Reconciling Vector Bundles

By Jordan Grant

May 6, 2020

1 History and Motivation

The purpose of this paper will be to define the concept of the vector bundle through two separate definitions, and then to prove that these definitions, and the intuitive reasoning behind them, are compatible. Before we do this however, I would like to give some motivation for our study of these objects, and just how their necessity arises.

When students first begin studying surfaces in $\mathbb{R}^3$, they are introduced to local invariants such as curvature by means of the tangent plane at a point.

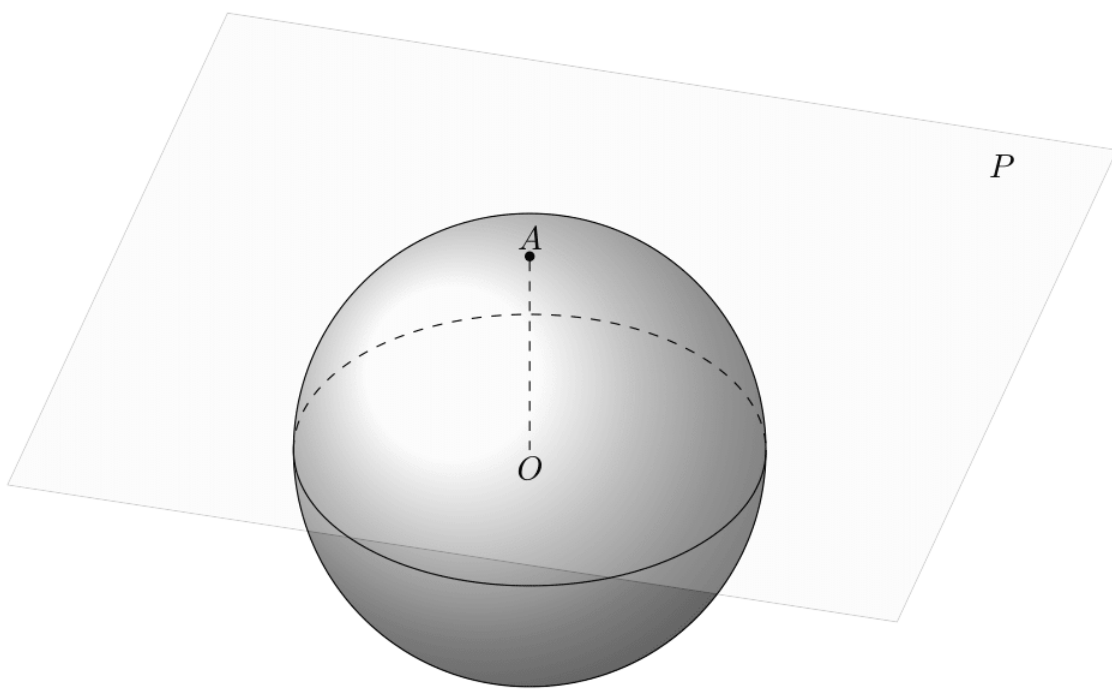

Figure 1: Tangent space P to the point A on S^2

By defining a tangent plane to a point P on our surface M, we can then use the Gauss Map to define a normal vector to the plane. We then define the Shape Operator,

$$S_A : T_A(M) \to T_A(M) \text{ where for } v \in T_A(M),$$
$$S_A(v) = -D_v n(A)$$
where $n(A)$ is the unit vector normal to the plane P.

which allows us to define curvatures in a specific direction around our point P. In order to see the intuition behind this, imagine an old man with a walking stick. The old man can't see very well, so when he's walking along he never knows whether he's on the edge of a hill, or at the bottom of a slope. One thing he does know however is his walking stick, whose orientation and speed he's always perfectly cognizant of.

So when the old man wants to know the shape of the ground beneath him, he simply balances his walking stick straight up on the ground, picks the direction he wants to go in, and lets the stick fall in that direction. Based on the speed at which it fell, and its orientation after falling, the old man knows the shape of the ground in that direction and whether or not it's safe to proceed.

This notion of relating the tangent plane to the directional derivative is indicative of a much more fundamental relationship. Instead of thinking of a tangent plane, think instead of a tangent space, T_pM, associated to each point on the surface. Then instead of vectors, T_pM can be conceived of as a vector space of derivations at p,

A linear map $\alpha : C^\infty(M) \to R$ is called a derivation at p if it satisfies

$$\alpha(fg) = f(p)\alpha(g) + g(p)\alpha(f) \ \forall f, g \in C^\infty(M)$$

These derivations generalize the concept of directional derivatives for more general spaces. Another advantage of this view of T_pM is that as an abstract vector space, T_pM isn't reliant on an embedding. Therefore the tools that we have developed generalize to an embedding free environment. This naturally leads us to asking what structure the space of all tangent spaces for a particular surface has.

