The Existence of Classifying Spaces for Principal G-Bundles

Deval Mehta

12/17/2020

Preface

The following has been submitted as a final project for MATH GU4051: Topology in the Fall 2020 Term at Columbia University, under the instruction of Dr. S. Michael Miller. The objective herein is to adequately develop the necessary notions to understand that each principal G-bundle has a classifying space, with the assumption that a prospective reader will have a strong foundation in linear algebra, and will have equivalent knowledge to that of the topics covered in MATH GU4041: Introduction to Modern Algebra I, MATH GU4051: Topology, and MATH GU4061: Introduction to Modern Analysis I.

I acknowledge that most of the proofs provided herein are not entirely original, but most expand upon arguments found in the sources referenced alongside each statement. In addition, some proofs are either entirely original or clarify an argument in the source material. If a proof is provided for a given statement in a cited source, it will be found either immediately before or after the relevant statement in that source. All figures and diagrams were made by me using MS PowerPoint or LaTeX.

Contents

I Introduction

In an interview with *Quanta Magazine* in 2017 [5], Levi L. Conant Prize recipient Dr. Alex Wright remarked, "Classification is like a definitive study for all time, like writing the final book so nobody has to write another book." Indeed, the history of mathematics and the inclination of mathematicians, spanning as far back as ancient Greece, is littered with attempts (both successful and not) to classify various types of objects. These classifications, when complete, can yield powerful results in their own right, propelling a field forward and leading to the construction of new objects which may have previously remained unstudied. Universally, however, they provide more insight into the structure of certain classes of objects and explain why apparently disparate objects may seem strikingly similar. Topology, of course, is no exception.

Topologists of various flavors have attempted to classify topological spaces for as long as we have had the concept of a homeomorphism, though a classification "up to homeomorphism" proved futile. Eventually, the field moved to other attempts, from classifying topological spaces by invariants, to identifying their homotopy groups. In 1935, Hassler Whitney attempted a new type of classification: that of which he then called "sphere spaces" by characteristic classes, an idea he defined in parallel with Eduard Stiefel. By 1940, Whitney referred to these objects as "sphere bundles" and we refer to these today as *fiber bundles*. Our objective is not quite as deep as these results. Instead, we seek to demonstrate that it is possible to classify a special type of fiber bundle, called a *principal G-bundle*, by their *classifying spaces*.

Theorem I.1. *For any topological space X, there exists a bijection $\Phi : [X, BG] \to \mathcal{P}_G(X)$ where G is a topological group and BG is a classifying space.*

Rather than jumping head-first into the formalism, we'll first develop an analogous result for a simpler object, the *vector bundle*, demonstrate how these relate to principal G-bundles, and finally prove Theorem I.1 (or at least an approximate case), before demonstrating the power of this formalism by considering a special case.

II A Quick Taste: Classifying Complex Vector Bundles of Spheres

2.1 Vector Bundles

We begin with something more tangible that will help us develop some intuition, particularly because of its connection to linear algebra and familiar surfaces in early examples. Let $\mathbb{F}$ be either $\mathbb{R}$ or $\mathbb{C}$ (or, in general, a field, for those readers who are familiar with the notion).

Definition II.1. Let E and B be topological spaces. A **vector bundle of dimension** n (also referred to as a **vector bundle**) is a triplet (p, E, B) consisting of E, B, and a map $p : E \to B$, such that

- for all $b \in B$, $p^{-1}(b)$ is a vector space
- there exists an open cover $\mathcal{U}$ of B where, for each $U_\alpha \in \mathcal{U}$, there is a homeomorphism (called a **local trivialization**) $h_\alpha : p^{-1}(U_\alpha) \to U_\alpha \times \mathbb{F}^n$ which maps $p^{-1}(b)$ to $\{b\} \times \mathbb{F}^n$.

We say that B is the **base space**, E is the **total space**, and each $p^{-1}(b)$ is a **fiber**. If $\mathbb{F} = \mathbb{R}$, we call (p, E, B) a **real vector bundle** and if $\mathbb{F} = \mathbb{C}$, (p, E, B) is a **complex vector bundle**. We will at times refer to the map p as a vector bundle for simplicity.

Let's examine this definition with some examples.

- If we simply let $E \cong B \times \mathbb{F}^n$ and $p : E \to B$ the projection map onto the first factor, then (p, E, B) is the **product bundle** or **trivial bundle** over B. Since this is globally the product $B \times \mathbb{F}^n$, it is trivially the local product as well and thus satisfies the local trivialization criterion.
- Let $I = [0, 1]$. If E is the quotient space given by $I \times \mathbb{R}/((0, t) \sim (1, -t))$ with base space $B = S^1$ and $p : E \to B$ the induced map given by $E \to I \to B$, then (p, E, B) is the **line bundle**. In addition, we note that E is homeomorphic to the Möbius band without its boundary circle, so the line bundle is also sometimes called **the Möbius bundle**.

- Consider the space
$$TS^2 = \{(x, v) \in S^2 \times \mathbb{R}^3\},$$
the disjoint union of tangent planes P_x to S^2 at each point $x \in S^2$, which we call the **tangent space** of S^2, and the map $p : TS^2 \to S^2$ defined by $p(x, v) = x$. We claim that (p, TS^2, S^2) is a vector bundle. First, we observe that, $p^{-1}(x) = P_x$. To satisfy the local trivialization condition, we choose an arbitrary point $x \in S^2$ and let $U_x \subset S^2$ be the open hemisphere containing x, bounded by the plane through the origin and orthogonal to x. Then, we define
$$h_x : p^{-1}(U_x) \to U_x \times P_x \cong U_x \times \mathbb{R}^2$$
by $h_x(y, v) = (y, \pi_x(v))$, where π_x describes orthogonal projection onto P_x. Since π_x restricts to an isomorphism of $p^{-1}(y)$ onto P_x for each $y \in U_x$, h_x is a local trivialization. We call a vector bundle for which the fibers are tangent spaces a **tangent bundle**.

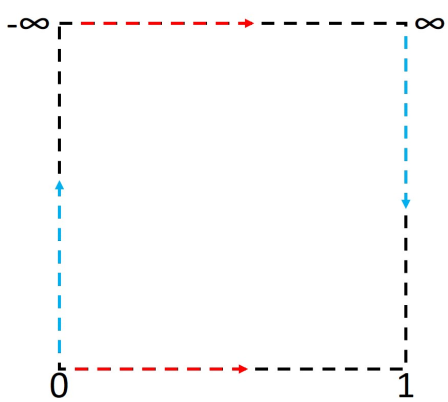

Figure 1: Mobius bundle "flattened" into a square

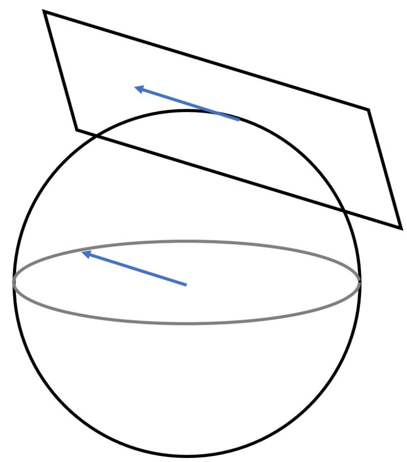

Figure 2: Sphere with a Tangent Plane, v_x, and $\tau(v_x)$

The involved construction of a local trivialization for the tangent bundle of S^2 may seem overly intricate, but this is done with purpose. The fact is that it is not always so easy to trivialize TS^n, but the construction we have provided *does* extend to any n. We can also generalize the line bundle to the **canonical line bundle** by taking $\mathbb{R}P^n$ to be our base space and
$$E = \{(\ell, v) \in \mathbb{R}P^n \times \mathbb{R}^{n-1} | v \perp \ell\}$$
with $p : E \to \mathbb{R}P^n$ defined by $p(\ell, v) = \ell$. The base space here may be further generalized to the **Grassmann manifold**, which we will encounter later. With all of these examples of vector bundles under our belts, a natural question to ask might be, "When are two vectors bundles, in some sense, the same?"

Definition II.2. If (p_1, E_1, B) and (p_2, E_2, B) are two vector bundles over the same base space B, then a homeomorphism $h : E_1 \to E_2$ which maps each fiber $p_1^{-1}(b)$ to $p_2^{-1}(b)$ by a vector space isomorphism is an **isomorphism of vector bundles**.

Before we proceed to some more involved theory, there is one more critical notion we must define in relation to vector bundles.

