CONNECTIONS

MATTHEW WOOLF

Abstract. Anyone who has ever seen any differential geometry in action has probably seen connections, and they may well have wondered what exactly they are. In this paper, I will try to give them an excruciatingly formal definition, as well as some intuition about why you should care.

Contents

1. Background

For us, everything in sight is a differential manifold. A differential manifold is a Hausdorff, second-countable, paracompact topological space X equipped with a sheaf of $\mathbb{R}$-algebras, $\mathcal{A}_X$, called the sheaf of smooth, or differentiable functions on X (which will be denoted $\mathcal{A}$ where there is no ambiguity), which is a subsheaf of the sheaf of continuous $\mathbb{R}$-valued functions, such that any point $x \in X$ has a neighborhood U homeomorphic to an open subset V of $\mathbb{R}^n$ (with $n < \infty$) for some n such that $\mathcal{A}$ restricted to U is isomorphic to the sheaf of C^∞ functions on V. A map between differentiable manifolds M and N is said to be differentiable if it is continuous, and the induced map on the sheaves of continuous functions restricts to a map on the sheaf of smooth functions.

A vector bundle over a manifold M is a manifold E and a surjective differentiable map $\pi : E \to M$ such that the fibers are isomorphic to $\mathbb{R}^r$ and for each point $x \in M$ there is a neighborhood U such that $\pi^{-1}(U)$ is diffeomorphic to $U \times \mathbb{R}^r$ over π. The manifold $M \times \mathbb{R}^r$ is called the trivial bundle. Many operations on vector spaces can also be performed on vector bundles, for example direct sum, tensor product, symmetric product, exterior product, and Hom (and hence also dualizing). We define p-forms to be sections of $\bigwedge^p(\mathcal{T}^*)$, and if E is a vector bundle, then E-valued p-forms are sections of $\bigwedge^p(\mathcal{T}^*) \otimes E$. If V is a vector space, then a V-values p-form is just a $M \times V$-valued p-form in the above sense.

We define a sheaf $\mathcal{T}_m$, called the tangent sheaf, by letting $\mathcal{T}_m(U)$ be the $\mathbb{R}$-linear derivations from $\mathcal{A}_M(U)$ to itself. Note that this is a sheaf of $\mathbb{R}$-Lie algebras, and of left $\mathcal{A}$-modules, but not of $\mathcal{A}$-Lie algebras. Since derivations on $\mathbb{R}^n$ are $\mathcal{A}$-linear combinations of the $\frac{\partial}{\partial x_i}$, this gives a locally free sheaf of rank n (where n is the dimension of the manifold). Let $\mathcal{M}_x$ be the unique maximal ideal of the stalk of $\mathcal{A}$ at x. We define $\mathcal{T}_x/\mathcal{M}_x\mathcal{T}_x$ to be the tangent space of M at x. One can think of the tangent space as consisting of derivations from $\mathcal{A}_x$ to $\mathbb{R}$. We can glue these together in a construction analogous to the étale space, to get TM, the tangent bundle of M, which is a vector bundle (note that this construction will work for any locally free sheaf, and conversely, the sheaf of sections of any vector bundle will be locally free). Sections of the tangent bundle (i.e. differentiable right inverses of the projection map) are called vector fields. Note that vector fields can also be thought of as global sections of $\mathcal{T}$, since given a function f and a vector field x, we look at the germ of f at x, and the tangent vector of X at x, say v, gives a well-defined real number. It is easy to verify that this gives a derivation, and that this map is in fact an isomorphism.

