorphic to E via the map $(v, w) \mapsto v + w$, using Lemma 1.1.

To see that $E_0^\perp$ satisfies the local triviality condition for a vector bundle, note first that we may assume E is the product $B \times \mathbb{R}^n$ since the question is local in B. Since E_0 is a vector bundle, of dimension m say, it has m independent local sections $b \mapsto (b, s_i(b))$ near each point $b_0 \in B$. We may enlarge this set of m independent local sections of E_0 to a set of n independent local sections $b \mapsto (b, s_i(b))$ of E by choosing $s_{m+1}, \cdots, s_n$ first in the fiber $p^{-1}(b_0)$, then taking the same vectors for all nearby fibers, since if $s_1, \cdots, s_m, s_{m+1}, \cdots, s_n$ are independent at b_0, they will remain independent for nearby b by continuity of the determinant function. Apply the Gram-Schmidt orthogonalization process to $s_1, \cdots, s_m, s_{m+1}, \cdots, s_n$ in each fiber, using the given inner product, to obtain new sections s_i'. The explicit formulas for the Gram-Schmidt process show the s_i''s are continuous, and the first m of them are a basis for E_0 in each fiber. The sections s_i' allow us to define a local trivialization $h : p^{-1}(U) \to U \times \mathbb{R}^n$ with $h(b, s_i'(b))$ equal to the i^{th} standard basis vector of $\mathbb{R}^n$. This h carries E_0 to $U \times \mathbb{R}^m$ and $E_0^\perp$ to $U \times \mathbb{R}^{n-m}$, so $h \,|\, E_0^\perp$ is a local trivialization of $E_0^\perp$. □

Note that when the subbundle E_0 is equal to E itself, the last part of the proof shows that for any vector bundle with an inner product it is always possible to choose local trivializations that carry the inner product to the standard inner product, so the local trivializations are by isometries.

We have seen several cases where the sum of two bundles, one or both of which may be nontrivial, is the trivial bundle. Here is a general result result along these lines:

Proposition 1.4. *For each vector bundle $E \to B$ with B compact Hausdorff there exists a vector bundle $E' \to B$ such that $E \oplus E'$ is the trivial bundle.*

This can fail when B is noncompact. An example is the canonical line bundle over $\mathbb{R}P^\infty$, as we shall see in Example 3.6.

Proof: To motivate the construction, suppose first that the result holds and hence that E is a subbundle of a trivial bundle $B \times \mathbb{R}^N$. Composing the inclusion of E into this product with the projection of the product onto $\mathbb{R}^N$ yields a map $E \to \mathbb{R}^N$ that is a linear injection on each fiber. Our strategy will be to reverse the logic here, first constructing a map $E \to \mathbb{R}^N$ that is a linear injection on each fiber, then showing that this gives an embedding of E in $B \times \mathbb{R}^N$ as a direct summand.

Each point $x \in B$ has a neighborhood U_x over which E is trivial. By Urysohn's Lemma there is a map $\varphi_x : B \to [0,1]$ that is 0 outside U_x and nonzero at x. Letting x vary, the sets $\varphi_x^{-1}(0,1]$ form an open cover of B. By compactness this has a finite subcover. Let the corresponding U_x's and φ_x's be relabeled U_i and φ_i. Define $g_i : E \to \mathbb{R}^n$ by $g_i(v) = \varphi_i(p(v))[\pi_i h_i(v)]$ where p is the projection $E \to B$ and $\pi_i h_i$ is the composition of a local trivialization $h_i : p^{-1}(U_i) \to U_i \times \mathbb{R}^n$ with the projection π_i to $\mathbb{R}^n$. Then g_i is a linear injection on each fiber over $\varphi_i^{-1}(0,1]$, so if we make the various g_i's the coordinates of a map $g : E \to \mathbb{R}^N$ with $\mathbb{R}^N$ a product of copies of $\mathbb{R}^n$, then g is a linear injection on each fiber.

The map g is the second coordinate of a map $f : E \to B \times \mathbb{R}^N$ with first coordinate p. The image of f is a subbundle of the product $B \times \mathbb{R}^N$ since projection of $\mathbb{R}^N$ onto the i^{th} $\mathbb{R}^n$ factor gives the second coordinate of a local trivialization over $\varphi_i^{-1}(0,1]$. Thus we have E isomorphic to a subbundle of $B \times \mathbb{R}^N$ so by preceding proposition there is a complementary subbundle E' with $E \oplus E'$ isomorphic to $B \times \mathbb{R}^N$. $\square$

Tensor Products

In addition to direct sum, a number of other algebraic constructions with vector spaces can be extended to vector bundles. One which is particularly important for K-theory is tensor product. For vector bundles $p_1 : E_1 \to B$ and $p_2 : E_2 \to B$, let $E_1 \otimes E_2$, as a set, be the disjoint union of the vector spaces $p_1^{-1}(x) \otimes p_2^{-1}(x)$ for $x \in B$. The topology on this set is defined in the following way. Choose isomorphisms $h_i : p_i^{-1}(U) \to U \times \mathbb{R}^{n_i}$ for each open set $U \subset B$ over which E_1 and E_2 are trivial. Then a topology $\mathcal{T}_U$ on the set $p_1^{-1}(U) \otimes p_2^{-1}(U)$ is defined by letting the fiberwise tensor product map $h_1 \otimes h_2 : p_1^{-1}(U) \otimes p_2^{-1}(U) \to U \times (\mathbb{R}^{n_1} \otimes \mathbb{R}^{n_2})$ be a homeomorphism. The topology $\mathcal{T}_U$ is independent of the choice of the h_i's since any other choices are obtained by composing with isomorphisms of $U \times \mathbb{R}^{n_i}$ of the form $(x, v) \mapsto (x, g_i(x)(v))$ for continuous maps $g_i : U \to GL_{n_i}(\mathbb{R})$, hence $h_1 \otimes h_2$ changes by composing with analogous isomorphisms of $U \times (\mathbb{R}^{n_1} \otimes \mathbb{R}^{n_2})$ whose second coordinates $g_1 \otimes g_2$ are continuous maps $U \to GL_{n_1 n_2}(\mathbb{R})$, since the entries of the matrices $g_1(x) \otimes g_2(x)$ are the products of the entries of $g_1(x)$ and $g_2(x)$. When we replace U by an open subset V, the topology on $p_1^{-1}(V) \otimes p_2^{-1}(V)$ induced by $\mathcal{T}_U$ is the same as the topology $\mathcal{T}_V$ since local trivializations over U restrict to local trivializations over V. Hence we get a well-defined topology on $E_1 \otimes E_2$ making it a vector bundle over B.

There is another way to look at this construction that takes as its point of departure a general method for constructing vector bundles we have not mentioned previously. If we are given a vector bundle $p : E \to B$ and an open cover $\{U_\alpha\}$ of B with local trivializations $h_\alpha : p^{-1}(U_\alpha) \to U_\alpha \times \mathbb{R}^n$, then we can reconstruct E as the quotient space of the disjoint union $\coprod_\alpha (U_\alpha \times \mathbb{R}^n)$ obtained by identifying $(x, v) \in U_\alpha \times \mathbb{R}^n$ with $h_\beta h_\alpha^{-1}(x, v) \in U_\beta \times \mathbb{R}^n$ whenever $x \in U_\alpha \cap U_\beta$. The functions $h_\beta h_\alpha^{-1}$ can be viewed as maps $g_{\beta\alpha} : U_\alpha \cap U_\beta \to GL_n(\mathbb{R})$. These satisfy the 'cocycle condition' $g_{\gamma\beta} g_{\beta\alpha} = g_{\gamma\alpha}$ on $U_\alpha \cap U_\beta \cap U_\gamma$. Any collection of 'gluing functions' $g_{\beta\alpha}$ satisfying this condition can be used to construct a vector bundle $E \to B$.

In the case of tensor products, suppose we have two vector bundles $E_1 \to B$ and $E_2 \to B$. We can choose an open cover $\{U_\alpha\}$ with both E_1 and E_2 trivial over each U_α, and so obtain gluing functions $g^i_{\beta\alpha} : U_\alpha \cap U_\beta \to GL_{n_i}(\mathbb{R})$ for each E_i. Then the gluing functions for the bundle $E_1 \otimes E_2$ are the tensor product functions $g^1_{\beta\alpha} \otimes g^2_{\beta\alpha}$ assigning to each $x \in U_\alpha \cap U_\beta$ the tensor product of the two matrices $g^1_{\beta\alpha}(x)$ and $g^2_{\beta\alpha}(x)$.

It is routine to verify that the tensor product operation for vector bundles over a fixed base space is commutative, associative, and has an identity element, the trivial line bundle. It is also distributive with respect to direct sum.

If we restrict attention to line bundles, then the set $\mathrm{Vect}^1(B)$ of isomorphism classes of one-dimensional vector bundles over B is an abelian group with respect to the tensor product operation. The inverse of a line bundle $E \to B$ is obtained by replacing its gluing matrices $g_{\beta\alpha}(x) \in GL_1(\mathbb{R})$ with their inverses. The cocycle condition is preserved since 1×1 matrices commute. If we give E an inner product, we may rescale local trivializations h_α to be isometries, taking vectors in fibers of E to vectors in $\mathbb{R}^1$ of the same length. Then all the values of the gluing functions $g_{\beta\alpha}$ are ± 1, being isometries of $\mathbb{R}$. The gluing functions for $E \otimes E$ are the squares of these $g_{\beta\alpha}$'s, hence are identically 1, so $E \otimes E$ is the trivial line bundle. Thus each element of $\mathrm{Vect}^1(B)$ is its own inverse. As we shall see in §3.1, the group $\mathrm{Vect}^1(B)$ is isomorphic to a standard object in algebraic topology, the cohomology group $H^1(B; \mathbb{Z}_2)$ when B is homotopy equivalent to a CW complex.

These tensor product constructions work equally well for complex vector bundles. Tensor product again makes the complex analog $\mathrm{Vect}^1_{\mathbb{C}}(B)$ of $\mathrm{Vect}^1(B)$ into an abelian group, but after rescaling the gluing functions $g_{\beta\alpha}$ for a complex line bundle E, the values are complex numbers of norm 1, not necessarily ± 1, so we cannot expect $E\otimes E$ to be trivial. In §3.1 we will show that the group $\mathrm{Vect}^1_{\mathbb{C}}(B)$ is isomorphic to $H^2(B;\mathbb{Z})$ when B is homotopy equivalent to a CW complex.

We may as well mention here another general construction for complex vector bundles $E\to B$, the notion of the *conjugate bundle* $\overline{E}\to B$. As a topological space, $\overline{E}$ is the same as E, but the vector space structure in the fibers is modified by redefining scalar multiplication by the rule $\lambda(v) = \overline{\lambda}v$ where the right side of this equation means scalar multiplication in E and the left side means scalar multiplication in $\overline{E}$. This implies that local trivializations for $\overline{E}$ are obtained from local trivializations for E by composing with the coordinatewise conjugation map $\mathbb{C}^n\to\mathbb{C}^n$ in each fiber. The effect on the gluing maps $g_{\beta\alpha}$ is to replace them by their complex conjugates as well. Specializing to line bundles, we then have $E\otimes\overline{E}$ isomorphic to the trivial line bundle since its gluing maps have values $z\overline{z} = 1$ for z a unit complex number. Thus conjugate bundles provide inverses in $\mathrm{Vect}^1_{\mathbb{C}}(B)$.