Definition II.3. A **section** of a vector bundle (p, E, B) is a map $s : B \to E$ sending each $b \in B$ to a vector in the fiber $p^{-1}(b)$.

Notice that each vector bundle admits at least a canonical section: the **zero section**, whose value is the zero vector on each fiber. We may identify the zero section with its image, which is the subspace of E which projects homeomorphically onto B by p. Sections are often useful in the study of vector bundles and their generalizations, as we shall see. In the case of vector bundles, we can use the zero section to check whether two vector bundles are nonisomorphic quickly. This is not necessarily a perfect test, but we can say definitively that if the complements of the zero sections of two vector bundles do not match, then they are not isomorphic, since we can see that any vector bundle isomorphism must carry the zero section of one to the zero section of the other. Though we do not make use of this here, it is valuable enough a piece of information that we would be remiss to not include it.

2.2 The General Linear Group

Before we continue discussing vector bundles and their related constructions, we take a quick detour to groups to study one group in particular which is ever-present in our discussion.

Definition II.4. The set of all invertible $n \times n$ matrices whose entries are elements of $\mathbb{F}$ becomes a group under the operation of matrix multiplication. We call this group the **general linear group of rank n over $\mathbb{F}$**, denoted $GL(n, \mathbb{F})$.

Not only is $GL(n, \mathbb{F})$ a group, but it is particularly relevant to us as a topological group.

Proposition II.1. (Theorem 1.16 in [1].) *$GL(n, \mathbb{F})$ is a topological group.*

Proof. To see that $GL(n, \mathbb{F})$ is a topological group, we must first demonstrate that it is a space, then that the multiplication and inversion maps are continuous, and the existence of the identity. Notice that $GL(n, \mathbb{F})$ is a subset of $M_n(\mathbb{F})$, the set of all $n \times n$ matrices over $\mathbb{F}$, so by endowing $M_n(\mathbb{F})$ with a topology, we can describe the topology on $GL(n, \mathbb{F})$. Recall that we define the 2-norm of a vector $\mathbf{x} \in \mathbb{F}^n$ by

$$|\mathbf{x}| = \sqrt{\sum_{i=1}^{n} |x_i|}.$$

From this, we can adequately define the notion of a norm for a matrix $\mathbf{M} \in M_n(\mathbb{F})$. Since the product $\mathbf{Mx}$ produces a vector in $\mathbb{F}^n$, we can find its 2-norm by the above formula. We then define a norm on $M_n(\mathbb{F})$ as follows. Consider the set $\mathcal{N}_{\mathbf{M}} = \{\frac{|\mathbf{Mx}|}{\mathbf{x}} | \mathbf{x} \in \mathbb{F}^n \backslash \{\mathbf{0}\}\}$. As a set, $\mathcal{N}_{\mathbf{M}}$ is identical to the set

$$\mathcal{N}_{\mathbf{M}}^1 := \{|\mathbf{Mx}| \, | \mathbf{x} \in \mathbb{F}^n, |\mathbf{x}| = 1\}$$

since $\frac{|\mathbf{Mx}|}{|\mathbf{x}|} = |\mathbf{Mx}'|$, where $\mathbf{x}' = \frac{1}{|\mathbf{x}|}\mathbf{x}$ is a vector of 2-norm 1 when $\mathbf{x} \neq \mathbf{0}$. Since the set $\{\mathbf{x} \in \mathbb{F}^n | |\mathbf{1}| = 1\}$ is closed and bounded, it is compact by the Heine-Borel theorem, thus the $\mathcal{N}_{\mathbf{M}}^1 = \mathcal{N}_{\mathbf{M}}$ is compact as the image of the continuous map which carries $\mathbf{x} \in \mathbb{F}^n$ to $|\mathbf{Mx}| \in \mathbb{R}$. This map is then bounded and attains its supremum, thus we define the **operator norm of M** by $\|\mathbf{M}\| = \max \mathcal{N}_{\mathbf{M}}$.

Checking that $\| \cdot \|$ satisfies the usual properties of a norm is an easy enough task, but rather laborious. We defer to the proof which begins on page 208 of [11] for this. The operator norm induces a metric on $M_n(\mathbf{F})$, given by $d(\mathbf{M}, \mathbf{N}) = \|\mathbf{M} - \mathbf{N}\|$, which is demonstrated in the same proof as above in [11]. With this, we have that $M_n(\mathbb{F})$ is a metric space and, thus, a topological space endowed with the metric topology. Since $GL(n, \mathbb{F})$ is a subset of $M_n(\mathbb{F})$, it is also a metric space when endowed with the subspace topology.

Now that we know that $GL(n, \mathbb{F})$ is a space, we confirm the group axioms and that the multiplication and inversion maps are continuous. The identity, rather obviously, is the identity matrix $\mathbf{1}_n$, since

$$\forall \mathbf{M} \in GL(n, \mathbb{F}),\ \mathbf{1}_n\mathbf{M} = \mathbf{M}\mathbf{1}_n = \mathbf{M}.$$

Multiplication on $GL(n, \mathbb{F})$ is a map $m : GL(n, \mathbb{F}) \times GL(n, \mathbb{F}) \to GL(n, \mathbb{F})$ defined by $m(\mathbf{X}, \mathbf{Y}) = \mathbf{XY}$. Likewise, inversion is a map $\mathrm{inv} : GL(n, \mathbb{F}) \to GL(n, \mathbb{F})$ defined by $\mathrm{inv}(\mathbf{X}) = \mathbf{X}^{-1}$. Matrix multiplication may be expressed in terms of polynomials of the entries in the involved matrices, which are continuous maps, and inversion may be expressed in terms of said polynomials and the continuous map $\det : GL(n, \mathbb{F}) \to \mathbb{F}$ (see page 9 of [1] for a proof of continuity of det), thus multiplication and inversion are continuous as compositions of continuous maps. Therefore, we conclude that $GL(n, \mathbb{F})$ is a topological group. $\square$

In particular, since $GL(n, \mathbb{F})$ is a metric space, it has all of the nice properties we love about metric spaces, namely that it is Hausdorff and normal. When we admit a choice of $\mathbb{F}$, we find some more properties of interest. Most importantly, we have the following:

Proposition II.2. (Theorem 3.1 in [2].) *$GL(n, \mathbb{C})$ is path-connected.*

Proof. Take two matrices $\mathbf{X}, \mathbf{Y} \in GL(n, \mathbb{C})$. We seek a continuous map $\gamma : [0, 1] \to GL(n, \mathbb{C})$ for which $\gamma(0) = X$ and $\gamma(1) = Y$. Recall the three elementary row operations: scaling a row, row exchange, and row replacement. Each of these is expressed as a path: if two matrices $\mathbf{X}, \mathbf{Y}$ are related by a row operation, we can construct a path between them by $\gamma(t) = (1 - t)\mathbf{X} + t\mathbf{Y}$. These are clearly continuous since row operations preserve the determinant. Now note that every matrix in $GL(n, \mathbb{C})$ is diagonalizable by way of row operations; indeed, the conjugation of a matrix by the eigenvector matrices amounts to a series of row operations. We see that this is true since the characteristic polynomial of each such matrix has no repeated roots. The set of all diagonal matrices in $GL(n, \mathbb{C})$ is path-connected since we can connect any diagonal matrix to the identity by a path connecting each eigenvalue to 1, explicitly avoiding 0. Such a path can be constructed by first mapping each eigenvalue λ onto the unit circle through $\frac{\lambda}{(1-t)+|\lambda|t}$ and then rotating it along the path $e^{i\theta(1-t)}$ where $e^{i\theta} = \frac{\lambda}{|\lambda|}$. The concatenation of all of these paths results in a path from each element of $GL(n, \mathbb{C})$ to the identity, thus we can connect any two such matrices by a path through the identity. □

It might not be immediately apparent how all of this is relevant, but we encourage the reader to proceed with these ideas in mind. Beyond the scope of our study, $GL(n, \mathbb{F})$ and its subgroups appear throughout much of modern mathematics, particularly in topology and geometry when symmetries are concerned. We point the reader to the book by Baker [1] for a more detailed introduction to $GL(n, \mathbb{F})$ and other **matrix groups**, should it pique their interest.

2.3 Clutching Functions

With our brief digression out of the way, we return to vector bundles. For our first major intermediate result, we return to the roots of this study, taking Whitney's original name (and concept) of a sphere bundle somewhat literally. We concern ourselves with spheres in particular because they are nice, familiar spaces. Perhaps by understanding a result analogous to our desired result over spheres, we might glean some information on how to approach the general theorem. To work with spheres in mind, we require some method of constructing vector bundles over spheres. The clutching function comes to our rescue here.