We define $\mathcal{D}^1$, the sheaf of first-order differential operators to be $\mathcal{A}\oplus\mathcal{T}$, as a sheaf of $\mathbb{R}$-vector spaces, with the $\mathcal{A}$-bimodule structure given by $f(g\oplus X) = fg\oplus fX$ and $(g\oplus X)f = (gf+Xf)\oplus fX$. We give it this module structure, because we want $(f(g\oplus X))(h) = f((g\oplus X)(h) = fgh + fXh$, and $((g\oplus X)f)h = (g\oplus X)(fh) = gfh + (Xf)h + f(Xh)$. We can in fact define an algebra of differential operators of which the first-order ones form a generating set, but this is not necessary for our purposes. We can also define first-order differential operators between locally free sheaves. By definition, $\mathcal{D}^1(\mathcal{E},\mathcal{F})$ will be $\mathcal{E}^*\otimes\mathcal{D}^1\otimes\mathcal{F}$, where both tensor products over $\mathcal{A}$, using the $\mathcal{A}$-bimodule structure of $\mathcal{D}^1$.D is a first-order differential operator in the above sense if and only if the map $f\mapsto\sigma(D(fs))$ is a first-order differential operator, where σ is a section of $\mathcal{E}^*$ and s is a section of $\mathcal{F}$.

Let G be a Lie group. Then G acts on the tangent bundle by left translation, and this preserves the Lie algebra structure of vector fields, so we can consider the Lie algebra of invariant vector fields, denoted $\mathfrak{g}$. Any such vector field is determined by its value at the identity of G, and any tangent vector at $e\in G$ gives rise to an invariant vector field, so there is an isomorphism between $\mathfrak{g}$ and the tangent space of G at e. Any homomorphism of Lie groups gives rise to a Lie algebra homomorphism of their invariant vector fields by taking the differential of the homomorphism at the identity.

2. Connections

Definition 2.1. Let E be a vector bundle and $\mathcal{E}$ be its sheaf of sections. A connection in E is an $\mathcal{A}$-linear map from $\mathcal{T}$ to $\mathcal{D}^1(\mathcal{E},\mathcal{E})$ (where the image of X in this map is denoted ∇_X) such that for all $f\in\mathcal{A}(M)$, $\nabla_X(fs) = f\nabla_X(s) + (Xf)s$.

We can think of a connection as a way of differentiating sections of a given vector bundle along a vector field. Sections of trivial vector bundles are just functions to the fiber, so a section of a vector bundle is given locally by a function. We can differentiate functions along vector fields, so we would like a way of differentiating sections of vector bundles along vector fields.

A homomorphism from $\mathcal{T}$ to $\mathcal{D}^1(\mathcal{E},\mathcal{E})$ can be thought of as an element of $\mathcal{T}^*\otimes\mathcal{E}^*\otimes\mathcal{D}^1\otimes\mathcal{E} \cong \mathcal{T}^*\otimes\mathcal{E}^*\otimes\mathcal{D}^1\otimes\mathcal{A}\otimes\mathcal{E} \cong (\mathcal{E}^*\otimes\mathcal{T})^*\otimes(\mathcal{E}^*\otimes\mathcal{D}^1\otimes\mathcal{A})$, which is the same as a homomorphism from $\mathrm{Hom}(\mathcal{E},\mathcal{T})$ to $\mathcal{D}^1(\mathcal{E},\mathcal{A})$. There is a map from $\mathcal{D}^1$ to $\mathcal{T}$ given by projection onto the vector field part, and so we have a map from $\mathcal{D}^1(\mathcal{E},\mathcal{A})$ to $\mathrm{Hom}(\mathcal{E},\mathcal{T})$. We know that there is a short exact sequence $0\to\mathcal{A}\to\mathcal{D}^1\to\mathcal{T}\to 0$, so applying $\mathrm{Hom}(\mathcal{E},\cdot)$, we get the short exact sequence $0\to\mathrm{Hom}(\mathcal{E},\mathcal{A})\to\mathcal{D}^1(\mathcal{E},\mathcal{A})\to\mathrm{Hom}(\mathcal{E},\mathcal{T})\to 0$, and hence the kernel of the most recent map is $\mathrm{Hom}(\mathcal{E}, \mathcal{A})$. By tracing out the definitions, we can see that those maps from $\mathrm{Hom}(\mathcal{E}, \mathcal{T})$ to $\mathcal{D}^1(\mathcal{E}, \mathcal{A})$ which come from connections are precisely those which split the above short exact sequence.