Besides tensor product of vector bundles, another construction useful in K–theory is the exterior power $\lambda^k(E)$ of a vector bundle E. Recall from linear algebra that the exterior power $\lambda^k(V)$ of a vector space V is the quotient of the k-fold tensor product $V\otimes\cdots\otimes V$ by the subspace generated by vectors of the form $v_1\otimes\cdots\otimes v_k - \mathrm{sgn}(\sigma)v_{\sigma(1)}\otimes\cdots\otimes v_{\sigma(k)}$ where σ is a permutation of the subscripts and $\mathrm{sgn}(\sigma) = \pm 1$ is its sign, $+1$ for an even permutation and -1 for an odd permutation. If V has dimension n then $\lambda^k(V)$ has dimension $\binom{n}{k}$. Now to define $\lambda^k(E)$ for a vector bundle $p: E\to B$ the procedure follows closely what we did for tensor product. We first form the disjoint union of the exterior powers $\lambda^k(p^{-1}(x))$ of all the fibers $p^{-1}(x)$, then we define a topology on this set via local trivializations. The key fact about tensor product which we needed before was that the tensor product $\varphi\otimes\psi$ of linear transformations φ and ψ depends continuously on φ and ψ. For exterior powers the analogous fact is that a linear map $\varphi:\mathbb{R}^n\to\mathbb{R}^n$ induces a linear map $\lambda^k(\varphi):\lambda^k(\mathbb{R}^n)\to\lambda^k(\mathbb{R}^n)$ which depends continuously on φ. This holds since $\lambda^k(\varphi)$ is a quotient map of the k-fold tensor product of φ with itself.

Associated Fiber Bundles

If we modify the definition of a vector bundle by dropping all references to vector spaces and replace the model fiber $\mathbb{R}^n$ by an arbitrary space F, then we have the more general notion of a *fiber bundle* with fiber F. This is a map $p: E\to B$ such that there is a cover of B by open sets U_α for each of which there exists a homeomorphism $h_\alpha: p^{-1}(U_\alpha)\to U_\alpha\times F$ taking $p^{-1}(b)$ to $\{b\}\times F$ for each $b\in U_\alpha$. We will describe now several different ways of constructing fiber bundles from vector bundles.

Having an inner product on a vector bundle E, lengths of vectors are defined, and we can consider the subspace $S(E)$ consisting of the unit spheres in all the fibers. The natural projection $S(E) \to B$ is a fiber bundle with sphere fibers since we have observed that local trivializations for E can be chosen to be isometries in each fiber, so these local trivializations restrict to local trivializations for $S(E)$. Similarly there is a disk bundle $D(E) \to B$ with fibers the disks of vectors of length less than or equal to 1. It is possible to describe $S(E)$ without reference to an inner product, as the quotient of the complement of the zero section in E obtained by identifying each nonzero vector with all positive scalar multiples of itself. It follows that $D(E)$ can also be defined without invoking a metric, namely as the mapping cylinder of the projection $S(E) \to B$. This is the quotient space of $S(E) \times [0,1]$ obtained by identifying two points in $S(E) \times \{0\}$ if they have the same image in B.

The canonical line bundle $E \to \mathbb{R}P^n$ has as its unit sphere bundle $S(E)$ the space of unit vectors in lines through the origin in $\mathbb{R}^{n+1}$. Since each unit vector uniquely determines the line containing it, $S(E)$ is the same as the space of unit vectors in $\mathbb{R}^{n+1}$, i.e., S^n.

Here are some more examples.

(1) Associated to a vector bundle $E \to B$ is the *projective bundle* $P(E) \to B$, where $P(E)$ is the space of all lines through the origin in all the fibers of E. We topologize $P(E)$ as the quotient of the sphere bundle $S(E)$ obtained by factoring out scalar multiplication in each fiber. Over a neighborhood U in B where E is a product $U \times \mathbb{R}^n$, this quotient is $U \times \mathbb{R}P^{n-1}$, so $P(E)$ is a fiber bundle over B with fiber $\mathbb{R}P^{n-1}$, with respect to the projection $P(E) \to B$ which sends each line in the fiber of E over a point $b \in B$ to b.

(2) For an n-dimensional vector bundle $E \to B$, the associated *flag bundle* $F(E) \to B$ has total space $F(E)$ the subspace of the n-fold product of $P(E)$ with itself consisting of n-tuples of orthogonal lines in fibers of E. The fiber of $F(E)$ is thus the flag manifold $F(\mathbb{R}^n)$ consisting of n-tuples of orthogonal lines through the origin in $\mathbb{R}^n$. Local triviality follows as in the preceding example. More generally, for any $k \le n$ one could take k-tuples of orthogonal lines in fibers of E and get a bundle $F_k(E) \to B$.

(3) As a refinement of the last example, one could form the *Stiefel bundle* $V_k(E) \to B$, where points of $V_k(E)$ are k-tuples of orthogonal unit vectors in fibers of E, so $V_k(E)$ is a subspace of the product of k copies of $S(E)$. The fiber of $V_k(E)$ is the Stiefel manifold $V_k(\mathbb{R}^n)$ of orthonormal k-frames in $\mathbb{R}^n$.

(4) Generalizing $P(E)$, there is the *Grassmann bundle* $G_k(E) \to B$ of k-dimensional linear subspaces of fibers of E. This is the quotient space of $V_k(E)$ obtained by identifying two k-frames in a fiber if they span the same subspace of the fiber. The fiber of $G_k(E)$ is the Grassmann manifold $G_k(\mathbb{R}^n)$ of k-planes through the origin in $\mathbb{R}^n$.

Exercises

1. Show that a vector bundle $E \to X$ has k independent sections iff it has a trivial k-dimensional subbundle.

2. For a vector bundle $E \to X$ with a subbundle $E' \subset E$, construct a quotient vector bundle $E/E' \to X$.

3. Show that the orthogonal complement of a subbundle is independent of the choice of inner product, up to isomorphism.

1.2 Classifying Vector Bundles

As was stated in the Introduction, it is usually a difficult problem to classify all the different vector bundles over a given base space. The goal of this section will be to rephrase the problem in terms of a standard concept of algebraic topology, the idea of homotopy classes of maps. This will allow the problem to be solved directly in a few very simple cases. Using machinery of algebraic topology, other more difficult cases can be handled as well, as is explained in §??. The reformulation in terms of homotopy also offers some conceptual enlightenment about the structure of vector bundles.

For the reader who is unfamiliar with the notion of homotopy we give the basic definitions in the Glossary [not yet written], and more details can be found in the introductory chapter of the author's book *Algebraic Topology*.

Pullback Bundles

We will denote the set of isomorphism classes of n-dimensional real vector bundles over B by $\mathrm{Vect}^n(B)$. The complex analogue, when we need it, will be denoted by $\mathrm{Vect}^n_{\mathbb{C}}(B)$. Our first task is to show how a map $f:A\to B$ gives rise to a function $f^*:\mathrm{Vect}(B)\to\mathrm{Vect}(A)$, in the reverse direction.

Proposition 1.5. *Given a map $f:A\to B$ and a vector bundle $p:E\to B$, then there exists a vector bundle $p':E'\to A$ with a map $f':E'\to E$ taking the fiber of E' over each point $a\in A$ isomorphically onto the fiber of E over $f(a)$, and such a vector bundle E' is unique up to isomorphism.*

From the uniqueness statement it follows that the isomorphism type of E' depends only on the isomorphism type of E since we can compose the map f' with an isomorphism of E with another vector bundle over B. Thus we have a function $f^*:\mathrm{Vect}(B)\to\mathrm{Vect}(A)$ taking the isomorphism class of E to the isomorphism class of E'. Often the vector bundle E' is written as $f^*(E)$ and called the bundle *induced* by f, or the *pullback* of E by f.

Proof: First we construct an explicit pullback by setting $E' = \{\,(a,v)\in A\times E \mid f(a) = p(v)\,\}$. This subspace of $A\times E$ fits into the diagram at the right where $p'(a,v) = a$ and $f'(a,v) = v$. Notice that the two compositions fp' and pf' from E' to B are equal since they both since a pair (a,v) to $f(a)$. The formula $fp' = pf'$ looks a bit like a commutativity relation, which may explain why the word *commutative* is used to describe a diagram like this one in which any two compositions of maps from one point in the diagram to another are equal.

$$\begin{tikzcd} E' \arrow[r, "\widetilde{f}'"] \arrow[d, "p'"] & E \arrow[d, "p"] \\ A \arrow[r, "f"] & B \end{tikzcd}$$

If we let Γ_f denote the graph of f, all points $(a,f(a))$ in $A\times B$, then p' factors as the composition $E'\to\Gamma_f\to A$, $(a,v)\mapsto(a,p(v)) = (a,f(a))\mapsto a$. The first of these two maps is the restriction of the vector bundle $1\!\!1\times p : A\times E\to A\times B$ over the graph Γ_f, so it is a vector bundle, and the second map is a homeomorphism, so their composition $p' : E' \to A$ is a vector bundle. The map f' obviously takes the fiber E' over a isomorphically onto the fiber of E over $f(a)$.

For the uniqueness statement, we can construct an isomorphism from an arbitrary E' satisfying the conditions in the proposition to the particular one just constructed by sending $v' \in E'$ to the pair $(p'(v'), f'(v'))$. This map obviously takes each fiber of E' to the corresponding fiber of $f^*(E)$ by a vector space isomorphism, so by Lemma 1.1 it is an isomorphism of vector bundles. □

One can be more explicit about local trivializations in the pullback bundle $f^*(E)$ constructed above. If E is trivial over a subspace $U \subset B$ then $f^*(E)$ is trivial over $f^{-1}(U)$ since linearly independent sections s_i of E over U give rise to independent sections $a \mapsto (a, s_i(f(a)))$ of $f^*(E)$ over $f^{-1}(U)$. In particular, the pullback of a trivial bundle is a trivial bundle. This can also be seen directly from the definition, which in the case $E = B\times\mathbb{R}^n$ just says that $f^*(E)$ consists of the triples (a, b, v) in $A\times B\times\mathbb{R}^n$ with $b = f(a)$, so b is redundant and we have just the product $A\times\mathbb{R}^n$.

Now let us give some examples of pullbacks.

(1) The restriction of a vector bundle $p : E\to B$ over a subspace $A \subset B$ can be viewed as a pullback with respect to the inclusion map $A \hookrightarrow B$ since the inclusion $p^{-1}(A) \hookrightarrow E$ is certainly an isomorphism on each fiber.

(2) Another very special case is when the map f is a constant map, having image a single point $b \in B$. Then $f^*(E)$ is just the product $A\times p^{-1}(b)$, a trivial bundle.

(3) The tangent bundle TS^n is the pullback of the tangent bundle $T\mathbb{R}P^n$ via the quotient map $S^n\to\mathbb{R}P^n$. Indeed, $T\mathbb{R}P^n$ was defined as a quotient space of TS^n and the quotient map takes fibers isomorphically to fibers.

(4) An interesting example which is small enough to be visualized completely is the pullback of the Möbius bundle $E\to S^1$ by the two-to-one map $f : S^1\to S^1$ given by $f(z) = z^2$ in complex notation. One can regard the Möbius bundle as the quotient of $S^1\times\mathbb{R}$ under the identifications $(z, t) \sim (-z, -t)$, and the quotient map for this identification is the map f' exhibiting the annulus $S^1\times\mathbb{R}$ as the pullback of the Möbius bundle. More concretely, the pullback bundle can be thought of as a coat of paint applied to 'both sides' of the Möbius bundle, and the map f' sends each molecule of paint to the point of the Möbius band to which it adheres. Since E has one half-twist, the pullback has two half-twists, hence is the trivial bundle. More generally, if E_n is the pullback of the Möbius bundle by the map $z \mapsto z^n$, then E_n is the trivial bundle for n even and the Möbius bundle for n odd. This can be seen by viewing the Möbius bundle as the quotient of a strip $[0, 1]\times\mathbb{R}$ obtained by identifying the two edges $\{0\}\times\mathbb{R}$ and $\{1\}\times\mathbb{R}$ by a reflection of $\mathbb{R}$, and then the bundle E_n can be constructed from n such strips by identifying the right edge of the i^{th} strip to the

left edge of the $i+1^{st}$ strip by a reflection, the number i being taken modulo n so that the last strip is glued back to the first.