Definition II.5. Write $S^k = D^k_+ \cup D^k_-$ so that $D^k_+ \cap D^k_- = S^{k-1}$. A map $f : S^{k-1} \to GL(n, \mathbb{C})$ is called a **clutching function** or **clutching construction** if for the space

$$E_f = (D^k_+ \times \mathbb{C}^n \sqcup D^k_- \times \mathbb{C}^n)/((x, v) \sim (x, f(x)(v))),$$

where the former pair lies in $\partial D^k \times \mathbb{C}^n$ and the latter in $\partial D^k_+ \times \mathbb{C}^n$, $p : E_f \to S^k$ is an n-dimensional complex vector bundle.

Observe that the same definition applies for n-dimensional real vector bundles if we take $\mathbb{R}$ in the place of $\mathbb{C}$. It is easy for the general idea to be obfuscated by details here. Let's quickly consider a couple of examples to see if we can walk our way through this construction.

- We return to our good friend, S^2 to see if we can use the clutching function to reconstruct the bundle (p, TS^2, S^2). Let's take a tangent vector at the north pole and use this to construct a vector field v_+ on D^2_+. We'll move this vector along each meridian circle in a manner in which it maintains a constant angle with the given meridian. We then obtain v_- on D^2_- by reflecting v_+ across the equatorial plane. We can now obtain new vector fields w_+, w_- by rotating v_+, v_- counterclockwise by $\pi/2$ when viewed from outside the sphere. Both $v_\pm$ and $w_\pm$ provide trivializations of TS^2 over $D^2_\pm$, thus we identify the two halfs of TS^2 with $D^2_\pm \times \mathbb{R}^2$. We then recover TS^2 as the disjoint union $D^2_+ \times \mathbb{R}^2 \sqcup_{S^1 \times \mathbb{R}^2} D^2_- \times \mathbb{R}^2$ by a clutching function f which reads off the coordinates of the v_-, w_- in the v_+, w_+ coordinate system, thus rotating (v_+, w_+) to (v_-, w_-). As we traverse the equator S^1 counterclockwise when viewed from above from a point where the two trivializations agree, the angle of rotation increases from 0 to 4π, so parameterizing S^1 by the angle θ, we find that $f(\theta)$ is a rotation by 2θ.

- We consider now the canonical complex line bundle over $\mathbb{C}P^1 = S^2$. Observe that $\mathbb{C}P^1$ is identified with the quotient of $\mathbb{C}^2\backslash\{0\}$ under the equivalence relation $(z_0, z_1) \sim \lambda(z_0, z_1)$ with $\lambda \in \mathbb{C}\backslash\{0\}$, in a manner analogous to what we know of the construction of $\mathbb{R}P^2$. We thus write the points of $\mathbb{C}P^1$ as ratios $z = z_0/z_1 \in \mathbb{C} \cup \{\infty\} = S^2$. In doing so, the points in the disc D^2_0 inside the unit circle $S^1 \subset \mathbb{C}$ may be expressed uniquely in the form $[z_0/z_1, 1] = [z, 1]$ such that $|z| \leqslant 1$. Furthermore, the points in the disc D^2_∞ outside S^1 may be expressed uniquely in the form $[1, z_1/z_0] = [1, 1/z]$ such that $|z| \geqslant 1$. Sections of the canonical line bundle over D^2_0 and D^2_∞ are then given by maps $[z, 1] \mapsto (z, 1)$ and $[1, 1/z] \mapsto (1, 1/z)$, respectively. These sections determine trivializations of the canonical line bundle over the aforementioned discs and, over their common boundary S^1, we pass from one to the other by multiplying or dividing by a factor of z. Taking D^2_∞ to be D^2_+ and D^2_0 to be D^2_-, we see that the canonical line bundle has the clutching function $f : S^1 \to GL(1, \mathbb{C})$ defined by $f(z) = z$. Had we chosen the opposite identification, we would have that $f(z) = 1/z$.

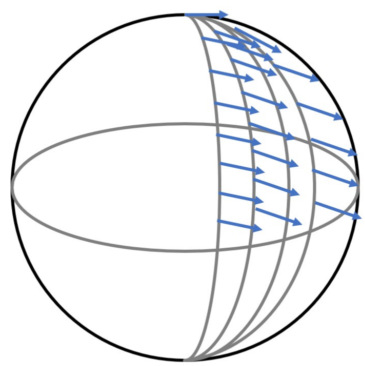

Figure 3: Vector Field v_+ on S^2

Suppose we have two clutching functions f, g for real vector bundles which were homotopic, that is, suppose there exists a map $F : S^{k-1} \times I \to GL(n, \mathbb{R})$ which is f at time 0 and g at time 1. We can then employ the same sort of clutching construction to produce a vector bundle $E_F \to S^k \times I$ which restricts to E_f over $S^k \times \{0\}$ and to E_g over $S^k \times \{1\}$. It can then be shown (and Hatcher does on page 20 of [7]), that $E_f \cong E_g$. This leads us to the next result.

2.4 Proof of the Result

We have now assembled the tools that we need in order to consider our first intermediary theorem. We can now classify, in some sense, complex vector bundles over spheres, an analogous, but more well-behaved result, than the real case described briefly above.

Theorem II.1. (Proposition 1.11 in [7].) *Denote by* $[S^{k-1}, GL(n, \mathbb{C})]$ *the* ***set of homotopy classes of clutching functions*** $S^{k-1} \to GL(n, \mathbb{C})$ *and by* $\mathrm{Vect}^n_{\mathbb{C}}(S^k)$ *the* ***set of isomorphism classes of*** n***-dimensional complex vector bundles over*** S^k*. The map* $\Phi : [S^{k-1}, GL(n, \mathbb{C})] \to \mathrm{Vect}^n_{\mathbb{C}}(S^k)$ *is a bijection.*

Proof. Showing that Φ is bijective amounts to constructing a well-defined inverse function Ψ. Consider the restrictions of an n-dimensional vector bundle $p : E \to S^k$ to the spaces E_+ and E_- over the hemispheres D^k_+ and D^k_- respectively. These restrictions are then trivial since $D^k_\pm$ are contractible. We'll accept this as fact here and refer the reader to pages 20 - 21 of [7]. We have yet to develop the necessary tools to prove this statement, but the technology will be made accessible by the end of our discussion. That said, we will not prove it for the sake of brevity. We now choose local trivializations $h_\pm : E_\pm \to D^k_\pm \times \mathbb{C}^n$, so that we have a map $h_+ h_-^{-1} : S^{k-1} \to GL_n(\mathbb{C})$ whose homotopy class is $\Psi(E) \in [S^{k-1}, GL(n, \mathbb{C})]$. Observe that any two choices of $h_\pm$ differ by a map $D^k_\pm \to GL(n, \mathbb{C})$. Since $D^k_\pm$ are contractible, such maps are homotopic to a constant map when we compose with a contraction on $D^k_\pm$. From this and the path-connectivity of $GL(n, \mathbb{C})$, it follows that h_+ and h_- are unique up to homotopy, thus $\Psi(E)$ is well-defined and it is evident by construction that Ψ is the inverse of Φ, thus Φ is a bijection. □

This provides immediately only one hint: we may need to extend the notion of a clutching function to abstractions of vector bundles in order to address the general problem. For now, we'll consider a different, but important, consequence.

2.5 Not Every Subgroup: Real Vector Bundles of Spheres

Unlike our discussion to this point, we cannot recover the same result by simply working with $\mathbb{R}$ instead of $\mathbb{C}$. A crucial part of the proof here is that $GL(n,\mathbb{C})$ is path-connected; this is not true for $GL(n,\mathbb{R})$. We can, however, identify a similar result once we introduce an additional condition. Let $GL^+(n,\mathbb{R})$ be the subgroup of $GL(n,\mathbb{R})$ of matrices with positive determinant.