Let $0 \to E' \to E \to E'' \to 0$ be any short exact sequence of vector bundles; we will show that it splits (though not necessarily canonically), and hence, by the previous discussion, that any vector bundle admits a connection. We note that it suffices to give a section $S^2(E)^*$ which is positive-definite on the fibers, since then we will be able to take orthogonal complements. We can certainly do this for a trivial bundle, so we can do this on the local trivializations, but since the space of symmetric positive-definite bilinear forms is convex, we can patch these functions together using partitions of unity to get such a function on the entire vector bundle.

If we are given a connection on a vector bundle, we can often use it to get connections on other related bundles. For the following, let ∇ be a connection on E and ∇' a connection on E'. Then we can form a connection on $E \oplus E'$ by letting $\nabla''_X(s_1 \oplus s_2) = \nabla_X(s_1) + \nabla'_X(s_2)$. We get a connection on $E \otimes E'$ by letting $\nabla''_X(s_1 \otimes s_2) = \nabla_X(s_1) \otimes s_2 + s_1 \otimes \nabla'_X(s_2)$. We get a connection on E^* by letting $\nabla''_X(f)(s) = X(f(s)) - f(\nabla_X(s))$. We can similarly get connections on symmetric and exterior powers. In particular, a connection on the tangent bundle, called a linear connection, induces a connection on many of the most important vector bundles of differential geometry.

Notice that we can think of connections as homomorphisms from $\mathrm{Hom}(\mathcal{E}, \mathcal{T}) \to \mathcal{D}^1(\mathcal{E}, \mathcal{A})$, and that any two differ by a homomorphism from $\mathrm{Hom}(\mathcal{E}, \mathcal{T}) \to \mathrm{Hom}(\mathcal{E}, \mathcal{A})$, which is the same as a homomorphism from $\mathcal{T}$ to $\mathrm{Hom}(\mathcal{E}, \mathcal{E})$, or equivalently, an $\mathrm{End}(V)$-valued 1-form, where V is the fiber of E. This means that the vector space of such homomorphisms acts simply transitively on the set of connections, which is *ipso facto* an affine space, i.e. a set on which a vector space acts simply transitively. This means that the set of all connections (since we have shown it is nonempty) is an affine space. On any trivial bundle, we have a natural connection given by $\nabla_X(s) = Xs$, and hence any connection can be written as $\nabla'_X(s) = Xs + \alpha(X)(S)$.

We define a gauge transformation on a vector bundle to be a self-diffeomorphism under which the fibers are invariant. A gauge transformation A has an induced affine action on the space of connections given by $A^*(\nabla_X(s)) = A(\nabla_X((A^{-1})^*S))$. Any compact Lie group acting by affine transformations of any affine space has a fixed point, as can be seen by looking at an orbit and averaging with respect to the Haar measure. In particular, any compact Lie group of gauge transformations has a connection invariant under it.

3. Vector Bundles and Principal G-Bundles

Definition 3.1. Let G be a Lie group and M a manifold. A principal G-bundle over M is a manifold P with a surjective map $\pi : P \to M$ such that G acts on P on the right, and this action descends to a simply transitive action of G on the fibers, and furthermore, P is locally trivial, in the sense that there is an open cover U_α of M such that $\pi^{-1}(U)$ is diffeomorphic to $G \times U$ over π. A bundle is called trivial if it is diffeomorphic to $G \times M$ over π.

There is a very close relation between vector bundles and principal G-bundles. Let E be a vector bundle over M with fiber V. Then there is an associated principal $\mathrm{GL}(V)$-bundle, the frame bundle, which as a set is the disjoint union of all the ordered bases of E_x. Conversely, given any principal $\mathrm{GL}(V)$-bundle, we can get a vector bundle, by looking at $P \times V/\mathrm{GL}(V)$, where $(p, v)g = (pg, g^{-1}v)$.