(5) At the end of the previous section we defined the flag bundle $F(E)$ associated to an n-dimensional vector bundle $E\to B$ to be the space of orthogonal direct sum decompositions of fibers of E into lines. The vectors in the i^{th} line form a line bundle $L_i\to F(E)$, and the direct sum $L_1\oplus\cdots\oplus L_n$ is nothing but the pullback of E with respect to the projection $F(E)\to B$ since a point in the pullback consists of an n-tuple of lines $\ell_1\perp\cdots\perp\ell_n$ in a fiber of E together with a vector v in this fiber, and v can be expressed uniquely as a sum $v=v_1+\cdots+v_n$ with $v_i\in\ell_i$. This construction for pulling an arbitrary bundle back to a sum of line bundles is a key ingredient in the so-called 'splitting principle' which is important in §2.3 and §3.1.

$$\begin{array}{ccc} L_1\oplus\cdots\oplus L_n & \longrightarrow & E \\ \downarrow & & \downarrow \\ F(E) & \longrightarrow & B \end{array}$$

Here are some elementary properties of pullbacks:

(i) $(fg)^*(E)\approx g^*(f^*(E))$.
(ii) $1\!\!1^*(E)\approx E$.
(iii) $f^*(E_1\oplus E_2)\approx f^*(E_1)\oplus f^*(E_2)$.
(iv) $f^*(E_1\otimes E_2)\approx f^*(E_1)\otimes f^*(E_2)$.

The proofs are easy applications of the preceding proposition. In each case one just checks that the bundle on the right satisfies the characteristic property of a pullback. For example in (iv) there is a natural map from $f^*(E_1)\otimes f^*(E_2)$ to $E_1\otimes E_2$ that is an isomorphism on each fiber, so $f^*(E_1)\otimes f^*(E_2)$ satisfies the condition for being the pullback $f^*(E_1\otimes E_2)$.

Now we come to the main technical result about pullbacks:

Theorem 1.6. *Given a vector bundle $p:E\to B$ and homotopic maps $f_0,f_1:A\to B$, then the induced bundles $f_0^*(E)$ and $f_1^*(E)$ are isomorphic if A is compact Hausdorff or more generally paracompact.*

Proof: Let $F:A\times I\to B$ be a homotopy from f_0 to f_1. The restrictions of $F^*(E)$ over $A\times\{0\}$ and $A\times\{1\}$ are then $f_0^*(E)$ and $f_1^*(E)$. So the theorem will be an immediate consequence of the following result:

Proposition 1.7. *The restrictions of a vector bundle $E\to X\times I$ over $X\times\{0\}$ and $X\times\{1\}$ are isomorphic if X is paracompact.*

Proof: We need two preliminary facts:

(1) A vector bundle $p:E\to X\times[a,b]$ is trivial if its restrictions over $X\times[a,c]$ and $X\times[c,b]$ are both trivial for some $c\in(a,b)$. To see this, let these restrictions be $E_1=p^{-1}(X\times[a,c])$ and $E_2=p^{-1}(X\times[c,b])$, and let $h_1:E_1\to X\times[a,c]\times\mathbb{R}^n$ and $h_2:E_2\to X\times[c,b]\times\mathbb{R}^n$ be isomorphisms. These isomorphisms may not agree on $p^{-1}(X\times\{c\})$, but they can be made to agree by replacing h_2 by its composition with

the isomorphism $X\times[c,b]\times\mathbb{R}^n \to X\times[c,b]\times\mathbb{R}^n$ which on each slice $X\times\{x\}\times\mathbb{R}^n$ is given by $h_1h_2^{-1}: X\times\{c\}\times\mathbb{R}^n \to X\times\{c\}\times\mathbb{R}^n$. Once h_1 and h_2 agree on $E_1\cap E_2$, they define a trivialization of E.

(2) For a vector bundle $p:E\to X\times I$, there exists an open cover $\{U_\alpha\}$ of X so that each restriction $p^{-1}(U_\alpha\times I)\to U_\alpha\times I$ is trivial. This is because for each $x\in X$ we can find open neighborhoods $U_{x,1},\cdots,U_{x,k}$ in X and a partition $0=t_0<t_1<\cdots<t_k=1$ of $[0,1]$ such that the bundle is trivial over $U_{x,i}\times[t_{i-1},t_i]$, using compactness of $[0,1]$. Then by (1) the bundle is trivial over $U_\alpha\times I$ where $U_\alpha = U_{x,1}\cap\cdots\cap U_{x,k}$.

Now we prove the proposition. By (2), we can choose an open cover $\{U_\alpha\}$ of X so that E is trivial over each $U_\alpha\times I$. Let us first deal with the simpler case that X is compact Hausdorff. Then a finite number of U_α's cover X. Relabel these as U_i for $i=1,2,\cdots,m$. As shown in Proposition 1.4 there is a corresponding partition of unity by functions φ_i with the support of φ_i contained in U_i. For $i\geq 0$, let $\psi_i=\varphi_1+\cdots+\varphi_i$, so in particular $\psi_0=0$ and $\psi_m=1$. Let X_i be the graph of ψ_i, the subspace of $X\times I$ consisting of points of the form $(x,\psi_i(x))$, and let $p_i:E_i\to X_i$ be the restriction of the bundle E over X_i. Since E is trivial over $U_i\times I$, the natural projection homeomorphism $X_i\to X_{i-1}$ lifts to a homeomorphism $h_i:E_i\to E_{i-1}$ which is the identity outside $p_i^{-1}(U_i)$ and which takes each fiber of E_i isomorphically onto the corresponding fiber of E_{i-1}. The composition $h=h_1h_2\cdots h_m$ is then an isomorphism from the restriction of E over $X\times\{1\}$ to the restriction over $X\times\{0\}$.

In the general case that X is only paracompact, Lemma 1.21 asserts that there is a countable cover $\{V_i\}_{i\geq 1}$ of X and a partition of unity $\{\varphi_i\}$ with φ_i supported in V_i, such that each V_i is a disjoint union of open sets each contained in some U_α. This means that E is trivial over each $V_i\times I$. As before we let $\psi_i=\varphi_1+\cdots+\varphi_i$ and let $p_i:E_i\to X_i$ be the restriction of E over the graph of ψ_i. Also as before we construct $h_i:E_i\to E_{i-1}$ using the fact that E is trivial over $V_i\times I$. The infinite composition $h=h_1h_2\cdots$ is then a well-defined isomorphism from the restriction of E over $X\times\{1\}$ to the restriction over $X\times\{0\}$ since near each point $x\in X$ only finitely many φ_i's are nonzero, so there is a neighborhood of x in which all but finitely many h_i's are the identity. □

Corollary 1.8. *A homotopy equivalence $f:A\to B$ of paracompact spaces induces a bijection $f^*:\mathrm{Vect}^n(B)\to\mathrm{Vect}^n(A)$. In particular, every vector bundle over a contractible paracompact base is trivial.*

Proof: If g is a homotopy inverse of f then we have $f^*g^* = \mathbb{1}^* = \mathbb{1}$ and $g^*f^* = \mathbb{1}^* = \mathbb{1}$. □

We might remark that Theorem 1.6 holds for fiber bundles as well as vector bundles, with the same proof.

Clutching Functions

Let us describe a way to construct vector bundles $E \to S^k$ with base space a sphere. Write S^k as the union of its upper and lower hemispheres D^k_+ and D^k_-, with $D^k_+ \cap D^k_- = S^{k-1}$. Given a map $f : S^{k-1} \to GL_n(\mathbb{R})$, let E_f be the quotient of the disjoint union $D^k_+ \times \mathbb{R}^n \amalg D^k_- \times \mathbb{R}^n$ obtained by identifying $(x, v) \in \partial D^k_- \times \mathbb{R}^n$ with $(x, f(x)(v)) \in \partial D^k_+ \times \mathbb{R}^n$. There is then a natural projection $E_f \to S^k$ and this is an n-dimensional vector bundle, as one can most easily see by taking an equivalent definition in which the two hemispheres of S^k are enlarged slightly to open balls and the identification occurs over their intersection, a product $S^{k-1} \times (-\varepsilon, \varepsilon)$, with the map f used in each slice $S^{k-1} \times \{t\}$. From this viewpoint the construction of E_f is a special case of the general construction of vector bundles by gluing together products, described earlier in the discussion of tensor products.

The map f is called a *clutching function* for E_f, by analogy with the mechanical clutch which engages and disengages gears in machinery. The same construction works equally well with $\mathbb{C}$ in place of $\mathbb{R}$, so from a map $f : S^{k-1} \to GL_n(\mathbb{C})$ one obtains a complex vector bundle $E_f \to S^k$.

Example 1.9. Let us see how the tangent bundle TS^2 to S^2 can be described in these terms. Consider the vector field v_+ on D^2_+ obtained by choosing a tangent vector at the north pole and transporting this along each meridian circle in such a way as to maintain a constant angle with the meridian. By reflecting across the plane of the equator we obtain a corresponding vector field v_- on D^2_-. Rotating each vector v_+ and v_- by 90 degrees counterclockwise when viewed from outside the sphere, we obtain vector fields w_+ and w_- on D^2_+ and D^2_-. The vector fields $v_\pm$ and $w_\pm$ give trivializations of TS^2 over $D^2_\pm$, allowing us to identify these two halves of TS^2 with $D^2_\pm \times \mathbb{R}^2$. We can then reconstruct TS^2 as a quotient of the disjoint union of $D^2_+ \times \mathbb{R}^2$ and $D^2_- \times \mathbb{R}^2$ by identifying along $S^1 \times \mathbb{R}^2$ using the clutching function f which reads off the coordinates of the v_-, w_- vectors in the v_+, w_+ coordinate system, since the definition of a clutching function was that it identifies $(x, v) \in \partial D^k_- \times \mathbb{R}^n$ with $(x, f(x)(v)) \in \partial D^k_+ \times \mathbb{R}^n$. Thus f rotates (v_+, w_+) to (v_-, w_-). As we go around the equator S^1 in the counterclockwise direction when viewed from above, starting from a point where the two trivializations agree, the angle of rotation increases from 0 to 4π. If we parametrize S^1 by the angle θ from the starting point, then $f(\theta)$ is rotation by 2θ.