Lemma II.1. (Page 24 of [7].) *$GL^+(n,\mathbb{R})$ is path-connected.*

Proof. The construction of a path between an arbitrary matrix in $GL(n,\mathbb{R})$ and a diagonal matrix is identical to the complex case, but there is no path from diagonal matrices of negative determinant to the identity. We can, however, move along a path from each diagonal matrix to either the identity or $-\mathbf{1}_n$, thus any element of $GL(n,\mathbb{R})$ is path-connected to any other element with a determinant of the same sign. to go continuously from a matrix of negative determinant to the identity or from a matrix of positive determinant to $-\mathbf{1}_n$, we must cross a matrix of determinant 0, thus leaving $GL(n,\mathbb{R})$. Therefore, $GL(n,\mathbb{R})$ is partitioned into two sets: $GL^+(n,\mathbb{R})$ and $GL^-(n,\mathbb{R})$. Notice that $GL^+(n,\mathbb{R})$ is a subgroup of index 2 and $GL^-(n,\mathbb{R})$ is its coset. We then have a homeomorphism between $GL^+(n,\mathbb{R})$ and $GL^-(n,\mathbb{R})$ given by $\mathbf{X} \mapsto \mathbf{YX}$, where $\mathbf{Y}$ is a matrix of determinant -1, with an inverse map $\mathbf{X} \mapsto \mathbf{Y}^{-1}\mathbf{X}$. As such, these are the two path components of $GL(n,\mathbb{R})$ and are thus each is path connected. □

With that in mind, we have essentially shown the following.

Proposition II.3. (Proposition 1.14 in [7].) *The map $\Phi : [S^{k-1}, GL(n,\mathbb{C})] \to \mathrm{Vect}^n_+(S^k)$ is a bijection, where $[S^{k-1}, GL(n,\mathbb{C})]$ is the set of homotopy classes of maps $S^{k-1} \to GL(n,\mathbb{C})$ and $\mathrm{Vect}^n_+(S^k)$ is the set of isomorphism classes of n-dimensional **oriented** real vector bundles over S^k.*

Definition II.6. We say that a real vector bundle (p, E, B) is **oriented** if there is a map, called the **orientation**, which assigns an orientation (in the sense of a vector space) to each fiber $p^{-1}(b)$ such that there is a local trivialization $h : p^{-1}(U) \to U \times \mathbb{R}^n$ around each point $b \in B$, which carries the orientation of the fibers in $p^{-1}(U)$ to the standard orientation of $\mathbb{R}^n$ in the fibers of $U \times \mathbb{R}^n$.

Proof. The clutching construction gives us a map $\Phi : [S^{k-1}, GL^+(n,\mathbb{R})] \to \mathrm{Vect}^n_+(S^k)$. The proof is then identical to the complex case. □

III Building Up: Complex General Linear Principal Bundles

3.1 Associated Bundles

Definition III.1. Having established a "base case" of sorts, we may now begin to generalize. To move from the result we have to result we want, we need to do a few different things: we must somehow find a way to extend vector bundles to principal G-bundles, generalize the sphere to a general space or a weaker condition that we can use to construct a general space, and generalize clutching functions. We'll begin with the first task.

A **fiber bundle** is a triplet (p, E, B) consisting of a base space, B, a total space E, and a map $p : E \to B$ such that there exists an open cover $\mathcal{U}$ of B where, for each $U_\alpha \in \mathcal{U}$, there is a homeomorphism $h_\alpha : p^{-1}(U_\alpha) \to U_\alpha \times F$ which maps $p^{-1}(b)$ to $\{b\} \times F$ for F a fiber.

Notice that the definition of a fiber bundle matches exactly that of a vector bundle, expect in that there are no references to a vector space structure on the fibers. In fact, the fibers in a fiber bundle are left completely arbitrary; they can be any kind of space! Specifying what type of space our fibers are will thus inform the structure of our particular choice of fiber bundle.

We can also now extend the structures we had associated to vector bundles with fiber bundles. In particular, we can discuss sections of fiber bundles:

Definition III.2. A **section** of a fiber bundle (p, E, B) is a map $s : B \to E$ such that $p(s(b)) = b$ for all $b \in B$.

If we consider the notion of a trivial fiber bundle in the same way as we did a trivial vector bundle, we can make progress toward the third of our stated goals: generalizing the clutching function. Replacing S^k with a space B, $GL(n, \mathbb{F})$ with a topological group G, $\mathbb{C}^n$ with fiber $p^{-1}(b)$ of a trivialized fiber bundle, and $D^k_\pm$ by closed sets X, Y such that $X \cup Y = B$, we can define a clutching function on $X \cap Y$ which yields a vector bundle on B.

Finally, we can construct many different fiber bundles, which we'll call **associated bundles**, from vector bundles.

- If the total space E of a vector bundle has an inner product, then we can consider the subspace $S(E)$ of all unit spheres in the fibers. The projection $S(E) \to B$ is then a fiber bundle with sphere fibers, with local trivializations being isometries in each fiber, restricted to $S(E)$ (Hatcher [7] demonstrates on page 12 that this is viable). We call this particular fiber bundle a **sphere bundle**. There is also an analogous notion of the **disc bundle** $D(E) \to B$.
- There is, associated to each vector bundle, a **projective bundle** $P(E) \to B$, the space of all lines through the origin in all fibers of the total space E. Taking inspiration from the definition of $\mathbb{R}P^n$, we assign a topology to the total space $P(E)$ by the quotient of the sphere bundle $S(E)$ by factoring out scalar multiplication in each fiber. If E is locally $U \times \mathbb{R}^n$ for the sphere bundle over an open neighborhood $U \subset B$, then this quotient is $U \times \mathbb{R}P^{n-1}$, thus $P(E)$ is the total space of a fiber bundle over B with fibers $\mathbb{R}P^{n-1}$.
- Perhaps most relevant example to us aside from what follows is the n**-frame bundle** or the **Stiefel bundle** $V_n(E) \to B$. Let (p, E, B) be a vector bundle of dimension n. We say that an ordered basis for $p^{-1}(b)$ where $b \in B$ is a **frame at** b. The points of $V_n(E)$ are n-tuples of orthogonal unit vectors in fibers of E, thus $V_n(E)$ is the subspace of the product of n copies of $S(E)$. The fiber of the Stiefel bundle is called the **Stiefel manifold** $V_n(\mathbb{R}^k)$, the space of orthonormal n-frames over $\mathbb{R}^k$, so we may think of the Stiefel bundle as the space of orthonormal n-frames over E.

3.2 Group Actions and Orbit Spaces

We hope at this point that it is obvious to the reader that the G in "principal G-bundle" comes from "group." As such, it necessary for us to quickly discuss some facts about topological groups in order to proceed. Recall that the **left action of a group** G on a space X is given by a continuous map $\rho : G \times X \to X$ for which

- $\rho(e, x) = x$ for all $x \in X$ with $e \in G$ the identity,
- $\rho(g, \rho(h, x)) = \rho(gh, x)$ for all $g, h \in G, x \in X$.

We often write $g \cdot x$ in lieu of $\rho(g, x)$. We can define in much the same way the notion of the **right action of a group** on a space X, denoted $x \cdot g$. The only difference here is that g acts on x from the right, rather than the left. Moving forward, we will only concern ourselves with topological groups, though many of the results demonstrated herein will hold regardless of whether G is endowed with a topology. It should be understood that when we say "group," we mean "topological group" unless otherwise specified. We say that a space X with a left action of G is a **left** G**-space** and, similarly, a space X with a right action of G is a **right** G**-space**. Since most statements will apply in the similar ways to both left and right actions, we will refer generally to an action unless we must specify the direction from which G acts. Defining an equivalence relation by $x \sim g \cdot x$ (or $x \sim x \cdot g$) for any $g \in G$ then allows us to define the **orbit space** $X/\sim = X/G$, for which each conjugacy class $G \cdot x$ is called the **orbit of** x. We are, in particular, interested in two specific types of group actions.

Definition III.3. We say that an action of G is **free** if for all $x \in X$, $g \cdot x = x$ (or $x \cdot g = x$ in the case of a right action) implies that $e = g$.

Definition III.4. An action of G is **transitive** if for every $x, y \in X$, there exists some $g \in G$ such that $g \cdot x = y$ (or $x \cdot g = y$).

Since we are working with homotopy classes of maps, it might be a good idea to consider the interaction between a group action and a homotopy, as well as that of a group action and a map between spaces in general.

Definition III.5. Let X, Y be G-spaces. A map $\phi : X \to Y$ is G**-equivariant** if $\phi(g \cdot x) = g \cdot \phi(x)$ (or $\phi(x \cdot g) = \phi(x) \cdot g$). We say that a G**-homotopy** between G-equivariant maps ϕ, ψ is a homotopy which is itself a G-equivariant map. For the sake of brevity, we refer to G-equivariant maps as G-maps.

Since products seem to be involved in describing local trivializations, it would also do us well to consider the interaction of group actions and product spaces.