There is another way of looking at bundles using transition functions. Let E be a vector bundle with fiber V and U_i, U_j two sets on which we have a local trivialization. Then we have maps $V \times (U_i \cap U_j) \to \pi^{-1}(U_i \cap U_j) \to V \times (U_i \cap U_j)$, where the first map comes from the trivialization over U_i and the second from the trivialization over U_j. This then gives us a map from $U_i \cap U_j \to \mathrm{GL}(V)$. Such maps (denoted m_{ij}) must satisfy $m_{ij}m_{jk} = m_{ik}$, called the cocyle condition. Conversely, given an open cover U_i of M and maps $m_{ij} : U_i \cap U_j \to \mathrm{GL}(V)$ satisfying the cocyle condition, we can look at $(\coprod U_i \times V)/\sim$, where for $x \in U_i \cap U_j$, $(x, v) \sim (x, m_{ij}(v))$.

We can also define transition functions for a principal G-bundle. Let P be a principal G-bundle, trivial on U_i and U_j. Then using the trivializations we get a map $G \times (U_i \cap U_j) \to G \times (U_i \cap U_j)$. The induced self-map of G commutes with right multiplication, so one can show that it must be left multiplication by an element of G, which is denoted $m_{ij}(g)$. Once again, the transition functions must satisfy the cocyle condition, and then may be called cocycles, and again such transition functions give rise to a principal G-bundle. Since vector bundles and principal $\mathrm{GL}(V)$-bundles are both determined by transition functions to $\mathrm{GL}(V)$, it seems eminently reasonable to suppose that they are in natural bijection. If you take different trivializations over an open cover for either principal G-bundles or vector bundles, then it turns out that the corresponding transition functions satisfy the same equivalence relation in either case (the cocycles are in this case called cohomologous), so this does lead to another proof of that fact. If $f : G \to H$ is a homomorphism of Lie groups, then we can compose with the cocycles of a principal G-bundle to get the cocycles for a principal H-bundle. This is called extension of the structure group. In particular, if $\rho : G \to GL(V)$ is a representation of a Lie group, then we obtain an associated vector bundle. If we have a homomorphism $f : G \to H$, and P is a principal H-bundle, then we say that P admits a reduction of the structure group to G if there is a principal G-bundle P' such that P is obtained from extending the structure group of P'. Reductions of the structure group are not necessarily unique, and they don't always exist. For example, a principal G-bundle admits a reduction of the structure group to the trivial group if and only if the bundle is trivial.

Let P be a principal G-bundle over M, and $H \subset G$. Then $P \to P/H$ is a principal H-bundle, so given a section s of $P/H \to M$, we can pull back P to M and extend the structure group to G to get a principal G-bundle isomorphic to P, and hence sections of P/H to M give rise to reductions of the structure group of G to H. Conversely, given P' a reduction of the structure group of P to H, then $P' \times G \to P$ is a principal H-bundle, so $(P' \times G)/H \to P$ is an isomorphism of principal G-bundles, and hence the quotients by H are isomorphic, but the quotient of $(P' \times G)/H$ by H is isomorphic to $M \times G/H$, which has a section, and hence so does P/H. We have thus proved that there is a one-to-one correspondence between reductions of the structure group of P to $H \subset G$ and sections of $P/H \to M$.

It is a theorem of linear algebra that there is a canonical homeomorphism from $\mathrm{GL}(V)/\mathrm{O}(V)$ to the set of symmetric positive-definite matrices, so reductions of the structure group of P to $\mathrm{O}(V)$ are the same as differentiable assignments of a symmetric positive-definite matrix to each point of M, but this is the same as a section of $S^2(E)^*$ which is positive-definite on fibers. Such a section is called a metric on E. If E is the tangent bundle of M, then this is called a Riemannian metric on M.