Here an interesting contention arises. We defined our tangent spaces in an embedding free way, and so the space of tangent spaces should be equivalent regardless of embedding. However, if we consider the sphere embedded inside of a space diffeomorphic to $\mathbb{R}^3$, then the space of tangent planes need not be the same as those we get when we use $\mathbb{R}^3$ as an embedding.

So if the spaces are not identical across embeddings, in what way are they equal? They would both be examples of tangent bundles of the sphere, and therefore be diffeomorphic as Vector Bundles. Therefore by studying these spaces as Vector Bundles we can encapsulate the common traits all of these spaces possess, irrespective of their particular embeddings.

2 Defining Vector Bundles

Broadly, a vector bundle E over a topological space M can be seen as a collection of vector spaces, called fibres, satisfying some local triviality conditions on M. A more useful intuitive notion is that a vector bundle encodes the most universal components of embedding, those that generalize to all spaces that you embed M inside of. This allows us to study M in an embedding free way, while still enjoying some of the perks of their structure.

I will now give two definitions of Vector Bundles, which I will prove equality for later.

Local Trivialization Definition:

Let M be a topological space. A (real) vector bundle of rank k over M is a topological space E together with a surjective continuous map $\pi : E \to M$ satisfying the following conditions:

(i) For each $p \in M$, the fiber $E_p = \pi^{-1}(p)$ over p is endowed with the structure of a k-dimensional vector space.

(ii) For each $p \in M$, there exists a neighborhood U of p in M and a homeomorphism $\Phi : \pi^{-1}(U) \to U \times \mathbb{R}^k$ (called a local trivialization of E over U) satisfying the following conditions:

$$\pi_U \circ \Phi = \pi$$
(where $\pi_U : U \times \mathbb{R}^k \to U$ is the projection on the first factor)

For each $q \in U$, the restriction of Φ to E_q is a vector space isomorphism from E_q to $\{q\} \times \mathbb{R}^k \cong \mathbb{R}^k$

Bundle Transition Definition:

Let M be a topological space. Then E is the (real) vector bundle of rank k over M is defined by the two following criteria:

A locally finite cover Ω of M

For each $U, U' \in \Omega$ an assignment of a map $g_{U,U'} \to Gl_n(\mathbb{R})$ satisfying the following constraints:

$$g_{U,U'} = g^{-1}_{U',U}$$

If $U, U', U'' \in \Omega$ and $U \cap U' \cap U'' \neq \emptyset$ then,
$$g_{U,U'} \circ g_{U',U''} \circ g_{U'',U} = id$$

The bundle E is then defined as the quotient of the disjoint union $\cup_{U \in \Omega}(U \times \mathbb{R}^k)$ by the equivalence relation that puts $(p', v') \in U' \times \mathbb{R}^k$ equivalent to $(p, v) \in U \times \mathbb{R}^k$ if and only if $p = p'$ and $v' = g_{U,U'}(p)v$. We then define a continuous map $\pi : E \to M$ as $\pi([p, v]) = p$.

The primary difference between these two definitions is one of focus. The Local Trivialization definition stresses global structures, splitting our underlying topological space into pseudo copies of $\mathbb{R}^n$ and imparting some the structure of euclidean space onto our topological space. The Bundle Transition definition on the other hand focuses on how we transition from one of these pseudo copies to the next, and is far more useful when dealing with problems that are local.

Thus when working on global problems we would like to be able to use the Local Trivialization definition, and when working locally we would like to use Bundle Transition definition. In order to do this we need to prove that these definitions are equivalent, and therefore interchangeable.

3 Equivalency of Definitions

Local Trivialization $\to$ Bundle Transition:

Suppose $\pi : E \to M$ is a vector bundle defined using local trivialization's.

Suppose $h \in U \cap U'$, and have φ and Φ be the local trivialization's on U and U' respectively.