Definition III.6. Let X be a right G-space and Y a left G-space. The **balanced product** of X and Y is the quotient space $X \times_G Y = X \times Y/\sim$ where we define equivalence by $(x \cdot g, y) \sim (x, g \cdot y)$ for $x \in X, y \in Y$.

Note that we can convert a right G-space to a left G-space (and vice versa) by setting $xg = g^{-1}x$. Once we do this, the balanced product becomes the usual notion of the orbit space $(X \times Y)/G$ of a product by the diagonal action $(x, y) \cdot g = (x \cdot g, g^{-1} \cdot y)$. With that, we consider some key properties of balanced products.

Proposition III.1. (Proposition 3.1 in [9].) *The balanced product is associative up to isomorphism. That is, if X is right G-space, Y is a left G-space and a right H-space where the actions of G, H commute (we'll call this a (G, H)-space henceforth), and Z is a left H-space, then there is a natural homeomorphism such that $(X \times_G Y) \times_H Z \cong X \times_G (Y \times_H Z)$, so we may write $X \times_G Y \times_H Z$ without ambiguity.*

Proof. Consider an element $\big[[x, y], z\big] \in X \times_G Y \times_H Z$. Then, choosing a representative of this equivalence class, we have

$$\big([x, y], h \cdot z\big) \sim (x, g \cdot y, h \cdot z) \sim (x, g \cdot y \cdot h, z) \sim (x \cdot g, y \cdot h, z) \sim \big(x \cdot g, [y, z]\big) \in \big[x, [y, z]\big],$$

so the quotient spaces are identified and we have that the balanced product is associative. In principle, we could also check that the map between the two quotient spaces and its inverse are continuous. □

We then have some special cases where H is a subgroup of G.

Corollary III.1.1. (Corollary 3.3 in [9].) *Let $H < G$. If X is a right G-space and Y is a left H-space, then $X \times_G G \times_H Y \cong X \times_H Y$.*

Proof. Since $H < G$, G may be considered a (G, H)-space, thus $X \times_G G$ is an H-space. More explicitly, $X \times_G G = (X \times G)/G$ which is homeomorphic to X, so this gives us that X is an H-space. We thus have that $(X \times_G G) \times_H Y \cong (X \times G)/G \times_H Y \cong X \times_H Y$. □

In other words, when $H < G$, we can "cancel" the balanced product relative to G. This is further exemplified in the following.

Corollary III.1.2. (Corollary 3.4 in [9].) *If X is a right G-space and $H \lhd G$, then $X \times_G (G/H) \cong X/H$.*

Proof. We are done once we take Y in the previous corollary to be a one-point space. □

3.3 Principal Bundles

We are finally in a position to discuss one of the central objects we wish to study: principal G-bundles.

Definition III.7. Let B be a space and G be a group whose action on B is free. A **principal G-bundle** is a fiber bundle (p, P, B) where P is a right G-space and $p : P \to B$ is a G-map such that there is an open cover $\mathcal{U}$ of B where, for each $U_\alpha \in \mathcal{U}$, there is a G-equivariant homeomorphism $h_\alpha : p^{-1}(U_\alpha) \to U_\alpha \times G$ which maps $p^{-1}(b)$ to $\{b\} \times G$, i.e. the action of G is transitive on each fiber.

Note that each $U_\alpha \times G$ has a right action by G defined by $(u, g) \cdot h = (u, gh)$ and the definition of a principal G-bundle implies that B is the orbit space P/G. Though it might be nice, we'll refrain from providing examples here, more for the purpose of novelty than anything else. A slight spoiler for the reader: we do so because our next major result is itself possibly the best example in some sense.

As with vector bundles, we'd like to know how to discuss relations between principal G-bundles in some respect.

Definition III.8. A G-map between the total spaces of principal G-bundles P, Q over the same base space is called a **morphism of principal G-bundles**.

We say that a principal G-bundle is **trivial** if the total space is isomorphic to $B \times G$. By this definition, every principal G-bundle is locally trivial. This might suggest that the total space of a principal G-bundle may be considered simply as a local product of B with the fiber G, but this is not the case; being the total space of a principal G-bundle is a much stronger condition, as demonstrated by the next two properties.

Proposition III.2. (Proposition 2.1 in [9].) *Every morphism of principal G-bundles is an isomorphism.*

Suppose first that $P \cong Q \cong B \times G$ and let $\sigma : P \to Q$ be a morphism. Then, for some continuous function $f : B \to G$, $\sigma(b, g) = (b, f(b)g)$, thus σ is an isomorphism with $\sigma^{-1}(b, g) = (b, f^{-1}(b)g)$ when P, Q are trivial. Since every principal G-bundle is locally trivial, we have that σ is bijective on each fiber. Further, since σ is a G-map, so is σ^{-1}, and we find that $p_1 \circ \sigma^{-1} = p_2$. If $U \subset P$ is open, then it follows that $V = p_1(U)$ is open since p_1 is an open map. We assume without loss of generality that V is a trivializing open set. Then, $U \cong V \times G$. By definition then, $\sigma(U) = p_2^{-1}(V) \cong V \times G$, thus σ^{-1} is well-defined and the morphism is an isomorphism.

Proposition III.3. (Proposition 2.2 in [9].) *A principal G-bundle is trivial if and only if it admits a section.*

Proof. Every local product admits a section, so if P is trivial, then it admits a section. Now suppose $s : B \to P$ is a nontrivial section. Then, the map $\phi : B \times G \to P$ defined by $\phi(b, g) = s(b)g$ is a morphism and, thus, an isomorphism, so P is trivial. □

The fact that being a principal G-bundle is a relatively strong condition means that we are not guaranteed that if $H < G$, then $P \to P/H$ is a principal H-bundle. We do, however, have a particular criterion for determining when this is possible.

Definition III.9. A normal subgroup H of G is **admissible** if $(q, G, G/H)$, where q is the quotient map, is a principal H-bundle.

Proposition III.4. (Proposition 3.5 in [9].) *If (p, P, B) is a principal G-bundle and H is an admissible subgroup of G, then $(q, P, P/H)$ is a principal H-bundle.*

Proof. We have for any normal subgroup H of G that $P/H = P \times_G G/H$. We may then identify the quotient map $P \to P/H$ with $P \times_G G \to P \times_G G/H$. Since H is admissible, we have $q : P \to P/H$ defined by composition, thus $(q, P, P/H)$ is a principal H-bundle. □

This leads us to a lemma that is of no use immediately, but will be helpful later on. In fact, this is exactly why we emphasized the importance of sections previously.

Lemma III.1. (Proposition 6.1 in [9].) *Let (p, P, B) be principal G-bundle over B and X a right G-space. We write $\mathrm{Hom}_G(P, X)$ for the **set of G-maps** $P \to X$ and $\Gamma(P \times_G X \to B)$ for the **set of sections of the fiber bundle** $(q, P \times_G X, B)$. There is a natural bijection $\Phi : \mathrm{Hom}_G(P, X) \to \Gamma(P \times_G X \to B)$ defined by $f \mapsto s_f$.*

Here we have a natural use of the conversion of a right G-space to a left G-space, as we previously discussed.

Proof. Consider first a trivial bundle $P \cong B \times G$. In this case,

$$\mathrm{Hom}_G(P, X) = \mathrm{Hom}_G(B \times G, X) = \mathrm{Hom}(B, X),$$

the set of continuous maps $B \to X$. By the same token, we have that

$$\Gamma(P \times_G X \to B) = \Gamma(B \times G \times_G X \to B) = \Gamma(B \times X \to B) = \mathrm{Hom}(B, X),$$

thus the statement holds trivially for trivial bundles. Since every principal G-bundle is locally trivial, as before, we have that Φ is bijective on each fiber. We should then check that Φ is well-defined when sections overlap, then construct a well-defined inverse function, but we'll omit this. □

3.4 From Vector Bundles to Principal G-Bundles

Now let's examine that example we promised by relating principal G-bundles to the vector bundles with which we've become familiar. In the ensuing discussion, we characterize the principal G-bundle in two distinct but related ways. Both require the notion of the Whitney sum.

Definition III.10. Let $(p_1, E_1, B), (p_2, E_2, B)$ be two vector bundles over B. The **Whitney sum** of E_1 and E_2 is given by $E_1 \oplus_B E_2 = \{(v_1, v_2) \in E_1 \times E_2 | p_1(E_1) = p_2(E_2)\}$. This generalizes naturally to a sum over n total spaces.