Let $P \to M$ be a principal G-bundle with projection map π. Then the pullback of P to P is canonically trivial. Let $f : G \to \mathrm{GL}(V)$ be a representation of G, and let E be the associated vector bundle. Let σ be a section of E. We can pull it back to a section of $\pi^*(E)$ (map $x \in P$ to $\pi(x)$, and look at $\sigma(x) \in E_{\pi(x)} \cong E_x$). But $\pi^*(E)$ is associated to $\pi^*(P)$, which is canonically trivial, and hence is canonically trivial itself, so this gives a function $g : P \to V$ such that $g(ph) = f(h)^{-1}g(p)$. Conversely, given such a function, by looking at the construction of associated vector bundles, we get a section of E.

4. Connection Forms on Principal G-Bundles

Let S be a manifold on which the Lie group G acts simply transitively. Then the tangent bundle of S is canonically isomorphic to $S \times \mathfrak{g}$ with the isomorphism sending (s, X) to the image of X under the differential at e of the map $G \to S$ given by $g \mapsto sg$. Projection onto $\mathfrak{g}$ hence gives a $\mathfrak{g}$-valued 1-form on S, called the Maurer-Cartan form, and denoted m_G.

Definition 4.1. Let P be a principal G-bundle over M. A connection form on P is a $\mathfrak{g}$-valued 1-form γ on P which is the Maurer-Cartan form on the fibers such that $\gamma(vg) = \mathrm{Ad}(g^{-1})\gamma(v)$, where $\mathrm{Ad}(g)$ is the action of G on $\mathfrak{g}$ given by the differential of the conjugation map.

Let E be a vector bundle. If E is trivial on U_i, then we have determined that we can write $\nabla_X(s) = Xs + \alpha_i(X)(s)$ there. Let m_{ij} be the transition function from U_i to U_j. Then on $U_i \cap U_j$, one can verify that we must have the condition $\alpha_i = m_{ij}\alpha_j m_{ij}^{-1} - (dm_{ij})m_{ij}^{-1}$. On the other hand, a similar argument shows that if P is a principal G-bundle with a connection form γ, then on $U \times G$, it is given by $(\mathrm{Ad}(g^{-1})\pi_1^*(\omega)) + \pi_2^*(m_G)$, where the π_i are the projections onto the first or second coordinate, and ω is a $\mathfrak{g}$-valued 1-form on U. Under a change of trivializations given by transition function m_{ij}, we require that $\omega_i = \mathrm{Ad}(m_{ij})\omega_j - (dm_{ij})m_{ij}^{-1}$. Since $\mathrm{GL}(V)$ is an open subset of $\mathrm{End}(V)$, which is a real vector space, we can canonically identify the tangent space at any point, and hence $\mathfrak{gl}(V)$ with $\mathrm{End}(V)$. This means that the data for a connection on a vector bundle and a connection form on the associated frame bundle are both given by $\mathrm{End}(V)$-valued 1-forms which must transform in exactly the same way, so we see that there is a canonical bijection between connections on a vector bundle and connections on the frame bundle. Notice also that as with connections, we can pull back connection forms, so on a vector bundle with a metric, we get a connection form on the associated $\mathrm{O}(V)$-bundle.

Let γ be a connection form on the principal G-bundle P. The vertical vectors in the tangent bundle of P are those which are mapped to 0 by the differential of the projection to M. The connection form γ maps the vertical space at each point isomorphically to $\mathfrak{g}$, so the kernel of this map, which will be called the horizontal space, is complementary to the vertical space. The differential of the projection map maps the horizontal space isomorphically to the tangent space at the corresponding point of M. Given a tangent vector on M, we can lift it uniquely to a tangent vector in the horizontal space of P at any point of P lying above the point of M.

In particular, we can use this to uniquely lift vector fields of M to horizontal vector fields of P.

Let E be a vector bundle with metric g, and ∇ a connection on E. ∇ is called a metric connection if it comes from a connection on the associated principal $\mathrm{O}(V)$-bundle. This condition is equivalent to the vanishing of $\nabla'_X(g)$ for all vector fields X, where ∇' is the associated connection to ∇ on $(S^2E)^*$. Expanding this out, we get $Xg(Y,Z) = g(\nabla_X Y, Z) + g(Y, \nabla_X Z)$. The group $\mathrm{O}(V)$ acts by gauge transformations on E, and is compact, so there must be a metric connection. In general, this metric connection will not be unique. However, if E is the tangent bundle, then we can impose another condition in order to make the metric connection unique.