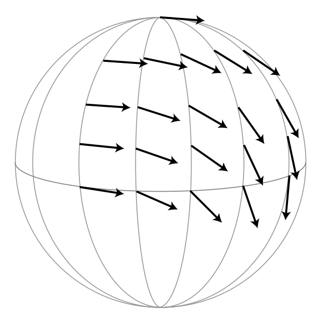

Example 1.10. Let us find a clutching function for the canonical complex line bundle over $\mathbb{C}\mathrm{P}^1 = S^2$. The space $\mathbb{C}\mathrm{P}^1$ is the quotient of $\mathbb{C}^2 - \{0\}$ under the equivalence relation $(z_0, z_1) \sim \lambda(z_0, z_1)$ for $\lambda \in \mathbb{C} - \{0\}$. Denote the equivalence class of (z_0, z_1) by $[z_0, z_1]$. We can also write points of $\mathbb{C}\mathrm{P}^1$ as ratios $z = z_0/z_1 \in \mathbb{C} \cup \{\infty\} = S^2$. If

we do this, then points in the disk D_0^2 inside the unit circle $S^1 \subset \mathbb{C}$ can be expressed uniquely in the form $[z_0/z_1, 1] = [z, 1]$ with $|z| \le 1$, and points in the disk D_∞^2 outside S^1 can be written uniquely in the form $[1, z_1/z_0] = [1, z^{-1}]$ with $|z^{-1}| \le 1$. Over D_0^2 a section of the canonical line bundle is given by $[z, 1] \mapsto (z, 1)$ and over D_∞^2 a section is $[1, z^{-1}] \mapsto (1, z^{-1})$. These sections determine trivializations of the canonical line bundle over these two disks, and over their common boundary S^1 we pass from the D_∞^2 trivialization to the D_0^2 trivialization by multiplying by z. Thus if we take D_∞^2 as D_+^2 and D_0^2 as D_-^2 then the canonical line bundle has the clutching function $f: S^1 \to GL_1(\mathbb{C})$ defined by $f(z) = (z)$. If we had taken D_+^2 to be D_0^2 rather than D_∞^2 then the clutching function would have been $f(z) = (z^{-1})$.

A basic property of the construction of bundles $E_f \to S^k$ via clutching functions $f: S^{k-1} \to GL_n(\mathbb{R})$ is that $E_f \approx E_g$ if f and g are homotopic. For if we have a homotopy $F: S^{k-1} \times I \to GL_n(\mathbb{R})$ from f to g, then the same sort of clutching construction can be used to produce a vector bundle $E_F \to S^k \times I$ that restricts to E_f over $S^k \times \{0\}$ and E_g over $S^k \times \{1\}$. Hence E_f and E_g are isomorphic by Proposition 1.7. Thus if we denote by $[X, Y]$ the set of homotopy classes of maps $X \to Y$, then the association $f \mapsto E_f$ gives a well-defined map $\Phi: [S^{k-1}, GL_n(\mathbb{R})] \longrightarrow \mathrm{Vect}^n(S^k)$.

There is a similarly-defined map Φ for complex vector bundles. This turns out to have slightly better behavior than in the real case, so we will prove the following basic result about the complex form of Φ before dealing with the real case.

Proposition 1.11. *The map $\Phi: [S^{k-1}, GL_n(\mathbb{C})] \to \mathrm{Vect}_{\mathbb{C}}^n(S^k)$ which sends a clutching function f to the vector bundle E_f is a bijection.*

Proof: We construct an inverse Ψ to Φ. Given an n-dimensional vector bundle $p: E \to S^k$, its restrictions E_+ and E_- over D_+^k and D_-^k are trivial since D_+^k and D_-^k are contractible. Choose trivializations $h_\pm: E_\pm \to D_\pm^k \times \mathbb{C}^n$. Then $h_+ h_-^{-1}$ defines a map $S^{k-1} \to GL_n(\mathbb{C})$, whose homotopy class is by definition $\Psi(E) \in [S^{k-1}, GL_n(\mathbb{C})]$. To see that $\Psi(E)$ is well-defined, note first that any two choices of $h_\pm$ differ by a map $D_\pm^k \to GL_n(\mathbb{C})$. Since $D_\pm^k$ is contractible, such a map is homotopic to a constant map by composing with a contraction of $D_\pm^k$. Now we need the fact that $GL_n(\mathbb{C})$ is path-connected. The elementary row operation of modifying a matrix in $GL_n(\mathbb{C})$ by adding a scalar multiple of one row to another row can be realized by a path in $GL_n(\mathbb{C})$, just by inserting a factor of t in front of the scalar multiple and letting t go from 0 to 1. By such operations any matrix in $GL_n(\mathbb{C})$ can be diagonalized. The set of diagonal matrices in $GL_n(\mathbb{C})$ is path-connected since it is homeomorphic to the product of n copies of $\mathbb{C} - \{0\}$.

From this we conclude that h_+ and h_- are unique up to homotopy. Hence the composition $h_+ h_-^{-1}: S^{k-1} \to GL_n(\mathbb{C})$ is also unique up to homotopy, which means that Ψ is a well-defined map $\mathrm{Vect}_{\mathbb{C}}^n(S^k) \to [S^{k-1}, GL_n(\mathbb{C})]$. It is clear that Ψ and Φ are inverses of each other. $\square$

Example 1.12. Every complex vector bundle over S^1 is trivial, since by the proposition this is equivalent to saying that $[S^0, GL_n(\mathbb{C})]$ has a single element, or in other words that $GL_n(\mathbb{C})$ is path-connected, which is the case as we saw in the proof of the proposition.

Example 1.13. Let us show that the canonical line bundle $H \to \mathbb{C}\mathrm{P}^1$ satisfies the relation $(H \otimes H) \oplus 1 \approx H \oplus H$ where 1 is the trivial one-dimensional bundle. This can be seen by looking at the clutching functions for these two bundles, which are the maps $S^1 \to GL_2(\mathbb{C})$ given by

$$z \mapsto \begin{pmatrix} z^2 & 0 \\ 0 & 1 \end{pmatrix} \qquad \text{and} \qquad z \mapsto \begin{pmatrix} z & 0 \\ 0 & z \end{pmatrix}$$

More generally, let us show the formula $E_{fg} \oplus n \approx E_f \oplus E_g$ for n-dimensional vector bundles E_f and E_g over S^k with clutching functions $f, g : S^{k-1} \to GL_n(\mathbb{C})$, where fg is the clutching function obtained by pointwise matrix multiplication, $fg(x) = f(x)g(x)$.

The bundle $E_f \oplus E_g$ has clutching function the map $f \oplus g : S^{k-1} \to GL_{2n}(\mathbb{C})$ having the matrices $f(x)$ in the upper left $n \times n$ block and the matrices $g(x)$ in the lower right $n \times n$ block, the other two blocks being zero. Since $GL_{2n}(\mathbb{C})$ is path-connected, there is a path $\alpha_t \in GL_{2n}(\mathbb{C})$ from the identity matrix to the matrix of the transformation which interchanges the two factors of $\mathbb{C}^n \times \mathbb{C}^n$. Then the matrix product $(f \oplus \mathbb{1})\alpha_t(\mathbb{1} \oplus g)\alpha_t$ gives a homotopy from $f \oplus g$ to $fg \oplus \mathbb{1}$, which is the clutching function for $E_{fg} \oplus n$.

The preceding analysis does not quite work for real vector bundles since $GL_n(\mathbb{R})$ is not path-connected. We can see that there are at least two path-components since the determinant function is a continuous surjection $GL_n(\mathbb{R}) \to \mathbb{R} - \{0\}$ to a space with two path-components. In fact $GL_n(\mathbb{R})$ has exactly two path-components, as we will now show. Just as in the complex case we can construct a path from an arbitrary matrix in $GL_n(\mathbb{R})$ to a diagonal matrix. Then by a path of diagonal matrices we can make all the diagonal entries $+1$ or -1. Two -1's represent a 180 degree rotation of a plane so they can be replaced by $+1$'s via a path in $GL_n(\mathbb{R})$. This shows that the subgroup $GL_n^+(\mathbb{R})$ consisting of matrices of positive determinant is path-connected. This subgroup has index 2, and $GL_n(\mathbb{R})$ is the disjoint union of the two cosets $GL_n^+(\mathbb{R})$ and $\alpha GL_n^+(\mathbb{R})$ for α a fixed matrix of determinant -1. The two cosets are homeomorphic since the map $\beta \mapsto \alpha\beta$ is a homeomorphism of $GL_n(\mathbb{R})$ with inverse the map $\beta \mapsto \alpha^{-1}\beta$. Thus both cosets are path-connected and so $GL_n(\mathbb{R})$ has two path-components.

The closest analogy with the complex case is obtained by considering oriented real vector bundles. Recall from linear algebra that an orientation of a real vector space is an equivalence class of ordered bases, two ordered bases being equivalent if the invertible matrix taking the first basis to the second has positive determinant.

An *orientation* of a real vector bundle $p:E \to B$ is a function assigning an orientation to each fiber in such a way that near each point of B there is a local trivialization $h:p^{-1}(U) \to U \times \mathbb{R}^n$ carrying the orientations of fibers in $p^{-1}(U)$ to the standard orientation of $\mathbb{R}^n$ in the fibers of $U \times \mathbb{R}^n$. Another way of stating this condition would be to say that the orientations of fibers of E can be defined by ordered n-tuples of independent local sections. Not all vector bundles can be given an orientation. For example the Möbius bundle is not orientable. This is because an oriented line bundle over a paracompact base is always trivial since it has a canonical section formed by the unit vectors having positive orientation.

Let $\mathrm{Vect}_+^n(B)$ denote the set of isomorphism classes of oriented n-dimensional real vector bundles over B, where isomorphisms are required to preserve orientations. The clutching construction defines a map $\Phi:[S^{k-1}, GL_n^+(\mathbb{R})] \longrightarrow \mathrm{Vect}_+^n(S^k)$, and since $GL_n^+(\mathbb{R})$ is path-connected, the argument from the complex case shows:

Proposition 1.14. *The map* $\Phi:[S^{k-1}, GL_n^+(\mathbb{R})] \longrightarrow \mathrm{Vect}_+^n(S^k)$ *is a bijection.* □

To analyze $\mathrm{Vect}^n(S^k)$ let us introduce an *ad hoc* hybrid object $\mathrm{Vect}_0^n(S^k)$, the n-dimensional vector bundles over S^k with an orientation specified in the fiber over one point $x_0 \in S^{k-1}$, with the equivalence relation of isomorphism preserving the orientation of the fiber over x_0. We can choose the trivializations $h_\pm$ over $D_\pm^k$ to carry this orientation to a standard orientation of $\mathbb{R}^n$, and then $h_\pm$ is unique up to homotopy as before, so we obtain a bijection of $\mathrm{Vect}_0^n(S^k)$ with the homotopy classes of maps $S^{k-1} \to GL_n(\mathbb{R})$ taking x_0 to $GL_n^+(\mathbb{R})$.

The map $\mathrm{Vect}_0^n(S^k) \to \mathrm{Vect}^n(S^k)$ that forgets the orientation over x_0 is a surjection that is two-to-one except on vector bundles that have an automorphism (an isomorphism from the bundle to itself) reversing the orientation of the fiber over x_0, where it is one-to-one.

When $k = 1$ there are just two homotopy classes of maps $S^0 \to GL_n(\mathbb{R})$ taking x_0 to $GL_n^+(\mathbb{R})$, represented by maps taking x_0 to the identity and the other point of S^0 to either the identity or a reflection. The corresponding bundles are the trivial bundle and, when $n = 1$, the Möbius bundle, or the direct sum of the Möbius bundle with a trivial bundle when $n > 1$. These bundles all have automorphisms reversing orientations of fibers. We conclude that $\mathrm{Vect}^n(S^1)$ has exactly two elements for each $n \geq 1$. The nontrivial bundle is nonorientable since an orientable bundle over S^1 is trivial, any two clutching functions $S^0 \to GL_n^+(\mathbb{R})$ being homotopic.