If we have (p, E, B) an n-dimensional complex vector bundle, we identify (q, P, B) with the Stiefel bundle $V_n(E)$. Observe that this is a subspace of the above defined Whitney sum of n copies of the total space E, and thus we endow it with the subspace topology inherited from nE. We justify this identification by noting that the right action of $GL(n, \mathbb{C})$ on nE is free, thus (q, P, B) is a principal $GL(n, \mathbb{C})$-bundle over B. Our second characterization begins with the vector bundle $Hom(\epsilon^n, E)$, the space of homomorphisms from the n-dimensional trivial bundle $\epsilon^n = B \times \mathbb{C}^n$ to the total space E. We, again, identify this with the Whitney sum nE and then identify the principal $GL(n, \mathbb{C})$-bundle (q, P, B) with the subspace $Iso(\epsilon^n, E)$, the space of linear isomorphisms from ϵ^n to E, justified by the fact that the right action of $GL(n, \mathbb{C})$ on ϵ^n induces a free right action on $Hom(\epsilon^n, E)$. We emphasize that the bundles $Hom(\epsilon^n, E)$ and $Iso(\epsilon^n, E)$ are *not* the same as the vector spaces often denoted as such. These vector spaces are, in fact, the spaces of sections of the Hom and Iso with which we are working. Both of these characterizations are related by a G-equivariant homeomorphism $Iso(\epsilon^n, E) \to V_n(E)$ defined by $(b, f : \mathbb{C}^n \to E_b) \mapsto (b, (f(e_1), ..., f(e_n)))$.

Theorem III.1. (Proposition 4.1 in [9].) *Denote by $\mathcal{P}_G(B)$ the set of isomorphism classes of principal G-bundles over B. The map $\phi : \mathcal{P}_{GL(n,\mathbb{C})}(B) \to \mathrm{Vect}^n_{\mathbb{C}}(B)$ is a bijection.*

Proof. Denote by ξ the vector bundle (p, E, B). There is then a clear map $\sigma : Iso(\epsilon^n, \xi) \to E$ defined by $\sigma([b, f, v]) = (b, f(v))$ where $f : \mathbb{C}^n \to E_b$ and $v \in \mathbb{C}^n$. If $g \in GL(n, \mathbb{C})$, then $f(g \cdot v) = g \cdot f(v)$, thus σ is well-defined. Now consider a trivializing neighborhood U of ξ. Over U, σ is an isomorphism

$$U \times GL(n, \mathbb{C}) \times_{GL(n,\mathbb{C})} \mathbb{C}^n \to U \times \mathbb{C}^n,$$

thus σ is a linear isomorphism on fibers and, therefore, an isomorphism of vector bundles. As such, we have that Ψ takes the isomorphism class of ξ to $[b, f, v]$.

Now consider a principle $GL(n, \mathbb{C})$-bundle (p, P, B) and let $\xi = P \times_{GL(n,\mathbb{C})} \mathbb{C}^n$. If we can show that P is isomorphic to $V_n(\xi)$, then we will have that $\Psi\Phi$ is the identity. Proposition III.2 reveals that it will suffice, then, to construct a G-map $\tau : P \to V_n(\xi)$. Let $\tau(x) = ([x, e_1], ...[x, e_n])$. It is clear that τ is continuous. Identifying $V_n(\xi)$ with $Iso(\epsilon^n, \xi)$, we have that $\tau(x) = (b, f_x)$ where $b = [x]$ and $f_x : \mathbb{C}^n \to E_b$ takes e_i to $[x, e_i]$. Therefore, if $g \in GL(n, \mathbb{C})$, then f_{xg} maps e_i to $[x \cdot g, e_i] = [x, g \cdot e_i]$, so $\tau(x \cdot g) = \tau(x) \cdot g$, as desired. $\square$

In other words, every vector bundle is principal G-bundle! More specifically, we have shown that every n-dimensional complex vector bundle is a principal $GL(n, \mathbb{C})$ bundle and, by replacing $\mathbb{C}$ with $\mathbb{R}$, we find that the same is true in the real case with structure group $GL(n, \mathbb{R})$. Combining this theorem with Theorem II.1, we now have a bijection between $[S^{k-1}, GL(n, \mathbb{C})]$ and $\mathcal{P}_{GL(n,\mathbb{C})}(S^k)$, so we are left to generalize spheres to spaces, which will go hand-in-hand with generalizing clutching functions.

IV The (Al)Most General Scenario

4.1 Pullback Bundles

To accomplish our remaining goals, we need just a few more constructions and concepts, starting with the pullback and pullback bundle.

Definition IV.1. Let $f : A \to B$ be a map between spaces and (p, E, B) a fiber bundle. There then exists a fiber bundle (p', E', A) and a map $f' : E' \to E$ which carries the fiber of E' over each $a \in A$ to the fiber of E over $f(a) \in B$. These fiber bundles are unique up to isomorphism. We call E' the **pullback of E by f** and (p', E', A) the **pullback bundle**. To emphasize the relation to the map f, E' is written as $f^*(E)$.

The relations between the spaces E', E, A, B are demonstrated in the following commuting diagram.

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

Though we have taken this as a definition, it is possible to prove the statements regarding the existence and uniqueness of pullback bundles. The enthused reader will find this stated as Proposition 1.15 in Hatcher's *Vector Bundles and K-Theory* [7], along with a proof.

There are a few properties of pullbacks and pullback bundles that we should note immediately:

- $\mathbb{1}^*(E) = E$.
- $(fg)^*(E) = g^*(f^*(E))$.
- If E is trivial, i.e. $E = B \times F$, and F is a fiber, then $f^*E \subset A \times B \times F$, consisting of triples (a, b, v) where $f(a) = b$, thus $E' = A \times F$.

Let's consider some examples of pullbacks to ground ourselves (and recover from the *certain* example withdrawal from the group action and principal G-bundle discussions).

- We may consider restrictions of a vector bundle (p, E, B) to a sub-base space $A \subset B$ to be pullbacks, since the inclusion map $p^{-1}(A) \hookrightarrow E$ is an isomorphism on each fiber.
- If f is the constant map, then f^*E is a trivial bundle with total space $A \times p^{-1}(b)$.
- The tangent bundle TS^n is the pullback of $T\mathbb{R}P^n$ by the quoitent map $S^n \to \mathbb{R}P^n$.

Now we move to a more general scenario. For the remainder of our discussion, we must make some assumptions about the spaces we consider. In particular, we will focus on **CW-complexes**, which, for all intents and purposes, are "nicely behaved" topological spaces. For a more rigorous definition and study of these spaces, we recommend the discussion beginning on page 5 of Hatcher's *Algebraic Topology* [6]. CW-complexes serve as an intermediate stage between spheres and general spaces, for spheres can be thought of as n-dimensional CW-complexes. Breaking down and gluing CW-complexes and their constituent parts, called **cells**, we can construct new CW-complexes. Once we see that our results hold for CW-complexes, we can, in principle, extend them to general spaces, but we will not approach topic formally. To see why this is possible, see Chapter 4 of [6], starting with page 348. An added benefit of the CW-complex formalism is that it grants us access to Serre fibrations:

Definition IV.2. For topological spaces E, B, the map p is a **Serre fibration** if it has the homotopy lifting property with respect to all CW-complexes X. That is, for any CW-complex X, there exists a homotopy $H : X \times I \to E$ such that the following diagram commutes. Here, $G : X \times I \to B$ is a homotopy and i_0 is the inclusion map $x \mapsto (x, 0)$.

$$\begin{tikzcd} X \arrow[r, "f"] \arrow[d, "i_0"'] & E \arrow[d, "p"] \\ X \times I \arrow[r, "G"] \arrow[ur, dashed, "H"] & B \end{tikzcd}$$

We direct the reader to Mitchell's *Notes on Serre Fibrations* [10] for more information on Serre fibrations and their study, as well as the discussion which begins on page 375 of [6] for a general study of fibrations. This is a wonderful topic upon which we would like to exposit, however we must leave it be in the interest of time. Henceforth, we'll freely use facts about Serre fibrations.

We've seen a few examples about pullbacks on vector bundles, but we have yet to discuss pullbacks on principal G-bundles; we'll address that now.