Then restricting Φ and φ to $\pi^{-1}(h)$ yields an isomorphism on $\{h\} \times \mathbb{R}^n$,

$$g_{U',U}(h) = (\varphi_h \circ \Phi_h^{-1}) : \{h\} \times \mathbb{R}^n \to \{h\} \times \mathbb{R}^n$$

Since $\{h\}$ x $R^n \cong R^n$ the matrix representation of $g_{U',U}$ will be an $n \times n$ invertible matrix and therefore an element of $GL_n(R)$. Using this we define a new function, $g_{U,U'} : U \cap U' \to GL_n(R)$.

Co-Cycle Constraint:

Suppose $h \in U \cap U' \cap U''$. Give U, U',and U'' the local trivialization's, $\Phi, \Phi^{'}, \Phi^{''}$ respectively.

Then,

$$\begin{aligned} & g_{U,U'} \circ g_{U',U''} \circ g_{U'',U} \\ &= \Phi \circ (\Phi')^{-1} \circ \Phi' \circ (\Phi'')^{-1} \circ \Phi'' \circ \Phi^{-1} \\ &= \Phi \circ ((\Phi')^{-1} \circ \Phi') \circ ((\Phi'')^{-1} \circ \Phi'') \circ \Phi^{-1} \\ &= \Phi \circ \Phi^{-1} = id \end{aligned}$$

Bundle Transition → Local Trivialization:

Suppose $\pi : E \to M$ is a vector bundle defined using bundle transition maps.

(i) $\pi^{-1}(p)$ is a vector space:

For $a \in R$ and $[p, v_1], [p, v_2] \in pi^{-1}(p)$ define vector addition as $[p, v_1] + [p, v_2] = [p, v_1 + v_2]$ and scalar addition as $r[p, v] = [p, vr]$. We will now confirm the vector space structure on $\pi^{-1}(p)$. For the sake of brevity, for the following proofs assume that $[p, v_1], [p, v_2], [p, v_3], [p, v] \in \pi^{-1}(p)$ and $a, b \in R$.

Additive Closure:
$[p, v_1] + [p, v_2] = [p, v_1 + v_2]$ since $v_1 + v_2 \in \mathbb{R}^n$, $[p, v_1] + [p, v_2] \in \pi^{-1}(p)$

Scalar Closure:
$a[p, v] = [p, av]$ since $av \in \mathbb{R}^n$, $[p, av] \in \pi^{-1}(p)$

Commutativity of Addition:
$[p, v_1] + [p, v_2] = [p, v_1 + v_2] = [p, v_2 + v_1] = [p, v_2] + [p, v_1]$ since $v_1, v_2 \in \mathbb{R}^n$.

Associativity of Addition:
$[p, v_1] + ([p, v_2] + [p, v_3]) = [p, v_1] + [p, v_2 + v_3] = [p, v_1 + (v_2 + v_3)] = [p, (v_1 + v_2) + v_3] = [p, v_1 + v_2] + [p, v_3] = ([p, v_1] + [p, v_2]) + [p, v_3]$ since addition in $\mathbb{R}^n$ is associative.

Zero Vector:
Have $\mathbf{0}$ be the zero vector in $\mathbb{R}^n$. Then $[p, \mathbf{0}] \in \pi^{-1}(p)$, and $[p, v] + [p, 0] = [p, v + \mathbf{0}] = [p, v]$. $[p, \mathbf{0}]$ can therefore act as the zero vector for $\pi^{-1}(p)$.

Additive Inverses:
If $[p, v] \in \pi^{-1}(p)$ then $[p, -v] \in \pi^{-1}(p)$, and $[p, v] + [p, -v] = [p, v + -v] = [p, \mathbf{0}]$.

Associativity of Scalar Multiplication:
$a(b[p, v]) = a[p, bv] = [p, (ab)v] = (ab)[p, v]$.

Distributivity of Vector Addition:
$a([p, v_1] + [p, v_2]) = a[p, v_1 + v_2] = [p, a(v_1 + v_2)] = [p, av_1 + av_2] = [p, av_1] + [p, av_2] = a[p, v_1] + a[p, v_2]$.

Distributivity of Scalar Addition:
$(a + b)[p, v] = [p, (a + b)v] = [p, av + bv] = [p, av] + [p, bv] = a[p, v] + b[p, v]$.

By the above proofs it follows that $\pi^{-1}(p)$ is a vector space for all $p \in M$. If $p \in U \cap U'$ for $U, U' \in \Omega$, then if we choose to change our representative for the equivalence classes the resultant space will be isomorphic to our original via the map that takes $[p, v]$ to $[p, g_{U,U'}(p)(v)]$.