When $k > 1$ the sphere S^{k-1} is path-connected, so maps $S^{k-1} \to GL_n(\mathbb{R})$ taking x_0 to $GL_n^+(\mathbb{R})$ take all of S^{k-1} to $GL_n^+(\mathbb{R})$. This implies that the natural map $\mathrm{Vect}_+^n(S^k) \to \mathrm{Vect}_0^n(S^k)$ is a bijection. Thus every vector bundle over S^k is orientable and has exactly two different orientations, determined by an orientation in one fiber. It follows that the map $\mathrm{Vect}_+^n(S^k) \to \mathrm{Vect}^n(S^k)$ is surjective and at most two-to-one.

It is one-to-one on bundles having an orientation-reversing automorphism, and two-to-one otherwise.

Inside $GL_n(\mathbb{R})$ is the orthogonal group $O(n)$, the subgroup consisting of isometries, or in matrix terms, the $n\times n$ matrices A with $AA^T = I$, that is, matrices whose columns form an orthonormal basis. The Gram-Schmidt orthogonalization process provides a deformation retraction of $GL_n(\mathbb{R})$ onto the subspace $O(n)$. Each step of the process is either normalizing the i^{th} vector of a basis to have length 1 or subtracting a linear combination of the preceding vectors from the i^{th} vector in order to make it orthogonal to them. Both operations are realizable by a homotopy. This is obvious for rescaling, and for the second operation one can just insert a scalar factor of t in front of the linear combination of vectors being subtracted and let t go from 0 to 1. The explicit algebraic formulas show that both operations depend continuously on the initial basis and have no effect on an orthonormal basis. Hence the sequence of operations provides a deformation retraction of $GL_n(\mathbb{R})$ onto $O(n)$. Restricting to the component $GL_n^+(\mathbb{R})$, we have a deformation retraction of this component onto the special orthogonal group $SO(n)$, the orthogonal matrices of determinant $+1$. This is the path-component of $O(n)$ containing the identity matrix since deformation retractions preserve path-components. The other component consists of orthogonal matrices of determinant -1.

In the complex case the same argument applied to $GL_n(\mathbb{C})$ gives a deformation retraction onto the unitary group $U(n)$, the $n\times n$ matrices over $\mathbb{C}$ whose columns form an orthonormal basis using the standard Hermitian inner product.

Composing maps $S^{k-1}\to SO(n)$ with the inclusion $SO(n)\hookrightarrow GL_n^+(\mathbb{R})$ gives a function $[S^{k-1},SO(n)]\to[S^{k-1},GL_n^+(\mathbb{R})]$. This is a bijection since the deformation retraction of $GL_n^+(\mathbb{R})$ onto $SO(n)$ gives a homotopy of any map $S^{k-1}\to GL_n^+(\mathbb{R})$ to a map into $SO(n)$, and if two such maps to $SO(n)$ are homotopic in $GL_n^+(\mathbb{R})$ then they are homotopic in $SO(n)$ by composing each stage of the homotopy with the retraction $GL_n^+(\mathbb{R})\to SO(n)$ produced by the deformation retraction.

The advantage of $SO(n)$ over $GL_n^+(\mathbb{R})$ is that it is considerably smaller. For example, $SO(2)$ is just a circle, since orientation-preserving isometries of $\mathbb{R}^2$ are rotations, determined by the angle of rotation, a point in S^1. Taking $k=2$ as well, we have for each integer $m\in\mathbb{Z}$ the map $f_m:S^1\to S^1$, $z\mapsto z^m$, which is the clutching function for an oriented 2-dimensional vector bundle $E_m\to S^2$. We encountered E_2 in Example 1.9 as TS^2, and E_1 is the real vector bundle underlying the canonical complex line bundle over $\mathbb{C}\mathrm{P}^1 = S^2$, as we saw in Example 1.10. It is a fact proved in every introductory algebraic topology course that an arbitrary map $f:S^1\to S^1$ is homotopic to one of the maps $z\mapsto z^m$ for some m uniquely determined by f, called the degree of f. Thus via the bundles E_m we have a bijection $\mathrm{Vect}_+^2(S^2)\approx\mathbb{Z}$. Reversing the orientation of E_m changes it to E_{-m}, so elements of $\mathrm{Vect}^2(S^2)$ correspond bijectively with integers $m \geq 0$, via the bundles E_m. Since $U(1)$ is also homeomorphic to S^1,we obtain a similar bijection $\mathrm{Vect}^2_{\mathbb{C}}(S^2) \approx \mathbb{Z}$.

Another fact proved near the beginning of algebraic topology is that any two maps $S^{k-1} \to S^1$ are homotopic if $k > 2$. (Such a map lifts to the covering space $\mathbb{R}$ of S^1, and $\mathbb{R}$ is contractible.) Hence every 2-dimensional vector bundle over S^k is trivial when $k > 2$, as is any complex line bundle over S^k for $k > 2$.

It is a fact that any two maps $S^2 \to SO(n)$ or $S^2 \to U(n)$ are homotopic, so all vector bundles over S^3 are trivial. This is not too hard to show using a little basic machinery of algebraic topology, as we show in §??, where we also compute how many homotopy classes of maps $S^3 \to SO(n)$ and $S^3 \to U(n)$ there are.

For any space X the set $[X, SO(n)]$ has a natural group structure coming from the group structure in $SO(n)$. Namely, the product of two maps $f, g : X \to SO(n)$ is the map $x \mapsto f(x)g(x)$. Similarly $[X, U(n)]$ is a group. For example the bijection $[S^1, SO(2)] \approx \mathbb{Z}$ is a group isomorphism since for maps $S^1 \to S^1$, the product of $z \mapsto z^m$ and $z \mapsto z^n$ is $z \mapsto z^m z^n$, and $z^m z^n = z^{m+n}$.

The groups $[S^i, SO(n)]$ and $[S^i, U(n)]$ are isomorphic to the homotopy groups $\pi_i SO(n)$ and $\pi_i U(n)$, special cases of the homotopy groups $\pi_i X$ defined for all spaces X and central to algebraic topology. To define $\pi_i X$ requires choosing a basepoint $x_0 \in X$, and then $\pi_i X$ is the set of homotopy classes of maps $S^i \to X$ taking a chosen basepoint $s_0 \in S^i$ to x_0, where homotopies are also required to take s_0 to x_0 at all times. This basepoint data is needed in order to define a group structure in $\pi_i X$ when $i > 0$. (When $i = 0$ there is no natural group structure in general.) However, the group structure on $\pi_i X$ turns out to be independent of the choice of x_0 when X is path-connected. Taking $X = SO(n)$, there is an evident map $\pi_i SO(n) \to [S^i, SO(n)]$ obtained by ignoring basepoint data. To see that this is surjective, suppose we are given a map $f : S^i \to SO(n)$. Since $SO(n)$ is path-connected, there is a path α_t in $SO(n)$ from the identity matrix to $f(s_0)$, and then the matrix product $\alpha_t^{-1} f$ defines a homotopy from f to a map $S^i \to SO(n)$ taking s_0 to the identity matrix, which we choose as the basepoint of $SO(n)$. For injectivity of the map $\pi_i SO(n) \to [S^i, SO(n)]$ suppose that f_0 and f_1 are two maps $S^i \to SO(n)$ taking s_0 to the identity matrix, and suppose f_t is a homotopy from f_0 to f_1 that may not take basepoint to basepoint at all times. Then $[f_t(s_0)]^{-1} f_t$ is a new homotopy from f_0 to f_1 that does take basepoint to basepoint for all t. The same argument also applies for $U(n)$.

Note to the reader: The rest of this chapter will not be needed until Chapter 3.

The Universal Bundle

We will show that there is a special n-dimensional vector bundle $E_n \to G_n$ with the property that all n-dimensional bundles over paracompact base spaces are obtainable as pullbacks of this single bundle. When $n = 1$ this bundle will be just the canonical line bundle over $\mathbb{R}P^\infty$, defined earlier. The generalization to $n > 1$ will consist in replacing $\mathbb{R}P^\infty$, the space of 1-dimensional vector subspaces of $\mathbb{R}^\infty$, by the space of n-dimensional vector subspaces of $\mathbb{R}^\infty$.

First we define the *Grassmann manifold* $G_n(\mathbb{R}^k)$ for nonnegative integers $n \le k$. As a set this is the collection of all n-dimensional vector subspaces of $\mathbb{R}^k$, that is, n-dimensional planes in $\mathbb{R}^k$ passing through the origin. To define a topology on $G_n(\mathbb{R}^k)$ we first define the *Stiefel manifold* $V_n(\mathbb{R}^k)$ to be the space of orthonormal n-frames in $\mathbb{R}^k$, in other words, n-tuples of orthonormal vectors in $\mathbb{R}^k$. This is a subspace of the product of n copies of the unit sphere S^{k-1}, namely, the subspace of orthogonal n-tuples. It is a closed subspace since orthogonality of two vectors can be expressed by an algebraic equation. Hence $V_n(\mathbb{R}^k)$ is compact since the product of spheres is compact. There is a natural surjection $V_n(\mathbb{R}^k) \to G_n(\mathbb{R}^k)$ sending an n-frame to the subspace it spans, and $G_n(\mathbb{R}^k)$ is topologized by giving it the quotient topology with respect to this surjection. So $G_n(\mathbb{R}^k)$ is compact as well. Later in this section we will construct a finite CW complex structure on $G_n(\mathbb{R}^k)$ and in the process show that it is Hausdorff and a manifold of dimension $n(k-n)$.

The inclusions $\mathbb{R}^k \subset \mathbb{R}^{k+1} \subset \cdots$ give inclusions $G_n(\mathbb{R}^k) \subset G_n(\mathbb{R}^{k+1}) \subset \cdots$, and we let $G_n(\mathbb{R}^\infty) = \bigcup_k G_n(\mathbb{R}^k)$. This is the set of all n-dimensional vector subspaces of the vector space $\mathbb{R}^\infty$ since if we choose a basis for a finite-dimensional subspace of $\mathbb{R}^\infty$, each basis vector will lie in some $\mathbb{R}^k$, hence there will be an $\mathbb{R}^k$ containing all the basis vectors and therefore the whole subspace. We give $G_n(\mathbb{R}^\infty)$ the weak or direct limit topology, so a set in $G_n(\mathbb{R}^\infty)$ is open iff it intersects each $G_n(\mathbb{R}^k)$ in an open set.

There are canonical n-dimensional vector bundles over $G_n(\mathbb{R}^k)$ and $G_n(\mathbb{R}^\infty)$. Define $E_n(\mathbb{R}^k) = \{(\ell, v) \in G_n(\mathbb{R}^k) \times \mathbb{R}^k \mid v \in \ell\}$. The inclusions $\mathbb{R}^k \subset \mathbb{R}^{k+1} \subset \cdots$ give inclusions $E_n(\mathbb{R}^k) \subset E_n(\mathbb{R}^{k+1}) \subset \cdots$ and we set $E_n(\mathbb{R}^\infty) = \bigcup_k E_n(\mathbb{R}^k)$, again with the weak or direct limit topology.

Lemma 1.15. *The projection $p : E_n(\mathbb{R}^k) \to G_n(\mathbb{R}^k)$, $p(\ell, v) = \ell$, is a vector bundle., both for finite and infinite k.*

Proof: First suppose k is finite. For $\ell \in G_n(\mathbb{R}^k)$, let $\pi_\ell : \mathbb{R}^k \to \ell$ be orthogonal projection and let $U_\ell = \{\ell' \in G_n(\mathbb{R}^k) \mid \pi_\ell(\ell') \text{ has dimension } n\}$. In particular, $\ell \in U_\ell$. We will show that U_ℓ is open in $G_n(\mathbb{R}^k)$ and that the map $h : p^{-1}(U_\ell) \to U_\ell \times \ell \approx U_\ell \times \mathbb{R}^n$ defined by $h(\ell', v) = (\ell', \pi_\ell(v))$ is a local trivialization of $E_n(\mathbb{R}^k)$.