Proposition IV.1. (Proposition 7.1 in [9].) *Let X be a space with (p, P, X) a principal G-bundle over X. If B is a CW-complex with maps $f, g : B \to X$ which are homotopic, then the pullbacks f^*P, g^*P are isomorphic as total spaces of principal G-bundles over B.*

Define $F : B \times I \to X$ to be a homotopy from f to g. We then reduce the proof to that of the following lemma once we consider the pullback F^*P:

Lemma IV.1. (Lemma 7.2 in [9].) *Let $Q \to B \times I$ be a principal G-bundle with Q_0 its restriction to $B \times \{0\}$. Then Q is isomorphic to $Q_0 \times I$ and, in particular, Q_0 is isomorphic to Q_1, the restriction to $B \times \{1\}$.*

Proof. Since every morphism between princpal G-bundles is an isomorphism by Proposition III.2, it suffices to construct a morphism $Q \to Q_0 \times I$. Lemma III.1 then tells us that this is equivalent to constructing a section s of $Q \times_G (Q_0 \times I) \to B \times I$, but the same lemma yields a section s_0 on B_0. We then apply Corollary 5.3 of [10] to extend s_0 to a section s over $B \times I$, thus concluding the proof. $\square$

Corollary IV.1.1. (Corollary 7.3 in [9].) *Every principal G-bundle over a contractible base space is trivial.*

Proof. Take $X = \{*\}$ in this proposition. $\square$

4.2 In (Almost) Full Generality

We are now nearly in a position to approximate our main result. Our approximate statement requires one last concept.

Definition IV.3. A space X is **weakly contractible** if for all $n \geqslant 0$, every map $S^n \to X$ extends to a map $D^{n+1} \to X$. By Whitehead's Theorem (see pages 346 - 347 of [6]), every weakly contractible CW-complex is contractible.

Theorem IV.1. (Theorem 7.4 in [9].) *Let (p, EG, BG) be a principal G-bundle with EG weakly contractible. For all CW-complexes X, there exists a bijection $\Phi : [X, BG] \to \mathcal{P}_G(X)$, defined by $f \mapsto f^*EG$, where BG is called a* ***classifying space*** *of G and (p, EG, BG) is called a* ***universal bundle*** *over the base space BG.*

Proof. Define a principal G-bundle $Q \to BG$. The Serre fibration $Q \times_G EG \to BG$ then has a weakly contractible fiber, thus it admits a section. Lemma III.1 tells us that this section determines a G-map $\tilde{f} : Q \to EG$. If $f : BG \to X$ is the induced map on orbit spaces, then $Q = f^*EG$, so Φ is surjective.

Now consider two maps $f_0, f_1 : BG \to X$ and an isomorphism $\Psi : f_0^*EG \to f_1^*EG$. Denote by Q the principal G-bundle $(f_0^*EG) \times I$ over $BG \times I$ and consider the local product $\rho : Q \times_G EG \to BG \times I$. The equivariant maps $Q_0 = f_0^*EG \to EG$ and $Q_1 = f_0^*EG \to f_1^*EG \to EG$ define a section of ρ over $BG \times \{0\} \cup BG \times \{1\}$. Since the fiber is weakly contractible, this section extends over all of $BG \times I$, thus determining a G-map $Q \to EG$. Passing to the orbit spaces then yields a homotopy $BG \times I \to X$ from f_0 to f_1, therefore Φ is injective, so it is bijective. $\square$

Notice that the goal of generalizing clutching functions has now been met. By elevating spheres to CW-complexes, we have implicitly accomplished the final minor goal we sought to achieve.

4.3 Special Case: Vector Bundles of Paracompact Spaces

To demonstrate the power of this result, let's compare two different approaches to the same problem.

Definition IV.4. A Hausdorff space X is **paracompact** if every open cover $\mathcal{U}$ of X has associated to it a set of maps $\{\varphi_\beta : X \to I\}$ which has support contained in $\mathcal{U}$ and $\sum_\beta \varphi_\beta = 1$. The support of a map is the closure of its zero set.

Everything we have done in terms of CW-complexes can be expressed in terms of paracompact spaces, since paracompactness is a weaker condition than being a CW-complex (see page 36 of [7] for a proof). This formalism is expressed in Dold's article, "Partitions of Unity in the Theory of Fibrations" [4], and relies on the notion of a **partition of unity**, which is a set of m maps $\{\varphi_j\}$ for which

$$\psi_j = \sum_{i=0}^{j} \varphi_i,$$

where $\psi_0 = \varphi_0 = 0$ and $\psi_m = 1$ such that the support of each φ_i is contained in an open set $U_\alpha \in \mathcal{U}$, an element of an open cover of the codomain. Let's use the paracompactness formalism to discuss a special case of Theorem II.1 for a particular space.

Definition IV.5. The **Grassmann manifold** $G_n(\mathbb{R}^k)$ (where $n \leqslant k$) is the collection of n-dimensional vector subspaces of $\mathbb{R}^k$.

Let $E_n(\mathbb{R}^k) = \{(\ell, v) \in G_n(\mathbb{R}^k) \times \mathbb{R}^k | v \in \ell\}$.

Lemma IV.2. (Lemma 1.15 in [7].) *$(p, E_n(\mathbb{R}^k), G_n(\mathbb{R}^k))$ is a vector bundle with $p(\ell, v) = \ell$.*

Proof. Let $\pi_\ell : \mathbb{R}^k \to \ell$ be the orthogonal projection map, where $\ell \in G_n(\mathbb{R}^k)$, and define

$$U_\ell = \{\ell' \in G_n(\mathbb{R}^k) | \dim(\pi_\ell(\ell')) = n\}.$$

Let's show that U_ℓ is open in $G_n(\mathbb{R}^k)$ and that the map

$$h : p^{-1}(U_\ell) \to U_\ell \times \ell \cong U_\ell \times \mathbb{R}^n$$

defined by $h(\ell, v) = (\ell', \pi_\ell(v))$ is a local trivialization of $E_n(\mathbb{R}^k)$.

Observe that the map which takes $V_n(\mathbb{R}^k)$ to $G_n(\mathbb{R}^k)$ is open, thus U_ℓ is open if its preimage in $V_n(\mathbb{R}^k)$ (consisting of orthonormal frames $v_1, ..., v_n$ such that $\pi_\ell(v_1), ..., \pi_\ell(v_n)$ are independent) is open. We write $\mathbf{A}$ for the matrix representation of π_ℓ in the standard basis of $\mathbb{R}^k$ and a fixed basis in ℓ. We then have that the column vectors $\mathbf{A}v_1, ...\mathbf{A}v_n$ have nonzero determinant. The fact that det is continuous implies that the frames $v_1, ..., v_n$ yielding a nonzero determinant form an open set in $V_n(\mathbb{R}^k)$.

We see that h is clearly a linear isomorphism on each fiber, so it is a bijection. It then remains to show that both h and h^{-1} are continuous. For each $\ell' \in U_\ell$, there is a unique invertible linear map $L_{\ell'} : \mathbb{R}^k \to \mathbb{R}^k$ which restricts π_ℓ to ℓ' and the identity on $\ell^\perp = \ker(\pi_\ell)$. We then write $L_{\ell'}$ as a matrix $L_{\ell'} = \mathbf{AB^{-1}}$ where $\mathbf{B}$ maps the standard basis onto $v_1, ...v_n, v_{n+1}, ...v_k$, with $v_1, ..., v_n$ an orthonormal basis of ℓ' and $v_{n+1}, ..., v_k$ a fixed basis for $\ell^\perp$, and $\mathbf{A}$ maps the standard basis to $\pi_\ell(v_1), ..., \pi_\ell(v_n), v_{n+1}, ..., v_k$. Since both $\mathbf{A}$ and $\mathbf{B}$ depend continuously on $v_1, ..., v_n$ and matrix multiplication and matrix inversion are continuous operations for matrices of nonzero determinant, $L_{\ell'}$ depends continuously on $v_1, ..., v_n$. $L_{\ell'}$, however, depends only on ℓ', not the basis $v_1, ..., v_n$ of ℓ', thus it follows that $L_{\ell'}$ depends *continuously* on ℓ' since $G_n(\mathbb{R}^k)$ has the quotient topology from $V_n(\mathbb{R}^k)$. By the same token, $L_{\ell'}^{-1}$ depends continuously on ℓ'.

We now have that h is continuous since $h(\ell', v) = (\ell', \pi_\ell(v)) = (\ell', L_{\ell'}(v))$ and, similarly, h^{-1} is continuous because $h^{-1}(\ell', w) = (\ell', L_{\ell'}^{-1}(w))$. ☐

From here on out, we will work with the space $\mathbb{R}^\infty = \prod_{k=1}^\infty \mathbb{R}^k$, so we will write

$$G_n = G_n(\mathbb{R}^\infty) = \bigcup_k G_n(\mathbb{R}^k)$$

$$E_n = E_n(\mathbb{R}^\infty) = \bigcup_k E_n(\mathbb{R}^k).$$

Observe that (p, E_n, G_n) is still a vector bundle in this case; simply take U_ℓ to be the union of the U_ℓs defined above for increasing values of k, so that the local trivializations h constructed for finite k give together a local trivialization over this new U_ℓ. To guarantee continuity, we note that we must use the weak topology, the topology in which a set $A \subset X$ is open if and only if $A \cap X^n$ is open in X^n for each n. Here, we define X^n recursively by

$$X^n = X^{n-1} \bigsqcup_\alpha e_\alpha^n,$$

where each e_α^n is an open n-disc and X^0 is a discrete space.