Definition 4.2. Let ∇ be a linear connection. The torsion of ∇ is $\nabla_X Y - \nabla_Y X - [X,Y]$.

Proposition 4.3. *On any Riemannian manifold there is a unique metric connection with vanishing torsion. This connection is called the Levi-Civita connection.*

This is proved by expanding out $g(\nabla_X Y, Z)$ using the fact that ∇ is a metric connection, summing over the cyclic permutations of X, Y, and Z, and then using the vanishing of the torsion to express $g(\nabla_X Y, Z)$ independently of ∇. This then shows that $\nabla_X Y$ must be a section of the tangent bundle, and one verifies that it is indeed a connection.

Definition 4.4. Let ∇ be a connection on a vector bundle E. The curvature form of ∇ is the $\mathrm{End}(E)$-valued 2-form R_∇ given by $R_\nabla(X,Y) = \nabla_X\nabla_Y - \nabla_Y\nabla_X - \nabla_{[X,Y]}$. If γ is a connection on a principal G-bundle given locally by $\mathfrak{g}$-valued 1-forms α_i, then the curvature is $d\alpha_i + [\alpha_i, \alpha_i]$, where $[\alpha_i, \alpha_i](X,Y)$ is by definition $[\alpha_i(X), \alpha_i(Y)]$. A connection is called flat if the curvature vanishes identically.

It is relatively easy to verify that these conditions are compatible. One can show that given any principal G-bundle with a flat connection, locally the bundle can be trivialized along with the connection, an if the trivialization is on a connected open set, then the trivialization is unique up to translation by an element of G. As a consequence, any principal G-bundle with a flat connection is given by locally constant transition functions. This means that the transition functions are continuous functions to G given the discrete topology, so they give a principal G-bundle where G has the discrete topology, but this is just a covering space.

5. Holonomy and Parallel Transport

Definition 5.1. A path in a manifold M from $a \in M$ to $b \in M$ is a differentiable map f from a small open interval containing $[0,1]$ (which we will henceforth call I for convenience) to M such that $f(0) = a$ and $f(1) = b$.

Notice that since I is a one-dimensional manifold, it cannot have any 2-forms, and hence all connections on it are flat. Since it is simply connected, the previous discussion shows that all principal G-bundles over it with a given connection must be canonically trivial. By multiplying by the suitable element of G, we can find a canonical section which has any desired value at 0.

Now let f be a path in M, and P a principal G-bundle over M with a connection γ. We can pull back P to I, where it becomes trivial, and take a section whose value at 0 is any point of P in the fiber over a. By looking at the value of the section at other points of the interval, we get a unique lift of the path to P. In effect, we are taking the unique path starting at a given point whose tangent vector at each point is horizontal and maps to the tangent vector of the path in M. If the path is a loop, then the section is in the same fiber at 0 and 1, so they differ by a group element.

Definition 5.2. Let f be a loop in M such that $f(0) = f(1) = p$. The group element defined by f in the above construction is called the holonomy of f. The holonomy group at p is the set of all elements of G arising from holonomy of loops which start and end at p. The restricted holonomy group is the subgroup arising from holonomy of null-homotopic loops.

The holonomy comes from two different sources: the homotopy class of the loop, and the connection. The first one can be seen even in covering space theory, where loops which are not null-homotopic need not lift to loops. The restricted holonomy group is basically determined by the curvature, however. The Ambrose-Singer theorem states that the structure group of any principal G-bundle P with a connection form can be reduced to the holonomy group at each point, and that the connection form on P comes from a connection form on the reduced bundle. Furthermore, the Lie algebra of the restricted holonomy group is the subalgebra of $\mathfrak{g}$ generated by the values of the curvature form on the reduced bundle.

References

[1] S. Ramanan Global Calculus. American Mathematical Society. 2005.