For U_ℓ to be open is equivalent to its preimage in $V_n(\mathbb{R}^k)$ being open. This preimage consists of orthonormal frames $v_1, \cdots, v_n$ such that $\pi_\ell(v_1), \cdots, \pi_\ell(v_n)$ are independent. Let A be the matrix of π_ℓ with respect to the standard basis in the domain $\mathbb{R}^k$ and any fixed basis in the range ℓ. The condition on $v_1, \cdots, v_n$ is then that the $n \times n$ matrix with columns $Av_1, \cdots, Av_n$ have nonzero determinant. Since the value of this determinant is obviously a continuous function of $v_1, \cdots, v_n$, it follows that the frames $v_1, \cdots, v_n$ yielding a nonzero determinant form an open set in $V_n(\mathbb{R}^k)$.

It is clear that h is a bijection which is a linear isomorphism on each fiber. We need to check that h and h^{-1} are continuous. For $\ell' \in U_\ell$ there is a unique invertible linear map $L_{\ell'} : \mathbb{R}^k \to \mathbb{R}^k$ restricting to π_ℓ on ℓ' and the identity on $\ell^\perp = \operatorname{Ker} \pi_\ell$. We claim that $L_{\ell'}$, regarded as a $k \times k$ matrix, depends continuously on ℓ'. Namely, we can write $L_{\ell'}$ as a product AB^{-1} where:

— B sends the standard basis to $v_1, \cdots, v_n, v_{n+1}, \cdots, v_k$ with $v_1, \cdots, v_n$ an orthonormal basis for ℓ' and $v_{n+1}, \cdots, v_k$ a fixed basis for $\ell^\perp$.

— A sends the standard basis to $\pi_\ell(v_1), \cdots, \pi_\ell(v_n), v_{n+1}, \cdots, v_k$.

Both A and B depend continuously on $v_1, \cdots, v_n$. Since matrix multiplication and matrix inversion are continuous operations (think of the 'classical adjoint' formula for the inverse of a matrix), it follows that the product $L_{\ell'} = AB^{-1}$ depends continuously on $v_1, \cdots, v_n$. But since $L_{\ell'}$ depends only on ℓ', not on the basis $v_1, \cdots, v_n$ for ℓ', it follows that $L_{\ell'}$ depends continuously on ℓ' since $G_n(\mathbb{R}^k)$ has the quotient topology from $V_n(\mathbb{R}^k)$. Since we have $h(\ell', v) = (\ell', \pi_\ell(v)) = (\ell', L_{\ell'}(v))$, we see that h is continuous. Similarly, $h^{-1}(\ell', w) = (\ell', L_{\ell'}^{-1}(w))$ and $L_{\ell'}^{-1}$ depends continuously on ℓ', matrix inversion being continuous, so h^{-1} is continuous.

This finishes the proof for finite k. When $k = \infty$ one takes U_ℓ to be the union of the U_ℓ's for increasing k. The local trivializations h constructed above for finite k then fit together to give a local trivialization over this U_ℓ, continuity being automatic since we use the weak topology. $\square$

We will mainly be interested in the case $k = \infty$ now, and to simplify notation we will write G_n for $G_n(\mathbb{R}^\infty)$ and E_n for $E_n(\mathbb{R}^\infty)$. As earlier in this section, we use the notation $[X, Y]$ for the set of homotopy classes of maps $f : X \to Y$.

Theorem 1.16. *For paracompact X, the map $[X, G_n] \to \operatorname{Vect}^n(X)$, $[f] \mapsto f^*(E_n)$, is a bijection.*

Thus, vector bundles over a fixed base space are classified by homotopy classes of maps into G_n. Because of this, G_n is called the *classifying space* for n-dimensional vector bundles and $E_n \to G_n$ is called the *universal bundle*.

As an example of how a vector bundle could be isomorphic to a pullback $f^*(E_n)$, consider the tangent bundle to S^n. This is the vector bundle $p : E \to S^n$ where $E = \{ (x, v) \in S^n \times \mathbb{R}^{n+1} \mid x \perp v \}$. Each fiber $p^{-1}(x)$ is a point in $G_n(\mathbb{R}^{n+1})$, so we have a map $S^n \to G_n(\mathbb{R}^{n+1})$, $x \mapsto p^{-1}(x)$. Via the inclusion $\mathbb{R}^{n+1} \hookrightarrow \mathbb{R}^\infty$ we can view this as a map $f : S^n \to G_n(\mathbb{R}^\infty) = G_n$, and E is exactly the pullback $f^*(E_n)$.

Proof of 1.16: The key observation is the following: For an n-dimensional vector bundle $p : E \to X$, an isomorphism $E \approx f^*(E_n)$ is equivalent to a map $g : E \to \mathbb{R}^\infty$ that is a linear injection on each fiber. To see this, suppose first that we have a map $f : X \to G_n$ and an isomorphism $E \approx f^*(E_n)$. Then we have a commutative diagram

$$\begin{array}{ccccc} E \approx f^*(E_n) & \xrightarrow{\widetilde{f}} & E_n & \xrightarrow{\pi} & \mathbb{R}^\infty \\ \downarrow & & \downarrow & & \\ X & \xrightarrow{f} & G_n & & \end{array}$$

where $\pi(\ell, v) = v$. The composition across the top row is a map $g : E \to \mathbb{R}^\infty$ that is a linear injection on each fiber, since both $\widetilde{f}$ and π have this property. Conversely, given a map $g : E \to \mathbb{R}^\infty$ that is a linear injection on each fiber, define $f : X \to G_n$ by letting $f(x)$ be the n-plane $g(p^{-1}(x))$. This clearly yields a commutative diagram as above.

To show surjectivity of the map $[X, G_n] \longrightarrow \mathrm{Vect}^n(X)$, suppose $p : E \to X$ is an n-dimensional vector bundle. Let $\{U_\alpha\}$ be an open cover of X such that E is trivial over each U_α. By Lemma 1.21 in the Appendix to this chapter there is a countable open cover $\{U_i\}$ of X such that E is trivial over each U_i, and there is a partition of unity $\{\varphi_i\}$ with φ_i supported in U_i. Let $g_i : p^{-1}(U_i) \to \mathbb{R}^n$ be the composition of a trivialization $p^{-1}(U_i) \to U_i \times \mathbb{R}^n$ with projection onto $\mathbb{R}^n$. The map $(\varphi_i p) g_i$, $v \mapsto \varphi_i(p(v)) g_i(v)$, extends to a map $E \to \mathbb{R}^n$ that is zero outside $p^{-1}(U_i)$. Near each point of X only finitely many φ_i's are nonzero, and at least one φ_i is nonzero, so these extended $(\varphi_i p) g_i$'s are the coordinates of a map $g : E \to (\mathbb{R}^n)^\infty = \mathbb{R}^\infty$ that is a linear injection on each fiber.

For injectivity, if we have isomorphisms $E \approx f_0^*(E_n)$ and $E \approx f_1^*(E_n)$ for two maps $f_0, f_1 : X \to G_n$, then these give maps $g_0, g_1 : E \to \mathbb{R}^\infty$ that are linear injections on fibers, as in the first paragraph of the proof. We claim g_0 and g_1 are homotopic through maps g_t that are linear injections on fibers. If this is so, then f_0 and f_1 will be homotopic via $f_t(x) = g_t(p^{-1}(x))$.

The first step in constructing a homotopy g_t is to compose g_0 with the homotopy $L_t : \mathbb{R}^\infty \to \mathbb{R}^\infty$ defined by $L_t(x_1, x_2, \cdots) = (1-t)(x_1, x_2, \cdots) + t(x_1, 0, x_2, 0, \cdots)$. For each t this is a linear map whose kernel is easily computed to be 0, so L_t is injective. Composing the homotopy L_t with g_0 moves the image of g_0 into the odd-numbered coordinates. Similarly we can homotope g_1 into the even-numbered coordinates. Still calling the new g's g_0 and g_1, let $g_t = (1-t)g_0 + t g_1$. This is linear and injective on fibers for each t since g_0 and g_1 are linear and injective on fibers. □

An explicit calculation of $[X, G_n]$ is usually beyond the reach of what is possible technically, so this theorem is of limited usefulness in enumerating all the different vector bundles over a given base space. Its importance is due more to its theoretical implications. Among other things, it can reduce the proof of a general statement to the special case of the universal bundle. For example, it is easy to deduce that vector bundles over a paracompact base have inner products, since the bundle $E_n \to G_n$ has an obvious inner product obtained by restricting the standard inner product in $\mathbb{R}^\infty$ to each n-plane, and this inner product on E_n induces an inner product on every pullback $f^*(E_n)$.

The preceding constructions and results hold equally well for vector bundles over $\mathbb{C}$, with $G_n(\mathbb{C}^k)$ the space of n-dimensional $\mathbb{C}$-linear subspaces of $\mathbb{C}^k$, and so on. In particular, the proof of Theorem 1.16 translates directly to complex vector bundles, showing that $\mathrm{Vect}^n_{\mathbb{C}}(X) \approx [X, G_n(\mathbb{C}^\infty)]$.

There is also a version of Theorem 1.16 for oriented vector bundles. Let $\widetilde{G}_n(\mathbb{R}^k)$ be the space of oriented n-planes in $\mathbb{R}^k$, the quotient space of $V_n(\mathbb{R}^k)$ obtained by identifying two n-frames when they determine the same oriented subspace of $\mathbb{R}^k$. Forming the union over increasing k we then have the space $\widetilde{G}_n(\mathbb{R}^\infty)$. The universal oriented bundle $\widetilde{E}_n(\mathbb{R}^\infty)$ over $\widetilde{G}_n(\mathbb{R}^\infty)$ consists of pairs $(\ell, v) \in \widetilde{G}_n(\mathbb{R}^\infty) \times \mathbb{R}^\infty$ with $v \in \ell$. In other words, $\widetilde{E}_n(\mathbb{R}^\infty)$ is the pullback of $E_n(\mathbb{R}^\infty)$ via the natural projection $\widetilde{G}_n(\mathbb{R}^\infty) \to G_n(\mathbb{R}^\infty)$. Small modifications in the proof that $\mathrm{Vect}^n(X) \approx [X, G_n(\mathbb{R}^\infty)]$ show that $\mathrm{Vect}^n_+(X) \approx [X, \widetilde{G}_n(\mathbb{R}^\infty)]$.

Both $G_n(\mathbb{R}^\infty)$ and $\widetilde{G}_n(\mathbb{R}^\infty)$ are path-connected since $\mathrm{Vect}^n(X)$ and $\mathrm{Vect}^n_+(X)$ have a single element when X is a point. The natural projection $\widetilde{G}_n(\mathbb{R}^\infty) \to G_n(\mathbb{R}^\infty)$ obtained by ignoring orientations is two-to-one, and readers familiar with the notion of a covering space will have no trouble in recognizing that this two-to-one projection map is a covering space, using for example the local trivializations constructed in Lemma 1.15. In fact $\widetilde{G}_n(\mathbb{R}^\infty)$ is the universal cover of $G_n(\mathbb{R}^\infty)$ since it is simply-connected, because of the triviality of $\mathrm{Vect}^n_+(S^1) \approx [S^1, \widetilde{G}_n(\mathbb{R}^\infty)]$. One can also observe that a vector bundle $E \approx f^*(E_n(\mathbb{R}^\infty))$ is orientable iff its classifying map $f : X \to G_n(\mathbb{R}^\infty)$ lifts to a map $\widetilde{f} : X \to \widetilde{G}_n(\mathbb{R}^\infty)$, and in fact orientations of E correspond bijectively with lifts $\widetilde{f}$.