Proposition IV.2. (Theorem 1.16 in [7].) *Let X be a paracompact space. Then G_n is the classifying space for n-dimensional vector bundles over X and $E_n \to G_n$ is the universal bundle, i.e. the map $\Phi : [X, G_n] \to Vect^n(X)$ is a bijection where $[X, G_n]$ is the set of isomorphism classes of maps $X \to G_n$.*

Proof. Let's take a map $f : X \to G_n$ and an isomorphism $E \cong f^*E_n$. We then have a commutative diagram

$$\begin{tikzcd}
E \arrow[r, "\cong"] \arrow[dr, "p"'] & f^*E_n \arrow[r, "\tilde{f}"] \arrow[d] & E_n \arrow[r, "\pi"] \arrow[d] & \mathbb{R}^\infty \\
& X \arrow[r, "f"] & G_n &
\end{tikzcd}$$

where $\pi(\ell, v) = v$. The composition of maps in the top row yields a map $g : E \to \mathbb{R}^\infty$ which is a linear injection on each fiber since $\tilde{f}$ and π are linear injections. Conversely, if we begin with a map $g : E \to \mathbb{R}^\infty$ which is a linear injection on each fiber, we define $f : X \to G_n$ by taking $f(x)$ to be the n-plane $g(p^{-1}(x))$, thus satisfying the above commutative diagram.

Now suppose that (p, E, X) is an n-dimensional vector bundle. Let $\mathcal{U}$ be an open cover of X such that E is trivial over each $U_\alpha \in \mathcal{U}$. We claim that there is a countable open cover $\{V_k\}$ such that each V_k is a disjoint union of open sets contained in some U_α and a partition of unity $\{\varphi_k\}$ supported in V_k.

Let $\{\varphi_\beta\}$ be a partition of unity with support in some element of $\mathcal{U}$. For a finite set S of functions φ_β, define V_S as the subset of X for which each $\varphi_\beta \in S$ is strictly greater than any $\varphi_\beta \notin S$. Observe that only finitely many φ_β are nonzero in a neighborhood of any $x \in X$, thus V_S is determined by finitely many inequalities among functions φ_β near x, which means that V_S is open. Furthermore, it follows that V_S is contained in some $U_\alpha \in \mathcal{U}$, namely any U_α which contains the support for any $\varphi_\beta \in S$, since $\varphi_\beta \in S$ implies that $\varphi_\beta > 0$ on V_S. Now we denote by V_k the union of all sets V_S which have k elements. Each of these is clearly disjoint by the definition of the support of a function. We then see that the set $\{V_k\}$ covers X since if $x \in X$, then $x \in V_S$ where $S = \{\varphi_\beta | \varphi_\beta(x) > 0\}$. We then define $\{\varphi_\gamma\}$ to be a partition of unity whose elements are supported in some element of $\{V_k\}$ and define φ_k to be the sum of those φ_γ which are supported in V_k, but not V_j when $j < k$.

Now let $g_k : p^{-1}(V_k) \to \mathbb{R}^n$ be the composition of a trivialization $p^{-1}(V_k) \to V_k \times \mathbb{R}^n$ and a projection onto $\mathbb{R}^n$. The map $(\varphi_k p) g_k$, which carries $v \mapsto \varphi_k(p(v)) g_k(v)$, then extends to a map $E \to \mathbb{R}^n$ which is zero outside of $p^{-1}(V_k)$. As we observed before, only finitely many of the φ_k are nonzero near a given point of X and we have by definition that at least one φ_k must be nonzero, so these φ_k are the coordinates of a map $g : E \to (\mathbb{R}^n)^\infty = \mathbb{R}^\infty$, which is a linear injection on each fiber, thus the map $[X, G_n] \to \mathrm{Vect}^n(X)$ is surjective.

To demonstrate injectivity, consider two isomorphisms $E \cong f_0^* E_n$ and $E \cong f_1^* E_n$ which are pullbacks of maps $f_0, f_1 : X \to G_n$. These then yield linear injections $g_0, g_1 : E \to \mathbb{R}^\infty$. In order for the Φ to be injective, we must find a homotopy between g_0 and g_1, for this will then translate to a homotopy between f_0 and f_1, defined by $f_t(x) = g_t(p^{-1}(x))$.

We construct a homotopy as follows: consider a homotopy $L_t : \mathbb{R}^\infty \to \mathbb{R}^\infty$ defined by

$$L_t(x_1, x_2, ...) = (1 - t)(x_1, x_2, ...) + t(x_1, 0, x_2, 0, ...).$$

This is a linear map for each t with kernel 0, thus L_t is injective. When we compose L_t with g_0, we have a map which moves the image of g_0 to into the odd-numbered coordinates. We can construct a similar homotopy which moves the image of g_1 into the even coordinates. Rename these compositions g_0 and g_1. Then, we define a new homotopy $g_t = (1-t)g_0 + t g_1$, which is a linear injection on the fibers for each t since g_0 and g_1 are linear injections on the fibers, thus Φ is injective and, thereby, bijective. □

In order to arrive at this result without many of the tools we've developed, we required a great deal of effort, time, and results (in fairness, we've spent equally as much time developing our own tools here). Now let's see what happens when we make use of our shiny, new toys. Recognizing that G_n is a CW-complex (more information about this can be found on pages 31 - 34 of [7]), let $V_n(\mathbb{R}^\infty) = V_n$ be the local product over G_n with fiber $GL(n, \mathbb{R})$. Then, the natural map $V_n \to G_n$ is a principal $GL(n, \mathbb{R})$-bundle.

Proposition IV.3. (Theorem 7.7 in [9].) *(q, V_n, G_n) is a principal $GL(n, \mathbb{R})$-bundle and, thus, G_n is the classifying space of $GL(n, \mathbb{R})$.*

Proof. Through the process of Gramm-Schmidt orthogonalization, we find that the subspace $V_n^O \subset V_n$ of orthonormal frames is a deformation retract of V_n. There is a proof with some additional technology we have not addressed on page 14 of [9] showing that the orthonormal frames are weakly contractible, thus V_n is weakly contractible, so G_n is a principal $GL(n, \mathbb{R})$-bundle. We therefore conclude that G_n is the classifying space of $GL(n, \mathbb{R})$. □

Proposition IV.2 then follows as a corollary by applying Theorem III.1.

V Conclusion

Of course, the classification story does not end here. As discussed in [7], there is a deeper story regarding the notion of characteristic classes of various types, in addition to a deeper theory of classifying spaces and classifying maps in the general literature. Furthermore, these results find some very interesting applications in algebraic topology, differential topology, gauge theory, and physics, though the background necessary to discuss this is beyond the scope of what we can reasonably develop here (after all, even the discussion surrounding G_n was more difficult to develop than our prior results). Should the reader be eager to learn more, we encourage them to continue to study topology, both algebraic and differential, and embed themselves deeper into this rich narrative.

VI References

[1] Andrew Baker, *Matrix Groups: An Introduction to Lie Group Theory*, 2002.

[2] Joseph Breen, *Path Connectedness and Invertible Matrices*, 2016.

[3] Jean Dieudonné, *A History of Algebraic and Differential Topology, 1900-1960*, 1989.

[4] Albrecht Dold, "Partitions of Unity in the Theory of Fibrations," *Annals of Mathematics*, Vol. 78, No.2, 1963.

[5] Kevin Hartnett, "Why Mathematicians Like to Classify Things," *Quanta Magazine: Abstractions*, 2017.

[6] Allen Hatcher, *Algebraic Topology*, 2001.

[7] Allen Hatcher, *Vector Bundles and K-Theory*, 2017.

[8] Max Karoubi, *K-Theory: An Introduction*, 1978.

[9] Stephen A. Mitchell, *Notes on Principal Bundles and Classifying Spaces*, 2001.

[10] Stephen A. Mitchell, *Notes on Serre Fibrations*, 2001.

[11] Walter Rudin, *Principles of Mathematical Analysis*, Third Edition, 1976.