(ii): Local Trivializations:

Consider any $c \in M$. Since Ω is a cover of M, there exists some open $U \in \Omega$ s.t. $c \in U$. We will now define a new map,

$$\lambda_U : \pi^{-1}(U) \to \mathbb{R}^n$$
$$\text{s.t. } \lambda_U([p, v]) = v$$

Note that λ_U relies on a particular choice of U. if for instance, $p \in U \cap U'$, then $[p, v] = [p, v']$ if $v = g_{U,U'}(p)(v)$ even though v might not necessarily equal v'. In order to avoid this contention we simply fix a representation for each equivalence class given to us by our choice of U. Once we've done this, for each $p \in U$ every equivalence class in $\pi^{-1}(p)$ receives a unique representative in $\mathbb{R}^n$, and λ_U becomes well defined.

To see that λ_U is continuous, suppose that $C \subset \mathbb{R}^n$ is open. Then $\lambda_U^{-1}(C) = \{p \in U, v \in C \mid [p, v]\}$.

Have $\tau : \cup_{U\in\Omega}(U \times \mathbb{R}^n) \to E$ be the natural projection for the equivalence relation that defines E. Then $\tau^{-1}(\lambda^{-1}(C)) = U \times C$ which is open. Therefore $\lambda_U(C)$ is open in the quotient topology on E, and λ_U is continuous.

Now suppose that $H \subset \pi^{-1}(U)$ is open. Then $\tau^{-1}(H) = C \times D$ where $C \subset U$, $D \subset \mathbb{R}^n$ are both open. Since λ_U simply projects the second representative of every equivalence class in E, the image of H under λ_U will be the collection of v's s.t. $[p, v] \in H$. Since this set is D, it follows that $\lambda_U(H) = D$ which is open. Therefore for any open set $H \in E$, $\lambda_U(H)$ is open. λ_U is therefore open.

Now consider the restriction of π to $\pi^{-1}(U)$,$\pi|_U : \pi^{-1}(U) \to U$. By construction π is continuous so its restriction to $\pi^{-1}(U)$ will be continuous as well. Similarly to λ_U, if $H \subset \pi^{-1}(U)$ is open, then $\tau^{-1}(H) = C \times D$ where $C \subset U$, $D \subset \mathbb{R}^n$ are both open, and $\pi|_U = C$ which is open in U. Therefore $\pi|_U$ is open. Now consider the map,

$$\Phi_U : \pi^{-1}(U) \to U \times \mathbb{R}^n$$
$$\Phi_U([p, v]) = (\pi|_U([p, v]), \lambda_U([p, v]))$$

By the previous proofs, $\pi|_U, \lambda_U$ are both continuous and open, so Φ_U is continuous and open. To see that Φ_U is surjective simply note that for any $p \in U$, $\pi^{-1}(p) \subset \pi^{-1}(U)$ and $\lambda_U(\pi^{-1}(p)) = \mathbb{R}^n$. Therefore $\Phi_U(\pi^{-1}(U)) = (\pi|_U(\pi^{-1}(U)), (\lambda_U(\pi^{-1}(U))) = U \times \mathbb{R}^n$.

To see that Φ_U is injective, consider $\pi^{-1}(p)$ for any $p \in U$. Every $[p, v] \in \pi^{-1}(p)$ has a single unique representative, v, and Φ_U will therefore take each $[p, v]$ to (p, v) injectively. Φ_U is therefore injective fiberwise, which implies that Φ_U is injective on $\pi^{-1}(U)$ seeing as $\Phi_U([p, v]) = \Phi_U([p', v'])$ iff $p = \pi|_U([p, v]) = \pi|_U([p', v']) = p'$ which implies that $[p, v]$ and $[p', v']$ belong to the same fiber.

Φ_U is a Local Trivialization:

Φ_U is bijective, continuous, and open, which implies that Φ_U is a homeomorphism.
Since Φ_U is injective fiberwise, and $\lambda_U(\pi^{-1}(U)) = \mathbb{R}^n$ it follows that $\Phi_U : \pi^{-1}(p) \to \{p\} \times \mathbb{R}^n \cong \mathbb{R}^n$ is bijective, and therefore a vector space isomorphism.
$\pi_U \circ \Phi_U = \pi|_U = \pi$ on U.

$\therefore$ Φ_U fulfills all of the criteria of a local trivialization.

$\therefore$ For every $c \in M$ there exists some open neighborhood U of c and a local trivialization $\Phi_U : \pi^{-1}(U) \to U \times \mathbb{R}^n$.