Cell Structures on Grassmannians

Since Grassmann manifolds play a fundamental role in vector bundle theory, it would be good to have a better grasp on their topology. Here we show that $G_n(\mathbb{R}^\infty)$ has the structure of a CW complex with each $G_n(\mathbb{R}^k)$ a finite subcomplex. We will also see that $G_n(\mathbb{R}^k)$ is a closed manifold of dimension $n(k-n)$. Similar statements hold in the complex case as well, with $G_n(\mathbb{C}^k)$ a closed manifold of dimension $2n(k-n)$.

For a start let us show that $G_n(\mathbb{R}^k)$ is Hausdorff, since we will need this fact later when we construct the CW structure. Given two n-planes ℓ and ℓ' in $G_n(\mathbb{R}^k)$, it suffices to find a continuous map $f : G_n(\mathbb{R}^k) \to \mathbb{R}$ taking different values on ℓ and ℓ'. For a vector $v \in \mathbb{R}^k$ let $f_v(\ell)$ be the length of the orthogonal projection of v onto ℓ. This is a continuous function of ℓ since if we choose an orthonormal basis $v_1, \cdots, v_n$ for ℓ then $f_v(\ell) = \left((v \cdot v_1)^2 + \cdots + (v \cdot v_n)^2\right)^{1/2}$, which is certainly continuous in $v_1, \cdots, v_n$ hence in ℓ since $G_n(\mathbb{R}^k)$ has the quotient topology from $V_n(\mathbb{R}^k)$. Now for an n-plane $\ell' \neq \ell$ choose $v \in \ell - \ell'$, and then $f_v(\ell) = |v| > f_v(\ell')$.

There is a nice description of the cells in the CW structure on $G_n(\mathbb{R}^k)$ in terms of the familiar concept of echelon form for matrices. Recall that any $n \times k$ matrix A can be put into an echelon form by a finite sequence of elementary row operations consisting either adding a multiple of one row to another, multiplying a row by a nonzero scalar, or permuting two rows. In the standard version of echelon form one strives to make zeros in the lower left corner of the matrix, but for our purposes it will be more convenient to make zeros in the upper right corner instead. Thus a matrix in echelon form will look like the following, with the asterisks denoting entries that are arbitrary numbers:

$$\begin{bmatrix} * & * & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ * & * & 0 & * & 1 & 0 & 0 & 0 & 0 & 0 \\ * & * & 0 & * & 0 & 1 & 0 & 0 & 0 & 0 \\ * & * & 0 & * & 0 & 0 & * & * & 1 & 0 \end{bmatrix}$$

We are assuming that our given $n \times k$ matrix A has rank n and that $n \le k$. The shape of the echelon form is specified by which columns contain the special entries 1, say in the columns numbered $\sigma_1 < \cdots < \sigma_n$. The n-tuple $(\sigma_1, \cdots, \sigma_n)$ is called the *Schubert symbol* $\sigma(A)$. To see that this depends only on the given matrix A and not any particular reduction of A to echelon form, consider the n-plane ℓ_A in $\mathbb{R}^k$ spanned by the rows of A. This is the same as the n-plane spanned by the rows of an echelon form of A since elementary row operations do not change the subspace spanned by the rows. Let $p_i : \mathbb{R}^k \to \mathbb{R}^{n-i}$ be projection onto the last $n-i$ coordinates. As i goes from 0 to k the dimension of the subspace $p_i(\ell_A)$ of $\mathbb{R}^k$ decreases from n to 0 in n steps and the decreases occur precisely when i hits one of the values σ_j, as is apparent from the echelon form since $p_i(\ell_A)$ has basis the rows of the matrix obtained by deleting the first i columns of the echelon form. Thus the Schubert symbol only depends on the n-plane spanned by the rows of A.

In fact the echelon form itself depends only on the n-plane spanned by the rows. To see this, consider projection onto the n coordinates of $\mathbb{R}^k$ given by the numbers σ_i. This projection is surjective when restricted to the n-plane, hence is also injective, and the n rows of the echelon form are the vectors in the n-plane that project to the standard basis vectors.

As noted above, the numbers σ_j in the Schubert symbol of an n-plane ℓ are the numbers i for which the dimension of the image $p_i(\ell)$ decreases as i goes from 0 to n. Equivalently, these are the i's for which the dimension of the kernel of the restriction $p_i | \ell$ increases. This kernel is just $\ell \cap \mathbb{R}^i$ where $\mathbb{R}^i$ is the kernel of p_i, the subspace of $\mathbb{R}^k$ spanned by the first i standard basis vectors. This gives an alternative definition of the Schubert symbol $\sigma(\ell)$.

Given a Schubert symbol σ one can consider the set $e(\sigma)$ of all n-planes in $\mathbb{R}^k$ having σ as their Schubert symbol. In terms of echelon forms, the various n-planes in $e(\sigma)$ are parametrized by the arbitrary entries in the echelon form. There are $\sigma_i - i$ of these entries in the i^{th} row, for a total of $(\sigma_1 - 1) + \cdots + (\sigma_n - n)$ entries. Thus

$e(\sigma)$ is homeomorphic to a euclidean space of this dimension, or equivalently an open cell.

Proposition 1.17. *The cells $e(\sigma)$ are the cells of a CW structure on $G_n(\mathbb{R}^k)$.*

For example $G_2(\mathbb{R}^4)$ has six cells corresponding to the Schubert symbols $(1,2)$, $(1,3)$, $(1,4)$, $(2,3)$, $(2,4)$, $(3,4)$, and these cells have dimensions 0, 1, 2, 2, 3, 4 respectively. In general the number of cells of $G_n(\mathbb{R}^k)$ is $\binom{k}{n}$, the number of ways of choosing the n distinct numbers $\sigma_i \le k$.

Proof: Our main task will be to find a characteristic map for $e(\sigma)$. This is a map from a closed ball of the same dimension as $e(\sigma)$ into $G_n(\mathbb{R}^k)$ whose restriction to the interior of the ball is a homeomorphism onto $e(\sigma)$. From the echelon forms described above it is not clear how to do this, so we will use a slightly different sort of echelon form. We allow the special 1's to be arbitrary nonzero numbers and we allow the entries below these 1's to be nonzero. Then we impose the condition that the rows be orthonormal and that the last nonzero entry in each row be positive. Let us call this an *orthonormal echelon form*. Once again there is a unique orthonormal echelon form for each n-plane ℓ since if we let ℓ_i denote the subspace of ℓ spanned by the first i rows of the standard echelon form, or in other words $\ell_i = \ell \cap \mathbb{R}^{\sigma_i}$, then there is a unique unit vector in ℓ_i orthogonal to ℓ_{i-1} and having positive ${\sigma_i}^{th}$ coordinate.

The i^{th} row of the orthonormal echelon form then belongs to the hemisphere H_i in the unit sphere $S^{\sigma_i-1} \subset \mathbb{R}^{\sigma_i} \subset \mathbb{R}^k$ consisting of unit vectors with non-negative ${\sigma_i}^{th}$ coordinate. In the Stiefel manifold $V_n(\mathbb{R}^k)$ let $E(\sigma)$ be the subspace of orthonormal frames $(v_1,\cdots,v_n) \in (S^{k-1})^n$ such that $v_i \in H_i$ for each i. We claim that $E(\sigma)$ is homeomorphic to a closed ball. To prove this the main step is to show that the projection $\pi : E(\sigma) \to H_1$, $\pi(v_1,\cdots,v_n) = v_1$, is a trivial fiber bundle. This is equivalent to finding a projection $p : E(\sigma) \to \pi^{-1}(v_0)$ which is a homeomorphism on fibers of π, where $v_0 = (0,\cdots,0,1) \in \mathbb{R}^{\sigma_1} \subset \mathbb{R}^k$, since the map $\pi \times p : E(\sigma) \to H_1 \times \pi^{-1}(v_0)$ is then a continuous bijection of compact Hausdorff spaces, hence a homeomorphism. The map $p : \pi^{-1}(v) \to \pi^{-1}(v_0)$ is obtained by applying the rotation ρ_v of $\mathbb{R}^k$ that takes v to v_0 and fixes the $(k-2)$-dimensional subspace orthogonal to v and v_0. This rotation takes H_i to itself for $i > 1$ since it affects only the first σ_1 coordinates of vectors in $\mathbb{R}^k$. Hence p takes $\pi^{-1}(v)$ onto $\pi^{-1}(v_0)$.

The fiber $\pi^{-1}(v_0)$ can be identified with $E(\sigma')$ for $\sigma' = (\sigma_2 - 1,\cdots,\sigma_n - 1)$. By induction on n this is homeomorphic to a closed ball of dimension $(\sigma_2 - 2) + \cdots + (\sigma_n - n)$, so $E(\sigma)$ is a closed ball of dimension $(\sigma_1 - 1) + \cdots + (\sigma_n - n)$. The boundary of this ball consists of points in $E(\sigma)$ having v_i in ∂H_i for at least one i. This too follows by induction since the rotation ρ_v takes ∂H_i to itself for $i > 1$.

The natural map $E(\sigma) \to G_n$ sending an orthonormal n-tuple to the n-plane it spans takes the interior of the ball $E(\sigma)$ to $e(\sigma)$ bijectively. Since G_n has the quotient topology from V_n, the map $\operatorname{int} E(\sigma) \to e(\sigma)$ is a homeomorphism. The boundary of $E(\sigma)$ maps to cells $e(\sigma')$ of G_n where σ' is obtained from σ by decreasing some σ_i's, so these cells $e(\sigma')$ have lower dimension than $e(\sigma)$.

To see that the maps $E(\sigma) \to G_n(\mathbb{R}^k)$ for the cells $e(\sigma)$ are the characteristic maps for a CW structure on $G_n(\mathbb{R}^k)$ we can argue as follows. Let X^i be the union of the cells $e(\sigma)$ in $G_n(\mathbb{R}^k)$ having dimension at most i. Suppose by induction on i that X^i is a CW complex with these cells. Attaching the $(i+1)$-cells $e(\sigma)$ of X^{i+1} to X^i via the maps $\partial E(\sigma) \to X^i$ produces a CW complex Y and a natural continuous bijection $Y \to X^{i+1}$. Since Y is a finite CW complex it is compact, and X^{i+1} is Hausdorff as a subspace of $G_n(\mathbb{R}^k)$, so the map $Y \to X^{i+1}$ is a homeomorphism and X^{i+1} is a CW complex, finishing the induction. Thus we have a CW structure on $G_n(\mathbb{R}^k)$. $\square$

Since the inclusions $G_n(\mathbb{R}^k) \subset G_n(\mathbb{R}^{k+1})$ for varying k are inclusions of subcomplexes and $G_n(\mathbb{R}^\infty)$ has the weak topology with respect to these subspaces, it follows that we have also a CW structure on $G_n(\mathbb{R}^\infty)$.

Similar constructions work to give CW structures on complex Grassmann manifolds, but here $e(\sigma)$ will be a cell of dimension $(2\sigma_1 - 2) + (2\sigma_2 - 4) + \cdots + (2\sigma_n - 2n)$. The hemisphere H_i is defined to be the subspace of the unit sphere $S^{2\sigma_i - 1}$ in $\mathbb{C}^{\sigma_i}$ consisting of vectors whose ${\sigma_i}^{th}$ coordinate is real and nonnegative, so H_i is a ball of dimension $2\sigma_i - 2$. The transformation $\rho_v \in SU(k)$ is uniquely determined by specifying that it takes v to v_0 and fixes the orthogonal $(k-2)$-dimensional complex subspace, since an element of $U(2)$ of determinant 1 is determined by where it sends one unit vector.

The cell of $G_n(\mathbb{R}^k)$ of largest dimension is $e(\sigma)$ for the Schubert symbol $\sigma = (k-n+1, k-n+2, \cdots, k)$. This cell has dimension $n(k-n)$, so this is the dimension of $G_n(\mathbb{R}^k)$. Near points in this top-dimensional cell $G_n(\mathbb{R}^k)$ is a manifold. But $G_n(\mathbb{R}^k)$ is homogeneous in the sense that given any two points in $G_n(\mathbb{R}^k)$ there is a homeomorphism $G_n(\mathbb{R}^k) \to G_n(\mathbb{R}^k)$ taking one point to the other, namely, the homeomorphism induced by an invertible linear map $\mathbb{R}^k \to \mathbb{R}^k$ taking one n-plane to the other. From this homogeneity it follows that $G_n(\mathbb{R}^k)$ is a manifold near all points. Since it is compact, it is a closed manifold.

There is a natural inclusion $i : G_n \hookrightarrow G_{n+1}$ defined by $i(\ell) = \mathbb{R} \times j(\ell)$ where $j : \mathbb{R}^\infty \to \mathbb{R}^\infty$ is the embedding $j(x_1, x_2, \cdots) = (0, x_1, x_2, \cdots)$. If $\sigma(\ell) = (\sigma_1, \cdots, \sigma_n)$ then $\sigma(i(\ell)) = (1, \sigma_1 + 1, \cdots, \sigma_n + 1)$, so i takes cells of G_n to cells of G_{n+1} of the same dimension, making $i(G_n)$ a subcomplex of G_{n+1}. Identifying G_n with the subcomplex $i(G_n)$, we obtain an increasing sequence of CW complexes $G_1 \subset G_2 \subset \cdots$ whose union $G_\infty = \bigcup_n G_n$ is therefore also a CW complex. Similar remarks apply as well in the complex case.

Appendix: Paracompactness

A Hausdorff space X is *paracompact* if for each open cover $\{U_\alpha\}$ of X there is a partition of unity $\{\varphi_\beta\}$ subordinate to the cover. This means that the φ_β's are maps $X \to I$ such that each φ_β has support (the closure of the set where $\varphi_\beta \neq 0$) contained in some U_α, each $x \in X$ has a neighborhood in which only finitely many φ_β's are nonzero, and $\sum_\beta \varphi_\beta = 1$. An equivalent definition which is often given is that X is Hausdorff and every open cover of X has a locally finite open refinement. The first definition clearly implies the second by taking the cover $\{\varphi_\beta^{-1}(0,1]\}$. For the converse, see [Dugundji] or [Lundell-Weingram]. It is the former definition which is most useful in algebraic topology, and the fact that the two definitions are equivalent is rarely if ever needed. So we shall use the first definition.

A paracompact space X is normal, for let A_1 and A_2 be disjoint closed sets in X, and let $\{\varphi_\beta\}$ be a partition of unity subordinate to the cover $\{X - A_1, X - A_2\}$. Let φ_i be the sum of the φ_β's which are nonzero at some point of A_i. Then $\varphi_i(A_i) = 1$, and $\varphi_1 + \varphi_2 \leq 1$ since no φ_β can be a summand of both φ_1 and φ_2. Hence $\varphi_1^{-1}(1/2, 1]$ and $\varphi_2^{-1}(1/2, 1]$ are disjoint open sets containing A_1 and A_2, respectively.

Most of the spaces one meets in algebraic topology are paracompact, including:

(1) compact Hausdorff spaces

(2) unions of increasing sequences $X_1 \subset X_2 \subset \cdots$ of compact Hausdorff spaces X_i, with the weak or direct limit topology (a set is open iff it intersects each X_i in an open set)

(3) CW complexes

(4) metric spaces

Note that (2) includes (3) for CW complexes with countably many cells, since such a CW complex can be expressed as an increasing union of finite subcomplexes. Using (1) and (2), it can be shown that many manifolds are paracompact, for example $\mathbb{R}^n$.

The next three propositions verify that the spaces in (1), (2), and (3) are paracompact.

Proposition 1.18. *A compact Hausdorff space X is paracompact.*

Proof: Let $\{U_\alpha\}$ be an open cover of X. Since X is normal, each $x \in X$ has an open neighborhood V_x with closure contained in some U_α. By Urysohn's lemma there is a map $\varphi_x : X \to I$ with $\varphi_x(x) = 1$ and $\varphi_x(X - V_x) = 0$. The open cover $\{\varphi_x^{-1}(0,1]\}$ of X contains a finite subcover, and we relabel the corresponding φ_x's as φ_β's. Then $\sum_\beta \varphi_\beta(x) > 0$ for all x, and we obtain the desired partition of unity subordinate to $\{U_\alpha\}$ by normalizing each φ_β by dividing it by $\sum_\beta \varphi_\beta$. □

Proposition 1.19. *If X is the direct limit of an increasing sequence $X_1 \subset X_2 \subset \cdots$ of compact Hausdorff spaces X_i, then X is paracompact.*

Proof: A preliminary observation is that X is normal. To show this, it suffices to find a map $f:X\to I$ with $f(A)=0$ and $f(B)=1$ for any two disjoint closed sets A and B. Such an f can be constructed inductively over the X_i's, using normality of the X_i's. For the induction step one has f defined on the closed set $X_i\cup(A\cap X_{i+1})\cup(B\cap X_{i+1})$ and one extends over X_{i+1} by Tietze's theorem.

To prove that X is paracompact, let an open cover $\{U_\alpha\}$ be given. Since X_i is compact Hausdorff, there is a finite partition of unity $\{\varphi_{ij}\}$ on X_i subordinate to $\{U_\alpha\cap X_i\}$. Using normality of X, extend each φ_{ij} to a map $\varphi_{ij}:X\to I$ with support in the same U_α. Let $\sigma_i=\sum_j\varphi_{ij}$. This sum is 1 on X_i, so if we normalize each φ_{ij} by dividing it by $\max\{1/2,\sigma_i\}$, we get new maps φ_{ij} with $\sigma_i=1$ in a neighborhood V_i of X_i. Let $\psi_{ij}=\max\{0,\varphi_{ij}-\sum_{k<i}\sigma_k\}$. Since $0\le\psi_{ij}\le\varphi_{ij}$, the collection $\{\psi_{ij}\}$ is subordinate to $\{U_\alpha\}$. In V_i all ψ_{kj}'s with $k>i$ are zero, so each point of X has a neighborhood in which only finitely many ψ_{ij}'s are nonzero. For each $x\in X$ there is a ψ_{ij} with $\psi_{ij}(x)>0$, since if $\varphi_{ij}(x)>0$ and i is minimal with respect to this condition, then $\psi_{ij}(x)=\varphi_{ij}(x)$. Thus when we normalize the collection $\{\psi_{ij}\}$ by dividing by $\sum_{i,j}\psi_{ij}$ we obtain a partition of unity on X subordinate to $\{U_\alpha\}$. □

Proposition 1.20. *Every CW complex is paracompact.*

Proof: Given an open cover $\{U_\alpha\}$ of a CW complex X, suppose inductively that we have a partition of unity $\{\varphi_\beta\}$ on X^n subordinate to the cover $\{U_\alpha\cap X^n\}$. For a cell e_γ^{n+1} with characteristic map $\Phi_\gamma:D^{n+1}\to X$, $\{\varphi_\beta\Phi_\gamma\}$ is a partition of unity on $S^n=\partial D^{n+1}$. Since S^n is compact, only finitely many of these compositions $\varphi_\beta\Phi_\gamma$ can be nonzero, for fixed γ. We extend these functions $\varphi_\beta\Phi_\gamma$ over D^{n+1} by the formula $\rho_\varepsilon(r)\varphi_\beta\Phi_\gamma(x)$ using 'spherical coordinates' $(r,x)\in I\times S^n$ on D^{n+1}, where $\rho_\varepsilon:I\to I$ is 0 on $[0,1-\varepsilon]$ and 1 on $[1-\varepsilon/2,1]$. If $\varepsilon=\varepsilon_\gamma$ is chosen small enough, these extended functions $\rho_\varepsilon\varphi_\beta\Phi_\gamma$ will be subordinate to the cover $\{\Phi_\gamma^{-1}(U_\alpha)\}$. Let $\{\psi_{\gamma j}\}$ be a finite partition of unity on D^{n+1} subordinate to $\{\Phi_\gamma^{-1}(U_\alpha)\}$. Then $\{\rho_\varepsilon\varphi_\beta\Phi_\gamma,(1-\rho_\varepsilon)\psi_{\gamma j}\}$ is a partition of unity on D^{n+1} subordinate to $\{\Phi_\gamma^{-1}(U_\alpha)\}$. This partition of unity extends the partition of unity $\{\varphi_\beta\Phi_\gamma\}$ on S^n and induces an extension of $\{\varphi_\beta\}$ to a partition of unity defined on $X^n\cup e_\gamma^{n+1}$ and subordinate to $\{U_\alpha\}$. Doing this for all $(n+1)$-cells e_γ^{n+1} gives a partition of unity on X^{n+1}. The local finiteness condition continues to hold since near a point in X^n only the extensions of the φ_β's in the original partition of unity on X^n are nonzero, while in a cell e_γ^{n+1} the only other functions that can be nonzero are the ones coming from $\psi_{\gamma j}$'s. After we make such extensions for all n, we obtain a partition of unity defined on all of X and subordinate to $\{U_\alpha\}$. □

Here is a technical fact about paracompact spaces that is occasionally useful:

Lemma 1.21. *Given an open cover $\{U_\alpha\}$ of the paracompact space X, there is a countable open cover $\{V_k\}$ such that each V_k is a disjoint union of open sets each contained in some U_α, and there is a partition of unity $\{\varphi_k\}$ with φ_k supported in V_k.*

Proof: Let $\{\varphi_\beta\}$ be a partition of unity subordinate to $\{U_\alpha\}$. For each finite set S of functions φ_β let V_S be the subset of X where all the φ_β's in S are strictly greater than all the φ_β's not in S. Since only finitely many φ_β's are nonzero near any $x \in X$, V_S is defined by finitely many inequalities among φ_β's near x, so V_S is open. Also, V_S is contained in some U_α, namely, any U_α containing the support of any $\varphi_\beta \in S$, since $\varphi_\beta \in S$ implies $\varphi_\beta > 0$ on V_S. Let V_k be the union of all the open sets V_S such that S has k elements. This is clearly a disjoint union. The collection $\{V_k\}$ is a cover of X since if $x \in X$ then $x \in V_S$ where $S = \{\,\varphi_\beta \mid \varphi_\beta(x) > 0\,\}$.

For the second statement, let $\{\varphi_\gamma\}$ be a partition of unity subordinate to the cover $\{V_k\}$, and let φ_k be the sum of those φ_γ's supported in V_k but not in V_j for $j < k$. □

Exercises

1. Show that the projection $V_n(R^k) \to G_n(\mathbb{R}^k)$ is a fiber bundle with fiber $O(n)$ by showing that it is the orthonormal n-frame bundle associated to the vector bundle $E_n(\mathbb{R}^k) \to G_n(\mathbb{R}^k)$.